Overabundant Information and Learning Traps*

Annie Liang† Xiaosheng Mu‡

June 1, 2018

Abstract

We develop a model of social learning from overabundant information: Agents have access to many sources of information, and observation of all sources is not necessary in order to learn the payoff-relevant state. Short-lived agents sequentially choose to acquire a signal realization from the best source for them. All signal realizations are public. Our main results characterize two starkly different possible long-run outcomes, and the conditions under which each obtains: (1) efficient information aggregation, where the community eventually achieves the highest possible speed of learning; (2) "learning traps," where the community gets stuck using a suboptimal set of sources and learns inefficiently slowly. A simple property of the correlation structure separates these two possibilities. 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—at any cost—access arbitrarily precise signals about the importance of genetic factors. He can however acquire many kinds of information related to this question; for example,

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

†University of Pennsylvania

‡Harvard University

he might acquire neurochemical and genetic data from affected individuals, or observe the incidence of depression within families.

These sources of information contribute to learning in different ways, and the value of new information depends on its relationship to to what is already known. Consequently, past information acquisitions can change the perception of which kinds of information are most useful. For example, if the relationship between a given neurochemical and depression is well-understood, then comparisons in the level of that neurochemical across individuals is more meaningful. Models of information often abstract away from explicitly describing these heterogeneities across kinds of information, because it complicates the analysis of information acquisition.[1] But the relationships across these 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 long-run acquisitions in two starkly different ways:

- *Efficient information aggregation:* Past information helps future 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 future decision-makers 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 reveal the 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 of which is associated with a signal about the 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.

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

In contrast to the classic sequential learning model (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992; Smith and Sorenson, 2000), we suppose that all signal realizations are public. 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 assume that each agent takes an action based on all prior information, and maximizes an individual objective that depends only on his action and the payoff-relevant state. We work with normal signals, which yields the additional tractability that each agent's information choice is the one that maximizes that period's reduction of uncertainty about the payoff-relevant state.

Our main interest is in settings with many sources of information, including some of which 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, corresponding to the choices that maximize the speed of information revelation about the payoff-relevant state.

We then show that whether society's acquisitions converge to this optimal long-run frequency depends critically on how many signals are needed to identify this 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 additionally reveals all of the confounding terms. Then endogenously, agents will acquire information that (collectively) reveals all of the unknowns. The aggregated information eventually overwhelms prior beliefs, so that agents come to evaluate *all* sources by an "objective" asymptotic criterion. This leads them to discover 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, independently of the prior belief.

In contrast, if it is possible to learn the payoff-relevant state without recovering all of the confounding terms, then long-run learning may be inefficient. This is because agents can persistently undervalue sources that provide information confounded by the remaining unknowns. 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 the optimal speed of learning to the achieved speed of learning—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. For example, better brain imaging technology may allow researchers to make significant advances towards understanding depression. If development of these technologies requires long-term investments, and researchers are rewarded only for their immediate contributions to knowledge, they may prefer instead to exploit existing technologies. 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' payoffs depend on information that they acquire over many periods. The success of these interventions does depend 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 the number of sources and states are the same. 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, while we fix error variances, so that 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; thus, our welfare comparisons are new.

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

2 Framework

There are K persistent unknown states: A payoff-relevant state ω, and $K-1$ additional confounding states $b_1, \ldots, b_{K-1}$.[3] We assume that the state vector $\theta := (\omega, b_1, \ldots, b_{K-1})'$ follows a multivariate normal distribution $\mathcal{N}(\mu^0, V^0)$,[4] where the prior covariance matrix V^0 has full rank.[5]

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 $c_i = (c_{i1}, \ldots, c_{iK})'$ 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 permitting signals to be of differing precision levels).[6] 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 $y = \omega c_i^1 + b_1 c_i^2 + \ldots b_{K-1} c_i^K + \epsilon_i$, 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 easy they are to measure (measurement error) and how well their role is understood (quality of interpretation). The confounding terms

[3]See Appendix B.4 for a partial extension to the case of multiple-payoff relevant states.

[4]All vectors in this paper are column vectors.

[5]The 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 satisfying the full rank condition with a lower dimensional state-space.

[6]Scaling up the coefficients is equivalent to scaling up the precision of the signal.

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 acquires an independent realization of one of the N signals, and then chooses an action $a \in A$ to maximize an individual objective $u_t(a, \omega)$. He bases his action on the realization of his own signal acquisition, as well as the history of signal acquisitions and realizations thus far. Thus, all signal realizations are public.

Payoff functions may differ across agents, but we assume that all decision problems depend only on the unknown state ω and the agent's own action, and are moreover non-trivial in the following way.

Assumption 1 (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, we require that 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.

Throughout, we use $[N] = \{1, \ldots, N\}$ to index the set of signals. We call a set of signals $\mathcal{S} \subset [N]$ *spanning* if the vectors $\{c_i : i \in \mathcal{S}\}$ span the coordinate vector $e_1 = (1, 0, \ldots, 0)'$, 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 moreover no proper subset 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.[7] 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

[7]This 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 V be the subspace of $\mathbb{R}^K$ spanned by $\tilde{c}_1, \ldots, \tilde{c}_N$. Then project u onto V: $u = v + w$ with $v \in V$ and w orthogonal to V. Thus $\langle u, \tilde{\theta} \rangle = \langle v, \tilde{\theta} \rangle + \langle w, \tilde{\theta} \rangle$. By assumption, the random variable $\langle w, \tilde{\theta} \rangle$ is independent from any random variable $\langle c, \tilde{\theta} \rangle$ with $c \in V$ (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 $\langle v, \tilde{\theta} \rangle$, returning to the case where the payoff-relevant parameter *is* identified.

recover ω by observing each information source infinitely often, but not by exclusively observing any proper subset of sources.

Our main interest is in settings of *informational overabundance*, where $[N]$ is spanning but not minimally spanning. In these cases, multiple different sets 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).[8]

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. The agent's beliefs about the state vector, prior to making his own signal choice, are $\theta \sim \mathcal{N}(\mu^{t-1}, V^{t-1})$. Given an observation of the signal i, his posterior beliefs become $\theta \sim \mathcal{N}(\mu^t, V^t)$ where V^t is a *deterministic* function of the prior covariance matrix V^{t-1} and the signal choice X_i, and the posterior expected value is the random variable $\mu^t \sim \mathcal{N}(\mu^{t-1}, V^{t-1} - V^t)$.

Agent t's posterior belief about the payoff-relevant state ω is given by $\omega \sim \mathcal{N}(\mu_1^t, V_{11}^t)$.[9] His maximum expected payoff (after observing his signal) is

$$\max_{a \in A} \mathbb{E}[u_t(a, \omega) \mid \omega \sim \mathcal{N}(\mu_1^t, V_{11}^t)]. \tag{1}$$

Each agent t chooses the signal that maximizes the expected value of (1), where the expectation is taken with respect to the random variable $\mu_1^t \sim \mathcal{N}(\mu_1^{t-1}, V_{11}^{t-1} - V_{11}^t)$. From this we see that the agent's expected payoff is measurable with respect to the posterior variance V_{11}^t.

The signal acquisition that maximizes agent t's payoffs is the one that minimizes his posterior variance about ω.[10] Thus, we can track society's acquisitions as a se-

[8] It 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 $N > K$.

[9] Subscripts indicate particular entries of a vector or matrix.

[10] Under our normality assumption, the signal that maximally reduces posterior variance about ω Blackwell dominates the remaining signals; see e.g. Hansen and Torgersen (1974). So the statement here is independent of the payoff function.

quence of division vectors $m(t) = (m_1(t), \ldots, m_N(t))$, where $m_i(t)$ is the number of times that signal i has been observed up to and including time t. Let $f(q_1, \ldots, q_N)$ denote the posterior variance about ω, given the initial prior covariance matrix V^0 and q_i observations of each signal i.[11] 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,

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

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 $m(t)$.

The long-run frequencies of observation are $\lim_{t\to\infty} m_i(t)/t$ for each signal i. Our subsequent results in Section 5 show these limits to be well-defined. In Section 4, we first characterize the "optimal" acquisitions that a social planner might impose, and identify the corresponding long-run observation frequencies. In Section 5, we characterize the actual signal acquisitions, and compare this to the optimal benchmark.

4 Optimal Information Revelation

Our optimal benchmark describes the *maximum possible information revelation* about ω. For each period t, define a t-*optimal* vector to be any allocation of t observations that minimizes posterior variance about ω:

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

Then, $f(n(t))$ is the lowest achievable posterior variance by period t. Generically, there is a unique t-optimal division vector for every t.[12]

We interpret each $n(t)$ as the optimal social benchmark for the finite horizon problem with final period t. Suppose a social planner takes an action a on behalf of the society at period t to satisfy a payoff criterion that depends on his action a

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

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

and the payoff-relevant state ω. Then, at every t, the social planner's payoffs are maximized if the history of t signal acquisitions corresponds to a t-optimal division.[13]

The limiting frequencies $\lim_{t\to\infty} n(t)/t$ are well-defined under a subsequent condition (Assumption 2), and we refer to these as the *optimal frequencies*. In Appendix B.1, we show that under conditions on agents' payoff functions, any strategy that maximizes a δ-discounted payoff objective also approximates these optimal frequencies as $\delta \to 1$.[14] This result further justifies the use of $\lim_{t\to\infty} n(t)/t$ as a benchmark.

