

Overabundant Information and Learning Traps*

Annie Liang[†] Xiaosheng Mu[‡]

June 18, 2018

Abstract

We develop a model of social learning from overabundant information: Short-lived agents sequentially choose from a large set of (flexibly correlated) information sources for prediction of an unknown state. Signal realizations are public. We demonstrate two starkly different long-run outcomes: (1) efficient information aggregation, where the community eventually learns as fast as possible; (2) “learning traps,” where the community gets stuck observing suboptimal sources and learns inefficiently. Our main results identify a simple property of the signal correlation structure that separates these outcomes. In both regimes, we characterize which sources are observed in the long run and how often.

1 Introduction

In many learning problems, agents cannot design their information in a completely flexible way. Instead, they choose from a given, finite (though often large) set of information sources. For instance, a researcher studying depression cannot access arbitrarily precise signals about the role of stress. He can however acquire many kinds of information related to this question; for example, he might acquire neurochemical and genetic data from individuals undergoing stressful life changes, or collect observational data about simultaneous occurrences of depression and stress.

The value to acquiring new information depends on how that information relates to what is already known. Thus, past information acquisitions can shape the perception of which

*We thank Vasilis Syrgkanis for insightful comments in early conversations about this project. We are also grateful to Aislinn Bohren, Ben Golub, Paul Milgrom, Andrew Postlewaite, Ilya Segal, Carlos Segura, Rajiv Sethi, Andrzej Skrzypacz, and Yuichi Yamamoto for comments that improved this paper.

[†]University of Pennsylvania

[‡]Harvard University

kinds of information are currently most useful. For example, a better understanding of the relationship between serotonin and depression increases the value to comparing serotonin levels across individuals (as opposed to other neurochemicals whose roles are less well-understood). Models of information acquisition often avoid explicitly describing these informational complementarities, as they complicate the analysis.¹ But relationships across diverse sources can have significant implications for behavior, in particular in dynamic learning environments where information is passed down over time.

The main contribution of this paper is to identify a new externality driven by complementarities across sources, and to characterize the consequences of this externality for long-run aggregation of information. We show that past information acquisitions have the potential to shape future acquisitions in two starkly different ways.

- *Efficient information aggregation:* Past information helps agents to identify the “best” kinds of information. At all sufficiently late periods, a social planner cannot improve on the history of acquisitions.
- *Learning traps:* Past information pushes agents to acquire information that leads to inefficiently slow learning. Early suboptimal choices propagate over time.

Relationships across the entire *set* of information sources are relevant to which of these outcomes emerges, and our main results identify a key property that determines the outcome.

In our model, agents are indexed by (discrete) time and sequentially choose from a large number of information sources, each associated with a signal about a payoff-relevant state. We allow for flexible correlation across the sources by modeling each kind of information as a (noisy) linear combination of the payoff-relevant state and a set of “confounding” variables. After acquiring information, each agent predicts the payoff-relevant state.

In contrast to the classic sequential learning model (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992; Smith and Sorenson, 2000), we suppose that signal realizations are public, so that predictions are based on the history of realizations thus far. This departure permits us to focus on the externalities created by agents’ choice of *kind* of information, as opposed to the more frequently studied frictions that emerge from inference.

We are mainly interested in settings with many sources of information, including some that are redundant. Formally, agents can completely learn the payoff-relevant state from exclusive (repeated) observation of signals from each of several possible subsets. As a benchmark, we first derive the *optimal* long-run frequency of signal acquisitions. These correspond to the choices that maximize information revelation about the payoff-relevant state, and also to the choices that maximize a discounted sum of agent payoffs (in a patient limit).

¹Exceptions include [Borgers, Hernando-Veciana and Krahmer \(2013\)](#), [Chen and Waggoner \(2016\)](#), and [Chade and Eeckhout \(2018\)](#) among others.

Our main results demonstrate that whether society’s acquisitions converge to this optimal long-run frequency depends critically on how many signals are needed to identify the payoff-relevant state. The key intuition refers back to an observation made in [Sethi and Yildiz \(2016\)](#): An agent who repeatedly observes a source confounded by an unknown parameter learns *both* about the payoff-relevant state and also about the confounding term, and hence improves his interpretation of *this* source over time. In our setting, where a single confounding term can affect multiple sources, there is a further spillover effect: Learning from one source helps agents to interpret information from *all* sources confounded by the same parameters.

Suppose that in order to learn the payoff-relevant state, agents must observe a set of sources that reveals all of the confounding terms as well. Then endogenously, agents will acquire information that (collectively) reveals all of the unknowns. This will lead agents to evaluate *all* sources by an “objective” asymptotic criterion, which then reveals the best set of sources. More formally, we obtain the following result: If K sources are required to recover the payoff-relevant state (where K is also the number of unknown states), then long-run acquisitions are optimal starting from any prior belief.

In contrast, if it is possible to learn the payoff-relevant state without recovering all of the confounding terms, then agents can persistently undervalue sources that provide information confounded by these remaining variables. Thus, long-run learning may be inefficient. Our second main result says that any set of fewer than K sources that recovers the payoff-relevant state creates a “learning trap” under some set of prior beliefs. We further show that the long-run inefficiency under a learning trap—measured as the ratio of optimal aggregated payoffs and achieved aggregated payoffs (under a social planner criterion)—can be arbitrarily large.

The basic friction here is that investment in learning about confounding terms is socially beneficial, but not necessarily optimal for individuals. Our main results show that this wedge between individual incentives and social objectives does not preclude long-run efficient learning. When the available kinds of information are related in certain ways, individual incentives will nevertheless endogenously drive individuals to acquire information in a way that is socially efficient.

In the remaining cases, interventions may be needed to transition agents towards better sets of sources. In the final part of our paper, we study possible such interventions. We show that policymakers can restore efficient information aggregation by providing certain kinds of free information (that we characterize), or by reshaping the reward structure so that agents’ predictions are based on information that they acquire over many periods. The success of these interventions depends on specific features of the informational environment, as we will discuss.

1.1 Related Literature

A recent literature considers choice from different kinds of information sources (Sethi and Yildiz, 2016; Che and Mierendorff, 2017; Fudenberg, Strack and Strzalecki, 2017; Liang, Mu and Syrgkanis, 2017; Mayskaya, 2017; Sethi and Yildiz, 2017). We build upon Liang, Mu and Syrgkanis (2017) in particular, which introduced the framework we use here (see Section 2) under the restriction that all of the sources are required to learn the payoff-relevant state. This restriction rules out the possibility of overabundant information, which is the focus of the present paper.

Sethi and Yildiz (2016, 2017) study long-run (myopic) acquisitions from a large number of Gaussian sources, as we do. Our model differs from this work in a few key ways: First, Sethi and Yildiz (2016, 2017) consider stochastic error variances, so that the “best” sources vary from period to period; in contrast, we fix error variances, so there is (generically) a unique “best” asymptotic set. Second, Sethi and Yildiz (2016, 2017) focus on correlation structures that fall under our “learning traps” result (part (a) of Theorem 2), while we explore arbitrary correlation structures and show that many lead to optimal learning.

Our model builds on the social learning and herding literatures (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992; Smith and Sorenson, 2000), which consider information aggregation by short-lived agents who sequentially acquire information. At a high level, the externality identified in our paper relates to the classic externality from this literature: In both settings, the precision of public information can grow inefficiently slowly because of endogenous information acquisitions driven by past choices. But in the present paper, all signal realizations are publicly and perfectly observed, which turns off the inference problem essential to the existence of cascades in standard herding models. Our focus is on a new mechanism, in which externalities arise through choice of *kind* of information; as we will see, this externality has a rather different structure.

Our setting with choice of information connects to Burguet and Vives (2000), Mueller-Frank and Pai (2016), and Ali (2018), which introduced endogenous information acquisition to social learning models. Relative to this work, our paper considers choice from a fixed set of information sources (with a capacity constraint), in contrast to choice from a flexible set of information sources (with a cost on precision). Our results focus on the *speed* of learning, as in Vives (1992), Golub and Jackson (2012), Hann-Caruthers, Martynov and Tamuz (2017), and Harel et al. (2018) among others.

Our motivation of the optimal frequency in Section 4 is related to the experimental design literature in statistics, and in particular to the notion of c -optimality (choice of t experiments to minimize the posterior variance of an unknown state). Chaloner (1984) showed that a c -optimal design exists on at most K points. Our Theorem 1 extends this result, supplying

a characterization of the optimal design itself and demonstrating uniqueness.² We further show that the optimal frequency maximizes a discounted objective in the patient limit. Finally, our main results admit a re-interpretation that connects to a literature on learning convergence in potential games, with more detail given in Section 6.

2 Framework

There are K persistent unknown states: a payoff-relevant state ω and $K - 1$ confounding states b_1, \dots, b_{K-1} . We assume that the state vector $\theta := (\omega, b_1, \dots, b_{K-1})'$ follows a multivariate normal distribution $\mathcal{N}(\mu^0, \Sigma^0)$,³ where the prior covariance matrix Σ^0 has full rank.⁴

Agents have access to N *kinds* or *sources* of information. Observation of source i produces an independent realization of the random variable

$$X_i = \langle c_i, \theta \rangle + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, 1)$$

where $c_i = (c_{i1}, \dots, c_{iK})'$ is a vector of constants, and the error terms ϵ_i^t are independent from each other and over time. It is without loss to normalize the error terms as we do above, since the coefficients c_i are unrestricted (thus, signals can be of differing precision levels). Throughout, we take C to be the $N \times K$ coefficient matrix whose i -th row is c_i' .

The payoff-irrelevant states produce correlations across the sources, and we can interpret these states for example as:

- *Confounding explanatory variables:* Observation of signal i produces the (random) outcome $y = \omega c_i^1 + b_1 c_i^2 + \dots b_{K-1} c_i^K + \epsilon_i$, which depends linearly on an observable characteristic vector c_i . For example, y might be the average incidence of depression in a group of individuals with characteristics c_i . The state of interest ω is the coefficient on a given characteristic c_i^1 , and the payoff-irrelevant states are the unknown coefficients on the auxiliary characteristics. Different sources represent subpopulations with different characteristics.
- *Knowledge and technologies that aid interpretation of information:* Interpret each signal as a measurement. For example, researchers studying depression can acquire measurements of various neurochemicals from lab subjects. Neurochemicals differ in how

²Another difference is that Chaloner (1984) studies the optimal continuous design, while we impose an integer constraint on signal counts.

³All vectors in this paper are column vectors.

⁴The full rank assumption is without loss of generality: If there is linear dependence across the states, the model can be mapped into an equivalent setting that has a lower dimensional state space and that satisfies the full rank condition.

precisely they can be measured (measurement error) and how well their role is understood (precision of interpretation). The confounding terms represent the quality of these different measurement technologies, and also the degree of background contextual knowledge.

Agents indexed by (discrete) time $t \in \mathbb{N}$ move sequentially. Each agent t chooses one of the N signals and observes an independent realization. He then predicts the state ω , receiving a payoff of $-(a - \omega)^2$. We assume throughout that all signal realizations are public; thus, predictions are based on the entire history of signal acquisitions and realizations.

Throughout, the set of signals is indexed by $[N] = \{1, \dots, N\}$. We call a subset of signals $\mathcal{S} \subset [N]$ *spanning* if the vectors $\{c_i : i \in \mathcal{S}\}$ span the coordinate vector $e_1 = (1, 0, \dots, 0)' \in \mathbb{R}^K$, so that it is possible to learn the payoff-relevant state ω by repeatedly observing signals from only \mathcal{S} . We call \mathcal{S} *minimally spanning* if it is spanning, and no proper subset of \mathcal{S} is spanning.

We assume in this paper that the complete set of signals $[N]$ is spanning, so that the payoff-relevant state can be recovered by observing all signals infinitely often.⁵ This assumption nests two interesting cases. Say that the informational environment has *exactly sufficient information* if $[N]$ is minimally spanning. Then, it is possible to recover ω by observing each information source infinitely often, but not by exclusively observing any proper subset of sources.

We are primarily interested in settings of *informational overabundance*, where $[N]$ is spanning but not minimally spanning. Multiple different subsets of signals allow for recovery of ω , and a key point of our analysis is to compare the set of sources that “should” be observed in the long run with the set of sources that is in fact observed in the long run. Except for trivial cases, informational overabundance corresponds to $N > K$ (more signals than states).⁶

⁵This assumption is without loss, and our results do extend to situations where ω is *not* identified from the available signals. To see this, we first take a linear transformation and work with the following equivalent model: The state vector $\tilde{\theta}$ is K -dimensional *standard Gaussian*, each signal $X_i = \langle \tilde{c}_i, \tilde{\theta} \rangle + \epsilon_i$, and the payoff-relevant parameter is $\langle u, \tilde{\theta} \rangle$ for some fixed vector u . Let R be the subspace of \mathbb{R}^K spanned by $\tilde{c}_1, \dots, \tilde{c}_N$. Then project u onto R : $u = r + w$ with $r \in R$ and w orthogonal to R . Thus $\langle u, \tilde{\theta} \rangle = \langle r, \tilde{\theta} \rangle + \langle w, \tilde{\theta} \rangle$. By assumption, the random variable $\langle w, \tilde{\theta} \rangle$ is independent from any random variable $\langle c, \tilde{\theta} \rangle$ with $c \in R$ (because they have zero covariance). Thus the uncertainty about $\langle w, \tilde{\theta} \rangle$ cannot be reduced upon any signal observation. Consequently, agents only seek to learn about $\langle r, \tilde{\theta} \rangle$, returning to the case where the payoff-relevant parameter *is* identified.

⁶It is possible for ω to be “overidentified” from a set of $N \leq K$ signals, e.g. $X_1 = \omega + \epsilon_1$, $X_2 = \omega + b_1 + b_2 + \epsilon_2$, and $X_3 = b_1 + b_2 + \epsilon_3$. In this case, the set $\{X_1, X_2, X_3\}$ is spanning, but not minimally spanning since both of its subsets $\{X_1\}$ and $\{X_2, X_3\}$ are also spanning. Although $N = K = 3$ in this example, it is equivalent to a model in which there is a single confounding term $\tilde{b}_1 = b_1 + b_2$, and the three signals are rewritten $X_1 = \omega + \epsilon_1$, $X_2 = \omega + \tilde{b}_1 + \epsilon_2$ and $X_3 = \tilde{b}_1 + \epsilon_3$. Then we do have $N > K$ in this equivalent model.

3 Preliminaries

Each agent t faces a history $h^{t-1} \in ([N] \times \mathbb{R})^{t-1} = H^{t-1}$ consisting of all past signal choices and their realizations. Write $\theta \sim \mathcal{N}(\mu^{t-1}, \Sigma^{t-1})$ for the agent's beliefs about the state vector, prior to making his own signal choice. Given an observation of signal i , his posterior beliefs become $\theta \sim \mathcal{N}(\mu^t, \Sigma^t)$. The posterior covariance matrix Σ^t is a *deterministic* function of the prior covariance matrix Σ^{t-1} and the signal choice X_i , and the posterior mean is the random vector $\mu^t \sim \mathcal{N}(\mu^{t-1}, \Sigma^{t-1} - \Sigma^t)$.

Agent t 's posterior belief about *the payoff-relevant state* ω (after observing his signal) is given by $\omega \sim \mathcal{N}(\mu_1^t, \Sigma_{11}^t)$,⁷ and his maximum expected payoff (corresponding to prediction of his posterior mean) is $-\Sigma_{11}^t$. Thus, the signal acquisition that maximizes agent t 's payoffs is the one that minimizes his posterior variance about ω .

We can track society's acquisitions as a sequence of division vectors

$$m(t) = (m_1(t), \dots, m_N(t))$$

where $m_i(t)$ is the number of times that signal i has been observed up to and including time t . Let $V(q_1, \dots, q_N)$ denote posterior variance about ω , given the initial prior covariance matrix Σ^0 and q_i observations of each signal i .⁸ Then, $m(t)$ evolves deterministically according to the following rule: $m(0)$ is the zero vector, and for each time $t \geq 0$ and signal i ,

$$m_i(t+1) = \begin{cases} m_i(t) + 1 & \text{if } V(m_i(t) + 1, m_{-i}(t)) \leq V(m_j(t) + 1, m_{-j}(t)) \ \forall j. \\ m_i(t) & \text{otherwise.} \end{cases}$$

That is, in each period t the division vector increases by 1 in exactly one coordinate, corresponding to the signal that allows for the greatest immediate reduction in posterior variance. We allow ties to be broken arbitrarily, so there may be multiple possible paths $m(t)$. We are interested in the *long-run frequencies* of observation $\lim_{t \rightarrow \infty} m_i(t)/t$ for each signal i . This describes the fraction of periods that is (eventually) devoted to each signal.⁹ Note the possibility for some signals to have zero long-run frequency.

4 Optimal Benchmark

We begin by characterizing the optimal frequency with which each signal *should* be viewed in the long run, both from a perspective of information revelation and also from the perspective

⁷Subscripts indicate particular entries of a vector or matrix.

