OPTIMAL BAYESIAN OPTIMIZATION WITH STEREOGRAPHIC FUNCTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

In Bayesian optimization, the reward-maximizing policy solves a Bellman equation whose state is an entire posterior distribution, and that recursion is intractable; practice substitutes acquisition functions that look a single query ahead. We give a setting in which the optimal policy can instead be written down. The unknown function is linear in the stereographic projection of its input into $\mathbb{R}^n$, which makes it a smooth function of the $n-1$ input coordinates with one peak and one valley; evaluations are priced, and probing one known direction costs more than probing any other; and the learner is scored on a single final query, against a point drawn at random from those it could afford. The terminal reward turns out to be a quadratic form in the posterior mean whose weight matrix annihilates the expensive direction, so information about that direction has no value. The Bellman equation then admits a closed-form solution, and the optimal query schedule is non-adaptive: it can be computed in full before any data is seen. Choosing that schedule reduces to water-filling over the eigenvalues of the prior, which we solve in $O(n^3)$ time, and the resulting design puts no weight on the expensive direction, so every query it issues costs only the base price. Experiments check the closed-form value function against Monte Carlo estimates and measure how much exact optimality is worth relative to structured baselines.

1 INTRODUCTION

Bayesian optimization is the standard approach to optimizing a function that is expensive to evaluate and available only through noisy queries (Jones et al., 1998; Garnett, 2023). One maintains a probability distribution over the unknown function and chooses each new evaluation point by maximizing an *acquisition function*, a cheap score computed from that distribution. An acquisition function is a stand-in for something harder to compute. Choosing queries so as to maximize the expected reward at the end is a sequential decision problem, and its exact solution is characterized by a *Bellman equation*: a recursion whose unknown is the *value function*, the best expected reward still achievable from a given state. Here the state is the whole posterior distribution over functions, and for the nonparametric priors used in practice the recursion cannot be solved. Expected improvement, the knowledge gradient and entropy search sidestep this by being *myopic*, scoring the gain from the very next query rather than from the whole remaining sequence. The gap this leaves is well recognized, and the usual response is to approximate it away with rollout, multi-step lookahead or approximate dynamic programming (Lam et al., 2016; Jiang et al., 2020), which approximates an object nobody can write down. We go the other way, restricting the problem until the optimal policy can be computed and then looking at what the solution does. Three restrictions are involved. *A parametric function class.* The unknown function is linear in a fixed feature map, $f(\theta) = \phi(\theta)^\top z$, where the features $\phi(\theta)$ are the *stereographic projection* of the input $\theta \in \mathbb{R}^{n-1}$ onto the unit sphere in $\mathbb{R}^n$ and $z \in \mathbb{R}^n$ is the unknown parameter vector. Inference is then ordinary Bayesian linear regression, but the induced functions on $\mathbb{R}^{n-1}$ are not trivial: each has exactly one peak and one valley, and learning f means learning which way the peak points and how tall it is (Figure 1). *Priced queries.* Every query costs a base price κ, plus a surcharge that depends on where we look: probing along a known direction c costs up to $\kappa + \|c\|^2$, while probing anywhere orthogonal to c costs only κ. We call c the *expensive direction*. A budget B pays for exploration and a separate budget $\hat{B}$ pays for one final query. *A terminal score.* Only the final query is rewarded, and we score it against a random baseline: how much better is the point we choose than a point drawn at random from those

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

we could just barely afford? Within this setting the optimal policy can be computed exactly. We do not offer it as a competitive method for practical Bayesian optimization; it is a solved instance of the problem that Bayesian optimization approximates, and what it provides is diagnostic: it identifies the structural features that give adaptivity its value, and since our results fail as soon as any one of them is dropped, it also says where the boundary lies (Section 6). It also gives a testbed on which a non-myopic method can be scored against the true optimum instead of against another heuristic. Our contributions are as follows:

- **The reward is a quadratic form in the posterior mean** (Section 3.1): the value of the final query is $\mu^\top P \mu$, with μ the posterior mean of z and P a fixed matrix that annihilates the expensive direction. What the final query gains by leaning into c is exactly what the random baseline gains for free, so knowing z along c is worth nothing.
- **Optimal querying is non-adaptive** (Section 4.1): the Bellman equation has a closed-form solution in which the value splits into a term depending on the posterior mean and one depending only on accumulated uncertainty. That uncertainty shrinks along a path fixed by the queries alone, so the entire schedule can be planned before any data arrives.
- **The design problem is a matrix optimization** (Section 4.2): the sequence matters only through the sum of the queries' outer products, and we characterize exactly which such matrices a whole number of queries can realize.
- **The solution is water-filling on the prior spectrum** (Section 4.4): the optimal design measures along the prior's own principal directions and equalizes posterior certainty across them. It spends nothing on the expensive direction, so every query costs only the base price, and is computable in $O(n^3)$ time (Algorithm 1).

2 Related Work

Bayesian optimization. One places a prior over the unknown function, updates it as evaluations arrive, and picks each evaluation by maximizing an acquisition function (Kushner, 1964; Jones et al., 1998; Garnett, 2023). Expected improvement, the knowledge gradient and entropy search all look one query ahead (Frazier et al., 2009; Hennig & Schuler, 2012; Wang & Jegelka, 2017), and non-myopic variants close the resulting gap approximately (Lam et al., 2016; Jiang et al., 2020). These approximate the optimal policy; we restrict the function class until it is computable.

Exactly solvable sequential problems. Such problems are rare and highly structured. The classical example is the Gittins index for discounted bandits with independent arms (Gittins, 1979; Whittle, 1980), which breaks down under the finite horizons, correlated arms and query costs present here. Ours comes from elsewhere: the terminal reward is quadratic in the posterior mean while the posterior uncertainty follows a path fixed by the queries alone, so the value function separates, much as under certainty equivalence in stochastic control (Bar-Shalom & Tse, 1974).

Bandits and pure exploration. Optimality in the bandit literature is asymptotic (Lai & Robbins, 1985; Lattimore & Szepesvári, 2019). Closest to us is *pure exploration* (Bubeck et al., 2009; Garivier & Kaufmann, 2016), where the learner explores freely and is graded once at the end, and where Soare et al. (2014) and Fiez et al. (2019) observe that in a linear model much of the measurement allocation can be computed offline. In our setting adaptivity has no value at all, but we trade asymptotic for exact optimality in one parametric family.

Experimental design and query costs. After the reduction of Section 4.2 our problem is L-optimal experimental design (Fedorov, 1972; Pukelsheim, 2006; Chaloner & Verdinelli, 1995), a theory built for *approximate* designs because fractional allocations make it convex. Our budget instead buys a whole number of queries, and Lemma 3 characterizes those *exact* designs directly, so no rounding is needed. Domain-varying query costs are usually handled heuristically (Snoek et al., 2012; Kandasamy et al., 2017; Astudillo et al., 2021); here the cost is known, and its interaction with the reward geometry produces our main structural conclusion: the expensive directions are exactly the ones whose values do not affect the reward. Appendix A treats these connections in detail.

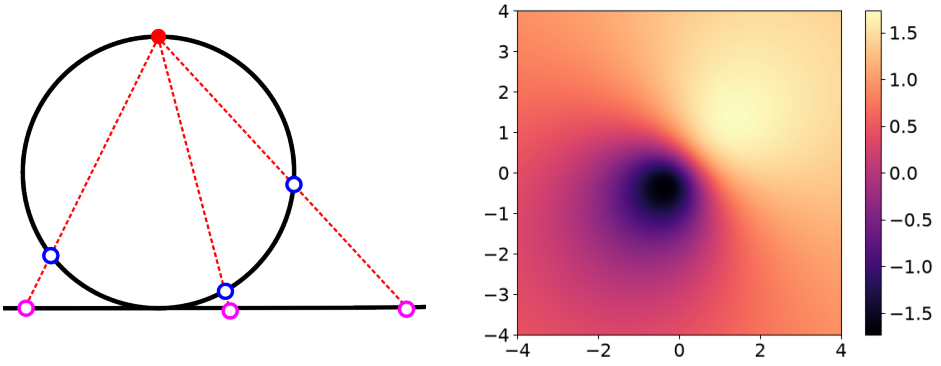

Figure 1: (left) Stereographic projection for $n = 2$: the line $\mathbb{R}^1$ maps onto the circle $\mathbb{S}^1$ along rays meeting at the north pole. (right) A stereographic function with $n = 3$ and $z = (1, 1, 1)$: a smooth function of $\theta \in \mathbb{R}^2$ with one peak, at $\theta = \frac{1}{2}[1 + \sqrt{3},\ 1 + \sqrt{3}]$, and one valley, at $\theta = \frac{1}{2}[1 - \sqrt{3},\ 1 - \sqrt{3}]$.

3 Problem Setup

Notation. $\|\cdot\|$ is the Euclidean norm and $\mathbb{S}^{n-1} = \{x \in \mathbb{R}^n : \|x\| = 1\}$ the unit sphere in $\mathbb{R}^n$; e_i is the i-th standard basis vector; $x_{2:n}$ the last $n - 1$ coordinates of x; tr, rank and diag are trace, rank and the diagonal matrix built from a vector; I is the identity. Superscripts in parentheses index time, subscripts index coordinates. For symmetric matrices, $X \succeq 0$ means X is positive semi-definite, $X \succ 0$ positive definite, and $X \succeq Y$ means $X - Y \succeq 0$; this is the *Loewner order*. Two symmetric matrices *commute* exactly when they share an orthonormal eigenbasis, and we use the two descriptions interchangeably.

