

Optimal Myopic Information Acquisition*

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

April 8, 2018

Abstract

A decision-maker repeatedly samples from many flexibly correlated information sources and takes actions. His payoff depends on the actions taken as well as an unknown state. *Myopic* information acquisition corresponds to (history-independent) acquisition of signals that maximize next-period precision of beliefs. We identify a canonical class of informational environments (normal-linear), in which the optimal dynamic information acquisition rule is myopic *from period 1* when agents sample sufficiently many signals each period. For all sample sizes, the optimal rule is generically eventually myopic. These results demonstrate the possibility of robust and “simple” optimal information acquisition, and simplify the analysis of dynamic information acquisition in an environment that is commonly used in economics.

JEL codes: C44, D81, D83

Keywords: Information Acquisition, Correlation, Endogenous Attention, Myopic Choice, Robustness, Value of Information

*This paper was previously circulated under the title of “Dynamic Information Acquisition from Multiple Sources.” We thank Yash Deshpande, Johannes Hörner, Michihiro Kandori, Elliot Lipnowski, George Mailath, Juuso Toikka, Anton Tsoy, and Muhamet Yildiz for helpful comments. We are especially grateful to Drew Fudenberg, Eric Maskin and Tomasz Strzalecki for their constant guidance and support.

[†]Department of Economics, University of Pennsylvania

[‡]Department of Economics, Harvard University

[§]Microsoft Research

Contents

1	Introduction	4
2	Preliminaries	7
2.1	Model	7
2.2	Interpretations	8
3	(Eventual) Optimality of Myopic Rule	10
3.1	Myopic Information Acquisition	10
3.2	Main Results	10
4	Discussion	13
4.1	Intuition for Main Results	13
4.2	Precision vs. Correlation	16
4.3	How Important is Normality?	17
5	Games with Dynamic Information Acquisition	18
6	Extensions	20
6.1	Endogenous Learning Intensities	20
6.2	Multiple Payoff-Relevant States	21
6.3	Continuous Time	21
7	Related Literature	22
8	Conclusion	23
9	Appendix	25
9.1	Preliminary Results	25
9.1.1	Posterior Variance Function	25
9.1.2	The Matrix Q_i	26
9.1.3	Order Difference Lemma	26
9.2	Dynamic Blackwell Comparison	27
9.2.1	The Lemma	27
9.2.2	Proof of Lemma 3	29
9.3	Proof of Theorem 1 (Large Block of Signals)	30

9.4	Proof of Theorem 2 (Separable Environments)	32
9.5	Preparation for the Proof of Theorem 3	32
9.5.1	Switch Deviations	32
9.5.2	Asymptotic Characterization of Optimal Strategy	34
9.6	Proof of Theorem 3 (Generic Eventual Myopia)	35
9.6.1	Outline of the Proof	35
9.6.2	Equivalence at Almost All Times	36
9.6.3	A Simultaneous Diophantine Approximation Problem	40
9.6.4	Monotonicity of t -Optimal Divisions	42
9.6.5	Completing the Proof of Theorem 3	43
9.7	Proof of Proposition 1 (Bound on B)	44
9.7.1	Preliminary Estimates	44
9.7.2	Refined Asymptotic Characterization of $n(t)$	44
10	Online Appendix	50
10.1	Applications of Results from Section 5 (Multi-player Games)	50
10.1.1	Beauty Contest	50
10.1.2	Strategic Trading	52
10.2	Example in Which $n(t)$ is Not Monotone Even for Large t	52
10.3	Additional Result for $K = 2$	53
10.4	Eventual Optimality of the Myopic Strategy	54

1 Introduction

Consider the following setting: A Bayesian decision-maker (DM) takes actions each period, and his payoff depends on the sequence of actions as well as an unknown state. He has access to many correlated information sources, and learns about the state via sequential sampling. What is the optimal rule for information acquisition?

Neglecting dynamic considerations, a simple strategy is to acquire at each period the signal (or the set of signals) that maximally reduces uncertainty about the payoff-relevant state. We refer to this as the *myopic* (or greedy) rule, as it is the optimal rule if the DM mistakenly believes each period to be the last possible period of information acquisition.

Surprisingly, this myopic rule turns out to possess strong optimality properties in a canonical class of informational environments:

Suppose that the DM’s prior over the states, and also the signals, are jointly normal. Our first result is that if signals are acquired in sufficiently large blocks each period, then myopic information acquisition is optimal from period 1. We provide a sufficient condition on the required size of the block; this condition depends on primitives of the informational environment but not on the payoff function. Our second result characterizes a condition on the prior and signal structure, given which myopic information acquisition is optimal from period 1 for *any* block size. In both of these cases, the optimal information acquisition strategy can be exactly and simply characterized. Additionally, instead of solving for the optimal decision rule and information acquisition strategy jointly (as would otherwise be required), our results show that one can separate these two problems, and solve for the optimal decision rule as if information acquisition were exogenous.

Our final main result says that for generic signal structures, for any block size, the optimal strategy proceeds by myopic acquisition after finitely many periods. We note that these results hold across all payoff functions (and in particular, independently of discounting); thus, myopic information acquisition is (eventually) “robustly” best.

Why does the myopic rule perform so well? The main inefficiency of myopic planning is that it neglects potential complementarities across signals. A signal that is individually uninformative can be very informative when paired with other signals; thus, repeated (greedy) acquisition of the best single signal need not result in the best sequence of signals. A key observation is that complementarities across signals are not intrinsic to the underlying signal structure; they depend on the agent’s beliefs over the unknown states. Thus, as the DM’s beliefs about the states evolve, so too do his perceptions of the correlations across signals. We show that as information accumulates, these perceived correlations wash out. Specifi-

cally, not only does the DM learn about each state individually (which by itself does *not* imply optimality of the myopic rule), but the signals effectively “de-correlate.” At the limit, all signals are independent and the dynamic problem is “separable”; at late periods, we have approximate separability, which is still sufficient for the myopic rule to be optimal.

This intuition does not rely on the assumption of a normal environment. Eventual de-correlation of the available signals—the core of our analysis—derives from a Bayesian version of the Central Limit Theorem, which holds for arbitrary signal distributions. This suggests the possibility of generalizing our eventual optimality result. Specifically, we conjecture that even with general signals, the optimal rule eventually proceeds myopically when we restrict to certain decision problems (e.g. prediction of the payoff-relevant state). However, the stronger statements that we make—that the optimal rule is immediately myopic given sufficiently many signals, and that optimality obtains for all payoff functions—do rely on properties of the normal environment.¹

This mechanism we identify is quite different from the one underlying a classic, and spiritually related, result from the experimentation literature. In general “learning by experimentation” settings, myopic behavior is eventually near-optimal: in the long run the DM’s beliefs converge, so the value of exploration (i.e. learning) becomes second-order relative to the value of exploitation of the perceived best arm.² In contrast, learning is never of second-order importance in our setting, as there is by design no exploitation incentive. To see this, recall that in the classic multi-armed bandit problem (Gittins, 1979; Bergemann and Välimäki, 2008), actions play the dual role of influencing the evolution of beliefs and also determining flow payoffs. In our setting (which does not fall into the multi-armed bandit framework), *signal choices* influence the evolution of beliefs, while *actions* separately determine (unobserved) payoffs. Myopic signal choices become optimal in our framework because *they maximize the speed of learning*, and not because they optimize a tradeoff between learning and payoff. We mention also that the myopic strategy is *immediately* optimal in multi-armed bandit problems only under very restrictive assumptions (Berry and Fristedt, 1988; Banks and Sundaram, 1992).

Our results simplify the analysis of optimal dynamic information acquisition in an informational environment (normal-linear) that is commonly used in economics. To further

¹Specifically, we use that posterior variances do not depend on signal realizations, and also that one-dimensional normal signals admit a Blackwell ordering.

²Easley and Kiefer (1988) and Aghion et al. (1991) show that if there is a unique myopically optimal policy at the limiting beliefs, then the optimal policies must converge to this policy. In our setting, every policy (signal choice) is trivially myopic at the limiting beliefs (a point mass at the true parameter), so we do not have uniqueness and cannot use this argument to identify long-run behavior.

facilitate application of our results, we extend our environment to a multi-player setting in which individuals privately acquire information before playing a one-shot game at a random final period. This extension connects our results to a literature on games with Gaussian information (Hellwig and Veldkamp, 2009; Myatt and Wallace, 2012; Colombo et al., 2014; Lambert et al., 2017).³ We present corollaries of our main results, and use these to extend results from Hellwig and Veldkamp (2009) and Lambert et al. (2017) in an online appendix.

We conclude by demonstrating extensions to environments with choice of information “intensity” (the number of signals to acquire each period), to multiple payoff-relevant states (for a specific decision problem of prediction) and to a continuous-time setting.

Our work primarily builds on a large literature about optimal dynamic information acquisition (Moscarini and Smith, 2001; Fudenberg et al., 2017; Che and Mierendorff, 2017; Mayskaya, 2017; Steiner et al., 2017; Hébert and Woodford, 2017; Zhong, 2017) as well as a related literature on sequential search (Wald, 1947; Arrow et al., 1949; Weitzman, 1979; Callander, 2011; Ke and Villas-Boas, 2017; Bardhi, 2017). In contrast to an earlier focus on optimal stopping and choice of signal precision, our framework characterizes choice between different *kinds* of information, as in the concurrent work of Fudenberg et al. (2017) (where the sources are two Brownian motions), and Che and Mierendorff (2017) and Mayskaya (2017) (where the sources are two Poisson signals).⁴ Compared to this work, we allow for many (i.e. more than two) sources with flexible correlation.⁵

Under another common approach, the DM chooses from completely flexible information structures at entropic (or more generally, “posterior-separable”) costs, such as in Steiner et al. (2017), Hébert and Woodford (2017) and Zhong (2017). Compared to these papers, our DM has access to a *prescribed* and limited set of signals.⁶

Finally, acquisition of Gaussian signals whose means are linear combinations of unknown

³For games of information acquisition beyond the Gaussian setting, see e.g. Persico (2000), Bergemann and Välimäki (2002), Yang (2015) and Denti (2017). All of these papers restrict to a single signal choice.

⁴Che and Mierendorff (2017) and Mayskaya (2017) consider choice between two Poisson signals, each of which provides evidence towards/against a binary-valued state. The Poisson model (with a finite state space) is more suited to applications such as learning about whether a defendant is guilty or innocent, while the Gaussian model describes for example learning about the (real-valued) return to an investment.

⁵Callander (2011) also emphasizes the correlation between different signals. But the signals in Callander (2011) are related by a Brownian motion path, which yields a special correlation structure. Similar models are studied in Garfagnini and Strulovici (2016) and in Bardhi (2017).

⁶In Section 6.1, we do allow the DM to also control the intensity of information acquisition by endogenously choosing how many signals to acquire in each period. But even in that extension, we assume that the incurred information cost is a function of the number of observations. This is analogous to Moscarini and Smith (2001) and is distinguished from the above papers that measure information cost based on belief changes.

states appears previously in the work of Meyer and Zwiebel (2007) and Sethi and Yildiz (2016). In particular, Sethi and Yildiz (2016) characterizes the long-run behavior of a DM who myopically acquires information from experts with independent biases.

2 Preliminaries

2.1 Model

Time is discrete. At each time $t = 1, 2, \dots$, the DM chooses from among K information sources, and additionally chooses an action a_t from a set A_t . We suppose that information choices occur before action choices, but this is not important for the results.

The DM's payoff is an arbitrary function $U(a_1, a_2, \dots; \theta_1)$, which depends on the sequence of action choices in addition to a payoff-relevant state $\theta_1 \in \mathbb{R}$. We assume that payoffs are realized only at an (exogenously or endogenously determined) end date; thus, the information sources described below are the only channel through which the DM learns. This assumption distinguishes our model from multi-armed bandit problems, see Section 7 for further discussion.

Stylized cases of such decision problems include:

Exogenous Final Date. An action is taken just once at a final period T that is determined by an arbitrary distribution over periods.⁷ The DM's payoff is $U(a_1, a_2, \dots; \theta_1) = u_T(a_T, \theta_1)$ where T is the (random) final time period, and a_T is the action chosen in that period. The time-dependent payoff function $u_T(a_T, \theta_1)$ may incorporate discounting.

Endogenous Stopping with Per-Period Costs. Take each action a_t to specify both the decision of whether to stop, and also the action to be taken if stopped. The DM's payoff is $U(a_1, a_2, \dots; \theta_1) = u_T(a_T, \theta_1)$ where T is the (endogenously chosen) final time period, and a_T is the action chosen in that period. The payoff-function $u_T(a_T, \theta_1)$ may incorporate discounting and/or a per-period cost to signal acquisition. Costs must be fixed across sources in a given period, but can vary across periods.⁸

Apart from the decision problem, there are K information sources, which depend on the unknown and persistent state vector $\theta = (\theta_1, \dots, \theta_K)' \sim \mathcal{N}(\mu^0, V^0)$.⁹ This vector includes the

⁷Special cases include *geometric discounting*, in which every period (conditional on being reached) has a constant probability of being final, as well as *Poisson arrival* of the final period.

⁸See e.g. Fudenberg et al. (2017) and Che and Mierendorff (2017) for recent models with constant waiting cost per period.

⁹Here and later, we exclusively use the apostrophe to denote vector or matrix transpose.

payoff-relevant unknown θ_1 and additionally $K - 1$ *payoff-irrelevant* unknowns $\theta_2, \dots, \theta_K$. The role of these auxiliary states is to permit correlations across the information sources conditional on θ_1 ; this allows, for example, for the sources to be afflicted by common and persistent biases. In each period, the DM chooses B sources (allowing for repetition), where $B \in \mathbb{N}^+$ is interpreted as a fixed time/attention constraint (see Section 6.1 for extension to endogenous choice of B). Choice of source $k = 1, \dots, K$ produces an observation of

$$X_k = \langle c_k, \theta \rangle + \epsilon_k, \quad \epsilon_k \sim \mathcal{N}(0, \sigma_k^2)$$

where the coefficient vectors $c_k = (c_{k1}, \dots, c_{kK})'$ and signal variances σ_k^2 are fixed (and known), but the Gaussian error terms are independent across realizations and sources. Throughout, we use C to denote the matrix of coefficients whose k -th row is c_k' .

We impose the following assumption on the informational environment:

Assumption 1 (Full Rank and Exact Identifiability). *The matrix C has full rank, and no proper subset of row vectors of C spans the coordinate vector e_1' . Equivalently, the inverse matrix C^{-1} exists, and its first row consists of non-zero entries.*

Heuristically, this means that the DM *can and must* observe each source infinitely often to recover the value of the payoff-relevant state θ_1 . In other words, θ_1 is identified from the set of all available sources, but not from any proper subset. Since in this paper we restrict the number of sources to be the same as the number of unknown states, the above assumption is generically satisfied.¹⁰ We refer the reader to Liang and Mu (2017) for additional results for the general case with an arbitrary number of available sources (possibly “overabundant information”).¹¹

2.2 Interpretations

We provide below several interpretations of this framework.

News Sources with Correlated Biases. On election day T , a DM will choose which of two candidates I and J to vote for, where $\theta_1 = v_I - v_J$ is the difference in their respective qualities. In each period up to and including T , he can acquire information from different news sources. All sources provide biased information, and moreover the biases are correlated

¹⁰Throughout, “generic” means with probability 1 for randomly drawn coefficient matrices C .

¹¹With more source than states, the DM can learn the underlying states by restricting attention to any subset of at least K sources. So a new question arises as to which subset the DM will focus on. We provide a complete characterization of asymptotic information acquisitions in Liang and Mu (2017).

across the sources.¹² As the DM acquires information, he learns not only about the payoff-relevant state θ_1 , but also how to de-bias (and aggregate) information from the various news sources.

Attribute Discovery. A consumer good or investment option has K unknown attributes $\tilde{\theta}_1, \dots, \tilde{\theta}_K$. Its value θ_1 is some arbitrary linear combination of these attribute values. For example, the DM may want to learn the value of a conglomerate composed of several companies, where each company i is valued at $\tilde{\theta}_i$ and the value of the conglomerate is $\theta_1 := \alpha_1 \tilde{\theta}_1 + \dots + \alpha_K \tilde{\theta}_K$. The DM has access to (noisy) observations of different linear combinations of the attributes; for example, he might have access to evaluations of each $\tilde{\theta}_i$ individually.¹³ At some endogenously chosen end time, the DM decides whether or not to invest in the conglomerate.

Sequential Polling. A polling organization wants to predict the average perspective in the population towards an issue. Perspectives differ across demographic groups, of which there are K in the population. Perspectives in demographic group k follow a normal distribution with unknown mean μ_k and known variance σ_k^2 . The average perspective is $\theta_1 := \sum_k p_k \mu_k$, where p_k is the fraction of the population in the k -th demographic group. It is not feasible to directly sample individuals according to p_k , but there exist other (non-representative) polling technologies. For example, the group can poll via a liberal website, which over-samples from demographic groups that are more predominantly liberal. Each period, the polling organization chooses one of these technologies for acquiring another opinion sample. The organization posts a prediction for θ_1 each period, and its payoff is the average prediction error across some fixed number of periods.

Intertemporal Investment. Each action a_t is a decision of how to allocate capital between consumption, and two investment possibilities: a liquid asset (bond), and an illiquid asset (pension fund). The return to the liquid asset is known: 1 dollar saved today is worth e^r dollars tomorrow. The return to the illiquid asset is unknown, and it is the payoff-relevant state in the worker's problem; that is, every dollar invested today in the pension fund deterministically yields e^{θ_1} dollar(s) tomorrow. The worker works for T periods, and in each of these periods he learns about θ_1 (from some information sources) and then allocates his wealth across consumption, saving and investment. In period $T + 1$, the worker retires and receives all the returns from his investments into the illiquid asset. His objective is to maximize the aggregated sum of his discounted consumption utilities and the payoff after

¹²For example, there may be some collection of relevant audiences and donors, where different media sources place (possibly different) weights on catering to each of these groups.

