Online Appendix

"Complementary Information and Learning Traps"

Annie Liang and Xiaosheng Mu

Appendices

A Proof of Proposition 2 Part (b)

We will first generalize the "if" part of Theorem 2 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 $\mathcal{S}^*$ starting from all priors.

A.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 $\lim_{t\to\infty} \frac{d_\delta(t)}{t}$ exists and is equal to $\lambda^{\mathcal{S}}$ for some strongly complementary set $\mathcal{S}$.*

Proof. We follow the proof of Theorem 2 "if" part 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 2 "if" part.

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

$$\frac{1}{W(t+1)} \geq \frac{1}{W(t)} + \frac{L}{(L+1)\phi(\mathcal{S}^*)^2}. \tag{13}$$

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

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

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

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

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

To prove (13), 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

$$\frac{1}{W(t+1)} - \frac{1}{W(t)} \geq \frac{1}{\sum_{j=0}^{\infty}(1-\delta)\delta^j \cdot V(\tilde{d}(t+1+j))} - \frac{1}{\sum_{j=0}^{\infty}(1-\delta)\delta^j \cdot V(\tilde{d}(t+1+j))} \tag{14}$$

By the assumption that $C'QC \succeq Lc_jc_j'$ after t periods, we can apply (11) 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 (14), the desired result (13) 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 14. *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

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

By the Cauchy-Schwarz inequality,

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

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 14 is proved, and so is Proposition 7. ☐

A.2 Efficiency as $\delta \to 1$

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.[51] 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 (12)) 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).

[51] Otherwise, $C'QC \succeq L_0c_jc_j'$ holds at large t, implying a contradicting upper bound on $V(d(t))$ (see the argument in the previous subsection).

B Strengthening of Theorem 2 "If" Part

Here we show the following result, which strengthens the restated Theorem 2 "if" part (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 2 "if" part. *Suppose that the signals observed infinitely often span* $\mathbb{R}^K$*. Then* $m_i(t) = \lambda_i^* \cdot t + O(1), \forall i$.

The proof is divided into two subsections below.

B.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 13 yields that for some constant L and every $t \geq L$,

$$V(m(t+1)) \leq V(m(t)) - \left(1 - \frac{L}{t}\right) \cdot \frac{V(m(t))^2}{\phi(\mathcal{S}^*)^2}. \tag{15}$$

This is because we may apply Lemma 13 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

$$\frac{1}{g(t+1)} \geq \frac{1}{g(t)} + \frac{1 - \frac{L}{t}}{1 - (1 - \frac{L}{t})g(t)} \geq \frac{1}{g(t)} + 1 - \frac{L}{t}. \tag{16}$$

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

$$V^*\left(\frac{m(t)}{t}\right) \leq t \cdot V(m(t)) + O\left(\frac{1}{t}\right) \leq \phi(\mathcal{S}^*)^2 + O\left(\frac{\ln t}{t}\right). \tag{17}$$

Suppose $\mathcal{S}^* = \{1, \ldots, k\}$. Then the above estimate together with (10) 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. We work with the transformed model, as in Appendix A.2.2. Let $\Sigma = \Sigma(t)$ be the posterior covariance matrix after observing $m_j(t)$ observations of each signal $j > k$. Then the posterior covariance matrix after additionally observing $m_i(t)$ observations of each signal $i \leq k$ can be written as $\left[\Sigma^{-1} + \mathrm{diag}(m_1(t), \ldots, m_k(t))\right]^{-1}$. It follows that

$$V(m(t)) = \beta' \cdot \left[\Sigma^{-1} + \mathrm{diag}(m_1(t), \ldots, m_k(t))\right]^{-1} \cdot \beta,$$

where $\beta \in \mathbb{R}^k$ is a shorthand for $\beta^{\mathcal{S}^*}$.

Using the formula for matrix derivatives, we have that for $1 \leq i \leq k$,

$$\partial_i V(m(t)) = -\beta' \cdot \left[\Sigma^{-1} + \mathrm{diag}(m_1(t), \ldots, m_k(t))\right]^{-1} \cdot \Delta_{ii} \cdot \left[\Sigma^{-1} + \mathrm{diag}(m_1(t), \ldots, m_k(t))\right]^{-1} \cdot \beta. \quad (18)$$

Recall that each $m_i(t)$ is on the order of t, whereas the precision matrix Σ^{-1} is given by

$$\Sigma^{-1} = (\Sigma^0)^{-1} + \sum_{j>k} m_j(t) \cdot c_j c_j' = O(\ln t).$$

Thus the matrix inverse $\left[\Sigma^{-1} + \mathrm{diag}(m_1(t), \ldots, m_k(t))\right]^{-1}$ can be approximated by $\mathrm{diag}(m_1(t), \ldots, m_k(t))^{-1}$ up to a factor of $O(\frac{\ln t}{t})$. Plugging into (18), we deduce

$$\begin{aligned}\partial_i V(m(t)) &= -\beta' \cdot \mathrm{diag}(1/m_1(t), \ldots, 1/m_k(t)) \cdot \Delta_{ii} \cdot \mathrm{diag}(1/m_1(t), \ldots, 1/m_k(t)) \cdot \beta \cdot (1 + O(\ln t/t)) \\ &= -\left(\frac{\beta_i}{m_i(t)}\right)^2 \cdot \left(1 + O\left(\frac{\ln t}{t}\right)\right).\end{aligned} \quad (19)$$

