

Complementary Information and Learning Traps^{*}

Annie Liang[†] Xiaosheng Mu[‡]

January 23, 2019

Abstract

We develop a model of social learning from complementary information: Short-lived agents sequentially choose from a large set of (flexibly correlated) information sources for prediction of an unknown state, and information is passed down across periods. Will the community collectively acquire the best kinds of information? Long-run outcomes fall into one of two cases: (1) efficient information aggregation, where the community eventually learns as fast as possible; (2) “learning traps,” where the community gets stuck observing suboptimal sources and information aggregation is inefficient. Our main results identify a simple property of the underlying informational complementarities that determines which occurs. In both regimes, we characterize which sources are observed in the long run and how often.

1 Introduction

Research is a central input into economic and social policies. As a society gains in its understanding of basic scientific facts, it can better tailor policy choices and investment decisions to the present world. The rate of accumulation of knowledge—as produced through research—is a key contributor to the rate of improvements in social welfare.

The process of research itself appears quite path-dependent—varying substantially across countries and time, for instance. This path-dependence arises due to interdependencies inherent to research: Past research often affects current researchers’ information and their

^{*}This paper was previously circulated under the title of “Overabundant Information and Learning Traps.” We are grateful to Nageeb Ali, Aislinn Bohren, Drew Fudenberg, Ben Golub, David Hirshleifer, Emir Kamenica, Paul Milgrom, Andrew Postlewaite, Ilya Segal, Carlos Segura, Rajiv Sethi, Andrzej Skrzypacz, and Yuichi Yamamoto for comments that improved this paper. We also thank Vasilis Syrgkanis for insightful comments in early conversations about this project.

[†]University of Pennsylvania

[‡]Cowles Foundation and Columbia University

incentives to undertake various investigations. Because research and its implications can be so consequential, it is important to understand the nature of this interdependence and how it affects the rate of learning. What helps a society efficiently refine its knowledge over time? What, in contrast, can account for research getting trapped in slower, less productive paths?

If research were centrally planned, then a planner’s directives would account for the externalities of researchers’ activities. But researchers—even in the academy and public sector—are often motivated by more immediate and individual goals, such as making the biggest possible discoveries now (rather than enabling others to make better discoveries later). So it is worth considering how the research process looks when these externalities are not internalized. In this paper, we develop a model to contrast the social planner’s research process with a decentralized one that results from the choices of short-lived researchers. We find that the structure of the aforementioned interdependencies—or, more precisely, *informational complementarities*—is crucial in determining whether the decentralized process is essentially efficient, or whether it gets trapped at a learning rate inferior to the optimal one.¹

In more detail, consider a sequence of researchers working on a single scientific problem, which we model as learning over time about the value of an unknown parameter. For a concrete example, we take the parameter to be the effect of carbon emissions on the global mean temperature a century from now. Each researcher can choose a study to learn more about this parameter. A study is modeled as a noisy Gaussian signal, whose realization is informative about the value of the parameter—for example, the researcher can measure the rise in the ocean’s temperature, or the shrinkage of Arctic ice sheets. We capture the limits on the researcher’s time, funding, and other resources by restricting him to choose one signal (study) from a finite but potentially large set. The researcher’s payoffs are determined by how well he can predict the parameter of interest—in the example, he seeks to maximally reduce uncertainty about the effect of carbon emissions on global warming.

The linkage between successive researchers comes from informational complementarities, which we now describe. In our model, each signal observation contains idiosyncratic noise but also some persistent confounds, which the researcher must learn in order to correctly de-bias the signal observation. Formally, each signal realization is a (noisy) combination of the parameter of interest, various confounding variables, and idiosyncratic noise. For example, measurements of ocean temperatures reflect the impact of carbon emissions, but also the impact of natural climate drivers (e.g., solar radiation and geological movements). A better understanding of these confounding variables increases the informativeness of ocean measurements about the parameter of interest (effect of carbon emissions on temperature);

¹Prior work studying the rate of learning include [Vives \(1992\)](#), [DeMarzo, Vayanos and Zwiebel \(2003\)](#), [Golub and Jackson \(2012\)](#), [Hann-Caruthers, Martynov and Tamuz \(2017\)](#), and [Harel et al. \(2018\)](#) among others.

in this sense, information about confounding variables, and observations of measurements confounded by those variables, are complementary.

Whether researchers end up acquiring information efficiently depends critically on how the available signals are related through these confounding variables. To describe the relevant relationships, we introduce a definition for *complementary sets of signals*, inspired by [Borgers, Hernando-Veciana and Krahmer \(2013\)](#). If the *smallest complementary set of signals is of size K* —where K is the number of unknown variables (including both the parameter of interest and all of the confounds)—then a decentralized research process nonetheless achieves efficient learning in the long run. In contrast, if the smallest complementary set of signals contains fewer than K sources, then early suboptimal information acquisitions can propagate across time, creating persistent inefficiencies in information gathering. We call these outcomes “learning traps.”

The key intuition is as follows. As [Sethi and Yildiz \(2016\)](#) point out, an agent who repeatedly observes a signal confounded by an unknown variable learns *both* about the payoff-relevant parameter and about the confounding variable. In the example above, measurements of ocean temperatures teach researchers about the impact of carbon emissions (the parameter of interest), but also about the impact of solar radiation. Thus, as researchers take more measurements of the ocean’s temperature, they learn more about the role of solar radiation, and are better able to interpret future ocean measurements.

In our setting, different signals may be confounded by the same variables. Thus, observations of a given signal not only help agents interpret future observations of signals of the same kind, but also potentially aid in the interpretation of *other* kinds of signals. We show that learning traps emerge exactly when the extent of these spillovers is limited—for example, if learning about the impact of solar radiation on ocean temperatures only helps with interpreting ocean measurements. Formally, this corresponds to our condition above regarding the existence of small complementary sets.

In contrast, suppose that solar radiation confounds many kinds of measurements—e.g., it also affects the shrinkage of ice caps. In this case, repeated measurements of ocean temperatures will allow future researchers to better interpret ocean measurements as well as measured shrinkage in ice mass. We show that with knowledge of all confounding variables that affect each of these two measurements, researchers can identify which of the two permits faster long-run learning. More generally, we show that if all complementary sets are of size K , then agents eventually learn all K unknown variables (including all confounds) and from this discover the best set of signals. Despite decentralized incentives, efficient information aggregation is achieved.

The main technical innovation behind these results is to establish a formal connection between society’s learning dynamics and a limiting dynamical system (see Section 6 for

details). We demonstrate that a set of signals is observed in the long-run (under some prior beliefs) if and only if it corresponds to a stable point in this system. This connection enables us to completely characterize the possible long-run observation sets.

We next study interventions for breaking learning traps. We show that policymakers can restore efficient information aggregation by providing sufficiently many kinds of free information about the relevant confounding variables (for example, a forward-looking funding agency can support research about confounding variables that are not of direct societal interest). Another successful intervention is to reshape the reward structure so that agents can be rewarded for information acquired over sufficiently many periods.

The final part of our paper considers the size of welfare losses associated with learning traps. We show that the rate of information aggregation can be arbitrarily slow (relative to the efficient benchmark), and payoffs can be arbitrarily inefficient when we consider the *ratio* of achieved and feasible payoffs in a patient limit. However, when the payoff-relevant parameter is persistent across time, agents eventually learn its value even while in a learning trap. Thus payoff losses are negligible when measured as the *difference* of achieved and feasible payoffs, both of which vanish in the patient limit. We next consider a generalization of the model where the unknown variables are changing over time.² We show that welfare losses in learning traps—measured either by the payoff ratio or payoff difference—can be arbitrarily large in nearby models with almost perfect persistence. This suggests that learning traps generally lead to long-run payoff inefficiencies.

1.1 Related Literature

The basic feature that information acquisitions are complementary appears in many settings besides research—for example, informational complementarities are relevant for auctions (Milgrom and Weber, 1982a,b), for firms (Athey and Schmutzler, 1995), for team composition (Chade and Eeckhout, 2018), and within markets (Goldstein and Yang, 2015; Chen and Waggoner, 2016). We build on this literature—which has primarily focused on *one-time* information acquisitions—by asking how informational complementarities affect learning in a *dynamic* setting.

This focus on dynamic implications connects to another body of work regarding investment in human capital (Jovanovic and Nyarko, 1996; Cunha and Heckman, 2007; Lizzeri and Siniscalchi, 2008). These papers typically consider complementarities in a production technology, in contrast to complementarities in information. At a high level certain effects are

²This is a technically challenging setting to analyze, and correspondingly prior work is limited: Jovanovic and Nyarko (1996), Moscarini, Ottaviani and Smith (1998), Frongillo, Schoenebeck and Tamuz (2011), Vivi Alatas and Olken (2016), and Dasaratha, Golub and Hak (2018) are the only social learning settings with a dynamic state that we are aware of.

similar: For example, early investment choices impact the incentives to invest in the future, potentially leading to worse long-run outcomes. And related to our Section 7, interventions can push agents onto better skill-acquisition paths. When the relevant complementarities are caused by informational correlations, the learning traps that emerge permit sharp characterizations, as we demonstrate.

The framework that we study is a social learning model, but we depart from the classic model (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992) in assuming all information to be public, thus turning off the inference problem essential to cascades. We instead allow for endogenous information acquisition—see Burguet and Vives (2000), Mueller-Frank and Pai (2016), and Ali (2018)—with the new feature that agents choose from a finite set of complementary signals.

Our paper is related to, but does not fall within, the classic multi-armed bandit framework (Gittins, 1979; Easley and Kiefer, 1988). The key difference is that in our model, agents’ choices of which signal to observe influence their learning, but do not directly influence their payoffs. Thus, agents engage in exploration but not exploitation, and our paper is what one might call a *pure learning* problem. This feature relates to a number of recent papers also studying dynamic information acquisition from different kinds of information—see for example Che and Mierendorff (2017) and Mayskaya (2017), who study learning from two Poisson signals, and Sethi and Yildiz (2016) and Fudenberg, Strack and Strzalecki (2018), who study learning from multiple Gaussian signals.

In earlier work (Liang, Mu and Syrgkanis, 2017), we studied dynamic learning from flexibly correlated Gaussian signals, but focused on a special case of the environment considered here, for which efficient information aggregation was guaranteed.³ This restriction allowed us to (under some conditions) fully characterize the optimal information acquisition path for any discount factor. But it precluded the welfare comparisons that are the focus here (indeed, learning traps could not emerge).

Finally, our setting is related to Sethi and Yildiz (2017): In that paper, people are sources of information, and they belong to disjoint communities with community-specific bias terms. Individuals choose whom to listen to over time. Sethi and Yildiz (2017) show that individuals can exhibit homophily in the long run, listening only to other individuals from the same community; this phenomenon is related to the observation of learning traps in the present paper.

³Formally, Liang, Mu and Syrgkanis (2017) assume that the set of available signals consists of a single complementary set (see Section 3).

2 Setup

Informational Environment. There are K persistent unknown states: a real-valued payoff-relevant state ω and $K - 1$ real-valued confounding variables b_1, \dots, b_{K-1} . We assume that the state vector $\theta := (\omega, b_1, \dots, b_{K-1})'$ follows a multivariate normal distribution $\mathcal{N}(\mu^0, \Sigma^0)$ where $\mu^0 \in \mathbb{R}^K$, and the prior covariance matrix Σ^0 has full rank.^{4,5}

There are N (fixed) *kinds* or *sources* of information available at each discrete period $t \in \mathbb{Z}_+$. Observation of source i in period t produces a realization of the random variable

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

where $c_i = (c_{i1}, \dots, c_{iK})'$ is a vector of constants, and the noise terms ϵ_i^t are independent from each other and across periods. Normalizing these noise terms to have unit variance is without loss of generality, since the coefficients c_i are unrestricted. We will often drop the time indices on the random variables, associating $X_i = \langle c_i, \theta \rangle + \epsilon_i$ with source i and understanding that the noise term is independently realized with each new observation.

The payoff-irrelevant states b_1, \dots, b_{K-1} produce correlations across the sources, and can be interpreted for example as:

- *Confounding explanatory variables:* Observation of source i produces the (random) outcome $y = \omega c_i^1 + b_1 c_i^2 + \dots b_{K-1} c_i^K + \epsilon_i$, which depends linearly on an observable characteristic vector c_i . For example, y might be the change in ocean temperature given physical measurements c_i . The state of interest ω is the coefficient on a given characteristic c_i^1 (i.e., increase in carbon emissions), and the payoff-irrelevant states are the unknown coefficients on the auxiliary characteristics c_i^2, \dots, c_i^K (i.e., change in solar radiation).
- *Knowledge and technologies that aid in information extraction:* Interpret the confounding variables as “disturbance” terms. For example, measurement of the rise in ocean temperatures may correspond to observations of the signal $X = \omega + b + \epsilon$, where the confounding variable b corresponds to a *systematic* bias in the sensors used to measure ocean temperatures. The difference between the term b and the term ϵ is that the former is persistent over time while the latter is i.i.d.—so the variance of b can be reduced over time, but the variance of ϵ is fixed. Separating these terms allows us to distinguish between reducible and irreducible noise in the signal.

⁴The full rank assumption is without loss of generality: If there is linear dependence across the states, the model can be reduced to a lower dimensional state space that satisfies full rank.

⁵Unless otherwise indicated, vectors in this paper are column vectors.

Decision Environment. Agents indexed by discrete time t move sequentially. Each agent chooses one of the N sources and observes a realization of the corresponding signal. He then predicts ω , selecting an action $a \in \mathbb{R}$ and receiving the payoff $-(a - \omega)^2$. We assume throughout that all signal realizations (but not payoffs) are public. Thus, 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, and his signal acquisition strategy is a function from histories to sources. The agent's optimal prediction of ω is his posterior mean, and his expected payoff is the negative of his posterior variance of ω . At every history h^{t-1} , the agent's expected payoffs are maximized by choosing the signal that minimizes his posterior variance of ω .

Since the environment is Gaussian, the posterior variance of ω is a deterministic function $V(q_1, \dots, q_N)$ of the number of times q_i that each signal i has been observed so far.⁶ Thus, each agent's signal acquisition is a function of past signal acquisitions only (and not of the signal realizations). This allows us to track society's acquisitions as deterministic *count vectors*

$$m(t) = (m_1(t), \dots, m_N(t))' \in \mathbb{Z}_+^N$$

where $m_i(t)$ is the number of times that signal i has been observed up to and including period t . The count vector $m(t)$ evolves according to the following rule: $m(0)$ is the zero vector, and for each time $t \geq 0$ there exists $i^* \in \operatorname{argmin}_i V(m_i(t) + 1, m_{-i}(t))$ such that

$$m_i(t+1) = \begin{cases} m_i(t) + 1 & \text{if } i = i^* \\ m_i(t) & \text{otherwise} \end{cases}$$

That is, the count vector increases by 1 in the coordinate corresponding to the signal that yields the greatest immediate reduction in posterior variance. We allow ties to be broken arbitrarily, and there may be multiple possible paths $m(t)$.

We are interested in the *long-run frequencies* of observation $\lim_{t \rightarrow \infty} \frac{m_i(t)}{t}$ for each source i —that is, the fraction of periods eventually devoted to each source. As we show later in Section 5, these limits exist under a weak technical assumption.

3 Complementary Information

In this section, we introduce a definition for *complementary sets* of sources.

Let $\tau(q_1, \dots, q_N) = 1/V(q_1, \dots, q_N)$ be the posterior precision about the payoff-relevant state ω given q_i observations of each source i , with $\tau_0 := \tau(0, \dots, 0)$ representing the prior precision. We define the *informational value* of a set of sources $\mathcal{S} \subseteq [N] := \{1, \dots, N\}$

⁶See Appendix A.1 for the complete closed-form expression for V .

to be the largest feasible improvement on precision (averaged across many periods), given observations of signals from \mathcal{S} alone.⁷

Definition 1. *The (asymptotic) informational value of the set \mathcal{S} is the maximal per-period increase in the precision about ω over a long horizon:*

$$\text{val}(\mathcal{S}) = \limsup_{t \rightarrow \infty} \left[\max_{q^t \in Q_{\mathcal{S}}^t} \left(\frac{\tau(q^t) - \tau_0}{t} \right) \right]$$

where

$$Q_{\mathcal{S}}^t = \left\{ q \in \mathbb{Z}_+^N : \sum_{i=1}^N q_i = t \text{ and } \text{supp}(q) \subset \mathcal{S} \right\}$$

is the set of all count vectors that allocate t observations across (only) the sources in \mathcal{S} .

We emphasize that the informational value is defined with respect to learning about ω . But since we fix the payoff-relevant state throughout this paper, we will omit the dependence of this value on ω . Separately, we highlight that a key property of the above definition is that informational value is *prior-independent* (see Proposition 1 below).

The following definition of a *complementary set* is based on [Borgers, Hernando-Veciana and Krahmer \(2013\)](#).

Definition 2. *The set \mathcal{S} is complementary if $\text{val}(\mathcal{S}) > \text{val}(\emptyset)$ and*

$$\text{val}(\mathcal{S}) - \text{val}(\mathcal{S} \setminus \mathcal{S}') > \text{val}(\mathcal{S}') - \text{val}(\emptyset)$$

for all nonempty proper subsets \mathcal{S}' of \mathcal{S} .