¹³This model can be rewritten in our framework above, where the state vector is $(\theta_1, \tilde{\theta}_2, \dots, \tilde{\theta}_K)$.

retirement.¹⁴

3 (Eventual) Optimality of Myopic Rule

3.1 Myopic Information Acquisition

We now define the notion of *myopic* information acquisition that is the focus of this paper.

In our setting, a strategy consists of an information acquisition strategy and a decision strategy. An *information acquisition strategy* captures the DM’s signal choices given any history of observations. Formally, it is a measurable map from possible histories of signal realizations to multi-sets of B signals. Similarly, a *decision strategy* maps histories to actions.

Definition 1. *Say that an information acquisition strategy is myopic, if at every next period, it prescribes choosing the B signals that (combined with the history of observations) lead to the lowest posterior variance about θ_1 .*

Note that the B signals which minimize the posterior variance also Blackwell dominate any other multi-set of B signals (see e.g. Hansen and Torgersen (1974)). Thus, myopic acquisition is optimal if the current period is the last chance for information acquisition, independent of the payoff function. Our results below characterize the optimal information acquisition strategy, and in particular its relation to the myopic information acquisition strategy. We do not pursue a characterization of the optimal decision strategy (which will in general depend on the payoff function), although we point out one application of our main results towards simplification of this characterization.

3.2 Main Results

We present three results regarding optimality of the myopic information acquisition rule: Theorem 1 says that myopic information acquisition is optimal from period 1 if B (the number of signals acquired each period) is sufficiently large. Our next two results hold for arbitrary B : Theorem 2 provides a sufficient condition on the prior and the coefficient

¹⁴An important assumption of this example is that the return to investment is deterministic and only observed at the end. However, our model and results extend to a situation where there are “free” signals arriving each period that do not count toward the capacity constraint B . By considering the realized log return as a particular free signal, the extension of our model covers the case where investment returns are stochastic and the DM observes past return realizations.

matrix C under which myopic information acquisition is optimal from period 1, and Theorem 3 states that the optimal rule is *eventually* myopic in generic environments.

Theorem 1 (Immediate Optimality under Many Observations). *Fix any prior and signal structure, and suppose B is sufficiently large. Then the DM has an optimal strategy that acquires information myopically.*¹⁵

Thus, it is without loss to assume that the DM chooses the most informative B signals in the first period based on his prior, then chooses the most informative B signals in the second period based on his updated posterior, and so on. Note that since posterior variances are independent of signal realizations, and we have assumed that there is no feedback from actions, the above myopic strategy is *history-independent*, and can be represented as a deterministic signal path. Consequently, instead of solving for the optimal decision strategy and information acquisition strategy jointly (as would otherwise be required), we can take information as exogenous, and then solve for the optimal decision strategy with respect to this exogenous stream of information.¹⁶

We additionally mention that Theorem 1 can be strengthened to optimality of myopic information acquisition at all histories, including those in which the DM has previously deviated from the myopic rule. Finally, see Section 4.2 for a precise bound for how large B must be.

Our next two results hold for arbitrary B . First, exact optimality again obtains in the following class of environments. Let $f(q_1, \dots, q_K)$ denote the DM’s posterior variance about θ_1 , updating from his prior and q_i observations of each signal X_i .

Definition 2. *The informational environment $(V^0, C, \{\sigma_i^2\})$ is separable if there exist convex functions g_1, \dots, g_K and a strictly increasing function F such that*

$$f(q_1, \dots, q_K) = F(g_1(q_1) + \dots + g_K(q_K)).$$

Note that separability is not defined directly on the primitives of the informational environment $(V^0, C, \{\sigma_i^2\})$, as they are based instead on the posterior variance function f .

¹⁵Without further assumptions on the payoff function U , we cannot assert *strict* optimality of the myopic information acquisition strategy. For instance, this would not be true if there exists a “dominant” action sequence that maximizes $U(a_1, a_2, \dots; \theta_1)$ for every value of θ_1 . But in most other cases, strictly more precise beliefs do lead to strictly higher expected payoffs, which implies unique optimality of the myopic rule.

¹⁶An application of this two-step approach (in continuous time) appears in the concurrent work of Fudenberg et al. (2017), Section 3.5. See Section 6.3 for a brief discussion of how our model extends to continuous time.

Nevertheless, f can be directly computed from these primitives (and thus is not itself an endogenous object). Intuitively, separability ensures that observing signal i does not change the relative value of other signals, but strictly decreases the marginal value of signal i relative to every other signal.

Theorem 2 (Immediate Optimality in Separable Informational Environments). *Suppose the informational environment is separable. Then for every $B \in \mathbb{N}^+$, the DM has an optimal strategy that acquires information myopically.*

Separability encompasses several classes of informational environments that are of independent interest. For example:

Example 1 (Orthogonal Signals). The DM’s prior is standard Gaussian ($V^0 = \mathbf{I}_K$), and the row vectors of C are orthogonal to one another.¹⁷

Example 2 (Multiple Biases). There is a single payoff-relevant state $\theta_1 \sim \mathcal{N}(0, v_1)$. The DM has access to observations of $X_1 = \theta_1 + \theta_2 + \dots + \theta_K + \epsilon_1$, where each θ_i ($i > 1$) is a persistent “bias” independently drawn from $\mathcal{N}(0, v_i)$, and $\epsilon_1 \sim \mathcal{N}(0, \sigma_1^2)$ is a noise term i.i.d. over time. Additionally, he can learn about each bias θ_i ($i > 1$) by observing $X_i = \theta_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$.¹⁸

In all remaining cases, optimal signal choices are eventually myopic.

Theorem 3 (Generic Eventual Myopia). *Fix any prior covariance V^0 and signal variances $\{\sigma_i^2\}_{i=1}^K$. For generic coefficient matrices C , there exists a time $T^* \in \mathbb{N}$ that depends only on the informational environment. For this T^* , and for any B and any decision problem, the DM has an optimal strategy that acquires information myopically after T^* periods.*

That is, at all late periods, the optimal signal acquisitions are those that maximally reduce posterior variance in the given period. We can further show that these eventual signal choices are the same as those acquired under the fully myopic signal path. In other words, it not only holds that the optimal information acquisition strategy is eventually myopic, but that acquiring signals myopically (from period 1) is also eventually optimal. See Appendix 10.4 for details.

¹⁷This is because the signals are independent from each other, see also Example 2 in Figure 1.

¹⁸The DM’s posterior variance about θ_1 is given by

$$f(q_1, \dots, q_K) = v_0 - \frac{v_0^2}{v_0 + \frac{\sigma_1^2}{q_1} + \sum_{i=2}^K \left(v_i - \frac{v_i^2}{v_i + \sigma_i^2/q_i} \right)}.$$

This can be written in the separable form.

4 Discussion

4.1 Intuition for Main Results

Define a t -optimal “division vector” $n(t)$ to be an allocation of t observations over the K signals that minimizes posterior variance about the payoff-relevant state; that is,

$$n(t) = (n_1(t), \dots, n_K(t)) \in \underset{(q_1, \dots, q_K): q_i \in \mathbb{Z}_+, \sum_{i=1}^K q_i = t}{\operatorname{argmin}} f(q_1, \dots, q_K)$$

where $n_i(t)$ is the number of observations allocated to signal i . Again applying Hansen and Torgersen (1974), this set of signals maximizes expected payoffs to any decision (among all one-shot allocations of t signals).

If the count $n_i(t)$ increases monotonically in t for all signals i , then the divisions $(n(t))_{t \geq 1}$ can be achieved under *some* sequential sampling rule, and this sampling rule corresponds to myopic information acquisition.¹⁹ Additionally, no trade-offs are necessary across periods, since uncertainty about the payoff-relevant state is minimized *at every period*. This suggests that the resulting path of signal acquisitions is optimal for any objective, although we need to argue also that this history-independent path cannot be improved upon by any history-dependent strategy. Towards this end, we establish a “dynamic Blackwell” lemma, which generalizes a result from Greenshtein (1996); see Section 7 for further discussion. Thus, optimality of myopic information acquisition directly follows from existence of a sequence $(n(t))_{t \geq 1}$ of monotonically increasing division vectors.

Example 2 of Figure 1 is a setting in which such a sequence exists: Any allocation that is as close to balanced across the signals as possible is t -optimal; thus, it is possible to find t -optimal division vectors that evolve monotonically.

Theorem 2 generalizes Example 2 to a *class* of environments in which $n(t)$ evolve monotonically. Informally, the dynamic problem is “separable” in this case, and the optimal choice at every period can be evaluated without regard to future periods.

In general, the division vectors $n(t)$ need not be monotonic, and Example 1 of Figure 1 is one such setting: $n(5) = (4, 1, 0)$ but $n(6) = (3, 2, 1)$, so the optimal division vector for $t = 6$ is not achievable from the optimal division vector for $t = 5$.²⁰ This occurs because

¹⁹Proceed by induction: in the first period the myopic rule chooses the signal that minimizes posterior variance. In the second period, he again wants to minimize posterior variance at the given period; since the division chosen by the totally optimal rule is best and feasible given the period 1 choice, this is what myopic information acquisition will yield. So on and so forth.

²⁰To see these are the t -optimal divisions, we calculate that at $t = 5$, $f(4, 1, 0) = \frac{11}{23} < \frac{14}{29} = f(3, 1, 1) = f(3, 2, 0)$. Whereas at $t = 6$, $f(3, 2, 1) = \frac{5}{11} < \frac{17}{37} = f(4, 1, 1) = f(4, 2, 0)$.

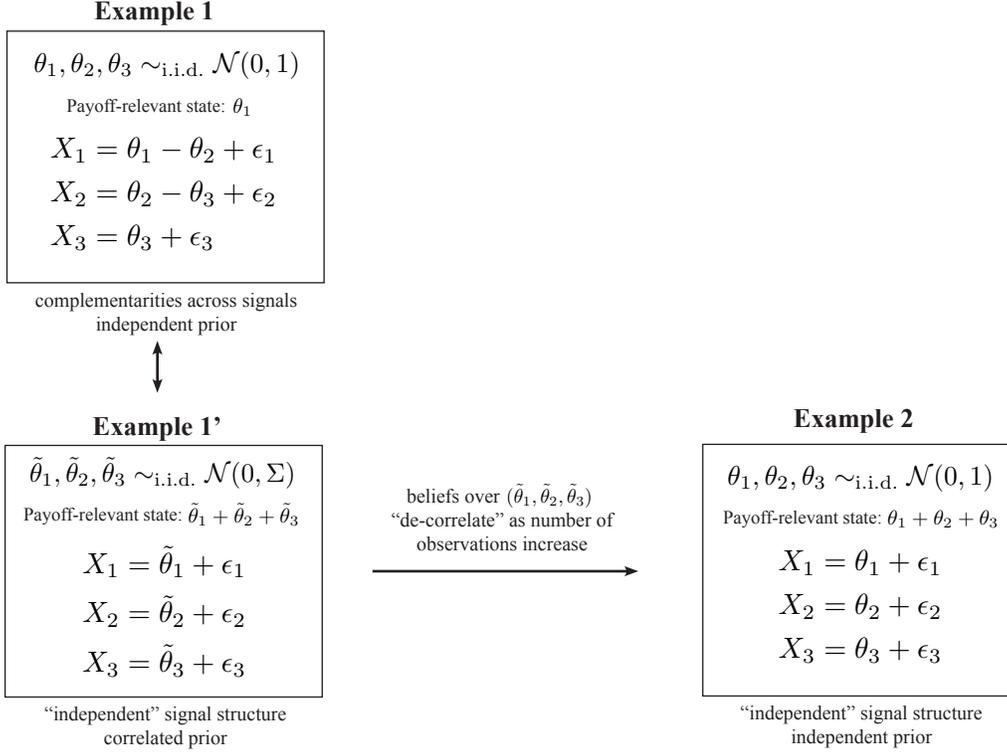


Figure 1

X_2 and X_3 are strong complements (observation of either increases the value of observation of the other). Observation of a signal can thus *reverse* the ranking of other signals: for example, X_1 may be more informative than X_2 at a given history, but less informative after an additional observation of X_3 . This reversal was not possible in Example 2.

A key step in our proofs is to show that environments which start off like Example 1 necessarily become “like” Example 2 (as observations of each signal accumulate). To facilitate this comparison, rewrite Example 1 in the following way: Define a new set of states $\tilde{\theta}_1 = \theta_1 - \theta_2$, $\tilde{\theta}_2 = \theta_2 - \theta_3$, and $\tilde{\theta}_3 = \theta_3$. The original prior over $(\theta_1, \theta_2, \theta_3)$ defines a new prior over $(\tilde{\theta}_1, \tilde{\theta}_2, \tilde{\theta}_3)$, and the original payoff-relevant state can be re-expressed in the new states as $\theta_1 = \tilde{\theta}_1 + \tilde{\theta}_2 + \tilde{\theta}_3$. The signal structure in Example 1' is the same as in Example 2, with the crucial difference that the prior is correlated in Example 1' and independent in Example 2. Given sufficiently many observations of each signal in Example 1', however, the DM's posterior beliefs over $(\tilde{\theta}_1, \dots, \tilde{\theta}_K)$ become *almost independent*. That is, not only does learning about each $\tilde{\theta}_k$ occur (which will happen naturally as information accumulates), but the states $\tilde{\theta}_k$ “de-correlate.” Thus, eventually the signals can be viewed as approximately

independent.

The above heuristic statements about de-correlation roughly follow from a Bayesian version of the Central Limit Theorem. We establish a technical lemma (Lemma 2 in Appendix 9.1) that strengthens this with a comparison of the value of different signals. Specifically, we characterize the externality that observation of a given signal X_i has on the marginal value of future observations. We show that the effect of observing X_i on the value of future observations of X_i (eventually) far outweighs its effect on future observations of any X_j , $j \neq i$. That is, like in Example 2, the only effect that observation of X_i can have on the ranking of signals is by reducing the position of signal i . This property is immediate when the covariance matrix is the identity, and holds also when the covariance matrix is close to diagonal (almost independent). Thus, in the limit in which signals are viewed as independent, we have a setting much like Example 2, and the problem is “near-separable.”

Now finally, observe that the transformation we used to rewrite Example 1 as Example 1’ was not special. Indeed, we can rewrite any signal structure using the following rule: For each signal k , define a new state $\tilde{\theta}_k = \langle c_k, \theta \rangle$, so that the signal X_k is simply $\tilde{\theta}_k$ plus independent Gaussian noise. The payoff-relevant state can be rewritten as a linear combination of the transformed states $\tilde{\theta}_1, \dots, \tilde{\theta}_K$, and the original prior defines a new prior over $(\tilde{\theta}_1, \dots, \tilde{\theta}_K)$.

The arguments above form the core of our proofs of Theorem 1 and Theorem 3. They establish that $n(t)$ eventually evolves approximately monotonically.²¹ Our main results, however, have that the optimal rule is (eventually) *exactly* myopic. We introduce two different conditions that imply this.

Our first approach is to allow for acquisition of a larger number of signals each period. We show that in transitioning from the t -optimal division to the $(t + d)$ -optimal division (where t is sufficiently large relative to d), if some signal count decreases (failing sequentiality), then every other signal count increases by at most one. Specifically, taking $d = K - 1$, we can show that $n_i(t + K - 1) \geq n_i(t)$ for every signal i at every period $t \geq T$, with T a large number depending on the informational environment. Thus, given block size $B \geq \max\{K - 1, T\}$, the division vectors $n(Bt)$ are attainable using a sequential rule from period 1. Applying the dynamic Blackwell lemma mentioned earlier, we have that the optimal strategy immediately follows $n(Bt)$, myopically. This is the intuition behind Theorem 1.

Our second approach quantifies the “typicality” of failures of monotonicity of $n(t)$. We show that at late periods t , these failures occur for a purely technical reason: The vectors $n(t)$ eventually seek to approximate an optimal limiting frequency $\lambda \in \Delta^{K-1}$ across signals,

²¹See Appendix 10.2 for an example in which $n(t)$ fails to be monotone even when we restrict to arbitrarily late periods.

but do so under an integer constraint; and the best integer approximations to λt may not be monotone.²² A key lemma demonstrates these integer approximations do (eventually) evolve monotonically for a “measure-1” set of coefficient matrices C .²³ Thus, if the optimal strategy coincides with $n(t)$ at some late period t (assuming $B = 1$ for illustration), it will follow these t -optimal divisions afterwards.

The last step is to verify that the optimal strategy will coincide with $n(t)$. We argue that if this were not true, then the DM could “deviate toward the t -optimal divisions” and (in generic environments) improve the posterior variance at every period, thus contradicting optimality of the original strategy. Hence generically, the optimal strategy will eventually coincide with $n(t)$ and subsequently follow it. This yields Theorem 3.

4.2 Precision vs. Correlation

With sufficiently many observations, the DM’s beliefs simultaneously become more precise and less correlated, and these two effects are confounded in our main results. It is tempting to think that Theorem 1 (or Theorem 3) follows from the eventual precision of beliefs. However, as our discussion above suggests, the key feature is not how precise beliefs are, but how correlated they are. Specifically, the block size B needed in Theorem 1 depends on how many observations are required for the transformed states $\tilde{\theta}_1, \dots, \tilde{\theta}_K$ to “de-correlate,” at which point complementarities across signals are weak.

Below we make this formal with a bound on B . To state our result, we define transformed states $\tilde{\theta}_k = \frac{1}{\sigma_k} \langle c_k, \theta \rangle$ (dividing through by σ_k normalizes all signal variances to 1), and let \tilde{V} denote the prior covariance matrix over these transformed states. The payoff-relevant state θ_1 can be rewritten as a linear combination of the transformed states: $\theta_1 = \langle w, \tilde{\theta} \rangle$ for some fixed payoff weight vector $w \in \mathbb{R}^K$. In the following result we assume for simplicity that $w = \mathbf{1}$ is the vector of ones, although our analysis can be easily adapted to any w .