We consider first a restricted version of the social planner problem, supposing that agents must 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. In such a case, it is possible to decompose the first coordinate vector e_1 as the following (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 under optimal sampling from $\mathcal{S}$, each signal $i \in \mathcal{S}$ should be observed (asymptotically) in proportion to its coefficient $\beta_i^{\mathcal{S}}$:

Proposition 1 (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*

$$n_i^{\mathcal{S}}(t) = \frac{|\beta_i^{\mathcal{S}}|}{\sum_{j\in\mathcal{S}}|\beta_j^{\mathcal{S}}|} \cdot t + O(1). \tag{2}$$

Throughout, $O(1)$ represents a residual term that remains bounded as $t \to \infty$.

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 the signals in $\mathcal{S}$ have full rank,[15] and we let $C_{\mathcal{S}}$ denote the matrix of these coefficient vectors. The

[13]The social planner's optimal strategy is to observe each signal i exactly $n_i(t)$ times, in an arbitrary order. In particular, such a strategy does not need to condition on signal realizations; see Liang, Mu and Syrgkanis (2017) for further discussion.

[14]It is clear that the strategy that samples signals (randomly) according to the optimal frequencies is best among *stationary* information acquisition strategies as $\delta \to 1$. However, our result in Appendix B.1 is stronger because the optimal strategy for any fixed δ may be far from stationary.

[15]Otherwise, $\mathcal{S}$ would not be a *minimal* spanning set.

(random) vector of realizations corresponding to one observation of each signal in this set can be written as

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

where ε is a $K \times 1$ vector of error terms. Given these realizations, the best linear unbiased estimate for ω is

$$\hat{\omega} = \left[C_{\mathcal{S}}^{-1} Y\right]_{11}. \tag{3}$$

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 the OLS estimate responds to changes in the realization of signal i. Proposition 1 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 that is achievable from signals in $\mathcal{S}$:

Corollary 1. *Suppose agents are constrained to a minimal spanning set $\mathcal{S}$. The minimum achievable posterior variance after t observations satisfies the following approximation:*

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

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

Thus, sampling according to the frequencies $n^{\mathcal{S}}(t)/t$ will approximate (at large period t) a posterior variance of

$$\text{AsympVar}_t(\mathcal{S}) = \left(\sum_{i \in S} |\beta_i^{\mathcal{S}}|\right)^2 / t := \frac{\phi(\mathcal{S})^2}{t}.$$

In what follows, we work with the simpler statistic $\phi(\mathcal{S})$ (roughly an *asymptotic standard deviation*), noting that the asymptotic variance is strictly increasing in $\phi(\mathcal{S})$. The smaller $\phi(\mathcal{S})$ is, the faster the community learns from $\mathcal{S}$, so ϕ establishes an ordering over minimal spanning sets.

We assume throughout that there is a *best* minimal spanning set according to this ordering:

Assumption 2 (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 signals are $X_1 = \omega + \epsilon_1$ and $X_2 = \omega + \epsilon_2$. Unique Minimizer fails, because learning occurs equally fast from either of the minimal spanning sets $\{X_1\}$ or $\{X_2\}$.

Example 2. The signals are $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, $X_3 = \omega + b_2 + \epsilon_3$, and $X_4 = b_2 + \epsilon_4$. Unique Minimizer fails, because learning occurs equally fast from either of the minimal spanning sets $\{X_1, X_2\}$ and $\{X_3, X_4\}$.

These examples are special, in the sense that Assumption 2 holds under arbitrarily small perturbations of the above environments.

If we restrict agents to sample exclusively from a single minimal spanning set, then the optimal sampling rule (under Assumption 2) is clearly the frequency vector $\lambda^* \in \Delta^{N-1}$ satisfying

$$\lambda_i^* = \begin{cases} \frac{|\beta_i^{\mathcal{S}^*}|}{\sum_{j \in \mathcal{S}^*} |\beta_j^{\mathcal{S}^*}|} & \forall\ i \in \mathcal{S}^* \\ 0 & \forall\ i \notin \mathcal{S}^* \end{cases} \tag{4}$$

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 1.

In principle, the community 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, then the best long-run strategy is to restrict to the best minimal spanning set, and to sample from that set as in Proposition 1.

Theorem 1. *Let λ^* be given by (4). Under Assumption 2, the optimal count $n_i(t) \sim \lambda_i^* \cdot t$ for each signal i.*[16]

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.[17]

[16]We conjecture that the stronger conclusion $n_i(t) = \lambda_i^* \cdot t + O(1)$ also holds. In Appendix A.2.5, we prove this conjecture assuming $|\mathcal{S}^*| = K$.

[17]Specifically, 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 $\beta_i = \beta_i^*$ (that is, when focusing on a single minimal spanning set).

We show in Appendix B.2 that Assumption 2 is necessary for the result: Indeed, in the environment described in Example 2, there are priors given which it is *strictly* optimal to observe all four available signals with positive frequency.

Theorem 1 directly implies the following comparative static: If signal i is viewed with positive long-run frequency in the social planner problem, then this frequency is (locally) decreasing in its precision.

Corollary 2. *Suppose the coefficient matrix C satisfies Unique Minimizer. Write each signal as $X_i = \alpha\langle c_i, \theta\rangle + \epsilon_i$, so that the precision of signal X_i is increasing in α. Then, either $\lambda_i^* = 0$ or λ_i^* is locally decreasing in α.*

Consider a problem complementary to ours, in which information sources choose the precision of their signals in order to maximize the long-run frequency λ_i^* with which they are viewed.[18] There is clearly an incentive to provide precise information, since sources are not viewed in the long-run unless they belong to the best set $\mathcal{S}^*$. But Corollary 2 implies that conditional on inclusion in this set, each source wants to provide signals as imprecise as possible. These conflicting forces suggest that characterization of the equilibrium provisions of information precision has interesting subtleties.

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. Below, we show that whether society's acquisitions $m(t)$ eventually approximate the optimal acquisitions $n(t)$ depends critically on how many signals are required to identify ω.

We first present our main results under the following (generic) technical assumption:

Assumption 3 (Strong Linear Independence)**.** *$N \geq K$ and every $K \times K$ submatrix of C is of full rank.*

Strong Linear Independence requires that every set of $k \leq K$ signals is linearly

[18] Similar comparative statics hold for society's long-run frequencies, which are characterized later.

independent.[19] We impose this restriction in Sections 5.1 and 5.2 to allow for a simpler exposition of the main forces. Our results extend beyond Strong Linear Independence, and we characterize the general setting 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.

Example 3. There are three available signals:

$$X_1 = \omega/2 + \epsilon_1$$
$$X_2 = \omega + b_1 + \epsilon_2$$
$$X_3 = \omega - b_1 + \epsilon_3$$

Both $\{X_1\}$ and $\{X_2, X_3\}$ are minimal spanning sets, but $\{X_2, X_3\}$ is optimal because $\phi(\{X_1\}) = 2 > 1 = \phi(\{X_2, X_3\})$.[20]

Consider a prior where ω and b_1 are independent, and the prior variance of b_1 is large (exceeds 3). In the first period, the precision of the first signal exceeds that of the latter two signals (interpreting all signals as noisy observations of ω).[21] 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.

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.

Corollary 3. *Under Strong Linear Independence, for every minimal spanning set $\mathcal{S}$ that contains fewer than K signals, there exists an open set of prior beliefs under which agents exclusively observe signals from $\mathcal{S}$.*

[19] Besides 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.

[20] 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.

[21] The signal $X_1 = \omega/2 + \epsilon_1$ is equivalent to $\omega + 2\epsilon_1$, which is distributed as $\mathcal{N}(\omega, 4)$. Each of the signals $\omega + b_1 + \epsilon_2$ and $\omega + b_2 + \epsilon_2$ has greater variance conditional on ω.

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 $1/(c^2 t)$.[22] Across all information environments, however, no uniform bound on inefficiency is possible. Specifically, for any positive number L, there exists an environment in 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 $X_2 = \alpha\omega + b_1 + \epsilon_2, X_3 = \alpha\omega - b_1 + \epsilon_3$ with α sufficiently large.[23] Thus, the size of inefficiency can be arbitrarily large.

5.2 Efficient Information Aggregation

Suppose in contrast 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 frequencies. Thus, even though agents are short-lived ("myopic"), they will end up acquiring information in a way that is eventually socially best.

Corollary 4. *Under Unique Minimizer and Strong Linear Independence, 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.*

A brief intuition for this result is as follows: If K signals must be observed in order to recover ω, then the incentive to learn ω will *endogenously* drive agents to sample from at least K different signals. Under the assumption of Strong Linear Independence, these K signals further reveal all K unknown states. Thus, as observations accumulate, agents not only learn about ω but about all of the confounding

[22]This 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 Andy Skrzypacz for this observation.

[23]The "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.

terms. This allows agents to eventually evaluate all signals according to an "objective" asymptotic value, and to identify the best set.

The condition that all minimal spanning sets have size K is generically satisfied.[24] 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 any source that directly reveals ω (that is, $X = \alpha\omega + \epsilon$) is non-generic in the probabilistic sense, 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 3 shows that inefficiency is a likely outcome in these cases.

5.3 General Setting

We now state a more general version of our results that does not require Strong Linear Independence. 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 $\mathcal{A}$. For example, if the available signals are $X_1 = \omega/2 + \epsilon_1$ and $X_2 = \omega + \epsilon_2$, and we define $\mathcal{S} = \{X_1\}$, then $\overline{\mathcal{S}} = \{X_1, X_2\}$. 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. For example, the set $\{X_1\}$ is minimally spanning but not subspace-optimal.

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

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 4 (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 $\overline{\mathcal{A}}$.*

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

[24] We point out that the set of coefficient matrices satisfying Unique Minimizer is "generic" in the following stronger sense: Fixing 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.