We now use this to show $m_i(t) \leq \lambda_i^* \cdot t + O(\ln t)$ for each $1 \leq i \leq k$. Suppose for the sake of contradiction that $m_1(t)$ exceeds $\lambda_1^* \cdot t$ by a (big) multiple of $\ln t$. Consider $\tau + 1 \leq t$ to be the last period in which signal 1 was observed. Then $m_1(\tau)$ is larger than $\lambda_1^* \cdot \tau$ by several $\ln \tau$. Since $\sum_{1 \leq i \leq k} m_i(\tau) \leq \tau$, there exists some other signal in the best set, say signal 2, with $m_2(\tau) < \lambda_2^* \cdot \tau$. This implies $\frac{\beta_2}{m_2(\tau)}$ exceeds $\frac{\beta_1}{m_1(\tau)}$ by a factor larger than $O(\frac{\ln \tau}{\tau})$. By (19), we then deduce that $\partial_2 V(m(\tau))$ is more negative than $\partial_1 V(m(\tau))$. But this suggests that the agent in period $\tau + 1$ should not have chosen signal 1, leading to a contradiction.

Hence $m_i(t) \leq \lambda_i^* \cdot t + O(\ln t)$ holds for all signals i, whether $i \leq k$ or $i > k$. Using $\sum_{i=1}^N m_i(t) = t$, we deduce that each $m_i(t)$ is also lower-bounded by $\lambda_i^* \cdot t - O(\ln t)$. This proves $m_i(t) = \lambda_i^* \cdot t + O(\ln t)$ as desired.

B.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 (15) and (16) 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, $t_l \geq 2^{\varepsilon \cdot l}$ 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 12, we can estimate this difference if the change were along the direction λ^*:

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1} - 1 - t_l))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l))} \geq t_{l+1} - 1 - t_l. \tag{20}$$

Now, the vector $m(t_{l+1} - 1)$ is not exactly equal to $m(t_l) + \lambda^*(t_{l+1} - 1 - t_l)$, so the above estimate is not directly applicable. However, by our definition of t_l and t_{l+1}, any difference between these vectors must be in the first k signals. In addition, the difference is bounded by $O(\ln t_{l+1})$ by what we have shown. This implies[52]

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

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

$$\frac{\phi(\mathcal{S}^*)^2}{V(m(t_{l+1} - 1))} - \frac{\phi(\mathcal{S}^*)^2}{V(m(t_l) + \lambda^*(t_{l+1} - 1 - t_l))} \geq -\frac{L \ln^2 t_{l+1}}{t_{l+1}}. \tag{21}$$

(20) and (21) 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 (16) 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 (16), 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

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

[52] By 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 $O(\frac{\ln t}{t^3})$.

This equation (22) improves upon the previously-derived (17). Hence by (10) again, $m_j(t) = O(1)$ for each signal $j > k$.

Once these signal counts are fixed, we can repeat the argument in the previous subsection to argue that $m_i(t) = \lambda_i^* \cdot t + O(1)$ for $i \leq k$. Indeed, the matrix Σ^{-1} is now bounded by $O(1)$, rather than $O(\ln t)$. Thus from (18) we can obtain

$$\partial_i V(m(t)) = -\frac{\beta_i}{m_i(t)}^2 \cdot \left(1 + O\left(\frac{1}{t}\right)\right) \tag{23}$$

Observe that the current estimate (23) improves upon the previous estimate (19). Hence by the same argument we used after (19), we can now deduce that

$$m_i(t) = \lambda_i^* \cdot t + O(1), \quad \forall i.$$

This completes the proof.

C Incomplete Learning

C.1 Dimension Reduction

Our results extend to situations where ω *cannot* be identified from the available sources. To see this, suppose that the signal coefficient vectors $c_1, \ldots, c_N$ span a $k-1$ dimensional subspace. Moreover, without loss assume that $c_1, \ldots, c_{k-1}$ are linearly independent (and thus span this subspace). Redefine the confounding variables to be the following linear combinations of the original states:

$$\tilde{b}_i = \langle c_i, \theta \rangle, \quad \forall 1 \leq i \leq k-1.$$

By assumption, each of the N signals can then be written as a linear combination of $\tilde{b}_1, \ldots, \tilde{b}_{k-1}$ plus noise. On the other hand, ω cannot be written as such a linear combination since it is not identified. We can thus work in this linearly transformed model, with new state vector $\tilde{\theta} = (\omega, \tilde{b}_1, \ldots, \tilde{b}_{k-1})'$ having dimension $k \leq K$. Note that the original prior covariance matrix $\Sigma^0 \in \mathbb{R}^{K \times K}$ induces a transformed prior covariance matrix $\tilde{\Sigma}^0 \in \mathbb{R}^{k \times k}$.