Proposition 1. *Let R denote the operator norm of the matrix $(\tilde{V})^{-1}$.²⁴ Suppose $w = \mathbf{1}$, then $B \geq 8(R + 1)K^{1.5}$ is sufficient for Theorem 1 to hold.*

Notice that this bound is increasing in the norm of \tilde{V}^{-1} . To interpret this, suppose first that we adjust the precision of the DM’s prior beliefs over $\tilde{\theta}_1, \dots, \tilde{\theta}_K$ but fix the degree of correlation, for example by scaling \tilde{V} by a factor less than 1. Then, the norm of \tilde{V}^{-1}

²²This is indeed the case for Example 2, see Appendix 10.2 for details.

²³The lemma follows from results in Diophantine approximation theory, which studies the extent to which an arbitrary real number can be approximated by rational numbers.

²⁴The operator norm of a matrix M is defined as $\|M\|_{op} = \sup \left\{ \frac{\|Mx\|}{\|x\|} : x \in \mathbb{R}^K \text{ with } x \neq \mathbf{0} \right\}$.

increases, and a larger number of signals B is needed. This is because a more precise prior can be understood as “re-scaling” the state space by shrinking all states towards zero. Since signal noise is not correspondingly rescaled, each signal now reveals less about the states, and de-correlation takes longer.

In contrast, suppose we hold prior precision fixed and increase the degree of prior correlation. This would correspond to fixing the diagonal entries of \tilde{V} and increasing the off-diagonal entries, so that the variances about individual states are unchanged but their covariances become larger in magnitude. Then, the entire matrix \tilde{V} becomes closer to being singular, the norm of its inverse increases and a larger B is required. That is, greater correlation in the prior requires more observations to de-correlate.

Thus, having arbitrarily precise beliefs does not guarantee (immediate) optimality of the myopic rule; on the contrary, this equivalence is less likely to obtain under precise but correlated prior beliefs. Finally, we highlight that Proposition 1 only provides a sufficient condition on the block size for the myopic rule to be optimal. However, the aforementioned comparative statics hold not just for the B that we identify, but also for the smallest B that produces the result. Indeed, these comparative statics are sharp in the continuous-time limit of our model (see Section 6.3), where doubling the prior precision would also double the capacity B needed for Theorem 1 to be true.

4.3 How Important is Normality?

De-correlation. The key part of our argument is that signals eventually de-correlate. This de-correlation derives from a Bayesian version of the Central Limit Theorem, and does not rely on special properties of normality. Consider a more general setting with signals $X_i = \tilde{\theta}_i + \epsilon_i$, where the noise term ϵ_i has an arbitrary distribution with zero mean and finite variance. Then, we have that the suitably normalized posterior distribution over $(\tilde{\theta}_1, \dots, \tilde{\theta}_K)$ converges towards a standard normal distribution (so that signals are approximately independent). We thus expect that given sufficiently many observations, our previous comparisons on the value of information extend beyond normal signals.

But if we drop normality, then our results weaken in the following ways. Specifically, in working with general signal distributions, we sacrifice the potential for immediate optimality of the myopic rule, and also the generality to all intertemporal decisions.

Immediate Optimality of the Myopic Rule. For normal signals, we established a T such that given T observations of each signal, the posterior covariance matrix (over the transformed states) is almost independent. Notably, this bound holds uniformly across all his-

tories of signal realizations, thanks to the fact that posterior variances do not depend on signal realizations under normality. As mentioned above, we can use a Bayesian Central Limit Theorem to argue a similar property for other signal distributions. The difference is that the CLT gives us (near) independence *almost surely*, so that at every period t , there is still positive probability (albeit vanishing as t increases) that the normalized posterior covariance matrix is very different from the identity. This precludes us from demonstrating a block size B given which the optimal rule would be myopic *from period 1* (Theorem 1). For general signal distributions, we thus conjecture that *almost surely* the optimal rule is eventually myopic, but do not know what conditions would produce immediate optimality of the myopic rule.

General Intertemporal Payoffs. The place where we rely most heavily on normality is the statement that our results hold for *all* payoff criteria that depend only on θ_1 (and actions). Indeed, when payoff-relevant uncertainty is one-dimensional (as it is here), then all normal signals can be Blackwell-ordered based on their posterior variances. We use this fact in Section 3.1 when we define the myopic rule to maximize reduction in posterior variance. We use this fact again in Section 4.1 when we define the t -optimal divisions $n(t)$ without explicit reference to the payoff function.

Finally, while the above uses of normality are concerned with static decisions (i.e. taking an action once), we also need normality to be able to compare *signal sequences*. Generalizing Greenshtein (1996), we show that the ranking of sequences of normal signals is the same whether we consider the class of static decision problems or the broader class of intertemporal decisions. This equivalence does not hold in general, see Greenshtein (1996) for a counterexample involving Bernoulli signals.

5 Games with Dynamic Information Acquisition

We now extend our results to a multi-player setting in which individuals privately acquire information before playing a normal-form game at a random (and exogenously determined) end date.²⁵

There are N players (indexed by i), each of whom has access to a set of K signals

$$X_k^i = \langle c_k^i, \theta^i \rangle + \epsilon_k^i.$$

²⁵Reinganum (1983) considers a similar multi-agent model with private information acquisition (specifically, firms engaging in R&D before competing in oligopoly). Her model is further developed by Taylor (1995) within the context of research tournaments. However, these papers assume perfectly revealing signals and are thus distinguished from our setting.

The state vector $\theta^i = (\theta_1^i, \theta_2^i, \dots, \theta_K^i)$ represents persistent unknown states particular to player i . Noise terms ϵ_k^{it} are (normalized to) standard normal and independent across signals, players and time.

In each period up to and including the final period, each player i acquires B independent observations of his signals described above, possibly obtaining multiple (independent) realizations of the same signal. Signal choices and their realizations are both *privately observed*. The final period is determined by an exogenous distribution π . At this period, agents play a one-shot incomplete information game, where each player i 's payoff $u_i(a, \omega)$ depends on actions $a = (a^1, \dots, a^N)$ in addition to a payoff-relevant state ω .

We require that the players share a common prior over all states (ω and the player-specific state vectors $(\theta^i)_{1 \leq i \leq N}$) with the following *conditional independence* property: for each player i , conditional on the value of θ_1^i , both the payoff-relevant state ω and also the *other* players' unknown states $(\theta^j)_{j \neq i}$ are conditionally independent from player i 's own states θ^i .²⁶ This ensures that no player i infers anything about ω or about any other player j 's information beyond what he (player i) learns about θ_1^i , which makes θ_1^i the only state of interest for player i .²⁷

For concreteness, we provide examples (adapted from Lambert et al. (2017)) that do and do not satisfy conditional independence.²⁸

Example 3 (Satisfies Conditional Independence). In addition to the payoff-relevant state ω , there is a common unknown state ξ , and two player-specific unknown states b_1 and b_2 . All states are independent. Player 1 has access to observations of $\omega + \rho_1 \xi + b_1 + \epsilon_1^1$ (where ρ_1 is a constant) and $b_1 + \epsilon_2^1$. Player 2 has access to observations of $\omega + \rho_2 \xi + b_2 + \epsilon_1^2$ (where ρ_2 is a constant) and $b_2 + \epsilon_2^2$. To see that this example satisfies Conditional Independence, define $\theta_1^1 = \omega + \rho_1 \xi$ and $\theta_1^2 = \omega + \rho_2 \xi$.

Example 4 (Fails Conditional Independence). The payoff-relevant state is ω , and there is additionally a common unknown state ξ . These states are independent. Player 1 has access to noisy observations of $\omega + \xi$ only. Player 2 has access to noisy observations of both ω and ξ . Because both states ω and ξ covary with $\omega + \xi$, there is no way to define the second player's "state of interest" θ_1^2 that would satisfy Conditional Independence.

We maintain Assumption 1, so that each θ_1^i is exactly identified by the signals available

²⁶Note that we do not impose conditional independence *between* ω and the other players' states.

²⁷Conditional independence is imposed on players' prior beliefs. However, this implies conditional independence for subsequent posterior beliefs; given the value of θ_1^i , each signal is simply a linear combination of player i 's other states plus noise, thus conditional independence is preserved under updating.

²⁸Example 3 is based on Example OA.3 in Lambert et al. (2017). Example 4 is based on their Example 1.

to player i . The following result generalizes Theorem 1 and 2. (Although we do not state it here, Theorem 3 also extends.)

Corollary 1. *Suppose B is sufficiently large or the informational environment is separable. Then there exists a Nash equilibrium of this model where each player acquires information myopically.*

In fact, we show that the myopic information acquisition strategy is *dominant* in the following sense: for arbitrary opponent strategies, player i 's best response consists of acquiring signals myopically.

In Appendix 10.1, we apply the above corollary to extend results from Hellwig and Veldkamp (2009) and Lambert et al. (2017) to a setting with sequential information acquisition.

6 Extensions

6.1 Endogenous Learning Intensities

The main model imposes an exogenous capacity constraint of B signals per period. Suppose now that in each period t , the DM can choose to observe any number $N_t \in \mathbb{Z}^+$ of signal realizations (which are then optimally allocated across signals). The DM incurs a flow cost of information acquisition, modeled as $\kappa(N_t)$ for some increasing cost function $\kappa(\cdot)$ with $\kappa(0) = 0$. This framework embeds our main model if we define $\kappa(N) = 0$ for $N \leq B$ and $\kappa(N) = \infty$ for $N > B$.

We assume that the DM's payoff is $U(a_1, a_2, \dots; \theta_1) - \sum_t \delta^{t-1} \cdot \kappa(N_t)$ for some discount factor δ .²⁹ For the special case of endogenous stopping, the payoff function simplifies to

$$\delta^\tau \cdot u(a_\tau; \theta_1) - \sum_{t=1}^{\tau} \delta^{t-1} \cdot \kappa(N_t)$$

whenever the DM stops after τ periods. This is a discrete-time generalization of the framework proposed in Moscarini and Smith (2001), although our focus is on allocation of the signals instead of choice of intensity level.³⁰

Theorems 1 and 2 generalize to this setting:

²⁹Our analysis can accommodate more general payoff functions of the form $U(N_1, a_1, N_2, a_2, \dots; \theta_1)$.

³⁰Moscarini and Smith (2001) has a single state and a single signal ($K = 1$), so the DM chooses only the learning intensities N_t . Unlike Moscarini and Smith (2001), we do not characterize the optimal sequence of intensity choices $(N_t)_{t \geq 1}$, but instead show how this problem can be separated from allocation of those signals across different kinds of sources.

Corollary 2. *Suppose B is sufficiently large or the informational environment is separable. Then, even with endogenous learning intensities, the DM has an optimal strategy that chooses signals myopically.*

In the above corollary, “myopic acquisition” means the following: In any period t , given (endogenous) intensity choice N_t , the optimal acquisitions are the N_t signals that minimize posterior variance about θ_1 . We emphasize that while myopic signal choices are optimal, myopic intensity choices need not be. However, knowing that the signal choices must follow the myopic path provides a simplifying first step towards the characterization of optimal intensity levels.

Generic eventual myopia (Theorem 3) also extends, but we omit the details.

6.2 Multiple Payoff-Relevant States

In the main model, the DM’s payoff function depends on a one-dimensional state. Our results do not extend in general to payoff functions that depend on the full state vector $(\theta_1, \dots, \theta_K)$. Loosely, this is because the Blackwell ordering over signals becomes incomplete once there is multi-dimensional uncertainty; thus, even the statement that the signal which maximally reduces posterior variance is best for *static* decision problems has no analogue.

However, our results do extend for the specific problem of prediction. Specifically, suppose that at an exogenous end date (determined by an arbitrary distribution over periods), the DM is asked to predict the state vector θ . At this time, he receives a payoff of

$$-(a - \theta)' \mathbf{W} (a - \theta),$$

where \mathbf{W} is a given positive semi-definite matrix and $a \in \mathbb{R}^K$ is the DM’s prediction.

In such a case, our main results and their proofs extend essentially without modification. To see this, note that when \mathbf{W} is diagonal, the DM simply minimizes a *weighted sum* of posterior variances about multiple states. Generalizing Lemma 2 in Appendix 9.1, we can show that any such objective function exhibits “eventual near-separability,” which is sufficient to derive Theorem 1 and 3. When \mathbf{W} is not diagonal, we can use the spectral theorem to write the DM’s objective function as a weighted sum of posterior variances about some linearly-transformed states. Our proofs still carry through.

6.3 Continuous Time

In a working paper, we analyze a continuous-time version of our problem. In that model, the DM has B units of attention in total at every point in time. He chooses attention

levels $\beta_1(t), \dots, \beta_K(t)$ subject to $\beta_i(t) \geq 0$ and $\sum_i \beta_i(t) \leq B$, and then observes K diffusion processes X_1, \dots, X_K , whose evolutions are affected by the attention rates in the following way:

$$dX_i(t) = \beta_i(t) \cdot \langle c_i, \theta \rangle dt + \sqrt{\beta_i(t)} dB_i,$$

where each B_i is an independent standard Brownian motion. This formulation can be seen as a limit of our discrete-time model in the current paper, where we take period length to zero and also “divide” the signals to hold constant the amount of information that can be gathered every instant.

We provide a brief review of the results here, and refer the reader to the working paper for more detail. In short, all results from this paper extend (and occasionally can be strengthened): We demonstrate that the optimal rule is eventually myopic in *all* informational environments (thus dropping the generic qualifier in Theorem 3). Additionally, we provide more permissive sufficient conditions on the informational environment under which an optimal strategy is myopic from period 1.

7 Related Literature

Besides the references mentioned in the introduction, our setting is related to a recent literature (Bubeck et al., 2009; Russo, 2016) regarding “best-arm identification” in a multi-armed bandit setting: a DM samples for a number of periods before selecting an arm and receiving its payoff. In Appendix 10.3, we characterize the optimal information acquisition strategy for the case of two states ($K = 2$), which exactly applies to the problem of identifying the better one of two correlated normal arms. However, due to our assumption of an one-dimensional payoff-relevant state, we are not able to handle more than two arms.³¹

We note that correlation is the key feature of our setting, and are not aware of many papers that study correlated bandits, either under the classical framework or under best-arm identification (see Rothschild (1974), Keener (1985) and Mersereau et al. (2009) for a few stylized cases).

³¹With two arms, the DM only cares about the difference in their expected payoffs. Choosing among more than two arms would involve *multi-dimensional payoff uncertainty* and a *decision problem that is not prediction*. As we discussed in 6.2, the lack of a complete Blackwell ordering limits the generalization of our argument. Incidentally, in related sequential search settings, Sanjurjo (2017), Ke and Villas-Boas (2017) and Chick and Frazier (2012) also highlight the challenge of characterizing the optimal strategy once there are at least three alternatives.

Our results on the comparison of sequential normal experiments (see the discussion in Section 4, and results in Appendix 9.2) generalize the main result in Greenshtein (1996). Greenshtein (1996) compares two *deterministic* (*i.e. history-independent*) sequences of signals, where each signal is θ_1 plus *independent* normal noise. His Theorem 3.1 implies that the former sequence is Blackwell-dominant if and only if its cumulative precision is higher at every time. Note that this statement does not refer to the prior beliefs, but if we impose a normal prior on θ_1 , then higher cumulative precision is equivalent to lower posterior variance. Thus, the result of Greenshtein (1996) coincides with ours when θ_1 is the only persistent state, and when all signals are independent conditional on θ_1 . Our setting features additional correlation across different signals through the persistent (payoff-irrelevant) states $\theta_2, \dots, \theta_K$. Consequently, the dynamic Blackwell comparison in our model depends on prior beliefs.³² This feature, together with the endogenous choice of signals (which may be history-dependent), complicates our problem relative to Greenshtein (1996).

Finally, our work is closely related to *optimal design*, a field initiated by the the early work of Robbins (1952) (see Chernoff (1972) for a survey). Specifically, the problem of one-shot allocation of t signals (our t -optimal criterion in Section 4) is equivalent to a Bayesian optimal design problem with respect to the “ c -optimality criterion”, which seeks to minimize the variance of an unknown parameter. Our analysis is however focused on dynamics, and we demonstrate here the optimality of “greedy design” for a broad class of (intertemporal) objectives.

8 Conclusion

A DM learns about a payoff-relevant state by sequentially sampling (batches of) signals from flexibly correlated Gaussian sources. Under conditions that we provide, myopic information acquisition is optimal and robust across all possible payoff functions. Generically, the optimal strategy eventually acquires signals myopically. These results are robust to extension to multi-player settings, to endogenous choice of the number of signals to acquire each period, and to multi-dimensional uncertainty for certain payoff functions.

We conclude with a re-interpretation of the main setting. Suppose there is a sequence of short-lived decision-makers indexed by time t , each of whom acquires information and then takes an action a_t to maximize some private objective $u_t(a_t, \theta_1)$. All information acquisition

³²This is already the case for static comparisons: as the prior beliefs vary, it is not always the same signal that leads to the lowest posterior variance about θ_1 .

is public.³³ A social planner has (an arbitrary) objective function $U(a_1, a_2, \dots; \theta_1)$; thus, his incentives are misaligned with those of the short-lived decision makers. Our main results demonstrate conditions under which this mis-alignment is of no consequence. If each DM acquires sufficiently many signals, or if the environment is separable, then each DM will acquire exactly the information that the social planner would have wanted. Generically, the social planner will not be able to improve (at late periods) upon the information that has been aggregated so far. We generalize this qualitative insight in our companion piece [Liang and Mu \(2017\)](#) and demonstrate also how it can fail.

³³This separates our model from classic social learning frameworks ([Banerjee, 1992](#); [Bikhchandani et al., 1992](#)), where decision-makers only observe coarse summary statistics of past information acquisitions.

9 Appendix

9.1 Preliminary Results

9.1.1 Posterior Variance Function

We begin by presenting basic results that are used throughout the appendix. The following lemma characterizes the posterior variance function f mentioned in the main text, which maps signal counts to the DM's posterior variance about the payoff-relevant state θ_1 .