Informally, a set \mathcal{S} is complementary if the set is informative (the set's informational value exceeds that of the empty set), and the marginal value of having access to the sources in any $\mathcal{S}' \subset \mathcal{S}$ is enhanced by also having access to sources $\mathcal{S} \setminus \mathcal{S}'$. We note that the first condition that $\text{val}(\mathcal{S}) > \text{val}(\emptyset)$ is implied by the second condition whenever \mathcal{S} is not a singleton.⁸

Rewriting the condition in Definition 2 as

$$\frac{1}{2} \text{val}(\mathcal{S}) + \frac{1}{2} \text{val}(\emptyset) > \frac{1}{2} \text{val}(\mathcal{S}') + \frac{1}{2} \text{val}(\mathcal{S} \setminus \mathcal{S}')$$

we can interpret the definition a second way. Suppose the set of available sources is to be determined by a lottery at time $t = 0$, after which the social planner acquires information

⁷This definition of informational value closely resembles the definition of the value of a team in [Chade and Eeckhout \(2018\)](#), although we consider belief precision instead of negative posterior variance. Using posterior variances in Definition 1 would yield a value given by $\text{val}(\mathcal{S}) = \limsup_{t \rightarrow \infty} \left[\max_{q^t \in Q_{\mathcal{S}}^t} (-tV(q^t)) \right]$. This together with Definition 2 would return a similar notion of complementarity, but would present the technical issue of evaluating $\infty - \infty$ since value as defined this way (is always negative and) can be $-\infty$.

⁸The first condition has bite when $|\mathcal{S}| = 1$. Specifically, it rules out all singleton sets \mathcal{S} , except those consisting of a single unbiased signal $X = c\omega + \epsilon$.

optimally from the (realized) available set of sources, valuing each set \mathcal{S} at $\text{val}(\mathcal{S})$. Compare the choice between: (1) access to all of the sources in \mathcal{S} with probability $1/2$, and otherwise no access to information; and (2) an equal probability of access to the sources in $\mathcal{S}' \subsetneq \mathcal{S}$ and access to the sources $\mathcal{S} \setminus \mathcal{S}'$. If (1) yields a higher value than (2) for all nonempty $\mathcal{S}' \subsetneq \mathcal{S}$, then we say that the set of sources \mathcal{S} is complementary. Thus, the condition above implies that there is extra value to having access to *all* of the sources in \mathcal{S} .

3.1 Discussion

Our definition and interpretations above mirror the constructions in [Borgers, Hernando-Veciana and Krahmer \(2013\)](#) for complementary pairs of signals, but differ in a few key aspects:

First, Definition 2 is for sets of signals, while [Borgers, Hernando-Veciana and Krahmer \(2013\)](#) focus on pairs. Our generalization to sets can be understood in either of two ways. On one hand, we might consider a set to be complementary if *all* of the pieces combine to enhance the whole. For example, the sources $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + b_2 + \epsilon_2$, and $X_3 = b_2 + \epsilon_3$ are complementary, since the presence of each is critical to enhancing the value of the others (ω can only be learned by observing all three sources).

Another interpretation of Definition 2 is that it requires each *pair of subsets* that partition the whole set to be complementary. From this perspective, the relevant generalization is not from “two complementary signals” to “many complementary signals,” but rather from “two complementary signals” to “two complementary sets.” Given the signals X_1, X_2, X_3 defined above, we might say that access to the set $\{X_1, X_2\}$ is complementary to access to $\{X_3\}$, and likewise for the other combinations.⁹

Another difference from [Borgers, Hernando-Veciana and Krahmer \(2013\)](#) is that we consider complementary *sources* as opposed to complementary *signal observations*. That is, our definition does not ask whether a *single observation* of some signal improves the marginal value of a single observation of another. Rather, we ask whether *access* to some source improves the marginal value of access to another, where the social planner can optimally allocate many observations across the available sources.

Finally, [Borgers, Hernando-Veciana and Krahmer \(2013\)](#) value information based on its contribution to decision problems, while the *val* function that we use is more statistical

⁹Our proposed Definition 2 is stated more explicitly in terms of this second perspective, although it turns out that complementary sets are also characterized by having an informational value strictly greater than all proper subsets, thus relating to the first perspective that each piece contributes to the whole. Formally, Proposition 1 and part (a) of Proposition 2 below imply that a set \mathcal{S} is complementary according to Definition 2 if and only if $\text{val}(\mathcal{S}) > \text{val}(\mathcal{S}')$ for every proper subset \mathcal{S}' of \mathcal{S} , including $\mathcal{S}' = \emptyset$.

in nature (based on asymptotic improvements to belief precision).¹⁰ Like the definition in [Borgers, Hernando-Veciana and Krahmer \(2013\)](#), however, our notion of complementarity does turn out to be uniform across prior beliefs.

3.2 Characterization

The following result characterizes complementary sets.

Proposition 1. *\mathcal{S} is a complementary set if and only if the first coordinate vector in \mathbb{R}^K admits a unique decomposition $(1, 0, \dots, 0)' = \sum_{i \in \mathcal{S}} \beta_i^{\mathcal{S}} \cdot c_i$, where all coefficients $\beta_i^{\mathcal{S}}$ are nonzero.*

Thus, a set \mathcal{S} is complementary if its signals uniquely combine to produce an unbiased signal about ω . This characterization allows us to easily identify complementary sets based on their signal coefficient vectors:

Example 1. The set $\{X_1, X_2, X_3\}$ consisting of signals $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + b_2 + \epsilon_2$, and $X_3 = b_2 + \epsilon_3$ is complementary. To see this, observe that $(1, 0, 0)' = c_1 - c_2 + c_3$ (where $c_1 = (1, 1, 0)'$ is the coefficient vector associated with X_1 , $c_2 = (0, 1, 1)'$ is the coefficient vector associated with X_2 , and $c_3 = (0, 0, 1)'$ is the coefficient vector associated with X_3). In contrast, the set of signals $\{X_4, X_5\}$ with $X_4 = \omega + \epsilon_1$ and $X_5 = 2\omega + \epsilon_2$ is not complementary, since many different linear combinations of c_4 and c_5 produce $(1, 0)$. The set $\{X_1, X_2, X_3, X_4\}$ is also not complementary, although it contains multiple complementary subsets.

The next claim characterizes the informational value of a complementary set.

Claim 1. *Let \mathcal{S} be a complementary set. Then, $\text{val}(\mathcal{S}) = 1 / (\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|)^2$, where $\beta_i^{\mathcal{S}}$ are the ones given in Proposition 1.*

More generally, $\text{val}(\mathcal{S})$ can be determined for an arbitrary set \mathcal{S} as follows: If \mathcal{S} contains at least one complementary subset, then its value is equal to the *highest value among its complementary subsets*; otherwise the value of \mathcal{S} is zero. This result will follow from Proposition 2 part (a) in the next section.

Throughout the paper, we assume that there is at least one complementary set, and also that complementary sets can be completely ordered based on their informational values.

Assumption 1. *There is at least one complementary set $\mathcal{S} \subseteq [N]$.*

Assumption 2. *Each complementary set has a distinct informational value; that is, $\text{val}(\mathcal{S}) \neq \text{val}(\mathcal{S}')$ for all complementary sets $\mathcal{S} \neq \mathcal{S}'$.*

¹⁰Our definition is robust to monotone transformations of the val function; this follows from the characterization in Proposition 1 below.

The first assumption is without loss,¹¹ while the second assumption is *generically* satisfied. In particular, Assumption 2 implies the existence of a “best” complementary set, whose informational value is largest among complementary sets. This set plays an important role, and we will denote it by \mathcal{S}^* in the remainder of this paper.

4 Optimal Long-Run Observations

In this section, we show that optimal information acquisitions eventually concentrate on the best complementary set \mathcal{S}^* . Specifically, consider the distribution

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

which assigns zero frequency to signals outside of the best set \mathcal{S}^* , and samples sources within \mathcal{S}^* proportionally to the magnitude of $\beta_i^{\mathcal{S}^*}$. That is, each signal in \mathcal{S}^* receives frequency proportional to its contribution to an unbiased signal about ω , as defined in Proposition 1. The result below shows two senses in which λ^* is the optimal long-run frequency over signals.

Proposition 2. (a) Optimal Information Aggregation: $\text{val}([N]) = \text{val}(\mathcal{S}^*)$. Moreover, for any sequence $q(t)$ such that $\lim_{t \rightarrow \infty} \frac{\tau(q^t) - \tau_0}{t} = \text{val}([N])$, it holds that $\lim_{t \rightarrow \infty} \frac{q(t)}{t} = \lambda^*$.

(b) Social Planner Problem: For any δ , let $d_\delta(t)$ be the vector of signal counts (up to period t) associated with any signal path that maximizes the δ -discounted average payoff

$$U_\delta := \mathbb{E} \left[- \sum_{t=1}^{\infty} (1 - \delta) \delta^{t-1} \cdot (a_t - \omega)^2 \right]$$

Then there exists $\underline{\delta} < 1$ such that $\lim_{t \rightarrow \infty} \frac{d_\delta(t)}{t} = \lambda^*$ for every $\delta \geq \underline{\delta}$.

Part (a) says that the informational value of \mathcal{S}^* is the same as the informational value of the entire set of available sources. In this sense, having access to all available sources does not improve upon the speed of learning achievable from the best complementary set \mathcal{S}^* alone. Moreover, this speed of learning is attainable *only if* the long-run frequency over sources is

¹¹Our results extend to situations where ω cannot be identified from the available sources. 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 constant vector u . Let R be the subspace of \mathbb{R}^K spanned by $\tilde{c}_1, \dots, \tilde{c}_N$. Then project u onto R : $u = r + w$ with $r \in R$ and w orthogonal to R . Thus $\langle u, \tilde{\theta} \rangle = \langle r, \tilde{\theta} \rangle + \langle w, \tilde{\theta} \rangle$. By assumption, the random variable $\langle w, \tilde{\theta} \rangle$ is independent from any random variable $\langle c, \tilde{\theta} \rangle$ with $c \in R$ (because they have zero covariance). Thus the uncertainty about $\langle w, \tilde{\theta} \rangle$ cannot be reduced upon any signal observation. Consequently, agents only seek to learn about $\langle r, \tilde{\theta} \rangle$, returning to the case where the payoff-relevant parameter *is* identified.

the distribution λ^* .¹² Part (b) of Proposition 2 says that a (patient) social planner—who maximizes a discounted average of agent payoffs—will eventually observe sources in the proportions described by λ^* . Based on these results, we subsequently use λ^* as the optimal benchmark against which to compare society’s long-run information acquisitions.

5 Main Results

We now ask whether society’s acquisitions converge to the optimal long-run frequencies λ^* characterized above. We show that informational environments can be classified into two kinds—those for which efficient information aggregation is guaranteed (long-run frequencies are λ^*), and those for which “learning traps” are possible (agents exclusively observe some set of sources different from the efficient set \mathcal{S}^*). Separation of these two classes depends critically on the size of the smallest complementary set.

5.1 Learning Traps vs. Efficiency

The following example demonstrates that efficient information aggregation need not occur. Indeed, the set of signals observed in the long run can be disjoint from the efficient set \mathcal{S}^* .

Example 2. There are three available sources:

$$X_1 = 3\omega + b_1 + \epsilon_1$$

$$X_2 = b_1 + \epsilon_2$$

$$X_3 = \omega + \epsilon_3$$

Both $\{X_1, X_2\}$ and $\{X_3\}$ are complementary sets, but optimal information acquisitions (as defined in Section 4) eventually concentrate on $\mathcal{S}^* = \{X_1, X_2\}$.¹³

Now suppose that agents’ prior beliefs are such that ω and b_1 are independent, and the prior variance of b_1 is large (exceeds 8). In the first period, observation of X_3 is most informative about ω , since X_1 is perceived as a noisier signal about ω than X_3 , and X_2 provides information only about the confounding variable b_1 (which is uncorrelated with ω). Agent 1’s acquisition of X_3 does not update the variance of b_1 , so the same argument shows

¹²This result builds on Chaloner (1984), who shows that a “ c -optimal simultaneous experiment design” exists on at most K points. Part (a) additionally supplies a characterization of the optimal design itself and demonstrates uniqueness, with a minor technical difference that we impose an integer constraint on signal counts. We are not aware of prior work on the discounted payoff criterion studied in Part (b).

¹³It is straightforward to verify that $\text{val}(\{X_3\}) = 1 < 9/4 = \text{val}(\{X_1, X_2\})$. Note also that $X_1 - X_2$ is an unbiased signal about ω , and it is more informative than two realizations of X_3 . This demonstrates $\{X_1, X_2\}$ is the better complementary set without direct computation of informational values.

that agent 2 acquires X_3 . Iterating, we have that *every* agent observes signal X_3 . In this way, the set $\{X_3\}$ represents a learning trap.¹⁴

Returning to researchers who sequentially acquire information to learn about the impact of carbon emissions on climate change, we can interpret the source X_1 as projections of future temperatures using a new kind of climate model. Crucially, this model is biased due to misspecification—for example, by assuming a particular ecosystem response to increased temperatures—and the size and direction of this bias are unknown. The source X_2 corresponds to research into the actual feedback between climate and the ecosystem, which can be used to correct for misspecifications in the underlying model. The new model, if correctly de-biased, can lead to much improved projections, and socially it is optimal for researchers to invest towards this path. But if research into the climate-ecosystem feedback is slow, each researcher may choose instead to exploit existing models for measurement (observation of X_3), maximizing the marginal value of their work but reinforcing the learning trap.

Generalizing Example 2, the result below (stated as a corollary, since it will follow from the subsequent Theorem 1) gives a sufficient condition for learning traps. We impose the following generic assumption on the signal structure, which requires that every set of K signals are linearly independent:

Assumption 3 (Strong Linear Independence). *$N \geq K$ and every K signal coefficient vectors $c_{i_1}, c_{i_2}, \dots, c_{i_K}$ are linearly independent.*

Corollary 1. *Assume Strong Linear Independence. Then for every complementary set \mathcal{S} with $|\mathcal{S}| < K$, there exists an open set of prior beliefs given which agents exclusively observe signals from \mathcal{S} .*

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

Now suppose in contrast that no complementary sets are smaller than size K .¹⁵ Our next result shows that a very different long-run outcome obtains: Starting from *any* prior, society’s information acquisitions eventually approximate the optimal frequency. Thus, even though agents are short-lived (“myopic”), they end up acquiring information efficiently.

Corollary 2. *If there are no complementary sets with fewer than K sources, then starting from any prior belief, $\lim_{t \rightarrow \infty} m_i(t)/t = \lambda_i^*$ holds for every signal i . Thus, efficient information aggregation is guaranteed.*

¹⁴The existence of learning traps is not special to the assumption of normality. We report a related example with non-normal signals in Appendix D.2.

¹⁵It follows from Proposition 1 that there are no complementary sets with more than K sources, so this is equivalent to assuming that all complementary sets are of size K .

This result, like Corollary 1 above, follows from the subsequent Theorem 1. (We note that Strong Linear Independence is not assumed in Corollary 2, or in any subsequent results.)

An implication of these corollaries is that adding sources can worsen overall learning:

Example 3. Suppose the available sources are: $X_1 = 3\omega + b_1 + \epsilon_1$ and $X_2 = b_1 + \epsilon_2$. Then, $\{X_1, X_2\}$ is the only complementary set, with size $K = 2$. It follows from Corollary 2 that agents will achieve the efficient benchmark with a value of $\text{val}(\{X_1, X_2\}) = 9/4$. Now suppose we add $X_3 = \omega + \epsilon_3$, returning Example 2. The efficient benchmark does not change, but now there are priors that lead to exclusive observation of X_3 , achieving $\text{val}(\{X_3\}) = 1$.¹⁶

We provide a brief intuition for the above results, in particular for the importance of the number K . Recall that each agent chooses the signal whose marginal value—reduction of posterior variance of ω —is highest. Thus, if signal acquisitions eventually concentrate on a set \mathcal{S} , two properties must hold. First, all signals in the set must repeatedly have their turn as “most valuable,” so that agents do not focus on a strict subset of \mathcal{S} . This condition requires that the long-run outcome is a complementary set, where (by definition) all sources are critical to the value of the set.

Second, the marginal values of signals in that set must be persistently higher than marginal values of other signals. Not all complementary sets satisfy this criterion, since the sources within a given complementary set \mathcal{S} could have even stronger complementaries with sources outside of the set. If that were the case, observation of sources within \mathcal{S} would eventually push agents to acquire information outside of \mathcal{S} .

Indeed, a complementary set \mathcal{S} consisting of K sources *cannot* be self-reinforcing unless it is the best set: As observations accumulate from such a set, agents learn about all of the confounding variables and come to evaluate all sources according to “objective” (i.e. prior-independent) asymptotic values. Repeated acquisitions of signals from \mathcal{S} improve the value of signals in \mathcal{S}^* over the value of signals in \mathcal{S} . Thus, agents eventually turn to the sources in \mathcal{S}^* , achieving efficient information aggregation as predicted by Corollary 2.

In contrast, if agents observe only $k < K$ sources, then they can have persistent uncertainty about some confounding variables. This may cause society to persistently undervalue those sources confounded by these variables and to continually observe signals from a small complementary set. We saw this already in Example 2 where agents failed to obtain any information about the confounding variable b_1 , and thus persistently undervalued the sources X_1 and X_2 . The same intuition applies to Corollary 1.

One may argue that the condition that no complementary set has fewer than K sources is generically satisfied. However, if we expect that sources are endogenous to design or strategic

¹⁶In contrast, if the signal $X_4 = \omega + b_1 + \epsilon_4$ were added instead of X_3 in this example, then the complementary sets would be $\{X_1, X_2\}$, $\{X_1, X_4\}$ and $\{X_2, X_4\}$. By Corollary 2, (from all priors) agents necessarily observe the efficient set $\{X_1, X_2\}$ in the long run.

motivations, the relevant informational environments may not fall under this condition. For example, signals that partition into different groups with group-specific confounding variables (as studied in [Sethi and Yildiz \(2017\)](#)) are economically interesting but non-generic. Corollary 1 shows that inefficiency is a possible outcome in these cases.