⁸For a normal prior and normal signals, the posterior covariance matrix does not depend on signal realizations. See Appendix A.1 for the complete (closed-form) expression for V .

⁹We show in Section 5 that the limit exists under a mild technical assumption.

of a social planner aggregating agent payoffs. These frequencies will serve as a benchmark against which to evaluate the actual social acquisitions.

4.1 Definition and Motivation

Consider the one-shot problem in which some t observations can be allocated across the available signals, and the goal is to minimize posterior variance about ω . Call any solution

$$n(t) \in \underset{(q_1, \dots, q_K): q_i \in \mathbb{Z}^+, \sum_i q_i = t}{\operatorname{argmin}} V(q_1, \dots, q_K)$$

a *t-optimal division vector*. Signal acquisitions prior to time t maximize information revelation about ω if and only if their composition is described by $n(t)$. Thus, we will refer to $n(t)$ as the *optimal count* over signals.

Define $\lambda^{OPT} = \lim_{t \rightarrow \infty} n(t)/t$ to be the limiting frequency over signals corresponding to these t -optimal division vectors; that is, λ_i^{OPT} is the fraction of periods that is (asymptotically) optimally devoted to signal i . These frequencies are well-defined and unique under a subsequent condition (Assumption 1), and we will refer to λ^{OPT} from here on as the *optimal (long-run) frequency vector*.

The above justification for λ^{OPT} is purely in terms of information revelation. The frequency vector λ^{OPT} also emerges, however, as an approximation of the (limiting patient) solution to a social planner problem, in which signals and actions are chosen to maximize a discounted aggregate of payoffs across individuals:

$$U_\delta := \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \cdot u(a_t, \omega) \right].$$

For fixed δ , let $d_\delta(t)$ be the vector of signal counts (up to period t) associated with any strategy that maximizes U_δ . The long-run frequency vector produced by this optimal sampling procedure is $\lim_{t \rightarrow \infty} d_\delta(t)/t$, and the result below says that λ^{OPT} approximates this long-run frequency vector in the patient limit $\delta \rightarrow 1$.

Proposition 1. *Under the subsequent Assumption 1, for any $\epsilon > 0$, there exists $\underline{\delta} < 1$ such that for any $\delta \geq \underline{\delta}$ it holds that*

$$\limsup_{t \rightarrow \infty} \left\| \frac{d_\delta(t)}{t} - \lambda^{OPT} \right\| \leq \epsilon.$$

Here $\|\cdot\|$ represents the Euclidean norm.

This result further justifies our interest in λ^{OPT} .¹⁰ Throughout, we will say that *efficient information aggregation* occurs if society's frequency of observations tends to λ^{OPT} .

¹⁰It is clear that the strategy that samples signals (randomly) according to λ^{OPT} is best among *stationary* information acquisition strategies as $\delta \rightarrow 1$. This is however weaker than our result, because the optimal strategy for any fixed δ may be far from stationary.

4.2 Characterization

The main difficulty in characterizing the optimal frequency vector λ^{OPT} is that we do not know its support. Because ω is recoverable from different proper subsets of signals, there are many signals that could conceivably receive zero frequency in the long-run.

As a first step, we thus consider a simpler version of the social planner problem in which agents can acquire signals from (only) some minimal spanning set \mathcal{S} . By definition, the setting is one of exactly sufficient information, where all signals must be observed in order to recover ω . It is straightforward to show that the first coordinate vector e_1 can be decomposed as a (unique) linear combination of signals in \mathcal{S} :

$$e_1 = \sum_{i \in \mathcal{S}} \beta_i^{\mathcal{S}} \cdot c_i,$$

where the coefficients $\beta_i^{\mathcal{S}}$ are non-zero.

We showed in prior work that each signal $i \in \mathcal{S}$ should be observed (asymptotically) in proportion to its coefficient $\beta_i^{\mathcal{S}}$:

Proposition 2 (Liang, Mu and Syrgkanis (2017)). *Suppose agents are constrained to a minimal spanning set \mathcal{S} . Then, for every signal $i \in \mathcal{S}$, the optimal count satisfies*

$$n_i^{\mathcal{S}}(t) = \frac{|\beta_i^{\mathcal{S}}|}{\sum_{j \in \mathcal{S}} |\beta_j^{\mathcal{S}}|} \cdot t + O(1) \quad (1)$$

where $O(1)$ represents a residual term that remains bounded as $t \rightarrow \infty$.

Thus, the optimal long-run frequency with which signal i is viewed is $|\beta_i^{\mathcal{S}}| / \left(\sum_{j \in \mathcal{S}} |\beta_j^{\mathcal{S}}| \right)$.

To understand these critical coefficients $\beta_i^{\mathcal{S}}$, consider (for simplicity) the case in which the set \mathcal{S} has size K . The coefficient vectors associated with signals in \mathcal{S} have full rank,¹¹ and we let $C_{\mathcal{S}}$ denote the matrix of these coefficient vectors. The (random) vector of realizations corresponding to one observation of each signal in this set can be written as

$$Y = (y_1, \dots, y_K)' = C_{\mathcal{S}}\theta + \varepsilon$$

where ε is a $K \times 1$ vector of error terms. Given these realizations, the best linear unbiased estimate for ω is $\hat{\omega} = [C_{\mathcal{S}}^{-1}Y]_{11}$. Perturbing the realization of signal i by δ_i changes this estimate by $[C_{\mathcal{S}}^{-1}]_{1i} \cdot \delta_i$. One can show that the coefficients $\beta_i^{\mathcal{S}} = |[C_{\mathcal{S}}^{-1}]_{1i}|$, so the larger $\beta_i^{\mathcal{S}}$ is, the more $\hat{\omega}$ responds to changes in the realization of signal i . Proposition 2 thus says that agents should observe more frequently those signals whose realizations more strongly influence the best linear estimate of ω .

This proposition additionally implies the following corollary regarding the speed of learning from \mathcal{S} :

¹¹Otherwise, \mathcal{S} would not be a *minimal* spanning set.

Corollary 1. *Under optimal sampling from any minimal spanning set \mathcal{S} , the minimum achievable posterior variance after t observations satisfies:*

$$V(n^{\mathcal{S}}(t)) \sim \phi(\mathcal{S})^2/t := \left(\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}| \right)^2 / t.$$

where the notation “ $F(t) \sim G(t)$ ” means $\lim_{t \rightarrow \infty} \frac{F(t)}{G(t)} = 1$.

In what follows, we work with the simpler statistic $\phi(\mathcal{S})$ (roughly an *asymptotic standard deviation*), noting that the smaller $\phi(\mathcal{S})$ is, the faster the community learns from \mathcal{S} . We assume throughout that there is a *best* minimal spanning set according to ϕ :

Assumption 1 (Unique Minimizer). $\phi(\mathcal{S})$ has a unique minimizer \mathcal{S}^* among minimal spanning sets $\mathcal{S} \subset [N]$.

This assumption is a restriction on the coefficient matrix C , and it rules out examples such as the following:

Example 1. The available signals are $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, $X_3 = \omega + b_2 + \epsilon_3$, and $X_4 = b_2 + \epsilon_4$. Assumption 1 fails, because learning occurs equally fast from either of the minimal spanning sets $\{X_1, X_2\}$ and $\{X_3, X_4\}$.

We note that Assumption 1 holds generically—for example, it holds given arbitrarily small perturbations of the above environment.¹²

If we restrict agents to sample exclusively from a single minimal spanning set, then the optimal sampling rule (under Assumption 1) is clearly the frequency vector $\lambda^* \in \Delta^{N-1}$ satisfying

$$\lambda_i^* = \begin{cases} \frac{|\beta_i^{\mathcal{S}^*}|}{\sum_{j \in \mathcal{S}^*} |\beta_j^{\mathcal{S}^*}|} & \forall i \in \mathcal{S}^* \\ 0 & \forall i \notin \mathcal{S}^* \end{cases} \quad (2)$$

This sampling rule assigns zero frequency to signals outside of the set \mathcal{S}^* , and samples signals within \mathcal{S}^* according to the frequencies given in Proposition 2.

In principle, a social planner may improve on λ^* by sampling from multiple spanning sets. Our first theorem shows to the contrary that λ^* remains optimal when arbitrary sampling procedures are permitted. So long as C satisfies Unique Minimizer, the best long-run strategy is to restrict to the best minimal spanning set and sample from that set as in Proposition 2.

¹²Throughout the paper, “generic” means with probability 1 for signal coefficients c_{ij} randomly drawn from a full support distribution on \mathbb{R}^{NK} .

Theorem 1. *Let λ^* be given by (2). Under Unique Minimizer, the optimal long-run frequency vector satisfies $\lambda^{OPT} = \lambda^*$.*¹³

The conclusion can be loosely interpreted as stating that λ^* is the “most efficient linear representation” of the payoff-relevant state in terms of the signal coefficients.¹⁴

We show in Appendix B.1 that Assumption 1 is necessary for this result: In the environment described in Example 1, there are priors given which it is strictly optimal to observe all four available signals with positive frequency.

5 Main Results

In general, we may expect a difference between the best one-shot allocation of t acquisitions, described in the previous section, and the set of t acquisitions chosen by sequential (short-lived) decision-makers. We turn now to characterizing the latter.

Sections 5.1 and 5.2 focus on welfare evaluation—that is, will society’s acquisitions $m(t)$ eventually approximate the optimal acquisitions $n(t)$? We show that signal correlation structures can be classified into two kinds—those for which efficient information aggregation is guaranteed (starting from all prior beliefs), and those for which “learning traps” are possible (depending on the prior belief). Separation of these two classes depends critically on how many signals are required to identify ω .

For a given signal structure, the set of signals that are observed in the long run potentially depends on the prior belief. A related question is then *which* long-run outcomes are possible for a given signal structure, when we allow for arbitrary priors. We provide a complete characterization in Section 5.3.

5.1 Learning Traps

The following simple example demonstrates that sequential information acquisition need not lead to efficient information aggregation. Indeed, the set of signals that are observed in the long run can be disjoint from the optimal set.

¹³This theorem equivalently states that $n_i(t) \sim \lambda_i^* \cdot t$. In Appendix A.2.7, we show the stronger result that $n_i(t) = \lambda_i^* \cdot t + O(1)$.

¹⁴Specifically, consider the following constrained minimization problem: $\min \sum_{i=1}^N |\beta_i|$ subject to $\sum_{i=1}^N \beta_i \cdot c_i = e_1$. It can be shown by linear programming that the minimum is attained exactly when $\beta_i = \beta_i^{\mathcal{S}^*}$ (that is, when focusing on a single minimal spanning set).

Example 2. There are three available signals:

$$\begin{aligned} X_1 &= \omega + \epsilon_1 \\ X_2 &= 3\omega + b_1 + \epsilon_2 \\ X_3 &= b_1 + \epsilon_3 \end{aligned}$$

Both $\{X_1\}$ and $\{X_2, X_3\}$ are minimal spanning sets, but $\{X_2, X_3\}$ is optimal.¹⁵

Consider a prior where ω and b_1 are independent, and the prior variance of b_1 is large (exceeds 8). In the first period, observation of X_1 is most informative about ω , since X_2 is perceived as a noisier signal *about* ω than X_1 , and observations of X_3 provide information only about the confounding term b_1 (which is uncorrelated with ω). Thus the best choice is to observe X_1 . This observation does not affect the variance of b_1 , so the same argument shows that every agent observes signal X_1 .¹⁶ We refer to $\{X_1\}$ in this example as a *learning trap*.

Generalizing this example, the result below (stated as a corollary, since it will follow from the subsequent Theorem 2) gives a sufficient condition for learning traps. We impose the following (generic) assumption on the informational environment, which requires that every set of $k \leq K$ signals is linearly independent:

Assumption 2 (Strong Linear Independence). $N \geq K$ and every $K \times K$ submatrix of C is of full rank.¹⁷

Corollary 2. Assume Strong Linear Independence. For every minimal spanning set \mathcal{S} with $|\mathcal{S}| < K$, there exists an open set of prior beliefs given which agents exclusively observe signals from \mathcal{S} .

Thus, every small set (fewer than K signals) that identifies ω is a candidate learning trap.

In special environments, simple bounds on the extent of inefficiency are possible. For example, if there is an unbiased signal $c\omega + \epsilon$, then the posterior variance at each time t cannot exceed $1/(c^2t)$, and so aggregated payoffs must be at least $-\sum_{t \geq 1} \delta^{t-1}/(c^2t)$.¹⁸ The

¹⁵It is straightforward to verify that $\phi(\{X_1\}) = 1 > 2/3 = \phi(\{X_2, X_3\})$. Note also that $X_2 + X_3$ is an unbiased signal about ω , and it is more informative than two realizations of X_1 ; this demonstrates that $\{X_2, X_3\}$ is the best minimal spanning set without direct computation of ϕ .

¹⁶Note that the existence of learning traps is not special to the assumption of normality. We report a related example with non-normal signals in Appendix B.3.

¹⁷Besides trivial cases with redundant signals, Strong Linear Independence also rules out settings such as the following: $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, $X_3 = 2\omega + b_2 + \epsilon_3$, $X_4 = b_2 + \epsilon_4$, $X_5 = 3\omega + b_3 + \epsilon_5$, and $X_6 = b_3 + \epsilon_6$. Then $K = 4$ but the four signals X_1, X_2, X_3, X_4 are *not* linearly independent.

¹⁸This is because each agent can at least sample this unbiased signal and improve the posterior precision (i.e., inverse of the posterior variance) by c^2 . We thank Andrzej Skrzypacz for this observation.

size of inefficiency cannot be uniformly bounded across environments, however. Specifically, for any positive number L , there exists a set of signals and a prior belief given which

$$\frac{\phi(\mathcal{S})}{\phi(\mathcal{S}^*)} > L$$

where \mathcal{S} is the set of signals observed in the long run with positive frequency, and \mathcal{S}^* is the optimal set. This can be shown by direct construction: Modify the example above so that $X_2 = \alpha\omega + b_1 + \epsilon_2$, with α sufficiently large. For every choice of α , there is a set of priors given which X_1 is again exclusively observed.¹⁹ Thus, society’s long-run speed of learning can be arbitrarily slower than the optimal speed. This also implies that—from the perspective of a social planner who maximizes δ -discounted payoffs—the payoff ratio between society’s sequential acquisitions and optimal acquisitions can be arbitrarily large as $\delta \rightarrow 1$.²⁰

Finally, note that learning traps can emerge even if the best spanning set has the smallest size among all minimal spanning sets. We demonstrate this below.

Example 3. The available signals are:

$$X_1 = 10\omega + b_1 + \epsilon_1$$

$$X_2 = b_1 + \epsilon_2$$

$$X_3 = \omega + b_2 + \epsilon_3$$

$$X_4 = b_2 + b_3 + \epsilon_4$$

$$X_5 = b_3 + \epsilon_5$$

Here, the best set is $\{X_1, X_2\}$, which is also the minimal spanning set of lowest cardinality. But suppose the prior belief is such that ω, b_1, b_2, b_3 are independent, and there is initially high uncertainty about b_1 and low uncertainty about b_2 and b_3 . Then, agents suboptimally acquire only observations of X_3, X_4 , and X_5 .

5.2 Efficient Information Aggregation

Suppose in contrast to the previous section that repeated observation of K sources is required to recover ω . Our next result shows that a very different long-run outcome obtains: Starting from *any* prior, information acquisition eventually approximates the optimal frequency. Thus, even though agents are short-lived (“myopic”), they will end up acquiring information in a way that is socially best.

¹⁹The region of inefficient priors (that result in suboptimal learning) does decrease in size as the level of inefficiency increases. As α increases, the prior variance of b_1 has to increase correspondingly in order for the first agent to choose X_1 .

²⁰To derive this payoff comparison, note that $\phi(\mathcal{S})/\phi(\mathcal{S}^*) > L$ implies the ratio of flow payoffs in any late period t is larger than L . Using the fact that the harmonic series diverges, we know that as $\delta \rightarrow 1$, these (later) payoffs dominate the total payoffs from the initial periods.

Corollary 3. *Under Unique Minimizer, if every minimal spanning set has size K , then starting from any prior belief, it holds that $m_i(t) \sim \lambda_i^* \cdot t$ for every signal i .*²¹

A detailed intuition for this result appears in the subsequent Section 6.

The condition that all minimal spanning sets have size K is generically satisfied.²² However, if we expect that sources are endogenous to design or strategic motivations, the relevant informational environments may not fall under this condition. For example, the existence of an unbiased signal about ω (that is, $X = c\omega + \epsilon$) is non-generic, but plausible in practice. Sets of signals that partition into different groups (with group-specific confounding terms) are also economically interesting but non-generic. The previous Corollary 2 shows that inefficiency is a possible outcome in these cases.