Lemma 1. *Given prior covariance matrix V^0 and q_i observations of each signal i , the DM's posterior variance about θ_1 is given by³⁴*

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

where $\Sigma = C V^0 C' + D^{-1}$ and $D = \text{diag} \left(\frac{q_1}{\sigma_1^2}, \dots, \frac{q_K}{\sigma_K^2} \right)$. The function f is decreasing and convex in each q_i whenever these arguments take non-negative extended real values: $q_i \in \overline{\mathbb{R}_+} = \mathbb{R}_+ \cup \{+\infty\}$.

Proof. The expression (1) comes directly from the conditional variance formula for multivariate Gaussian distributions. To prove $\frac{\partial f}{\partial q_i} \leq 0$, consider the partial order \succeq on positive semi-definite matrices so that $A \succeq B$ if and only if $A - B$ is positive semi-definite. As q_i increases, the matrices D^{-1} and Σ decrease in this order. Thus Σ^{-1} increases in this order, which implies that $V^0 - V^0 C' \Sigma^{-1} C V^0$ decreases in this order. In particular, the diagonal entries of $V^0 - V^0 C' \Sigma^{-1} C V^0$ are uniformly smaller, so that f becomes smaller. Intuitively, more information always improves the decision-maker's estimates.

To prove f is convex, it suffices to prove f is *midpoint-convex* since the function is clearly continuous. Take $q_1, \dots, q_K, r_1, \dots, r_K \in \overline{\mathbb{R}_+}$ and let $s_i = \frac{q_i + r_i}{2}$.³⁵ Define the corresponding diagonal matrices to be D_q, D_r, D_s . We need to show $f(q_1, \dots, q_K) + f(r_1, \dots, r_K) \geq 2f(s_1, \dots, s_K)$. For this, we first use the Woodbury inversion formula to write

$$\Sigma^{-1} = (C V^0 C' + D^{-1})^{-1} = J - J(J + D)^{-1} J,$$

with $J = (C V^0 C')^{-1}$. Plugging this back into (1), we see that it suffices to show the following matrix order:

$$\frac{(J + D_q)^{-1} + (J + D_r)^{-1}}{2} \succeq (J + D_s)^{-1}.$$

³⁴When M is a matrix, we let M_{ij} denote its (i, j) -th entry.

³⁵We allow the function f to take $+\infty$ as arguments. This relaxation does not affect the properties of f , and it is convenient for our future analysis.

Inverting both sides, we need to show $2((J + D_q)^{-1} + (J + D_r)^{-1})^{-1} \preceq J + D_s$. From definition, $D_q + D_r = \text{diag}(\frac{q_1+r_1}{\sigma_1^2}, \dots, \frac{q_K+r_K}{\sigma_K^2}) = 2D_s$. Thus the above follows from the AM-HM inequality for positive definite matrices, see for instance [Ando \(1983\)](#). \square

9.1.2 The Matrix Q_i

Let us define for each $1 \leq i \leq K$,

$$Q_i = C^{-1} \Delta_{ii} C'^{-1} \quad (2)$$

where Δ_{ii} is the matrix with ‘1’ in the (i, i) -th entry, and zeros elsewhere. We note that $[Q_i]_{11} = ([C^{-1}]_{1i})^2$, which is strictly positive under Assumption 1. These matrices Q_i will be repeatedly used in our proofs.

9.1.3 Order Difference Lemma

Here we establish the asymptotic order for the second derivatives of f .

Lemma 2. *As $q_1, \dots, q_K \rightarrow \infty$, $\frac{\partial^2 f}{\partial q_i^2}$ is positive with order $\frac{1}{q_i^3}$, whereas $\frac{\partial^2 f}{\partial q_i \partial q_j}$ has order at most $\frac{1}{q_i^2 q_j^2}$ for any $j \neq i$. Formally, there is a positive constant L depending on the informational environment, such that $\frac{\partial^2 f}{\partial q_i^2} \geq \frac{1}{L q_i^3}$ and $|\frac{\partial^2 f}{\partial q_i \partial q_j}| \leq \frac{L}{q_i^2 q_j^2}$.*

To interpret, the second derivative $\partial^2 f / \partial q_i^2$ is the effect of observing signal i on the marginal value of the next observation of signal i . Our lemma says that this second derivative is always eventually positive, so that each observation of signal i makes the next observation of signal i less valuable. The cross-partial $\partial^2 f / \partial q_i \partial q_j$ is the effect of observing signal i on the marginal value of the next observation of a different signal j , and its sign is ambiguous.

The key content of the lemma is that regardless of the sign of the cross partial, it is always of lower order compared to the second derivative. In words, the effect of observing a signal on the marginal value of other signals (as quantified by the cross-partial) is eventually second-order to its effect on the marginal value of further realizations of the same signal (as quantified by the second derivative). This is true for any signal path in which the signal counts q_1, \dots, q_K go to infinity proportionally, which we will justify later.

Proof. Recall from Lemma 1 that $f(q_1, \dots, q_K) = [V^0 - V^0 C' \Sigma^{-1} C V^0]_{11}$ and therefore

$$\frac{\partial^2 f}{\partial q_i \partial q_j} = [\partial_{ij}(V^0 - V^0 C' \Sigma^{-1} C V^0)]_{11} \quad \frac{\partial^2 f}{\partial q_i^2} = [\partial_{ii}(V^0 - V^0 C' \Sigma^{-1} C V^0)]_{11} \quad (3)$$

Using properties of matrix derivatives,

$$\partial_{ii}(\Sigma^{-1}) = \Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1} - \Sigma^{-1}(\partial_{ii} \Sigma)\Sigma^{-1} + \Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1}. \quad (4)$$

The relevant derivatives of the covariance matrix Σ are

$$\partial_i \Sigma = -\frac{\sigma_i^2}{q_i^2} \Delta_{ii} \quad \partial_{ii} \Sigma = \frac{2\sigma_i^2}{q_i^3} \Delta_{ii}$$

Plugging these into (4), we obtain $\partial_{ii}(\Sigma^{-1}) = -\frac{2\sigma_i^2}{q_i^3}(\Sigma^{-1} \Delta_{ii} \Sigma^{-1}) + O\left(\frac{1}{q_i^4}\right)$. Thus by (3),

$$\frac{\partial^2 f}{\partial q_i^2} = \left[-V^0 C' \cdot \frac{\partial^2(\Sigma^{-1})}{\partial q_i^2} \cdot C V^0 \right]_{11} = \frac{2\sigma_i^2}{q_i^3} \cdot [V^0 C' \Sigma^{-1} \Delta_{ii} \Sigma^{-1} C V^0]_{11} + O\left(\frac{1}{q_i^4}\right). \quad (5)$$

As $q_1, \dots, q_k \rightarrow \infty$, $\Sigma \rightarrow C V^0 C'$ which is symmetric and non-singular. Thus the matrix $V^0 C' \Sigma^{-1} \Delta_{ii} \Sigma^{-1} C V^0$ converges to the matrix Q_i defined earlier in (2). From (5) and $[Q_i]_{11} > 0$, we conclude that $\frac{\partial^2 f}{\partial q_i^2}$ is positive with order $\frac{1}{q_i^3}$. Similarly, for $i \neq j$, we have

$$\partial_{ij}(\Sigma^{-1}) = \Sigma^{-1}(\partial_j \Sigma)\Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1} - \Sigma^{-1}(\partial_{ij} \Sigma)\Sigma^{-1} + \Sigma^{-1}(\partial_i \Sigma)\Sigma^{-1}(\partial_j \Sigma)\Sigma^{-1}.$$

The relevant derivatives of the covariance matrix Σ are

$$\partial_i \Sigma = -\frac{\sigma_i^2}{q_i^2} \Delta_{ii} \quad \partial_j \Sigma = -\frac{\sigma_j^2}{q_j^2} \Delta_{jj} \quad \partial_{ij} \Sigma = \mathbf{0}$$

From this it follows that $\partial_{ij}(\Sigma^{-1}) = O\left(\frac{1}{q_i^2 q_j^2}\right)$. The same holds for $\frac{\partial^2 f}{\partial q_i \partial q_j}$ because of (3), completing the proof of the lemma. \square

9.2 Dynamic Blackwell Comparison

9.2.1 The Lemma

This subsection establishes a dynamic version of Blackwell dominance for sequences of normal signals. As an overview, we first generalize Greenshtein (1996) and show that a *deterministic* (i.e. history-independent) signal sequence yields higher expected payoff than another in every intertemporal decision problem if (and only if) the former sequence induces lower posterior variances about θ_1 at every period. This will be a corollary of the lemma below, which also covers strategies that may condition on signal realizations.

We introduce some notation: Since θ_1 is the only payoff-relevant state, the DM in our model only needs to remember the expected value of θ_1 and the covariance matrix over all of the states (that is, expected values of the other states do not matter). Thus, we can

summarize any history of beliefs by $h^T = (\mu_1^0, V^0; \dots, \mu_1^T, V^T)$, with μ_1^t representing the posterior expected value of θ_1 after t periods and V^t the posterior covariance matrix. Since the posterior covariance matrix is a function of signal counts, we can also keep track of the evolution of posterior covariance matrices by a sequence of division vectors. That is, we will write the history as $h^T = (\mu_1^0, d(0); \dots; \mu_1^T, d(T))$, where each $d(t) = (d_1(t), \dots, d_K(t))$ counts the number of each signal acquired by time t . We can then view any information acquisition strategy S as a mapping from such sequences of expected values and division vectors to signal choices.

Consider a mapping \tilde{G} from possible sequences of divisions to these sequences themselves: For each $(d(0), \dots, d(T))$, \tilde{G} maps to another sequence $(\tilde{d}(0), \dots, \tilde{d}(T))$, subject to the following ‘‘consistency’’ requirements. First, $\sum_i \tilde{d}_i(t) = t$, meaning that each $\tilde{d}(t)$ must be a possible division at time t . Second, $\tilde{d}_i(t) \geq \tilde{d}_i(t-1)$, meaning that the sequence \tilde{d} can be attained via a sequential sampling rule. Lastly, we require

$$(\tilde{d}(0), \dots, \tilde{d}(T-1)) = \tilde{G}(d(0), \dots, d(T-1))$$

so that nesting sequences are mapped to nesting sequences.

The following lemma says that if $d(\cdot)$ represents the division vectors under an information acquisition strategy S , and if \tilde{G} is a consistent mapping that uniformly reduces the posterior variance, then we can find another information acquisition strategy \tilde{S} whose division vectors are given by $\tilde{d}(\cdot)$. Moreover, our construction ensures that \tilde{S} leads to more dispersed posterior beliefs than S at every period, so that in any decision problem, acquiring signals according to \tilde{S} is weakly better than S (when actions are taken optimally).

Lemma 3. *Fix any information acquisition strategy S and any consistent mapping \tilde{G} defined above. Suppose that for every sequence of divisions $(d(0), \dots, d(T))$ realized under S , it holds that*

$$f(\tilde{d}(T)) \leq f(d(T))$$

Then there exists deviation strategy \tilde{S} such that, at every period T , any history $h^T = (\mu_1^0, d(0); \dots; \mu_1^T, d(T))$ under S can be ‘‘associated with’’ a distribution of histories $\tilde{h}^T = (\nu_1^0, \tilde{d}(0); \dots; \nu_1^T, \tilde{d}(T))$ with the following properties:

1. *the probability of h^T occurring under S is the same as the probability of its associated \tilde{h}^T (integrated with respect to the probability of ‘‘association’’) occurring under \tilde{S} ;*
2. *the total probability that any \tilde{h}^T is associated to (integrated with respect to different possible h^T) is 1;*

3. under the association, the distribution of ν_1^t is normal with mean μ_1^t and variance $f(d(t)) - f(\tilde{d}(t))$ for each t .

Consequently, for any decision strategy A , there exists another decision strategy \tilde{A} such that the expected payoff under (\tilde{S}, \tilde{A}) is no less than the expected payoff under (S, A) .

To interpret, the first two properties require that the association relation is a *Markov kernel* between histories under S and histories under \tilde{S} ; this enables us to compare payoffs under \tilde{S} to those under S . The third property guarantees that the alternative strategy \tilde{S} is more informative than S .

We note the following corollary, which is obtained from the previous lemma by considering a constant mapping \tilde{G} .

Corollary 3. *Define the t -optimal division vectors as in Section 4.1. Suppose each coordinate of $n(Bt)$ increases in t . Then it is optimal for the DM to achieve $n(Bt)$ at every period.*

9.2.2 Proof of Lemma 3

We construct \tilde{S} iteratively as follows. In the first period, consider the signal choice under S (given the null history). This signal leads to the division $d(1)$. Let \tilde{S} observe the unique signal that would achieve the division $\tilde{d}(1)$.

After the first observation, the DM's *distribution of posterior beliefs* about θ_1 under S is $\theta_1 \sim \mathcal{N}(\mu_1^1, f(d(1)))$ with μ_1^1 a normal random variable with mean μ_1^0 and variance $f(\mathbf{0}) - f(d(1))$. By comparison, the distribution of posterior beliefs under \tilde{S} is $\theta_1 \sim \mathcal{N}(\nu_1^1, f(\tilde{d}(1)))$ with ν_1^1 drawn from $\mathcal{N}(\mu_1^0, f(\mathbf{0}) - f(\tilde{d}(1)))$. Since $f(\tilde{d}(1)) \leq f(d(1))$, the latter distribution of beliefs (under \tilde{S}) is more informative a la Blackwell. Thus, we can associate each belief $\theta_1 \sim \mathcal{N}(\mu_1^1, f(d(1)))$ under S with a more informative distribution of beliefs $\mathcal{N}(\nu_1^1, f(\tilde{d}(1)))$ under \tilde{S} . To be more specific, for fixed μ_1^1 , the associated ν_1^1 is distributed normally with mean μ_1^1 and variance $f(d(1)) - f(\tilde{d}(1))$. Thus by construction, all three properties are satisfied at period 1. To facilitate the discussion below, we say this distribution of beliefs under \tilde{S} “imitates” the belief $(\mu_1^1, f(d(1)))$ under S .

In the second period, the deviation strategy \tilde{S} takes the current belief $(\nu_1^1, f(\tilde{d}(1)))$ and randomly selects some μ_1^1 (with conditional probabilities under the Markov kernel) to “imitate.” That is, given any selection of μ_1^1 , find the signal that S would observe in period 2 given belief $(\mu_1^1, f(d(1)))$. This signal choice under S leads to the division sequence $(d(0), d(1), d(2))$, which is mapped to $(\tilde{d}(0), \tilde{d}(1), \tilde{d}(2))$. Naturally, we let \tilde{S} observe the signal that would lead to the division $\tilde{d}(2)$. Such a signal is well-defined due to our consistency requirements on \tilde{G} .

To proceed with the analysis, let us fix μ_1^1 and study the distribution of posterior beliefs about θ_1 after two observations. Under S , the distribution of posterior beliefs is $\theta_1 \sim \mathcal{N}(\mu_1^2, f(d(2)))$ with μ_1^2 normally distributed with mean μ_1^1 and variance $f(d(1)) - f(d(2))$. While under \tilde{S} , the distribution of posterior beliefs is $\theta_1 \sim (\nu_1^2, f(\tilde{d}(2)))$ with ν_1^2 drawn from $\mathcal{N}(\mu_1^1, f(d(1)) - f(\tilde{d}(2)))$.³⁶

Since $f(\tilde{d}(2)) \leq f(d(2))$, the distribution of beliefs under \tilde{S} Blackwell-dominates the distribution under S , for each μ_1^1 . We can thus associate each history $(\mu_1^1, d(1); \mu_1^2, d(2))$ under S with a distribution of histories $(\nu_1^1, \tilde{d}(1); \nu_1^2, \tilde{d}(2))$ under \tilde{S} , such that the corresponding beliefs under \tilde{S} are more informative at both periods. Repeating this procedure completes the construction of \tilde{S} , which satisfies all three properties stated in the lemma.

Finally, suppose A is any decision strategy that maps histories to actions. We need to find \tilde{A} such that the pair (\tilde{S}, \tilde{A}) does no worse than (S, A) . This is straightforward given what we have done: at any history \tilde{h}^T under \tilde{S} , let \tilde{h}^T randomly select h^T to imitate, and define $\tilde{A}(\tilde{h}^T) = A(h^T)$. Then we see that a DM who follows the decision strategy A obtains the same payoff along any belief history h as another DM who uses the decision strategy \tilde{A} and faces the distribution of belief histories \tilde{h} . Integrating over h , we have shown that (\tilde{S}, \tilde{A}) achieves the same payoff as (S, A) . The lemma is proved.

9.3 Proof of Theorem 1 (Large Block of Signals)

By Corollary 3, it suffices to show that for sufficiently large B , each coordinate $n(Bt)$ is increasing in t . To do this, we first argue that the signal counts grow to infinity (roughly) proportionally. In more detail, define

$$\lambda_i = \frac{|[C^{-1}]_{1i}| \cdot \sigma_i}{\sum_{j=1}^K |[C^{-1}]_{1j}| \cdot \sigma_j}. \quad (6)$$

Then we will show that for each signal i , $n_i(t) - \lambda_i \cdot t$ remains bounded even as $t \rightarrow \infty$.

Indeed, we must at least have $n_i(t) \rightarrow \infty$; otherwise the posterior variance $f(n(t))$ would be bounded away from zero, which would contradict the optimality of $n(t)$ since $f(t/K, \dots, t/K) \rightarrow 0$. Additionally, we compute from (1) that

$$\partial_i f(n(t)) = -\frac{\sigma_i^2}{n_i^2} \cdot [V^0 C' \Sigma^{-1} \Delta_{ii} \Sigma^{-1} C V^0]_{11}. \quad (7)$$

³⁶Here we use the following technical result: suppose the DM is endowed with a *distribution of prior beliefs* $\theta \sim \mathcal{N}(\mu, V)$, with μ_1 normally distributed with mean y and variance σ^2 , then upon observing signal i and performing Bayesian updating, his distribution of posterior beliefs is $\theta \sim \mathcal{N}(\hat{\mu}, \hat{V})$, with $\hat{\mu}_1$ normally distributed with mean y and variance $\sigma^2 + [V - \hat{V}]_{11}$. This is proved by noting that the DM's distribution of beliefs about θ_1 must integrate to the same ex-ante distribution of θ_1 .