5.2 General Characterization of Long-run Outcomes

Fixing an arbitrary signal structure, we now provide a complete characterization of the possible long-run observation sets as the prior belief varies. We introduce a new definition, which strengthens the notion of a complementary set:

Definition 3. \mathcal{S} is a strongly complementary set if it is complementary, and $\text{val}(\mathcal{S}) > \text{val}(\mathcal{S}')$ for all sets \mathcal{S}' that differ from \mathcal{S} in exactly one source ($|\mathcal{S} - \mathcal{S}'| = |\mathcal{S}' - \mathcal{S}| = 1$).¹⁷

The property of *strongly complementary* can be understood as requiring that the set is complementary and also something more: These complementarities are “locally best” in the sense that it is not possible to obtain stronger complementarities by swapping out just one source. We point out that while the definition of complementary sets does not depend on the ambient set (i.e., $[N]$) of available sources, the notion of strong complementarity does.

Example 4. Suppose the available signals are $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, and $X_3 = 2b_1 + \epsilon_3$. Then the set $\{X_1, X_2\}$ is complementary but not strongly complementary, as $\text{val}(\{X_1, X_3\}) > \text{val}(\{X_1, X_2\})$.

Our main result generalizes both the learning traps result and the efficient information aggregation result from the previous subsection. Theorem 1 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 strongly complementary.

Theorem 1. *The set \mathcal{S} is strongly complementary \iff there exists an open set of prior beliefs given which agents eventually exclusively observe signals from \mathcal{S} (that is, long-run frequencies exist and have support in \mathcal{S}).*¹⁸

When there is a single strongly complementary set, then all priors must lead to this set. Our previous Corollary 2 provides a sufficient condition that implies this, and moreover gives that the unique strongly complementary set is the *best* complementary set. When

¹⁷In fact, the requirement that \mathcal{S} is complementary is extraneous: One can show using Proposition 1 that if $\text{val}(\mathcal{S}) > \text{val}(\mathcal{S}')$ for all sets \mathcal{S}' differing from \mathcal{S} in exactly one source, then \mathcal{S} must be complementary.

¹⁸The “if” part of this statement can be strengthened as follows: The set \mathcal{S} is strongly complementary if there exists *any* prior belief given which agents eventually choose from \mathcal{S} . Thus, the regions of prior beliefs associated with different strongly complementary sets cover the whole space.

there are multiple strongly complementary sets, then different priors lead to different long-run outcomes, some of which are inefficient. Our previous Corollary 1 describes a sufficient condition for such multiplicity.

Obtaining a complete characterization of the sets of prior beliefs associated with different long-run outcomes is challenging, since society’s signal path may in general exhibit complex dynamics—for example, switching multiple times between different complementary sets (including the efficient set). This makes it difficult to relate the initial prior to eventual learning behavior. In the next section we discuss the technical details that go into the proof of Theorem 1. In particular, we will explain how we are able to determine the *range* of possible long-run outcomes despite incomplete knowledge about which occurs under a given prior.¹⁹

6 Proof Outline for Theorem 1

6.1 Limiting Dynamical System

We first introduce the following *normalized asymptotic* posterior variance function V^* , which takes frequency vectors $\lambda \in \Delta^{N-1}$ as input:

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

This function is convex in λ and its unique minimizer is the optimal frequency vector λ^* (see Lemma 7 in the appendix).

For simplicity of explanation, we will assume throughout this section that at large t , the signal choice that minimizes V also minimizes V^* .²⁰ Then, the frequency vector $\lambda(t) := \frac{m(t)}{t}$ evolves in the coordinate direction that minimizes V^* . We will refer to this as *coordinate descent*. Unlike the usual gradient descent, coordinate descent is restricted to move in coordinate directions. This restriction corresponds to our assumption that each agent can only acquire a discrete signal (rather than a mixture of signals).

One case where coordinate descent coincides with gradient descent is when V^* is everywhere differentiable: Differentiability ensures that all directional derivatives can be written

¹⁹In some cases, such as the signal structure in Example 2, a partial characterization is feasible. We showed previously that if ω and b_1 are independent under the prior, and the prior variance of b_1 exceeds 8, then every agent finds it optimal to observe X_3 , leading to a learning trap. If instead the prior variance of b_1 is smaller than 8, then we can show that every agent chooses from the efficient set $\{X_1, X_2\}$. See Appendix D.3 for the analysis as well as a related example.

²⁰This is not in fact generally correct, and the potential gap is one of the technical challenges in the proof. Nevertheless, we do show that at large t , the signal choice that minimizes V *approximately* minimizes V^* .

as convex combinations of partial derivatives along coordinate directions. In that case, evolution of $\lambda(t)$ would necessarily end at the global minimizer λ^* , implying efficient information aggregation.

6.2 Differentiability of V^*

The function V^* , however, is *not* guaranteed to be differentiable everywhere. Consider our learning traps example with signals $X_1 = 3\omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, $X_3 = \omega + \epsilon_3$, and fix the frequency vector to be $\lambda = (0, 0, 1)$. It is easy to verify that the asymptotic posterior variance $V^*(\lambda)$ is increased if we perturb λ by re-assigning weight from X_3 to X_1 , or from X_3 to X_2 . But V^* is reduced if we re-assign weight from X_3 to *both* X_1 and X_2 in an even manner.²¹ So V^* is not differentiable at λ . Such points of non-differentiability are exactly why learning traps are possible: Coordinate descent can become stuck at these λ , so that agents repeatedly sustain an inefficient frequency vector.²²

A sufficient condition for V^* to be differentiable at some frequency vector turns out to be that the signals receiving positive frequencies span all of \mathbb{R}^K . This explains the result in Corollary 2: When each complementary set consists of K signals, society *has to* observe K signals in order to learn the payoff-relevant state ω . Thus, driven by learning about ω , agents end up observing signals that span \mathbb{R}^K , which leads to efficient information aggregation.

6.3 Generalization to Arbitrary Subspaces

Now observe that our arguments above were not special to considering the whole space \mathbb{R}^K . If we restrict the available sources to some subset of $[N]$, and look at the subspace of \mathbb{R}^K spanned by these sources, our previous analysis applies to this restricted space.

Specifically, given any prior belief, define \mathcal{S} to be the set of sources that agents eventually observe. Let $\bar{\mathcal{S}}$ be the available signals that can be reproduced as a linear combination of signals from \mathcal{S} —these sources belong to the “subspace spanned by \mathcal{S} .” We can consider the restriction of the function V^* to all frequency vectors with support in $\bar{\mathcal{S}}$. Parallel to

²¹This follows from the formula $V^*(\lambda_1, \lambda_2, \lambda_3) = \lambda_3 + \frac{9}{1/\lambda_1 + 1/\lambda_2}$. The derivative of V^* in either direction $(1, 0, -1)$ or $(0, 1, -1)$ is positive, while its derivative in the direction $(\frac{1}{2}, \frac{1}{2}, -1)$ is in fact negative.

²²The above intuition connects to a literature on learning convergence in potential games (Monderer and Shapley, 1996; Sandholm, 2010). Define an N -player game where each player i chooses a number $\lambda_i \in \mathbb{R}_+$ and receives payoff $-\left(\sum_{j=1}^N \lambda_j\right) \cdot V^*(\lambda) = -V^*\left(\lambda / \sum_{j=1}^N \lambda_j\right)$. Then, we have a potential game with (exact) potential function $-V^*$, and our long-run observation sets correspond to equilibria of this game. This 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 observation of learning traps. Nonetheless, we note that the connection to potential games is not sufficient to derive our main results, since our agents receive payoff $-V(\lambda_t)$ rather than its asymptotic variant V^* .

the discussion above, the restricted version of V^* is both convex and differentiable in this subspace (at frequency vectors that assign positive weights to signals in \mathcal{S}). Thus, coordinate descent must lead to the minimizer of V^* in this subspace.

Recall that the overall optimal frequency vector λ^* is supported on the best complementary set \mathcal{S}^* . Analogously, the frequency vector that minimizes V^* in the restricted subspace is supported on the best complementary set within $\overline{\mathcal{S}}$. So our assumption that agents eventually concentrate signal acquisitions on the set \mathcal{S} is valid only if \mathcal{S} is *best in its subspace*; that is, $\text{val}(\mathcal{S}) = \text{val}(\overline{\mathcal{S}})$.

6.4 An Equivalence Result

The lemma below relates the property of “best in its subspace” to the notion of “strongly complementary.”

Lemma 1. *The following conditions are equivalent for a complementary set \mathcal{S} :*

- (a) $\text{val}(\mathcal{S}) = \text{val}(\overline{\mathcal{S}})$.
- (b) \mathcal{S} is strongly complementary.
- (c) For any $i \in \mathcal{S}$ and $j \notin \mathcal{S}$, $\partial_i V^*(\lambda^{\mathcal{S}}) < \partial_j V^*(\lambda^{\mathcal{S}})$, where $\lambda^{\mathcal{S}}$ (proportional to $|\beta^{\mathcal{S}}|$) is the optimal frequency vector supported on \mathcal{S} .

This lemma states that a strongly complementary set \mathcal{S} is “locally best” in three different senses. Part (a) says such a set has the highest informational value in its subspace. Part (b) says its informational value is higher than any set obtained by swapping out one source. Part (c) says that starting from the optimal sampling rule over \mathcal{S} , re-allocating frequencies from signals in \mathcal{S} to any other signal increases posterior variance and reduces speed of learning. The rest of this subsection is devoted to the proof of Lemma 1.

The implication from part (a) to part (b) is straightforward: Suppose \mathcal{S} is best in its subspace, and \mathcal{S}' is obtained from \mathcal{S} by removing signal i and adding signal j . Then the informational value of \mathcal{S}' is either zero, or equal to the highest value among its complementary subsets. In the latter case, such a complementary subset necessarily includes signal j , and Proposition 1 implies that j belongs to the subspace spanned by \mathcal{S} . Thus $\mathcal{S}' \subset \overline{\mathcal{S}}$, and $\text{val}(\mathcal{S}') \leq \text{val}(\overline{\mathcal{S}}) = \text{val}(\mathcal{S})$. The inequality is in fact strict, because complementary sets have different values (see Assumption 2).

We next show that part (b) implies part (c). Suppose part (c) fails, so some perturbation that shifts weight from source $i \in \mathcal{S}$ to source $j \notin \mathcal{S}$ decreases V^* . Then, by definition of informational value, we would have $\text{val}(\mathcal{S} \cup \{j\}) > \text{val}(\mathcal{S})$. But as Proposition 2 part (a) suggests, the value of $\mathcal{S} \cup \{j\}$ is equal to the highest value among its complementary subsets.

Strong complementarity of \mathcal{S} ensures that \mathcal{S} is the best complementary subset of $\mathcal{S} \cup \{j\}$. Thus we obtain $\text{val}(\mathcal{S} \cup \{j\}) = \text{val}(\mathcal{S})$, leading to a contradiction.

Finally, part (c) implies that $\lambda^{\mathcal{S}}$ is a *local* minimizer of V^* in the subspace spanned by \mathcal{S} (where the restriction of V^* is differentiable). Since V^* is convex, the frequency vector $\lambda^{\mathcal{S}}$ must in fact be a *global* minimizer of V^* in this subspace. Hence \mathcal{S} is best in its subspace and part (a) holds.

6.5 Completing the Argument

The arguments above tell us that information acquisitions eventually concentrate on a strongly complementary set, delivering one direction of Theorem 1: \mathcal{S} is a long-run outcome only if \mathcal{S} is strongly complementary.

To prove the “if” direction, we directly construct priors such that a given strongly complementary set \mathcal{S} is the long-run outcome. The construction generalizes the idea in Example 2, where we assign high uncertainty to those confounding variables that do not afflict signals in \mathcal{S} (as well as those in the same subspace $\bar{\mathcal{S}}$), and low uncertainty to those that do. This asymmetry guarantees that signals in $\bar{\mathcal{S}}$ have persistently higher marginal values than the remaining signals. Lastly, we use part (c) of the above Lemma 1 to show that agents focus on observing from \mathcal{S} , rather than the potentially larger set $\bar{\mathcal{S}}$. Indeed, if the historical frequency of acquisitions is close to $\lambda^{\mathcal{S}}$, then signals in \mathcal{S} have higher marginal values than the remaining signals in their subspace; and as these signals in \mathcal{S} continue to be chosen, society’s frequency vector remains close to $\lambda^{\mathcal{S}}$. This completes the proof of Theorem 1.

7 Interventions

The previous sections demonstrate the possibility for agents to persistently acquire suboptimal sources of information. This naturally suggests a question of what kinds of policies might free agents from these learning traps. We compare several possible policy interventions: increasing the *quality* of information acquisition (so that each signal realization is more informative); restructuring incentives so that agents’ payoffs are based on information obtained over several periods (equivalent to acquisition of *multiple signals* each period); and providing a one-shot release of *free information*, which can guide subsequent acquisitions.

7.1 More Precise Information

Consider first an intervention in which the precision of each signal draw is uniformly increased. We model this intervention by supposing that each signal acquisition now produces B independent observations from that source (with the main model corresponding to $B = 1$).

The result below shows that providing more informative signals is of limited effectiveness: Any set of signals that is a potential learning trap given $B = 1$ remains a potential learning trap under arbitrary improvements to signal precision.

Corollary 3. *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 these signals are exclusively viewed in the long run.*

This corollary follows directly from Theorem 1.²³

7.2 Batches of Signals

Another possibility is to restructure the incentive scheme so that agents' payoffs are based on information acquired from multiple signals. In practice, this might mean that payoffs are determined after a given time interval: For example, researchers may be evaluated based on a set of papers, so that they want to maximize the impact of the entire set. Alternatively, agents might be given the means to acquire multiple signals each period: For example, researchers may be arranged in labs, with a principal investigator directing the work of multiple individuals at once.

Formally, we suppose here that each agent can allocate B observations across the sources (where $B = 1$ returns the main model). Note the key difference from the previous intervention: It is now possible for the B observations to be allocated across *different* signals. This distinction enables agents to take advantage of the presence of complementarities, and we show that efficient information aggregation can be guaranteed in this case:

Proposition 3. *For sufficiently large B , if each agent acquires B signals every period, then society's long-run frequency vector is λ^* starting from every prior belief.*

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

The number of observations needed, however, depends on details of the informational environment. In particular, the required B cannot be bounded as a function of the number of states K and number of signals N .²⁴ See Appendix A.7.1 for further details.

²³However, the *set of* prior beliefs that yield \mathcal{S} as a long-run outcome need not be the same as B varies. For a fixed prior belief, subsidizing higher quality acquisitions may or may not move society out of a learning trap. To see this, consider first the signal structure and prior belief from Example 2. Increasing the precision of signals is ineffective there: As long as the prior variance on b is larger than 8, each agent still chooses signal X_3 regardless of signal precision. In Appendix D.4, we provide a contrasting example in which increasing the precision of signals indeed breaks agents out of a learning trap from a specified prior.

²⁴The required B depends on two properties: first, on how well the optimal frequency λ^* can be approximated by B (discrete) observations; second, on the difference in learning speed between the best set and the next best complementary set, which determines the slack that is permitted in the approximation of λ^* .

7.3 Free Information

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

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

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

Proposition 4. *Let $k := |\mathcal{S}^*|$. There exists a $\gamma < \infty$, and $k - 1$ signals $Y_j = \langle p_j, \theta \rangle + \mathcal{N}(0, 1)$ with $\|p_j\|_2 \leq \gamma$, such that with these free signals provided at $t = 0$, society's long-run frequency vector is λ^* starting from every prior belief.*

The proof is by construction. We show that as long as agents understand those confounding variables that appear in the best set of signals (these parameters have dimension $k - 1$), they will come to acquire information from this set.²⁵

We point out the following converse to Proposition 4: Whenever agents begin with sufficiently low prior uncertainty about the confounding variables that afflict the signals in \mathcal{S}^* , it is *impossible* for a malevolent third-party to provide free information and induce a learning trap as the long-run outcome.

8 Welfare Loss Under Learning Traps

We conclude with a more detailed analysis of the size of welfare losses under learning traps, and a generalization of our model in which the unknown states evolve over time.

8.1 Welfare Criteria

We focus on two classic welfare criteria: the speed of information aggregation and discounted average payoffs.

²⁵This intervention requires knowledge of the full correlation structure as well as which set \mathcal{S}^* is best. An alternative intervention, with higher demands on information provision but lower demands on knowledge of the environment, is to provide $K - 1$ (sufficiently precise) signals about all of the confounding variables.

Information Aggregation. A small modification of Example 2 shows that society’s long-run speed of learning can be arbitrarily slower than the optimal speed. Specifically, the informational value of the best complementary set can be arbitrarily large relative to the value of the set that agents eventually observe.

Example 5. There are three available sources: $X_1 = \omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, and $X_3 = \frac{1}{L}\omega + \epsilon_3$, where $L > 0$ is a constant. In this example, the ratio $\text{val}(\{X_1, X_2\}) / \text{val}(\{X_3\}) = \frac{L^2}{4}$ increases without bound as $L \rightarrow \infty$. But for every choice of L , there is a set of priors given which X_3 is exclusively observed.²⁶

Discounted Average Payoffs. Define

$$U_\delta^M = \mathbb{E}_M \left[- \sum_{t=1}^{\infty} (1 - \delta) \delta^{t-1} \cdot (a_t - \omega)^2 \right]$$

to be the δ -discounted average payoff across agents, who follow a “myopic” signal acquisition strategy with optimal predictions a_t . Also define U_δ^{SP} to be the maximum δ -discounted average payoff, where the social planner can use any signal acquisition strategy. Note that both payoff sums are negative, since flow payoffs are quadratic loss at every period.

Again from Example 5, we see that for every constant $c > 0$, there is a signal structure and prior belief such that the limiting payoff ratio satisfies²⁷

$$\liminf_{\delta \rightarrow 1} U_\delta^M / U_\delta^{SP} > c.$$

Thus, the *payoff ratio* can be arbitrarily large. Note that because payoffs are negative, larger values of the ratio $U_\delta^M / U_\delta^{SP}$ correspond to greater payoff inefficiencies.