Our next result generalizes both the learning trap 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 under which long-run frequencies are strictly positive for signals in $\mathcal{S}$, and zero everywhere else.*

(b) Under Assumption 4, long-run frequencies exist for every signal. Moreover, if $\mathcal{S}$ denotes the signals viewed with positive long-run frequencies, then $\mathcal{S}$ is a minimal spanning set that is subspace-optimal.

This theorem directly implies our previous Corollaries 3 and 4. To see this, note that under Strong Linear Independence, $\overline{\mathcal{S}} = \mathcal{S}$ for every minimal spanning set $\mathcal{S}$ with fewer than K signals. This implies that every minimal spanning set with fewer than K signals is (trivially) optimal in its subspace, producing Corollary 3 from part (a) of Thereom 2.

On the other hand, if every minimal spanning set has size K and Strong Linear Independence is satisfied, 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 4 from part (b) of the theorem above.

6 Intuitions for Main Results

6.1 High-Level Argument

Agents compare the marginal value of observations of different signals. 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 allows them to identify the best set of signals and approximate the optimal frequencies.

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 prior. It is exactly this difference that leads to our learning trap result (Corollary 3): 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, our analysis reveals that our previous arguments for the exactly sufficient information case hold "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 (which are independent of the prior beliefs) *for 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 Proof Sketch for Theorem 2

First observe that society's acquisitions follow a procedure of "pseudo"-gradient descent, where the frequency vector $\lambda(t) = m(t)/t$ evolves according to

$$\lambda(t+1) = \frac{t}{t+1}\lambda(t) + \frac{1}{t+1}e_i$$

and e_i represents the coordinate vector that yields the greatest (immediate) reduction in the posterior variance function f.

Instead of working directly with posterior variance, we define the following related function, which takes as input frequency vectors $\lambda \in \Delta^K$ and describes a "normalized" asymptotic posterior variance:

$$f^*(\lambda_1, \ldots, \lambda_N) = \lim_{t\to\infty} t \cdot f(\lambda_1 t, \ldots, \lambda_N t).$$

We establish the following relationships between f^* and f. First, signal acquisitions chosen according to a frequency vector that minimizes f^* will asymptotically also minimize the posterior variance function f (Lemma 3); this justifies our study of f^*. Second, $f^*(\lambda)$ 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 frequencies $\lambda(t)$ come to minimize f^*. Third, under a condition (that we show will be met at late periods), the signal that achieves the greatest reduction in f also roughly achieves the greatest reduction in f^* (Lemma 9). This allows us to consider (pseudo-)gradient descent in terms of f^*.

While the convexity of f^* ensures that standard gradient descent is well-behaved, the process of descent in our problem can only occur along a finite set of feasible directions (indexed by the available signals). This constraint corresponds to our assumption that each agent acquires a single, discrete, observation of a chosen signal. The limitation is without loss whenever f^* is differentiable, since all directional derivatives can then be rewritten as convex combinations of the partial derivatives along basis vectors.[25] The function f^*, however, is *not* differentiable everywhere. Consider our learning trap example with signals

$$X_1 = \omega/2 + \epsilon_1$$
$$X_2 = \omega + b_1 + \epsilon_2$$
$$X_3 = \omega - b_1 + \epsilon_3$$

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.[26] This means that the derivative of f^* 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, f^* is not differentiable at λ.

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 f^*. This is reflected in our learning trap results (Corollary 3 and part (a) of Theorem 2). A key lemma towards efficient information aggregation shows that f^* is differentiable whenever λ places nonzero weight on a spanning set of K signals. The condition in Corollary 4 guarantees that the community will repeatedly observe such a set, since ω (by assumption) can only be recovered from a spanning set of K signals. Thus, the sufficient condition for differentiability is eventually satisfied,

[25] The limitation also goes away if each agent is allowed to acquire sufficiently many observations of different signals; see Proposition 2 below.

[26] That is, $f^*(1, 0, 0)$ is strictly smaller than both $f^*(1-\epsilon, \epsilon, 0)$ and $f^*(1-\epsilon, 0, \epsilon)$, but it is strictly larger than $f^*\left(1-\epsilon, \frac{\epsilon}{2}, \frac{\epsilon}{2}\right)$.

and descent (from then on) is well-behaved, ending at the global minimum λ^*. This explains the result in Corollary 4. Part (b) of Theorem 2 follows from a similar (albeit slightly more involved) argument.

Remark 1. *Our analysis 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 $-f^*(\lambda)$. Then, we have a potential game with (exact) potential function $-f^*$, 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.*

The connection to potential games is not sufficient to derive Theorem 2, since we have a setting in which play changes the payoff function $-f(\lambda t)$. This contrasts with most of the potential games literature, which studies convergence of play in a fixed *game. We use the potential function $-f^*$ as an approximation for agents' behavior given $-f(\lambda t)$ at late periods t, and a key step in the proofs is to demonstrate the validity of this approximation (which depends on the prior belief as well as the endogenous path of signal acquisitions).*[27]

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 nudge agents towards efficient learning. Naturally, this question applies only when agents would otherwise achieve a suboptimal speed of learning (with conditions given in part (a) of Theorem 2).

We compare several possible policy interventions. One is to increase the quality of information acquisition, so that each individual observation is more informative. For example, if we interpret signal acquisitions as experiments, policymakers can provide funding that allows for collection of more observations per experiment. We show

[27] Note also that the prior belief—which is outside of the description of the asymptotic potential game—influences which outcome agents will end up converging to. This is a feature special to our setting.

that this intervention is of limited effectiveness: Any set of signals that is a potential learning trap remains a potential learning trap under arbitrary improvements to signal precision.

Another possibility is to restructure the incentive structure so that agents' payoffs are based on information obtained over several periods (equivalent to acquisition of a block of signals each period). This would correspond, for example, to evaluation of researchers based on contributions developed over several studies, or a lab structure where payoffs to individual researchers determined by the lab's aggregate output. We show that it is possible to guarantee efficient information aggregation in this case, but the number of periods (or observations) that is needed varies across informational environments.

Finally, we consider one-shot provision of free information, which allows the policy-maker to release information about confounding terms that would not otherwise be of direct interest. Intuitively, this information can have long-term effects for what kind of information is subsequently acquired. For example, government-funded advancements in brain imaging technology may induce researchers to take more brain scans, which subsequently increases the value to advancing this imaging technology. In this sense, auxiliary information can push agents onto the right path. We show that given sufficiently many kinds of information (which we characterize), efficient learning can be restored.

7.1 More Precise Information

Consider an intervention that increases the information value of any signal draw. Formally, we suppose that each signal acquisition produces B independent observations from that source.

The result below (which follows as a corollary from Theorem 2) says that every set of signals that is a potential learning trap given $B = 1$ (as in our main model) remains a potential learning trap for every choice of B.

Corollary 5. *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 our previous learning traps example (Example 3). This is an environment in which increasing the precision of signals is ineffective: At the first period, the best choice is X_1 regardless of the value of B. Our previous logic again implies that each subsequent agent will also choose signal X_1, so that $\{X_1\}$ remains a learning trap. In Appendix B.3, 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 and the prior.

7.2 More Kinds of Information

Suppose that the policymaker can restructure the incentive structure, so that each agent has B periods to acquire information prior to taking an action. This is equivalent to supposing that each agent receives a block of B signals to allocate across sources. Then:

Proposition 2. *Under Unique Minimizer, there is a B such that given acquisition of B signals every period, long-run frequencies are 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 frequencies.

The number of observations needed, however, depends on subtle details of the informational environment. In particular, there is no uniform bound for B over all environments of fixed size; that is, there does not exist a function $B(K, N)$ such that acquisition of $B(K, N)$ signals by each agent ensures efficient information aggregation for all environments with K states and N signals.

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.[28] 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.

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

7.3 Free Information

Finally, suppose that the policy-maker can either directly provide free information, or can subsidize investment in this information by an external party. Formally, the policy-maker can provide (independent realizations of) 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 $t = 0$, this information is made public. All subsequent agents update their prior beliefs based on this free information in addition to the history of signal acquisitions thus far.

Is there a sufficient number of (kinds of) signals, such that efficient learning can be guaranteed under such an intervention? We answer in the affirmative below: Let $k \leq K$ be the size of the optimal set $\mathcal{S}^*$. Then $k-1$ precise signals are sufficient to guarantee efficient learning:

Proposition 3. *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 frequencies are λ^* starting from every prior belief.*

When designing these signals, the policy-maker does not need to teach directly about the payoff-relevant state ω, which the agents will learn on their own. Rather, auxiliary information should be provided to help agents better interpret the confounding terms. Our proof shows 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.[29]

8 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. We compare their information acquisitions with an optimal benchmark, under which the speed of information revelation is maximized.

In general, because agents do not internalize the impact of their information acquisitions on later decision-makers, inefficient information acquisition may obtain.

[29] This 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 $K-1$ signals about all of the confounding terms.

Specifically, past information acquisitions can increase the value of "low-quality" sources relative to "high-quality" sources, pushing future agents to acquire information from a set of sources that yields inefficiently slow learning. We show however that inefficiency is not guaranteed: Depending on the correlation structure, myopic concerns can endogenously push agents to identify and observe only the most informative sources. Our main results separate these outcomes, and fully characterize the set of possible long-run outcomes.

Our framework and results highlight some of the forces that are important for the design of incentives for information acquisition. In particular, do the kinds of information that are of immediate societal interest also have spillovers for knowledge that is only of indirect value? When such spillovers are present, simple incentive schemes—in which agents care about immediate contributions to knowledge—are sufficient to enable 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 learning about unknowns that are not directly of interest, but which are useful as intermediate steps. Alternatively, agents can be rewarded for advancements developed across several contributions. These observations are consistent with practices that have arisen in academic research, including evaluation of the contribution of a body of work, and the establishment of third-party funding agencies to support methodological research and technical advances.