As each $n_i \rightarrow \infty$, the matrix $\Sigma = CV^0C' + D^{-1}$ (see Lemma 1) converges to CV^0C' . So $V^0C'\Sigma^{-1}\Delta_{ii}\Sigma^{-1}CV^0$ converges to the matrix Q_i defined in (2). It follows from (7) that $\partial_i f \sim \frac{-\sigma_i^2}{n_i^2} \cdot [Q_i]_{11}$ (ratio converges to 1). Since a t -optimal division must satisfy $\partial_i f \sim \partial_j f$ (because we are doing discrete optimization, $\partial_i f$ and $\partial_j f$ need only be approximately equal), we deduce that n_i and n_j must grow proportionally. Using $[Q_i]_{11} = ([C^{-1}]_{1i})^2$, we have $n_i(t) \sim \lambda_i t$.

Next, note that because $n_i(t) \sim \lambda_i t$, $\Sigma = CV^0C' + D^{-1} = CV^0C' + O(\frac{1}{t})$.³⁷ Thus in fact $V^0C'\Sigma^{-1}\Delta_{ii}\Sigma^{-1}CV^0$ converges to Q_i at the rate of $\frac{1}{t}$. From (7), we obtain $\partial_i f = \frac{-\sigma_i^2 \cdot [Q_i]_{11} + O(\frac{1}{t})}{n_i^2}$. t -optimality gives us the first-order condition $\partial_i f = \partial_j f + O(\frac{1}{t^3})$.³⁸ So

$$\frac{\lambda_i^2 + O(\frac{1}{t})}{n_i^2} = \frac{\lambda_j^2 + O(\frac{1}{t})}{n_j^2}.$$

This is equivalent to $\lambda_i^2 n_j^2 - \lambda_j^2 n_i^2 = O(t)$, which yields $\lambda_i n_j - \lambda_j n_i = O(1)$ after factorization. Hence $n_i(t) = \lambda_i \cdot t + O(1)$ as we claimed.

Having completed this asymptotic characterization of the t -optimal division vectors, we will now show that $n(t+K-1) \geq n(t)$ (in each coordinate) whenever t is sufficiently large. Theorem 1 will follow once this is proved.³⁹

Suppose for the sake of contradiction that $n_1(t+K-1) \leq n_1(t) - 1$. Note we have $\sum_{i=1}^K (n_i(t+K-1) - n_i(t)) = K - 1$. So $\sum_{i=2}^K (n_i(t+K-1) - n_i(t)) \geq K$, and we can without loss of generality assume $n_2(t+K-1) \geq n_2(t) + 2$. To summarize, when transitioning from t -optimality to $t+K-1$ -optimality, signal 1 is acquired at least once less and signal 2 at least twice more. Below we will obtain a contradiction by arguing that at period $t+K-1$, the posterior variance could be further reduced by observing signal 1 once more and signal 2 once less.

Indeed, let us write $n_i = n_i(t)$ and $\tilde{n}_i = n_i(t+K-1)$. Then t -optimality of $n(t)$ gives us

$$f(n_1 - 1, n_2 + 1, \dots, n_K) \geq f(n_1, n_2, \dots, n_K).$$

With a slight abuse of notation, we let $\partial_i f$ to denote the *discrete* partial derivative of f : $\partial_i f(q) = f(q_i + 1, q_{-i}) - f(q)$. Then the above display is equivalent to

$$\partial_2 f(n_1 - 1, n_2, \dots, n_K) \geq \partial_1 f(n_1 - 1, n_2, \dots, n_K). \quad (8)$$

³⁷“Big O” notation has the usual meaning.

³⁸That is, error terms due to discreteness are no larger than $\frac{1}{t^3}$. We omit the details.

³⁹To be fully rigorous, this only proves Theorem 1 when B is sufficiently large and is a multiple of $K-1$. However, we can similarly show $n(t+K) \geq n(t)$ for sufficiently large t . The two inequalities $n(t+K-1) \geq n(t)$ and $n(t+K) \geq n(t)$ together are sufficient to deduce Theorem 1 for all large B .

We claim this implies the following:

$$\partial_2 f(\tilde{n}_1, \tilde{n}_2 - 1, \dots, \tilde{n}_K) > \partial_1 f(\tilde{n}_1, \tilde{n}_2 - 1, \dots, \tilde{n}_K). \quad (9)$$

This would lead to

$$f(\tilde{n}_1, \tilde{n}_2, \dots, \tilde{n}_K) > f(\tilde{n}_1 + 1, \tilde{n}_2 - 1, \dots, \tilde{n}_K),$$

which would be our desired contradiction.

It remains to show (8) \implies (9). By assumption, we have $\tilde{n}_1 \leq n_1 - 1$, $\tilde{n}_2 \geq n_2 + 2$ and the difference between any \tilde{n}_j and n_j is bounded uniformly over t . Thus the LHS of (9) exceeds the LHS of (8) by (at least) a second derivative ∂_{22} minus a finite number of cross partial derivatives ∂_{2j} . By Lemma 2, this difference on the LHS is positive with order $\frac{1}{t^3}$. The difference between the RHS of (9) and the RHS of (8) can be positive or negative, but either way it has order $O(\frac{1}{t^4})$. This shows (9) is a consequence of (8), and the theorem follows.

9.4 Proof of Theorem 2 (Separable Environments)

Suppose the informational environment is separable. We will show $n(t)$ increases in t , which implies the theorem via Corollary 3.

Note that in a separable environment, the definition of t -optimality reduces to the following.

$$n(t) = (n_1(t), \dots, n_K(t)) \in \underset{(q_1, \dots, q_K): q_i \in \mathbb{Z}_+, \sum_{i=1}^K q_i = t}{\operatorname{argmin}} \sum_{i=1}^K g_i(q_i)$$

where g_1, \dots, g_K are convex functions.

In this setting, the myopic information acquisition strategy sequentially chooses the signal i that minimizes the difference $g_i(q_i + 1) - g_i(q_i)$, given the current division vector q . But since the g -functions are convex, the outcome under the myopic strategy coincides with t -optimality at every period t (This is fairly well known, and it can be proved quickly by induction.) Hence $n(t)$ can and will be sequentially attained.

9.5 Preparation for the Proof of Theorem 3

9.5.1 Switch Deviations

We now introduce preliminary results that will be used to show that the optimal rule eventually proceeds myopically in generic environments. Relative to the proofs of Theorems 1

and 2, the new difficulty that arises is that in general, the optimal information acquisition strategy conditions on signal realizations. As a result, the induced division vectors $d(\cdot)$ are stochastic, and we will need the full power of our dynamic Blackwell lemma.

In what follows, we will apply Lemma 3 using a particular class of mappings \tilde{G} .

Definition 3. Fix a particular sequence of divisions $(d^*(0), d^*(1), \dots, d^*(t_0))$. Let i be the signal observed in period t_0 and j be any other signal. An (i, j) -switch mapping \tilde{G} specifies the following:

1. Suppose $T < t_0$ or $d(t) \neq d^*(t)$ for some $t \leq t_0$, then let $\tilde{G}(d(0), \dots, d(T))$ be itself.
2. Otherwise $T \geq t_0$ and $d(t) = d^*(t), \forall t \leq t_0$. If $d_j(T) = d_j(t_0)$, then let $\tilde{d}(T) = (d_i(T) - 1, d_j(T) + 1, d_{-ij}(T))$. If $d_j(T) > d_j(t_0)$, then let $\tilde{d}(T) = d(T)$.

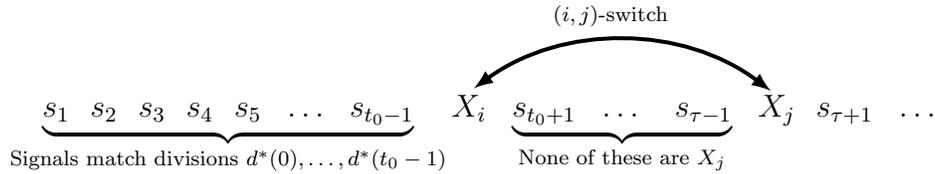


Figure 2: Pictorial representation of an (i, j) -switch based on a sequence of divisions $d^*(0), \dots, d^*(t_0)$.

Let us interpret this definition by relating to the resulting deviation strategy \tilde{S} constructed in Lemma 3. The first case above says that \tilde{S} only deviates when the history of divisions is $d^*(0), \dots, d^*(t_0 - 1)$ and S is about to observe signal i in period t_0 . The second case says that \tilde{S} dictates observing signal j instead at that history; subsequently, \tilde{S} observes the same signal as S (at the imitated belief) until the first period at which S is about to observe signal j . If that period exists, the deviation strategy \tilde{S} switches back to observing signal i and coincides with S afterwards.

The benefit of these “switch deviations” is that their posterior variances can be easily compared to the original strategy. Specifically, $\tilde{d}(t) = d(t)$ except at those histories that begin with $d^*(0), d^*(1), \dots, d^*(t_0 - 1)$ (and before signal j is observed again under S). At such histories, the posterior variance is strictly lower under \tilde{S} if and only if

$$f(d_i(t) - 1, d_j(t) + 1, d_{-ij}(t)) < f(d(t)).$$

Using (absolute values of) the discrete partial derivatives, we can rewrite this condition as

$$|\partial_i f(d_i(t) - 1, d_j(t), d_{-ij}(t))| < |\partial_j f(d_i(t) - 1, d_j(t), d_{-ij}(t))|. \quad (*)$$

We can thus obtain the following corollary:

Corollary 4. *Suppose we can find a history of divisions $d(0), \dots, d(t_0)$ realized under S such that $d_i(t_0) = d_i(t_0 - 1) + 1$ and moreover $(*)$ holds for all divisions $d(t)$ with $d_j(t) = d_j(t_0)$ and $d_k(t) \geq d_k(t_0), \forall k$. Then the switch deviation \tilde{S} constructed above improves upon S .*

Note that the condition $d_j(t) = d_j(t_0)$ captures the fact that $\tilde{d}(t)$ differs from $d(t)$ only until signal j is chosen again by S . Meanwhile, $d_k(t) \geq d_k(t_0), \forall k$ holds because we only compare posterior variances after t_0 periods.

9.5.2 Asymptotic Characterization of Optimal Strategy

Below we will use the contrapositive of Corollary 4 to argue that if S is the optimal information acquisition strategy, then we cannot find any history of realized divisions such that $(*)$ always holds. Technically speaking, we might worry that although \tilde{S} strictly improves upon S in terms of posterior variances, it might achieve the same expected payoff as S (for instance, when the DM faces a constant payoff function). Nonetheless, by Zorn's lemma we can choose S to be an optimal strategy that is additionally “un-dominated” in terms of posterior variances. With that choice, the deviation \tilde{S} cannot exist, and our arguments remain valid. (Note that our theorems only state that the DM *has an optimal strategy* . . .)

To illustrate, we now derive the asymptotic signal proportions for the optimal information acquisition strategy S .

Lemma 4. *Suppose S is the optimal information acquisition strategy, and $d(\cdot)$ is its induced divisions. Let λ_k be defined as in (6). In generic informational environments, the difference $d_k(T) - \lambda_k \cdot T$ remains bounded as $T \rightarrow \infty$, for any realized division $d(T)$ and each signal k .*

Proof. For this proof, we only need the informational environment to be such that each signal has *strictly positive marginal value*. That is, for any signal k and any possible division q , we require

$$f(q_k + 1, q_{-k}) < f(q).$$

This is “generically” satisfied because any equality $f(q_k + 1, q_{-k}) = f(q)$ would impose a non-trivial polynomial equation over the signal linear coefficients, and the number of such constraints is at most countable.

Under this genericity assumption, let us first show $d_k(T) \rightarrow \infty$ holds for each signal k , and the speed of divergence depends only on the informational environment. For contradiction, suppose this is not true. Then we can find a sequence of histories $\{h^{T_m}\}$ such that $T_m \rightarrow \infty$ but $d_1(T_m)$ remains bounded (these histories need not nest one another). By passing to a subsequence, we may assume $q_k = \lim_{m \rightarrow \infty} d_k(T_m)$ exists for every signal k , where this limit

may be infinity. Define I to be the non-empty subset of signals (not including signal 1) with $q_k = \infty$. Furthermore, we assume that the signal observed in the last period of each of these histories h^{T_m} is the same signal i . We also assume $i \in I$; otherwise just truncate the histories by finitely many periods.

Take any signal $j \notin I$ (for instance, $j = 1$ works). Choose T_m sufficiently large and consider the (i, j) -switch deviation \tilde{S} that deviates from h^{T_m} by observing signal j instead of i in period T_m . We will verify (*) for all possible divisions $d(t)$ with $d_j(t) = d_j(T_m)$ and $d_i(t) \geq d_i(T_m)$, which will contradict the optimality of S via Corollary 4. Indeed, note that as $T_m \rightarrow \infty$, $d_i(T_m) \rightarrow \infty$ because $i \in I$. Since $d_i(t) \geq d_i(T_m)$, the LHS of (*) approaches zero as T_m increases. By comparison, the RHS of (*) is bounded away from zero because $d_j(t) = d_j(T_m)$ is bounded, and we assume each signal has strictly positive marginal value. Hence (*) holds and we have shown that $d(T) \rightarrow \infty$ in each coordinate.

Next, from (7), we have the following approximations for the partial derivatives:

$$|\partial_i f(d_i(t) - 1, d_j(t), d_{-ij}(t))| \sim \frac{\sigma_i^2 \cdot [Q_i]_{11}}{d_i(t)^2} \quad |\partial_j f(d_i(t) - 1, d_j(t), d_{-ij}(t))| \sim \frac{\sigma_j^2 \cdot [Q_j]_{11}}{d_j(t)^2}$$

If $\limsup_{t_0 \rightarrow \infty} \frac{d_i(t_0)}{d_j(t_0)} > \frac{\lambda_i}{\lambda_j}$ (recall that λ_i is proportional to $\sigma_i \cdot \sqrt{[Q_i]_{11}}$), then the above estimates would imply (*) whenever $d_i(t) \geq d_i(t_0)$ (because $t \geq t_0$) and $d_j(t) = d_j(t_0)$. That would contradict the optimality of S . Hence, $\limsup_{t_0 \rightarrow \infty} \frac{d_i(t_0)}{d_j(t_0)} \leq \frac{\lambda_i}{\lambda_j}$ for every pair of signals i and j . It follows that $d_k(t_0) \sim \lambda_k \cdot t_0, \forall k$.

Once these asymptotic proportions are proved, we know that the matrix $\Sigma = CV^0C' + D^{-1}$ converges to CV^0C' at the rate of $\frac{1}{t}$. By (7), we can deduce more precise approximations:

$$|\partial_i f(d_i(t) - 1, \dots)| = \frac{\sigma_i^2 \cdot [Q_i]_{11} + O(\frac{1}{t})}{d_i(t)^2} \quad |\partial_j f(d_i(t) - 1, \dots)| = \frac{\sigma_j^2 \cdot [Q_j]_{11} + O(\frac{1}{t})}{d_j(t)^2}$$

If $\frac{d_i(t_0)}{d_j(t_0)} > \frac{\lambda_i}{\lambda_j} + O(\frac{1}{t_0})$, then these refined estimates would again imply (*) whenever $d_i(t) \geq d_i(t_0)$ and $d_j(t) = d_j(t_0)$. To avoid the resulting contradiction, we must have $\frac{d_i(t_0)}{d_j(t_0)} \leq \frac{\lambda_i}{\lambda_j} + O(\frac{1}{t_0})$ for every pair of signals. This enables us to conclude $d_k(t_0) = \lambda_k \cdot t_0 + O(1)$ as desired. \square

9.6 Proof of Theorem 3 (Generic Eventual Myopia)

9.6.1 Outline of the Proof

To guide the reader through this appendix, we begin by outlining the proof of the theorem, which is broken down into several steps. Throughout, we focus on the case of $B = 1$ (one

signal each period), but our proof easily extends to arbitrary B . We will first show a simpler (and weaker) result that in generic environments, the number of periods in which the optimal strategy differs from the t -optimal division has natural density 1. Our proof of this result is based on the observation that if equivalence does not hold at some time t , there must be *two different divisions* over signals for which the resulting posterior variances about θ_1 are within $O(\frac{1}{t^4})$ from each other. This leads to a Diophantine approximation inequality, which we can show only occurs at a vanishing fraction of periods t .

To improve the result and demonstrate equivalence at *all late periods*, we show that the number of “exceptional periods” t is generically finite if there are *three different divisions* over signals whose posterior variances are within $O(\frac{1}{t^4})$ from each other. This allows us to conclude that in generic environments, the t -optimal divisions eventually monotonically increase in t .

In such environments, t -optimality can be achieved at every late period. Thus, whenever t -optimality obtains in *some* late period, it will be sustained in all future periods. Since we have already established that the optimal strategy achieves t -optimality infinitely often, we conclude equivalence at all large t .

We highlight that in this appendix, we use a slightly different notion of “generic” where we fix the signal coefficient matrix C and instead (randomly) vary the signal variances $\{\sigma_i^2\}$. This concept implies (and is stronger than) the previous genericity concept defined on C .

9.6.2 Equivalence at Almost All Times

We begin by proving a weaker result, that the optimal strategy induces the t -optimal division $n(t)$ at almost all periods t .