3.1 Stereographic functions, costs, and the terminal score

The function class. We wish to learn an unknown scalar function of an input $\theta \in \mathbb{R}^{n-1}$ from noisy evaluations at points of our choosing. We assume it is linear in the *stereographic projection* $\phi(\theta) \in \mathbb{R}^n$ of the input,

$$f(\theta; z) = \phi(\theta)^\top z, \qquad \phi(\theta) = \left[\frac{\|\theta\|^2 - 1}{\|\theta\|^2 + 1},\ \frac{2\theta}{\|\theta\|^2 + 1}\right], \qquad \phi^{-1}(m) = \frac{m_{2:n}}{1 - m_1}, \tag{1}$$

where $z \in \mathbb{R}^n$ collects the unknown parameters. Stereographic projection maps $\mathbb{R}^{n-1}$ bijectively onto the unit sphere with the north pole $e_1 = [1, 0, \ldots, 0]$ removed: one checks directly that $\|\phi(\theta)\| = 1$ for every θ and that $\phi(\theta) \to e_1$ as $\|\theta\| \to \infty$, which is why the north pole is the one point never reached. We call functions of the form equation 1 *stereographic functions*; written out they are ratios of quadratics, $f(\theta; z) = \big[(\|\theta\|^2 - 1)z_1 + 2\theta^\top z_{2:n}\big]/(\|\theta\|^2 + 1)$. We use this class for two reasons. Inference is ordinary Bayesian linear regression in the features $\phi(\theta)$, so posteriors are available in closed form. And the geometry of optimization is simple without being trivial: on the sphere $f(\cdot\,; z)$ is the linear map $m \mapsto m^\top z$, whose only critical points for $z \neq 0$ are a maximum at $z/\|z\|$ and a minimum at $-z/\|z\|$. Each f is therefore a smooth function on $\mathbb{R}^{n-1}$ with a single peak and a single valley (Figure 1), and learning f amounts to learning which way the peak points and how tall it is. Since everything of interest happens on the sphere, we work throughout with the projected input $m := \phi(\theta) \in \mathbb{S}^{n-1} \setminus \{e_1\}$ and translate back at the end using ϕ^{-1}.

Query costs. We select queries $m^{(1)}, m^{(2)}, \ldots$, and the response to query t is a noisy observation $y^{(t)}$ of $f(\theta^{(t)}; z)$. Each query is priced by

$$k(m) \ := \ \kappa + (m^\top c)^2, \tag{2}$$

where $\kappa > 0$ is a base price paid for every query and the known vector $c \in \mathbb{R}^n$, the *expensive direction*, makes some directions costlier than others: querying near $\pm c/\|c\|$ costs up to $\kappa + \|c\|^2$,

whereas querying on the "free equator" $\{m : m^\top c = 0\}$ costs only κ, and since the surcharge depends on $(m^\top c)^2$ the points m and $-m$ cost the same. Unlike z, the quantities κ and c are known in advance. We are given a budget B for exploratory queries and a separate budget $\hat{B}$ for one final *optimizing* query, at which we try to attain an extremum of f and collect our reward; $B^{(t)}$ is the budget remaining after t exploratory queries, so $B^{(0)} = B$.

The terminal score. We score the optimizing query against a random baseline: how much better is the point we choose than a point picked at random from those we can just barely afford? Write $\beta := \sqrt{\hat{B} - \kappa}$ and $\hat{c} := c/\|c\|$, so that affordability, $k(m) \le \hat{B}$, is exactly the condition $|m^\top c| \le \beta$. Let the baseline point m' be drawn uniformly from the slice of the sphere cut out by the budget boundary,

$$\mathcal{B} := \left\{ m \in \mathbb{S}^{n-1} : m^\top c = \beta \right\}, \qquad \text{so that} \qquad \mathbb{E}[m'] = \frac{\beta}{\|c\|^2} c, \tag{3}$$

the mean following by symmetry, since $\mathcal{B}$ is itself a sphere centred at $\frac{\beta}{\|c\|^2}c$. Writing $\mu := \mathbb{E}[z \mid y^{(1)}, y^{(2)}, \dots]$ for the posterior mean of the parameter vector given the data so far, the objective is to choose the optimizing query m so as to

$$\max_{m \in \mathbb{S}^{n-1},\, k(m) \le \hat{B}} \; \min_{s \in \{-1,1\}} \left(\mathbb{E}_{z,m'} \left[m^\top z - s\, m'^\top z \;\middle|\; y^{(1)}, y^{(2)}, \dots \right] \right)^2, \tag{4}$$

subject to the exploration budget $\sum_t k(m^{(t)}) \le B$. The outer square rewards us for finding *either* extremum, peak or valley, since a large negative gap scores as well as a large positive one; the adversarial sign s then flips the baseline so that we are always measured against whichever extremum we are chasing, which also makes it harmless to pick the positive one of the two slices $m'^\top c = \pm\beta$ forming the budget boundary. The baseline must be a law on the projected point m' rather than on an input θ', because the inputs mapping into $\mathcal{B}$ have infinite volume.

Assumption 1. Throughout we assume

(A1) $n \ge 3$, $\kappa > 0$, $B \ge \kappa$, and the prior precision (defined in Section 3.2) satisfies $\Lambda^{(0)} \succ 0$;

(A2) $\beta \le \|c\|$, so that the slice $\mathcal{B}$ is nonempty and

$$p \;:=\; 1 - \frac{\hat{B} - \kappa}{\|c\|^2} \;=\; 1 - \frac{\beta^2}{\|c\|^2} \;\in\; [0, 1]; \tag{5}$$

we assume moreover $p > 0$, since $p = 0$ makes the objective identically zero;

(A3) the final-query budget binds and the informative branch dominates, i.e.

$$\beta \;\le\; \frac{|c^\top \mu|}{\|\mu\|} \qquad \text{and} \qquad (\hat{B} - \kappa) + \frac{(c^\top \mu)^2}{\|\mu\|^2} \;\le\; \|c\|^2, \tag{6}$$

at every posterior mean μ arising along the query path.

(A1) is bookkeeping. (A2) says the optimizing budget does not reach the whole sphere, whose most expensive point costs $\kappa + \|c\|^2$; here p is the squared radius of the affordable slice $\mathcal{B}$, and $p = 1$ confines the optimizing query to the free equator. (A3) is different in kind: it constrains the posterior mean rather than the problem data. Its first part says the optimizing query genuinely wants to move in the expensive direction, so the budget binds; its second rules out a degenerate regime (Remark 1). Splitting the posterior mean along the expensive direction and the rest, $\mu = \mu_\parallel + \mu_\perp$ with $\mu_\parallel := \hat{c}\,\hat{c}^\top \mu$ and $\mu_\perp := (I - \hat{c}\,\hat{c}^\top)\mu$, the terminal score has a closed form; we prove the following in Appendix D.1.

Proposition 1 (Terminal objective)**.** *Under Assumption 1, the value of equation 4 is*

$$\max_{\substack{m \in \mathbb{S}^{n-1} \\ |m^\top c| \le \beta}} \; \min_{s \in \{-1,1\}} \left(\left(m - s \tfrac{\beta}{\|c\|^2} c \right)^\top \mu \right)^2 \;=\; p\, \|\mu_\perp\|^2 \;=\; \mu^\top P \mu, \qquad P \;:=\; p\left(I - \hat{c}\,\hat{c}^\top\right) \;\succeq\; 0, \tag{7}$$

and the maximum is attained at

$$m^\star = \alpha\, \hat{c} + \sqrt{1 - \alpha^2}\; \frac{\mu_\perp}{\|\mu_\perp\|}, \qquad \alpha = \operatorname{sign}(c^\top \mu)\, \frac{\beta}{\|c\|}. \tag{8}$$

The matrix P is p times the orthogonal projection onto the directions perpendicular to c, so only the component of the posterior mean orthogonal to the expensive direction contributes to the reward, and even that is discounted by p. The reason is visible in equation 8: what the optimizing query earns by leaning into the expensive direction, namely $\beta|c^\top \mu|/\|c\|^2$, is exactly what the random baseline earns for free, and the two cancel. What we learn about z along c is therefore worthless, and the rest of the analysis is a consequence of this.

3.2 Bayesian updates

We take a Gaussian prior $\mathcal{N}(\mu^{(0)}, \Lambda^{(0)})$ over z and Gaussian noise of variance σ^2, so every posterior is Gaussian; $\mathcal{N}(\mu, \Lambda)$ denotes the density proportional to $e^{-\frac{1}{2}(z-\mu)^\top \Lambda (z-\mu)}$, parameterized by its mean and its *inverse* covariance Λ, the *precision*. We work with the precision because independent observations simply add it: one query contributes the rank-one term $mm^\top/\sigma^2$, information aligned with the direction probed, and moves the mean along $\Lambda^{-1}m$ in proportion to the surprise $y - m^\top \mu$, which before the data is seen has mean 0 and variance $\sigma^2 + m^\top \Lambda^{-1} m$. Writing the update with a standard normal ε and iterating, for any $t \geq 0$,

$$\mu^{(t)} = \mu^{(0)} + \sum_{j=1}^{t} \frac{\Lambda^{(j-1)-1} m^{(j)}}{\sqrt{\sigma^2 + m^{(j)\top} \Lambda^{(j-1)-1} m^{(j)}}}\, \varepsilon^{(j)}, \qquad \Lambda^{(t)} = \Lambda^{(0)} + \sum_{j=1}^{t} \frac{m^{(j)} m^{(j)\top}}{\sigma^2}, \tag{9}$$

with $\varepsilon^{(1)}, \ldots, \varepsilon^{(t)}$ independent standard normals. The posterior mean is thus a random walk whose step directions and sizes are known in advance, and in particular a *martingale*: its expected future value equals its current value. The second identity is what the rest of the paper leans on: once the queries are fixed, the precision path is deterministic. Information accumulates according to where we look, not according to what we see.

4 The Optimal Design

4.1 The dynamic program

Because every query costs at least $\kappa > 0$, at most $\lfloor B/\kappa \rfloor$ queries can ever be made and the problem has a finite horizon. The state is the triple (posterior mean, posterior precision, remaining budget); let $V(\mu, \Lambda, B)$ denote the *value function*, the largest expected terminal reward achievable from that state, and let $\mathcal{D}(B)$ be the set of finite, possibly empty, affordable query sequences, i.e. those $(m^{(1)}, m^{(2)}, \ldots)$ in $\mathbb{S}^{n-1}$ with $\sum_t k(m^{(t)}) \leq B$. The value function satisfies the Bellman equation

$$V(\mu^{(t)}, \Lambda^{(t)}, B^{(t)}) = \max \Bigg\{ \mu^{(t)\top} P \mu^{(t)}, \max_{m \in \mathbb{S}^{n-1} : k(m) \leq B^{(t)}} \mathbb{E}_\varepsilon \Bigg[V\Bigg(\mu^{(t)} + \frac{\Lambda^{(t)-1} m}{\sqrt{\sigma^2 + m^\top \Lambda^{(t)-1} m}} \varepsilon, \Lambda^{(t)} + \frac{mm^\top}{\sigma^2},\ B^{(t)} - k(m) \Bigg) \Bigg] \Bigg\}. \tag{10}$$

The first branch is the option to stop exploring and cash out with the optimizing query, worth $\mu^\top P \mu$ by Proposition 1; the second is the option to spend $k(m)$ on one more exploratory query and continue. Including the stopping branch pins the recursion down: with a finite horizon and a specified stopping value, backward induction shows that equation 10 has a unique solution, which we exhibit in closed form below and prove in Appendix D.2.