On the other hand, the *payoff difference* vanishes in the patient limit; that is,

$$\lim_{\delta \rightarrow 1} (U_\delta^{SP} - U_\delta^M) = 0$$

in all environments. To see this, note that agents eventually learn ω even while in a learning trap, albeit slowly. Thus flow payoffs converge to zero at large periods, implying $\lim_{\delta \rightarrow 1} U_\delta^{SP} = \lim_{\delta \rightarrow 1} U_\delta^M = 0$.

In what follows, we show that this conclusion critically depends on the assumption that unknown states are perfectly persistent. We outline a sequence of autocorrelated models that converge to our main model (with perfect state persistence). At near perfect persistence, welfare losses under learning traps can be large according to all of the above measures.

²⁶The region of inefficient priors (that result in suboptimal learning) does decrease in size as the level of inefficiency increases. Specifically, as L increases, the prior variance of b_1 has to increase correspondingly in order for the first agent to choose X_3 .

²⁷Example 5 implies the ratio of flow payoffs at late periods can be arbitrarily large. As $\delta \rightarrow 1$, these later payoffs dominate the total payoffs from the initial periods (since the harmonic series diverges). So the ratio of aggregate discounted payoffs is also large.

8.2 Extension: Autocorrelated Model

In our main model, the state vector $\theta = (\omega, b_1, \dots, b_{K-1})'$ is persistent across time. We now consider a state vector θ^t that evolves according to the following law:

$$\theta^1 \sim \mathcal{N}(0, \Sigma^0); \quad \theta^{t+1} = \sqrt{\alpha} \cdot \theta^t + \sqrt{1 - \alpha} \cdot \eta^t, \quad \text{where } \eta^t \sim \mathcal{N}(0, M).$$

Above, means are normalized to zero, and the prior covariance matrix of the state vector at time $t = 1$ is Σ^0 . We restrict the autocorrelation coefficient $\sqrt{\alpha}$ to belong to $(0, 1)$. Choice of $\alpha = 1$ returns our main model, and we will be interested in approximations where α is close to but strictly less than 1. The *innovation* $\eta^t \sim \mathcal{N}(0, M)$ captures the additional noise terms that emerge under state evolution, which we assume to be i.i.d. across time.²⁸ Fixing signal coefficients $\{c_i\}$, every autocorrelated model is indexed by the triple (M, Σ^0, α) .

In each period, the available signals are

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

The signal noises ϵ_i^t are i.i.d. and further independent from the innovations in state evolution. The agent in period t chooses the signal that minimizes the posterior variance of ω^t , while the social planner seeks to minimize a discounted sum of such posterior variances.

We have the following result:

Theorem 2. *Suppose \mathcal{S} is strongly complementary. Then there exist M and Σ^0 such that for every $\epsilon > 0$, there is an $\underline{\alpha}(\epsilon) < 1$ with the following property: In each autocorrelated model (M, Σ^0, α) with $\alpha > \underline{\alpha}(\epsilon)$,*

1. *all agents only observe signals in \mathcal{S} ;*
2. *the resulting discounted average payoff satisfies*

$$\limsup_{\delta \rightarrow 1} U_\delta^M \leq -(1 - \epsilon) \cdot \sqrt{(1 - \alpha) \left(\frac{M_{11}}{\text{val}(\mathcal{S})} \right)},$$

while it is feasible to achieve a patient payoff of

$$\liminf_{\delta \rightarrow 1} U_\delta^{SP} \geq -(1 + \epsilon) \cdot \sqrt{(1 - \alpha) \left(\frac{M_{11}}{\text{val}(\mathcal{S}^*)} \right)}$$

by sampling from \mathcal{S}^ .*

²⁸The coefficient $\sqrt{1 - \alpha}$ in front of η^t is chosen so that when no signals are observed, society's posterior covariance matrix about θ^t will converge to M . This allows us to meaningfully consider the limit as $\alpha \rightarrow 1$ while keeping M fixed.

Part (1) generalizes Theorem 1, showing that every strongly complementary set is a potential long-run observation set given imperfect persistence. This suggests that the notion of strong complementarity and its importance extend beyond our main model.

Part (2) shows that whenever \mathcal{S} is different from the best complementary set \mathcal{S}^* , then social acquisitions result in significant payoff inefficiency as measured by the payoff ratio. Indeed, for α close to 1 the ratio $\lim_{\delta \rightarrow 1} U_\delta^M / U_\delta^{SP}$ is at least $\sqrt{\text{val}(\mathcal{S}^*) / \text{val}(\mathcal{S})}$, which can be arbitrarily large depending on the signal structure.

The following proposition strengthens this statement, using Example 5 to show that the payoff difference between optimal and social acquisitions can also be arbitrarily large:

Proposition 5. *For every $\epsilon > 0$, there exists a signal structure as in Example 5 and a corresponding autocorrelated model (M, Σ^0, α) such that*

$$\liminf_{\delta \rightarrow 1} U_\delta^{SP} \geq -\epsilon;$$

$$\limsup_{\delta \rightarrow 1} U_\delta^M \leq -\frac{1}{\epsilon}.$$

From this analysis, we take away that learning traps in general result in average payoff losses (and potentially large losses) so long as unknown states are not perfectly persistent over time.

9 Other Extensions

General Payoff Functions. Our main results extend when each agent t chooses an action to maximize an arbitrary individual payoff function $u_t(a_t, \omega)$ (we have so far restricted to $u_t(a_t, \omega) = -(a_t - \omega)^2$). Since in the Gaussian environment the signal that minimizes the posterior variance of ω Blackwell-dominates every other signal (Hansen and Torgersen, 1974), each agent’s signal acquisition remains unchanged.²⁹

However, the interpretation of the optimal benchmark (that we defined in Section 4) is more limited. Specifically, while the optimal frequency vector λ^* can still be interpreted as maximizing information revelation (part (a) of Proposition 2), the relationship to the social planner problem (part (b) of Proposition 2) may fail. See Appendix A.8 for further discussion.

²⁹To be fully rigorous, we also need a mild regularity condition on individual payoff functions to ensure that the signal minimizing the posterior variance yields *strictly* higher payoffs than other signals. The following condition is sufficient: For every t , any variance $\sigma^2 > 0$ and any action 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, we require that for every belief variance, the expected value of ω should affect the optimal action to take. This rules out cases where some actions are dominant regardless of how much is learned.

Low Altruism. We have assumed that each agent is short-lived and cares only to maximize the accuracy of his own prediction of the payoff-relevant state. Now suppose that agents are slightly altruistic; that is, each agent t chooses a signal as well as an action a_t to maximize discounted payoffs $\mathbb{E} [\sum_{t' \geq t} \rho^{t'-t} \cdot (a_{t'} - \omega)^2]$, assuming that future agents will behave similarly. ($\rho = 0$ returns our main model.) We show in Appendix D.5 that for ρ sufficiently small, part (a) of Theorem 1 continues to hold in every equilibrium of this game. So the existence of learning traps is robust to a small degree of altruism. By Proposition 7 in the appendix, part (b) of Theorem 1 also extends, showing that strongly complementary sets are the only possible long-run outcomes.

Multiple Payoff-Relevant States. In our main model, only one of the K persistent states is payoff-relevant. Consider a generalization in which each agent predicts (the same) $r \leq K$ unknown states and his payoff is determined via a weighted sum of quadratic losses. We show in Appendix D.6 that our main results extend to this setting.

10 Conclusion

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

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

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

We conclude below with brief mention of additional directions and interpretations of the model. First, we have focused on a sequence of decision-makers with a common prior. We might alternatively consider parallel sequences of decision-makers, each seeded with a different prior belief—for example, in the absence of a global research community, researchers in different countries may share different prior beliefs (and pass down their information within their country only). Using this setup, our results can be interpreted as answering the question: *Will individuals from different communities end up observing the same (best) set of sources, or will they persistently acquire information from different sources?* Our main results show that when there is a unique strongly complementary set of sources, then different priors wash out; otherwise, different priors can result in persistent differences across communities in what sources are listened to and consequently differences in beliefs.³⁰

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

³⁰In our main model with persistent states, beliefs about ω end up converging across the population to the truth. However, beliefs about other confounding variables need not converge. And when states are not fully persistent (as in Section 8.2), even beliefs about ω can diverge across communities.

A Proofs for the Main Model

The structure of the appendix follows that of the paper. In this appendix we provide proofs for the results in our main model, where states are perfectly persistent. The next appendix provides proofs for the autocorrelated model as discussed in Section 8.2. The only exception is that the proof of part (b) of Proposition 2 is more technical, so it is given in a separate Online Appendix, which also contains additional results and examples.

A.1 Preliminaries

A.1.1 Posterior Variance Function

Throughout, let C denote the $N \times K$ matrix of signal coefficients, whose i -th row is the vector c'_i associated with signal i . Here we review and extend a basic result from Liang, Mu and Syrgkanis (2017). Specifically, we show that the posterior variance of ω weakly decreases over time, and the marginal value of any signal decreases in its signal count.

Lemma 2. *Given prior covariance matrix Σ^0 and q_i observations of each signal i , society's posterior variance of ω is*

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

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

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

To prove that V is convex, it suffices to prove that V is *midpoint-convex* since the function is clearly continuous.³¹ Take $q_1, \dots, q_N, r_1, \dots, r_N \in \mathbb{R}_+$ and let $s_i = \frac{q_i + r_i}{2}$. Define the corresponding diagonal matrices to be Q, R, S . Note that $Q + R = 2S$. Thus by the AM-HM inequality for positive-definite matrices, we have

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

Using (1), we conclude that $V(q_1, \dots, q_N) + V(r_1, \dots, r_N) \geq 2V(s_1, \dots, s_N)$. This proves the (midpoint) convexity of V . \square

³¹A function V is midpoint-convex if the inequality $V(a) + V(b) \geq 2V(\frac{a+b}{2})$ always holds. Every continuous function that is midpoint-convex is also convex.

A.1.2 Inverse of Positive Semi-definite Matrices

For future use, we provide a definition of $[X^{-1}]_{11}$ for positive *semi-definite* matrices X . When X is positive definite, its eigenvalues are strictly positive, and its inverse matrix is defined as usual. In general, we can apply the Spectral Theorem to write $X = UDU'$, where U is a $K \times K$ orthogonal matrix whose columns are eigenvectors of X , and $D = \text{diag}(d_1, \dots, d_K)$ is a diagonal matrix consisting of non-negative eigenvalues. When these eigenvalues are strictly positive, we have

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

where u_j is the j -th column vector of U . In this case

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

is well-defined. Even if some d_j are zero, we can still use the RHS above to define $[X^{-1}]_{11}$, applying the convention that $\frac{0}{0} = 0$ and $\frac{z}{0} = \infty$ for any $z > 0$. Note that by this definition,

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

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

A.1.3 Asymptotic Posterior Variance

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

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

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

Lemma 3. *Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$. Then³²*

$$V^*(\lambda) = [(C' \Lambda C)^{-1}]_{11} \quad (3)$$

The value of $[(C' \Lambda C)^{-1}]_{11}$ is well-defined, see (2).

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

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

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

³²Note that $C' \Lambda C$ is the Fisher Information Matrix when signals are observed according to frequencies λ . So this lemma can also be seen as an application of the Bayesian Central Limit Theorem.

A.2 Key Object ϕ

We now define an object that will play a central role in the proofs. For each set of signals \mathcal{S} , consider writing the first coordinate vector $e_1 \in \mathbb{R}^K$ (corresponding to the payoff-relevant state ω) as a linear combination of signals in \mathcal{S} :

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

Definition 4. $\phi(\mathcal{S}) := \min_{\beta} \sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|$.

That is, $\phi(\mathcal{S})$ measures the size of the “smallest” (in the l_1 norm) linear combination of the signals in \mathcal{S} to produce an unbiased estimate of the payoff-relevant state. In case ω is not spanned by \mathcal{S} , this definition sets $\phi(\mathcal{S}) = \infty$.

When \mathcal{S} *minimally spans* ω (so that no subset spans), the coefficients $\beta_i^{\mathcal{S}}$ are unique and nonzero. In this case $\phi(\mathcal{S})$ is easy to compute. In general, we have the following characterization:

Lemma 4. *For any set \mathcal{S} that spans ω , $\phi(\mathcal{S}) = \min_{\mathcal{T} \subset \mathcal{S}} \phi(\mathcal{T})$ where the minimum is over subsets \mathcal{T} that “minimally span” ω .*

This lemma is a standard result linear programming, so we omit the proof. As a corollary, $\phi([N]) = \phi(\mathcal{S}^*)$ where \mathcal{S}^* is the set of signals that minimize ϕ among all sets that minimally span ω . The following proposition makes clear that this set \mathcal{S}^* is exactly the best complementary set defined in the main text.

Proposition 6. *For any set of signals \mathcal{S} , $\text{val}(\mathcal{S}) = \frac{1}{\phi(\mathcal{S})^2}$.*

[Liang, Mu and Syrgkanis \(2017\)](#) proves this proposition for sets \mathcal{S} that minimally span ω . We will prove this result in its general form shortly.

A.3 Proof of Proposition 1 and Claim 1

To see why Proposition 6 is useful, let us use it to show the results in Section 3 and 4. Indeed, Claim 1 directly follows from Proposition 6 and the definition of ϕ .

As for Proposition 1, take any set of signals \mathcal{S} . If these signals do not span ω , then Proposition 6 implies $\text{val}(\mathcal{S}) = 0$ and \mathcal{S} is not complementary by Definition 2. If a *proper* subset of \mathcal{S} spans ω , then Proposition 6 together with Lemma 4 implies that the informational value of \mathcal{S} is equal to the highest value among its subsets that minimally span ω . Let \mathcal{S}' denote this subset that achieves this highest value. For this \mathcal{S}' the inequality in Definition 2 is violated, and \mathcal{S} is again not complementary.

Finally, suppose \mathcal{S} itself minimally spans ω . In this case any nonempty proper subset of \mathcal{S} does not span ω and have zero informational value, whereas \mathcal{S} has positive value. So Definition 2 is satisfied and such sets \mathcal{S} are complementary, as described in Proposition 1.

A.4 Proof of Proposition 6 and Proposition 2 Part (a)

We will focus on proving $\text{val}([N]) = \frac{1}{\phi(\mathcal{S}^*)^2}$, which is thus equal to $\frac{1}{\phi([N])^2}$ by Lemma 4. Once this is proved, it is a direct generalization that $\text{val}(\mathcal{S}) = \frac{1}{\phi(\mathcal{S})^2}$ whenever \mathcal{S} spans ω . And in case \mathcal{S} does not span ω , the posterior variance of ω is bounded away from zero when agents are constrained to observe from \mathcal{S} . Thus $\tau(q^t)$ is bounded and $\text{val}(\mathcal{S}) = \limsup_{t \rightarrow \infty} \frac{\tau(q^t) - \tau_0}{t} = 0$, which is also equal to $\frac{1}{\phi(\mathcal{S})^2}$.

A.4.1 Reduction to Study of V^*

Consider the asymptotic variance function V^* introduced previously. We claim that $\text{val}([N]) = \frac{1}{\phi(\mathcal{S}^*)^2}$ will follow from the fact that λ^* is the (unique) frequency vector that minimizes V^* .

Lemma 5. *Suppose λ^* uniquely minimizes $V^*(\lambda)$ for $\lambda \in \Delta^{N-1}$. Then $\text{val}([N]) = \frac{1}{\phi(\mathcal{S}^*)^2}$, and Proposition 6 as well as part (a) of Proposition 2 holds.*

Proof. By definition, we always have $\text{val}([N]) \geq \text{val}(\mathcal{S}^*) = \frac{1}{\phi(\mathcal{S}^*)^2}$. In the opposite direction, take any sequence q^t with $t \rightarrow \infty$ with $\limsup_t \frac{\tau(q^t) - \tau_0}{t} = \text{val}([N])$. Since τ_0 is a constant, we equivalently have $\limsup_t \frac{\tau(q^t)}{t} = \text{val}([N])$, which gives

$$\liminf_{t \rightarrow \infty} t \cdot V(q^t) = \frac{1}{\text{val}([N])}$$

using the fact that the precision $\tau(q^t)$ is just the inverse of the variance $V(q^t)$.

By passing to a subsequence if necessary, we may assume the frequency vector $\lambda := \lim_{t \rightarrow \infty} \frac{q^t}{t}$ exists. Then by definition of V^* , the LHS of the above display is simply $V^*(\lambda)$. We therefore deduce $\text{val}([N]) = \frac{1}{V^*(\lambda)}$ for some $\lambda \in \Delta^{N-1}$. Since λ^* minimizes V^* , we conclude that $\text{val}([N]) \leq \frac{1}{V^*(\lambda^*)} = \frac{1}{\phi(\mathcal{S}^*)^2}$.

Combined with the earlier analysis, $\text{val}([N]) = \frac{1}{\phi(\mathcal{S}^*)^2}$ must hold with equality. Moreover, since λ^* is the unique minimizer, equality can only hold when $\lambda = \lambda^*$, so that $\frac{q^t}{t}$ necessarily converges to λ^* . This is what we desire to prove for Proposition 2. \square

A.4.2 Lemma on Linear Combinations

To show λ^* uniquely minimizes V^* , we need the following technical lemma.