Proposition 2. *Suppose the informational environment $(V^0, C, \{\sigma_i^2\})$ has the property that for any $i \neq j$, the ratio $\frac{\lambda_i}{\lambda_j}$ is an irrational number. Then, at a set of times with natural density 1,⁴⁰ $d(t) = n(t)$ (which is unique) holds for every decision problem. In particular, the optimal strategy induces a deterministic division vector at such times.*

Proof of Proposition 2. Suppose that $d_1(T) \geq n_1(T) + 1$ and $d_2(T) \leq n_2(T) - 1$. Consider the last period $t_0 \leq T$ in which the optimal strategy observed signal 1. Then

$$d_1(t_0) = d_1(T) \geq n_1(T) + 1; \quad d_2(t_0) \leq d_2(T) \leq n_2(T) - 1.$$

⁴⁰Formally, for any set of positive integers A , let $A(N)$ count the number of integers in A no greater than N . Then we define the natural density of A to be $\lim_{N \rightarrow \infty} \frac{A(N)}{N}$, when this limit exists.

Using the contrapositive of Corollary 4 with the (1, 2)-switch, we know that (*) cannot always hold. Thus there exists a division $d(t)$ such that the inequality (*) is reversed. That is, we can find a division $d(t)$ with $d_1(t) \geq d_1(t_0)$ and $d_2(t) = d_2(t_0)$ such that (getting rid of the absolute values)

$$\partial_1 f(d_1(t) - 1, d_2(t), d_{-12}(t)) \leq \partial_2 f(d_1(t) - 1, d_2(t), d_{-12}(t)). \quad (10)$$

On the other hand, t -optimality of $n(t)$ gives us

$$\partial_1 f(n_1(T), n_2(T) - 1, n_{-12}(T)) \geq \partial_2 f(n_1(T), n_2(T) - 1, n_{-12}(T)). \quad (11)$$

Note that $d_2(t) = d_2(t_0)$ implies $t - t_0$ is bounded (due to Lemma 4). On the other hand, we have $d_1(t) = d_1(T_0)$ by construction (t_0 is the last period signal 1 was observed). Hence $t_0 - T$ is also bounded. Combining both, we deduce $t - T$ must be bounded. Applying Lemma 4 again, we know that any difference $d_i(t) - n_i(T)$ is bounded.

Now because $d_1(t) - 1 \geq d_1(t_0) - 1 \geq n_1(T)$, the LHS of (10) has size at least the LHS of (11) minus a finite number of cross partial derivatives ∂_{1j} . Similarly, the RHS of (10) is at most bigger than the RHS of (11) by a number of cross partials. Together with the order difference lemma, these imply that the only way (10) and (11) can both hold is if the two sides of (11) differ by at most $O\left(\frac{1}{T^4}\right)$.

To summarize: a necessary condition for $d_1(T) \geq n_1(T) + 1$ and $d_2(T) \leq n_2(T) + 1$ to occur is that

$$|f(n_1(T) + 1, n_2(T) - 1, \dots, n_K(T)) - f(n(T))| = O\left(\frac{1}{T^4}\right). \quad (12)$$

Hence, to prove the Proposition we only need to show that (12) holds at a set of times with natural density 0. The following lemma proves exactly this property. \square

Lemma 5. *Suppose $\frac{\lambda_1}{\lambda_2}$ is an irrational number. For positive constants c_0, c_1 , define $\mathcal{A}(c_0, c_1)$ to be the following set of positive integers:*

$$\begin{aligned} \{t : \exists q_1, q_2, \dots, q_K \in \mathbb{Z}^+, \text{ s.t. } |q_i - \lambda_i \cdot t| \leq c_0, \forall i \\ \wedge |f(q_1, q_2 + 1, \dots, q_K) - f(q_1 + 1, q_2, \dots, q_K)| \leq c_1/t^4\}. \end{aligned}$$

Then $\mathcal{A}(c_0, c_1)$ has natural density zero.

Proof of Lemma 5. The proof relies on the following technical result, which gives a precise approximation of the *discrete partial derivatives* of f :

Lemma 6. *Fix the informational environment. There exists a constant a_j such that*

$$f(q_j, q_{-j}) - f(q_j + 1, q_{-j}) = \frac{\sigma_j^2 \cdot [Q_j]_{11}}{(q_j - a_j)^2} + O\left(\frac{1}{t^4}; c_0\right) \quad (13)$$

holds for all q_1, \dots, q_K with $|q_i - \lambda_i t| \leq c_0, \forall i$. The notation $O\left(\frac{1}{t^4}; c_0\right)$ means an upper bound of $\frac{L}{t^4}$, where the constant L may depend on the informational environment as well as on c_0 .⁴¹

Assuming (13), we see that the condition

$$|f(q_1, q_2 + 1, \dots, q_K) - f(q_1 + 1, q_2, \dots, q_K)| \leq \frac{c_1}{t^4}$$

implies $\left| \frac{\sigma_1^2 \cdot [Q_1]_{11}}{(q_1 - a_1)^2} - \frac{\sigma_2^2 \cdot [Q_2]_{11}}{(q_2 - a_2)^2} \right| \leq \frac{c_2}{t^4}$ and thus $\left| \left(\frac{\lambda_1}{q_1 - a_1}\right)^2 - \left(\frac{\lambda_2}{q_2 - a_2}\right)^2 \right| \leq \frac{c_3}{t^4}$ for some larger positive constants c_2, c_3 . This further implies $\left| \frac{\lambda_1}{q_1 - a_1} - \frac{\lambda_2}{q_2 - a_2} \right| \leq \frac{c_4}{t^3}$, which reduces to

$$\left| q_2 - a_2 - \frac{\lambda_2}{\lambda_1}(q_1 - a_1) \right| \leq \frac{c_5}{t}. \quad (14)$$

This inequality says that the fractional part of $\frac{\lambda_2}{\lambda_1} q_1$ is very close to the fractional part of $\frac{\lambda_2}{\lambda_1} a_1 - a_2$. But since $\frac{\lambda_2}{\lambda_1}$ is an irrational number, the fractional part of $\frac{\lambda_2}{\lambda_1} q_1$ is ‘‘equi-distributed’’ in $(0,1)$ as q_1 ranges in the positive integers.⁴² Thus the Diophantine approximation (14) only has solution at a set of times t with natural density 0, proving Lemma 5. Below we supply the technically involved proof of (13). \square

Proof of Lemma 6. Fix q_1, \dots, q_K and the signal j . Recall the diagonal matrix $D = \text{diag}\left(\frac{q_1}{\sigma_1^2}, \dots, \frac{q_K}{\sigma_K^2}\right)$. Consider any $\hat{q}_j \in [q_j, q_j + 1]$ and let \hat{D} be the analogue of D for the division (\hat{q}_j, q_{-j}) . That is, $\hat{D} = D$ except that $[\hat{D}]_{jj} = \frac{\hat{q}_j}{\sigma_j^2}$. Let $\hat{\Sigma} = CV^0C' + \hat{D}^{-1}$. From (7), we have

$$\partial_j f(\hat{q}_j, q_{-j}) = -\frac{\sigma_j^2}{\hat{q}_j^2} \cdot \left[V^0 C' \hat{\Sigma}^{-1} \Delta_{jj} \hat{\Sigma}^{-1} C V^0 \right]_{11}. \quad (15)$$

Here and later in this proof, $\partial_j f$ represents the usual continuous derivative rather than the discrete derivative.

Let $D_0 = \text{diag}\left(\frac{\lambda_1 t}{\sigma_1^2}, \dots, \frac{\lambda_K t}{\sigma_K^2}\right)$ and $\Sigma_0 = CV^0C' + D_0^{-1}$. For $|q_i - \lambda_i t| \leq c_0, \forall i$ we have $\hat{D} - D_0 = O(c_0)$, where the Big O notation applies entry-wise. It follows that

$$\hat{\Sigma} = CV^0C' + \hat{D}^{-1} = CV^0C' + D^{-1} + O\left(\frac{1}{t^2}; c_0\right) = \Sigma_0 + O\left(\frac{1}{t^2}; c_0\right).$$

⁴¹In applying Lemma 5 to prove Proposition 2, c_0 is taken to be the bound on $n_i - \lambda_i \cdot t$.

⁴²The Equi-distribution Theorem states that for any irrational number α and any sub-interval $(a, b) \subset (0, 1)$, the set of positive integers n such that the fractional part of αn belongs to (a, b) has natural density $b - a$. It is a special case of the Ergodic Theorem.

Observe that the matrix inverse is a differentiable mapping at Σ_0 (which is $CV^0C' + D_0^{-1} \succeq CV^0C'$ and thus positive definite). Thus we have

$$\hat{\Sigma}^{-1} = \Sigma_0^{-1} + O\left(\frac{1}{t^2}; c_0\right).$$

Plugging this into (15) and using $\hat{q}_j \sim \lambda_j t$, we obtain that

$$\partial_j f(\hat{q}_j, q_{-j}) = -\frac{\sigma_j^2}{\hat{q}_j^2} \cdot [V^0 C' \Sigma_0^{-1} \Delta_{jj} \Sigma_0^{-1} C V^0]_{11} + O\left(\frac{1}{t^4}; c_0\right). \quad (16)$$

Since $\Sigma_0 = CV^0C' + \frac{1}{t} \cdot \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right)$, we can apply Taylor expansion (to the matrix inverse map) and write

$$\Sigma_0^{-1} = (CV^0C')^{-1} - \frac{1}{t} (CV^0C')^{-1} \cdot \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right) \cdot (CV^0C')^{-1} + O\left(\frac{1}{t^2}\right). \quad (17)$$

This implies

$$\begin{aligned} V^0 C' \Sigma_0^{-1} \Delta_{jj} \Sigma_0^{-1} C V^0 &= V^0 C' (CV^0C')^{-1} \Delta_{jj} (CV^0C')^{-1} C V^0 - \frac{M_j}{t} + O\left(\frac{1}{t^2}\right) \\ &= Q_j - \frac{M_j}{t} + O\left(\frac{1}{t^2}\right), \end{aligned} \quad (18)$$

where M_j is a fixed $K \times K$ matrix depending only on the informational environment. For future use, we note that

$$\begin{aligned} M_j &= V^0 C' (CV^0C')^{-1} \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right) (CV^0C')^{-1} \Delta_{jj} (CV^0C')^{-1} C V^0 \\ &\quad + V^0 C' (CV^0C')^{-1} \Delta_{jj} (CV^0C')^{-1} \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right) (CV^0C')^{-1} C V^0 \\ &= C^{-1} \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right) (CV^0C')^{-1} \Delta_{jj} C'^{-1} \\ &\quad + C^{-1} \Delta_{jj} (CV^0C')^{-1} \text{diag}\left(\frac{\sigma_1^2}{\lambda_1}, \dots, \frac{\sigma_K^2}{\lambda_K}\right) C'^{-1}. \end{aligned} \quad (19)$$

Using (18), we can simplify (16) to

$$\partial_j f(\hat{q}_j, q_{-j}) = -\frac{\sigma_j^2}{\hat{q}_j^2} \cdot \left[Q_j - \frac{M_j}{t}\right]_{11} + O\left(\frac{1}{t^4}; c_0\right). \quad (20)$$

Integrating this over $\hat{q}_j \in [q_j, q_j + 1]$, we conclude that

$$f(q_j, q_{-j}) - f(q_j + 1, q_{-j}) = \frac{\sigma_j^2}{q_j(q_j + 1)} \cdot \left[Q_j - \frac{M_j}{t}\right]_{11} + O\left(\frac{1}{t^4}; c_0\right). \quad (21)$$

We set $a_j = -\left(\frac{\lambda_j \cdot [M_j]_{11}}{2[Q_j]_{11}} + \frac{1}{2}\right)$. Then,

$$\frac{\sigma_j^2}{q_j(q_j+1)} \cdot \left[Q_j - \frac{M_j}{t}\right]_{11} = (\sigma_j^2 \cdot [Q_j]_{11}) \cdot \frac{1 + \frac{2a_j+1}{\lambda_j t}}{q_j(q_j+1)} = \frac{\sigma_j^2 \cdot [Q_j]_{11}}{(q_j - a_j)^2} + O\left(\frac{1}{t^4}; c_0\right),$$

implying the desired approximation (13). The last equality above uses $\frac{1 + \frac{2a_j+1}{\lambda_j t}}{q_j(q_j+1)} = \frac{1}{(q_j - a_j)^2} + O\left(\frac{1}{t^4}; c_0\right)$, which is because

$$\frac{q_j(q_j+1)}{(q_j - a_j)^2} = 1 + \frac{2(a_j+1)}{q_j - a_j} + O\left(\frac{1}{(q_j - a_j)^2}\right) = 1 + \frac{2a_j+1}{\lambda_j t} + O\left(\frac{1}{t^2}; c_0\right)$$

dividing through by $q_j(q_j+1)$. □

9.6.3 A Simultaneous Diophantine Approximation Problem

The above Lemma 5 tells us that at most times t , there do not exist a *pair* of divisions (differing minimally on two signal counts) that lead to posterior variances close to each other (with a difference of $\frac{c_1}{t^4}$). We obtain a stronger result if a *triple* of such divisions were to exist.

Lemma 7. *Fix V^0 and C , and let signal variances vary. For positive constants c_0, c_1 , define $\mathcal{A}^*(c_0, c_1)$ to be the following set of positive integers:*

$$\begin{aligned} \{t : \exists q_1, q_2, q_3, \dots, q_K \in \mathbb{Z}^+, \text{ s.t. } & |q_i - \lambda_i t| \leq c_0, \forall i \\ & \wedge |f(q_1, q_2 + 1, q_3, \dots, q_K) - f(q_1 + 1, q_2, q_3, \dots, q_K)| \leq c_1/t^4 \\ & \wedge |f(q_1, q_2, q_3 + 1, \dots, q_K) - f(q_1 + 1, q_2, q_3, \dots, q_K)| \leq c_1/t^4\} \end{aligned}$$

Then, for generic signal variances, $\mathcal{A}^(c_0, c_1)$ has finite cardinality.*

Proof. So far we have been dealing with fixed informational environments. However, a number of parameters defined above depend on the signal variances $\sigma = \{\sigma_i^2\}_{i=1}^K$. Specifically, while the matrix $Q_i = C^{-1} \Delta_{ii} C'^{-1}$ is independent of σ , the asymptotic proportions $\lambda_i \propto \sigma_i \cdot [Q_i]_{11}$ do vary with σ . In this proof, we write $\lambda_i(\sigma)$ to highlight this dependence.

Next, we recall the matrix M_j introduced earlier in (19). We note that for fixed matrices V^0 and C , each entry of $M_j(\sigma)$ is a fixed linear combination of $\frac{\sigma_1^2}{\lambda_1(\sigma)}, \dots, \frac{\sigma_K^2}{\lambda_K(\sigma)}$.

Then, the parameter $a_j(\sigma)$ in (13) is given by (see the previous proof)

$$a_j(\sigma) = -\frac{1}{2} - \frac{\lambda_j(\sigma) \cdot [M_j(\sigma)]_{11}}{2[Q_j]_{11}} = -\frac{1}{2} + \lambda_j(\sigma) \sum_{i=1}^K \tilde{b}_{ji} \frac{\sigma_i^2}{\lambda_i(\sigma)} = -\frac{1}{2} + \sum_{i=1}^K b_{ji} \sigma_i \sigma_j \quad (22)$$

for some constants \tilde{b}_{ji}, b_{ji} independent of σ . In the last equality above, we used the fact that $\frac{\lambda_j(\sigma)}{\lambda_i(\sigma)}$ equals a constant times $\frac{\sigma_j}{\sigma_i}$.

Thus Lemma 6 gives

$$f(q_j, q_{-j}) - f(q_j + 1, q_{-j}) = \frac{\sigma_j^2 \cdot [Q_j]_{11}}{(q_j - a_j(\sigma))^2} + O\left(\frac{1}{t^4}; c_0\right)$$

whenever $|q_i - \lambda_i(\sigma) \cdot t| \leq c_0, \forall i$. We comment that the Big O constant here may depend on σ . However, a single constant suffices if we restrict each σ_i to be bounded above and bounded away from zero. Since measure-zero sets are closed under countable unions, this restriction does not affect the result we want to prove.

By the above approximation, a necessary condition for $t \in \mathcal{A}^*(c_0, c_1)$ is that q_1, q_2, q_3 satisfy

$$\left| (q_2 - a_2(\sigma)) - \frac{\eta \cdot \sigma_2}{\sigma_1} (q_1 - a_1(\sigma)) \right| \leq \frac{c_6}{q_1} \quad (23)$$

as well as

$$\left| (q_3 - a_3(\sigma)) - \frac{\kappa \cdot \sigma_3}{\sigma_1} (q_1 - a_1(\sigma)) \right| \leq \frac{c_6}{q_1} \quad (24)$$

for some constant c_6 independent of σ (c_6 may depend on c_0, c_1 stated in the lemma). The constant η is given by $\eta = \sqrt{[Q_2]_{11}/[Q_1]_{11}}$, and similarly for κ .