5.3 Characterization of Long-Run Outcomes

We now provide a complete characterization of the possible long-run observation sets for any environment. Here we need to consider *subspaces spanned by different signal sets*. Formally, for any spanning set of signals \mathcal{A} , let $\overline{\mathcal{A}} \subseteq [N]$ be the set of available signals whose coefficient vectors belong to the subspace spanned by signals in \mathcal{A} . We say that a minimal spanning set \mathcal{S} is *subspace-optimal* if it uniquely maximizes the speed of learning among feasible sets of signals within its subspace.

Definition 1. *A minimal spanning set \mathcal{S} is subspace-optimal if it uniquely minimizes ϕ among all subsets of $\overline{\mathcal{S}}$ that are minimally spanning.*

Example 4. Suppose the available signals are $X_1 = \omega + \epsilon_1$ and $X_2 = 2\omega + \epsilon_2$, and define $\mathcal{S} = \{X_1\}$. Then, $\overline{\mathcal{S}} = \{X_1, X_2\}$. Since $\{X_2\}$ permits faster speed of learning than $\{X_1\}$, the set \mathcal{S} is not subspace-optimal.

We introduce one final assumption, which strengthens Unique Minimizer to require the existence of a best minimal spanning set \mathcal{S} within every subspace.

Assumption 3 (Unique Minimizer in Every Subspace). *For every $\mathcal{A} \subset [N]$, there exists a unique minimal spanning set \mathcal{S} that minimizes ϕ among subsets of $\overline{\mathcal{A}}$.*

²¹In Appendix B.2, we further show $m_i(t) = \lambda_i^*(t) + O(1)$ holds in this case. That result implies the difference between society's acquisitions $m(t)$ and optimal acquisitions $n(t)$ remains bounded as t increases.

²²We point out that the set of coefficient matrices satisfying Unique Minimizer is “generic” in the following stronger sense: Fix the *directions* of coefficient vectors, and suppose that the *precisions* are drawn at random; then, generically different minimal spanning sets correspond to different speed of learning. In contrast, whether every minimal spanning set has size K is a condition on the *directions* themselves.

This assumption is guaranteed if different minimal spanning sets correspond to different ϕ -values, and thus holds generically.

Our next result generalizes both the learning traps result and also the efficient information aggregation result from the previous sections. Theorem 2 says that long-run information acquisitions eventually concentrate on a set \mathcal{S} (starting from some prior belief) *if and only if* \mathcal{S} is a subspace-optimal minimal spanning set.

Theorem 2. (a) Suppose \mathcal{S} is a subspace-optimal minimal spanning set. Then, there exists an open set of prior beliefs given which agents exclusively observe signals from \mathcal{S} . Long-run frequencies are positive precisely for those signals in \mathcal{S} , and they are given by Proposition 2.

(b) Under Assumption 3, long-run frequencies exist given any prior belief. Moreover, if \mathcal{S} denotes the signals viewed with positive long-run frequency, then \mathcal{S} is a minimal spanning set that is subspace-optimal, and the long-run frequencies are given by Proposition 2.

This theorem directly implies our previous Corollaries 2 and 3. To see this, observe that under Strong Linear Independence, $\bar{\mathcal{S}} = \mathcal{S}$ for every minimal spanning set \mathcal{S} with fewer than K signals.²³ This implies that every minimal spanning set with fewer than K signals is (trivially) optimal in its subspace, producing Corollary 2 from part (a) of the theorem.

On the other hand, if every minimal spanning set has size K , then all minimal spanning sets belong to the same subspace. Under Unique Minimizer, there can only be one minimal spanning set that is optimal in this subspace, and this must also be the best set overall (in the sense of Section 4). This yields Corollary 3 from part (b) of the theorem above.

We collect below a few additional implications of Theorem 2:

Corollary 4. Under Unique Minimizer, if the best set \mathcal{S}^* is of size 1 (equivalently, if \mathcal{S}^* consists of a single unbiased signal $\alpha\omega + \epsilon$), then learning traps cannot emerge.²⁴

Corollary 5. Under Unique Minimizer, learning traps of size K cannot emerge.²⁵

6 Intuitions for Main Results

We provide intuitions for Corollaries 2-3 and Theorem 2 together in the sections below.

²³Suppose $|\mathcal{S}| < K$, then by assumption of Strong Linear Independence, every signal not in \mathcal{S} is linearly independent from the signals in \mathcal{S} . Hence $\bar{\mathcal{S}}$ cannot contain any other signal.

²⁴This is because the unbiased signal belongs to every subspace spanned by a minimal spanning set.

²⁵If \mathcal{S} is a learning trap of size K , then it spans the whole space. But part (b) of Theorem 2 shows \mathcal{S} must be the best set, leading to a contradiction.

6.1 High-Level Argument

Agents choose signals by comparing the marginal value of different observations. Thus, signal acquisitions eventually concentrate on a set \mathcal{S} if and only if the marginal values of signals in that set become persistently higher than those of other signals.

In settings of *exactly sufficient information*, in which agents must observe all available signals in order to learn ω , it can be shown that agents will eventually observe all signals and learn all states (including all of the confounding terms). Thus, agents will come to evaluate signals by a prior-independent “asymptotic” valuation, which also allows them to identify the best set of signals and approximate the optimal frequency.

When information is overabundant, agents can learn ω from many different (proper) subsets of signals, and there is no guarantee that agents will observe signals in the best set at all. This complicates the analysis, since the marginal value of any given signal depends critically on which signals have been observed previously. It is exactly this difference that leads to our learning traps result (Corollary 2): Observation of different minimal spanning sets in the long run can be sustained by prior beliefs (and resulting posterior beliefs) that overvalue the signals within the set relative to signals outside of the set.

However, we show that our previous argument for the exactly sufficient information case holds “subspace by subspace.” That is, as agents repeatedly acquire signals from any fixed subspace of signals, they will eventually discover the asymptotic marginal values of each signal *in that subspace*. In the long run, agents identify and choose from the best set of signals within that subspace. Thus, only those sets of signals that are best in their subspace are potentially “self-sustaining.” And if all sets of signals that reveal ω span the entire space, agents will identify the best set of signals overall and achieve efficient information aggregation.

6.2 Gradient Descent Dynamics

To provide further intuition, we introduce the following “normalized” asymptotic posterior variance function V^* (which takes as input frequency vectors $\lambda \in \Delta^{N-1}$):

$$V^*(\lambda) = \lim_{t \rightarrow \infty} t \cdot V(\lambda t).$$

We establish the following relationships between V^* and V . First, signal acquisitions chosen according to a frequency vector that minimizes V^* will asymptotically also minimize the posterior variance function V (Lemma 3); this justifies our interest in V^* . Second, $V^*(\lambda)$ is convex in λ and its unique minimum is the optimal frequency vector λ^* (Lemma 5). So the question of whether efficient information aggregation obtains is equivalent to the question of whether the frequency vector $m(t)/t$ comes to minimize V^* .

By construction, each agent chooses the signal that minimizes the value of V . Under a certain condition, this is also the signal that roughly achieves the greatest reduction in V^* . Thus, we can think of society's frequency vector $\lambda(t) := m(t)/t$ as evolving according to a gradient descent dynamic: At each (large) t , $\lambda(t)$ moves in a direction that minimizes V^* . Since V^* is a convex function, we might expect gradient descent to be well-behaved, eventually converging to the global minimizer λ^* .

While the above argument appears to be in conflict with our learning traps result, our problem has the special feature that descent can only occur along N directions (indexed by the available signals).²⁶ This limitation is without loss whenever V^* is differentiable, since all directional derivatives can then be rewritten as convex combinations of the partial derivatives along basis vectors. The function V^* , however, is *not* differentiable everywhere. Consider our learning trap example with signals

$$\begin{aligned} X_1 &= \omega + \epsilon_1 \\ X_2 &= 3\omega + b_1 + \epsilon_2 \\ X_3 &= b_1 + \epsilon_3 \end{aligned}$$

and set the frequency vector to be $\lambda = (1, 0, 0)$. It is easy to verify that beliefs are made less precise if we re-assign weight from X_1 to X_2 , or from X_1 to X_3 . But beliefs are made more precise if we simultaneously re-assign weight from X_1 to *both* X_2 and X_3 . This means that the derivative of V^* in either direction $(-1, 1, 0)$ or $(-1, 0, 1)$ is positive, while its derivative in the direction $(-1, \frac{1}{2}, \frac{1}{2})$ is in fact negative. Hence, V^* is not differentiable at λ .

Our constrained version of gradient descent can become stuck at vectors λ such as this, so that agents repeatedly sustain the frequency vector λ instead of moving to another frequency vector with smaller V^* . This is reflected in our learning traps result (Corollary 2 and part (a) of Theorem 2). On the other hand, it can be shown that V^* is differentiable at λ whenever λ places nonzero weight on a spanning set of K signals. This explains why efficient information aggregation occurs under the assumption of Corollary 3, where the need to recover ω requires society to learn all the states (and hence to observe a set of signals that spans all of \mathbb{R}^K). Repeating this argument for each subspace yields part (b) of Theorem 2.

Remark 1. *The above intuition connects to a literature on learning convergence in potential games (Monderer and Shapley, 1996; Sandholm, 2010). Define an N -player game where each player i chooses a number $\lambda_i \in \mathbb{R}_+$ and receives payoff $-\left(\sum_{j=1}^N \lambda_j\right) \cdot V^*(\lambda) = -V^*\left(\lambda / \sum_{j=1}^N \lambda_j\right)$.²⁷ Then, we have a potential game with (exact) potential function $-V^*$,*

²⁶This constraint corresponds to our assumption that each agent acquires a single, discrete, observation of a chosen signal.

²⁷Thus players' actions are normalized to a frequency vector $\lambda / \sum_{j=1}^N \lambda_j$.

and our long-run observation sets correspond to equilibria of this game. In finite potential games and infinite games with a differentiable potential function, pure-strategy Nash equilibria can be related to the extreme points of the potential function. However, our game described above is an infinite potential game with a non-differentiable potential function. It is known that Nash equilibria in such games need not occur at extreme points, and this is consistent with our result on learning traps.

We note that the connection to potential games is not sufficient to derive our main results, since our agents receive payoff $-V(\lambda t)$ rather than its asymptotic variant V^* . This difference is nontrivial, because whether we can substitute analysis of V^* for analysis of V depends on the prior belief as well as the endogenous path of signal acquisitions.²⁸

7 Interventions

Section 5 demonstrated the possibility for sequential information acquisition to result in inefficient learning. We ask now whether it is possible for a policymaker to influence agents towards efficient learning. Naturally, this question applies only when agents would otherwise potentially achieve a suboptimal speed of learning (with conditions given in part (a) of Theorem 2).

We compare several possible policy interventions: Increasing the *quality* of information acquisition (so that each signal acquisition is more informative); restructuring incentives so that agents' payoffs are based on information obtained over several periods (equivalent to acquisition of *multiple signals* each period); and providing a one-shot release of *free information*, which can then guide subsequent acquisitions. These interventions have different levels of effectiveness, as we explain below.

7.1 More Precise Information

Consider first an intervention in which the precision of signal draws is increased uniformly over signals. For example, if different signals correspond to measurement of different neurochemicals in a group of lab subjects, a government agency can provide researchers with funding that permits recruitment of more subjects. This improves the quality of the estimate regardless of which neurochemical the researcher chooses to measure.

We model this intervention by supposing that each signal acquisition produces B independent observations from that source (with the main model corresponding to $B = 1$). Our result below, which follows from part (a) of Theorem 2, says that providing more informative

²⁸Note that the prior belief—which is outside of the description of the asymptotic potential game—influences which outcome agents will converge to. This is also a feature special to our setting.

signals is of limited effectiveness: Any set of signals that is a potential learning trap given $B = 1$ remains a potential learning trap under arbitrary improvements to signal precision.

Corollary 6. *Suppose that for $B = 1$, there is a set of priors given which signals in \mathcal{S} are (exclusively) viewed in the long run. Then, for every $B \in \mathbb{Z}_+$, there is a set of priors given which \mathcal{S} is exclusively viewed in the long run.*

However, the sets of prior beliefs corresponding to different values of B need not be the same. For a *fixed* prior belief, subsidizing higher quality acquisitions may or may not move the community out of a learning trap. To see this, consider first the informational environment and prior belief from Example 2. Increasing the precision of signals is ineffective here: The first agent chooses X_1 regardless of the value of B , and our previous logic again implies that each subsequent agent also chooses signal X_1 . Thus, the set $\{X_1\}$ remains a learning trap. In Appendix B.4, we provide a contrasting example in which increasing the precision of signals can indeed break agents out of a learning trap from a specified prior belief. Which of these examples is relevant depends on fine details of the informational environment as well as the prior, which the policymaker may not know in practice.

7.2 Batches of Signals

Another possibility is to restructure the incentive scheme so that agents' payoffs are based on information obtained over several periods, equivalent to acquisition of a batch of signals each period. For example, evaluation of researchers can be based on a set of papers, or researchers can be given labs and permitted to direct the work of multiple individuals simultaneously.

Both of these approaches for restructuring the environment can be modeled as permitting each agent to allocate B observations across the sources (where $B = 1$ returns the main model). Note the key difference from the previous intervention: Here it is possible for the B observations to be allocated across *different* signals. We show that it is possible to guarantee efficient information aggregation in this case:

Proposition 3. *Under Unique Minimizer, there is a B such that given acquisition of B signals every period, long-run frequency is λ^* starting from every prior belief.*

Thus, given sufficiently many observations each period, agents will allocate observations in a way that eventually approximates the optimal frequency.

The number of observations needed, however, depends on subtle details of the informational environment. In particular, the required B cannot be uniformly bounded over all environments of fixed size (number of states K and number of signals N). The required B instead depends on two properties: First, it depends on how well the optimal frequency λ^*

can be approximated via allocation of B observations.²⁹ Second, it depends on the difference in learning speed between the best set and the next best minimal spanning set; this difference determines the “slack” that is permitted in the approximation of λ^* . Thus, small batch sizes B are sufficient when the optimal frequency λ^* can be well-approximated using a small number of observations, or when there are large efficiency gains from observing the best set. See Appendix A.5 for further details.

7.3 Free Information

Finally, we consider provision of free information to the community. We can think of this as releasing information that a policymaker knows, or as a reduced form for funding specific kinds of research, the results of which are made public.

Formally, the policymaker chooses M signals $\langle p_j, \theta \rangle + \mathcal{N}(0, 1)$, where each $\|p_j\|_2 \leq \gamma$, so that signal precisions are bounded by γ^2 . At time $t = 0$, independent realizations of these signals are made public. All subsequent agents update their prior beliefs based on this free information in addition to the history of signal acquisitions thus far.

We show that given a sufficient number of (kinds of) signals, of sufficiently high precision, efficient learning can be guaranteed. Specifically, if $k \leq K$ is the size of the optimal set \mathcal{S}^* , then $k - 1$ precise signals are sufficient to guarantee efficient learning:

Proposition 4. *Let $k := |\mathcal{S}^*|$. Under Unique Minimizer, there exists a $\gamma < \infty$, and $k - 1$ signals with $\|p_j\|_2 \leq \gamma$, such that with these free signals provided at $t = 0$, society’s long-run frequency is λ^* starting from every prior belief.*

The proof is by construction. We show that as long as agents understand those confounding terms that appear in the best set of signals (these parameters have dimension $k - 1$), they will come to evaluate the signals in the best set according to their asymptotic marginal values.³⁰

This intervention is most relevant in settings in which a technological advance could greatly speed up progress, but development of the technology is slow and tedious. For example, suppose that high-resolution brain scans would allow for rapid understanding of depression, but the current imaging technology is very poor. Researchers working to understand depression may prefer to exploit existing technologies, rather than contribute to development of this new technology. The government can intervene by funding preliminary

²⁹For example, $\lambda^* = (1/2, 1/2)$ can be achieved exactly using two observations, while $\lambda^* = (3/8, 5/8)$ cannot.

³⁰This intervention requires knowledge of the full correlation structure, and also which set \mathcal{S}^* is best. An alternative intervention, with higher demands on information provision but lower demands on knowledge of the environment, is to provide $K - 1$ (sufficiently precise) signals about all of the confounding terms.

development of brain imaging, which then encourages researchers to begin using brain scans. Once low-resolution brain scans are common, the payoff to advancing the imaging technology increases, and even short-sighted researchers may contribute to this agenda. In this way, provision of free information can nudge agents onto the right path of learning.

8 Extensions

8.1 General Payoff Functions

All of our main results extend when each agent t chooses an action to maximize an arbitrary individual payoff function $u_t(a_t, \omega)$ (recall that previously we restricted to $u_t(a_t, \omega) = -(a_t - \omega)^2$). We require only that these payoff functions are nontrivial in the following sense:

Assumption 4 (Payoff Sensitivity to Mean). *For every t , any variance $\sigma^2 > 0$ and any action $a^* \in A$, there exists a positive Lebesgue measure of μ for which a^* does not maximize $\mathbb{E}[u_t(a, \omega) \mid \omega \sim \mathcal{N}(\mu, \sigma^2)]$.*

That is, for every belief variance, the expected value of ω affects the optimal action to take. This rules out cases with a “dominant” action and ensures that each agent *strictly* prefers to choose the most informative signal. Since the signal that minimizes the posterior variance about ω Blackwell-dominates every other signal,³¹ each agent’s information acquisition remains unchanged.