Lemma 6. *Suppose \mathcal{S}^* (which uniquely minimizes ϕ) involves exactly K signals, and without loss let $\mathcal{S}^* = \{1, \dots, K\}$. Let C^* be the $K \times K$ submatrix of C corresponding to the first K signals. Further suppose $\beta_i^{\mathcal{S}^*} = [(C^*)^{-1}]_{1i}$ is positive for $1 \leq i \leq K$. Then for any signal $j > K$, if we write $c_j = \sum_{i=1}^K \alpha_i \cdot c_i$ (which is a unique representation), then $|\sum_{i=1}^K \alpha_i| < 1$.*

Proof. By assumption, we have the vector identity

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

Suppose for contradiction that $\sum_{i=1}^K \alpha_i \geq 1$ (the opposite case where the sum is ≤ -1 can be similarly treated). Then some α_i 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

$$e_1 = \sum_{i=1}^K \beta_i \cdot c_i = \left(\sum_{i=2}^K \left(\beta_i - \frac{\beta_1}{\alpha_1} \cdot \alpha_i \right) \cdot c_i \right) + \frac{\beta_1}{\alpha_1} \cdot \left(\sum_{i=1}^K \alpha_i \cdot c_i \right)$$

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

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

$$\sum_{i=2}^K \left(\beta_i - \frac{\beta_1}{\alpha_1} \cdot \alpha_i \right) + \frac{\beta_1}{\alpha_1} = \sum_{i=1}^K \beta_i + \frac{\beta_1}{\alpha_1} \cdot \left(1 - \sum_{i=1}^K \alpha_i \right) \leq \sum_{i=1}^K \beta_i.$$

But then $\phi(\{2, \dots, K, j\}) \leq \phi(\{1, 2, \dots, K\})$, contradicting the unique minimality of $\phi(\mathcal{S}^*)$. Hence the lemma must be true. \square

A.4.3 Case 1: $|\mathcal{S}^*| = K$

In this section, we prove that λ^* is indeed the unique minimizer of V^* whenever the set \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. As discussed, this result will imply Proposition 6 and part (a) of Proposition 2.

Lemma 7. *The function $V^*(\lambda)$ is uniquely minimized at $\lambda = \lambda^*$.*

Proof. First, we assume $\mathcal{S}^* = \{1, \dots, K\}$ and that $[(C^*)^{-1}]_{1i}$ is positive for $1 \leq i \leq K$. This is without loss because we can replace c_i with $-c_i$ without affecting the model.

Since $V(q_1, \dots, q_N)$ is convex in its arguments, $V^*(\lambda) = \lim_{t \rightarrow \infty} t \cdot V(\lambda_1 t, \dots, \lambda_N t)$ is also convex in λ . To show λ^* uniquely minimizes V^* , we only need to show λ^* is a *local minimum*. In other words, it suffices to show $V^*(\lambda^*) < V^*(\lambda)$ for any λ that belongs to an ϵ -neighborhood of λ^* . By definition, \mathcal{S}^* minimally spans ω and so its signals are linearly independent. Under the additional assumption that \mathcal{S}^* has size K , we deduce that its signals span the entire space \mathbb{R}^K . From this it follows that the $K \times K$ matrix $C' \Lambda^* C$ is positive definite, and by (3) the function V^* is differentiable near λ^* .

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

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

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

$$V^*(\lambda_1, \dots, \lambda_K, \lambda_{K+1}, \dots, \lambda_N) \geq V^* \left(\lambda_1, \dots, \lambda_{K-1}, \sum_{k=K}^N \lambda_k, 0, \dots, 0 \right) \geq V^*(\lambda^*). \quad (4)$$

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

To prove (*), we recall that

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

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

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

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

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

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

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

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

$$\partial_i V^*(\lambda^*) = -\gamma_i^2, \quad 1 \leq i \leq N. \quad (5)$$

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

$$\gamma_i = e_1' \cdot (C^*)^{-1} \cdot \text{diag}\left(\frac{1}{\lambda_1^*}, \dots, \frac{1}{\lambda_K^*}\right) \cdot e_i = \frac{[(C^*)^{-1}]_{1i}}{\lambda_i^*} = \phi(\mathcal{S}^*). \quad (6)$$

On the other hand, choosing any signal $j > K$, we can uniquely write the vector c_j as a linear combination of c_1, \dots, c_K . By Lemma 6,

$$\gamma_j = \langle v, c_j \rangle = \sum_{i=1}^K \alpha_i \cdot \langle v, c_i \rangle = \sum_{i=1}^K \alpha_i \cdot \gamma_i = \phi(\mathcal{S}^*) \cdot \sum_{i=1}^K \alpha_i, \quad (7)$$

where the last equality uses (6). Since $|\sum_{i=1}^K \alpha_i| < 1$, the absolute value of γ_j is strictly smaller than the absolute value of γ_K for any $j > K$. This together with (5) proves the desired inequality (*), and Lemma 7 follows. \square

A.4.4 A Perturbation Argument

To summarize, we have shown that when ϕ is uniquely minimized by a set \mathcal{S} containing exactly K signals,

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

We now use a perturbation argument to show this equality holds more generally.

Lemma 8. For any coefficient matrix C ,

$$\min_{\lambda \in \Delta^{N-1}} V^*(\lambda) = \phi([N])^2. \quad (8)$$

Proof. For general coefficient matrix C , the set \mathcal{S} that minimizes ϕ may not be unique or involve K signals. However, since society can choose to focus on \mathcal{S} , we always have

$$\min_{\lambda} V^*(\lambda) \leq V^*(\lambda^*) = \phi(\mathcal{S}^*)^2 = \phi([N])^2.$$

It remains to prove $V^*(\lambda) \geq \phi([N])^2$ for every $\lambda \in \Delta^{N-1}$. By Lemma 3, we need to show $[(C' \Lambda C)^{-1}]_{11} \geq \phi([N])^2$.

Note that we already proved this inequality for *generic* coefficient matrices C : specifically, those for which $\phi(\mathcal{S})$ is uniquely 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 analogue of ϕ for the coefficient matrix C_m .

As $m \rightarrow \infty$, the LHS above approaches $[(C' \Lambda C)^{-1}]_{11}$. We will show that on the RHS

$$\limsup_{m \rightarrow \infty} \phi_m([N]) \geq \phi([N]),$$

which then implies $[(C' \Lambda C)^{-1}]_{11} \geq \phi([N])^2$ and the lemma. Indeed, suppose $e_1 = \sum_i \beta_i^{(m)} \cdot c_i^{(m)}$ along the convergent sequence, then $e_1 = \sum_i \beta_i \cdot c_i$ for any limit point β of $\beta^{(m)}$. Using the definition of ϕ , this enables us to conclude $\liminf_{m \rightarrow \infty} \phi_m([N]) \geq \phi([N])$, which is more than sufficient. \square

A.4.5 Case 2: $|\mathcal{S}^*| < K$

We now consider the case where $\mathcal{S}^* = \{1, \dots, k\}$ with $k < K$. We will show that λ^* is still the unique minimizer of $V^*(\cdot)$. Since $V^*(\lambda^*) = \phi(\mathcal{S}^*)^2 = \phi([N])^2$ by definition, we know from Lemma 8 that λ^* does minimize V^* . It remains to show λ^* is the *unique* minimizer.

To do this, we will consider a perturbed informational environment in which signals $k+1, \dots, N$ are made slightly more precise. Specifically, let $\eta > 0$ be a small positive number. Consider an alternative signal coefficient matrix \tilde{C} with $\tilde{c}_i = c_i$ for $i \leq k$ and $\tilde{c}_i = (1 + \eta)c_i$ for $i > k$. Let $\tilde{\phi}(\mathcal{S})$ be the analogue of ϕ for this alternative environment. It is clear that $\tilde{\phi}(\mathcal{S}^*) = \phi(\mathcal{S}^*)$, while $\tilde{\phi}(\mathcal{S})$ is slightly smaller than $\phi(\mathcal{S})$ for $\mathcal{S} \neq \mathcal{S}^*$. Thus with sufficiently small η , the set \mathcal{S}^* remains the unique minimizer of ϕ (among sets that minimally span ω) in this perturbed environment, and the definition of λ^* is also maintained.

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

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

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

The equality uses (3) and $C'AC = \sum_i \lambda_i c_i c_i' = \sum_{i \leq k} \lambda_i c_i c_i' + \sum_{i > k} \frac{\lambda_i}{(1+\eta)^2} \tilde{c}_i \tilde{c}_i'$. The inequality follows from the homogeneity of \tilde{V}^* .

The above display implies that any frequency vector λ ,

$$V^*(\lambda) \geq \frac{\phi(\mathcal{S}^*)^2}{1 - \frac{2\eta + \eta^2}{(1+\eta)^2} \sum_{i > k} \lambda_i} \geq \frac{\phi(\mathcal{S}^*)^2}{1 - \eta \sum_{i > k} \lambda_i} \quad \text{for some } \eta > 0. \quad (9)$$

Hence $V^*(\lambda) > \phi(\mathcal{S}^*)^2 = V^*(\lambda^*)$ whenever λ puts positive weight outside of \mathcal{S}^* . But it is easily checked that $V^*(\lambda)$ is uniquely minimized at λ^* when λ is supported on \mathcal{S}^* . Hence λ^* is the unique minimizer of V^* over the whole simplex. This proves Lemma 7, which completes the proof of the propositions via Lemma 5.

A.5 Proof of Theorem 1 Part (a)

Let signals $1, \dots, k$ (with $k \leq K$) be a strongly complementary set; by Lemma 1 in the main text, these signals are best in their subspace. We will demonstrate an open set of prior beliefs given which *all agents* observe these k signals. Since these signals are complementary, Proposition 1 implies they must be linearly independent. Thus we can consider linearly transformed states $\tilde{\theta}_1, \dots, \tilde{\theta}_K$ such that these k signals are simply $\tilde{\theta}_1, \dots, \tilde{\theta}_k$ plus standard Gaussian noise. This linear transformation is invertible, so any prior over the original states is bijectively mapped to a prior over the transformed states. Thus it is without loss to work with the transformed model and look for prior beliefs over the transformed states.

The payoff-relevant state ω becomes a linear combination $\lambda_1^* \tilde{\theta}_1 + \dots + \lambda_k^* \tilde{\theta}_k$ (up to a scalar multiple). Since the first k signals are best in their subspace, Lemma 6 before implies that any other signal belonging 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 any prior belief with precision matrix P ; the inverse of P is the prior covariance matrix (in terms of the transformed states). Suppose ϵ is a very small positive number, and P satisfies the following conditions:

1. For $1 \leq i \leq k$, $P_{ii} \geq \frac{1}{\epsilon^2}$;
2. For $1 \leq i \neq j \leq k$, $\frac{P_{ij}}{\lambda_i^*} \leq (1 + \epsilon) \cdot \frac{P_{jj}}{\lambda_j^*}$;

3. For $k+1 \leq i \leq K$, $P_{ii} \in [\epsilon, 2\epsilon]$;
4. For $1 \leq i \neq j \leq K$, $|P_{ij}| \leq \epsilon^2$.

It is clear that any such P is positive definite, since on each row the diagonal entry has dominant size. Moreover, the set of P is open. Below we show that given any such prior, the myopic signal choice is among the first k signals, and that the posterior precision matrix also satisfies the same four conditions. As such, *all* agents would choose from the first k signals.

Let $V = P^{-1}$ be the prior covariance matrix. Applying Cramer's rule for the matrix inverse, the above conditions on P imply the following conditions on V :

1. For $1 \leq i \leq k$, $V_{ii} \leq 2\epsilon^2$;
2. For $1 \leq i \neq j \leq k$, $V_{ii}\lambda_i^* \leq (1 + L\epsilon) \cdot V_{jj}\lambda_j^*$;
3. For $k+1 \leq i \leq K$, $V_{ii} \in [\frac{1}{4\epsilon}, \frac{2}{\epsilon}]$;
4. For $1 \leq i \neq j \leq K$, $|V_{ij}| \leq L\epsilon \cdot V_{ii}$.

Here L is a constant depending only on K (but not on ϵ). For example, the last condition is equivalent to $\det(P_{-ij}) \leq L\epsilon \cdot \det(P_{-ii})$. This is proved by expanding both determinants into multilinear sums, and using the fact that on each row of P the off-diagonal entries are at most ϵ -fraction of the diagonal entry.

Given this matrix V , the variance reduction of $\omega = \sum_{i=1}^k \lambda_i^* \tilde{\theta}_i$ by any signal $\sum_{i=1}^k \alpha_i \tilde{\theta}_i + \mathcal{N}(0, 1)$ can be computed as

$$\frac{(\sum_{i,j=1}^k \alpha_i \lambda_j^* V_{ij})^2}{1 + \sum_{i,j=1}^k \alpha_i \alpha_j V_{ij}},$$

where the denominator is the variance of the signal and the numerator is the covariance between the signal and ω . By the first and last conditions on V , the denominator here is $1 + O(\epsilon^2)$. By the second and last condition, the numerator is

$$\left(\left(\sum_{i=1}^k \alpha_i + O(\epsilon) \right) \cdot \lambda_1^* V_{11} \right)^2.$$

Since $|\sum_{i=1}^k \alpha_i| < 1$, we deduce that any other signal belonging to the subspace of the first k signals is myopically worse than signal 1, whose variance reduction is $\frac{(\lambda_1^* V_{11})^2}{1 + V_{11}}$.

Meanwhile, take any signal outside of the subspace. The variance reduction by such a signal $\sum_{i=1}^K \beta_i \tilde{\theta}_i + \mathcal{N}(0, 1)$ is

$$\frac{(\sum_{i=1}^K \sum_{j=1}^k \beta_i \lambda_j^* V_{ij})^2}{1 + \sum_{i,j=1}^K \beta_i \beta_j V_{ij}}$$

By the second and last condition on V , the numerator here is $O((\lambda_1^* V_{11})^2)$. If we can show that the denominator is very large, then such a signal would also be myopically worse than signal 1. Indeed,

since $V_{ij} = O(\epsilon^2)$ whenever $i \leq k$ or $j \leq k$, it is sufficient to show $\sum_{i,j>k} \beta_i \beta_j V_{ij}$ is large. This holds by the last two conditions on V and the assumption that $\beta_{k+1}, \dots, \beta_K$ are not all zero.³⁴

Hence, we have shown that given any prior precision matrix P satisfying the above conditions, the myopic signal choice is among the first k signals. It remains to check the resulting posterior precision matrix \hat{P} also satisfies those four conditions. If the signal acquired is signal i ($1 \leq i \leq k$), then $\hat{P} = P + \Delta_{ii}$. Therefore we only need to show the second condition holds for \hat{P} ; that is, $\frac{P_{ii}+1}{\lambda_i^*} \leq (1+\epsilon) \cdot \frac{P_{jj}}{\lambda_j^*}$ for each $1 \leq j \leq k$. To this end, we note that since signal i is myopically best given V , the following must hold:

$$\frac{(\lambda_i^* V_{ii})^2}{1 + V_{ii}} \geq \frac{(\lambda_j^* V_{jj})^2}{1 + V_{jj}}.$$

As $0 \leq V_{ii}, V_{jj} \leq 2\epsilon^2$, this implies $\lambda_i^* V_{ii} \geq (1 - \epsilon^2) \lambda_j^* V_{jj}$. Now applying Cramer's rule to $V = P^{-1}$ again, we can deduce $V_{ii} = \frac{1+O(\epsilon^2)}{P_{ii}}$. So for ϵ small it holds that $\frac{P_{ii}}{\lambda_i^*} \leq (1 + \frac{\epsilon}{2}) \cdot \frac{P_{jj}}{\lambda_j^*}$. As $P_{ii} \geq \frac{1}{\epsilon^2}$, we also have $\frac{1}{\lambda_i^*} \leq \frac{\epsilon}{2} \cdot \frac{P_{jj}}{\lambda_j^*}$. Adding up these two inequalities yields the second condition for \hat{P} and completes the proof.

A.6 Proof of Theorem 1 Part (b)

A.6.1 Restated Version

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

Indeed, by definition we can find some period t after which agents *exclusively* observe signals from \mathcal{A} . Note that the variance reduction of any signal approaches zero as its signal count gets large. Thus, along society's signal path, the variance reduction is close to zero at sufficiently late periods. If \mathcal{A} does not span ω , 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 continuity, we deduce that at any sufficiently late period, observing signal j is better than observing any signal in \mathcal{A} . This contradicts our assumption that later agents only observe signals in \mathcal{A} .

³⁴Formally, we can without loss assume $\beta_K^2 V_{KK}$ is largest among $\beta_i^2 V_{ii}$ for $i > k$. Then for any $i \neq j$, the last condition implies

$$\beta_i \beta_j V_{ij} \geq -L\epsilon \cdot \beta_i \beta_j \sqrt{V_{ii} V_{jj}} \geq -L\epsilon \cdot \beta_K^2 V_{KK}.$$

This trivially also holds for $i = j \neq K$. Summing across all pairs $(i, j) \neq (K, K)$ yields $\sum_{i,j>k} \beta_i \beta_j V_{ij} > (1 - K^2 L\epsilon) \beta_K^2 V_{KK}$, which must be large by the third condition on V .

³⁵To see this, let s_1, \dots, s_N denote the limit signal counts, where $s_i = \infty$ if and only if $i \in \mathcal{A}$. We need to find some signal j such that $V(s_j + 1, s_{-j}) < V(s_j, s_{-j})$. If such a signal does not exist, then all partial derivatives of V at s are zero. Since V is always differentiable (unlike V^*), this would imply that all directional derivatives of V are also zero. By the convexity of V , V must be minimized at s . However, the minimum value of V is zero because there exists a complementary set. This contradicts $V(s) > 0$.