C.2 "Learnable" Component

In order to reduce this problem into one where the payoff-relevant state is identified, we will decompose ω into its *learnable and un-learnable components*. That is, we will write

$$\omega = \tilde{\omega} + \omega^{\perp}$$

where $\tilde{\omega}$ is a linear combination of $\tilde{b}_1, \ldots, \tilde{b}_{k-1}$ (thus "learnable"), and $\omega^{\perp}$ is orthogonal to $\tilde{b} := (\tilde{b}_1, \ldots, \tilde{b}_{k-1})'$ according to the prior $\tilde{\Sigma}^0$ (thus "un-learnable"). Such a decomposition exists and is unique. Formally, we seek $\tilde{\omega} = \gamma' \cdot \tilde{b}$ for some $\gamma \in \mathbb{R}^{k-1}$, such that $\text{Cov}(\tilde{b}, \omega - \tilde{\omega}) = 0$. Such a vector γ is uniquely determined by

$$\text{Cov}(\tilde{b}, \tilde{b}) \cdot \gamma = \text{Cov}(\tilde{b}, \omega). \tag{24}$$

In this equation, $\text{Cov}(\tilde{b}, \tilde{b})$ represents the $(k-1) \times (k-1)$ prior covariance matrix of $\tilde{b}$ (i.e., the bottom-right submatrix of $\tilde{\Sigma}^0$). Similarly, $\text{Cov}(\tilde{b}, \omega)$ is the $(k-1) \times 1$ vector of covariances between $\tilde{b}_i$ and ω (i.e., the bottom-left submatrix of $\tilde{\Sigma}^0$).

Since the random variable $\omega^{\perp} = \omega - \tilde{\omega}$ is independent from $\tilde{b}_1, \ldots, \tilde{b}_{k-1}$, it is also independent from any linear combination of these variables. Hence, uncertainty about $\omega^{\perp}$ cannot be reduced upon observation of any of the available signals. It follows that minimizing the variance of ω (the sum of the variances of $\tilde{\omega}$ and of $\omega^{\perp}$) is equivalent to minimizing the variance of $\tilde{\omega}$. Thus agents only seek to learn about $\tilde{\omega}$, and the problem is as if $\tilde{\omega}$ were the payoff-relevant state. This returns us to the case where the payoff-relevant state *is* identified. We can then define (strongly) complementary sets with respect to learning about $\tilde{\omega}$, as well as the efficient set $\mathcal{S}^*$.

C.3 Generalization of Results

We emphasize that the payoff-relevant state $\tilde{\omega}$ obtained in this way depends on the prior covariance matrix $\tilde{\Sigma}^0$. As a result, which sets are complementary in this more general setting are dependent on the prior belief over ω and $\tilde{b}$, since they are defined w.r.t. $\tilde{\omega}$.

Nonetheless, our results directly generalize for any fixed $\tilde{\omega}$, in the following way.

Definition 5. *Fix nonzero $\gamma \in \mathbb{R}^{k-1}$ and $\tilde{\omega} = \gamma' \cdot \tilde{b}$.*[53] *Define the* set of $\tilde{\omega}$-learnable priors *to include all prior covariance matrices over $(\omega, \tilde{b})$ that satisfy (24).*

Then Theorem 2 asserts: *Fix $\tilde{\omega}$ such that different complementary sets (w.r.t. $\tilde{\omega}$) have distinct informational values (Assumption 2). Then as the prior belief varies within the set of $\tilde{\omega}$-learnable priors, strongly complementary sets (w.r.t. $\tilde{\omega}$) are the only possible long-run observation sets.*

Theorem 1 also generalizes, so long as we replace"K" by "$k-1$" in the theorem statement and in the Strong Linear Independence assumption (Assumption 3). This modification reflects the dimension reduction we performed earlier.

We note that Assumption 2 is more than necessary for these results to hold. As the proof of Theorem 2 shows, it is sufficient to assume that *within each subspace, a unique complementary set maximizes* val (i.e., minimizes ϕ). Using the derivations in the proof of Lemma 8, this latter assumption is satisfied (for all γ and $\tilde{\omega} = \gamma' \cdot \tilde{b}$) whenever the signal coefficient vectors $c_1, \ldots, c_N$ have the following *generic* property:

Assumption 4. *For every set of linearly independent signal coefficient vectors $c_{i_1}, \ldots, c_{i_m}$ and every other c_j that can be (uniquely) written as $c_j = \sum_{l=1}^{m} \alpha_l \cdot c_{i_l}$, it holds that*

$$\pm\alpha_1 \pm \alpha_2 \pm \cdots \pm \alpha_m \neq 1$$

for all (2^m) choices of pluses and minuses.