However, the interpretation of the optimal benchmark (that we defined in Section 4) is more limited. Specifically, while the optimal frequency can still be interpreted as maximizing information revelation, the relationship to the social planner problem (Proposition 1) may fail. A detailed discussion is relegated to Appendix B.5.1.

8.2 Low Altruism

So far we have assumed that agents care only to maximize the accuracy of their own prediction of the payoff-relevant state. Consider a generalization in which agents are slightly altruistic; that is, each agent t chooses a signal as well as an action a_t to maximize discounted payoffs $\mathbb{E}[\sum_{t' \geq t} \delta^{t'-t} \cdot u(a_{t'}, \omega)]$, assuming that future agents will behave similarly. Note that $\delta = 0$ returns our main model. We show in Appendix B.5.2 that for δ sufficiently small, part (a) of Theorem 2 continues to hold (in every equilibrium of this game). So the existence of learning traps is robust to a small degree of altruism.

³¹See, e.g., Hansen and Torgersen (1974).

8.3 Multiple Payoff-Relevant States

In our main model, only one of the K persistent states is payoff-relevant. Consider a generalization in which each agent predicts (the same) $r \leq K$ unknown states and his payoff is determined via a weighted sum of quadratic losses. We show in Appendix B.5.3 that all of our main results extend to this setting. The possibility for agents to have payoffs that depend on *heterogeneous* states is also interesting, and we leave this for future work.

9 Conclusion

We study a model of sequential learning, where short-lived agents choose what kind of information to acquire from a large set of available information sources. Because agents do not internalize the impact of their information acquisitions on later decision-makers, they may acquire information inefficiently (from a social perspective). Inefficiency is not guaranteed, however: Depending on the informational environment, myopic concerns can endogenously push agents to identify and observe only the most informative sources.

Our main results separate these possibilities, and reveal that the extent of *learning spillovers* is essential to determining which outcome emerges. Specifically, does information about unknowns of immediate societal interest (i.e., the payoff-relevant state) also teach about unknowns that are only of indirect value (i.e., the confounding terms)?

When such spillovers are present, simple incentive schemes for information acquisition—in which agents care only about immediate contributions to knowledge—are sufficient for efficient long-run learning. When these spillovers are not built into the environment, other incentives are needed. For example, forward-looking funding agencies can encourage investment in the confounding terms (our “free information” intervention). Alternatively, agents can be evaluated on the basis of a body of work (our “multiple signal” intervention). These observations are consistent with practices that have arisen in academic research, including the establishment of third-party funding agencies (e.g. the NSF) to support basic science and methodological research, and the evaluation of researchers based on advancements developed across several papers (e.g. tenure and various prizes).

We conclude below with brief mention of additional directions. So far we have considered the demand for information given an exogenous set of information sources. In a complementary problem to ours, information sources choose the information they provide in order to maximize demand. Our Theorem 1 implies the following comparative static: If signal i is viewed with positive frequency in the optimal benchmark, then this frequency is (locally) decreasing in its precision. Thus, if demand is interpreted as λ_i^* (the long-run frequency with which source i is optimally viewed), sources face conflicting incentives: They want to provide

information sufficiently precise to be included in the best set and receive viewership at all, but subject to this, they want to provide signals as imprecise as possible. These conflicting forces suggest that characterization of the equilibrium provisions of information precision is subtle.

Finally, while we have described our setting as choice between information sources, our model may apply more generally to choice between actions with complementarities. For example, suppose a sequence of managers take actions that have externalities for future managers, and each manager seeks to maximize performance of the company during his tenure. The concepts we have developed here of *efficient information aggregation* and *learning traps* have natural analogues in that setting (actions that maximize the company’s long-term welfare, versus those that do not). Relative to the general setting, we study here a class of complementarities that are micro-founded in correlated signals. It is an interesting question of whether and how the forces we find here generalize to other sorts of complementarities.

References

- Ali, Nageeb.** 2018. “Herding with Costly Information.” *Journal of Economic Theory*.
- Banerjee, Abhijit.** 1992. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, 107(3): 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch.** 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades.” *Journal of Political Economy*, 100(5): 992–1026.
- Borgers, Tilman, Angel Hernando-Veciana, and Daniel Krahmer.** 2013. “When Are Signals Complements Or Substitutes.” *Journal of Economic Theory*, 148(1): 165–195.
- Burguet, Roberto, and Xavier Vives.** 2000. “Social Learning and Costly Information.” *Economic Theory*.
- Chade, Hector, and Edward E. Schlee.** 2002. “Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information.” *Journal of Economic Theory*, 107: 421–452.
- Chade, Hector, and Jan Eeckhout.** 2018. “Matching Information.” *Theoretical Economics*.
- Chaloner, Kathryn.** 1984. “Optimal Bayesian Experimental Design for Linear Models.” *The Annals of Statistics*, 12(1): 283–300.
- Chen, Yiling, and Bo Waggoner.** 2016. “Informational Substitutes.”
- Che, Yeon-Koo, and Konrad Mierendorff.** 2017. “Optimal Sequential Decision with Limited Attention.” Working Paper.
- Fudenberg, Drew, Philip Strack, and Tomasz Strzalecki.** 2017. “Stochastic Choice and Optimal Sequential Sampling.” Working Paper.

- Golub, Benjamin, and Matthew Jackson.** 2012. “How Homophily Affects the Speed of Learning and Best-Response Dynamics.” *The Quarterly Journal of Economics*.
- Hann-Caruthers, Wade, Vadim Martynov, and Omer Tamuz.** 2017. “The Speed of Sequential Asymptotic Learning.” Working Paper.
- Hansen, Ole Havard, and Eric N. Torgersen.** 1974. “Comparison of Linear Normal Experiments.” *The Annals of Statistics*, 2: 367–373.
- Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz.** 2018. “Group-think and the Failure of Information Aggregation in Large Groups.” Working Paper.
- Liang, Annie, Xiaosheng Mu, and Vasilis Syrgkanis.** 2017. “Optimal and Myopic Information Acquisition.” Working Paper.
- Mayskaya, Tatiana.** 2017. “Dynamic Choice of Information Sources.” Working Paper.
- Monderer, Dov, and Lloyd Shapley.** 1996. “Potential Games.” *Games and Economic Behavior*, 14(1): 124–143.
- Mueller-Frank, Manuel, and Mallesh Pai.** 2016. “Social Learning with Costly Search.” *American Economic Journal: Microeconomics*.
- Sandholm, William.** 2010. *Population Games and Evolutionary Dynamics*. MIT Press.
- Sethi, Rajiv, and Muhamet Yildiz.** 2016. “Communication with Unknown Perspectives.” *Econometrica*, 84(6): 2029–2069.
- Sethi, Rajiv, and Muhamet Yildiz.** 2017. “Culture and Communication.” Working Paper.
- Smith, Lones, and Peter Sorenson.** 2000. “Pathological Outcomes of Observational Learning.” *Econometrica*.
- Vives, Xavier.** 1992. “How Fast do Rational Agents Learn?” *Review of Economic Studies*.

A Appendix

The structure of the appendix follows that of the paper. We provide proofs for the results listed in the main text, in the order in which they appeared; the only exception is that the proof of Proposition 1 relies on tools we develop in the other proofs, and so it is given at the end. Other results and examples are deferred to a separate Online Appendix.

A.1 Posterior Variance Function

A.1.1 Monotonicity and Convexity

Here we review and extend a basic result from [Liang, Mu and Syrgkanis \(2017\)](#). Specifically, we show that the posterior variance about ω weakly decreases over time, and the marginal value of any signal decreases in its signal count.

Lemma 1. *Given prior covariance matrix Σ^0 and q_i observations of each signal i , society's posterior variance about ω is*

$$V(q_1, \dots, q_N) = [((\Sigma^0)^{-1} + C'QC)^{-1}]_{11} \quad (3)$$

where $Q = \text{diag}(q_1, \dots, q_N)$. The function V is decreasing and convex in each q_i whenever these arguments take non-negative real values.

Proof. Note that $(\Sigma^0)^{-1}$ is the prior precision matrix and $C'QC = \sum_{i=1}^N q_i \cdot [c_i c_i']$ is the total precision from the observed signals. Thus (3) simply represents the fact that for Gaussian prior and signals, the posterior precision matrix is the sum of the prior and signal precision matrices. To prove the monotonicity of V , consider the partial order \succeq on positive semi-definite matrices where $A \succeq B$ if and only if $A - B$ is positive semi-definite. As q_i increases, the matrix Q and $C'QC$ increase in this order. Thus the posterior covariance matrix $((\Sigma^0)^{-1} + C'QC)^{-1}$ decreases in this order, which implies that the posterior variance about ω decreases.

To prove that V is convex, it suffices to prove that V is *midpoint-convex* since the function is clearly continuous.³² Take $q_1, \dots, q_N, r_1, \dots, r_N \in \mathbb{R}_+$ and let $s_i = \frac{q_i + r_i}{2}$. Define the corresponding diagonal matrices to be Q, R, S . Observe that $Q + R = 2S$. Thus by the AM-HM inequality for positive-definite matrices, we have

$$((\Sigma^0)^{-1} + C'QC)^{-1} + ((\Sigma^0)^{-1} + C'RC)^{-1} \succeq 2((\Sigma^0)^{-1} + C'SC)^{-1}.$$

Using (3), we conclude that

$$V(q_1, \dots, q_N) + V(r_1, \dots, r_N) \geq 2V(s_1, \dots, s_N).$$

This proves the (midpoint) convexity of V . □

A.1.2 Inverse of Positive Semi-definite Matrices

For future use, we provide a definition of $[X^{-1}]_{11}$ for positive *semi-definite* matrices X . When X is positive definite, its eigenvalues are strictly positive, and its inverse matrix is defined as usual. In general, we can apply the Spectral Theorem to write

$$X = UDU',$$

where U is a $K \times K$ orthogonal matrix whose columns are eigenvectors of X , and $D = \text{diag}(d_1, \dots, d_K)$ is a diagonal matrix consisting of non-negative eigenvalues. Even if some of

³²A function V is midpoint-convex if the inequality $V(a) + V(b) \geq 2V(\frac{a+b}{2})$ always holds. Every continuous function that is midpoint-convex is also convex.

these eigenvalues are zero, we can think of X^{-1} as

$$X^{-1} = (UDU')^{-1} = UD^{-1}U' = \sum_{j=1}^K \frac{1}{d_j} \cdot [u_j u_j']$$

where u_j is the j -th column vector of U . We thus define

$$[X^{-1}]_{11} := \sum_{j=1}^K \frac{(\langle u_j, e_1 \rangle)^2}{d_j}, \quad (4)$$

with the convention that $\frac{0}{0} = 0$ and $\frac{z}{0} = \infty$ for any $z > 0$. Note that by this definition,

$$[X^{-1}]_{11} = \lim_{\epsilon \rightarrow 0+} \left(\sum_{j=1}^K \frac{(\langle u_j, e_1 \rangle)^2}{d_j + \epsilon} \right) = [(X + \epsilon I_K)^{-1}]_{11},$$

since the matrix $X + \epsilon I_K$ has the same set of eigenvectors as X , with eigenvalues increased by ϵ . Hence our definition of $[X^{-1}]_{11}$ is a continuous extension of the usual definition to positive semi-definite matrices.

A.2 Proof of Theorem 1

A.2.1 Asymptotic Posterior Variance Function

We first approximate the posterior variance as a function of the frequencies with which each signal is observed. Specifically, for any $\lambda \in \mathbb{R}_+^N$, define

$$V^*(\lambda) := \lim_{t \rightarrow \infty} t \cdot V(\lambda t).$$

The following result shows V^* to be well-defined and computes its value:

Lemma 2. *Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$. Then*

$$V^*(\lambda) = [(C' \Lambda C)^{-1}]_{11} \quad (5)$$

The value of $[(C' \Lambda C)^{-1}]_{11}$ is well-defined, see (4).

Proof. Recall that $V(q_1, \dots, q_N) = [((\Sigma^0)^{-1} + C' Q C)^{-1}]_{11}$ with $Q = \text{diag}(q_1, \dots, q_N)$. Thus

$$t \cdot V(\lambda_1 t, \dots, \lambda_N t) = \left[\left(\frac{1}{t} (\Sigma^0)^{-1} + C' \Lambda C \right)^{-1} \right]_{11}.$$

Hence the lemma follows from the continuity of $[X^{-1}]_{11}$ in the matrix X . \square

We note that $C' \Lambda C$ is the Fisher Information Matrix when signals are observed according to frequencies λ . Thus the above lemma can also be seen as an application of the Bayesian Central Limit Theorem.

A.2.2 Reduction to the Study of V^*

The development of the function V^* is useful for the following reason:

Lemma 3. *Suppose $\tilde{\lambda}$ uniquely minimizes $V^*(\cdot)$ subject to $\lambda \in \Delta^{N-1}$ (the $(N-1)$ -dimensional simplex). Then, the t -optimal divisions satisfy $n_i(t) \sim \tilde{\lambda}_i \cdot t$ for all i .*

Proof. Fix any increasing sequence of times t_1, t_2, \dots . It suffices to show that whenever the limit $\lambda_i := \lim_{m \rightarrow \infty} \frac{n_i(t_m)}{t_m}$ exists for each i , this limit λ must be $\tilde{\lambda}$. Suppose not, then by assumption $V^*(\lambda) > V^*(\tilde{\lambda})$. For $\epsilon > 0$, define another vector $\hat{\lambda} \in \mathbb{R}_+^N$ with $\hat{\lambda}_i = \lambda_i + \epsilon, \forall i$. By the continuity of V^* , it holds that $V^*(\hat{\lambda}) > V^*(\tilde{\lambda})$ for sufficiently small ϵ .

Since $\lambda_i = \lim_{m \rightarrow \infty} \frac{n_i(t_m)}{t_m}$, there exists M sufficiently large such that $n_i(t_m) \leq (\lambda_i + \epsilon) \cdot t_m$ for each i and $m \geq M$. Hence, for $m \geq M$,

$$t_m \cdot V(n_1(t_m), \dots, n_N(t_m)) \geq t_m \cdot V(\hat{\lambda}_1 \cdot t_m, \dots, \hat{\lambda}_N \cdot t_m) \rightarrow V^*(\hat{\lambda})$$

where the inequality uses the monotonicity of V . On the other hand,

$$t_m \cdot V(\tilde{\lambda}_1 \cdot t_m, \dots, \tilde{\lambda}_N \cdot t_m) \rightarrow V^*(\tilde{\lambda}).$$

Comparing the above two displays, we see that for sufficiently large m ,

$$V(n_1(t_m), \dots, n_N(t_m)) > V(\tilde{\lambda}_1 \cdot t_m, \dots, \tilde{\lambda}_N \cdot t_m).$$

But this contradicts the t -optimality of the division $n(t_m)$, as society could do better by following frequencies $\tilde{\lambda}$. The lemma is thus proved. \square

A.2.3 Crucial Lemma

We pause to demonstrate the following technical lemma:

Lemma 4. *Suppose $\mathcal{S}^* = \{1, \dots, K\}$ uniquely minimizes $\phi(\cdot)$ and let C^* be the $K \times K$ submatrix of C corresponding to the first K signals. Further suppose $\beta_j^{\mathcal{S}^*} = [(C^*)^{-1}]_{1j}$ is positive for $1 \leq j \leq K$. Then for any signal $i > K$, if we write $c_i = \sum_{j=1}^K \alpha_j \cdot c_j$ (which is a unique representation), then $|\sum_{j=1}^K \alpha_j| < 1$.*

Proof. By assumption, we have the vector identity

$$e_1 = \sum_{j=1}^K \beta_j \cdot c_j \quad \text{with } \beta_j = [(C^*)^{-1}]_{1j} > 0.$$

Suppose for contradiction that $\sum_{j=1}^K \alpha_j \geq 1$ (the opposite case where the sum is ≤ -1 can be similarly treated). Then some α_j must be positive. Without loss of generality, we assume $\frac{\alpha_1}{\beta_1}$ is the largest among such ratios. Then $\alpha_1 > 0$ and

$$e_1 = \sum_{j=1}^K \beta_j \cdot c_j = \left(\sum_{j=2}^K \left(\beta_j - \frac{\beta_1}{\alpha_1} \cdot \alpha_j \right) \cdot c_j \right) + \frac{\beta_1}{\alpha_1} \cdot \left(\sum_{j=1}^K \alpha_j \cdot c_j \right)$$

This represents e_1 as a linear combination of the vectors c_2, \dots, c_K and c_i , with coefficients $\beta_2 - \frac{\beta_1}{\alpha_1} \cdot \alpha_2, \dots, \beta_K - \frac{\beta_1}{\alpha_1} \cdot \alpha_K$ and $\frac{\beta_1}{\alpha_1}$. Observe that these coefficients are non-negative: For each $2 \leq j \leq K$, $\beta_j - \frac{\beta_1}{\alpha_1} \cdot \alpha_j$ is clearly positive if $\alpha_j \leq 0$ (since $\beta_j > 0$). And if $\alpha_j > 0$, then by assumption $\frac{\alpha_j}{\beta_j} \leq \frac{\alpha_1}{\beta_1}$ and $\beta_j - \frac{\beta_1}{\alpha_1} \cdot \alpha_j$ is again non-negative.