It remains to show that for generic σ , there are only finitely many positive integer triples (q_1, q_2, q_3) satisfying the *simultaneous Diophantine approximation* (23) and (24). To prove this, we assume that each σ_i is i.i.d. drawn from the uniform distribution on $[\frac{1}{L}, L]$, where L is a large constant. Denote by $F(q_1, q_2, q_3)$ the event that (23) and (24) hold simultaneously. We claim that there exists a constant c_7 such that $\mathbb{P}(F(q_1, q_2, q_3)) \leq \frac{c_7}{q_1^4}$ holds for all q_1, q_2, q_3 . Since $F(q_1, q_2, q_3)$ cannot occur for $q_2, q_3 > c_8 q_1$, this claim will imply

$$\sum_{q_1, q_2, q_3} \mathbb{P}(F(q_1, q_2, q_3)) < \sum_{q_1} \sum_{q_2, q_3 \leq c_8 q_1} \frac{c_7}{q_1^4} < \sum_{q_1} \frac{c_7 c_8^2}{q_1^2} < \infty. \quad (25)$$

Generic finiteness of tuples (q_1, q_2, q_3) will then follow from the Borel-Cantelli Lemma.⁴³

To prove this claim, it suffices to show that if $\sigma = (\sigma_1, \sigma_2, \sigma_3, \sigma_4, \dots, \sigma_K)$ and $\sigma' = (\sigma_1, \sigma'_2, \sigma'_3, \sigma_4, \dots, \sigma_K)$ both satisfy (23) and (24), then $|\sigma_2 - \sigma'_2|, |\sigma_3 - \sigma'_3| \leq \frac{c}{q_1^2}$ for some

⁴³Because of the use of Borel-Cantelli Lemma, this proof (unlike Lemma 5 above) does not allow us to effectively determine for given σ whether (23) and (24) only have finitely many integer solutions. Nonetheless, a modification of this proof does imply the following finite-time probabilistic statement: when $\sigma_1, \dots, \sigma_K$ are independently drawn, the probability that the optimality strategy coincides with t -optimality at every period $t \geq T$ is at least $1 - O(\frac{1}{T})$, where the constant involved only depends on the distribution of σ .

constant c .⁴⁴ Without loss, we assume $|\sigma_2 - \sigma'_2| \geq |\sigma_3 - \sigma'_3|$. Using (22), we can rewrite the condition (23) as

$$\left| \underbrace{\left(q_2 + \frac{1}{2} \right) - \frac{\eta \cdot \sigma_2}{\sigma_1} \left(q_1 + \frac{1}{2} \right)}_A + \underbrace{\sum_i \beta_i \sigma_2 \sigma_i}_B \right| \leq \frac{c_6}{q_1}$$

for some constants β_i independent of σ . A similar inequality holds at σ' :

$$\left| \underbrace{\left(q_2 + \frac{1}{2} \right) - \frac{\eta \cdot \sigma'_2}{\sigma_1} \left(q_1 + \frac{1}{2} \right)}_{A'} + \underbrace{\sum_i \beta_i \sigma'_2 \sigma'_i}_{B'} \right| \leq \frac{c_6}{q_1}$$

It follows from the above two inequalities that $|A + B - A' - B'| \leq \frac{2c_6}{q_1}$. Furthermore, since $|A - A'| \leq |A + B - A' - B'| + |B - B'|$ (by triangle inequality), we deduce

$$\left| \frac{\eta \cdot (\sigma'_2 - \sigma_2)}{\sigma_1} \cdot \left(q_1 + \frac{1}{2} \right) \right| \leq \frac{2c_6}{q_1} + \left| \sum_i \beta_i (\sigma'_2 \sigma'_i - \sigma_2 \sigma_i) \right| \quad (26)$$

Because $\sigma'_i = \sigma_i$ for $i \neq 2, 3$, we have

$$\begin{aligned} \left| \sum_i \beta_i (\sigma'_2 \sigma'_i - \sigma_2 \sigma_i) \right| &= \left| \sum_i \beta_i (\sigma'_2 - \sigma_2) \sigma_i + \sum_i \beta_i \sigma'_2 (\sigma'_i - \sigma_i) \right| \\ &= \left| \left(\sum_i \beta_i (\sigma'_2 - \sigma_2) \sigma_i \right) + \beta_2 \sigma'_2 (\sigma'_2 - \sigma_2) + \beta_3 \sigma'_2 (\sigma'_3 - \sigma_3) \right| \\ &\leq (K + 2)L \cdot \max_i |\beta_i| \cdot |\sigma'_2 - \sigma_2| \end{aligned}$$

Plugging this estimate into (26), we obtain the desired result $|\sigma_2 - \sigma'_2| \leq \frac{c}{q_1^2}$. This completes the proof of the lemma. \square

9.6.4 Monotonicity of t -Optimal Divisions

We apply Lemma 7 to prove the eventual monotonicity of t -optimal divisions in generic informational environments.

Proposition 3. *Fix V^0 and C . For generic signal variances $\{\sigma_i^2\}_{i=1}^K$, there exists T_0 such that for $t \geq T_0$, the t -optimal division $n(t)$ is unique, and it satisfies $n_i(t+1) \geq n_i(t), \forall i$.*

⁴⁴This implies that the probability of the event $F(q_1, q_2, q_3)$ conditional on any value of $\sigma_1, \sigma_4, \dots, \sigma_K$ is bounded by $\frac{c_7}{q_1^4}$, which is stronger than the claim.

Proof. Uniqueness follows from the stronger fact that in generic informational environments, $f(q_1, \dots, q_K)$ differs from $f(q'_1, \dots, q'_K)$ whenever $q \neq q'$. Below we focus on monotonicity.

Using the order difference lemma, we can already deduce the difference $|n_i(t+1) - n_i(t)|$ is no more than 1 at sufficiently late periods t . Suppose that $n_1(t+1) = n_1(t) - 1$. Then because $\sum_i (n_i(t+1) - n_i(t)) = 1$, we can without loss assume $n_2(t+1) = n_2(t) + 1$ and $n_3(t+1) = n_3(t) + 1$.

For notational ease, write $n_i = n_i(t), n'_i = n_i(t+1)$. By t -optimality, we have

$$f(n_1, n_2, n_3, \dots, n_K) \leq f(n_1 - 1, n_2 + 1, n_3, \dots, n_K)$$

$$f(n'_1, n'_2, n'_3, \dots, n'_K) \leq f(n'_1 + 1, n'_2 - 1, n'_3, \dots, n'_K)$$

These inequalities are equivalent to

$$\partial_2 f(n_1 - 1, n_2, n_3, \dots, n_K) \geq \partial_1 f(n_1 - 1, n_2, n_3, \dots, n_K) \quad (27)$$

$$\partial_2 f(n'_1, n'_2 - 1, n'_3, \dots, n'_K) \leq \partial_1 f(n'_1, n'_2 - 1, n'_3, \dots, n'_K) \quad (28)$$

with $\partial_i f$ representing the *discrete partial derivative*.

Since $n'_2 - 1 = n_2$, the LHS of (28) is at least the LHS of (27) minus a number of cross partials. Similarly, the RHS of (28) is at most bigger than the RHS of (27) by a number of cross partials. Thus the only way (27) and (28) can both hold is if the two sides of (27) differ by no more than $O(\frac{1}{t^4})$. That is, for some absolute constant c_1 ,⁴⁵ we have

$$|f(n_1 - 1, n_2 + 1, n_3, \dots, n_K) - f(n_1, n_2, n_3, \dots, n_K)| \leq \frac{c_1}{t^4}. \quad (29)$$

An analogous argument yields

$$|f(n_1 - 1, n_2, n_3 + 1, \dots, n_K) - f(n_1, n_2, n_3, \dots, n_K)| \leq \frac{c_1}{t^4}. \quad (30)$$

But now we can apply Lemma 7 to show that in generic environments, there are only finitely many integer tuples (n_1, \dots, n_K) that satisfy both (29) and (30). This proves the result. \square

9.6.5 Completing the Proof of Theorem 3

By Proposition 3, generically there exists T_0 such that $n(t)$ is monotonic in t after T_0 periods. Thus, using our dynamic Blackwell lemma, if the DM achieves t -optimality at some period $t \geq T_0$, he will continue to do so in the future. By Proposition 2, such a time t does exist. This proves Theorem 3.

⁴⁵As discussed in the proof of Lemma 7, we can find a single constant c_1 that works for all σ bounded above and bounded away from zero.

9.7 Proof of Proposition 1 (Bound on B)

9.7.1 Preliminary Estimates

Throughout, we work with the linearly-transformed model, where each signal X_i is simply $\tilde{\theta}_i$ plus standard Gaussian noise, and the DM's prior covariance matrix over the transformed states is \tilde{V} . Let $\gamma = \gamma(q_1, \dots, q_K)$ represent the following $K \times 1$ vector:

$$\gamma = (\tilde{V} + E)^{-1} \cdot \tilde{V} \cdot w \quad (31)$$

with $E = \text{diag}(\frac{1}{q_1}, \dots, \frac{1}{q_K})$. For $1 \leq i \leq K$, γ_i denotes the i -th coordinate of γ .

Here we re-derive the posterior variance function f , its (usual continuous) derivatives and second derivatives. Our formulae below take as primitives \tilde{V} and w , but they are equivalent to those presented in Appendix 9.1 (for the original model).

Fact 1 (Posterior Variance).

$$f(q_1, \dots, q_K) = w'(\tilde{V} - \tilde{V}(\tilde{V} + E)^{-1}\tilde{V})w.$$

Fact 2 (Partial Derivatives of Posterior Variance).

$$\partial_i f(q_1, \dots, q_K) = -\frac{1}{q_i^2} \cdot w' \tilde{V} (\tilde{V} + E)^{-1} \Delta_{ii} (\tilde{V} + E)^{-1} \tilde{V} w = -\frac{\gamma_i^2}{q_i^2}.$$

Fact 3 (Second-Order Partial Derivatives of Posterior Variance).

$$\begin{aligned} \partial_{ii} f(q_1, \dots, q_K) &= \frac{2 \cdot w' \tilde{V} (\tilde{V} + E)^{-1} \Delta_{ii} (\tilde{V} + E)^{-1} \tilde{V} w}{q_i^3} - \frac{2 \cdot w' \tilde{V} (\tilde{V} + E)^{-1} \Delta_{ii} (\tilde{V} + E)^{-1} \Delta_{ii} (\tilde{V} + E)^{-1} \tilde{V} w}{q_i^4} \\ &= \frac{2\gamma_i^2}{q_i^3} \cdot \left(1 - \frac{[(\tilde{V} + E)^{-1}]_{ii}}{q_i} \right) \end{aligned}$$

Fact 4 (Cross-Partial Derivatives of Posterior Variance).

$$\begin{aligned} \partial_{ij} f(q_1, \dots, q_K) &= \frac{-2}{q_i^2 q_j^2} \cdot w' \tilde{V} (\tilde{V} + E)^{-1} \Delta_{ii} (\tilde{V} + E)^{-1} \Delta_{jj} (\tilde{V} + E)^{-1} \tilde{V} w \\ &= \frac{-2\gamma_i \gamma_j}{q_i^2 q_j^2} \cdot [(\tilde{V} + E)^{-1}]_{ij}. \end{aligned}$$

All of the above facts can be proved by simple linear algebra, so we omit the details.

9.7.2 Refined Asymptotic Characterization of $n(t)$

We now specialize to $w = \mathbf{1}$ and establish the next lemma, which refines our asymptotic characterization of $n(t)$ in Appendix 9.3.⁴⁶ Proposition 1 will immediately follow.

⁴⁶Easy to see that in the transformed model, $\lambda_i \propto |w_i|$. So $\lambda_i = \frac{1}{K}$ here.

Lemma 8. For $t \geq 8(R+1)K\sqrt{K}$, it holds that $|n_i(t) - \frac{t}{K}| \leq 4(R+1)\sqrt{K}$.

Proof. Note from (31) that $(\tilde{V} + E)\gamma = \tilde{V}w$. So $\tilde{V}(w - \gamma) = E\gamma = (\frac{\gamma_1}{q_1}, \dots, \frac{\gamma_K}{q_K})'$, and

$$w - \gamma = (\tilde{V})^{-1} \cdot \left(\frac{\gamma_1}{q_1}, \dots, \frac{\gamma_K}{q_K} \right)'.$$

From the definition of the operator norm, we deduce

$$\sum_{i=1}^K (1 - \gamma_i)^2 = \|w - \gamma\|^2 \leq R^2 \cdot \left(\sum_{j=1}^K \frac{\gamma_j^2}{q_j^2} \right). \quad (32)$$

This holds for any division vector q and the corresponding γ (which is a function of q).

Now suppose without loss of generality that $n_1(t) \geq \frac{t}{K}$. Let $q = (n_1(t) - 1, n_2(t), \dots, n_K(t))$ and consider the corresponding γ . Then from t -optimality we have

$$|f(q_1 + 1, q_{-1}) - f(q)| \geq |f(q_j + 1, q_{-j}) - f(q)|, \quad \forall j.$$

Note that the discrete partial derivatives above are related to the usual continuous partials by the following inequalities:⁴⁷

$$\frac{\gamma_j^2}{q_j(q_j + 1)} \leq |f(q_j + 1, q_{-j}) - f(q)| \leq \frac{\gamma_j^2}{q_j^2}.$$

We therefore deduce

$$\frac{\gamma_1^2}{q_1^2} \geq \frac{\gamma_j^2}{q_j(q_j + 1)}, \quad \forall j. \quad (33)$$

Combining (32) and (33) and using $\frac{1}{q_j^2} \leq \frac{2}{q_j(q_j+1)}$, we see that

$$\sum_{i=1}^K (1 - \gamma_i)^2 \leq 2R^2K \cdot \frac{\gamma_1^2}{q_1^2}. \quad (34)$$

In particular, we know that $\gamma_1 - 1 \leq R\sqrt{2K} \cdot \frac{\gamma_1}{q_1}$. Easy to see this implies

$$\gamma_1 \leq 1 + \frac{2R\sqrt{K}}{q_1} \leq \sqrt{2} \quad (35)$$

whenever $q_1 = n_1(t) - 1 \geq \frac{t}{K} - 1 \geq (2\sqrt{2} + 2)R\sqrt{K}$. Plugging this back into the RHS of (34), we then obtain

$$\gamma_j \geq 1 - \frac{2R\sqrt{K}}{q_1} \geq 2 - \sqrt{2} \quad (36)$$

⁴⁷The RHS follows from the convexity of f ; the LHS can be proved by using Fact 2, Fact 3 and noting that γ_j^2 is an increasing function in q_j , because $\frac{\partial \gamma_j(q)}{\partial q_j} = \frac{\gamma_j}{q_j^2} \cdot [(V + E)^{-1}]_{jj}$ has the same sign of γ_j .

Now use (33), (35) and (36) to deduce that

$$q_j + 1 \geq \frac{\gamma_j}{\gamma_1} \cdot q_1 \geq \frac{1 - \frac{2R\sqrt{K}}{q_1}}{1 + \frac{2R\sqrt{K}}{q_1}} \cdot q_1 \geq \left(1 - \frac{4R\sqrt{K}}{q_1}\right) \cdot q_1 = q_1 - 4R\sqrt{K}.$$

Recall $q_j = n_j(t)$ for $j > 1$ and $q_1 = n_1(t) - 1$. We thus have

$$n_j(t) \geq n_1(t) - 4R\sqrt{K} - 2. \tag{37}$$

Since $n_1(t) \geq \frac{t}{K}$, the above implies $n_j(t) \geq \frac{t}{K} - 4(R+1)\sqrt{K}$ for each signal j . This proves half of the lemma.

For the other half, note that $n_j(t) \leq \frac{t}{K}$ must hold for *some* signal j . Thus (37) yields $n_1(t) \leq \frac{t}{K} + 4(R+1)\sqrt{K}$. This is not just true for signal 1, but in fact for any signal i with $n_i(t) \geq \frac{t}{K}$. So we conclude $n_i(t) \leq \frac{t}{K} + 4(R+1)\sqrt{K}$ for each signal i . The proof of the lemma is complete. \square

References

- AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- ANDO, T. (1983): “On the Arithmetic-Geometric-Harmonic-Mean Inequalities for Positive Definite Matrices,” *Linear Algebra and its Applications*, 52, 31–37.
- ARROW, K. J., D. BLACKWELL, AND M. A. GIRSHICK (1949): “Bayes and Minimax Solutions of Sequential Decision Problems,” *Econometrica*, 17, 213–244.
- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- BANKS, J. S. AND R. K. SUNDARAM (1992): “A Class of Bandit Problems Yielding Myopically Optimal Strategies,” *Journal of Applied Probability*, 29, 625–632.
- BARDHI, A. (2017): “Optimal Discovery and Influence Through Selective Sampling,” Working Paper.
- BERGEMANN, D. AND J. VÄLIMÄKI (2002): “Information Acquisition and Efficient Mechanism Design,” *Econometrica*, 70, 1007–1033.
- BERGEMANN, D. AND J. VÄLIMÄKI (2008): “Bandit Problems,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Basingstoke: Palgrave Macmillan.
- BERRY, D. AND B. FRISTEDT (1988): “Optimality of Myopic Stopping Times for Geometric Discounting,” *Journal of Applied Probability*, 25, 437–443.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades,” *Journal of Political Economy*, 100, 992–1026.
- BUBECK, S., R. MUNOS, AND G. STOLTZ (2009): “Pure Exploration in Multi-armed Bandits Problems,” in *Algorithmic Learning Theory. ALT 2009. Lecture Notes in Computer Science, vol 5809*, Springer, Berlin, Heidelberg.
- CALLANDER, S. (2011): “Searching and Learning by Trial and Error,” *American Economic Review*, 101, 2277–2308.
- CHE, Y.-K. AND K. MIERENDORFF (2017): “Optimal Sequential Decision with Limited Attention,” Working Paper.
- CHERNOFF, H. (1972): *Sequential Analysis and Optimal Design*, Society for Industrial and Applied Mathematics.
- CHICK, S. E. AND P. FRAZIER (2012): “Sequential Sampling with Economics of Selection Procedures,” *Management Science*, 58, 550–569.