Proposition 2. *Let $P \succeq 0$ and define*

$$J(\Lambda, B) \ := \min_{(m^{(t)}) \in \mathcal{D}(B)} \operatorname{tr}\left(P \Big(\Lambda + \sum_t \tfrac{m^{(t)} m^{(t)\top}}{\sigma^2} \Big)^{-1} \right), \tag{11}$$

with the convention $\min \emptyset = +\infty$. Then $V(\mu, \Lambda, B) = \mu^\top P \mu + \operatorname{tr}(P\Lambda^{-1}) - J(\Lambda, B)$.

The quantity $\operatorname{tr}(P\Lambda^{-1})$ is the posterior variance of what we care about, aggregated by the weighting P; call it the *residual uncertainty*. Then $J(\Lambda, B)$ is the smallest residual uncertainty the budget B can buy, so the value reads

$$V(\mu, \Lambda, B) \ = \ \underbrace{\mu^\top P \mu}_{\text{what we could collect now}} \ + \ \underbrace{\operatorname{tr}(P\Lambda^{-1}) - J(\Lambda, B)}_{\text{uncertainty we can still buy away}}.$$

The second term is nonnegative and does not involve μ. The reason is that the posterior mean is a martingale: future data does not change its expected value, but it does spread the mean out, and since the reward is convex in the mean that spread is worth something, namely exactly the uncertainty that remains to be resolved. When the budget is exhausted only the empty sequence survives, $J = \mathrm{tr}(P\Lambda^{-1})$, and the value reduces to $\mu^\top P\mu$, as Proposition 1 requires. The consequence that matters is that the optimization in equation 11 does not involve μ: it depends only on Λ and B. Combined with the fact that precisions evolve deterministically once the queries are chosen equation 9, this means the whole query sequence can be computed at time 0, so adaptivity confers no advantage. Nothing is gained by looking at the responses before deciding where to look next, and it suffices to solve the static problem

$$\min_{(m^{(t)})\in\mathcal{D}(B)} \mathrm{tr}\left(P\Big(\Lambda^{(0)} + \sum_t \tfrac{m^{(t)}m^{(t)\top}}{\sigma^2}\Big)^{-1}\right). \tag{12}$$

4.2 From queries to an information matrix

The objective in equation 12 depends on the query sequence only through the matrix $M := \sum_t m^{(t)}m^{(t)\top}$, since by equation 9 the design contributes exactly M/σ^2 to the precision. Following the experimental-design literature we call M the *information matrix* of the design. The total cost also depends only on M: because each $m^{(t)}$ is a unit vector, $\mathrm{tr}\, M = \sum_t \|m^{(t)}\|^2$ counts the queries, while $\sum_t (c^\top m^{(t)})^2 = c^\top Mc$, so

$$\sum_t \Big(\kappa + (c^\top m^{(t)})^2\Big) \;=\; \kappa\,\mathrm{tr}(M) + c^\top Mc \tag{13}$$

holds identically, for every way of decomposing M into unit outer products. We may therefore forget the individual queries, optimize over M, and reconstruct a query sequence at the end. To do so we need to know which matrices are achievable.

Lemma 3 (Realizability)**.** *Let M be a symmetric $n \times n$ matrix and T a nonnegative integer. There exist unit vectors $m^{(1)}, \dots, m^{(T)} \in \mathbb{S}^{n-1}$ with $M = \sum_{t=1}^T m^{(t)}m^{(t)\top}$ if and only if $M \succeq 0$, $\mathrm{tr}\, M = T$ and $\mathrm{rank}\, M \le T$. Moreover the vectors may be chosen so that none equals the excluded north pole e_1, and hence so that each $m^{(t)} = \phi(\theta^{(t)})$ for some $\theta^{(t)} \in \mathbb{R}^{n-1}$.*

The three conditions are clearly necessary: a sum of outer products is positive semi-definite, its trace is the number of summands since each summand has trace 1, and its rank is at most that number. Less obvious is that these conditions are also sufficient, and in particular that no constraint on the largest eigenvalue is needed. The proof, in Appendix D.3, rests on the Schur-Horn theorem (Schur, 1923; Horn, 1954), restated as Lemma 7: it fixes exactly which diagonals a symmetric matrix of given spectrum can have, and the flat diagonal we need, with all query norms equal to one, is always among them. Combining Lemma 3 with equation 13, problem equation 12 is equivalent to the matrix problem

$$\begin{aligned} &\min_{M\in\mathcal{M}(B)} G(M), \qquad G(M) := \mathrm{tr}\left(P\Big(\Lambda^{(0)} + \tfrac{M}{\sigma^2}\Big)^{-1}\right), \\ &\mathcal{M}(B) := \big\{M \succeq 0 : \kappa\,\mathrm{tr}\, M + c^\top Mc \le B,\ \mathrm{tr}\, M \in \mathbb{Z},\ \mathrm{rank}\, M \le \mathrm{tr}\, M\big\}. \end{aligned} \tag{14}$$

The whole sequential problem is now encoded here: the budget is a linear constraint on M, and the requirement of a whole number of unit-norm queries is the integrality and rank condition.

4.3 Simultaneous diagonalization

Problem equation 14 is still an optimization over matrices, and the rank constraint makes it nonconvex. We now assume that $\Lambda^{(0)}$ and P commute. Since $P = p(I - \hat{c}\hat{c}^\top)$ has only the two eigenspaces $\mathrm{span}(\hat{c})$ and its orthogonal complement, this says exactly that c is an eigenvector of the prior precision: the expensive direction is one of the prior's principal directions. The assumption is substantive. It describes the regime in which a single set of axes captures both the cost and the prior geometry, and under it the problem collapses to an optimization over n numbers.

Proposition 4. *Suppose $\Lambda^{(0)}$ and P commute. Then the minimum in equation 14 is attained, and there is a minimizer $M^\star$ with $M^\star c = 0$ and $M^\star \Lambda^{(0)} = \Lambda^{(0)} M^\star$. Such an $M^\star$ commutes with P as well, so $M^\star$, $\Lambda^{(0)}$ and P are simultaneously diagonalizable by an orthogonal matrix Q one of whose columns is $\hat{c}$.*

That $M^\star c = 0$ says the optimal design puts no weight at all on the expensive direction: every query is orthogonal to c and so costs only the base price κ. This is the design-side consequence of Proposition 1. We prove the proposition in Appendix D.5; the commutation claim rests on Lemma 8 (Appendix D.4): once we fix how much information a design collects and how it is spread across its own axes, the best orientation of those axes relative to the prior is the aligned one.

4.4 Water-filling on the prior spectrum

Let Q be an orthogonal matrix whose columns are the shared eigenvectors, indexed so that column k equals $\hat{c}$; we call k the *expensive coordinate*. Writing $\Lambda^{(0)} = Q\,\mathrm{diag}(\lambda_1, \ldots, \lambda_n)Q^\top$ and $M = Q\,\mathrm{diag}(d_1, \ldots, d_n)Q^\top$, the number $\lambda_i > 0$ is how much the prior already knows about direction i and $d_i \geq 0$ is how much information the design buys there. Then $c^\top M c = \|c\|^2 d_k$, while P has eigenvalue 0 in coordinate k and p in every other coordinate, so equation 14 becomes a scalar optimization,

$$\min_{d \in \mathcal{D}_B} \sum_{i \neq k} \frac{p}{\lambda_i + \frac{d_i}{\sigma^2}}, \qquad \mathcal{D}_B := \Big\{ d \in \mathbb{R}^n_{\geq 0} : \|c\|^2 d_k + \kappa \textstyle\sum_i d_i \leq B,\ \big|\{i : d_i > 0\}\big| \leq \sum_i d_i \in \mathbb{Z} \Big\}. \tag{15}$$

Each term is the posterior variance along one prior eigendirection and d_i/σ^2 is the precision the design adds there; the expensive coordinate has dropped out of the objective entirely, as it must. Appendix D.6 proves the solution below, together with Corollary 6.

Proposition 5. *Let $T := \lfloor B/\kappa \rfloor \geq 1$ be the largest number of affordable queries. For an integer $1 \leq N \leq \min(n-1, T)$, let $\mathcal{A}_N \subseteq \{1, \ldots, n\} \setminus \{k\}$ collect the N coordinates with the smallest λ_i, ties broken arbitrarily, and call N* admissible *if*

$$\max_{i \in \mathcal{A}_N} \lambda_i \ \leq \ \frac{1}{N} \left(\frac{T}{\sigma^2} + \sum_{j \in \mathcal{A}_N} \lambda_j \right). \tag{16}$$

Note $N = 1$ is always admissible, so admissible values exist. Let

$$N^\star \in \operatorname*{arg\,min}_{N \text{ admissible}} \left[\frac{pN^2}{\frac{T}{\sigma^2} + \sum_{j \in \mathcal{A}_N} \lambda_j} + \sum_{i \notin \mathcal{A}_N \cup \{k\}} \frac{p}{\lambda_i} \right]. \tag{17}$$

Then equation 15 is solved by

$$d_i = \frac{T}{N^\star} + \frac{\sigma^2}{N^\star} \sum_{j \in \mathcal{A}_{N^\star}} \lambda_j - \sigma^2 \lambda_i \quad \text{for } i \in \mathcal{A}_{N^\star}, \qquad d_i = 0 \quad \text{otherwise}. \tag{18}$$

Interpreting the solution. This has the classical water-filling form. Picture one vessel per coordinate $i \neq k$, with the floor of vessel i at height λ_i: the higher the floor, the more the prior already knows about that direction. We pour a fixed volume of water, T/σ^2, into a chosen subset of vessels; the water in vessel i is the added precision d_i/σ^2, the resulting level is the posterior precision $\lambda_i + d_i/\sigma^2$, and the reciprocal of that level is what we pay. Equation equation 18 equalizes the level across the chosen vessels, and equation 16 asks that it clear every floor in that set, so that no vessel holds a negative amount of water. Two features differ from textbook water-filling. The objective sums reciprocals rather than logarithms, and the rank constraint caps the number of filled vessels at the number of queries, so we cannot simply fill every vessel below the water line; sweeping over N in equation 17 handles the cap.