References

Ali, Nageeb. 2018. "Herding with Costly Information." *Journal of Economic Theory.*

Banerjee, Abhijit. 1992. "A Simple Model of Herd Behavior." *Quaterly 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 Substitues."

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. "Groupthink 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

A.1 Posterior Variance Function

A.1.1 A Basic Lemma

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 V^0 and q_i observations of each signal i, society's posterior variance about ω is given by*

$$f(q_1, \dots, q_N) = \left[((V^0)^{-1} + C'QC)^{-1}\right]_{11} \tag{5}$$

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

Proof. Note that $(V^0)^{-1}$ 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 (5) simply represents the fact that for a Gaussian prior and Gaussian signals, the posterior precision matrix is the sum of the prior and signal precision matrices. To prove the monotonicity of f, 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 $((V^0)^{-1} + C'QC)^{-1}$ decreases in this order, which implies that the posterior variance about ω decreases. Intuitively, more information always improves the decision-maker's estimates.

To prove that f is convex, it suffices to prove that f is *midpoint-convex* since the function is clearly continuous.[30] Take $q_1, \dots, q_N, r_1, \dots, r_N \in \mathbb{R}_+$ and let $s_i = \frac{q_i + r_i}{2}$.

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

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

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

Using (5), we conclude that

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

This proves the convexity of f. □

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

$$X=UDU'$$

with U being a $K\times K$ orthogonal matrix whose columns are eigenvectors of X, and D being a $K\times K$ diagonal matrix consisting of non-negative eigenvalues. Even if some of these eigenvalues are zero, we can think of X^{-1} as

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

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

$$[X^{-1}]_{11}=\sum_{j=1}^{K}\frac{(\langle u_j,e_1\rangle)^2}{d_j}, \tag{6}$$

with the convention that $\frac{0}{0}=0$. Note that by this definition,

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

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. Note that we allow $[X^{-1}]_{11}$ to be infinite.

A.2 Proof of Theorem 1

A.2.1 Asymptotic Behavior of Posterior Variance

We first approximate the posterior variance as a function of the frequencies with which each signal is observed. Specifically,

Lemma 2. *For any $\lambda_1, \dots, \lambda_N \geq 0$, let $\Lambda = \operatorname{diag}(\lambda_1, \dots, \lambda_N)$. Then*

$$\begin{aligned} f^*(\lambda_1, \dots, \lambda_N) &:= \lim_{t \to \infty} t \cdot f(\lambda_1 t, \dots, \lambda_N t) \\ &= [(C' \Lambda C)^{-1}]_{11} \end{aligned} \tag{7}$$

Note that the matrix $C'\Lambda C$ is positive semi-definite. So the value of $[(C'\Lambda C)^{-1}]_{11}$ is well-defined, see (6).

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

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

Hence by continuity of $[X^{-1}]_{11}$ in the matrix X, we obtain the lemma. □

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 f^*

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

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

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

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

$$t_m \cdot f(n_1(t_m), \dots, n_N(t_m)) \geq t_m \cdot f(\tilde{\lambda}_1 \cdot t_m, \dots, \tilde{\lambda}_N \cdot t_m) \to f^*(\tilde{\lambda}_1, \dots, \tilde{\lambda}_N)$$

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

$$t_m \cdot f(\hat{\lambda}_1 \cdot t_m, \dots, \hat{\lambda}_N \cdot t_m) \to f^*(\hat{\lambda}_1, \dots, \hat{\lambda}_N).$$

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

$$f(n_1(t_m), \dots, n_K(t_m)) > f(\hat{\lambda}_1 \cdot t_m, \dots, \hat{\lambda}_N \cdot t_m).$$

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

A.2.3 Crucial Lemma

We pause to demonstrate the following technical lemma:

Lemma 4. *Suppose $\mathcal{S}^* = \{1, \dots, K\}$ uniquely minimizes $\phi(\mathcal{S})$ 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 $1 \le j \le K$. 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 $\sum_{j=1}^K \alpha_j \ge 1$ (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 $2 \le j \le K$, $\beta_j - \frac{\beta_1}{\alpha_1} \cdot \alpha_j$ is clearly positive if $\alpha_j \le 0$ (since $\beta_j > 0$). And if $\alpha_j > 0$, then by assumption $\frac{\alpha_j}{\beta_j} \le \frac{\alpha_1}{\beta_1}$ 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

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

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

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 $f^*(\lambda)$ 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. *Suppose $\mathcal{S}^* = \{1,\dots,K\}$ is the unique minimizer of $\phi(\mathcal{S})$ over minimal spanning sets. Define $\lambda^* \in \Delta^{N-1}$ by*

$$\lambda_i^* = \frac{|[(C^*)^{-1}]_{1i}|}{\sum_{j=1}^K |[(C^*)^{-1}]_{1j}|}, 1 \leq i \leq K$$

with $C^* = C_{[K][K]}$,[31] *and $\lambda_i^* = 0, \forall i > K$. Then $f^*(\lambda^*) < f^*(\lambda)$ for any $\lambda \in \Delta^{N-1}, \lambda \neq \lambda^*$.*

Proof. First, we will 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 $f(q_1,\dots,q_N)$ is convex in its arguments, $f^*(\lambda) = \lim_{t\to\infty} t \cdot f(\lambda_1 t,\dots,\lambda_N t)$ is also convex in λ. To show $f^*(\lambda^*) < f^*(\lambda)$, we only need to show $f^*(\lambda^*) < f^*((1-\epsilon)\lambda^* + \epsilon\lambda)$ for some $\epsilon > 0$. In other words, it suffices to show $f^*(\lambda^*) < f^*(\lambda)$ for λ in 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 (7) the function f^* is differentiable near λ^* (see Remark 2 below).

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

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

[31] For any subset $\mathcal{I} \subset [N]$ and $\mathcal{J} \subset [K]$, write $C_{\mathcal{I}\mathcal{J}}$ for the sub-matrix of C with row indices in $\mathcal{I}$ and column indices in $\mathcal{J}$. Likewise, let $C_{-\mathcal{I}\mathcal{J}}$ be the sub-matrix of C after deleting rows in $\mathcal{I}$ and columns in $\mathcal{J}$.

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

$$f^*(\lambda_1,\ldots,\lambda_K,\lambda_{K+1},\ldots,\lambda_N) \geq f^*\left(\lambda_1,\ldots,\lambda_{K-1},\sum_{k=K}^{N}\lambda_k,0,\ldots,0\right) \geq f^*(\lambda^*). \quad (8)$$

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

To prove (**), we recall that

$$f^*(\lambda_1,\ldots,\lambda_N) = e_1'(C'\Lambda C)^{-1}e_1.$$

Since $\Lambda = \text{diag}(\lambda_1,\ldots,\lambda_N)$, 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

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

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

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

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

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

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

$$\partial_i f = -\gamma_i^2,\ 1\leq i\leq N. \quad (9)$$

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

$$\gamma_i = e_1'\cdot(C^*)^{-1}\cdot\text{diag}\left(\frac{1}{\lambda_1^*},\ldots,\frac{1}{\lambda_K^*}\right)\cdot e_i = \frac{[(C^*)^{-1}]_{1i}}{\lambda_i^*} = \sum_{j=1}^{K}[(C^*)^{-1}]_{1j} = \phi(\mathcal{S}^*) \quad (10)$$

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

$$\gamma_i = \langle v, c_i \rangle = \sum_{j=1}^{K} \alpha_j \cdot \langle v, c_j \rangle = \sum_{j=1}^{K} \alpha_j \cdot \gamma_j = \phi(\mathcal{S}^*) \cdot \sum_{j=1}^{K} \alpha_j. \tag{11}$$

The last equality uses (10). 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 (9) proves the desired inequality (**), and the lemma follows. □

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

$$\phi(\{1, \ldots, K\}) < \phi(\{1, \ldots, K\} \cup \{i\} \backslash \{j\}), \ \forall 1 \leq j \leq K < i \leq N.$$

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 6.2 of the main text for an example. In that case, even though the partial derivatives satisfy (**), we cannot deduce that any directional *derivative similarly satisfies (**). To put it differently, f^* fails 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 later.*

A.2.5 Stronger Statement of Theorem 1 when $|\mathcal{S}^*| = K$

Still assuming that the "best" subset $\mathcal{S}^*$ contains exactly K signals, we now show $n_i(t) = \lambda_i^* \cdot t + O(1), \forall i$, which improves upon the conclusion of Theorem 1. First, we can apply Lemma 4 to find a positive constant $\eta < 1$ such that for *each* $i > K$, if $c_i = \sum_{j=1}^{K} \alpha_j c_j$ then $|\sum_{j=1}^{K} \alpha_j| \leq 1 - \eta$. By (9), (10) and (11), we have

$$\partial_1 f(\lambda^*) = \cdots = \partial_K f(\lambda^*) = -\phi(\mathcal{S}^*)^2; \quad \partial_i f(\lambda^*) \geq -(1-\eta)^2 \cdot \phi(\mathcal{S}^*)^2, \forall i > K. \tag{12}$$

For any $\lambda \in \Delta^{N-1}$, the convexity of f^* implies[32]

$$\begin{aligned} f^*(\lambda) &\geq f^*(\lambda^*) + \sum_{i=1}^{N}(\lambda_i - \lambda_i^*) \cdot \partial_i f^*(\lambda^*) \\ &= f^*(\lambda^*) + \sum_{i=1}^{N}(\lambda_i - \lambda_i^*) \cdot (\partial_i f^*(\lambda^*) + \phi(\mathcal{S}^*)^2) \\ &\geq f^*(\lambda^*) + (2\eta - \eta^2) \cdot \phi(\mathcal{S}^*)^2 \cdot \sum_{i=K+1}^{N} \lambda_i. \end{aligned} \tag{13}$$

The second line uses $\sum_{i=1}^{N}(\lambda_i - \lambda_i^*) = 0$ and the last inequality is due to (12).