Now that \mathcal{A} spans ω , we can take \mathcal{S} to be the best complementary set in the subspace spanned by \mathcal{A} ; \mathcal{S} is strongly complementary by Lemma 1. To prove Theorem 1 part (b), we will show that long-run frequencies are positive precisely for the signals in \mathcal{S} . By ignoring the initial periods, we can assume without loss that *only signals in $\overline{\mathcal{A}}$ are available*. It thus suffices to show that whenever the signals observed infinitely often span a subspace, agents eventually focus on the best complementary set \mathcal{S} in that subspace. To ease notation, we assume this subspace is the entire \mathbb{R}^K , and prove the following result:

Theorem 1 part (b) Restated. *Suppose that the signals observed infinitely often span \mathbb{R}^K . Then society's long-run frequency vector is λ^* .*

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

A.6.2 Estimates of Derivatives

We introduce a few technical lemmata:

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

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

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

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

and

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

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

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

Note that $(\Sigma^0)^{-1} + C'QC \succeq q_j \cdot c_j c_j'$ in matrix norm. Thus the number $c_j' \cdot [(\Sigma^0)^{-1} + C'QC]^{-1} \cdot c_j$ is bounded above by $\frac{1}{q_j}$.³⁶ This proves the lemma. \square

³⁶Formally, we need to show that for any $\epsilon > 0$, the number $c_j' [c_j c_j' + \epsilon I_K]^{-1} c_j$ is at most 1. Using the identity $\text{Trace}(AB) = \text{Trace}(BA)$, we can rewrite this number as

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

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

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

Lemma 10. *For any q_1, \dots, q_N , we have³⁷*

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

Proof. We will show the more general result:

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

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

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

Equivalently, we need to show

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

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

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

This is exactly the result of Lemma 9. □

A.6.3 Lower Bound on Variance Reduction

Our next result lower bounds the directional derivative of V along the “optimal” direction λ^* :

Lemma 11. *For any q_1, \dots, q_N , we have $|\partial_{\lambda^*} V(q)| \geq \frac{V(q)^2}{\phi(\mathcal{S}^*)^2}$.*

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

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

This is equivalent to

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

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

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

³⁷Note that the convexity of V gives $V(q) - V(q_j + 1, q_{-j}) \leq |\partial_j V(q)|$. This lemma provides a converse that we need for the subsequent analysis.

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

We can now bound the variance reduction at late periods:

Lemma 12. *Fix any q_1, \dots, q_N . Suppose L is a positive number such that $(\Sigma^0)^{-1} + C'QC \succeq Lc_j c_j'$ holds for each signal $j \in \mathcal{S}^*$. Then we have*

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

Proof. Fix any signal $j \in \mathcal{S}^*$. Using the condition $(\Sigma^0)^{-1} + C'QC \succeq Lc_j c_j'$, we can deduce the following variant of Lemma 10:³⁸

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

Since V is always differentiable, $\partial_{\lambda^*} V(q)$ is a convex combination of the partial derivatives of V .³⁹ Thus $\max_{j \in \mathcal{S}^*} |\partial_j V(q)| \geq |\partial_{\lambda^*} V(q)|$. These inequalities, and Lemma 11, complete the proof. \square

A.6.4 Proof of the Restated Theorem 1 Part (b)

We will show $t \cdot V(m(t)) \rightarrow \phi(\mathcal{S}^*)^2$, so that society eventually approximates the optimal speed of learning. Since λ^* is the unique minimizer of V^* , this will imply the desired conclusion $\frac{m(t)}{t} \rightarrow \lambda^*$ (e.g., via the second half of Proposition 2 part (a)).

To estimate $V(m(t))$, we note that for any fixed L , society's acquisitions $m(t)$ eventually satisfy the condition $(\Sigma^0)^{-1} + C'QC \succeq Lc_j c_j'$. This is due to our assumption that the signals observed infinitely often span \mathbb{R}^K , which implies that $C'QC$ becomes arbitrarily large in matrix norm. Hence, we can apply Lemma 12 to find that

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

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

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

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

³⁸Even though we are not guaranteed $q_j \geq L$, we can modify the prior and signal counts such that the precision matrix $(\Sigma^0)^{-1} + C'QC$ is unchanged, and signal j has been observed at least L times. This is possible thanks to the condition $(\Sigma^0)^{-1} + C'QC \succeq Lc_j c_j'$. Then, applying Lemma 10 to this modified problem yields the result here.

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

Inverting both sides, we have

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)(1 - \frac{L}{L+1}g(t))} = \frac{1}{g(t)} + \frac{\frac{L}{L+1}}{1 - \frac{L}{L+1}g(t)} \geq \frac{1}{g(t)} + \frac{L}{L+1}. \quad (10)$$

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

$$V(m(t)) \leq \frac{L+1}{L} \cdot \frac{\phi(\mathcal{S}^*)^2}{t - t_0}. \quad (11)$$

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

We comment that the above argument leaves open the possibility that some signals outside of \mathcal{S}^* are observed *infinitely often*, yet with *zero long-run frequency*. In Appendix D.1, we show this does not happen.

A.7 Proofs for Interventions (Section 7)

A.7.1 Proof of Proposition 3

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

Fix any $\epsilon > 0$. If B is sufficiently large, then the above strategy acquires at least $(1 - \epsilon) \cdot B \cdot \lambda_i^*$ observations of each signal i . Adapting the proof of Lemma 11, we see that the agent's posterior precision about ω must increase by $\frac{(1-\epsilon)B}{\phi(\mathcal{S}^*)^2}$ under this strategy. Thus the same must hold for his optimal strategy, so that society's posterior precision at time t is at least $\frac{(1-\epsilon)Bt}{\phi(\mathcal{S}^*)^2}$. This implies that average precision per signal is at least $\frac{1-\epsilon}{\phi(\mathcal{S}^*)^2}$, which can be arbitrarily close to the optimal precision $\text{val}([N]) = \frac{1}{\phi(\mathcal{S}^*)^2}$ with appropriate choice of ϵ .

Since λ^* is the unique minimizer of V^* , society's long-run frequencies must be close to λ^* . In particular, with ϵ sufficiently small, we can ensure that each signal in \mathcal{S}^* are observed with positive frequencies. The restated Theorem 1 part (b) extends to the current setting and implies that society's long-run frequency vector must be λ^* . This yields the proposition.⁴⁰

⁴⁰This 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 complementary set. Intuitively, in order to achieve long-run efficient learning, agents need to allocate B observations in the best set to approximate the optimal frequencies. If another set of signals offers a speed of learning that is only slightly worse, we will need B sufficiently large for the approximately optimal frequencies in the best set to beat this other set.

A.7.2 Proof of Proposition 4

Suppose without loss that the best complementary set \mathcal{S}^* is $\{1, \dots, k\}$. By taking a linear transformation, we further assume each of the first k signals only involves ω and the first $k-1$ confounding terms b_1, \dots, b_{k-1} . We will show that whenever $k-1$ sufficiently precise signals are provided about each of these confounding terms, the long-run frequency vector converges to λ^* regardless of the prior.

Fix any positive real number L . Since the $k-1$ free signals are very precise, it is as if the prior precision matrix (after taking into account these free signals) satisfies

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

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

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

Adding up the above two displays, we have

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

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

Now this is exactly the condition we need in order to apply Lemma 12: Crucially, whether or not the condition is met for signals j outside of \mathcal{S}^* does not affect the argument there. Thus we can follow the proof of the restated Theorem 1 part (b) to deduce (11). That is, for fixed L and corresponding free information, society's long-run precision per signal is at least $\frac{L}{(L+1)\phi(\mathcal{S}^*)^2}$. This can be made arbitrarily close to the optimal average precision. Identical to the previous proof, we deduce that for large L , society's long-run frequency vector must be close to λ^* . The restated Theorem 1 part (b) allows us to conclude that the frequency is exactly λ^* .

A.8 General Payoff Functions

We comment here on the possibilities for (and limitations to) generalizing Proposition 2 beyond the quadratic loss payoff function. As discussed in the main text, the property that one-dimensional normal signals are Blackwell-ordered implies that part (a) of the proposition extends to general payoff functions. In other words, the frequency λ^* always maximizes long-run average precision.

On the other hand, at least for *some* other "prediction problems," λ^* continues to be the optimal frequency vector for maximizing the patient discounted average payoff. In a prediction problem, every agent's payoff function $u(a, \omega)$ is the same and depends only on $|a - \omega|$. For example, we can show that part (b) of Proposition 2 holds for $u(a, \omega) = |a - \omega|^\gamma$ with any exponent $\gamma \in (0, 2]$.

Nonetheless, even restricting to prediction problems, that part of the proposition does *not* hold in general. For a counterexample, consider $u(a, \omega) = -\mathbf{1}_{\{|a - \omega| > 1\}}$, which punishes the agent for

any prediction that differs from the true state by more than 1.⁴¹ Intuitively, the payoff gain from further information decreases sharply (indeed, exponentially) with the amount of information that has already been acquired. Thus, even with a forward-looking objective function, the range of future payoffs is limited and each agent cares mostly to maximize his own payoff. This results in an optimal sampling strategy that resembles myopic behavior, and differs from the rule that would maximize speed of learning.

The above counterexample illustrates the difficulty in estimating the value of information with an arbitrary payoff function. In order to make intertemporal payoff comparisons, we need to know how much payoff is gained/lost when the posterior variance is decreased/increased by a certain amount. This can be challenging in general, see [Chade and Schlee \(2002\)](#) for a related discussion.⁴²

Finally, while it may not be necessary to assume that agents have the same payoff function, part (b) of Proposition 2 can only hold under some restrictions on how the payoff functions differ. Otherwise, suppose for example that payoffs take the form $-\alpha_t(a_t - \omega)^2$, where α_t decreases exponentially fast. Then even with the δ -discounted objective, the social planner puts most of the weight on earlier agents, making it optimal to acquire signals myopically.

B Proofs for the Autocorrelated Model (Section 8.2)

B.1 Proof of Theorem 2

We work with the transformed model such that the signals in \mathcal{S} become the first k transformed states $\tilde{\theta}_1, \dots, \tilde{\theta}_k$. The payoff-relevant state becomes a certain linear combination $w_1\tilde{\theta}_1 + \dots + \tilde{\theta}_k$ with positive weights w_1, \dots, w_k . Choose M so that the innovations corresponding to the transformed states are independent from each other. In other words, \tilde{M} (the transformed version of M) is given by $\text{diag}(\frac{x}{w_1}, \dots, \frac{x}{w_k}, y_{k+1}, \dots, y_K)$. Here x is a *small* positive number, while y_{k+1}, \dots, y_K are *large* positive numbers. We further choose $\Sigma^0 = M$, which is the stable belief without learning.

With these choices, it is clear that if all agents only sample from \mathcal{S} , society's beliefs about the transformed states remain independent at every period. Let v_i^{t-1} denote the prior variance of $\tilde{\theta}_i^t$ at the beginning of period t (before the signal acquisition in that period). Then as long as agent t would continue to sample a signal $\tilde{\theta}_j + \mathcal{N}(0, 1)$ in \mathcal{S} , these prior variances would evolve as follows: $v_i^0 = \frac{x}{w_i}$ for $1 \leq i \leq k$ and $v_i^0 = y_i$ for $i > k$. And for $t \geq 1$,

$$v_i^t = \begin{cases} \alpha \cdot v_i^{t-1} + (1 - \alpha)\tilde{M}_{ii}, & \text{if } i \neq j; \\ \alpha \cdot \frac{v_i^{t-1}}{1 + v_i^{t-1}} + (1 - \alpha)\tilde{M}_{ii} & \text{if } i = j. \end{cases}$$

The particular signal j maximizes the reduction in the posterior variance of $\omega^t = \sum_{i=1}^k w_i \tilde{\theta}_i^t$. That is, $j \in \text{argmax}_{1 \leq i \leq k} \frac{(w_i \cdot v_i^{t-1})^2}{1 + v_i^{t-1}}$.

⁴¹We thank Alex Wolitzky for this example.

⁴²This difficulty becomes more salient if we try to go beyond prediction problems: The value of information in that case will depend on signal realizations.

By induction, it is clear that $v_i^t \leq \tilde{M}_{ii}$ holds for all pairs i, t , with equality for $i > k$. Thus at the beginning of each period t , assuming that all previous agents have sampled from \mathcal{S} , agent t 's prior uncertainties about $\tilde{\theta}_1, \dots, \tilde{\theta}_k$ are small while his uncertainties about $\tilde{\theta}_{k+1}, \dots, \tilde{\theta}_K$ are large. As such, our previous proof for the existence of learning traps with persistent states carries over, and we deduce that agent t continues to observe from \mathcal{S} .

Note that for α close to 1, agents will sample each signal i with frequency close to $\frac{w_i}{w_1 + \dots + w_k}$. It follows that the prior variances v_i^t approximately satisfy the following fixed-point equation:

$$v_i = \alpha \cdot \left(v_i - \frac{w_i}{w_1 + \dots + w_k} \cdot \frac{v_i^2}{1 + v_i} \right) + (1 - \alpha) \cdot \frac{x}{w_i}.$$

This yields the first-order approximation $v_i^t \sim \frac{\sqrt{(1-\alpha)x \cdot (w_1 + \dots + w_k)}}{w_i}$ as $\alpha \rightarrow 1$ and $t \rightarrow \infty$. The posterior variance of ω^t is therefore approximated as

$$\sum_{i=1}^k w_i^2 \cdot v_i^t \sim \sqrt{(1-\alpha)x \cdot \left(\sum_{i=1}^k w_i \right)^3}.$$

This is exactly $\sqrt{(1-\alpha) \left(\frac{M_{11}}{\text{val}(\mathcal{S})} \right)}$ since $M_{11} = \sum_{i=1}^k w_i^2 \cdot \tilde{M}_{ii} = x \cdot \sum_{i=1}^k w_i$ and $\text{val}(\mathcal{S}) = \frac{1}{\phi(\mathcal{S})^2} = \frac{1}{(\sum_{i=1}^k w_i)^2}$. We thus deduce the payoff estimate in part (1) of the theorem.

It remains to prove part (2). For that we just need to show society can achieve posterior variances smaller than $(1+\epsilon) \cdot \sqrt{(1-\alpha) \left(\frac{M_{11}}{\text{val}(\mathcal{S}^*)} \right)}$ at *every* late period. In fact, we show below that myopically choosing from the best set \mathcal{S}^* achieves this.

Let V^t denote society's prior covariance matrix at the beginning of period $t+1$, under this alternative sampling strategy. Note that for α close to 1, each signal in \mathcal{S}^* is observed with positive frequency. Thus, for any $L > 0$ we have

$$[V^t]^{-1} \succeq L \cdot c_j c_j', \quad \forall j \in \mathcal{S}^*$$

for α close to 1 and t large.

Now, take $v(t) := [V^t]_{11}$ to be the prior variance of ω^{t+1} . Then myopic sampling together with Lemma 12 implies the *posterior* variance of ω^{t+1} is bounded above by

$$v(t) - \frac{L}{L+1} \cdot \frac{v(t)^2}{\phi(\mathcal{S}^*)^2}.$$

Together with the innovation terms given by M , the next prior variance $v(t+1)$ admits the following upper bound:

$$v(t+1) \leq \alpha \cdot \left(v(t) - \frac{L}{L+1} \cdot \frac{v(t)^2}{\phi(\mathcal{S}^*)^2} \right) + (1-\alpha) \cdot M_{11}.$$

Clearly, this implies

$$v(t+1) < v(t) + (1-\alpha) \cdot M_{11} - \frac{L}{L+1} \cdot \frac{v(t)^2}{\phi(\mathcal{S}^*)^2}.$$

As a consequence, $\limsup_{t \rightarrow \infty} v(t) \leq \sqrt{1+1/L} \cdot \sqrt{(1-\alpha)M_{11} \cdot \phi(\mathcal{S}^*)^2}$. Since $\text{val}(\mathcal{S}^*) = \frac{1}{\phi(\mathcal{S}^*)^2}$, this yields the desired estimate if we choose $L > \frac{1}{\epsilon}$ in the first place. We have thus completed the proof of Theorem 2.

B.2 Proof of Proposition 5

The environment in Example 5 is equivalent to one with three signals $\frac{\omega+b}{2}$, $\frac{\omega-b}{2}$ and $\frac{1}{L}\omega$, each with standard Gaussian noise (just let $b = \omega + 2b_1$). We assume L is large, so that the best complementary set consists of the latter two signals.

For the autocorrelated model, we choose $M = \Sigma^0 = \text{diag}(x, x)$ with $x \geq L^2$ (this is the covariance matrix for the innovations associated with ω and b). Then assuming that all previous agents have chosen the third (unbiased) signal, agent t 's prior variance of b^t remains $x \geq L^2$. As such, he (and in fact each agent) continues to observe the third signal. In this case the prior variance v^t about ω^{t+1} evolves according to

$$v^t = \alpha \cdot \frac{L^2 \cdot v^{t-1}}{L^2 + v^{t-1}} + (1 - \alpha)x.$$

It is not difficult to show that v^t must converge to the (positive) fixed point of the above equation. Let us in particular take $\alpha = 1 - \frac{1}{L^3}$ and $x = L^2$, then the long-run prior variance v solves $v = \frac{(L^2 - \frac{1}{L})v}{L^2 + v} + \frac{1}{L}$. This yields exactly that $v = \sqrt{L}$. Hence long-run posterior variance is $\frac{L^2 \cdot v}{L^2 + v} > \sqrt{L}/2$, which implies $\limsup_{\delta \rightarrow 1} U_\delta^M \leq -\sqrt{L}/2$.

Let us turn to the optimal sampling strategy. Write $\tilde{\theta}_1 = \frac{\omega+b}{2}$ and $\tilde{\theta}_2 = \frac{\omega-b}{2}$. In this transformed model, $\tilde{M} = \tilde{\Sigma}^0 = \text{diag}(\frac{x}{2}, \frac{x}{2})$, and the payoff-relevant state is the sum of $\tilde{\theta}_1$ and $\tilde{\theta}_2$. Consider now a strategy that observes the first two signals alternatively. Then the beliefs about $\tilde{\theta}_1$ and $\tilde{\theta}_2$ remain independent (as in \tilde{M} and $\tilde{\Sigma}^0$), and their variances evolve as follows: $v_1^0 = v_2^0 = \frac{x}{2}$; in odd periods t

$$v_1^t = \alpha \cdot \frac{v_1^{t-1}}{1 + v_1^{t-1}} + (1 - \alpha)\frac{x}{2} \text{ and } v_2^t = \alpha \cdot v_2^{t-1} + (1 - \alpha)\frac{x}{2},$$

and symmetrically for even t .