By definition, $\phi(\{2, \dots, K, i\})$ is the sum of the absolute value of these coefficients. This sum is

$$\sum_{j=2}^K \left(\beta_j - \frac{\beta_1}{\alpha_1} \cdot \alpha_j \right) + \frac{\beta_1}{\alpha_1} = \sum_{j=1}^K \beta_j + \frac{\beta_1}{\alpha_1} \cdot \left(1 - \sum_{j=1}^K \alpha_j \right) \leq \sum_{j=1}^K \beta_j.$$

But then $\phi(\{2, \dots, K, i\}) \leq \phi(\{1, 2, \dots, K\})$, leading to a contradiction. Hence the lemma must be true. \square

A.2.4 Proof of Theorem 1 when $|\mathcal{S}^*| = K$

Given Lemma 3, Theorem 1 will follow once we show that λ^* uniquely minimizes $V^*(\cdot)$ over the simplex; recall that λ^* denotes the optimal frequencies for the minimal spanning set \mathcal{S}^* that minimizes ϕ . In this section, we prove that λ^* is indeed the unique minimizer whenever this best subset \mathcal{S}^* contains exactly K signals. Later on we will prove the same result even when $|\mathcal{S}^*| < K$, but that proof will require additional techniques.

Lemma 5. *For $\lambda \in \Delta^{N-1}$, the function $V^*(\lambda)$ is uniquely minimized at $\lambda = \lambda^*$.*

Proof. First, we assume that $[(C^*)^{-1}]_{1i}$ is positive for $1 \leq i \leq K$. This is without loss because we can always work with the “negative” of any signal (replace c_i with $-c_i$), which does not affect agents’ behavior.

Since $V(q_1, \dots, q_N)$ is convex in its arguments, $V^*(\lambda) = \lim_{t \rightarrow \infty} t \cdot V(\lambda_1 t, \dots, \lambda_N t)$ is also convex in λ . To show λ^* uniquely minimizes V^* , we only need to show λ^* is a *local minimum*. In other words, it suffices to show $V^*(\lambda^*) < V^*(\lambda)$ for any λ that belongs to an ϵ -neighborhood of λ^* . By assumption, \mathcal{S}^* is minimally-spanning and so its signals are linearly independent. Thus its signals must span all of the K states. From this it follows that the $K \times K$ matrix $C' \Lambda^* C$ is positive definite, and by (5) the function V^* is differentiable near λ^* (see Remark 2 below).

We claim that the partial derivatives of V^* satisfy the following inequality:

$$\partial_K V^*(\lambda^*) < \partial_i V^*(\lambda^*) \leq 0, \forall i > K. \quad (*)$$

Once this is proved, we will have, for λ close to λ^* ,

$$V^*(\lambda_1, \dots, \lambda_K, \lambda_{K+1}, \dots, \lambda_N) \geq V^* \left(\lambda_1, \dots, \lambda_{K-1}, \sum_{k=K}^N \lambda_k, 0, \dots, 0 \right) \geq V^*(\lambda^*). \quad (6)$$

The first inequality is based on (*) and differentiability of V^* , while the second inequality is because λ^* uniquely minimizes V^* when only the first K signals are observed. Moreover, when $\lambda \neq \lambda^*$, one of these inequalities is strict so that $V^*(\lambda) > V^*(\lambda^*)$ holds strictly.

To prove (*), we recall that

$$V^*(\lambda) = e_1'(C'\Lambda C)^{-1}e_1.$$

Since $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$, its derivative is $\partial_i \Lambda = \Delta_{ii}$, which is an $N \times N$ matrix whose (i, i) -th entry is 1 with all other entries equal to zero. Using properties of matrix derivatives, we obtain

$$\partial_i V^*(\lambda) = -e_1'(C'\Lambda C)^{-1}C'\Delta_{ii}C(C'\Lambda C)^{-1}e_1.$$

As the i -th row vector of C is c_i' , $C'\Delta_{ii}C$ is the $K \times K$ matrix $c_i c_i'$. The above simplifies to

$$\partial_i V^*(\lambda) = -[e_1'(C'\Lambda C)^{-1}c_i]^2.$$

At $\lambda = \lambda^*$, the matrix $C'\Lambda C$ further simplifies to $(C^*)' \cdot \text{diag}(\lambda_1^*, \dots, \lambda_K^*) \cdot (C^*)$, which is a product of $K \times K$ invertible matrices. We thus deduce that

$$\partial_i V^*(\lambda^*) = -\left[e_1' \cdot (C^*)^{-1} \cdot \text{diag}\left(\frac{1}{\lambda_1^*}, \dots, \frac{1}{\lambda_K^*}\right) \cdot ((C^*)')^{-1} \cdot c_i\right]^2.$$

Crucially, note that the term in the brackets is a linear function of c_i . To ease notation, we write $v' = e_1' \cdot (C^*)^{-1} \cdot \text{diag}\left(\frac{1}{\lambda_1^*}, \dots, \frac{1}{\lambda_K^*}\right) \cdot ((C^*)')^{-1}$ and $\gamma_i = \langle v, c_i \rangle$. Then

$$\partial_i V^*(\lambda^*) = -\gamma_i^2, \quad 1 \leq i \leq N. \quad (7)$$

For $1 \leq i \leq K$, $((C^*)')^{-1} \cdot c_i$ is just e_i . Thus, using the assumption $[(C^*)^{-1}]_{1j} > 0, \forall j$, we have

$$\gamma_i = e_1' \cdot (C^*)^{-1} \cdot \text{diag}\left(\frac{1}{\lambda_1^*}, \dots, \frac{1}{\lambda_K^*}\right) \cdot e_i = \frac{[(C^*)^{-1}]_{1i}}{\lambda_i^*} = \sum_{j=1}^K [(C^*)^{-1}]_{1j} = \phi(\mathcal{S}^*) \quad (8)$$

On the other hand, choosing any $i > K$, we can uniquely write the vector c_i as a linear combination of c_1, \dots, c_K . By Lemma 4, for any $i > K$ we have

$$\gamma_i = \langle v, c_i \rangle = \sum_{j=1}^K \alpha_j \cdot \langle v, c_j \rangle = \sum_{j=1}^K \alpha_j \cdot \gamma_j = \phi(\mathcal{S}^*) \cdot \sum_{j=1}^K \alpha_j, \quad (9)$$

where the last equality uses (8). Since $|\sum_{j=1}^K \alpha_j| < 1$, the absolute value of γ_i for any $i > K$ is strictly smaller than the absolute value of γ_K . This together with (7) proves the desired inequality (*), and Lemma 5 follows. \square

Remark 2. The essence of this proof is the following nontrivial property: The subset $\{1, \dots, K\}$ uniquely minimizes ϕ among all subsets of size K if and only if

$$\phi(\{1, \dots, K\}) < \phi(\{1, \dots, K\} \cup \{i\} \setminus \{j\}), \quad \forall 1 \leq j \leq K < i \leq N.$$

That is, if a set of K signals does not minimize ϕ , then we can improve the speed of learning by adding one signal to replace one existing signal. This property enables us to reduce the general problem with N signals to a much simpler problem with $K + 1$ signals. We are then able to use calculus to resolve the latter problem, see (*).

This argument breaks down if we start with a set of less than K signals; see Section 6 of the main text for an example in which the reduction above is not possible. In that example, even though the partial derivatives satisfy (*), it does not hold that every directional derivative similarly satisfies (*). Thus V^* can fail to be differentiable at the frequency vector of interest. It is for this reason that we need a different proof of Lemma 5 when $|\mathcal{S}^*| < K$, which we present next.

A.2.5 A Perturbation Argument

To handle the case of $|\mathcal{S}^*| < K$, we first extend the definition of $\phi(\cdot)$ to arbitrary sets of signals (not necessarily minimally-spanning) as follows. For any set \mathcal{A} that contains a minimal spanning set, define $\phi(\mathcal{A}) = \min_{\mathcal{S} \subset \mathcal{A}} \phi(\mathcal{S})$, where the minimum is taken over all minimal spanning sets \mathcal{S} contained in \mathcal{A} . If such \mathcal{S} does not exist (i.e., \mathcal{A} is not itself spanning), we let $\phi(\mathcal{A}) = \infty$. In particular,

$$\phi([N]) = \min_{\mathcal{S} \subset [N]} \phi(\mathcal{S})$$

represents the minimum asymptotic standard deviation achievable by only observing the signals in *some* minimal spanning set.

Our previous analysis shows that whenever $\phi(\mathcal{S})$ is uniquely minimized by a set \mathcal{S} containing exactly K signals,

$$\min_{\lambda \in \Delta^{N-1}} V^*(\lambda) = V^*(\lambda^*) = \min_{\mathcal{S} \subset [N]} \phi(\mathcal{S})^2 = \phi([N])^2$$

We now show this equality holds more generally.

Lemma 6. For any coefficient matrix C ,

$$\min_{\lambda \in \Delta^{N-1}} V^*(\lambda) = \phi([N])^2. \quad (10)$$

Proof. Because society can choose to focus on any minimal spanning set, it is clear that $\min_{\lambda} V^*(\lambda) \leq \phi([N])^2 = \min_{\mathcal{S}} (\phi(\mathcal{S}))^2$. It remains to prove $V^*(\lambda) \geq \phi([N])^2$ for any $\lambda \in \Delta^{N-1}$. By Lemma 2, we need to show $[(C' \Lambda C)^{-1}]_{11} \geq \phi([N])^2$.

This was already proved for *generic* coefficient matrices C (specifically, those for which $\phi(\mathcal{S})$ is minimized by a set of K signals). But even if C is “non-generic”, we can approximate it by a sequence of “generic” matrices C_m .³³ Along this sequence, we have

$$[(C'_m \Lambda C_m)^{-1}]_{11} \geq \phi_m([N])^2$$

where ϕ_m is the speed of learning from the N signals given by coefficient matrix C_m . As $m \rightarrow \infty$, the LHS above approaches $[(C' \Lambda C)^{-1}]_{11}$. Thus the lemma will follow once we show that $\limsup_{m \rightarrow \infty} \phi_m([N]) \geq \phi([N])$.

For this we invoke the following characterization

$$\phi([N]) = \min_{\beta \in \mathbb{R}^N} \sum_{i=1}^N |\beta_i| \quad \text{s.t.} \quad e_1 = \sum_{i=1}^N \beta_i \cdot c_i.$$

If $e_1 = \sum_i \beta_i^{(m)} \cdot c_i^{(m)}$ along the convergent sequence, then $e_1 = \sum_i \beta_i \cdot c_i$ for any limit point β of $\beta^{(m)}$. This enables us to conclude $\liminf_{m \rightarrow \infty} \phi_m([N]) \geq \phi([N])$, which is more than what we need. \square

A.2.6 Proof of Theorem 1 when $|\mathcal{S}^*| < K$

Let $\mathcal{S}^* = \{1, \dots, k\}$. We will now show that even if $k < K$, λ^* is still the unique minimizer of $V^*(\cdot)$. This will imply Theorem 1 via Lemma 3. Since $V^*(\lambda^*) = \phi(\mathcal{S}^*)^2 = \phi([N])^2$ by definition, we know from Lemma 6 that λ^* does minimize V^* . It remains to show that λ^* is the *unique* minimizer.

To do this, we will consider a perturbed informational environment in which signals $k+1, \dots, N$ are made slightly more precise. Specifically, let $\eta > 0$ be a small positive number. Consider an alternative signal coefficient matrix \tilde{C} with $\tilde{c}_i = c_i$ for $i \leq k$ and $\tilde{c}_i = (1+\eta)c_i$ for $i > k$. Let $\tilde{\phi}(\mathcal{S})$ represent the speed of learning function for this alternative problem. Then, it is clear that $\tilde{\phi}(\mathcal{S}^*) = \phi(\mathcal{S}^*)$, while $\tilde{\phi}(\mathcal{S})$ is slightly smaller than $\phi(\mathcal{S})$ for $\mathcal{S} \neq \mathcal{S}^*$. Thus with sufficiently small η , the subset \mathcal{S}^* remains the best set in this perturbed environment, and λ^* remains the optimal frequency vector.

Let \tilde{V}^* be the asymptotic posterior variance function here, then our previous analysis shows that \tilde{V}^* has minimum value $\phi(\mathcal{S}^*)^2$ on the simplex. Taking advantage of the connection

³³First, we may add repetitive signals to ensure $N \geq K$. This does not affect the value of $\min V^*(\lambda)$ or $\phi([N])$. Whenever $N \geq K$, it is generically true that every minimal spanning set contains exactly K signals. Moreover, the equality $\phi(\mathcal{S}) = \phi(\tilde{\mathcal{S}})$ for $\mathcal{S} \neq \tilde{\mathcal{S}}$ induces a non-trivial polynomial equation over the entries in C . This means we can always find $C^{(m)}$ close to C such that for each coefficient matrix $C^{(m)}$, different subsets \mathcal{S} (of size K) attain different values of ϕ .

between V^* and \tilde{V}^* , we thus have

$$\begin{aligned} V^*(\lambda_1, \dots, \lambda_N) &= \tilde{V}^* \left(\lambda_1, \dots, \lambda_k, \frac{\lambda_{k+1}}{(1+\eta)^2}, \dots, \frac{\lambda_N}{(1+\eta)^2} \right) \\ &\geq \frac{\phi(\mathcal{S}^*)^2}{\sum_{i \leq k} \lambda_i + \frac{1}{(1+\eta)^2} \sum_{i > k} \lambda_i}. \end{aligned}$$

The equality uses (5) and $C' \Lambda C = \sum_i \lambda_i c_i c_i' = \sum_{i \leq k} \lambda_i c_i c_i' + \sum_{i > k} \frac{\lambda_i}{(1+\eta)^2} \tilde{c}_i \tilde{c}_i'$. The inequality follows from the homogeneity of \tilde{V}^* .

The above display implies

$$\forall \lambda \in \Delta^{N-1}, \quad V^*(\lambda) \geq \frac{\phi(\mathcal{S}^*)^2}{1 - \frac{2\eta+\eta^2}{(1+\eta)^2} \sum_{i > k} \lambda_i} \geq \frac{\phi(\mathcal{S}^*)^2}{1 - \eta \sum_{i > k} \lambda_i} \quad \text{for some } \eta > 0. \quad (11)$$

Hence $V^*(\lambda) > \phi(\mathcal{S}^*)^2 = V^*(\lambda^*)$ whenever λ puts positive weight outside of the best set. But we already know that $V^*(\lambda)$ is uniquely minimized at λ^* when λ is restricted to the best set. Hence λ^* is the unique minimizer of V^* over the whole simplex. This completes the proof of Theorem 1.

A.2.7 Strengthening of Theorem 1

In this section we show $n_i(t) = \lambda_i^* \cdot t + O(1), \forall i$, which improves upon the conclusion of Theorem 1. Note that for any (q_1, \dots, q_N) to be a t -optimal division, it is necessary that $tV(q_1, \dots, q_N) \leq tV(\lambda^* t)$. A straightforward refinement of Lemma 2 gives that whenever $V^*(\lambda)$ is finite, $t \cdot V(\lambda t)$ approaches $V^*(\lambda)$ at the rate of $\frac{1}{t}$. Thus we must have

$$V^* \left(\frac{q_1}{t}, \dots, \frac{q_N}{t} \right) \leq V^*(\lambda^*) + O \left(\frac{1}{t} \right).$$

Together with (11), this implies $\sum_{i > k} \frac{q_i}{t} = O(\frac{1}{t})$, so that signals outside of the best set are only observed finitely many times. Conditional on *these* signal counts, Proposition 2 shows that society's optimal allocation over the first k signals must satisfy $n_i(t) = \lambda_i^* \cdot t + O(1), \forall 1 \leq i \leq k$. This proves what we want.

A.3 Proof of Theorem 2 Part (a)

Let signals $1, \dots, k$ (with $k \leq K$) be a minimally spanning set that is optimal in its subspace. We will demonstrate an open set of prior beliefs given which *all agents* observe these k signals. Since these signals are minimally spanning, they must be linearly independent. Thus we can consider linearly transformed states $\tilde{\theta}_1, \dots, \tilde{\theta}_K$ such that these k signals are simply $\tilde{\theta}_1, \dots, \tilde{\theta}_k$ plus standard Gaussian noise. This linear transformation is invertible, so any prior over the

original states is bijectively mapped to a prior over the transformed states. Thus it is without loss to work with the transformed model and look for prior beliefs over the transformed states.