Consider any division $(q_1, \dots, q_N)$ at time t. A straightforward refinement of Lemma 2 gives that whenever $f^*(\lambda)$ is finite, $t \cdot f(\lambda t)$ approaches $f^*(\lambda)$ at the rate of $\frac{1}{t}$. In particular $f(\lambda^* \cdot t) = \frac{1}{t} \cdot f^*(\lambda^*) + O(\frac{1}{t^2})$. For $(q_1, \dots, q_N)$ to be a t-optimal division, it is necessary that $f(q_1, \dots, q_N) \leq f(\lambda^* \cdot t)$. Thus

$$f^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right) \leq f^*(\lambda^*) + O\left(\frac{1}{t}\right). \tag{14}$$

By (13) and (14), any t-optimal division $n(t)$ must satisfy $n_i(t) = O(1)$ for each signal $i > K$. Conditional on these signal counts, society's optimal choice over signals 1 through K must satisfy $n_i(t) = \lambda_i^* \cdot t + O(1), \forall 1 \leq i \leq K$, as shown in Proposition 1. This is what we desire to prove here.

A.2.6 A Perturbation Argument

First, we extend the definition of $\phi(\cdot)$ to arbitrary sets (not necessarily minimally-spanning) of signals 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.

[32]As mentioned in Remark 2, it is crucial that f^* is differentiable at λ^*. The argument here relies on the directional derivative in the direction $\lambda - \lambda^*$ being well-defined and equal to a linear sum of partial derivatives.

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

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

We now show this equality holds more generally.

Lemma 6. *For any coefficient matrix C,*

$$\min_{\lambda \in \Delta^{N-1}} f^*(\lambda) = \phi([N])^2. \tag{15}$$

Proof. Because society can choose to focus on any minimal spanning set, it is clear that $\min_\lambda f^*(\lambda) \leq \phi([N])^2 = \min_{\mathcal{S}}(\phi(\mathcal{S}))^2$. It remains to prove $f^*(\lambda) \geq \phi([N])^2$ for any fixed $\lambda \in \Delta^{N-1}$. 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 $m \to \infty$, the LHS above approaches $[(C'\Lambda C)^{-1}]_{11}$. Thus the lemma will follow once we show that $\limsup_{m\to\infty} \phi_m([N]) \geq \phi([N])$.

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 $\liminf_{m\to\infty} \phi_m([N]) \geq \phi([N])$, which is more than what we need. □

[33] First, we may add repetitive signals to ensure $N \geq K$. This does not affect the value of $\min f^*(\lambda)$ 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 the coefficient matrix $C^{(m)}$, different subsets $\mathcal{S}$ (of size K) attain different values of $\phi(\mathcal{S})$.

A.2.7 Proof of Theorem 1 when $|\mathcal{S}^*| < K$

We now complete the proof of Theorem 1 for the case where the "best" subset $\mathcal{S}^*$ contains fewer than K signals. To be precise, let $\mathcal{S}^* = \{1, \ldots, k\}$ and define $\lambda^* \in \Delta^{N-1}$ to be the optimal frequencies when only the first k signals are observed. We will show $n_i(t) \sim \lambda_i^* \cdot t, \forall i$. By Lemma 3, we only need to show that λ^* uniquely minimizes $f^*(\lambda)$ over the simplex. Since $f^*(\lambda^*) = \phi(\mathcal{S}^*)^2 = \phi([N])^2$ by definition, we know from Lemma 6 that λ^* does minimize $f^*(\lambda)$.

It remains to show that λ^* is the *unique* minimizer. Suppose for contradiction that $f^*(\lambda^*) = f^*(\tilde{\lambda})$ for some $\tilde{\lambda} \in \Delta^{N-1}$ distinct from λ^*. For $\eta \in \mathbb{R}$, define $\lambda^\eta = \lambda^* + \eta \cdot (\tilde{\lambda} - \lambda^*)$, so that $\lambda^0 = \lambda^*, \lambda^1 = \tilde{\lambda}$. Observe that when $\eta \in (0,1)$, λ^η is a convex combination between λ^* and $\tilde{\lambda}$. Thus the convexity of f^* implies

$$f^*(\lambda^\eta) \leq (1-\eta) f^*(\lambda^*) + \eta f^*(\tilde{\lambda}) = f^*(\lambda^*)$$

Since $f^*(\lambda^*)$ is minimal, we must then have $f^*(\lambda^\eta) = f^*(\lambda^*)$ for $\eta \in (0,1)$. But for fixed λ^* and λ, (7) shows that the value of $f^*(\lambda^\eta)$ is a rational function (quotient of two polynomials) of η. Thus this rational function is a constant function. Consequently, $f^*(\lambda^\eta) = f^*(\lambda^*)$ for all η (not just those in the unit interval) such that $\lambda^\eta \in \Delta^{N-1}$.

Because $\tilde{\lambda} \neq \lambda^*$, there exists some $j \in \{1, \ldots, k\}$ such that $\tilde{\lambda}_j < \lambda_j^*$. Without loss, we assume $\frac{\tilde{\lambda}_1}{\lambda_1^*}$ is the smallest among such ratios. Let $\eta = \frac{\lambda_1^*}{\lambda_1^* - \tilde{\lambda}_1}$, then the vector λ^η has first-coordinate 0 and all other coordinates non-negative. By our preceding analysis, $f^*(\lambda^\eta) = f^*(\lambda^*)$ for this η. However, since λ^η "ignores" signal 1, Lemma 6 implies that

$$f^*(\lambda^\eta) \geq \min_{\lambda \in \Delta^{N-1},\ \lambda_1 = 0} f^*(\lambda) = \phi([N] \backslash \{1\})^2.$$

By assumption, $\mathcal{S}^* = \{1, \ldots, k\}$ is the *unique* minimal spanning set that minimizes ϕ. Thus the RHS above is strictly larger than $\phi(\mathcal{S}^*)^2 = f^*(\lambda^*)$, leading to the contradictory result $f^*(\lambda^\eta) > f^*(\lambda^*)$.

This contradiction shows λ^* must uniquely minimize $f^*(\lambda)$. Theorem 1 follows.

A.3 Proof of Theorem 2 Part (a)

Let signals $1, \ldots, k$ (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 $\tilde{\theta}_1, \ldots, \tilde{\theta}_K$

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 $w_1\tilde{\theta}_1 + \cdots + w_k\tilde{\theta}_k$. We may without loss assume the weights w_i are all positive. Moreover, 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 $w_1\tilde{\theta}_1 + \cdots + w_k\tilde{\theta}_k$ by any signal $\sum_{i=1}^{k} \alpha_i\tilde{\theta}_i + \mathcal{N}(0,1)$ is

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

If $v_1, \dots, v_k$ are small positive numbers and if the product $w_i v_i$ is approximately constant across $1 \leq i \leq k$, then the above is approximately $(\sum_{i=1}^{k} \alpha_i)^2 w_1^2 v_1^2$. 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 $\frac{w_1^2 v_1^2}{1+v_1}$.

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

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

As $\beta_{k+1}, \dots, \beta_K$ are not all zero, the denominator above can be 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 3.

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. $v_1, \ldots, v_k$ are close to 0;

2. $w_1 v_1, \ldots, w_k v_k$ have pairwise ratios close to 1;

3. $v_{k+1}, \ldots, v_K$ are large.[34]

To show that *every agent* chooses among the first k signals, it suffices to check that starting from any prior beliefs 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 $\tilde{\theta}_1, \ldots, \tilde{\theta}_k$ does not affect the variances of the remaining states. So $v_{k+1}, \ldots, v_K$ are unchanged, and the third condition is verified. Finally, the second condition holds for the posterior beliefs because the signal i that is chosen has the greatest value of $\frac{w_i^2 v_i^2}{1+v_i}$. This choice ensures that $v_i \propto \frac{1}{w_i}$, as shown also in Liang, Mu and Syrgkanis (2017). Hence Theorem 2 is proved.

Strictly 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 $\tilde{\theta}_1, \ldots, \tilde{\theta}_K$. But it can be shown that the argument extends to mild correlation across states. We omit the somewhat cumbersome details, which do not add any further intuition.

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 $\mathcal{A}$. Also 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.[35] By

[34]Formally, we require that for some $\xi > 0$, it holds that $v_1, \ldots, v_k < \xi$; $\max_{1 \le i \le k} w_i v_i \le (1+\xi) \cdot \min_{1 \le i \le k} w_i v_i$; and $v_{k+1}, \ldots, v_K > \frac{1}{\xi}$.

[35]To see this, let $s_1, \ldots, s_N$ denote the limit signal counts, where $s_i = \infty$ if and only if $i \in \mathcal{A}$. Then we need to find a signal j such that $f(s_j + 1, s_{-j}) < f(s_j, s_{-j})$. This is because if $f(s_j + 1, s_{-j}) =$

continuity, at sufficiently late periods, observing signal j would reduce the variance by a positive amount. This is a profitable deviation from observing some signal in $\mathcal{A}$, leading to a contradiction!