[53] If $\gamma = 0$, $\omega = \omega^{\perp}$ cannot be learned at all. We will henceforth rule out such priors.

Below we maintain this assumption.

While the above generalization of Theorem 2 considers prior beliefs that are $\tilde{\omega}$-learnable for a fixed $\tilde{\omega}$, we can take a step further and ask which sets can be observed in the long run *across all prior beliefs over* $(\omega, \tilde{b})$ (so that $\tilde{\omega}$ also varies). The answer builds on the following definition inspired by Lemma 8:

Definition 6. *The set $\mathcal{S}$ is* locally best *if its signal coefficient vectors $c_{i_1}, \ldots, c_{i_m}$ are linearly independent and have the following property: There exists a set of choices of pluses and minuses $\delta_1, \ldots, \delta_m \in \{-1, 1\}$, such that for all signal coefficient vectors $c_j \notin \mathcal{S}$ that can be written as $c_j = \sum_{l=1}^m \alpha_l \cdot c_{i_l}$ for some (unique) $(\alpha_1, \ldots, \alpha_m) \in \mathbb{R}^m$, it holds that*

$$|\delta_1 \alpha_{j1} + \cdots + \delta_m \alpha_{jm}| < 1. \tag{25}$$

In words: any signal outside of $\mathcal{S}$ that is spanned by signals in $\mathcal{S}$ (i.e., can be written as a linear combination of signals in S) must be expressible via "small" weights $(\alpha_1, \ldots, \alpha_m)$, in the sense described by (25). The rough intuition is that the "smaller" are these weights (as if the noise ϵ_j is scaled up), the less informative is the signal $X_j \notin \mathcal{S}$ compared to the signals in $\mathcal{S}$. As formalized in the result below, this is precisely what guarantees $\mathcal{S}$ to have the highest informational value in its subspace (thus strongly complementary by Lemma 1).

Proposition 8. *Under Assumption 4, a set of sources $\mathcal{S}$ are eventually exclusively observed starting from some set of prior beliefs over $(\omega, \tilde{b})$ if and only if it is "locally best" as defined above.*

Proof. It is equivalent to prove that $\mathcal{S}$ is strongly complementary w.r.t. *some* $\tilde{\omega}$ if and only if it is locally best. Suppose $\mathcal{S}$ is strongly complementary w.r.t. $\tilde{\omega}$, then it is at least complementary. So some linear combination $\sum_{l=1}^m \beta_l X_{i_l}$ of the signals in $\mathcal{S}$ produces an unbiased estimate of $\tilde{\omega}$. Let δ_l be the sign of β_l for $1 \le l \le m$, then the fact that $\mathcal{S}$ is strongly complementary together with Lemma 8 (when applied to the subspace spanned by $\mathcal{S}$) yields the above inequality (25). So $\mathcal{S}$ is locally best.

Conversely, suppose $\mathcal{S}$ is locally best. Let $\delta_1, \ldots, \delta_m \in \{-1, 1\}$ be given as in the definition. We now define $\tilde{\omega}$ to be $\sum_{l=1}^m \delta_l X_{i_l}$ with noise terms removed. Reversing the proof of Lemma 8, we deduce that $\phi(\mathcal{S}) < \phi(\mathcal{S}')$ whenever $|\mathcal{S} - \mathcal{S}'| = |\mathcal{S}' - \mathcal{S}| = 1$ and $\mathcal{S}'$ belongs to the subspace $\overline{\mathcal{S}}$. Note that $\phi(\mathcal{S}) < \phi(\mathcal{S}')$ also holds if such a set $\mathcal{S}'$ is not in this subspace, since in that case $\phi(\mathcal{S}') = \infty$.[54] Hence $\phi(\mathcal{S}) < \phi(\mathcal{S}')$ always holds, and it follows from Proposition 6 that $\mathrm{val}(\mathcal{S}) > \mathrm{val}(\mathcal{S}')$. $\mathcal{S}$ is then strongly complementary w.r.t. $\tilde{\omega}$ by definition. □

C.4 Example

We provide an example to illustrate the above analysis.

[54] Note that $c_{i_1}, c_{i_2}, \ldots, c_{i_l}$ uniquely span $\tilde{\omega}$. So $c_{i_2}, \ldots, c_{i_l}$ and another vector c_j (replacing c_{i_1}) span $\tilde{\omega}$ only if c_j is spanned by $c_{i_2}, \ldots, c_{i_l}$ and $\tilde{\omega}$. Thus c_j must be a linear combination of $c_{i_1}, c_{i_2}, \ldots, c_{i_l}$.