These imply that for odd t , v_1^t converges to v_1 and v_2^t converges to v_2 below (while for even t $v_1^t \rightarrow v_2$ and $v_2^t \rightarrow v_1$):

$$v_1 = \alpha \cdot \frac{\alpha v_1 + (1 - \alpha)x/2}{1 + \alpha v_2 + (1 - \alpha)x/2} + (1 - \alpha)x/2;$$

$$v_2 = \alpha^2 \cdot \frac{v_2}{1 + v_2} + (1 - \alpha^2) \cdot \frac{x}{2}.$$

From the second equation, we obtain $(1 - \alpha^2)(\frac{x}{2} - v_2) = \alpha^2 \cdot \frac{(v_2)^2}{1 + v_2}$. With $\alpha = 1 - \frac{1}{L^3}$ and $x = L^2$, it follows that

$$v_2 = (1 + o(1))\frac{1}{\sqrt{L}}.$$

where $o(1)$ is a term that vanishes as $L \rightarrow \infty$. Thus we also have

$$v_1 = \alpha \frac{v_2}{1 + v_2} + (1 - \alpha)\frac{x}{2} = (1 + o(1))\frac{1}{\sqrt{L}}.$$

Hence under this alternating sampling strategy, long-run posterior variances about $\tilde{\theta}_1$ and $\tilde{\theta}_2$ are both bounded above by $\frac{2}{\sqrt{L}}$. Since $\omega = \tilde{\theta}_1 + \tilde{\theta}_2$, we conclude that $\liminf_{\delta \rightarrow 1} U_\delta^{SP} \geq -\frac{4}{\sqrt{L}}$. Choosing L large proves the proposition.

References

- Ali, Nageeb.** 2018. “Herding with Costly Information.” *Journal of Economic Theory*, 175: 713–720.
- Athey, Susan, and Armin Schmutzler.** 1995. “Product and Process Flexibility in an Innovative Environment.” *RAND Journal of Economics*, 26(4): 557–574.
- Banerjee, Abhijit.** 1992. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, 107(3): 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch.** 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades.” *Journal of Political Economy*, 100(5): 992–1026.
- Borgers, Tilman, Angel Hernando-Veciana, and Daniel Krahmer.** 2013. “When Are Signals Complements Or Substitutes.” *Journal of Economic Theory*, 148(1): 165–195.
- Burguet, Roberto, and Xavier Vives.** 2000. “Social Learning and Costly Information.” *Economic Theory*, 15(1): 185–205.
- 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*, 13: 377–414.
- Chaloner, Kathryn.** 1984. “Optimal Bayesian Experimental Design for Linear Models.” *The Annals of Statistics*, 12(1): 283–300.
- Chen, Yiling, and Bo Waggoner.** 2016. “Informational Substitutes.” Working Paper.
- Che, Yeon-Koo, and Konrad Mierendorff.** 2017. “Optimal Sequential Decision with Limited Attention.” Working Paper.
- Cunha, Flavio, and James Heckman.** 2007. “The Technology of Skill Formation.” *American Economic Review*, 97(2): 31–47.
- Dasaratha, Krishna, Ben Golub, and Nir Hak.** 2018. “Social Learning in a Dynamic Environment.” Working Paper.
- DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel.** 2003. “Persuasion Bias, Social Influence, and Unidimensional Opinions.” *The Quarterly Journal of Economics*.
- Easley, David, and Nicolas M. Kiefer.** 1988. “Controlling a Stochastic Process with Unknown Parameters.” *Econometrica*, 56: 1045–1064.
- Frongillo, Rafael, Grant Schoenebeck, and Omer Tamuz.** 2011. “Social Learning in a Changing World.” *WINE’11 Proceedings of the 7th international conference on Internet and Network Economics*.
- Fudenberg, Drew, Philip Strack, and Tomasz Strzalecki.** 2018. “Speed, Accuracy, and the Optimal Timing of Choices.” *American Economic Review*. Forthcoming.
- Gittins, J. C.** 1979. “Bandit processes and dynamic allocation indices.” *Journal of the Royal Statistical Society, Series B*, 148–177.

- Goldstein, Itay, and Liyan Yang.** 2015. “Information Diversity and Complementarities in Trading and Information Acquisition.” *Journal of Finance*, 70(4): 1723–1765.
- Golub, Benjamin, and Matthew Jackson.** 2012. “How Homophily Affects the Speed of Learning and Best-Response Dynamics.” *The Quarterly Journal of Economics*, 127(3): 1287–1338.
- 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.
- Jovanovic, Boyan, and Yaw Nyarko.** 1996. “Learning by Doing and the Choice of Technology.” *Econometrica*, 64(6): 1299–1310.
- Liang, Annie, Xiaosheng Mu, and Vasilis Syrgkanis.** 2017. “Optimal and Myopic Information Acquisition.” Working Paper.
- Lizzeri, Alessandro, and Marciano Siniscalchi.** 2008. “Parental Guidance and Supervised Learning.” *The Quarterly Journal of Economics*.
- Mayskaya, Tatiana.** 2017. “Dynamic Choice of Information Sources.” Working Paper.
- Milgrom, Paul, and Rober J. Weber.** 1982*a*. “A Theory of Auctions and Competitive Bidding.” *Econometrica*, 50(5): 1089–1122.
- Milgrom, Paul, and Robert J. Weber.** 1982*b*. “The Value of Information in a Sealed-Bid Auction.” *Journal of Mathematical Economics*, 10(1): 105–114.
- Monderer, Dov, and Lloyd Shapley.** 1996. “Potential Games.” *Games and Economic Behavior*, 14(1): 124–143.
- Moscarini, Giuseppe, Marco Ottaviani, and Lones Smith.** 1998. “Social Learning in a Changing World.” *Economic Theory*, 11(3): 657–665.
- Mueller-Frank, Manuel, and Mallesh Pai.** 2016. “Social Learning with Costly Search.” *American Economic Journal: Microeconomics*, 8(1): 83–109.
- 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.
- Vives, Xavier.** 1992. “How Fast do Rational Agents Learn?” *Review of Economic Studies*, 60(2): 329–347.
- Vivi Alatas, Abhijit Banerjee, Arun Chandrasekhar Rema Hanna, and Ben Olken.** 2016. “Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia.” *American Economic Review*, 106(7): 1663–1704.

For Online Publication

C Proof of Proposition 2 Part (b)

We will first generalize part (b) of Theorem 1 to show that for *any* $\delta \in (0, 1)$ and any prior belief, the social planner's sampling strategy that maximizes δ -discounted payoff yields frequency vectors that converge over time. Moreover, the limit is the optimal frequency vector associated with some strongly complementary set. Later we will argue that that for δ close to 1, this long-run outcome must be the best complementary set \mathcal{S}^* from all priors.

C.1 Long-run Characterization for All δ

Here we prove the following result:

Proposition 7. *Suppose $\delta \in (0, 1)$. Given any prior, let $d_\delta(t)$ denote the vector of signals counts associated with any signal acquisition strategy that maximizes the δ -discounted average payoff. Then $\lim_{t \rightarrow \infty} \frac{d_\delta(t)}{t}$ exists and is equal to $\lambda^{\mathcal{S}}$ for some strongly complementary set \mathcal{S} .*

Proof. We follow the proof of Theorem 1 part (b) in Appendix A.6. The same argument there shows that for any $\delta < 1$, any strategy that maximizes δ -discounted payoff must infinitely observe a set of signals that span ω . Therefore it remains to prove the analogue of the restated version of Theorem 1 part (b).

To do that, let

$$W(t) = (1 - \delta) \sum_{t' \geq t} \delta^{t'-t} \cdot V(d(t'))$$

denote the expected discounted loss from period t onwards; henceforth we fix δ and use $d(t)$ as shorthand for $d_\delta(t)$. Suppose signal acquisitions in the first t periods satisfy $C'QC \succeq Lc_jc_j'$ for each signal $j \in \mathcal{S}^*$, where L is some positive constant. Then we are going to show that

$$\frac{1}{W(t+1)} \geq \frac{1}{W(t)} + \frac{L}{(L+1)\phi(\mathcal{S}^*)^2}. \quad (12)$$

Once this is proved, we can choose L large to show $W(t) \leq \frac{(1+\epsilon)\phi(\mathcal{S}^*)^2}{t}$ for any $\epsilon > 0$ and all sufficiently large t . Pick m so that $\delta^m \leq \epsilon$. Then for $t' \in (t, t+m)$ we have $V(d(t')) \geq \frac{(1-\epsilon/2)\phi(\mathcal{S}^*)^2}{t'} \geq \frac{(1-\epsilon)\phi(\mathcal{S}^*)^2}{t}$, so that

$$(1 - \delta) \sum_{t'=t+1}^{t+m-1} \delta^{t'-t} \cdot V(d(t')) \geq (\delta - \delta^m) \cdot \frac{(1 - \epsilon)\phi(\mathcal{S}^*)^2}{t} \geq \frac{(\delta - \epsilon)(1 - \epsilon)\phi(\mathcal{S}^*)^2}{t}.$$

Subtracting this from $W(t)$, we obtain

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

again for t sufficiently large depending on ϵ . Since ϵ is arbitrary, we would be able to conclude $t \cdot V(d(t)) \rightarrow \phi(\mathcal{S}^*)^2$, and $\frac{d(t)}{t} \rightarrow \lambda^*$ would follow.

To prove (12), we consider a deviation strategy that chooses signals myopically in every period $t' \geq t + 1$. Let the resulting signal count vectors be $\tilde{d}(t')$, and define $\tilde{d}(t) = d(t)$. This deviation provides an upper bound on $W(t + 1)$, given by

$$W(t + 1) \leq (1 - \delta) \sum_{t' \geq t+1} \delta^{t'-t-1} \cdot V(\tilde{d}(t')).$$

Since $W(t) = (1 - \delta) \cdot V(d(t)) + \delta \cdot W(t + 1)$, we have

$$\frac{1}{W(t + 1)} - \frac{1}{W(t)} = \frac{(1 - \delta) \cdot (V(d(t)) - W(t + 1))}{W(t + 1) \cdot ((1 - \delta) \cdot V(d(t)) + \delta \cdot W(t + 1))},$$

which is decreasing in $W(t + 1)$ (holding $V(d(t))$ equal). Thus from the previous upper bound on $W(t + 1)$, we obtain that

$$\frac{1}{W(t + 1)} - \frac{1}{W(t)} \geq \frac{1}{\sum_{j=0}^{\infty} (1 - \delta) \delta^j \cdot V(\tilde{d}(t + 1 + j))} - \frac{1}{\sum_{j=0}^{\infty} (1 - \delta) \delta^j \cdot V(\tilde{d}(t + 1 + j))} \quad (13)$$

By the assumption that $C'QC \succeq Lc_jc'_j$ after t periods, we can apply (10) to deduce that for each $j \geq 0$,

$$\frac{1}{V(\tilde{d}(t + 1 + j))} - \frac{1}{V(\tilde{d}(t + j))} \geq \frac{L}{(L + 1)\phi(\mathcal{S}^*)^2}.$$

Given this and (13), the desired result (12) follows from the technical lemma below (with $a = \frac{L}{(L+1)\phi(\mathcal{S}^*)^2}$, $x_j = V(\tilde{d}(t + 1 + j))$, $y_j = V(\tilde{d}(t + j))$ and $\beta_j = (1 - \delta)\delta^j$):

Lemma 13. *Suppose a is a positive number. $\{x_j\}_{j=0}^{\infty}, \{y_j\}_{j=0}^{\infty}$ are two sequences of positive numbers such that $\frac{1}{x_j} \geq \frac{1}{y_j} + a$ for each j . Then for any sequence of positive numbers $\{\beta_j\}_{j=0}^{\infty}$ that sum to 1, it holds that*

$$\frac{1}{\sum_{j=0}^{\infty} \beta_j x_j} \geq \frac{1}{\sum_{j=0}^{\infty} \beta_j y_j} + a.$$

To see why this lemma holds, note that it is without loss to assume $\frac{1}{x_j} = \frac{1}{y_j} + a$ holds with equality. Then

$$1 - a \sum_j \beta_j x_j = \sum_j \beta_j (1 - a x_j) = \beta_j \frac{x_j}{y_j}$$

By the Cauchy-Schwarz inequality,

$$\sum_j \beta_j \frac{x_j}{y_j} \geq \frac{1}{\sum_j \beta_j \frac{y_j}{x_j}} = \frac{1}{\sum_j \beta_j (1 + a y_j)} = \frac{1}{1 + a \sum_j \beta_j y_j}.$$

So $1 - a \sum_j \beta_j x_j \geq \frac{1}{1 + a \sum_j \beta_j y_j}$, which is easily seen to be equivalent to $\frac{1}{\sum_j \beta_j x_j} \geq \frac{1}{\sum_j \beta_j y_j} + a$.

Hence Lemma 13 is proved, and so is Proposition 7. \square

C.2 Efficiency as $\delta \rightarrow 1$

We now prove that for δ close to 1, the sampling strategy that maximizes δ -discounted payoff must eventually focus on the best complementary set \mathcal{S}^* . Recall that V^* is uniquely maximized at λ^* . Thus there exists positive η such that $V^*(\lambda) > (1 + \eta)V^*(\lambda^*)$ whenever λ puts zero frequency on at least one signal in \mathcal{S}^* .

Suppose for contradiction that sampling eventually focuses on a strongly complementary set \mathcal{S} different from \mathcal{S}^* . Then at large periods t we must have $V(d(t)) > \frac{(1+\eta)\phi(\mathcal{S}^*)^2}{t}$, using the fact that V^* is the asymptotic version of V . As a result, there exists sufficiently large L_0 such that some signal in \mathcal{S}^* is observed less than L_0 times under the optimal strategy for maximizing δ -discounted payoff.⁴³ Crucially, this L_0 can be chosen independently of δ . As a consequence, under the hypothesis of inefficient long-run outcome, $V(d(t)) > \frac{(1+\eta)\phi(\mathcal{S}^*)^2}{t}$ in fact holds for all $t > \underline{t}$ where \underline{t} is also independent of δ .

Now we fix a positive integer $L > \frac{2}{\eta}$, and consider the following deviation strategy starting in period $\underline{t} + 1$:

1. In periods $\underline{t} + 1$ through $\underline{t} + Lk$, observe each signal in the best set \mathcal{S}^* (of size k) exactly L times, in any order.
2. Starting in period $\underline{t} + Lk + 1$, sample myopically.

Let us study the posterior variance after period $\underline{t} + j$ under such a deviation. For $j \geq Lk + 1$, note that each signal $j \in \mathcal{S}^*$ has been observed at least L times before the period $\underline{t} + Lk + 1$. So $C'QC \succeq Lc_jc'_j$ holds, and we can deduce (similar to (11)) that the posterior variance is at most $(1 + \frac{1}{L}) \cdot \frac{\phi(\mathcal{S}^*)^2}{j - Lk}$. Since $\frac{1}{L} < \frac{\eta}{2}$, there exists \underline{j} (depending on $\eta, \underline{t}, L, k$) such that the posterior variance after period $\underline{t} + j$ is at most $(1 + \eta/2) \frac{\phi(\mathcal{S}^*)^2}{\underline{t} + j}$ for $j > \underline{j}$. Thus the flow payoff *gain* in each such period is at least

$$\frac{\eta}{2} \cdot \frac{\phi(\mathcal{S}^*)^2}{\underline{t} + j}, \quad \forall j > \underline{j}$$

under this deviation strategy.

On the other hand, for $j \leq \underline{j}$ we can trivially bound the posterior variance from above by the prior variance V_0 . This V_0 also serves as an upper bound on the flow payoff *loss* in these periods.

Combining both estimates, we find that the deviation strategy achieves payoff gain of at least

$$\delta^{\underline{t}} \cdot \left(\sum_{j > \underline{j}} \delta^{j-1} \cdot \frac{\eta}{2} \cdot \frac{\phi(\mathcal{S}^*)^2}{\underline{t} + j} - \sum_{j=1}^{\underline{j}} \delta^{j-1} \cdot V_0 \right).$$

Importantly, all other parameters in the above are constants independent of δ . As δ approaches 1, the sum $\sum_{j > \underline{j}} \frac{\delta^{j-1}}{\underline{t} + j}$ approaches a harmonic sum which diverges. Thus for all δ close to 1 the above display is strictly positive, suggesting that the constructed deviation is profitable. This contradiction completes the proof of Proposition 2 part (b).

⁴³Otherwise, $C'QC \succeq L_0c_jc'_j$ holds at large t , implying a contradicting upper bound on $V(d(t))$ (see the argument in the previous subsection).

D Other Results and Examples

D.1 Strengthening of Theorem 1 part (b)

Here we show the following result, which strengthens the restated Theorem 1 part (b) (see Appendix A.6). It says that any signal observed with zero long-run frequency must in fact be observed only finitely often.

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

The proof is divided into two subsections below.

D.1.1 Log Residual Term

Recall that we have previously shown $m_i(t) \sim \lambda_i^* \cdot t$. We can first improve the estimate of the residual term to $m_i(t) = \lambda_i^* \cdot t + O(\ln t)$. Indeed, Lemma 12 yields that for some constant L and every $t \geq L$,

$$V(m(t+1)) \leq V(m(t)) - \left(1 - \frac{L}{t}\right) \cdot \frac{V(m(t))^2}{\phi(\mathcal{S}^*)^2}. \quad (14)$$

This is because we may apply Lemma 12 with $M = \min_{j \in \mathcal{S}^*} m_j(t)$, which is at least $\frac{t}{L}$.