- COLOMBO, L., G. FEMMINIS, AND A. PAVAN (2014): “Information Acquisition and Welfare,” *Review of Economic Studies*, 81, 1438–1483.
- DENTI, T. (2017): “Unrestricted Information Acquisition,” Working Paper.
- DEWAN, T. AND D. P. MYATT (2008): “The Qualities of Leadership: Direction, Communication and Obfuscation,” *American Political Science Review*, 102, 351–368.
- EASLEY, D. AND N. M. KIEFER (1988): “Controlling a Stochastic Process with Unknown Parameters,” *Econometrica*, 56, 1045–1064.
- FUDENBERG, D., P. STRACK, AND T. STRZALECKI (2017): “Stochastic Choice and Optimal Sequential Sampling,” Working Paper.
- GARFAGNINI, U. AND B. STRULOVICI (2016): “Social Experimentation with Interdependent and Expanding Technologies,” *Review of Economic Studies*, 83, 1579–1613.
- GITTINS, J. C. (1979): “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society, Series B*, 148–177.
- GREENSHTEIN, E. (1996): “Comparison of Sequential Experiments,” *The Annals of Statistics*, 24, 436–448.
- HANSEN, O. H. AND E. N. TORGERSEN (1974): “Comparison of Linear Normal Experiments,” *The Annals of Statistics*, 2, 367–373.
- HÉBERT, B. AND M. WOODFORD (2017): “Rational Inattention with Sequential Information Sampling,” Working Paper.
- HELLWIG, C. AND L. VELDKAMP (2009): “Knowing What Others Know: Coordination Motives in Information Acquisition,” *The Review of Economic Studies*, 76, 223–251.
- KE, T. T. AND J. M. VILLAS-BOAS (2017): “Optimal Learning Before Choice,” Working Paper.
- KEENER, R. (1985): “Further Contributions to the Two-armed Bandit Problem,” *Annals of Statistics*, 13, 418–422.
- LAMBERT, N., M. OSTROVSKY, AND M. PANOV (2017): “Strategic Trading in Informationally Complex Environments,” Working Paper.
- LIANG, A. AND X. MU (2017): “Overabundant Information and Learning Traps,” Working Paper.
- MAYSKAYA, T. (2017): “Dynamic Choice of Information Sources,” Working Paper.
- MERSEREAU, A. J., P. RUSMEVICHIENTONG, AND J. N. TSITSIKLIS (2009): “A Structured Multiarmed Bandit Problem and the Greedy Policy,” *IEEE Transactions On Automatic Control*, 54, 2787–2802.
- MEYER, M. AND J. ZWIEBEL (2007): “Learning and Self-Reinforcing Behavior,” Working Paper.

- MOSCARINI, G. AND L. SMITH (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69, 1629–1644.
- MYATT, D. P. AND C. WALLACE (2012): “Endogenous Information Acquisition in Coordination Games,” *The Review of Economic Studies*, 79, 340–374.
- PERSICO, N. (2000): “Information Acquisition in Auctions,” *Econometrica*, 68, 135–148.
- REINGANUM, J. (1983): “Nash Equilibrium Search for the Best Alternative,” *Journal of Economic Theory*, 30, 139–152.
- ROBBINS, H. (1952): “Some aspects of the sequential design of experiments,” *Bull. Amer. Math. Soc.*, 58, 527–535.
- ROTHSCHILD, M. (1974): “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 9, 185–202.
- RUSSO, D. (2016): “Simple Bayesian Algorithms for Best-Arm Identification,” Working Paper.
- SANJURJO, A. (2017): “Search with Multiple Attributes: Theory and Empirics,” *Games and Economic Behavior*, 535–562.
- SETHI, R. AND M. YILDIZ (2016): “Communication with Unknown Perspectives,” *Econometrica*, 84, 2029–2069.
- STEINER, J., C. STEWART, AND F. MATĚJKA (2017): “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” *Econometrica*, 85, 521–553.
- TAYLOR, C. R. (1995): “Digging for Golden Carrots: An Analysis of Research Tournaments,” *American Economic Review*, 85, 872–890.
- WALD, A. (1947): “Foundations of a General Theory of Sequential Decision Functions,” *Econometrica*, 15, 279–313.
- WEITZMAN, M. L. (1979): “Optimal Search for the Best Alternative,” *Econometrica*, 47, 641–654.
- YANG, M. (2015): “Coordination with Flexible Information Acquisition,” *Journal of Economic Theory*, 158, 721–738.
- ZHONG, W. (2017): “Optimal Dynamic Information Acquisition,” Working Paper.

10 Online Appendix

10.1 Applications of Results from Section 5 (Multi-player Games)

10.1.1 Beauty Contest

Hellwig and Veldkamp (2009) introduced a beauty contest game with endogenous one-shot information acquisition. We build on this by modifying the information acquisition stage so that players *sequentially* acquire information over many periods (rather than once), and face a *capacity constraint* each period (rather than costly signals). We show that the basic insights of Hellwig and Veldkamp (2009) hold in this setting.

Specifically, suppose that at an unknown final period, a unit mass of players simultaneously chooses prices $p_i \in \mathbb{R}$ to minimize the (normalized) squared distance between their price and an unknown target price p^* , which depends on the unknown state ω and also on the average price $\bar{p} = \int p_i di$:

$$u_i(p_i, \bar{p}, \omega) = -\frac{1}{(1-r)^2} \cdot (p_i - p^*)^2 \quad \text{where } p^* = (1-r) \cdot \omega + r \cdot \bar{p} \quad (38)$$

The constant $r \in (-1, 1)$ determines whether pricing decisions are complements or substitutes.⁴⁸

In every period up until the final period, each player acquires B signals from the set (X_k^i) , as in the framework we have developed. To closely mirror the setup in Hellwig and Veldkamp (2009), we set each $\theta_1^i = \omega$. Assuming “conditional independence” of players’ signals, we can directly apply Corollary 1 and conclude that in every equilibrium, players choose a deterministic (myopic) sequence of information acquisitions. This result echoes Hellwig and Veldkamp (2009), who show that equilibrium is unique when players choose from *private* signals (see their Subsection 1.3.4).⁴⁹ Our extension is to introduce dynamics and show how the dynamic problem can be reduced into a static one.

Let $\Sigma(t)$ be the posterior variance about ω after the first t myopic observations. Since the players in our model acquire B signals each period, their (common) posterior variance at the

⁴⁸When $r > 0$, best responses are increasing in the prices set by other players, thus decisions are complements. Conversely, $r < 0$ implies decisions are substitutes.

⁴⁹Hellwig and Veldkamp (2009) also study a case in which players observe signals that are distorted by a common noise (which violates conditional independence). They show that multiple equilibria generally arise with such “public signals”. Dewan and Myatt (2008), Myatt and Wallace (2012) and Colombo et al. (2014) restore a unique linear symmetric equilibrium by assuming perfectly divisible signals, similar to the continuous-time variant of our model. In contrast, our equilibrium analysis relies on the informational environment (i.e. conditional independence), but not on symmetry or linearity of the strategy.

end of t periods is given by $\Sigma(Bt)$. Thus, conditional on period t being the final period, our game is as if the players acquire a batch of Bt signals and then choose prices. This means that equilibrium prices are determined in the same way as in [Hellwig and Veldkamp \(2009\)](#):

$$p(\mathcal{I}_{\leq Bt}^i) = \frac{1-r}{1-r+r \cdot \Sigma(Bt)} \cdot \mathbb{E}(\omega | \mathcal{I}_{\leq Bt}^i) \quad (39)$$

where $\mathcal{I}_{\leq Bt}^i$ represents player i 's information set, consisting of Bt signal realizations.

We can use this characterization of equilibrium to re-evaluate the main insight in [Hellwig and Veldkamp \(2009\)](#): the incentive to acquire more informative signals is increasing in aggregate information acquisition if decisions are complements and decreasing if decisions are substitutes. For this purpose, we augment the model with a period 0, in which each player i invests in a capacity level B_i at some cost. Afterwards, players acquire information myopically (under possibly differential capacity constraints) and participate in the beauty contest game.

Let $\mu \in \Delta(\mathbb{Z}^+)$ be the distribution over capacity levels chosen by player i 's opponents. Then, player i 's expected utility from choosing capacity B_i is given by

$$EU(B_i, \mu) = -\mathbb{E}_{t \sim \pi} \left[\frac{\Sigma(B_i t)}{(1-r+r \cdot \int_B \Sigma(Bt) d\mu(B))^2} \right]. \quad (40)$$

Above, the expectation is taken with respect to the random final period t distributed according to π , while inside the expectation, the term $\int_B \Sigma(Bt) d\mu(B)$ is the average posterior variance among the players. Similar to Proposition 1 in [Hellwig and Veldkamp \(2009\)](#), we have the following result:

Corollary 5. *Suppose $\hat{B}_i > B_i$ and $\hat{\mu} > \mu$ in the sense of first-order stochastic dominance. Then the sign of the difference $EU(B_i, \mu) + EU(\hat{B}_i, \hat{\mu}) - EU(B_i, \hat{\mu}) - EU(\hat{B}_i, \mu)$ is*

- (a) zero, if there is no strategic interaction ($r = 0$);
- (b) positive, if decisions are complementary ($r > 0$);
- (c) negative, if decisions are substitutes ($r < 0$).

When decisions are complements, the value of additional information is increasing in the amount of *aggregate* information. Thus player i has a stronger incentive to choose a higher signal capacity if his opponents (on average) acquire more signals. This incentive goes in the opposite direction when decisions are substitutes, which confirms the main finding of [Hellwig and Veldkamp \(2009\)](#).

10.1.2 Strategic Trading

We consider the strategic trading game introduced in [Lambert et al. \(2017\)](#), in which individuals trade given asymmetric information about the value of an asset. We endogenize the information available to traders by adding a pre-trading stage in which traders sequentially acquire signals. As before, we suppose that trading occurs at a final time period that is determined according to an arbitrary full-support distribution.

In more detail: at the final time period, a security with unknown value v is traded in a market, and each of n traders submits a demand d_i . There are additionally liquidity traders who generate exogenous random demand u . A market-maker privately observes a signal θ_M (possibly multi-dimensional) and the total demand $D = \sum_i d_i + u$. He sets the price $P(\theta_M, D)$, which in equilibrium equals $\mathbb{E}[v \mid \theta_M, D]$. Each strategic trader then obtains profit $\Pi_i = d_i \cdot (v - P(\theta_M, D))$.

We suppose that in each period up to and including the final time period, each trader i chooses to observe a signal from his set (X_k^i) (described in Section 5). The requirement of conditional independence is strengthened to apply to a payoff-relevant *vector* $\omega = (v, \theta_M, u)$ (instead of a real-valued unknown): that is, for each player i , conditional on the value of θ_1^i , the payoff-relevant vector ω and the other players' unknown states $(\theta^j)_{j \neq i}$ are assumed to be conditionally independent from player i 's states θ^i . Relative to the fully general setting considered in [Lambert et al. \(2017\)](#), this assumption allows for flexible correlation *within* a player's signals, but places a strong restriction on the correlation *across* different players' signals. Applying Corollary 1, we can conclude that:

Corollary 6. *Under the above assumptions, there is an essentially unique linear NE in which the on-path signal acquisitions are myopic, and in the final period, players play the unique linear equilibrium described in [Lambert et al. \(2017\)](#).*

Thus, the closed-form solutions that are a key contribution of [Lambert et al. \(2017\)](#) extend to our dynamic setting with endogenous information.

10.2 Example in Which $n(t)$ is Not Monotone Even for Large t

Here we continue to study Example 2 presented in Figure 1 of the main text. We will show that the t -optimal division vectors $n(t)$ fail to be monotone, even when we consider only periods after some large T .

The posterior variance function is

$$f(q_1, q_2, q_3) = 1 - \frac{1}{1 + \frac{1}{q_1} + 1 - \frac{1}{1 + \frac{1}{q_2} + \frac{1}{1 + q_3}}}$$

This suggests that the t -optimal problem can be separated into two parts: choosing q_1 , and allocating the remaining observations between q_2 and q_3 . The latter allocation problem is simple: an optimal division satisfies $q_3 = q_2 - 1$ or $q_3 = q_2$. With some extra algebra, we obtain that for $N \geq 1$:

1. If $t = 3N + 1$, then the unique t -optimal division is $(N + 2, N, N - 1)$;
2. If $t = 3N + 2$, then the unique t -optimal division is $(N + 3, N, N - 1)$;
3. If $t = 3N + 3$, then the unique t -optimal division is $(N + 2, N + 1, N)$.

Crucially, note that when transitioning from $t = 3N + 2$ to $t = 3N + 3$, the t -optimal number of X_1 signals is decreased. This reflects the complementarity between signals X_2 and X_3 , which causes the DM to observe them in pairs. Due to this failure of monotonicity, a sequential rule cannot achieve the t -optimal division vectors for all large t .

10.3 Additional Result for $K = 2$

When there are only two states and two signals, we can show that for a broad class of environments, the myopic information acquisition strategy is optimal from period 1.⁵⁰

Proposition 4. *Suppose $K = 2$, the prior is standard Gaussian ($V^0 = \mathbf{I}_2$), and both signals have variance 1.⁵¹ Write $C = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and assume without loss that $|ad| \geq |bc|$. Then the optimal information acquisition strategy is myopic whenever the following inequality holds:*

$$(1 + 2b^2) \cdot |ad - bc| \geq |ad + bc|. \quad (41)$$

In particular, this is true whenever $abcd \leq 0$.

To interpret, (41) requires that the determinant of the matrix C , $ad - bc$, is not too small (holding other terms constant). Equivalently, the two vectors (in \mathbb{R}^2) defining the signals should not be close to collinear. This rules out situations where the two signals provide such similar information in the initial periods that they substitute one another.

⁵⁰In the proposition below, if the linear coefficients a, b, c, d were picked at random, then with probability $\frac{1}{2}$ we would have $abcd \leq 0$.

⁵¹We make these simplifying assumptions so that the condition for immediate equivalence is easy to state and interpret.

Proof. Under the assumptions, the DM's posterior variance about θ_1 is computed to be

$$f(q_1, q_2) = \frac{1 + b^2 q_1 + d^2 q_2}{1 + (a^2 + b^2)q_1 + (c^2 + d^2)q_2 + (ad - bc)^2 q_1 q_2}$$

Given q_i observations of each signal i in the past, the myopic strategy chooses signal 1 if and only if $f(q_1 + 1, q_2) < f(q_1, q_2 + 1)$, which reduces to

$$\begin{aligned} & (ad - bc)^2 b^2 q_1^2 + (1 + b^2)(ad - bc)^2 q_1 - (a^2 d^2 - b^2 c^2) q_1 + c^2(1 + b^2) \\ < & (ad - bc)^2 d^2 q_2^2 + (1 + d^2)(ad - bc)^2 q_2 + (a^2 d^2 - b^2 c^2) q_2 + a^2(1 + d^2) \end{aligned} \quad (42)$$

The condition $|ad| \geq |bc|$ ensures that the RHS is an increasing function of q_2 , because the coefficients in front of q_2^2 and q_2 are both positive. Meanwhile, the condition $(1 + 2b^2)|ad - bc| \geq |ad + bc|$ implies the LHS is larger when $q_1 = 1$ than when $q_1 = 0$, so that the LHS is also increasing in (integer values of) q_1 .

Even if f may not be written into separable form, (42) suggests that the comparison between the marginal values of signal 1 and 2 “is separable.” It follows that the t -optimal division vectors are increasing in t . Proposition 4 is proved. \square

10.4 Eventual Optimality of the Myopic Strategy

Below, write $m(t)$ for the division vector at time t achieved under the (history-independent) myopic rule.⁵² We have discussed that when Theorems 1 or 2 apply, the myopic division vector $m(t)$ is t -optimal at every period t . In this appendix, we argue that generically, division vectors $m(t)$ at late periods are t -optimal. This result complements our Theorem 3, and suggests that a DM who naively follows the myopic rule all the way cannot do very poorly.

To avoid repetition, here we only sketch the core argument. The main new step is to show that the division vectors $m(t)$ under the myopic rule grow to infinity in each coordinate; that is, a myopic DM would *not get stuck* observing a subset of signals. Once this is shown, we can repeat the (rest of the) proof of Lemma 4 and deduce that $m_i(t) - \lambda_i \cdot t$ remains bounded. And with these asymptotic characterizations, we can reproduce the proof of Theorem 3 (now for the myopic strategy instead of the optimal strategy) without trouble.⁵³

To see myopic signal choices never get stuck, we establish the following lemma.

⁵²That is, $m(t) = (m_1(t), \dots, m_K(t))$ where $m_i(t)$ is the number of times signal i has been observed under myopic information acquisition prior to and including period t .

⁵³These latter steps are actually simpler to carry out for the myopic strategy. This is because in constructing a deviation from the myopic strategy, we only need to look for a lower posterior variance at a single period. So we no longer need to make use of switch deviations.

Lemma 9. Fix an arbitrary division vector $q \in \mathbb{R}_+^K$ (need not be integral). The partial derivatives of f at q are all zero if and only if $q_1 = \dots = q_K = \infty$.

Intuitively, this holds because for normal linear signals, the posterior variance is *globally convex*. So if each signal has zero marginal value relative to the division q , then q must be a global minimizer of posterior variance.

We conclude by mentioning that similar result (i.e. myopic information acquisition does not get stuck) would not in general be true for other signal structures. The following is a counterexample with normal but non-linear signals.

Example 5. Consider three states $\theta_1, \theta_2, \theta_3$ drawn independently. The DM has access to these three signals:

$$\begin{aligned} X_1 &= \theta_1 + \text{sign}(\theta_2) + \epsilon_1 \\ X_2 &= \text{sign}(\theta_2\theta_3) + \epsilon_2 \\ X_3 &= \theta_3 + \epsilon_3 \end{aligned}$$

where $\epsilon_1, \epsilon_2, \epsilon_3$ are Gaussian noise terms. We focus on the prediction problem, in which (at a random time) the DM makes a prediction about θ_1 and receives negative of the squared prediction error.

Note that prior to the first observation of X_2 , signal X_3 is completely uninformative about the payoff-relevant state θ_1 (even when combined with previous observations of X_1). Similarly, signal X_2 is individually uninformative about θ_2 ,⁵⁴ and thus about θ_1 . These imply that the DM's uncertainty about θ_1 is not reduced upon the first observation of either X_2 or X_3 . Hence, the myopic rule in this example is to always observe X_1 , contrary to Lemma 9.

Thus, if the DM acquires information myopically, he will never completely learn the value of θ_1 . By contrast, if the DM is sufficiently patient, then his optimal strategy will observe each signal infinitely often and identify the value of θ_1 in the long run. Thus, in this example the myopic signal path does not eventually agree with the optimal path.

⁵⁴This is because the sign of $\theta_2\theta_3$ does not contain any new information about θ_2 when θ_3 is equally likely to be positive or negative.