Example 7. The available sources are $X_1 = \omega + b_1 + b_2 + \varepsilon_1$, $X_2 = b_1 + \varepsilon_2$, and $X_3 = \omega + b_1 + b_3 + \varepsilon_3$. Note that ω cannot be completely learned, even given infinite observations of all three signals. Define $\tilde{b}_1 = \omega + b_1 + b_2, \tilde{b}_2 = b_1$ and $\tilde{b}_3 = \omega + b_1 + b_3$.

(a) First consider the simplest prior belief over (ω, b_1, b_2, b_3): $\Sigma^0 = I$. Decomposing ω into its learnable and un-learnable components, we get

$$\tilde{\omega} = \frac{2\omega + b_2 + b_3}{3} = \frac{\tilde{b}_1 - \tilde{b}_2 + \tilde{b}_3}{3},$$

which can be completely learned, and

$$\omega^{\perp} = \omega - \tilde{\omega} = \frac{\omega - b_2 - b_3}{3},$$

which is orthogonal to $\tilde{b}$ according to the prior. Because the prior Σ^0 is very simple, we can check orthogonality directly without having to compute $\tilde{\Sigma}^0$ and γ (which, as shown above, is equal to $(1/3, -1/3, 1/3)$).

Since agents cannot reduce uncertainty about $\omega^{\perp}$ by observing any of the available signals, the problem is equivalent to one in which the payoff-relevant state were given instead by $\tilde{\omega}$. In this case there is a unique complementary set w.r.t. learning about $\tilde{\omega}$, and it is the whole set $\{X_1, X_2, X_3\}$. Thus, starting from this prior belief (or any $\tilde{\omega}$-learnable prior), agents eventually observe all three signals with equal frequencies.

(b) As the prior belief Σ^0 varies, however, different long-run outcomes can emerge. Indeed, as the three available signals are linearly independent, condition (25) is trivially satisfied because $c_j \notin \mathcal{S}$ *cannot* be written as a linear combination of $c_{i_1}, \ldots, c_{i_m}$. Thus every non-empty subset is "locally best," and can be exclusively observed in the long run under some prior.

To illustrate, suppose we want to look for priors such that $\{X_1, X_2\}$ will be eventually observed with equal frequencies, whereas X_3 is not eventually observed. This suggests[55]

$$\tilde{\omega} = \tilde{b}_1 - \tilde{b}_2 = \omega + b_1 + b_2 - b_1 = \omega + b_2.$$

If this $\tilde{\omega}$ were the learnable component, then $\{X_1, X_2\}$ would be the only complementary set, and agents would indeed achieve the desired long-run frequencies.

For completeness, we provide an example of such a $\tilde{\omega}$-learnable prior. For this construction, note that $\omega^{\perp}$ should be $\omega - \tilde{\omega} = -b_2$. So the set of $\tilde{\omega}$-learnable priors are those covariance matrices Σ^0 such that b_2 is orthogonal to $\omega + b_1 + b_2$, b_1, as well as $\omega + b_1 + b_3$. Specifically, one such prior is

$$\begin{pmatrix} \omega \\ b_1 \\ b_2 \\ b_3 \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix}, \begin{pmatrix} 3 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 3 \end{pmatrix} \right).$$

for arbitrary $\mu_1, \mu_2, \mu_3, \mu_4$.

[55] Another possibility is to have $\tilde{\omega} = \tilde{b}_1 + \tilde{b}_2$, which involves a slightly more complex computation.

D Other Extensions

D.1 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.[56]

However, the interpretation of the optimal benchmark (that we defined in Section 5) 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. We comment on this possibility below.

We first note that the quadratic loss payoff assumed in our main model belongs to a class of "prediction problems", where every agent's payoff function $u(a, \omega)$ is the same and depends only on $|a - \omega|$. Part (b) of Proposition 2 does generalize to some other prediction problems; for example, our proof applies to any payoff function of the form $u(a, \omega) = |a - \omega|^\gamma$, with exponent $\gamma \in (0, 2]$.

Nonetheless, even restricting to prediction problems, part (b) of Proposition 2 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.[57] 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.[58]

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 ex-

[56] To 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.

[57] We thank Alex Wolitzky for this example.

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

ponentially 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.

D.2 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 $\mathbb{E}\left[\sum_{t'\geq t}\rho^{t'-t}\cdot(a_t-\omega)^2\right]$, assuming that future agents will behave similarly. Note that $\rho=0$ returns our main model. Below we show that for ρ sufficiently small, the existence of learning traps extends to this setting.[59]

Suppose signals $1,\dots,k$ are strongly complementary. We want to show that for sufficiently small $\rho>0$, there exist priors given which agents with 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 $O(\frac{m}{t^2})$.[60] 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.3 Multiple Payoff-Relevant States

In our main model, only one of the K persistent states is payoff-relevant. Consider a generalization in which each agent predicts (the same) $r\leq K$ unknown states and his payoff is determined via a

[59] The other half of Theorem 2 also extends: Proposition 7 shows that strongly complementary sets are the only possible long-run outcomes (for any ρ).