Corollary 6 (Ordering)**.** *The optimal design spends its budget on the directions the prior knows least about: among the coordinates $i \neq k$, the largest d_i is paired with the smallest λ_i, the second largest with the second smallest, and so on.*

4.5 RECONSTRUCTING THE QUERY SEQUENCE

We now have the optimal information matrix $M^\star = Q\,\mathrm{diag}(d)Q^\top$ but not the queries that produce it. Stack them as the rows of $A \in \mathbb{R}^{T\times n}$: then $\sum_t m^{(t)}m^{(t)\top} = M^\star$ says $A^\top A = M^\star$, for which any matrix square root will do, while unit-norm queries say that every row of A has norm 1. In Appendix C the row norms of a square root are repaired by rotations that leave $A^\top A$ untouched while trading squared length between two rows, and the pipeline is assembled into Algorithm 1, which runs in $O(n^3 + TN^\star)$ time, dominated by the eigendecomposition of $\Lambda^{(0)}$. Since $M^\star c = 0$, every query it returns satisfies $c^\top m^{(t)} = 0$ and costs exactly κ, for a total outlay of $\kappa T \le B$. Neither the observations nor the prior mean appears anywhere in the algorithm, which is the practical content of the non-adaptivity result. Once the responses arrive, the posterior is assembled in one shot,

$$\Lambda^{(T)} = Q\big(\,\mathrm{diag}(\lambda) + \mathrm{diag}(d)/\sigma^2\big)Q^\top, \qquad \mu^{(T)} = \Lambda^{(T)-1}\Big(\Lambda^{(0)}\mu^{(0)} + \sigma^{-2}\textstyle\sum_t y^{(t)}m^{(t)}\Big), \quad (19)$$

and the final optimizing query is $m^\star$ from equation 8, evaluated at $\mu = \mu^{(T)}$: lean as far into the expensive direction as the budget allows, and spend the rest of the sphere's reach pointing at the posterior mean.

5 EXPERIMENTS

Algorithm 1 is optimal by construction, so no experiment can strengthen that claim. We run them for two other reasons: Proposition 2 makes a sharp, falsifiable prediction that checks the implementation, and we want to know what optimality is worth against rules that each discard one part of the structure the theory exploits (details in Appendix E).

Instances and scoring. At each dimension $n = 3, \ldots, 16$ we draw 2000 problems: a Haar-random basis Q, a prior precision $\Lambda^{(0)} = Q\,\mathrm{diag}(\lambda)Q^\top$ whose eigenvalues span two decades, an expensive direction c equal to a column of Q so that $\Lambda^{(0)}$ and P commute, and z drawn from that prior. We set $\kappa = 1$, $\|c\| = 2$, $p = 0.91$ and $B = 2n\kappa$, buying $T = 2n$ queries, and put the prior mean along the expensive direction, $\mu^{(0)} = \hat{c}$: the peak is believed to lie toward the direction we cannot afford to probe. Since $P\hat{c} = 0$ the objective is then exactly zero before any query, and the score, which is the exact value of equation 4 normalised as in Appendix E, measures only what exploration bought.

Baselines. Each keeps part of the structure and discards the rest. *Cost-blind A-optimality* replaces P by I, so every direction looks equally valuable, and ignores the price when planning, paying it at execution and exhausting its budget early; this isolates the reward geometry. *Uncertainty sampling* queries the direction of largest posterior variance, blind to both reward weighting and price; *random* draws affordable queries uniformly.

Results. Our design scores highest at every dimension (Figure 2); since all methods face the same instances the comparison is paired, and ours wins on all 28 000. The shortfall against it falls from 27.7%, 22.2% and 34.7% at $n = 3$ to 3.1%, 2.7% and 12.8% at $n = 16$, in the order just listed. The two structured baselines fail as Proposition 1 predicts: they put 0.5 to 1.6 units of squared weight on $\hat{c}$, the one direction whose values do not affect the reward, costing three to seven of their $2n$ queries; ours puts none. That waste is $\Theta(1)$ against a budget growing like $2n$, so the penalty should decay as $\Theta(1/n)$, and n times the shortfall does settle near 0.5 out to $n = 128$. Random sampling instead flattens near 8%: ignoring Corollary 6, it spreads error over all directions, and dimension cannot dilute it.

Checking the value function. Proposition 2 predicts the expected objective of any design in closed form, as $\mu^{(0)\top}P\mu^{(0)} + \mathrm{tr}(P\Lambda^{(0)-1}) - \mathrm{tr}(P\Lambda^{(T)-1})$, computable before any data is seen; comparing it with the Monte Carlo mean of $\mu^{(T)\top}P\mu^{(T)}$ is a paired test with no free parameters. The largest discrepancy over all four methods and fourteen dimensions is 2.5 standard errors, what noise over 56 comparisons would give. This tests the value function itself, not only our implementation.

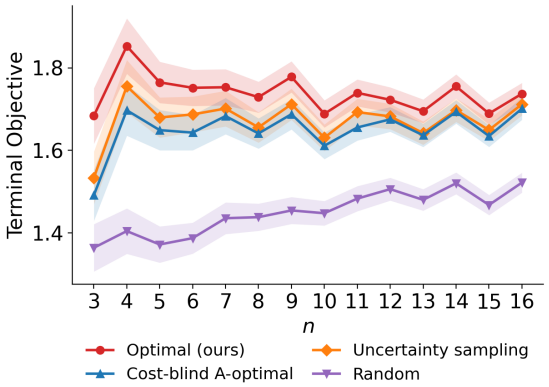

Figure 2: Mean terminal objective against n, over 2000 paired instances per n; bands are ± 1 standard error. The structured baselines' shortfall closes as $\Theta(1/n)$; random sampling's does not.

6 Discussion

Why the problem is solvable. Three properties are responsible, and none can be dropped. The terminal reward is quadratic in the posterior mean, which with the martingale property makes any continuation worth the current reward plus the variance future data will inject. The posterior precision of a Gaussian linear model depends on the queries and not the responses, so the uncertainty half of the value function is deterministic. And a query's price depends only on the query. Remove any one and adaptivity regains its value: a non-quadratic reward leaves a value function that does not separate, a likelihood whose informativeness depends on the observed value makes the precision path random, and random costs (Astudillo et al., 2021) make the budget a stochastic state. Seen from the other side, the result marks the edge of that structure, much as certainty equivalence does for control (Bar-Shalom & Tse, 1974).

Non-adaptivity does not make the problem easy. The four methods of Section 5 are all non-adaptive in effect, and they differ by up to a third of the objective. The theory removes the need to react to data, but not the need to aim queries well, nor the need to recognize that information about c is worthless.

The commuting assumption. Proposition 4 requires c to be an eigenvector of the prior precision, so that one set of axes describes both the reward and the prior geometry. Without it we have no solution, and we would not expect one to exist: if the prior correlates c with the rewarded directions, then measuring along c reduces uncertainty about directions that are rewarded, so spending on c may be optimal even though $Pc = 0$. That case is a nonconvex exact design problem, and we suspect it has no closed form.

The role of the stereographic projection. After Proposition 1, very little depends on it. From Section 4.2 onwards the only properties of ϕ used are that the features are unit vectors, which makes $\operatorname{tr} M$ count the queries and the cost a function of M alone, and that f is linear in them, so any feature map onto $\mathbb{S}^{n-1}$ inherits the value function, the water-filling solution and the reconstruction unchanged. What ϕ contributes is at the two ends of the pipeline: an interpretable function class on $\mathbb{R}^{n-1}$, and a guarantee, since ϕ is a bijection onto the sphere minus a point, that every realizable design corresponds to actual inputs.

Limitations. (A3) constrains realized posterior means rather than problem data, so the quadratic terminal value describes a regime rather than every run; it holds on about 95% of ours. Also, the reward comes from one final query and querying costs are known, which can be a restrictive set of assumptions. Another important option question is whether similar techniques can be used when multiple queries are rewarded instead of only the final one.

AI USE STATEMENT

In this work, we used generative AI tools for assisting with proofs, reviewing literature, and designing numerical experiments. We have not used generative AI tools for developing the core theoretical ideas. Additionally, we used generative AI tools for paper writing and creating code. We have reviewed all AI-assisted work. All text and code outputted by generative AI was carefully checked by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *NeurIPS*, 24, 2011.

Rajeev Agrawal. Sample mean based index policies by o (log n) regret for the multi-armed bandit problem. *Advances in applied probability*, 27(4):1054–1078, 1995.

Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT*, pp. 39–1. JMLR Workshop and Conference Proceedings, 2012.

Raul Astudillo, Daniel Jiang, Maximilian Balandat, Eytan Bakshy, and Peter Frazier. Multi-step budgeted bayesian optimization with unknown evaluation costs. *NeurIPS*, 34:20197–20209, 2021.

Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *COLT*, pp. 13–p, 2010.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.

Yaakov Bar-Shalom and Edison Tse. Dual effect, certainty equivalence, and separation in stochastic control. *IEEE Transactions on Automatic Control*, 19(5):494–500, 1974.

Richard Bellman. A problem in the sequential design of experiments. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 16(3/4):221–229, 1956.

Donald A Berry and Bert Fristedt. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). *London: Chapman and Hall*, 5(71-87):7–7, 1985.

Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.

Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *ALT*, pp. 23–37. Springer, 2009.

Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.

Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback-leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, pp. 1516–1541, 2013.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical science*, pp. 273–304, 1995.

Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, 2008.

Valerii Fedorov. *Theory of Optimal Experiments*. Academic Press, 1972.

Tanner Fiez, Lalit Jain, Kevin Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. *NeurIPS*, 32, 2019.

Peter Frazier, Warren Powell, and Savas Dayanik. The knowledge-gradient policy for correlated normal beliefs. *INFORMS journal on Computing*, 21(4):599–613, 2009.

Peter I Frazier, Warren B Powell, and Savas Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.

Aurélien Garivier and Olivier Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *COLT*, pp. 359–376. JMLR Workshop and Conference Proceedings, 2011.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *COLT*, pp. 998–1027. PMLR, 2016.

Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.

David Ginsbourger and Rodolphe Le Riche. Towards gaussian process-based optimization with finite time horizon. In *mODa 9–Advances in Model-Oriented Design and Analysis: Proceedings of the 9th International Workshop in Model-Oriented Design and Analysis held in Bertinoro, Italy, June 14-18, 2010*, pp. 89–96. Springer, 2010.

John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 41(2):148–164, 1979.

Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(6), 2012.

José M Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. *NeurIPS*, 27, 2014.

Alfred Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *American Journal of Mathematics*, 76(3):620–630, 1954.

Shali Jiang, Daniel Jiang, Maximilian Balandat, Brian Karrer, Jacob Gardner, and Roman Garnett. Efficient nonmyopic bayesian optimization via one-shot multi-step trees. *NeurIPS*, 33:18039–18049, 2020.

Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabás Póczos. Multi-fidelity bayesian optimisation with continuous approximations. In *ICML*, pp. 1799–1808. PMLR, 2017.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On bayesian upper confidence bounds for bandit problems. In *AISTATS*, pp. 592–600. PMLR, 2012a.

Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *ALT*, pp. 199–213. Springer, 2012b.

Jack Kiefer and Jacob Wolfowitz. The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366, 1960.

Johannes Kirschner and Andreas Krause. Information directed sampling and bandits with heteroscedastic noise. In *COLT*, pp. 358–384. PMLR, 2018.

Harold J Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 1964.

Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.

Remi Lam, Karen Willcox, and David H Wolpert. Bayesian optimization with a finite budget: An approximate dynamic programming approach. *NeurIPS*, 29, 2016.

Tor Lattimore and Csaba Szepesvári. Bandit algorithms. *Cambridge University*, 355, 2019.

Jonas Močkus. On bayesian methods for seeking the extremum. In *IFIP Technical Conference on Optimization Techniques*, pp. 400–404. Springer, 1974.

Friedrich Pukelsheim. *Optimal design of experiments*. SIAM, 2006.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 1952.

Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.

Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *NeurIPS*, 27, 2014.

Issai Schur. Über eine klasse von mittelbildungen mit anwendungen auf die determinantentheorie. *Sitzungsberichte der Berliner Mathematischen Gesellschaft*, 22:9–20, 1923.

Jasper Snoek, Hugo Larochelle, and Ryan Adams. Practical bayesian optimization of machine learning algorithms. *NeurIPS*, 25, 2012.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. *NeurIPS*, 27, 2014.

Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *ICML*, 2010.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509–534, 2009.

Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *ICML*, pp. 3627–3635. PMLR, 2017.

Peter Whittle. Multi-armed bandits and the gittins index. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(2):143–149, 1980.

A Extended related work

Bayesian optimization. Bayesian optimization places a prior over the unknown function, updates it as evaluations arrive, and picks each evaluation by maximizing an acquisition function (Kushner, 1964; Močkus, 1974; Jones et al., 1998; Garnett, 2023). The classical rules all look one query ahead, scoring a candidate by how much it should beat the best value seen so far (Močkus, 1974; Jones et al., 1998), by the expected increase in the best *believed* value (Frazier et al., 2008; 2009), or by how much it reveals about the optimum (Villemonteix et al., 2009; Hennig & Schuler, 2012; Hernández-Lobato et al., 2014; Wang & Jegelka, 2017); non-myopic variants close the resulting gap approximately, with rollout, approximate dynamic programming, or differentiable multi-step lookahead (Ginsbourger & Le Riche, 2010; Lam et al., 2016; Jiang et al., 2020). Those methods approximate the optimal policy; we restrict the function class until it is computable.

Tractable optimal policies. Exact solutions to sequential design problems are rare and structural. Gittins (1979) showed that for the bandit problem of Bellman (1956), with an infinite horizon, geometric discounting and independent arms, the optimal policy is an *index policy*, meaning that one assigns each arm a number computed from that arm alone and pulls the largest (Whittle, 1980; Berry & Fristedt, 1985). That tractability breaks under finite horizons, correlated arms, or query costs, all present here. Ours has a different source: the terminal reward is quadratic in the posterior mean, and in a Gaussian linear model the posterior uncertainty follows a path fixed by the queries alone, so the value function separates. This is a close relative of the certainty-equivalence and no-dual-effect conditions in stochastic control, under which a linear-Gaussian system with quadratic cost is controlled optimally without probing (Bar-Shalom & Tse, 1974); here the statement concerns the experimenter rather than the controller.

Bandits. The bandit literature offers the sharpest available notion of optimality for adaptive data collection, and it is asymptotic: regret must grow at least like $\log T$ over a horizon of T rounds (Lai & Robbins, 1985), a rate attained, sometimes with its exact constant, by upper confidence bound methods (Robbins, 1952; Agrawal, 1995; Auer et al., 2002; Burnetas & Katehakis, 1996; Garivier & Cappé, 2011; Cappé et al., 2013) and by Bayesian rules such as Bayes-UCB, Thompson sampling and information-directed sampling (Kaufmann et al., 2012a; Thompson, 1933; Agrawal & Goyal, 2012; Kaufmann et al., 2012b; Russo & Van Roy, 2014; Kirschner & Krause, 2018); see Bubeck & Cesa-Bianchi (2012); Lattimore & Szepesvári (2019). Since our function class is linear in a fixed feature map, the closest setting is the linear bandit (Dani et al., 2008; Rusmevichientong & Tsitsiklis, 2010; Abbasi-Yadkori et al., 2011), with Rusmevichientong & Tsitsiklis (2010) treating a spherical action set and Srinivas et al. (2010) extending optimism to Gaussian process priors. Scoring only the final query places us instead in the *pure exploration* branch (Bubeck et al., 2009; Audibert & Bubeck, 2010; Garivier & Kaufmann, 2016), where Soare et al. (2014) and Fiez et al. (2019) observe that in a linear model much of the measurement allocation can be computed offline. In our setting adaptivity has no value at all, but we obtain that in exchange for exact rather than asymptotic optimality, in one parametric family with one cost structure.

Experimental design and query costs. After the reduction of Section 4.2 our problem is a design problem: choose which directions to measure so as to minimize a weighted sum of the remaining posterior variances, which is Bayesian A-optimality for equal weights and L-optimality in general (Fedorov, 1972; Pukelsheim, 2006; Chaloner & Verdinelli, 1995). That theory, including the Kiefer-Wolfowitz equivalence theorem (Kiefer & Wolfowitz, 1960), is developed for *approximate* designs, that is, probability distributions over the design space, because fractional allocations make the problem convex. Our budget instead buys a whole number of queries, making ours an *exact* design problem, normally handled by rounding; Lemma 3 characterizes the exact designs precisely, so no rounding is needed and no optimality is lost. Query costs that vary over the domain, finally, are usually handled heuristically, by dividing an acquisition function by the expected cost (Snoek et al., 2012) or by trading cheap-but-crude against expensive-but-accurate evaluations (Kandasamy et al., 2017), while Astudillo et al. (2021) treat budget exhaustion as a dynamic-programming constraint and solve it approximately. Here the cost is known in advance, and the interaction between cost geometry and reward geometry is what produces our main structural conclusion: the expensive directions are exactly the ones whose values do not affect the reward.

B MATRIX FACTS USED IN THE PROOFS

We use three elementary facts about the Loewner order and the trace. For symmetric matrices X, Y and $P \succeq 0$,

$$X \succeq Y \succ 0 \implies Y^{-1} \succeq X^{-1} \qquad \text{(more information means less uncertainty)}, \tag{20}$$

$$X \succeq Y \implies \operatorname{tr}(PX) \geq \operatorname{tr}(PY) \qquad \text{(the order survives a PSD-weighted trace)}, \tag{21}$$

$$\operatorname{tr}(PX) = \operatorname{tr}(XP) \qquad \text{(cyclicity of the trace)}. \tag{22}$$

Together, equation 20 and equation 21 say that $X \mapsto \operatorname{tr}(PX^{-1})$ reverses the Loewner order on positive definite matrices. We also use the Schur-Horn theorem (Schur, 1923; Horn, 1954), which characterizes the diagonals a symmetric matrix can have once its eigenvalues are fixed.

Lemma 7 (Schur-Horn, as used here)**.** *Let $\lambda \in \mathbb{R}^T$ and $d \in \mathbb{R}^T$ have the same sum. There exists a symmetric $T \times T$ matrix with spectrum λ and diagonal d if and only if d is* majorized *by λ, written $d \prec \lambda$: for every j, the sum of the j largest entries of d is at most the sum of the j largest entries of λ.*

Majorization in words. Two vectors with the same total can differ in how concentrated they are, and $d \prec \lambda$ says d is the more evenly spread of the two. Among all nonnegative vectors with a given sum T, the most concentrated is $(T, 0, \ldots, 0)$ and the most balanced is the constant vector $\mathbf{1}_T = (1, \ldots, 1)$, which is therefore majorized by every one of them. The Schur-Horn theorem says that for a symmetric matrix the eigenvalues are the most concentrated summary of its entries, and the achievable diagonals are precisely the vectors they majorize: rotating a matrix redistributes its "energy" along the diagonal without changing the spectrum, and one can always rotate all the way to a flat diagonal. A 2×2 example makes this concrete. The matrix $\operatorname{diag}(2, 0)$ has spectrum $(2, 0)$,

and conjugating it by a 45° rotation gives $\left(\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, which has the same spectrum and the flat diagonal $(1, 1)$.

C Reconstructing the query sequence

This appendix expands Section 4.5, which turns the optimal information matrix $M^\star$ into the T unit vectors that realize it. By construction $\operatorname{tr} M^\star = T$ and $\operatorname{rank} M^\star = N^\star \le T$, so Lemma 3 guarantees that such vectors exist; the sweep below, due to Bendel and Mickey, is the constructive content of that lemma in the case where the target diagonal is flat, and Algorithm 1 puts it together with the water-filling step of Proposition 5 and the change of coordinates back to input space. The sweep over N costs $O(n)$ after sorting the prior eigenvalues, since then $\mathcal{A}_N$ is a prefix of that order and the admissibility test equation 16 is $O(1)$; the rotation loop does at most $T - 1$ rotations of $O(N^\star)$ work each; and the eigendecomposition of $\Lambda^{(0)}$ dominates, giving $O(n^3 + TN^\star)$. Applied on the left of A, a Givens rotation replaces two rows $A_{r,:}$ and $A_{s,:}$ by

$$\begin{pmatrix} A_{r,:} \\ A_{s,:} \end{pmatrix} \longleftarrow \begin{pmatrix} \gamma & \varsigma \\ -\varsigma & \gamma \end{pmatrix} \begin{pmatrix} A_{r,:} \\ A_{s,:} \end{pmatrix}, \qquad \gamma^2 + \varsigma^2 = 1, \tag{23}$$

and leaves all other rows untouched. Start from any $A \in \mathbb{R}^{T \times n}$ with $A^\top A = M^\star$, for instance the matrix square root padded with zero rows. The total squared row norm is $\operatorname{tr}(AA^\top) = \operatorname{tr} M^\star = T$, so the row norms already average to one and the only problem is that they are unequal. Repeat the following. Pick a row r that is too long, $\|A_{r,:}\|^2 > 1$, and a row s that is too short, $\|A_{s,:}\|^2 < 1$; such a pair exists as long as some row is not yet of unit norm, because the squared norms sum to the number of rows. Choose the angle of the Givens rotation equation 23 so that row r comes out with squared norm exactly 1: writing $\tau = \varsigma/\gamma$ and $g_r = \|A_{r,:}\|^2$, $g_s = \|A_{s,:}\|^2$, $\rho = A_{r,:}A_{s,:}^\top$, the requirement is the quadratic

$$(g_s - 1)\,\tau^2 + 2\rho\,\tau + (g_r - 1) = 0, \tag{24}$$

whose discriminant $\rho^2 + (1 - g_s)(g_r - 1)$ is positive whenever $g_r > 1 > g_s$, so a real solution always exists. Row r is now finished and is never touched again; row s may incidentally reach unit norm too, in which case it is retired as well. Each rotation therefore retires at least one row, and the last one retires two, so after at most $T - 1$ rotations every row has unit norm. Left multiplication by an orthogonal matrix leaves $A^\top A = M^\star$ invariant throughout, so every iterate remains a valid design. Finally, flipping the sign of a row leaves both $A^\top A$ and the costs unchanged, so we may assume no row equals the excluded north pole e_1 and every $m^{(t)}$ lies in the image of ϕ; the queries are then $\theta^{(t)} = m^{(t)}_{2:n}/(1 - m^{(t)}_1)$. Since $M^\star c = 0$ we have $\sum_t (c^\top m^{(t)})^2 = c^\top M^\star c = 0$, so every query lies on the free equator and costs exactly κ, for a total of $\kappa T \le B$.