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

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

Inverting both sides, we have

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)} + \frac{1 - \frac{L}{t}}{1 - (1 - \frac{L}{t})g(t)} \geq \frac{1}{g(t)} + 1 - \frac{L}{t}. \quad (15)$$

This enables us to deduce

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

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

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

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

$$V^*\left(\frac{m(t)}{t}\right) \leq t \cdot V(m(t)) + O\left(\frac{1}{t}\right) \leq \phi(\mathcal{S}^*)^2 + O\left(\frac{\ln t}{t}\right). \quad (16)$$

Suppose $\mathcal{S}^* = \{1, \dots, k\}$. Then the above estimate together with (9) implies $\sum_{j>k} \frac{m_j(t)}{t} = O(\frac{\ln t}{t})$. Hence $m_j(t) = O(\ln t)$ for each signal j outside of the best set.

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

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

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

D.1.2 Getting Rid of the Log

In order to remove the $\ln t$ residual term, we need a refined analysis. The reason we ended up with $\ln t$ is because we used (14) and (15) at *each* period t ; the “ $\frac{L}{t}$ ” term in those equations adds up to $\ln t$. In what follows, instead of quantifying the variance reduction in each period (as we did), we will lower-bound the variance reduction over multiple periods. This will lead to better estimates and enable us to prove $m_i(t) = \lambda_i^* \cdot t + O(1)$.

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

Ignoring period t_{l+1} for the moment, we are interested in $\frac{\phi(\mathcal{S}^*)^2}{V(m(t_{l+1}-1))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l))}$, which is just the difference in the *precision* about ω when the division vector changes from $m(t_l)$ to $m(t_{l+1}-1)$. From the proof of Lemma 11, we can estimate this difference if the change were along the direction λ^* :

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1} - 1 - t_l))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l))} \geq t_{l+1} - 1 - t_l. \quad (17)$$

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

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

we have shown. This implies⁴⁵

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

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

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_{l+1} - 1))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1} - 1 - t_l))} \geq -\frac{L \ln^2 t_{l+1}}{t_{l+1}}. \quad (18)$$

(17) and (18) together imply

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

Finally, we can apply (15) to $t = t_{l+1} - 1$. Altogether we deduce

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

Now observe that $\sum_l \frac{2L \ln^2 t_{l+1}}{t_{l+1}}$ converges (this is the sense in which our estimates here improve upon (15), where $\frac{L}{t}$ leads to a divergent sum). Thus we are able to conclude

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

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

$$V^*\left(\frac{m(t)}{t}\right) \leq t \cdot V(m(t)) + O\left(\frac{1}{t}\right) \leq \phi(\mathcal{S}^*)^2 + O\left(\frac{1}{t}\right). \quad (19)$$

This equation (19) improves upon the previously-derived (16). Hence by (9) again, $m_j(t) = O(1)$ for each signal $j > k$. And once these signal counts are fixed, $m_i(t) = \lambda_i^* \cdot t + O(1)$ also holds for signals i in the best set, as already proved in Liang, Mu and Syrgkanis (2017). This completes the proof.

D.2 Example of a Learning Trap with Non-Normal Signals

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

	θ_1	θ_2
s_1	p	$1 - p$
s_2	$1 - p$	p

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

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

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

D.3 Regions of Prior Beliefs

Here we analyze a pair of examples to illustrate which prior beliefs would lead to which strongly complementary set as the long-run outcome.

D.3.1 Revisiting Example 2

First consider Example 2 in the main text, with signals $X_1 = 3\omega + b_1 + \epsilon_1$, $X_2 = b_1 + \epsilon_2$, $X_3 = \omega + \epsilon_3$. For simplicity, we restrict attention to priors beliefs that are independent across ω and b_1 . We have shown that if the prior variance of b_1 is larger than 8, every agent observes the last signal and society gets stuck in a learning trap.

As a converse, we now show that whenever the prior variance of b_1 is *smaller* than 8, every agent chooses from the first two signals and learns efficiently in the long run. It is clear that the first agent observes X_1 . The rest of the argument is based on the following claim: After any number $k \geq 1$ of observations of X_1 and any number $l \geq 0$ of observations of X_2 , the next agent finds *either* signal X_1 *or* signal X_2 to be more valuable than X_3 .

We first prove this claim assuming $l = 0$. Let τ_ω and τ_b be the prior precisions about ω and b , with $\tau_b > \frac{1}{8}$ by assumption. After $k + 1$ observations of X_1 , society's posterior precision matrix is

$$P = \begin{pmatrix} \tau_\omega + 9(k+1) & 3(k+1) \\ 3(k+1) & \tau_b + k + 1 \end{pmatrix}$$

So the posterior variance of ω is $[P^{-1}]_{11}$, which is

$$f(k) = \frac{\tau_b + k + 1}{\tau_\omega \tau_b + 9(k+1)\tau_b + (k+1)\tau_\omega}.$$

If society had observed X_1 k times and X_2 once, posterior variance is

$$g(k) = \frac{\tau_b + k + 1}{\tau_\omega \tau_b + 9k\tau_b + (k+1)\tau_\omega + 9k}.$$

If society had observed X_1 k times and X_3 once, posterior variance would be

$$h(k) = \frac{\tau_b + k}{\tau_\omega \tau_b + (9k+1)\tau_b + k\tau_\omega + k}.$$

⁴⁶We thank Andrew Postlewaite for this example.

Now we need to check $\min\{f(k), g(k)\} < h(k)$. The comparison reduces to

$$\frac{\tau_b + k + 1}{\tau_b + k} < \frac{\max\{\tau_\omega \tau_b + 9(k+1)\tau_b + (k+1)\tau_\omega, \tau_\omega \tau_b + 9k\tau_b + (k+1)\tau_\omega + 9k\}}{\tau_\omega \tau_b + (9k+1)\tau_b + k\tau_\omega + k},$$

which further simplifies to

$$\frac{1}{\tau_b + k} < \frac{\max\{\tau_\omega + 8\tau_b - k, \tau_\omega - \tau_b + 8k\}}{\tau_\omega \tau_b + (9k+1)\tau_b + k\tau_\omega + k}.$$

Note that $\max\{\tau_\omega + 8\tau_b - k, \tau_\omega - \tau_b + 8k\} > \frac{2}{9}(\tau_\omega + 8\tau_b - k) + \frac{7}{9}(\tau_\omega - \tau_b + 8k) = \tau_\omega + \tau_b + 6k$. The above inequality (in the last display) thus follows from $(\tau_b + k)(\tau_\omega + \tau_b + 6k) > \tau_\omega \tau_b + (9k+1)\tau_b + k\tau_\omega + k$, which is equivalent to the simple inequality $\tau_b^2 + 6k^2 - k > (2k+1)\tau_b$.

Next we consider the case with $l \geq 1$ and show that whenever society has observed signals X_1 and X_2 both at least once, no future agent will choose X_3 . In this case we can apply Lemma 12 with $L = 1$ and $\mathcal{S}^* = \{X_1, X_2\}$. Note that $\phi(\mathcal{S}^*) = \frac{2}{3}$, so myopic sampling from this set leads to posterior variance at most

$$V(m(t)) - \frac{L}{L+1} \cdot \frac{V(m(t))^2}{\phi(\mathcal{S}^*)^2} = V(m(t)) - \frac{9}{8} \cdot V(m(t))^2 < V(m(t)) - V(m(t))^2.$$

On the other hand, observing signal X_3 would lead to posterior variance $\frac{V(m(t))}{1+V(m(t))}$, which is larger than $V(m(t)) \cdot (1 - V(m(t)))$. This proves the claim.

D.3.2 More Complex Dynamics

Here we study a more involved example, with signals

$$X_1 = 3\omega + b_1 + \epsilon_1$$

$$X_2 = b_1 + \epsilon_2$$

$$X_3 = 2\omega + b_2 + \epsilon_3$$

$$X_4 = b_2 + \epsilon_4$$

Note that the first two signals are the same as before, but X_3 and X_4 now involve another confounding term b_2 . The (strongly) complementary sets are $\{X_1, X_2\}$ and $\{X_3, X_4\}$, which partition the whole set as in Sethi and Yildiz (2017). We restrict attention to prior beliefs such that b_1 is independent from ω and b_2 .

We will show that learning trap is possible if and only if the prior variance of b_1 is larger than 8. On the one hand, if prior uncertainty about b_1 is large, we can specify prior beliefs over ω and b_2 such that $2\omega + b_2$ and b_2 are *independent and have the same small variance*. Given such a prior, the reduction in posterior variance of ω by either signal X_3 or X_4 is close to the reduction by the unbiased signal $\omega + \epsilon$ (see the proof of Theorem 1 part (a)). This allows us to show that agents alternate between observing X_3 and X_4 , leading to a learning trap.

On the other hand, if the prior variance of b_1 is smaller than 8, we claim that society eventually focuses on the efficient set $\{X_1, X_2\}$. To see this, we assume for contradiction that the long-run

outcome is $\{X_3, X_4\}$. Then at late periods, posterior beliefs have the property that $2\omega + b_2$ and b_2 are approximately independent and have approximately the same small variance. This means both signals X_3 and X_4 have approximately the same marginal value as the unbiased signal $\omega + \epsilon$. But then we return to the situation in Example 2, where future agents find it optimal to choose from $\{X_1, X_2\}$ instead (as analyzed in the previous section). This contradiction proves the claim.

Note however that unlike in Example 2, it may take a long time for society to eventually switch to the best set. For example, with prior variances $1, 7, \frac{1}{100}$ of ω, b_1, b_2 respectively, the first 89 agents observe X_3 and signal X_1 or X_2 is only observed afterwards. Moreover, switching between different complementary sets can also happen more than once: With a standard Gaussian prior of (ω, b_1, b_2) , society's signal path begins with $X_1 X_3 X_1 X_2 \dots$. These more complex dynamics make it difficult to pinpoint the long-run outcome as a function of the prior.

D.4 Example Mentioned in Section 7.1

Suppose the available signals are

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

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

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

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

D.5 Low Altruism

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

We follow the construction in Appendix A.5. The added difficulty here is to show that if any agent ever chooses a signal $j > k$, the payoff loss in that period (relative to myopically choosing

among the first k signals) is at least a constant fraction of possible payoff gains in future periods. Once this is proved, then for sufficiently small ρ such a deviation is not profitable.

Suppose that agents sample only from the first k signals in the first $t-1$ periods, with frequencies close to λ^* . Then, the posterior variances V_{11}, \dots, V_{kk} (which are also the prior for period t) are on the order of $\frac{1}{t}$. Thus, following the computation in Appendix A.5, we can show that for some positive constant ξ (independent of t), the variance reduction of ω by any signal $j > k$ is at least $\frac{\xi}{t^2}$ smaller than the variance reduction by signal 1. This is the amount of payoff loss in period t under a deviation to signal j .

Such a deviation could improve the posterior variance in future periods. But even for the best continuation strategy, the posterior variance in period $t+m$ could at most be reduced by $O(\frac{m}{t^2})$.⁴⁷ Thus if we choose ξ to be small enough, the payoff gain in each period $t+m$ is bounded above by $\frac{m}{\xi t^2}$. Note that for ρ sufficiently small,

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

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

D.6 Multiple Payoff-Relevant States

Here we consider an extension of our main model, where each agent chooses the signal that minimizes a weighted sum of posterior variances about some r payoff-relevant states. These states and their weights are the same across agents. As before, let $V(q_1, \dots, q_N)$ denote this weighted posterior variance as a function of the signal counts. V^* is the normalized, asymptotic version of V .

We assume that V^* is uniquely minimized at some frequency vector λ^* . Part (a) of Proposition 2 extends and implies that λ^* maximizes speed of learning. Unlike the case of $r = 1$, this optimal frequency vector generally involves more than K signals if $r > 1$.⁴⁸ We are not aware of any simple method to characterize λ^* .

Nonetheless, We can generalize the notion of “complementary sets” as follows: A set of signals \mathcal{S} is complementary if both of the following properties hold:

1. each of the r payoff-relevant states is spanned by \mathcal{S} ;
2. the optimal frequency vector supported on \mathcal{S} puts positive weight on *each* signal in \mathcal{S} .

Similarly, we say that a complementary set \mathcal{S} is “strongly complementary” if it is best in its subspace: the optimal frequency vector supported on $\bar{\mathcal{S}}$ only puts positive weights on signals in \mathcal{S} . When $r = 1$, these definitions agree with our main model.

⁴⁷This is because over m periods, the increase in the precision matrix is at most linear in m .

⁴⁸A theorem of Chaloner (1984) shows that λ^* is supported on at most $\frac{r(2K+1-r)}{2}$ signals.

By this definition, the existence of learning traps readily extends: For suitable prior beliefs, the marginal value of each signal in \mathcal{S} persistently exceeds the marginal value of each signal in $\overline{\mathcal{S}} - \mathcal{S}$. Since the marginal values of the remaining signals (those outside of the subspace) can be made very low by imposing large prior uncertainty on the relevant confounding terms, we deduce that society exclusively observes from the strongly complementary set \mathcal{S} .

We mention that part (b) of Theorem 1 also generalizes. For that we need a different proof, since there is no obvious analogue of Lemma 11 (and thus of Lemma 12) when $r > 1$. Instead, we prove the restated Theorem 1 part (b) in Appendix A.6 as follows: When society infinitely samples a set that spans \mathbb{R}^K , the marginal value of each signal j can be approximated by its asymptotic version:

$$\partial_i V(q_1, \dots, q_N) \sim \frac{1}{t^2} \cdot \partial_i V^*\left(\frac{q_1}{t}, \dots, \frac{q_N}{t}\right).$$

Together with Lemma 10, this shows that the myopic signal choice j in any sufficient late period must *almost* minimize the partial derivative of V^* , in the following sense:

Lemma 14. *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*

$$\partial_j V^*\left(\frac{m(t)}{t}\right) \leq (1 - \epsilon) \min_{1 \leq i \leq N} \partial_i V^*\left(\frac{m(t)}{t}\right).$$

Consider society's frequency vectors $\lambda(t) = \frac{m(t)}{t} \in \Delta^{N-1}$. Then they evolve according to

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

whenever j is the signal choice in period $t + 1$. So the frequencies $\lambda(t)$ move in the direction of e_j , which is the direction where V^* decreases almost the fastest. This suggests that the evolution of $\lambda(t)$ over time resembles the gradient descent dynamics. As such, we can expect that the value of $V^*(\lambda(t))$ roughly decreases over time, and that eventually $\lambda(t)$ approaches $\lambda^* = \operatorname{argmin} V^*$.

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

$$\begin{aligned} V^*(\lambda(t+1)) &= V^*\left(\frac{t}{t+1} \lambda(t) + \frac{1}{t+1} e_j\right) \\ &= V^*\left(\frac{t}{t+1} \lambda(t)\right) + \frac{1}{t+1} \cdot \partial_j V^*\left(\frac{t}{t+1} \lambda(t)\right) + O\left(\frac{1}{(t+1)^2} \cdot \partial_{jj} V^*\left(\frac{t}{t+1} \lambda(t)\right)\right) \\ &\leq V^*\left(\frac{t}{t+1} \lambda(t)\right) + \frac{1-\epsilon}{t+1} \cdot \partial_j V^*\left(\frac{t}{t+1} \lambda(t)\right) \\ &= \frac{t+1}{t} \cdot V^*(\lambda(t)) + \frac{(1-\epsilon)(t+1)}{t^2} \cdot \partial_j V^*(\lambda(t)) \\ &\leq V^*(\lambda(t)) + \frac{1}{t} \cdot V^*(\lambda(t)) + \frac{1-2\epsilon}{t} \cdot \min_{1 \leq i \leq N} \partial_i V^*(\lambda(t)). \end{aligned} \tag{20}$$

The first inequality uses Lemma 9, the next equality uses the homogeneity of V^* , and the last inequality uses Lemma 14.

Write $\lambda = \lambda(t)$ for short. Note that V^* is differentiable at λ , since $\lambda_i(t) > 0$ for a set of signals that spans the entire space. Thus the convexity of V^* yields

$$V^*(\lambda^*) \geq V^*(\lambda) + \sum_{i=1}^N (\lambda_i^* - \lambda_i) \cdot \partial_i V^*(\lambda).$$

The homogeneity of V^* implies $\sum_{i=1}^N \lambda_i \cdot \partial_i V^*(\lambda) = -V^*(\lambda)$. This enables us to rewrite the preceding inequality as

$$\sum_{i=1}^N \lambda_i^* \cdot \partial_i V^*(\lambda) \leq V^*(\lambda^*) - 2V^*(\lambda).$$

Thus, in particular,

$$\min_{1 \leq i \leq N} \partial_i V^*(\lambda(t)) \leq V^*(\lambda^*) - 2V^*(\lambda). \quad (21)$$

Combining (20) and (21), we have for all large t :

$$V^*(\lambda(t+1)) \leq V^*(\lambda(t)) + \frac{1}{t} \cdot [(1-2\epsilon) \cdot V^*(\lambda^*) - (1-4\epsilon) \cdot V^*(\lambda(t))]. \quad (22)$$

Now, suppose (for contradiction) that $V^*(\lambda(t)) > (1+4\epsilon) \cdot V^*(\lambda^*)$ holds for all large t . Then (22) would imply $V^*(\lambda(t+1)) \leq V^*(\lambda(t)) - \frac{\epsilon \cdot V^*(\lambda^*)}{t}$. But since the harmonic series diverges, $V^*(\lambda(t))$ would eventually decrease to be negative, which is impossible. Thus

$$V^*(\lambda(t)) \leq (1+4\epsilon) \cdot V^*(\lambda^*)$$

must hold for *some* large t . By (22), the same is true at all future periods. But since ϵ is arbitrary, the above inequality proves that $V^*(\lambda(t)) \rightarrow V^*(\lambda^*)$. Hence $\lambda(t) \rightarrow \lambda^*$, completing the proof of Theorem 1 for multiple payoff-relevant states.