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

weighted sum of quadratic losses. We show here that our main results extend to this setting. As before, let $V(q_1, \ldots, q_N)$ 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 $r > 1$.[61] 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 $\overline{\mathcal{S}}$ only puts positive weights on signals in $\mathcal{S}$. When $r = 1$, these definitions agree with our main model.

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 the "if" part of Theorem 2 also generalizes. For that we need a different proof, since there is no obvious analogue of Lemma 12 (and thus of Lemma 13) when $r > 1$. Instead, we prove the restated Theorem 2 "if" part 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:

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

Together with Lemma 11, 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 15. *For any $\varepsilon > 0$, there exists sufficiently large $t(\varepsilon)$ such that if signal j is observed in any period $t+1$ later than $t(\varepsilon)$, then*

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

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

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

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

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 $\varepsilon > 0$ and sufficiently large t)

$$\begin{aligned}
V^*(\lambda(t+1)) &= V^*\left(\frac{t}{t+1}\lambda(t) + \frac{1}{t+1}e_j\right) \\
&= V^*\left(\frac{t}{t+1}\lambda(t)\right) + \frac{1}{t+1}\cdot \partial_j V^*\left(\frac{t}{t+1}\lambda(t)\right) + O\left(\frac{1}{(t+1)^2}\cdot \partial_{jj} V^*\left(\frac{t}{t+1}\lambda(t)\right)\right) \\
&\leq V^*\left(\frac{t}{t+1}\lambda(t)\right) + \frac{1-\varepsilon}{t+1}\cdot \partial_j V^*\left(\frac{t}{t+1}\lambda(t)\right) \\
&= \frac{t+1}{t}\cdot V^*(\lambda(t)) + \frac{(1-\varepsilon)(t+1)}{t^2}\cdot \partial_j V^*(\lambda(t)) \\
&\leq V^*(\lambda(t)) + \frac{1}{t}\cdot V^*(\lambda(t)) + \frac{1-2\varepsilon}{t}\cdot \min_{1\leq i\leq N} \partial_i V^*(\lambda(t)).
\end{aligned} \tag{26}$$

The first inequality uses Lemma 10, the next equality uses the homogeneity of V^*, and the last inequality uses Lemma 15.

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,

$$\min_{1\leq i\leq N} \partial_i V^*(\lambda(t)) \leq V^*(\lambda^*) - 2V^*(\lambda). \tag{27}$$

Combining (26) and (27), we have for all large t:

$$V^*(\lambda(t+1)) \leq V^*(\lambda(t)) + \frac{1}{t}\cdot [(1-2\varepsilon)\cdot V^*(\lambda^*) - (1-4\varepsilon)\cdot V^*(\lambda(t))]. \tag{28}$$

Now, suppose (for contradiction) that $V^*(\lambda(t)) > (1+4\varepsilon)\cdot V^*(\lambda^*)$ holds for all large t. Then (28) would imply $V^*(\lambda(t+1)) \leq V^*(\lambda(t)) - \frac{\varepsilon\cdot V^*(\lambda^*)}{t}$. But since the harmonic series diverges, $V^*(\lambda(t))$ would eventually decrease to be negative, which is impossible. Thus

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

must hold for *some* large t. By (28), the same is true at all future periods. But since ε is arbitrary, the above inequality proves that $V^*(\lambda(t)) \to V^*(\lambda^*)$. Hence $\lambda(t) \to \lambda^*$, completing the proof of Theorem 2 for multiple payoff-relevant states.

E 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.

E.0.1 Revisiting Section 2.1

First consider the example in Section 2.1, with signals $X_1 = 3\omega + b_1 + \varepsilon_1, X_2 = b_1 + \varepsilon_2, X_3 = \omega + \varepsilon_3$. 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

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

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

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

which further simplifies to

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

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 13 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.

E.0.2 More Complex Dynamics

Here we study a more involved example, with signals

$$\begin{aligned}
X_1 &= 3\omega + b_1 + \varepsilon_1 \\
X_2 &= b_1 + \varepsilon_2 \\
X_3 &= 2\omega + b_2 + \varepsilon_3 \\
X_4 &= b_2 + \varepsilon_4
\end{aligned}$$

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 (2019). 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 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 $\omega + \varepsilon$ (see the proof of Theorem 2). 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 the unbiased signal $\omega + \varepsilon$. But then we return to the situation in Section 2.1, 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 Section 2.1, 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 $X_1 X_3 X_1 X_2 \cdots$. These more complex dynamics make it difficult to pinpoint the long-run outcome as a function of the prior.

F Example Mentioned in Footnote 34 (Section 9.1)

Suppose the available signals are

$$
\begin{aligned}
X_1 &= 10x + \varepsilon_1 \\
X_2 &= 10y + \varepsilon_2 \\
X_3 &= 4x + 5y + 10b \\
X_4 &= 8x + 6y - 20b
\end{aligned}
$$

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 10 independent realizations, agents starting from the above prior would observe only the signals X_3 and X_4. This breaks the learning trap.