D Proofs

D.1 Proof of Proposition 1

Proof. Since the model is linear in z, the conditional expectation in equation 4 is $m^\top \mu - s\,\mathbb{E}[m']^\top \mu$, and by equation 3 the inner objective is $(x - sa)^2$ with

$$x := m^\top \mu, \qquad a := \frac{\beta}{\|c\|^2}\, c^\top \mu. \tag{25}$$

Minimizing over $s \in \{-1, 1\}$ picks whichever sign brings sa closer to x, giving $(|x| - |a|)^2$. It remains to understand which values x can take. Using the decomposition $\mu = \mu_\parallel + \mu_\perp$, write any feasible query as $m = \alpha \hat{c} + \sqrt{1 - \alpha^2}\, u$ with u a unit vector orthogonal to c. The budget constraint is then $|\alpha| \le \beta/\|c\|$, and

$$x = \alpha \frac{c^\top \mu}{\|c\|} + \sqrt{1 - \alpha^2}\; u^\top \mu_\perp. \tag{26}$$

Maximizing over u sets $u = \mu_\perp/\|\mu_\perp\|$. Maximizing the result over $\alpha \in \mathbb{R}$ *without* the budget constraint would give $|\alpha| = |c^\top \mu|/(\|c\|\|\mu\|)$, i.e. point straight at μ; so the first inequality in equation 6

Algorithm 1 Computing the optimal query sequence

Require: prior precision $\Lambda^{(0)} \succ 0$; noise variance σ^2; cost vector c and base cost $\kappa > 0$; budgets $B \geq \kappa$ and $\hat{B}$
Ensure: queries $\theta^{(1)}, \dots, \theta^{(T)}$ attaining the minimum in equation 14

$\hat{c} \leftarrow c/\|c\|$; $\quad p \leftarrow 1 - (\hat{B} - \kappa)/\|c\|^2$; $\quad T \leftarrow \lfloor B/\kappa \rfloor$ ▷ T = number of affordable queries
assert $p > 0$ **and** $\Lambda^{(0)}\hat{c} = (\hat{c}^\top \Lambda^{(0)} \hat{c})\,\hat{c}$ ▷ (A2), and c is an eigenvector of $\Lambda^{(0)}$
$Q\,\mathrm{diag}(\lambda)Q^\top \leftarrow \Lambda^{(0)}$, with $\hat{c}$ placed in column k of Q
$\pi \leftarrow$ ordering of $\{1, \dots, n\} \setminus \{k\}$ with $\lambda_{\pi(1)} \leq \cdots \leq \lambda_{\pi(n-1)}$ ▷ least-known directions first

$S \leftarrow 0$; $\quad R \leftarrow \sum_{i \neq k} p/\lambda_i$; $\quad J^\star \leftarrow +\infty$ ▷ water-filling sweep, cf. equation 17
for $N = 1$ **to** $\min(n-1, T)$ **do**
  $\ell \leftarrow \lambda_{\pi(N)}$; $\quad S \leftarrow S + \ell$; $\quad R \leftarrow R - p/\ell$ ▷ S = floors of chosen vessels; R = untouched tail
  if $\ell \leq \dfrac{T}{\sigma^2 N} + \dfrac{S}{N}$ **then** ▷ level clears the highest chosen floor, cf. equation 16
    $J \leftarrow \dfrac{pN^2}{T/\sigma^2 + S} + R$
    if $J < J^\star$ **then**
      $J^\star \leftarrow J$; $\quad N^\star \leftarrow N$; $\quad S^\star \leftarrow S$
    end if
  end if
end for
$i_j \leftarrow \pi(j)$ for $j = 1, \dots, N^\star$; $\quad d_j \leftarrow \dfrac{T}{N^\star} + \dfrac{\sigma^2 S^\star}{N^\star} - \sigma^2 \lambda_{i_j}$ ▷ water levels, from equation 18
$A \in \mathbb{R}^{T \times N^\star} \leftarrow 0$; $\quad A_{jj} \leftarrow \sqrt{d_j}$ for $j = 1, \dots, N^\star$ ▷ $A^\top A = \mathrm{diag}(d)$; rows not yet unit norm
$g_r \leftarrow \|A_{r,:}\|^2$ for $r = 1, \dots, T$; $\quad \mathcal{F} \leftarrow \{r : g_r = 1\}$ ▷ rows already finished
while $\mathcal{F} \neq \{1, \dots, T\}$ **do**
  pick $r, s \notin \mathcal{F}$ with $g_r > 1 > g_s$ ▷ both exist since $\sum_{r \notin \mathcal{F}} g_r = T - |\mathcal{F}|$
  $\rho \leftarrow A_{r,:} A_{s,:}^\top$; $\quad \tau \leftarrow \dfrac{-\rho + \sqrt{\rho^2 + (1 - g_s)(g_r - 1)}}{g_s - 1}$; $\quad \gamma \leftarrow 1/\sqrt{1 + \tau^2}$; $\quad \varsigma \leftarrow \tau\gamma$ ▷ root of equation 24
  $\begin{pmatrix} A_{r,:} \\ A_{s,:} \end{pmatrix} \leftarrow \begin{pmatrix} \gamma & \varsigma \\ -\varsigma & \gamma \end{pmatrix} \begin{pmatrix} A_{r,:} \\ A_{s,:} \end{pmatrix}$ ▷ Givens rotation equation 23: $A^\top A$ fixed
  $g_s \leftarrow g_r + g_s - 1$; $\quad g_r \leftarrow 1$; $\quad \mathcal{F} \leftarrow \{r : g_r = 1\}$ ▷ squared norms are traded, not created
end while

for $t = 1$ **to** T **do**
  $\tilde{m}^{(t)} \leftarrow \sum_{j=1}^{N^\star} A_{tj}\, e_{i_j}$; $\quad m^{(t)} \leftarrow Q\tilde{m}^{(t)}$; $\quad$ **if** $m^{(t)} = e_1$ **then** $m^{(t)} \leftarrow -m^{(t)}$
  $\theta^{(t)} \leftarrow m^{(t)}_{2:n} / \big(1 - m^{(t)}_1\big)$ ▷ inverse stereographic projection
end for
return $\theta^{(1)}, \dots, \theta^{(T)}$

is exactly the condition that the budget constraint binds. It then binds at $|\alpha| = \beta/\|c\|$, which is equation 8, and gives

$$\max_m |x| = \frac{\beta |c^\top \mu|}{\|c\|^2} + \sqrt{1 - \frac{\beta^2}{\|c\|^2}}\, \|\mu_\perp\| = |a| + \sqrt{p}\, \|\mu_\perp\|. \tag{27}$$

At the other extreme, since $n \geq 3$ there is a unit vector orthogonal to both μ and c, which is feasible and gives $x = 0$; and x varies continuously over the connected feasible set, so $|x|$ sweeps out all of $[0,\ |a| + \sqrt{p}\|\mu_\perp\|]$. The map $|x| \mapsto (|x| - |a|)^2$ is U-shaped, so it is maximized at one end of that

interval or the other:

$$\max_m \min_s (x - sa)^2 = \max \left\{ \underbrace{p\|\mu_\perp\|^2}_{\text{informative branch}} , \underbrace{a^2}_{\text{degenerate branch}} \right\}. \tag{28}$$

To compare the two, normalize $\|\mu\| = 1$ and write $u_0 = \beta/\|c\|$ and $v_0 = |c^\top \mu|/(\|c\|\|\mu\|)$, so that $p = 1 - u_0^2$, $\|\mu_\perp\|^2 = 1 - v_0^2$ and $a^2 = u_0^2 v_0^2$. The inequality $p\|\mu_\perp\|^2 \geq a^2$ then reads $(1 - u_0^2)(1 - v_0^2) \geq u_0^2 v_0^2$, i.e. $u_0^2 + v_0^2 \leq 1$, which is the second condition in equation 6. Hence the informative branch attains the maximum, and $p\|\mu_\perp\|^2 = \mu^\top P \mu$. □

Remark 1. The second condition in equation 6 cannot be dropped: without it the maximum in equation 7 is attained instead by a query orthogonal to μ, which carries no information about where the optimum lies but inherits the magnitude of the adversarially signed baseline, and is worth $a^2 = (\hat{B} - \kappa)(c^\top \mu)^2/\|c\|^4$. In the notation of the proof the two conditions of equation 6 read $u_0 \leq v_0$ and $u_0^2 + v_0^2 \leq 1$, which are compatible: $u_0 = v_0 = 1/2$ satisfies both. But both constrain the realized posterior mean, which is random; they delimit the regime in which the quadratic terminal value, and hence Section 4.1, applies.