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 the long-run frequencies are positive precisely for the signals in $\mathcal{S}$. Ignoring the initial periods, it is without loss to assume that only signals in $\overline{\mathcal{A}}$ are available. It suffices to show that whenever the signals observed infinitely often *span that subspace*, agents eventually sample from the optimal subset $\mathcal{S}$. 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 Controlling the Derivatives

To study the posterior variance function f, it will be convenient to instead work with the homogenous function f^* we introduced in Lemma 2. We formalize this connection as follows:

Lemma 7. *Suppose that signals in $\mathcal{A}$ span $\mathbb{R}^K$. Then, as $q_i \to \infty$ for each $i \in \mathcal{A}$,*

$$f(q_1, \ldots, q_N) \sim \frac{1}{t} \cdot f^* \left(\frac{q_1}{t}, \ldots, \frac{q_N}{t} \right) \quad \text{with} \quad t = \sum_{i=1}^{N} q_i$$

The partial derivatives and second partial derivatives also satisfy the approximations

$$\partial_j f(q_1, \ldots, q_N) \sim \frac{1}{t^2} \cdot \partial_j f^* \left(\frac{q_1}{t}, \ldots, \frac{q_N}{t} \right)$$

$$\partial_{jj} f(q_1, \ldots, q_N) \sim \frac{1}{t^3} \cdot \partial_{jj} f^* \left(\frac{q_1}{t}, \ldots, \frac{q_N}{t} \right)$$

$f(s_j, s_{-j})$ for each j, then the partial derivatives of f at s are all zero. Since f is differentiable, this would imply all directional derivatives of f are also zero. By the convexity of f, $f(s)$ must achieve minimum value. But by assumption there exists a spanning set, so $f(q) = 0$ if $q_1, \ldots, q_N$ are all infinite. This contradicts $f(s) > 0$.

Proof. Recall that

$$f(q_1, \dots, q_N) = \left[((V^0)^{-1} + C'QC)^{-1}\right]_{11}.$$

Since $q_i \to \infty$ for $i \in \mathcal{A}$, the least eigenvalue of the matrix $C'QC$ approaches infinity. That is, for any $\epsilon > 0$, it holds eventually that $(V^0)^{-1} \preceq \epsilon \cdot C'QC$ in matrix order. Then

$$\frac{1}{1+\epsilon} \cdot [(C'QC)^{-1}]_{11} \leq f(q_1, \dots, q_N) \leq [(C'QC)^{-1}]_{11}.$$

Equivalently, this shows

$$\frac{1}{(1+\epsilon)t} \cdot f^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right) \leq f(q_1, \dots, q_N) \leq \frac{1}{t} \cdot f^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right).$$

A similar approximation holds for the derivatives, proving the lemma. □

Lemma 8. *For any $q_1, \dots, q_N$, we always have*

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

And under the same assumptions as in Lemma 7, it holds that

$$\frac{\partial_{jj} f^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right)}{t \cdot \partial_j f^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right)} \to 0$$

Proof. It suffices to prove the first result. From $f(q_1, \dots, q_N) = e_1' \cdot [(V^0)^{-1} + C'QC]^{-1} \cdot e_1$ we compute that

$$\partial_j f = -e_1' \cdot [(V^0)^{-1} + C'QC]^{-1} \cdot c_j \cdot c_j' \cdot [(V^0)^{-1} + C'QC]^{-1} \cdot e_1$$

and

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

Let $\gamma_j = e_1' \cdot [(V^0)^{-1} + C'QC]^{-1} \cdot c_j$, which is a number. Then the above becomes

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

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

The above lemmata imply that at sufficiently late periods along society's signal path, the variance reduction of any *discrete* signal can be approximated by the continuous partial derivative of f (or f^*). A direct corollary is the following:

Lemma 9. *For any $\epsilon > 0$, there exists sufficiently large $t(\epsilon)$ such that if signal j is observed in any period $t+1$ later than $t(\epsilon)$, then*

$$\partial_j f^* \left(\frac{m(t)}{t} \right) \leq (1-\epsilon) \min_{1 \leq l \leq N} \partial_l f^* \left(\frac{m(t)}{t} \right).$$

That is, the signal choice in any sufficiently late period *almost* minimizes the directional derivative of f^*.

A.4.3 (Pseudo) Gradient Descent of f^*

We define $\lambda(t) = \frac{m(t)}{t} \in \Delta^{N-1}$. If j is the signal choice in period $t+1$, then it is easily checked that

$$\lambda(t+1) = \frac{t}{t+1}\lambda(t) + \frac{1}{t+1}e_j.$$

The frequencies $\lambda(t)$ move in the direction of e_j, which is the direction where f^* decreases almost the fastest (by Lemma 9). Thus, the evolution of $\lambda(t)$ over time resembles gradient descent dynamics—the value of $f^*(\lambda(t))$ roughly decreases over time, and we can expect that eventually $\lambda(t)$ approaches the unique minimizer λ^* of f^*.

To formalize this intuition, we consider (for fixed $\epsilon > 0$ and sufficiently large t)

$$\begin{aligned} f^*(\lambda(t+1)) &= f^* \left(\frac{t}{t+1}\lambda(t) + \frac{1}{t+1}e_j \right) \\ &= f^* \left(\frac{t}{t+1}\lambda(t) \right) + \frac{1}{t+1} \cdot \partial_j f^* \left(\frac{t}{t+1}\lambda(t) \right) \\ &\qquad + O\left(\frac{1}{(t+1)^2} \cdot \partial_{jj} f^* \left(\frac{t}{t+1}\lambda(t) \right) \right) \\ &\leq f^* \left(\frac{t}{t+1}\lambda(t) \right) + \frac{1-\epsilon}{t+1} \cdot \partial_j f^* \left(\frac{t}{t+1}\lambda(t) \right) \\ &= \frac{t+1}{t} \cdot f^*(\lambda(t)) + \frac{(1-\epsilon)(t+1)}{t^2} \cdot \partial_j f^*(\lambda(t)) \\ &\leq f^*(\lambda(t)) + \frac{1}{t} \cdot f^*(\lambda(t)) + \frac{1-2\epsilon}{t} \cdot \min_{1 \leq l \leq N} \partial_l f^*(\lambda(t)). \end{aligned} \tag{16}$$

The first inequality uses Lemma 8, the next equality uses the homogeneity of f^*, and the last inequality uses Lemma 9.

Write $\lambda = \lambda(t)$ for short. Observe that f^* is differentiable at λ, since $\lambda_i(t) > 0$ for $i \in \mathcal{A}$, which spans the entire space. Thus the convexity of f^* yields

$$f^*(\lambda^*) \geq f^*(\lambda) + \sum_{j=1}^{N} (\lambda_j^* - \lambda_j) \cdot \partial_j f^*(\lambda).$$

The homogeneity of f^* implies $\sum_{j=1}^N \lambda_j \cdot \partial_j f^*(\lambda) = -f^*(\lambda)$. This enables us to rewrite the above display as

$$\sum_{j=1}^{N} \lambda_j^* \cdot \partial_j f^*(\lambda) \leq f^*(\lambda^*) - 2f^*(\lambda).$$

Thus, in particular,

$$\min_{1 \leq l \leq N} \partial_l f^*(\lambda(t)) \leq f^*(\lambda^*) - 2f^*(\lambda). \tag{17}$$

Combining (16) and (17), we have for all large t:

$$f^*(\lambda(t+1)) \leq f^*(\lambda(t)) + \frac{1}{t} \cdot [(1-2\epsilon) \cdot f^*(\lambda^*) - (1-4\epsilon) \cdot f^*(\lambda(t))]. \tag{18}$$

We claim this implies $f^*(\lambda(t)) \leq (1+4\epsilon) \cdot f^*(\lambda^*)$ holds for all large t. Indeed, if this holds for *some* t, then (18) implies the same is true at future periods. It thus suffices to show the opposite inequality $f^*(\lambda(t)) > (1+4\epsilon) \cdot f^*(\lambda^*)$ cannot hold at every large t. By (18), that would give $f^*(\lambda(t+1)) \leq f^*(\lambda(t)) - \frac{\epsilon \cdot f^*(\lambda^*)}{t}$. But since the harmonic series diverges, $f^*(\lambda(t))$ would then decrease without bound, leading to a contradiction!

Hence we have shown that for any fixed ϵ, $f^*(\lambda(t)) \leq (1+4\epsilon) \cdot f^*(\lambda^*)$ holds eventually. As λ^* is the unique minimizer of f^*, this implies $\lambda(t) \to \lambda^*$. Theorem 2 Part (b) follows.

Remark 3. *The above argument leaves open the possibility that some signals outside of $\mathcal{S}^*$ are observed* infinitely often*, yet with* zero long-run frequency*. We conjecture this cannot happen, but we are only able to show this when $|\mathcal{S}^*| = K$.*

Specifically, suppose $|\mathcal{S}^*| = K$ and $m_i(t) \sim \lambda_i^* \cdot t, \forall i$, then we claim that the stronger conclusion $m_i(t) = \lambda_i^* \cdot t + O(1)$ also holds.[36] *Together with the conclusions of Appendix A.2.5, this suggests that the difference between $m_i(t)$ and the optimal $n_i(t)$ remains bounded.*

To prove this claim, we assume without loss that $\mathcal{S}^* = \{1, \ldots, K\}$ is the first K signals. By the previously established (**), the first K partial derivatives of f^* are equal at λ^* and they are strictly smaller (i.e., more negative) than the other partial derivatives. Since these partial derivatives are continuous, we can find $\epsilon > 0$ such that whenever λ is within ϵ distance from λ^*, it holds that

$$\partial_i f^*(\lambda) < (1+\epsilon) \cdot \partial_j f^*(\lambda), \quad \forall 1 \leq i \leq K < j$$

[36]Thus, the conclusion of Corollary 3 can be strengthened.

By assumption we have $\lambda(t) = \frac{m(t)}{t} \to \lambda^*$. Thus at sufficiently late periods, Lemma 9 implies that the signal choice must be within the first K signals. This shows signals outside of $\mathcal{S}^*$ are observed finitely often, as desired. And for any signal i in $\mathcal{S}^*$, its signal count satisfies $m_i(t) = \lambda_i^* \cdot t + O(1)$ by Proposition 1.