The payoff-relevant state ω becomes a linear combination $\lambda_1^* \tilde{\theta}_1 + \dots + \lambda_k^* \tilde{\theta}_k$ (after scaling). Since the first k signals are optimal in their subspace, Lemma 4 implies that any other signal that belongs to this subspace can be written as

$$\sum_{i=1}^k \alpha_i \tilde{\theta}_i + \mathcal{N}(0, 1)$$

with $|\sum_{i=1}^k \alpha_i| < 1$. On the other hand, if a signal does not belong to this subspace, it must take the form of

$$\sum_{i=1}^K \beta_i \tilde{\theta}_i + \mathcal{N}(0, 1)$$

with $\beta_{k+1}, \dots, \beta_K$ not all equal to zero.

Now consider a prior belief such that $\tilde{\theta}_1, \dots, \tilde{\theta}_K$ are *independent* from each other. Given prior variances v_1, \dots, v_K , the reduction in the variance of $\lambda_1^* \tilde{\theta}_1 + \dots + \lambda_k^* \tilde{\theta}_k$ by any signal $\sum_{i=1}^k \alpha_i \tilde{\theta}_i + \mathcal{N}(0, 1)$ is

$$\frac{(\sum_{i=1}^k \alpha_i \lambda_i^* v_i)^2}{1 + \sum_{i=1}^k \alpha_i^2 v_i}$$

If v_1, \dots, v_k are small positive numbers and *if the product $\lambda_i^* v_i$ is approximately constant across $1 \leq i \leq k$* , then the above is approximately $(\sum_{i=1}^k \alpha_i)^2 (\lambda_1^* v_1)^2$. Since $|\sum_{i=1}^k \alpha_i| < 1$, we deduce that any other signal belonging to the subspace of the first k signals is worse than signal 1 (in the first period), whose variance reduction is $\frac{(\lambda_1^* v_1)^2}{1 + v_1}$.

Meanwhile, take any signal that does not belong to the subspace. The variance reduction by such a signal $\sum_{i=1}^K \beta_i \tilde{\theta}_i + \mathcal{N}(0, 1)$ is

$$\frac{(\sum_{i=1}^k \beta_i \lambda_i^* v_i)^2}{1 + \sum_{i=1}^K \beta_i^2 v_i}$$

As $\beta_{k+1}, \dots, \beta_K$ are not all zero, the denominator above is arbitrarily large if v_{k+1}, \dots, v_K are chosen to be large. Then, this signal is again worse than signal 1 for the first agent, similar to the situation in Example 2.

To summarize, we have shown that whenever the prior variances v_1, \dots, v_K satisfy the following three conditions, the first agent chooses among the first k signals:

1. v_1, \dots, v_k are close to 0;
2. $\lambda_1^* v_1, \dots, \lambda_k^* v_k$ have pairwise ratios close to 1;

3. v_{k+1}, \dots, v_K are large.³⁴

To show *every agent* chooses among the first k signals, it suffices to check that starting from any prior belief satisfying the above conditions, the posterior beliefs after observing a signal continue to satisfy these conditions. Since variances decrease over time, the first condition is obviously satisfied. By independence, learning about $\tilde{\theta}_1, \dots, \tilde{\theta}_k$ does not affect the variances of the remaining states. So v_{k+1}, \dots, v_K are unchanged, and the third condition is verified. Finally, the second condition holds by Proposition 2: Since each signal $i \leq k$ is sampled according to λ_i^* , the variance v_i after t periods is approximately $\frac{1}{\lambda_i^* t}$. Hence part (a) of Theorem 2 is proved.³⁵

A.4 Proof of Theorem 2 Part (b)

A.4.1 Preliminary Steps

Given any prior, let $\mathcal{A} \subset [N]$ be the set of all signals that are observed by infinitely many agents. We first show that \mathcal{A} is a spanning set.

Indeed, by definition we can find some period t after which agents only observe signals in \mathcal{A} . Note that the variance reduction of any signal approaches zero as its signal count gets large. Thus, along society's signal path, the variance reduction is close to zero at sufficiently late periods.

If \mathcal{A} is not spanning, society's posterior variance remains bounded away from zero. Thus in the limit where each signal in \mathcal{A} has infinite signal counts, there still exists some signal j outside of \mathcal{A} whose variance reduction is strictly positive.³⁶ By continuity, we deduce that at any sufficiently late period, observing signal j is better than observing any signal in \mathcal{A} . This contradicts our assumption that later agents only observe signals in \mathcal{A} .

Now that \mathcal{A} is spanning, we can take \mathcal{S} to be the optimal minimal spanning set in the subspace spanned by \mathcal{A} . To prove Theorem 2 part (b), we will show that long-run frequencies are positive precisely for the signals in \mathcal{S} . By ignoring the initial periods, we assume without

³⁴Formally, we require that for some fixed constant $\epsilon > 0$, it holds that $v_1, \dots, v_k < \epsilon$; $\max_{1 \leq i \leq k} \lambda_i^* v_i \leq (1 + \epsilon) \cdot \min_{1 \leq i \leq k} \lambda_i^* v_i$; and $v_{k+1}, \dots, v_K > \frac{1}{\epsilon}$.

³⁵Strictly speaking, the above construction does not provide an *open set* of prior beliefs given which agents always observe the first k signals. This is because we restricted attention to priors that are independent over $\tilde{\theta}_1, \dots, \tilde{\theta}_K$. But it can be shown that the argument extends to mild prior correlation across the transformed states. We omit the somewhat cumbersome details, which do not add further intuition.

³⁶To see this, let s_1, \dots, s_N denote the limit signal counts, where $s_i = \infty$ if and only if $i \in \mathcal{A}$. We need to find some signal j such that $V(s_j + 1, s_{-j}) < V(s_j, s_{-j})$. If such a signal does not exist, then all partial derivatives of V at s are zero. Since V is differentiable, this would imply that all directional derivatives of V are also zero. By the convexity of V , V must be minimized at s . However, the minimum value of V is zero because there exists a spanning set. This contradicts $V(s) > 0$.

loss that *only signals in $\overline{\mathcal{A}}$ are available*. It suffices to show that whenever the signals observed infinitely often span a subspace, agents eventually sample from the optimal subset \mathcal{S} in that subspace. To ease notation, we assume this subspace is the entire \mathbb{R}^K , and prove the following result:

Theorem 2 part (b) Restated. *Suppose that the signals observed infinitely often span \mathbb{R}^K . Then society eventually observes signals in \mathcal{S}^* with frequencies λ^* .*

The next sections are devoted to the proof of this restatement.

A.4.2 Estimates of Derivatives

We introduce a few technical lemmata:

Lemma 7. *For any q_1, \dots, q_N , we have*

$$\left| \frac{\partial_{jj} V(q_1, \dots, q_N)}{\partial_j V(q_1, \dots, q_N)} \right| \leq \frac{2}{q_j}.$$

Proof. Recall that $V(q_1, \dots, q_N) = e'_1 \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot e_1$. Thus

$$\partial_j V = -e'_1 \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j \cdot c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot e_1,$$

and

$$\partial_{jj} V = 2e'_1 \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j \cdot c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j \cdot c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot e_1.$$

Let $\gamma_j = e'_1 \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j$, which is a number. Then the above becomes

$$\partial_j f = -\gamma_j^2; \quad \partial_{jj} f = 2\gamma_j^2 \cdot c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j.$$

Note that $(\Sigma^0)^{-1} + C'QC \succeq q_j \cdot c_j c'_j$ in matrix norm. Thus the number $c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j$ is bounded above by $\frac{1}{q_j}$.³⁷ This proves the lemma. \square

Since the second derivative is small compared to the first derivative, we deduce that the variance reduction of any *discrete* signal can be approximated by the partial derivative of f . This property is summarized in the following lemma:

³⁷Formally, we need to show that for any $\epsilon > 0$, the number $c'_j [c_j c'_j + \epsilon I_K]^{-1} c_j$ is at most 1. Using the identity $\text{Trace}(AB) = \text{Trace}(BA)$, we can rewrite this number as

$$\text{Trace}([c_j c'_j + \epsilon I_K]^{-1} c_j c'_j) = \text{Trace}(I_K - [c_j c'_j + \epsilon I_K]^{-1} \epsilon I_K) = K - \epsilon \cdot \text{Trace}([c_j c'_j + \epsilon I_K]^{-1}).$$

The matrix $c_j c'_j$ has rank 1, so $K-1$ of its eigenvalues are zero. Thus the matrix $[c_j c'_j + \epsilon I_K]^{-1}$ has eigenvalue $1/\epsilon$ with multiplicity $K-1$, and the remaining eigenvalue is positive. This implies $\epsilon \cdot \text{Trace}([c_j c'_j + \epsilon I_K]^{-1}) > K-1$, and then the above display yields $c'_j \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j < 1$ as desired.

Lemma 8. For any q_1, \dots, q_N , we have³⁸

$$V(q) - V(q_j + 1, q_{-j}) \geq \frac{q_j}{q_j + 1} |\partial_j V(q)|.$$

Proof. We will show the more general result:

$$V(q) - V(q_j + x, q_{-j}) \geq \frac{q_j x}{q_j + x} \cdot |\partial_j V(q)|, \forall x \geq 0.$$

This clearly holds at $x = 0$. Differentiating with respect to x , we only need to show

$$-\partial_j V(q_j + x, q_{-j}) \geq \frac{q_j^2}{(q_j + x)^2} |\partial_j V(q)|, \forall x \geq 0.$$

Equivalently, we need to show

$$-(q_j + x)^2 \cdot \partial_j V(q_j + x, q_{-j}) \geq -q_j^2 \cdot \partial_j V(q), \forall x \geq 0.$$

Again, this inequality holds at $x = 0$. Differentiating with respect to x , it becomes

$$-2(q_j + x) \cdot \partial_j V(q_j + x, q_{-j}) - (q_j + x)^2 \cdot \partial_{jj} V(q_j + x, q_{-j}) \geq 0.$$

This is exactly the result of Lemma 7. □

A.4.3 Lower Bound on Variance Reduction

Our next result gives a lower bound on the directional derivative of V along the “optimal” direction λ^* :

Lemma 9. For any q_1, \dots, q_N , we have

$$|\partial_{\lambda^*} V(q)| \geq \frac{V(q)^2}{\phi(\mathcal{S}^*)^2}.$$

Proof. To compute this directional derivative, we think of agents acquiring signals in fractional amounts, where a fraction of a signal is just the same signal with precision multiplied by that fraction. Consider an agent who draws λ_i^* realizations of each signal i . Then he essentially obtains the following signals:

$$Y_i = \langle c_i, \theta \rangle + \mathcal{N}\left(0, \frac{1}{\lambda_i^*}\right), \forall i.$$

This is equivalent to

$$\lambda_i^* Y_i = \langle \lambda_i^* c_i, \theta \rangle + \mathcal{N}(0, \lambda_i^*), \forall i.$$

³⁸Note that the convexity of V gives $V(q) - V(q_j + 1, q_{-j}) \leq |\partial_j V(q)|$. This lemma provides a converse that we need for the subsequent analysis.

Such an agent receives at least as much information as the sum of these signals:

$$\sum_i \lambda_i^* Y_i = \sum_i \langle \lambda_i^* c_i, \theta \rangle + \sum_i \mathcal{N}(0, \lambda_i^*) = \frac{\omega}{\phi(\mathcal{S}^*)} + \mathcal{N}(0, 1).$$

Hence the agent's posterior precision about ω (which is the inverse of his posterior variance V) must increase by at least $\frac{1}{\phi(\mathcal{S}^*)^2}$ along the direction λ^* . The chain rule of differentiation yields the lemma. \square

We can now bound the variance reduction at late periods:

Lemma 10. *Fix any q_1, \dots, q_N . Suppose M is a positive real number such that $(\Sigma^0)^{-1} + C'QC \succeq Mc_j c_j'$ holds for each signal $j \in \mathcal{S}^*$. Then we have*

$$\min_{j \in \mathcal{S}^*} V(q_j + 1, q_{-j}) \leq V(q) - \frac{M}{M+1} \cdot \frac{V(q)^2}{\phi(\mathcal{S}^*)^2}.$$

Proof. Fix any signal $j \in \mathcal{S}^*$. Using the condition $(\Sigma^0)^{-1} + C'QC \succeq Mc_j c_j'$, we can deduce the following variant of Lemma 8:³⁹

$$V(q) - V(q_j + 1, q_{-j}) \geq \frac{M}{M+1} |\partial_j V(q)|.$$

Since V is always differentiable, $\partial_{\lambda^*} V(q)$ is a convex combination of the partial derivatives of V .⁴⁰ Thus

$$\max_{j \in \mathcal{S}^*} |\partial_j V(q)| \geq |\partial_{\lambda^*} V(q)|$$

These inequalities together with Lemma 9 complete the proof. \square

A.4.4 Proof of the Restatement of Theorem 2 Part (b)

We will show $V(m(t)) \sim \frac{\phi(\mathcal{S}^*)^2}{t}$, so that society eventually approximates the optimal speed of learning. Since λ^* is the unique minimizer of V^* , this will imply the desired conclusion $m(t)/t \rightarrow \lambda^*$.

To estimate $V(m(t))$, we note that for any fixed M , society's acquisitions $m(t)$ eventually satisfy the condition $(\Sigma^0)^{-1} + C'QC \succeq Mc_j c_j'$. This is due to our assumption that the signals

³⁹Even though we are not guaranteed $q_j \geq M$, we can modify the prior and signal counts such that the precision matrix $(\Sigma^0)^{-1} + C'QC$ is unchanged, and signal j has been observed at least M times. This is possible thanks to the condition $(\Sigma^0)^{-1} + C'QC \succeq Mc_j c_j'$. Then, applying Lemma 8 to this modified problem yields the result here.

⁴⁰While this may be a surprising contrast with V^* , the difference arises because the formula for V always involves a full-rank prior covariance matrix, whereas its asymptotic variant V^* corresponds to a flat prior.

observed infinitely often span \mathbb{R}^K , which implies that $C'QC$ becomes arbitrarily large in matrix norm. Hence, we can apply Lemma 10 to find that

$$V(m(t+1)) \leq V(m(t)) - \frac{M}{M+1} \cdot \frac{V(m(t))^2}{\phi(\mathcal{S}^*)^2}$$

for all $t \geq t_0$, where t_0 depends only on M .

We introduce the auxiliary function $g(t) = V(m(t)) \cdot \frac{M}{(M+1)\phi(\mathcal{S}^*)^2}$. Then the above simplifies to

$$g(t+1) \leq g(t) - g(t)^2.$$

Inverting both sides, we have

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)(1-g(t))} = \frac{1}{g(t)} + \frac{1}{1-g(t)} \geq \frac{1}{g(t)} + 1. \quad (12)$$

This holds for all $t \geq t_0$. Thus by induction, $\frac{1}{g(t)} \geq t - t_0$ and so $g(t) \leq \frac{1}{t-t_0}$. Going back to the posterior variance function V , this implies

$$V(m(t)) \leq \frac{M+1}{M} \cdot \frac{\phi(\mathcal{S}^*)^2}{t-t_0}. \quad (13)$$

Hence, by choosing M sufficiently large in the first place and then considering large t , we find that society's speed of learning is arbitrarily close to the optimal speed $\phi(\mathcal{S}^*)^2$. This completes the proof.

We comment that the above argument leaves open the possibility that some signals outside of \mathcal{S}^* are observed *infinitely often*, yet with *zero long-run frequency*. Similar to Appendix A.2.7, we can show this does not happen. The proof is involved, and we defer it to Appendix B.2.

A.5 Proof of Proposition 3

Given any history of observations, an agent can always allocate his B observations as follows: He draws $\lfloor B \cdot \lambda_i^* \rfloor$ realizations of each signal i , and samples arbitrarily if there is any capacity remaining. Here $\lfloor \cdot \rfloor$ denotes the floor function.

Fix any $\epsilon > 0$. If B is sufficiently large, then the above strategy acquires at least $(1-\epsilon) \cdot B \cdot \lambda_i^*$ observations of each signal i . Adapting the proof of Lemma 9, we see that the agent's posterior precision about ω must increase by $\frac{(1-\epsilon)B}{\phi(\mathcal{S}^*)^2}$ under this strategy. Thus the same must hold for his optimal strategy, so that society's posterior precision at time t is at least $\frac{(1-\epsilon)Bt}{\phi(\mathcal{S}^*)^2}$. This implies that society's speed of learning (per signal) is at least $\frac{\phi(\mathcal{S}^*)^2}{(1-\epsilon)}$, which can be arbitrarily close to the optimal speed $\phi(\mathcal{S}^*)^2$ with appropriate choice of ϵ .