D.2 Proof of Proposition 2

Proof. Write $W(\mu, \Lambda, B) := \mu^\top P \mu + \operatorname{tr}(P\Lambda^{-1}) - J(\Lambda, B)$ and evaluate the right-hand side of equation 10 with V replaced by W, omitting superscripts for clarity. Consider first the continuation branch and fix a feasible m, i.e. $k(m) \leq B$. Abbreviate the step direction by $v := \Lambda^{-1} m / \sqrt{\sigma^2 + m^\top \Lambda^{-1} m}$. Only the first term of W depends on μ, and it is quadratic, so $\mathbb{E}_\varepsilon[(\mu + v\varepsilon)^\top P (\mu + v\varepsilon)] = \mu^\top P \mu + v^\top P v$, and the expectation evaluates to

$$\frac{m^\top \Lambda^{-1} P \Lambda^{-1} m}{\sigma^2 + m^\top \Lambda^{-1} m} + W\left(\mu, \Lambda + \frac{mm^\top}{\sigma^2}, B - k(m)\right). \tag{29}$$

Substituting the form of W and applying the Sherman-Morrison identity, which describes how a rank-one increase in precision decreases the covariance,

$$\left(\Lambda + \tfrac{mm^\top}{\sigma^2}\right)^{-1} = \Lambda^{-1} - \frac{\Lambda^{-1} m m^\top \Lambda^{-1}}{\sigma^2 + m^\top \Lambda^{-1} m}, \tag{30}$$

to the second term gives

$$\begin{aligned} &\frac{m^\top \Lambda^{-1} P \Lambda^{-1} m}{\sigma^2 + m^\top \Lambda^{-1} m} + \mu^\top P \mu + \operatorname{tr}\big(P\Lambda^{-1}\big) - \operatorname{tr}\left(P \frac{\Lambda^{-1} m m^\top \Lambda^{-1}}{\sigma^2 + m^\top \Lambda^{-1} m}\right) \\ &\qquad - \min_{(m^{(t)}) \in \mathcal{D}(B - k(m))} \operatorname{tr}\left(P\Big(\Lambda + \tfrac{mm^\top}{\sigma^2} + \sum_t \tfrac{m^{(t)} m^{(t)\top}}{\sigma^2}\Big)^{-1}\right). \end{aligned} \tag{31}$$

By cyclicity of the trace, equation 22, the fourth term equals the first and they cancel, which is the formal counterpart of the observation that the immediate gain from one query is exactly the uncertainty it removes. What remains is

$$\mu^\top P \mu + \operatorname{tr}\big(P\Lambda^{-1}\big) - \min_{(m^{(t)}) \in \mathcal{D}(B - k(m))} \operatorname{tr}\left(P\Big(\Lambda + \tfrac{mm^\top}{\sigma^2} + \sum_t \tfrac{m^{(t)} m^{(t)\top}}{\sigma^2}\Big)^{-1}\right). \tag{32}$$

For infeasible m we have $B - k(m) < 0$, so $\mathcal{D}(B - k(m)) = \emptyset$ and equation 32 equals $-\infty$; infeasible queries are thus excluded automatically, and the explicit restriction in equation 10 is only for emphasis. Now maximize over m. Prepending m to a sequence in $\mathcal{D}(B - k(m))$ produces exactly the *nonempty* sequences in $\mathcal{D}(B)$, so the maximum of equation 32 over feasible m is $\mu^\top P \mu + \operatorname{tr}(P\Lambda^{-1})$ minus the minimum over nonempty sequences. The stopping branch $\mu^\top P \mu$ accounts for the *empty* sequence, for which the bracketed trace is $\operatorname{tr}(P\Lambda^{-1})$ and the two extra terms cancel. Taking the maximum of the two branches therefore yields the minimum over all of $\mathcal{D}(B)$, that is $W(\mu, \Lambda, B)$ itself. Since the empty sequence is always available when $B \geq 0$, we have $J(\Lambda, B) \leq \operatorname{tr}(P\Lambda^{-1})$ and hence $W \geq \mu^\top P \mu$, as it must be. Uniqueness of the solution to equation 10 was noted in Section 4.1, so $W = V$. □

D.3 Proof of Lemma 3

Proof. Necessity was argued in Section 4.2. For sufficiency, let $r = \operatorname{rank} M \le T$ and let $M = U\Lambda_M U^\top$ be an eigendecomposition. Deleting the vanishing rows of $\Lambda_M^{1/2}U^\top$ and appending $T-r$ zero rows produces a matrix $A \in \mathbb{R}^{T\times n}$ with $A^\top A = M$. Writing $m^{(t)}$ for the t-th row of A, we have $M = \sum_t m^{(t)}m^{(t)\top}$ already; what may fail is that the rows are not unit vectors. Their squared norms are the diagonal entries of the Gram matrix $AA^\top$, so we need to modify A so that $\operatorname{diag}(AA^\top) = \mathbf{1}_T$ while keeping $A^\top A = M$. The two quantities behave differently under an orthogonal change of basis $A \mapsto RA$ with R a $T\times T$ orthogonal matrix: this leaves $A^\top A = A^\top R^\top RA$ untouched, but conjugates $AA^\top \mapsto R(AA^\top)R^\top$, changing its diagonal while preserving its spectrum. So the question is whether the flat diagonal $\mathbf{1}_T$ is reachable. Now $AA^\top$ is positive semi-definite with spectrum equal to the nonzero eigenvalues of M padded by zeros, and with trace $\operatorname{tr} M = T$. Since $\mathbf{1}_T$ is majorized by every nonnegative vector summing to T, Lemma 7 supplies a symmetric H with the same spectrum as $AA^\top$ and $\operatorname{diag}(H) = \mathbf{1}_T$. Two symmetric matrices with the same spectrum are orthogonally conjugate, so $H = R(AA^\top)R^\top$ for some orthogonal R, and $\tilde{A} := RA$ satisfies $\tilde{A}^\top\tilde{A} = M$ with all rows of unit norm. Appendix C gives an explicit sweep of Givens rotations that produces such an R. Finally, $m^{(t)}$ and $-m^{(t)}$ have the same outer product and the same cost by equation 2, so any row that happens to equal e_1 may be replaced by its negative. □

D.4 Alignment on a spectral orbit

Proposition 4 rests on the following lemma, which says that misaligning a design with the prior is always weakly harmful.

Lemma 8 (Alignment on a spectral orbit)**.** *Let L be symmetric positive definite on $\mathbb{R}^d$, let $\nu \in \mathbb{R}^d_{\ge 0}$ be a fixed vector of eigenvalues, and let $\mathcal{O}_\nu := \{R\operatorname{diag}(\nu)R^\top : R \text{ orthogonal}\}$ be the set of all symmetric positive semi-definite matrices with that spectrum, i.e. all the rotations of* $\operatorname{diag}(\nu)$*. Then $g(N) := \operatorname{tr}\left((L+N)^{-1}\right)$ attains its minimum over $\mathcal{O}_\nu$, and every minimizer commutes with L.*

In other words, fix how much information the design will collect and how it will be spread across its own principal axes; the only remaining freedom is the orientation of those axes relative to the prior, and the best orientation is the aligned one. The proof is a first-order argument: rotating a design infinitesimally changes the objective at a rate proportional to a commutator, and that rate vanishes in every direction only when the commutator itself is zero, which is alignment.

Proof. $\mathcal{O}_\nu$ is compact, being a continuous image of the compact group of orthogonal matrices, and g is continuous on it since $L + N \succ 0$ for every $N \in \mathcal{O}_\nu$; hence a minimizer $N^\star$ exists. To take derivatives along the set, recall that the curves staying inside $\mathcal{O}_\nu$ are exactly the conjugations by rotations. Fix an antisymmetric S, so that $e^{\epsilon S}$ is orthogonal for every ϵ, and consider $N(\epsilon) = e^{\epsilon S}N^\star e^{-\epsilon S} \in \mathcal{O}_\nu$, which satisfies $N(0) = N^\star$ and $\dot{N}(0) = SN^\star - N^\star S =: [S, N^\star]$. Writing $A := (L+N^\star)^{-1}$ and using $\frac{d}{d\epsilon}X^{-1} = -X^{-1}\dot{X}X^{-1}$ together with cyclicity of the trace,

$$\frac{d}{d\epsilon}g(N(\epsilon))\bigg|_{\epsilon=0} = -\operatorname{tr}\left(A[S,N^\star]A\right) = -\operatorname{tr}\left(A^2(SN^\star - N^\star S)\right) = -\operatorname{tr}\left(S\left[N^\star, A^2\right]\right). \tag{33}$$

Since $N^\star$ is a minimizer, this vanishes for every antisymmetric S. The commutator of two symmetric matrices is antisymmetric, so $[N^\star, A^2]$ is antisymmetric; and an antisymmetric matrix orthogonal to every antisymmetric matrix under the trace inner product must be zero. Hence $[N^\star, A^2] = 0$. Since $A \succ 0$, A is the unique positive square root of A^2 and so shares its eigenspaces; therefore $N^\star$ commutes with A, hence with $A^{-1} = L + N^\star$, hence with L. □

D.5 Proof of Proposition 4

The rank constraint is what makes Proposition 4 nontrivial. The obvious argument, which averages a minimizer over the rotations fixing $\Lambda^{(0)}$ and P and then invokes convexity of $M \mapsto \operatorname{tr}(P(\Lambda + M)^{-1})$, does not work: averaging can raise the rank while holding the trace fixed, pushing the average outside $\mathcal{M}(B)$ and producing a "design" that no whole number of queries can realize. We instead rotate within a set of matrices of fixed spectrum, along which trace and rank are constant by construction.