A.5 Proof of Proposition 2

Without loss of generality we assume that the best set consists of the first k signals. 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 $i \in \{1, \dots, k\}$, and samples arbitrarily if there is any capacity remaining. $\lfloor \ \rfloor$ 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. Adding up the precisions of these i.i.d. draws, the agent essentially obtains the following signal:

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

This is in turn equivalent to

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

Since the agent acquires such a signal for each $1 \le i \le k$, he receives at least as much information as the sum of these signals:

$$\begin{aligned}\sum_{i=1}^{k} \lambda_i^* Y_i &= \sum_{i=1}^{k} \langle \lambda_i^* c_i, \theta \rangle + \mathcal{N}\left(0, \frac{\sum_{i=1}^{k} \lambda_i^*}{(1-\epsilon)B}\right) \\ &= \frac{\omega}{\phi(\mathcal{S}^*)} + \mathcal{N}\left(0, \frac{1}{(1-\epsilon)B}\right)\end{aligned}$$

where we use $\lambda_i^* = \frac{\beta_i^*}{\sum_{j=1}^{k} \beta_j^*}$, $\sum_{i=1}^{k} \beta_i^* c_i = e_1$ and the fact that signal error terms are independent from each other.

If the agent follows *this* strategy, then his posterior precision (which is the inverse of his posterior variance) about ω must increase by at least $\frac{(1-\epsilon)B}{\phi(\mathcal{S}^*)^2}$. Clearly the same conclusion holds for his optimal strategy. Thus, we have shown that for any ϵ and sufficiently large B, each agent improves the posterior precision by $\frac{(1-\epsilon)B}{\phi(\mathcal{S}^*)^2}$. Hence society's posterior variance at time t is bounded above by $\frac{\phi(\mathcal{S}^*)^2}{(1-\epsilon)Bt}$.

Note that by choosing ϵ very small in the first place, society's speed of learning $\frac{\phi(\mathcal{S}^*)^2}{(1-\epsilon)}$ can be made arbitrarily close to the optimal speed $\phi(\mathcal{S}^*)^2$. Meanwhile, Part (b) of Theorem 2 extends to the current setting (with the same proof), and it implies that society eventually focuses on *some* minimal spanning set. Under Unique Minimizer, the only minimal spanning set that can approximate the optimal speed of learning is the best set. This yields the proposition.[37]

A.6 Proof of Proposition 3

Suppose without loss of generality that the best set is $\{1, \ldots, k\}$. 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 $b_1, \ldots, b_{k-1}$. We will show that whenever $k-1$ sufficiently precise signals are provided about each of these confounding terms, long-run frequencies must converge to λ^*.

Let D be the $K \times K$ diagonal matrix whose (j, j) entry is some large number L (to be chosen) for $2 \leq j \leq k$ and is 0 otherwise. Then, providing precise signals about $b_1, \ldots, b_{k-1}$ is equivalent to adding D to the prior precision matrix (inverse of the prior covariance matrix). Hence, we just need to show that if agents' prior precision matrix $(V^0)^{-1}$ is larger than D in matrix norm, long-run frequencies will be λ^*. This restatement of the proposition allows us to think about the problem in terms of prior beliefs rather than free signals.

We now perform two further restatements to simplify the analysis. First, recall that we showed society must eventually learn ω. This means at some late period t, the posterior precision matrix $(V^t)^{-1}$ is larger than $L \cdot \Delta_{11}$. Since the posterior precision matrix is also larger than the prior precision matrix, it must actually be larger than $\hat{D} = D + L \cdot \Delta_{11}$, which is also the diagonal matrix whose first k diagonal entries are L. Ignoring the first t periods, we therefore only need to show the result assuming that the prior precision matrix exceeds $\hat{D}$.

For the second simplification, note that for fixed M and sufficiently large L, we have $\hat{D} \succeq M \sum_{i=1}^{k} c_i c_i'$ because each c_i only involves the first k states. Because of this, we can also interpret $\hat{D}$ as the posterior precision matrix (given some prior) after M observations of each signal $1 \sim k$. Hence, the following is what we need to prove:

[37]This 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.

Proposition 3 Restated. *Under Unique Minimizer, suppose that at $t = 0$, agents commonly observe M realizations of each signal in the best set. Then when M is sufficiently large, society's subsequent acquisitions converge to the optimal frequencies.*

To prove this restated version, we note that each agent can choose from the first k signals and improve the posterior variance by

$$\max_{1 \le i \le k} \; |f(q_i + 1, q_{-i}) - f(q)|.$$

Using (the first half of) Lemma 8, we can relate the discrete partial derivative in the above expression to the usual continuous partial derivative:[38]

$$\frac{q_i}{q_i + 1} \cdot |\partial_i f(q)| \le |f(q_i + 1, q_{-i}) - f(q)|.$$

Thus each agent can improve the posterior variance by at least

$$\max_{1 \le i \le k} \frac{q_i}{q_i + 1} \cdot |\partial_i f(q)| \ge \frac{M}{M+1} \cdot \max_{1 \le i \le k} \; |\partial_i f(q)|,$$

where the inequality uses $q_1, \dots, q_k \ge M$ thanks to the free signals at $t = 0$.

Since f is always differentiable,[39] $\max_{1 \le i \le k} \; |\partial_i f(q)|$ is no less than the directional derivative in the direction λ^* (which is a convex combination of these partial derivatives). Next, we can use the same argument as in the proof of Proposition 2 to find this directional derivative. Specifically, that proof shows that along the direction λ^*, the instantaneous change of the posterior *precision* is at least $\frac{1}{\phi(\mathcal{S}^*)^2}$. Since posterior variance is the inverse of posterior precision, the chain rule of differentiation yields

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

Combining all of the above (and using the fact that each agent's optimal strategy cannot do worse), we conclude that at each time t,

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

[38] This follows from the more general result: $\frac{q_i x}{q_i + x} \cdot |\partial_i f(q)| \le |f(q_i + x, q_{-i}) - f(q)|$ for every $x \ge 0$, which becomes Lemma 8 after differentiating with respect to x twice.

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

To pin down the asymptotic behavior of f, we introduce the auxiliary function $g(t) = f(m(t)) \cdot \frac{M}{(M+1)\phi(\mathcal{S}^*)^2}$. Then the above simplifies to

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

Inverting both sides, we have

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)(1-g(t))} = \frac{1}{g(t)} + \frac{1}{1-g(t)} \geq \frac{1}{g(t)} + 1. \tag{19}$$

Thus by induction, $\frac{1}{g(t+1)} \geq t$ and so $g(t+1) \leq \frac{1}{t}$. Going back to the posterior variance function f, this implies

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

Therefore, by choosing M sufficiently large in the first place, we can ensure that society's speed of learning is arbitrarily close to the optimal speed $\phi(\mathcal{S}^*)^2$. Hence the long-run frequencies must be the optimal ones, proving the proposition.

B Online Appendix

B.1 δ-discounted Objective and Optimal Frequencies

In the main text, we introduced the optimal frequency vector λ^* by considering a sequence of finite-horizon problems. As an alternative, we consider in this appendix a social planner who aggregates payoffs across individuals using geometric discounting. We will show that under certain conditions, λ^* also emerges in the solution to the objective of maximizing δ-discounted payoffs, as $\delta \to 1$.

Formally, assume that all agents have the same payoff function $u(a, \omega)$,[40] and the social planner chooses signals and actions to maximize

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

For fixed δ, we denote by $d_\delta(t)$ the vector of signal counts (up to period t) associated with any strategy that maximizes U_δ. The following result characterizes the asymptotic behavior of $d_\delta(t)$ as $\delta \to 1$:

Proposition 4. *Suppose all agents have the same quadratic loss payoff function; that is, suppose $u(a, \omega) = (a - \omega)^2$. Then for any $\epsilon > 0$, there exists $\underline{\delta} < 1$ such that for any $\delta \geq \underline{\delta}$ it holds that*

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

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

Note that we have stated the result for a particular decision problem: Predicting the payoff-relevant state ω subject to quadratic penalty. Our proof below extends (with minor changes) to a class of "prediction" problems (e.g. including $u(a, \omega) = (a - \omega)^4$). On the other hand, going beyond prediction problems involves a number of technical challenges. While we conjecture that the conclusion of Proposition 4 holds more generally, verification of this conjecture is left for future work.[41]

[40] Assuming that agents have the same payoff function is more than necessary for our result below; however, restrictions on how the payoff functions differ are indeed required. Otherwise, suppose for example that payoffs are of the form $\alpha_t u(a, \omega)$, where α_t is much larger for earlier agents. Then even with the δ-discounted objective, the planner would care most to maximize the payoffs to the first agent, resembling the myopic incentive.

[41] When working with an arbitrary payoff function, the main difficulty lies in estimating the value

Proof of Proposition 4. Recall that λ^* uniquely minimizes the function f^*. Thus the proposition is equivalent to the following: Fix $\epsilon > 0$, then for any δ close to 1, $f(d_\delta(t)) \leq \frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t}$ holds for sufficiently large t. That is, we only need to show that as $\delta \to 1$, the achieved speed of learning is close to the optimal speed.