Since λ^* is the unique minimizer of V^* , society's long-run frequencies must be close to λ^* . In particular, with ϵ sufficiently small, we can ensure that each signal in \mathcal{S}^* are observed with positive frequencies. The restated Theorem 2 part (b) extends to the current setting and implies that society eventually samples from \mathcal{S}^* . This yields the proposition.⁴¹

A.6 Proof of Proposition 4

Suppose without loss that the best set is $\{1, \dots, k\}$. By taking a linear transformation, we can further assume each signal i with $1 \leq i \leq k$ only involves ω and the first $k-1$ confounding terms b_1, \dots, b_{k-1} . We claim that whenever $k-1$ sufficiently precise signals are provided about each of these confounding terms, long-run frequencies must converge to λ^* .

Fix any positive real number M . Since the $k-1$ free signals are very precise, it is as if the prior precision matrix satisfies

$$(\Sigma^0)^{-1} \succeq M^2 \sum_{i=2}^k \Delta_{ii}$$

where Δ_{ii} be the $K \times K$ matrix that has one at the (i, i) entry and zero otherwise. Recall also that society eventually learns ω . Thus at some late period t_0 , society's acquisitions must satisfy

$$C'QC \succeq M^2 \Delta_{11}.$$

Adding up the above two displays, we have

$$(\Sigma^0)^{-1} + C'QC \succeq M^2 \sum_{i=1}^k \Delta_{ii} \succeq M c_j c_j', \forall 1 \leq j \leq k.$$

The last inequality uses the fact that each c_j only involves the first k coordinates.

Now this is exactly the condition we need in order to apply Lemma 10: Whether or not the condition is met for $j \notin \mathcal{S}^*$ does not affect the argument. Thus we can follow the proof of the restated Theorem 2 part (b) to deduce (13). That is, for fixed M and corresponding free information, society's long-run speed of learning cannot be slower than $(1 + 1/M) \cdot \phi(\mathcal{S}^*)^2$. This can be made arbitrarily close to the optimal speed, in which case we use Theorem 2 part (b) to conclude that society eventually samples according to λ^* .

⁴¹This proof also suggests that how small ϵ (and how large B) need to be depends on the distance between the optimal speed of learning and the "second-best" speed of learning from any other minimal spanning set. Intuitively, in order to achieve long-run efficient learning, agents need to allocate B observations in the best set to approximate the optimal frequencies. If another set of signals offers a speed of learning that is only slightly worse, we will need B sufficiently large for the approximately optimal frequencies in the best set to beat this other set.

A.7 Proof of Proposition 1

Recall that λ^* uniquely minimizes the function V^* . Thus the proposition is equivalent to the following: Fix $\epsilon > 0$, then for any δ close to 1, $V(d_\delta(t)) \leq \frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t}$ holds for sufficiently large t . That is, we only need to show that as $\delta \rightarrow 1$, the achieved speed of learning is close to the optimal speed.

Suppose for contradiction that $V(d_\delta(t)) > \frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t}$ at some large t . Let $\tau < t$ be the last period with $V(d_\delta(\tau)) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}$. Below we first assume such a period τ exists; later we will show how to modify the proof when it does not. Consider the following deviation:

1. Agents $i \leq \tau$ choose signals according to d_δ (i.e., they do not deviate);
2. Starting in period $\tau + 1$, the next Mk agents sample each signal in the best set (of size k) exactly M times, in an arbitrary order;
3. Starting in period $\tau + Mk + 1$, each future agent chooses the signal that maximizes his own expected payoff, as in our main model.

In what follows we will show that for appropriately chosen M as well as sufficiently large δ and t , this deviation yields a higher δ -discounted payoff than the original strategy d_δ .

By construction, the deviation strategy achieves the same payoff as the original strategy in the first τ periods. Next we consider those periods $\tau + 1$ through t . For $1 \leq j \leq t - \tau$, let $\tilde{V}(\tau + j)$ denote the posterior variance at time $\tau + j$ under the deviation strategy. We can bound it from above as follows: Our previous analysis in Appendix A.6 (specifically (12)) gives that for $j > Mk$,

$$\frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \tilde{V}(\tau + j)} \geq \frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \tilde{V}(\tau + Mk)} + j - Mk, \quad (14)$$

Using $\tilde{V}(\tau + Mk) \leq \tilde{V}(\tau) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}$, the above inequality further yields

$$\frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \tilde{V}(\tau + j)} \geq \frac{(M+1)\tau}{M(1+\epsilon/2)} + j - Mk. \quad (15)$$

With slight algebra, we obtain from the above

$$\frac{\tilde{V}(\tau + j)}{\phi(\mathcal{S}^*)^2} \leq \frac{1}{\frac{\tau}{1+\epsilon/2} + \frac{j-Mk}{1+1/M}}. \quad (16)$$

Fixing ϵ , we now choose M so that $\frac{1}{M} < \frac{\epsilon}{4}$. Then there exists \underline{j} (depending only on ϵ , M and K) such that for $j > \underline{j}$, it holds that

$$\frac{j - Mk}{1 + 1/M} \geq \frac{j + 1}{1 + \epsilon/2}.$$

Thus, (16) implies

$$\tilde{V}(\tau + j) \leq \frac{(1 + \epsilon/2)\phi(\mathcal{S}^*)^2}{\tau + j + 1}, \quad \forall \underline{j} + 1 \leq j \leq t - \tau. \quad (17)$$

On the other hand, for small j we have the following crude estimate:

$$\tilde{V}(\tau + j) \leq \tilde{V}(\tau) \leq \frac{(1 + \epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}, \quad \forall 1 \leq j \leq \underline{j}. \quad (18)$$

Now we go back to the original strategy and make payoff comparisons. Our choice of τ ensures that posterior variance under the *original* strategy is at least $\frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau+j}$, for $1 \leq j \leq t - \tau$. Hence by deviating, the payoff gain in periods $\tau + 1 \sim t$ is at least⁴²

$$\delta^\tau \cdot \underbrace{\left(\sum_{j=\underline{j}+1}^{t-\tau} \delta^{j-1} \left[\frac{(1 + \epsilon/2)\phi(\mathcal{S}^*)^2}{(\tau + j)(\tau + j + 1)} \right] - \sum_{j=1}^{\underline{j}} \delta^{j-1} \left[\frac{(1 + \epsilon/2)\phi(\mathcal{S}^*)^2 j}{\tau(\tau + j)} \right] \right)}_{(**)}.$$

Note that \underline{j} has already been fixed. So as $\delta \rightarrow 1$ and $t - \tau \rightarrow \infty$,⁴³ the term $(**)$ above converges to

$$(1 + \epsilon/2)\phi(\mathcal{S}^*)^2 \cdot \left[\frac{1}{\tau + \underline{j} + 1} - \sum_{j=1}^{\underline{j}} \frac{j}{\tau(\tau + j)} \right].$$

For large τ , the above expression is larger than $\phi(\mathcal{S}^*)^2/\tau$.

Summarizing the above, we have shown that whenever $\tau > \underline{\tau}$, the deviation strategy achieves payoff gain in periods $\tau + 1$ through t of at least $\delta^\tau \phi(\mathcal{S}^*)^2/\tau$ (for δ close to 1). Although the deviation strategy might do worse in periods $t+1$ onwards, the potential payoff loss is at most $O(\frac{\delta^t}{1-\delta})$, which is smaller than the aforementioned payoff gain $\delta^\tau \phi(\mathcal{S}^*)^2/\tau$ as $t - \tau \rightarrow \infty$ (since $\tau \leq \frac{1+\epsilon/2}{1+\epsilon}t$). Hence whenever $\tau > \underline{\tau}$, the deviation we constructed is a profitable deviation, and the proposition holds in these situations.

Finally, we need to address the case where the previously-defined τ is weakly less than $\underline{\tau}$. This covers the case in which τ does not exist according to our earlier definition (simply let $\tau = 0$). Instead of (15), we use the following weaker inequality

$$\frac{(M + 1)\phi(\mathcal{S}^*)^2}{M \cdot \tilde{V}(\tau + j)} \geq j - Mk.$$

⁴²In this derivation we use (17), (18) as well as the identities $\frac{1}{\tau+j} - \frac{1}{\tau+j+1} = \frac{1}{(\tau+j)(\tau+j+1)}$ and $\frac{1}{\tau} - \frac{1}{\tau+j} = \frac{j}{\tau(\tau+j)}$.

⁴³Since $\frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t} \leq V(d_\delta(t)) \leq V(d_\delta(\tau)) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}$, we have $\tau \leq \frac{1+\epsilon/2}{1+\epsilon}t$. So as t becomes large, the difference $t - \tau$ also necessarily becomes large.

That is, $\tilde{V}(\tau + j) \leq \frac{(1+\frac{1}{M})\phi(\mathcal{S}^*)^2}{j-Mk}$. Recall that $\frac{1}{M} < \frac{\epsilon}{4}$ and $\tau \leq \underline{\tau}$ is now bounded. Thus for $j > \bar{j}$ (where \bar{j} may need to be larger than \underline{j}), we would have

$$\tilde{V}(\tau + j) \leq \frac{(1 + \epsilon/4)\phi(\mathcal{S}^*)^2}{\tau + j}, \quad \forall \bar{j} + 1 \leq j \leq t - \tau. \quad (19)$$

And for small j we can simply bound the posterior variance by the prior:

$$\tilde{V}(\tau + j) \leq c, \quad \forall 1 \leq j \leq \bar{j}. \quad (20)$$

Using the estimates (19) and (20) in place of (17) and (18), we find that the deviation strategy achieves payoff gain in periods $\tau + 1$ through t of at least

$$\delta^\tau \cdot \left(\sum_{j=\bar{j}+1}^{t-\tau} \delta^{j-1} \left[\frac{\epsilon/4 \cdot \phi(\mathcal{S}^*)^2}{\tau + j} \right] - \sum_{j=1}^{\bar{j}} \delta^{j-1} c \right).$$

Importantly, because (19) improves upon (17), we now have a harmonic sum in (the first part of) the parenthesis, which becomes arbitrarily large for δ close to 1. Hence the above payoff gain is at least δ^τ as $\delta \rightarrow 1$ and $t \rightarrow \infty$. Once again, this payoff gain dominates any potential loss after period t , showing that the deviation strategy is profitable. The proof of Proposition 1 is complete.

For Online Publication

B Online Appendix

B.1 Example Failing Unique Minimizer

There are $K = 3$ states ω, b_1, b_2 independently drawn with prior variances $\frac{1}{\alpha}, \frac{1}{\beta}, \frac{1}{\gamma}$. $N = 4$ signals are available, and they are respectively

$$X_1 = \omega + b_1 + \epsilon_1$$

$$X_2 = b_1 + \epsilon_2$$

$$X_3 = \omega + b_2 + \epsilon_3$$

$$X_4 = b_2 + \epsilon_4$$

with standard normal errors. Note that the former two signals and the latter two signals are both spanning, and these two sets generate the same asymptotic variance. Thus Assumption 1 is not satisfied.

The posterior variance about ω as a function of the number of observations q_1, q_2, q_3, q_4 of each signal type can be derived as follows. First, given q_2 observations of signal X_2 and q_4 observations of signal X_4 , posterior variance about θ_2 and θ_3 are $1/(q_2 + \beta)$ and $1/(q_4 + \gamma)$ respectively. Consider now q_1 additional observations of X_1 ; this provides the same information about the payoff-relevant state ω as the signal $\omega + \epsilon'$, where ϵ' is an independent noise term with variance $\frac{1}{q_1} + \frac{1}{q_2 + \beta}$. Similarly, q_3 additional observations of X_3 are equivalent to the signal $\omega + \epsilon''$, where ϵ'' is an independent noise term with variance $\frac{1}{q_3} + \frac{1}{q_4 + \gamma}$. From this we deduce that posterior variance about ω is

$$V(q_1, q_2, q_3, q_4) = 1 \left/ \left(\alpha + \frac{1}{\frac{1}{q_1} + \frac{1}{q_2 + \beta}} + \frac{1}{\frac{1}{q_3} + \frac{1}{q_4 + \gamma}} \right) \right.$$

The optimal division vector thus seeks to *maximize*

$$\frac{1}{\frac{1}{q_1} + \frac{1}{q_2 + \beta}} + \frac{1}{\frac{1}{q_3} + \frac{1}{q_4 + \gamma}} \quad (21)$$

It is useful to rewrite (21) in the following way:

$$\frac{1}{4} \left(q_1 + q_2 + \beta + q_3 + q_4 + \gamma - \frac{(q_1 - q_2 - \beta)^2}{q_1 + q_2 + \beta} - \frac{(q_3 - q_4 - \gamma)^2}{q_3 + q_4 + \gamma} \right).$$

Then, since $q_1 + q_2 + \beta + q_3 + q_4 + \gamma = t + \beta + \gamma$ is fixed at any time t , it is equivalent to choose q_1, q_2, q_3, q_4 to minimize the sum of ratios

$$\frac{(q_1 - q_2 - \beta)^2}{q_1 + q_2 + \beta} + \frac{(q_3 - q_4 - \gamma)^2}{q_3 + q_4 + \gamma}.$$

Ideally, if signals were perfectly divisible, the optimum would be to choose $q_1 = q_2 + \beta$ and $q_3 = q_4 + \gamma$. But as each q_i is restricted to integer values, this continuous optimum is not feasible whenever β and γ are not integers.

The solution to this integer optimization problem is involved, and we need some additional notation. Let r be the integer that minimizes $|r - \beta|$ (the distance to β) and let s be the integer that minimizes $|s - \gamma|$. Further, let $\langle \beta \rangle$ and $\langle \gamma \rangle$ be the value of these absolute differences.

Claim 1. *When the period t has the same parity as $r + s$, the t -optimal (q_1, q_2, q_3, q_4) satisfy*

$$q_1, q_2 \approx \frac{\langle \beta \rangle}{2\langle \beta \rangle + 2\langle \gamma \rangle} \cdot t; \quad q_3, q_4 \approx \frac{\langle \gamma \rangle}{2\langle \beta \rangle + 2\langle \gamma \rangle} \cdot t.$$

Otherwise the t -optimal (q_1, q_2, q_3, q_4) satisfy

$$q_1, q_2 \approx \frac{\langle \beta \rangle}{2\langle \beta \rangle + 2 - 2\langle \gamma \rangle} \cdot t; \quad q_3, q_4 \approx \frac{1 - \langle \gamma \rangle}{2\langle \beta \rangle + 2 - 2\langle \gamma \rangle} \cdot t.$$

Thus, all four signals are observed with positive frequency in the long run according to the optimal criterion.

Proof. To solve the integer maximization problem (21), let r be the integer that minimizes $|r - \beta|$ (the distance to β) and let s be the integer that minimizes $|s - \gamma|$. Further, let $\langle \beta \rangle$ and $\langle \gamma \rangle$ be the value of these absolute differences. We assume $2\beta, 2\gamma$ are not integers, so that $0 < \langle \beta \rangle, \langle \gamma \rangle < \frac{1}{2}$. We also assume $\langle \beta \rangle \neq \langle \gamma \rangle$, and by symmetry focus on the case of $\langle \beta \rangle < \langle \gamma \rangle$.

With these assumptions, it is clear that when q_1, q_2 are integers, the minimum value of $|q_1 - q_2 - \beta|$ is $\langle \beta \rangle$, achieved if and only if $q_1 = q_2 + r$. Similarly the minimum value of $|q_3 - q_4 - \gamma|$ is $\langle \gamma \rangle$, achieved when $q_3 = q_4 + s$. Now if the total number of observations t has the *same parity* as $r + s$, it is possible to choose q_1, q_2, q_3, q_4 such that their sum is t and $q_1 = q_2 + r$, $q_3 = q_4 + s$ —any pair q_2, q_4 with sum $\frac{t-r-s}{2}$ leads to such a solution. Given these constraints, then, the optimum is to choose q_2, q_4 to minimize $\frac{\langle \beta \rangle^2}{2q_2 + r + \beta} + \frac{\langle \gamma \rangle^2}{2q_4 + s + \gamma}$. The optimal q_2 and q_4 satisfy $q_2/q_4 \approx \langle \beta \rangle/\langle \gamma \rangle$, which together with $q_2 + q_4 = \frac{t-r-s}{2}$ implies

$$q_1, q_2 \approx \frac{\langle \beta \rangle}{2\langle \beta \rangle + 2\langle \gamma \rangle} \cdot t; \quad q_3, q_4 \approx \frac{\langle \gamma \rangle}{2\langle \beta \rangle + 2\langle \gamma \rangle} \cdot t.$$

On the other hand, suppose t has the *opposite parity* to $r + s$. In this case $q_1 = q_2 + r$ and $q_3 = q_4 + s$ cannot both hold, thus $|q_1 - q_2 - \beta|$ and $|q_3 - q_4 - \gamma|$ cannot both take their minimum values $\langle \beta \rangle$ and $\langle \gamma \rangle$. It turns out that the best one can do is choose $q_1 = q_2 + r$ and

$q_3 = q_4 + s \pm 1$ so that $|q_1 - q_2 - \beta| = \langle \beta \rangle$ and $|q_3 - q_4 - \gamma| = 1 - \langle \gamma \rangle$. Then, the optimal choice of q_2, q_4 with sum $\frac{t-r-s\pm 1}{2}$ to minimize $\frac{\langle \beta \rangle^2}{2q_2+r+\beta} + \frac{(1-\langle \gamma \rangle)^2}{2q_4+s+\gamma\pm 1}$. This yields

$$q_1, q_2 \approx \frac{\langle \beta \rangle}{2\langle \beta \rangle + 2 - 2\langle \gamma \rangle} \cdot t; \quad q_3, q_4 \approx \frac{1 - \langle \gamma \rangle}{2\langle \beta \rangle + 2 - 2\langle \gamma \rangle} \cdot t$$

as desired. \square

The intuition for the conclusion above is simple: We would most prefer to set $q_1 = q_2 + \beta$ and $q_3 = q_4 + \gamma$, but this is not feasible when β and γ are not integers. Thus, there is inevitably some loss from the ideal case where signals are perfectly divisible. This loss turns out to be convex in signal counts, so both groups of signals are observed infinitely often to minimize total loss.