G Example of a Learning Trap beyond Normality

The existence of learning traps is not special to the assumption of normally-distributed states and signals. We provide here a simple example demonstrating a learning trap in an environment with binary states and signals, and leave to future work the interesting question of what general conditions are sufficient.

Suppose there is a payoff-relevant state $\omega \in \{\omega_1, \omega_2\}$, which agents can learn about from any of three available signals. The first signal X_1 is described by the following information structure:

	$X_1 = a$	$X_1 = b$
$\omega = \omega_1$	p	$1 - p$
$\omega = \omega_2$	$1 - p$	p

where p itself is unknown and equals either $p_1 = 0.9$ or $p_2 = 0.1$. A second signal X_2 provides information about p, and is described by

	$X_2 = a$	$X_2 = b$
$p = p_1$	0.9	0.1
$p = p_2$	0.1	0.9

Finally, agents have access to a signal with information structure

	$X_3 = a$	$X_3 = b$
$\omega = \omega_1$	2/3	1/3
$\omega = \omega_2$	1/3	2/3

As in our earlier example from Section 2.1, repeated acquisition of X_3 is sufficient for agents to learn ω, but this path of acquisitions is *inefficient* in the sense that it leads to a sub-optimal long-run speed of learning. Specifically, agents can learn faster by alternating between acquisition of X_1 and X_2, as the following result shows (see the next subsection for its proof):

Claim 2. *The pair of signals (X_1, X_2) strictly Blackwell-dominates the pair (X_3, X_3), so that observing each of X_1 and X_2 once is more informative than two independent realizations of X_3.*

Recall from Blackwell (1951) that if an experiment P is more informative another experiment Q, then for every $n \geq 1$, drawing n conditionally independent samples from P is also more informative than n samples from Q. Thus Claim 2 implies that with $2n$ total observations to allocate across the sources X_1, X_2 and X_3, an even division between X_1 and X_2 is better than repeated acquisition of X_3. And if the total number of observations is an odd number, the optimal allocation should at most acquire one observation of X_3. These imply that repeated and exclusive observation of X_3 is not an efficient long-run outcome.

Nonetheless, we show next that there is an open set of prior beliefs given which all agents acquire signal X_3, so that the set $\{X_3\}$ constitutes a "learning trap." The logic is similar to the example in Section 2.1. Suppose the prior belief is such that ω and p are independent, and that p is equally likely to be $p_1 = 0.9$ or $p_2 = 0.1$. Due to independence, signal X_2 alone is uninformative about ω. Moreover, the uniform prior on p implies that observation of X_1 alone does not change the agent's belief about ω. Thus the first agent myopically acquires X_3. But since this acquisition maintains independence and the uniform prior on p, *every* agent acquires X_3 starting from this prior (regardless of signal realizations).

The result below generalizes this argument to an open set of priors. To state the result, we note that the belief over (ω, p) can be summarized by four numbers:

$$u = \mathbb{P}\{\omega_1, p_1\}, \quad v = \mathbb{P}\{\omega_1, p_2\}, \quad y = \mathbb{P}\{\omega_2, p_1\}, \quad z = \mathbb{P}\{\omega_2, p_2\}.$$

Claim 3. *Suppose the prior belief satisfies*

$$\frac{1}{2} < \frac{u}{v} < 2 \quad \text{and} \quad \frac{1}{2} < \frac{y}{z} < 2. \tag{29}$$

Then every agent chooses to acquire X_3.

We conjecture that from all prior beliefs, agents will almost surely eventually concentrate their acquisitions on the set $\{X_3\}$ or $\{X_1, X_2\}$. These long-run outcomes constitute "complementary sets" in the sense that repeated acquisition of signals from these sets allows for complete learning of ω (almost surely), and moreover all signals in these sets are required to learn ω.[62] Indeed, the sets are "strongly complementary" in the sense that swapping out a single signal cannot improve the long-run rate of learning. These observations are suggestive that some of our key notions and results may be generalized beyond the normal informational environment that we consider in the main text.