Suppose for contradiction that $f(d_\delta(t)) > \frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t}$ at some large t. Let $\tau < t$ be the last period with $f(d_\delta(\tau)) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}$. 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 $\mathrm{Var}(\tau + j)$ 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 (19)) gives that for $j > Mk$,

$$\frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \mathrm{Var}(\tau + j)} \geq \frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \mathrm{Var}(\tau + Mk)} + j - Mk, \tag{20}$$

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

$$\frac{(M+1)\phi(\mathcal{S}^*)^2}{M \cdot \mathrm{Var}(\tau + j)} \geq \frac{(M+1)\tau}{M(1+\epsilon/2)} + j - Mk. \tag{21}$$

of information. Specifically, in order to make intertemporal payoff comparisons, we need to know how much expected payoff is gained/lost when the posterior variance is decreased/increased by a certain amount. This can be challenging in general, see Chade and Schlee (2002) for related discussion. Another difficulty is that the optimal strategy for fixed δ may choose signals depending on previous signal realizations. That means $d_\delta(t)$ is in general a stochastic vector and much less tractable (beyond prediction problems).

With slight algebra, we obtain from the above

$$\frac{\mathrm{Var}(\tau+j)}{\phi(\mathcal{S}^*)^2} \leq \frac{1}{\frac{\tau}{1+\epsilon/2}+\frac{j-Mk}{1+1/M}}. \tag{22}$$

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, (22) implies

$$\mathrm{Var}(\tau+j) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau+j+1}, \quad \forall \underline{j}+1 \leq j \leq t-\tau. \tag{23}$$

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

$$\mathrm{Var}(\tau+j) \leq \mathrm{Var}(\tau) \leq \frac{(1+\epsilon/2)\phi(\mathcal{S}^*)^2}{\tau}, \quad \forall 1 \leq j \leq \underline{j}. \tag{24}$$

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 least[42]

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

Note that $\underline{j}$ has already been fixed. So as $\delta \to 1$ and $t-\tau \to \infty$,[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

[42] In this derivation we use (23), (24) as well as the identities $\frac{1}{\tau+j} - \frac{1}{\tau+j+1} = \frac{1}{(\tau+j)(\tau+j+1)}$ and $\frac{1}{\tau} - \frac{1}{\tau+j} = \frac{j}{\tau(\tau+j)}$.

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

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 $t-\tau \to \infty$ (since $\tau <= \frac{1+\epsilon/2}{1+\epsilon}t$). 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 (21), we use the following weaker inequality

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

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

$$\mathrm{Var}(\tau+j) \leq \frac{(1+\epsilon/4)\phi(\mathcal{S}^*)^2}{\tau+j}, \quad \forall \overline{j}+1 \leq j \leq t-\tau. \tag{25}$$

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

$$\mathrm{Var}(\tau+j) \leq c, \quad \forall 1 \leq j \leq \overline{j}. \tag{26}$$

Using the estimates (25) and (26) in place of (23) and (24), we find that the deviation strategy achieves payoff gain in periods $\tau+1 \sim t$ of at least

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

Importantly, because (25) improves upon (23), 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 $\delta \to 1$ and $t \to \infty$. Once again, this payoff gain dominates any potential loss after period t, showing that the deviation strategy is profitable. The proof of Proposition 4 is complete. □

B.2 Examples Failing Unique Minimizer

We provide two examples below of environments which do not satisfy Unique Minimizer. The conclusion of Theorem 1 fails to hold in the first example: We show that *all* signals are observed infinitely often. In the second example, the conclusion of Theorem 1 does extend. The difference in these two examples, and in addition the complexity of the derivation of the optimal frequencies, suggest that characterization of optimal acquisitions is generally difficult without the Unique Minimizer assumption.

B.2.1 First Example

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

$$\begin{aligned} X_1 &= \omega + b_1 + \epsilon_1 \\ X_2 &= b_1 + \epsilon_2 \\ X_3 &= \omega + b_2 + \epsilon_3 \\ X_4 &= b_2 + \epsilon_4 \end{aligned}$$

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 2 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

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

The optimal division vector thus seeks to *maximize*

$$\frac{1}{\frac{1}{q_1} + \frac{1}{q_2+\beta}} + \frac{1}{\frac{1}{q_3} + \frac{1}{q_4+\gamma}} \tag{27}$$

It is useful to rewrite (27) 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 parity of t and $r + s$ are the same, the optimal (q_1, q_2, q_3, q_4) satisfy*

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

and otherwise the 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 (27), 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 $\frac{t-r-s\mp 1}{2}$ 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. □

Although the example is involved, its intuition is simple: We would most like 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 continuously 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.2.2 Second Example

In the following example, Unique Minimizer is violated. However, the qualitative conclusion of Theorem 1 still holds. Namely, as $t \to \infty$, at most K signals are observed with positive frequency under the t-optimal division.

Consider state ω and confounding term b_1 (prior beliefs will be specified shortly). There are three signals $\omega+b_1+\epsilon_1, \omega-b_1+\epsilon_2$ and $\omega+\epsilon_3$, where each noise term is standard normal. We assume the prior beliefs are such that $\omega+b_1$ and $\omega-b_1$ are independent, with variances $\frac{1}{\alpha}$ and $\frac{1}{\beta}$. Observe that $\phi(\{1,2\}) = 1 = \phi(\{3\})$, so Unique Minimizer fails.

We claim that whenever $\alpha-\beta$ is not an integer, t-optimal divisions choose the third signal only a bounded number of times. Intuitively, this is because one observation of $\omega+b_1+\epsilon_1$ combined with one observation of $\omega-b_1+\epsilon_2$ contain at least as much information as their sum $2\omega+\epsilon_1+\epsilon_2$, which is equivalent to two observations of $\omega+\epsilon_3$.

Thus, devoting any level of attention to the third signal is *weakly* worse than splitting that attention evenly between the first two signals. Moreover, the combination of the first two signals also informs about b_1, which is correlated with the payoff-relevant state ω whenever $\alpha \neq \beta$. Thus, society optimally "ignores" the third signal if its (prior and posterior) beliefs about $\omega + b_1$ and $\omega - b_1$ are *asymmetric*. As we show below, this occurs precisely when $\alpha - \beta$ is not an integer.

To formalize the above intuition, we observe that given q_1 observations of signal 1 and q_2 observations of signal 2, society's posterior variance about ω is $\left(\frac{1}{q_1+\alpha} + \frac{1}{q_2+\beta}\right)/4$. Thus, with q_3 additional observations of the third signal, society's posterior variance becomes

$$f(q_1, q_2, q_3) = 1 \bigg/ \left(\frac{4}{\frac{1}{q_1+\alpha} + \frac{1}{q_2+\beta}} + q_3 \right).$$

The optimal division at time t thus maximizes

$$\max_{q_1,q_2,q_3 \in \mathbb{Z}^+, q_1+q_2+q_3=t} \frac{4}{\frac{1}{q_1+\alpha} + \frac{1}{q_2+\beta}} + q_3.$$

The maximand can be rewritten as

$$\frac{4}{\frac{1}{q_1+\alpha} + \frac{1}{q_2+\beta}} + q_3 = q_1 + \alpha + q_2 + \beta + q_3 - \frac{(q_1 + \alpha - q_2 - \beta)^2}{q_1 + \alpha + q_2 + \beta}.$$

Note that $q_1 + \alpha + q_2 + \beta + q_3 = t + \alpha + \beta$ is fixed, so society chooses q_1, q_2 to *minimize* the ratio $\frac{(q_1+\alpha-q_2-\beta)^2}{q_1+\alpha+q_2+\beta}$.

Suppose $\alpha - \beta$ is not an integer, let $\langle \alpha - \beta \rangle$ denote its distance to the nearest integer. Then, as q_1, q_2 are restricted to integers, the difference $|q_1 + \alpha - q_2 - \beta|$ takes minimum value $\langle \alpha - \beta \rangle > 0$. It follows that $\frac{(q_1+\alpha-q_2-\beta)^2}{q_1+a+q_2+b}$ is uniquely minimized by choosing q_1, q_2 such that $|q_1 + \alpha - q_2 - \beta| = \langle \alpha - \beta \rangle$ and $q_1 + q_2$ is as large as possible. Hence, both q_1 and q_2 are close to $\frac{t}{2}$. As we claimed, t-optimal division eventually focuses on the first two signals.

B.3 Supplementary Material to Section 7

Suppose the available signals are

$$X_1 = 10x + \epsilon_1$$
$$X_2 = 10y + \epsilon_2$$
$$X_3 = 4x + 5y + 10b$$
$$X_4 = 8x + 6y - 20b$$

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.

B.4 Multiple Payoff-Relevant States (Prediction Problem)

In the main text, we assumed that only one of the K persistent states is payoff-relevant. In this case, each individual agent simply chooses the signal that minimizes posterior variance about the payoff-relevant state. Consider a generalization in which agents seek to minimize a (weighted) *sum* of posterior variances about multiple payoff-relevant states.[44]

For ease of comparison, we still let θ denote the $K \times 1$ persistent state vector. Now, however, the first r states $\theta_1, \dots, \theta_r$ are all payoff-relevant ($1 \leq r \leq K$). In more detail, we assume that each agent's expected payoff is equal to the sum of his posterior variances about these r payoff-relevant states. This sum is a function of signal counts, which we still denote by $f(q_1, \dots, q_N)$. As before, define f^* to be a normalized, asymptotic version of f.

Let $n(t)$ continue to represent any allocation of t observations that minimizes f. Then, under a modification of the Unique Minimizer assumption—we require f^* to be *uniquely* minimized at λ^*—the optimal frequencies $\lambda^* := \lim_{t \to \infty} n(t)/t$ are well-defined. Nonetheless, we emphasize that with $r > 1$, these optimal allocations generally require more than K signals. A theorem of Chaloner (1984) shows that the optimal frequencies are 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

[44] This corresponds to a setting in which each agent seeks to predict the K unknown states and his payoff is determined via a weighted sum of quadratic losses.

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 frequencies are supported on $\mathcal{S}$. With these definitions, our main result (Theorem 2) and its proof generalize without modification.