B.2 Strengthening of Theorem 2 part (b)

Here we show the following result, which strengthens the restated Theorem 2 part (b) (see Appendix A.4). It implies that under Assumption 3, any signal that is observed with zero long-run frequency must in fact be observed only finitely often.

Stronger Version of Theorem 2 part (b). *Suppose that the signals observed infinitely often span \mathbb{R}^K . Then $m_i(t) = \lambda_i^* \cdot t + O(1), \forall i$.*

The proof is divided into two subsections below.

B.2.1 A Weaker Result

Recall that we have previously shown $m_i(t) \sim \lambda_i^* \cdot t$. We can first improve the estimate of the residual term to $m_i(t) = \lambda_i^* \cdot t + O(\ln t)$. Indeed, Lemma 10 yields that for some constant L and every $t \geq L$,

$$V(m(t+1)) \leq V(m(t)) - \left(1 - \frac{L}{t}\right) \cdot \frac{V(m(t))^2}{\phi(\mathcal{S}^*)^2}. \quad (22)$$

This is because we may apply Lemma 10 with $M = \min_{j \in \mathcal{S}^*} m_j(t)$, which is at least $\frac{t}{L}$.

Let $g(t) = \frac{V(m(t))}{\phi(\mathcal{S}^*)^2}$. Then the above simplifies to

$$g(t+1) \leq g(t) - \left(1 - \frac{L}{t}\right) g(t)^2.$$

Inverting both sides, we have

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)} + \frac{1 - L/t}{1 - (1 - L/t)g(t)} \geq \frac{1}{g(t)} + 1 - \frac{L}{t}. \quad (23)$$

This enables us to deduce

$$\frac{1}{g(t)} \geq \frac{1}{g(L)} + \sum_{x=L}^{t-1} \left(1 - \frac{L}{x}\right) \geq t - O(\ln t).$$

Thus $g(t) \leq \frac{1}{t - O(\ln t)} \leq \frac{1}{t} + O(\frac{\ln t}{t^2})$. That is,

$$V(m(t)) \leq \frac{\phi(\mathcal{S}^*)^2}{t} + O\left(\frac{\ln t}{t^2}\right).$$

Since $t \cdot V(\lambda t)$ approaches $V^*(\lambda)$ at the rate of $\frac{1}{t}$, we have

$$V^*\left(\frac{m(t)}{t}\right) \leq t \cdot V(m(t)) + O\left(\frac{1}{t}\right) \leq \phi(\mathcal{S}^*)^2 + O\left(\frac{\ln t}{t}\right). \quad (24)$$

Suppose $\mathcal{S}^* = \{1, \dots, k\}$. Then the above estimate together with (11) implies $\sum_{j>k} \frac{m_j(t)}{t} = O(\frac{\ln t}{t})$. Hence $m_j(t) = O(\ln t)$ for each signal j outside of the best set.

Now we turn attention to those signals in the best set. If these were the only available signals, then the analysis in [Liang, Mu and Syrgkanis \(2017\)](#) gives $\partial_i V(m(t)) = -\left(\frac{\beta_i^{\mathcal{S}^*}}{m_i(t)}\right)^2$. In our current setting, signals $j > k$ affect this marginal value of signal i , but the influence is limited because $m_j(t) = O(\ln t)$. Specifically, we can show that

$$\partial_i V(m(t)) = -\left(\frac{\beta_i^{\mathcal{S}^*}}{m_i(t)}\right)^2 \cdot \left(1 + O\left(\frac{\ln t}{t}\right)\right).$$

This then implies $m_i(t) \leq \lambda_i^* \cdot t + O(\ln t)$.⁴⁴ Using $\sum_{i \leq k} m_i(t) = t - O(\ln t)$, we deduce that $m_i(t) \geq \lambda_i^* \cdot t - O(\ln t)$ must also hold. Hence $m_i(t) = \lambda_i^* \cdot t + O(\ln t)$ for each signal i .

B.2.2 Getting Rid of the Log

In order to remove the $\ln t$ residual term, we need a refined analysis. The reason we ended up with $\ln t$ is because we used (22) and (23) at *each* period t ; the “ $\frac{L}{t}$ ” term in those equations adds up to $\ln t$. In what follows, instead of quantifying the variance reduction in each period (as we did), we will lower-bound the variance reduction over multiple periods. This will lead to better estimates and enable us to prove $m_i(t) = \lambda_i^* \cdot t + O(1)$.

To give more detail, let $t_1 < t_2 < \dots$ denote the periods in which some signal $j > k$ is chosen. Since $m_j(t) = O(\ln t)$ for each such signal j , $t_l \geq 2^{\epsilon \cdot l}$ holds for some positive constant ϵ and each positive integer l . Continuing to let $g(t) = \frac{V(m(t))}{\phi(\mathcal{S}^*)^2}$, our goal is to estimate the difference between $\frac{1}{g(t_{l+1})}$ and $\frac{1}{g(t_l)}$.

⁴⁴Otherwise, consider $\tau + 1 \leq t$ to be the last period in which signal i was observed. Then $m_i(\tau)$ is larger than $\lambda_i^* \cdot \tau$ by several $\ln(\tau)$, while there exists some other signal \hat{i} in the best set with $m_{\hat{i}}(\tau) < \lambda_{\hat{i}}^* \cdot \tau$. But then $|\partial_i V(m(\tau))| < |\partial_{\hat{i}} V(m(\tau))|$, meaning that the agent in period $\tau + 1$ should not have chosen signal i .

Ignoring period t_{l+1} for the moment, we are interested in $\frac{\phi(\mathcal{S}^*)^2}{V(m(t_{l+1}-1))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l))}$, which is just the difference in the *precision* about ω when the division vector changes from $m(t_l)$ to $m(t_{l+1}-1)$. From the proof of Lemma 9, we can estimate this difference if the change were along the direction λ^* :

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1} - 1 - t_l))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l))} \geq t_{l+1} - 1 - t_l. \quad (25)$$

Now, the vector $m(t_{l+1}-1)$ is not exactly equal to $m(t_l) + \lambda^*(t_{l+1}-1-t_l)$, so the above estimate is not directly applicable. However, by our definition of t_l and t_{l+1} , any difference between these vectors must be in the first k signals. In addition, the difference is bounded by $O(\ln t_{l+1})$ by what we have shown. This implies⁴⁵

$$V(m(t_{l+1}-1)) - V(m(t_l) + \lambda^*(t_{l+1}-1-t_l)) = O\left(\frac{\ln^2 t_{l+1}}{t_{l+1}^3}\right).$$

Since $V(m(t_{l+1}-1))$ is on the order of $\frac{1}{t_{l+1}}$, we thus have (if the constant L is large)

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_{l+1}-1))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1}-1-t_l))} \geq -\frac{L \ln^2 t_{l+1}}{t_{l+1}}. \quad (26)$$

(25) and (26) together imply

$$\frac{1}{g(t_{l+1}-1)} \geq \frac{1}{g(t_l)} + (t_{l+1} - 1 - t_l) - \frac{L \ln^2 t_{l+1}}{t_{l+1}}.$$

Finally, we can apply (23) to $t = t_{l+1} - 1$. Altogether we deduce

$$\frac{1}{g(t_{l+1})} \geq \frac{1}{g(t_l)} + (t_{l+1} - t_l) - \frac{2L \ln^2 t_{l+1}}{t_{l+1}}.$$

Now observe that $\sum_l \frac{2L \ln^2 t_{l+1}}{t_{l+1}}$ converges (this is the sense in which our estimates here improve upon (23), where $\frac{L}{t}$ leads to a divergent sum). Thus we are able to conclude

$$\frac{1}{g(t_l)} \geq t_l - O(1), \quad \forall l.$$

In fact, this holds also at periods $t \neq t_l$. Therefore $V(m(t)) \leq \frac{\phi(\mathcal{S}^*)^2}{t} + O(\frac{1}{t^2})$, and

$$V^*\left(\frac{m(t)}{t}\right) \leq t \cdot V(m(t)) + O\left(\frac{1}{t}\right) \leq \phi(\mathcal{S}^*)^2 + O\left(\frac{1}{t}\right). \quad (27)$$

This equation (27) improves upon the previously-derived (24). Hence by (11) again, $m_j(t) = O(1)$ for each signal $j > k$. And once these signal counts are fixed, Proposition 2 implies $m_i(t) = \lambda_i^* \cdot t + O(1)$ also holds for signal $i \leq k$. This completes the proof.

⁴⁵By the mean-value theorem, the difference can be written as $O(\ln t_{l+1})$ multiplied by a certain directional derivative. Since the coordinates of $m(t_{l+1}-1)$ and of $m(t_l) + \lambda^*(t_{l+1}-1-t_l)$ both sum to $t_{l+1}-1$, this directional derivative has a direction vector whose coordinates sum to zero. Combined with $\partial_i V(m(t)) = -(\frac{\phi(\mathcal{S}^*)^2}{t}) \cdot (1 + O(\frac{\ln t}{t}))$ (which we showed before), this directional derivative has size $O(\frac{\ln t}{t^3})$.

B.3 Example of a Learning Trap with Non-Normal Signals

The payoff-relevant state $\theta \in \{\theta_1, \theta_2\}$ is binary and agents have a uniform prior. There are three available information sources. The first, X_1 , is described by the information structure

$$\begin{array}{cc} & \theta_1 & \theta_2 \\ s_1 & p & 1-p \\ s_2 & 1-p & p \end{array}$$

with $p > 1/2$. Information sources 2 and 3 provide perfectly correlated signals (conditional on θ) taking values in $\{a, b\}$: In state θ_1 , there is an equal probability that $X_2 = a$ and $X_3 = b$ or $X_2 = b$ and $X_3 = a$. In state θ_2 , there is an equal probability that $X_2 = X_3 = a$ and $X_2 = X_3 = b$.

In this environment, every agent chooses to acquire the noisy signal X_1 , even though one observation of each of X_2 and X_3 would perfectly reveal the state.⁴⁶

B.4 Example Mentioned in Section 7.1

Suppose the available signals are

$$\begin{aligned} X_1 &= 10x + \epsilon_1 \\ X_2 &= 10y + \epsilon_2 \\ X_3 &= 4x + 5y + 10b \\ X_4 &= 8x + 6y - 20b \end{aligned}$$

where $\omega = x + y$ and b is a payoff-irrelevant unknown. Set the prior to be

$$(x, y, b)' \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.1 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 0.039 \end{pmatrix} \right).$$

It can be computed that agents observe only the signals X_1 and X_2 , although the set $\{X_3, X_4\}$ is optimal with $\phi(\{X_1, X_2\}) = 1/5 > 3/16 = \phi(\{X_3, X_4\})$. Thus, the set $\{X_1, X_2\}$ constitutes a learning trap for this problem. But if each signal choice were to produce ten independent realizations, agents starting from the above prior would observe only the signals X_3 and X_4 . This breaks the learning trap.

⁴⁶We thank Andrew Postlewaite for this example.

B.5 Supplementary Material to Section 8

B.5.1 General Payoff Functions

We comment here on the possibilities for (and limitations to) generalizing Proposition 1 beyond the quadratic loss payoff function. Proposition 1 does extend to some other “prediction problems,” in which every agent’s payoff function $u(a, \omega)$ is the same and depends only on $|a - \omega|$. For example, the proposition holds for $u(a, \omega) = |a - \omega|^\gamma$ with any exponent $\gamma \in (0, 2]$; such an extension only requires minor changes to our proof in Appendix A.7.

Nonetheless, even restricting to prediction problems, Proposition 1 does *not* hold in general. For a counterexample, consider $u(a, \omega) = -\mathbf{1}_{\{|a - \omega| > 1\}}$, which punishes the agent for any prediction that differs from the true state by more than 1.⁴⁷ Intuitively, the payoff gain from further information decreases sharply (indeed, exponentially) with the amount of information that has already been acquired. Thus, even with a forward-looking objective function, the range of future payoffs is limited and each agent cares mostly to maximize his own payoff. This results in an optimal procedure that resembles myopic behavior (and differs from the procedure that would maximize speed of learning).

The above counterexample illustrates the difficulty in estimating the value of information when working with an arbitrary payoff function. In order to make intertemporal payoff comparisons, we need to know how much payoff is gained/lost when the posterior variance is decreased/increased by a certain amount.⁴⁸ This can be challenging in general, see Chade and Schlee (2002) for a related discussion.

Finally, while it is more than necessary to assume that agents have the same payoff function, the truth of Proposition 1 does require some restrictions on how the payoff functions differ. Otherwise, suppose for example that payoffs take the form $-\alpha_t(a_t - \omega)^2$, where α_t decreases exponentially fast. Then even with the δ -discounted objective, the social planner puts most of the weight on earlier agents, resembling the myopic behavior of individual agents.

B.5.2 Low Altruism

Here we argue that part (a) of Theorem 2 generalizes to agents who are not completely myopic, but are sufficiently impatient. That is, we will show that if signals $1, \dots, k$ are subspace-optimal, then there exist priors given which agents with low δ always observe these signals in equilibrium.

⁴⁷We thank Alex Wolitzky for this example.

⁴⁸This difficulty becomes more salient if we try to go beyond prediction problems: The value of information in that case will depend on signal realizations, which brings another challenge of stochasticity.

We follow the construction in Appendix A.3. The added difficulty here is to show that if any agent ever chooses a signal $j > k$, the payoff loss in that period (relative to myopically choosing among the first k signals) is of the same magnitude as possible payoff gains in future periods. Then, for sufficiently small δ , such a deviation is not profitable.

Suppose that agents sample only from the first k signals in the first $t - 1$ periods, with frequencies close to λ^* . Then, the posterior variances v_1, \dots, v_k (which are also the prior for period t) are on the order of $\frac{1}{t}$. Thus any signal acquisition in period t leads to a variance reduction on the order of $\frac{1}{t^2}$. However, using the computation in Appendix A.3, we can show that for some positive constant ξ (independent of t), the variance reduction of any signal $j > k$ is smaller than the variance reduction of any of the first k signals by $\frac{\xi}{t^2}$. This is the amount of payoff loss in period t under a deviation to signal j .

Of course, this deviation could improve the posterior variance in future periods. But even for the best continuation strategy, the posterior variance in period $t + m$ can at most be reduced by $O(\frac{m}{t^2})$.⁴⁹ Thus if we choose ξ to be small enough, the payoff gain in period $t + m$ is bounded above by $\frac{m}{\xi t^2}$. Note that for δ sufficiently small,

$$-\frac{\xi}{t^2} + \sum_{m \geq 1} \delta^m \cdot \frac{m}{\xi t^2} < 0.$$

Hence the deviation is not profitable and the proof is complete.

B.5.3 Multiple Payoff-Relevant States

Let $V(q_1, \dots, q_N)$ be a weighted sum of posterior variances about the r payoff-relevant states. As before, define V^* to be a normalized, asymptotic version of V . Let $n(t)$ continue to represent any allocation of t observations that minimizes V . Then, under a modification of the Unique Minimizer assumption—we require V^* to be *uniquely* minimized—the optimal frequency vector $\lambda^{OPT} := \lim_{t \rightarrow \infty} n(t)/t$ is well-defined. Nonetheless, we emphasize that with $r > 1$, these optimal allocations generally involve more than K signals. A theorem of Chaloner (1984) shows that λ^{OPT} is supported on at most $\frac{r(2K+1-r)}{2}$ signals.

We can generalize the notion of “minimal spanning sets” as follows: A set of signals \mathcal{S} is minimally-spanning if optimal sampling from \mathcal{S} puts positive frequency on *every* signal in \mathcal{S} . When $r = 1$, this definition agrees with the definition in our main model. But for $r > 1$, we no longer know of a simple method for checking whether a set is minimally-spanning.

Similarly, we say that a minimal spanning set \mathcal{S} is “subspace-optimal” if, when agents are constrained to choose from $\overline{\mathcal{S}}$, the optimal frequency vector is supported on \mathcal{S} . With these definitions, Theorem 2 and its proof generalize without change.

⁴⁹This is because over m periods, the increase in the *precision* matrix is $O(m)$.