G.1 Proofs of the Claims

Proof of Claim 2. Note that both pairs of signals share the signal space $\{aa, ab, ba, bb\}$. Their signal structures can be summarized as the following row-stochastic matrices, where the first one corresponds to (X_1, X_2) and the second one corresponds to (X_3, X_3):

	aa	ab	ba	bb
ω_1, p_1	0.81	0.09	0.09	0.01
ω_1, p_2	0.01	0.09	0.09	0.81
ω_2, p_1	0.09	0.01	0.81	0.09
ω_2, p_2	0.09	0.81	0.01	0.09

	aa	ab	ba	bb
ω_1, p_1	4/9	2/9	2/9	1/9
ω_1, p_2	4/9	2/9	2/9	1/9
ω_2, p_1	1/9	2/9	2/9	4/9
ω_2, p_2	1/9	2/9	2/9	4/9

Let A denote the first matrix and B denote the second matrix. By direct computation,

$$A^{-1} \cdot B = \begin{array}{c|cccc} & aa & ab & ba & bb \\ aa & \frac{155}{288} & 2/9 & 2/9 & \frac{5}{288} \\ ab & \frac{5}{288} & 2/9 & 2/9 & \frac{155}{288} \\ ba & \frac{5}{288} & 2/9 & 2/9 & \frac{155}{288} \\ bb & \frac{155}{288} & 2/9 & 2/9 & \frac{5}{288} \end{array}$$

Crucially, all entries are positive. Thus this row-stochastic matrix specifies a garbling that takes the pair (X_1, X_2) to (X_3, X_3). □

Proof of Claim 3. Consider observing the realization $X_3 = a$. By Bayes' rule, the posterior belief is summarized by the four numbers

$$(u_{3a}, v_{3a}, y_{3a}, z_{3a}) = \frac{1}{2u + 2v + y + z} \cdot (2u, 2v, y, z).$$

Similarly the realization $X_3 = b$ would lead to the posterior belief

$$(u_{3b}, v_{3b}, y_{3b}, z_{3b}) = \frac{1}{u + v + 2y + 2z} \cdot (u, v, 2y, 2z).$$

[62] We believe that the sets are also complementary as per our Definition 2, using a suitably generalized notion of informational value. This is left for future work.

Note that the ratios $\frac{u}{v}$ and $\frac{y}{z}$ are unchanged after either signal realization; that is, the condition (29) is preserved after acquisitions of X_3. Thus, it suffices to show that under this condition, the *first* agent prefers signal X_3 to both X_1 and X_2.

To do this, we will compare the expected payoffs from acquiring each of the three signals. Note that each signal has two possible realizations and leads to two possible posterior beliefs about ω. We show below that the two posterior beliefs induced by X_3 are *more extreme* than the posterior beliefs induced by any realization of X_1 or X_2. This implies that X_3 is better than X_1 and X_2 for *every* decision problem based on the binary state ω (since the indirect utility function is convex in the posterior belief).

More formally, observe from the preceding calculations that the signal X_3 induces two possible likelihood ratios between $\omega = \omega_1$ and $\omega = \omega_2$, which are $\frac{2u+2v}{y+z}$ and $\frac{u+v}{2y+2z}$. On the other hand, the different realizations of signal X_1 or X_2 induce likelihood ratios $\frac{9u+v}{y+9z}$, $\frac{u+9v}{9y+z}$, $\frac{9u+v}{9y+z}$ and $\frac{u+9v}{y+9z}$. We will show that all four of these likelihood ratios lie strictly between $\frac{u+v}{2y+2z}$ and $\frac{2u+2v}{y+z}$.

We focus on the first of the four, $\frac{9u+v}{y+9z}$, although essentially the same argument applies to the others. The lower bound

$$\frac{9u+v}{y+9z} > \frac{u+v}{2y+2z}$$

is equivalent to $17uy + vy + 9uz > 7vz$. This follows easily from the assumption that $u > v/2$ and $y > z/2$. The upper bound

$$\frac{9u+v}{y+9z} < \frac{2u+2v}{y+z}$$

is equivalent to $7uy < 17vz + vy + 9uz$, which similarly follows from $v > u/2$ and $z > y/2$. This completes the proof. □

H Comparison with Börgers, Hernando-Veciana and Krahmer (2013)

Our definition of complementary sets mirrors the constructions in Börgers, 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 Börgers, Hernando-Veciana and Krahmer (2013) focus on pairs. Our generalization to sets can be understood in either of two ways. The proposed Definition 2 requires each *pair of subsets* that partition the whole set to be complementary. In this way, we generalize from "two complementary signals" to "two complementary sets." For example, the sources $X_1 = \omega + b_1 + \varepsilon_1$, $X_2 = b_1 + b_2 + \varepsilon_2$, and $X_3 = b_2 + \varepsilon_3$ are complementary, since access to the set $\{X_1, X_2\}$ is complementary to access to $\{X_3\}$ (likewise for the other combinations).

Alternatively, one might consider a set to be complementary if *all* of the pieces combine to enhance the whole. For example, the above sources X_1, X_2, X_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). Proposition 1 shows these two perspectives are formally equivalent.

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

Finally, Börgers, Hernando-Veciana and Krahmer (2013) value information based on its contribution to decision problems, while the *val* function that we use is more statistical in nature (based on asymptotic improvements to belief precision).[63] Like the definition in Börgers, Hernando-Veciana and Krahmer (2013), however, our notion of complementarity does turn out to be uniform across prior beliefs.

[63]That being said, Proposition 1 shows that our definition is robust to monotone transformations of the *val* function.