Proof. Step 0 (attainment and block structure). The constraint $\kappa \operatorname{tr} M \le \kappa \operatorname{tr} M + c^\top M c \le B$ bounds $\operatorname{tr} M$, so $\mathcal{M}(B)$ is bounded; and it is the union over the finitely many integers $T \in \{0, 1, \dots, \lfloor B/\kappa \rfloor\}$ of the sets $\{M \succeq 0 : \operatorname{tr} M = T,\ \operatorname{rank} M \le T,\ c^\top M c \le B - \kappa T\}$, each of which is closed, the rank condition because rank can only drop in a limit. Hence $\mathcal{M}(B)$ is compact, and G is continuous on it because $\Lambda^{(0)} \succ 0$, so a minimizer exists. Since $p > 0$, the eigenspaces of P are $\operatorname{span}(\hat{c})$, with eigenvalue 0, and $W := \hat{c}^\perp$, with eigenvalue p. Commuting matrices preserve one another's eigenspaces, so $\Lambda^{(0)}\hat{c} = \lambda_k \hat{c}$ for some $\lambda_k > 0$, and $\Lambda^{(0)}$ maps W into W. Splitting $\mathbb{R}^n = \operatorname{span}(\hat{c}) \oplus W$ into "the expensive direction" and "everything else", we may write in block form

$$\Lambda^{(0)} = \begin{pmatrix} \lambda_k & 0 \\ 0 & L \end{pmatrix}, \qquad P = \begin{pmatrix} 0 & 0 \\ 0 & pI_W \end{pmatrix}, \qquad M = \begin{pmatrix} a & b^\top \\ b & N \end{pmatrix}, \tag{34}$$

with $L \succ 0$ on W. Here $a \ge 0$ is the information the design spends on the expensive direction, b describes how that spending is correlated with the rest, and $N \succeq 0$ is the information collected inside W. In this notation $c^\top M c = \|c\|^2 a$ and $\operatorname{tr} M = a + \operatorname{tr} N$. Since P annihilates the first block, only the W block of the inverse matters, and the block-inversion formula (the Schur complement of the $(1,1)$ entry) gives

$$G(M) = p \operatorname{tr}\left[\left(L + \frac{N}{\sigma^2} - \frac{bb^\top}{\sigma^4(\lambda_k + a/\sigma^2)}\right)^{-1}\right]. \tag{35}$$

The subtracted term is the price of correlation: information about W that is entangled with the effectively unobserved direction $\hat{c}$ is discounted. *Step 1 (an optimal design annihilates c).* Let $M \in \mathcal{M}(B)$ with $T := \operatorname{tr} M$. If $T = 0$ then $M = 0$ and there is nothing to prove, so assume $T \ge 1$. We move all of the mass a off the expensive direction and into W. Choose a unit vector $u \in W$ lying in the range of N if $N \ne 0$, and arbitrarily otherwise, and set

$$M' := \begin{pmatrix} 0 & 0 \\ 0 & N' \end{pmatrix}, \qquad N' := N + a\, uu^\top. \tag{36}$$

Feasibility. $M' \succeq 0$ and $\operatorname{tr} M' = \operatorname{tr} N + a = T \in \mathbb{Z}$, while $\kappa \operatorname{tr} M' + c^\top M' c = \kappa T \le \kappa T + \|c\|^2 a \le B$, so the budget constraint is only relaxed. The rank condition needs the choice of u: since N is a principal submatrix of the positive semi-definite M, we have $\operatorname{rank} N \le \operatorname{rank} M \le T$; if $N \ne 0$ then $\operatorname{range}(N') \subseteq \operatorname{range}(N)$ by construction, while $N' \succeq N$ forces $\ker N' \subseteq \ker N$, so $\operatorname{rank} M' = \operatorname{rank} N' = \operatorname{rank} N \le T$. If instead $N = 0$, then $b = 0$ by positive semi-definiteness, so $a = T \ge 1$ and $\operatorname{rank} M' = 1 \le T$. (This care is needed because piling mass onto a fresh direction could otherwise push the rank above the number of queries, making the design unrealizable.) *Improvement.* The subtracted term in equation 35 is positive semi-definite, so the bracketed matrix is $\preceq L + N/\sigma^2 \preceq L + N'/\sigma^2$. By equation 20 and equation 21, $X \mapsto \operatorname{tr}(X^{-1})$ reverses the Loewner order, so

$$G(M) \ \ge\ p \operatorname{tr}\left[\left(L + \tfrac{N}{\sigma^2}\right)^{-1}\right] \ \ge\ p \operatorname{tr}\left[\left(L + \tfrac{N'}{\sigma^2}\right)^{-1}\right] \ =\ G(M'). \tag{37}$$

Applying this to a minimizer, we may assume from now on that $a = 0$ and $b = 0$, i.e. $Mc = 0$. *Step 2 (reduction to W).* For designs of the form $M = 0 \oplus N$ the objective is $G(M) = p \operatorname{tr}((L + N/\sigma^2)^{-1})$, and membership in $\mathcal{M}(B)$ reduces to

$$N \succeq 0, \qquad \operatorname{tr} N = T \in \mathbb{Z}_{\ge 0}, \qquad \kappa T \le B, \qquad \operatorname{rank} N \le T. \tag{38}$$

All of these depend on N only through its spectrum, so the feasible set is a union of the sets $\mathcal{O}_\nu$ of Lemma 8. Writing $\mathcal{S}_T$ for the set of spectra $\nu \in \mathbb{R}^{n-1}_{\ge 0}$ with $\sum_i \nu_i = T$ and at most T nonzero entries,

$$\min_{\substack{M \in \mathcal{M}(B) \\ Mc = 0}} G(M) = \min_{0 \le T \le \lfloor B/\kappa \rfloor} \ \min_{\nu \in \mathcal{S}_T} \ \min_{N \in \mathcal{O}_\nu} \ p \operatorname{tr}\left[\left(L + \tfrac{N}{\sigma^2}\right)^{-1}\right]. \tag{39}$$

Step 3 (alignment). Apply Lemma 8 with L and N/σ^2: the innermost minimum is attained at some $N^\star$ commuting with L. Combining with Steps 0–2, the design $M^\star = 0 \oplus N^\star$ minimizes G over $\mathcal{M}(B)$, satisfies $M^\star c = 0$, and commutes with $\Lambda^{(0)} = \lambda_k \oplus L$. Finally, $M^\star \hat{c} = 0$ gives $PM^\star = pM^\star = M^\star P$, so $M^\star$ commutes with P too. A commuting family of symmetric matrices is simultaneously diagonalizable, and since $\hat{c}$ is a common eigenvector it may be taken as a column of Q. □

D.6 Proof of Proposition 5 and Corollary 6

Proof of Proposition 5. By Proposition 4 we may take $d_k = 0$; alternatively this follows directly here, since moving mass from d_k to any d_i with $i \neq k$ preserves the total, does not increase the cost, does not increase the number of nonzero entries, and weakly decreases the objective, index k contributing nothing to equation 15. Discard index k from now on. The objective is strictly decreasing in each d_i, and increasing $\sum_i d_i$ relaxes the rank constraint as well, so we should buy as many queries as we can afford: the constraints $\kappa \sum_i d_i \leq B$ and $\sum_i d_i \in \mathbb{Z}$ bind at $\sum_{i \neq k} d_i = T$. The problem reduces to

$$\min \left\{ \sum_{i \neq k} \frac{p}{\lambda_i + \frac{d_i}{\sigma^2}} \;:\; d_i \geq 0, \;\; \big|\{i \neq k : d_i > 0\}\big| \leq \sum_{i \neq k} d_i = T \right\}, \tag{40}$$

whose feasible set is compact and whose objective is continuous, so a minimizer exists. *Which coordinates are active.* Call the coordinates with $d_i > 0$ *active*, and suppose a minimizer has N of them; feasibility forces $1 \leq N \leq \min(n-1, T)$. Take an active i and an inactive j and swap their values, which keeps the point feasible. The objective changes by

$$p\,[h(\lambda_i) - h(\lambda_j)]\,, \qquad h(x) := \frac{1}{x} - \frac{1}{x + \frac{d_i}{\sigma^2}}, \tag{41}$$

where $h(x)$ is the improvement obtained by pouring the amount d_i into a vessel with floor x. Since $h'(x) = -x^{-2} + (x + d_i/\sigma^2)^{-2} < 0$, the function h is strictly decreasing: a given amount of water helps more in a lower vessel. Optimality requires the swap not to help, i.e. the displayed change to be nonnegative, which gives $\lambda_i \leq \lambda_j$. Hence at least one minimizer with N active coordinates has its active set equal to $\mathcal{A}_N$; ties among equal λ_i are immaterial. *How much goes into each active coordinate.* Fix N and its active set $\mathcal{A}_N$. On the affine set $\sum_{i \in \mathcal{A}_N} d_i = T$ the objective is convex in the active variables, so the first-order conditions are sufficient, not merely necessary. Ignoring the nonnegativity constraints for the moment, the Lagrangian is

$$\sum_{i \notin \mathcal{A}_N \cup \{k\}} \frac{p}{\lambda_i} + \sum_{i \in \mathcal{A}_N} \frac{p}{\lambda_i + \frac{d_i}{\sigma^2}} - \eta\Big(T - \sum_{i \in \mathcal{A}_N} d_i\Big), \tag{42}$$

and setting the derivative with respect to d_i to zero gives

$$-\frac{p}{\sigma^2\big(\lambda_i + \frac{d_i}{\sigma^2}\big)^2} + \eta = 0 \qquad \Longrightarrow \qquad d_i = \sigma\sqrt{\frac{p}{\eta}} - \sigma^2\lambda_i. \tag{43}$$

The right-hand side depends on i only through $-\sigma^2\lambda_i$: this is precisely the equalization of water levels. Imposing $\sum_{i \in \mathcal{A}_N} d_i = T$ pins down the multiplier and hence the level,

$$\eta = \frac{p}{\Big(\frac{T}{N\sigma} + \frac{\sigma}{N}\sum_{j \in \mathcal{A}_N} \lambda_j\Big)^2}, \qquad \text{hence} \qquad d_i = \frac{T}{N} + \frac{\sigma^2}{N}\sum_{j \in \mathcal{A}_N} \lambda_j - \sigma^2\lambda_i, \tag{44}$$

which is equation 18. This candidate satisfies the omitted constraints $d_i \geq 0$ exactly when N is admissible in the sense of equation 16. If N is not admissible, then by convexity the minimizer over the closed set $\{d_i \geq 0,\ i \in \mathcal{A}_N\}$ has a vanishing coordinate, so it does not in fact have N active coordinates, and no minimizer with exactly N active coordinates exists. *Choosing N.* For admissible N, all active coordinates share the level $\frac{1}{N}(T/\sigma^2 + \sum_{j \in \mathcal{A}_N} \lambda_j)$, so the N active terms of the objective contribute $pN^2/(T/\sigma^2 + \sum_{j \in \mathcal{A}_N} \lambda_j)$ in total, which is the first term of the bracket in equation 17; the inactive coordinates contribute the second. Sweeping over the admissible N, a nonempty set since $N = 1$ satisfies equation 16 trivially, therefore gives the global optimum. □

Proof of Corollary 6. Immediate from equation 18, since d_i is decreasing in λ_i across the active set. Alternatively, and without reference to the explicit solution: $(x, y) \mapsto (x + y)^{-1}$ has mixed partial derivative $2(x+y)^{-3} > 0$, so it is strictly supermodular, meaning that increasing one argument raises the marginal effect of the other; and a sum of a supermodular function over a pairing is minimized by pairing the largest with the smallest. □

E EXPERIMENTAL DETAILS

Generative model. For each dimension n and each trial we draw a Haar-random orthogonal matrix $Q \in \mathbb{R}^{n\times n}$, as the orthogonal factor of the QR decomposition of a matrix of independent standard normals, and take the expensive direction to be $c = \|c\|\, Q_{:,k}$ with $k = 1$, which makes $\hat{c}$ an eigenvector of $\Lambda^{(0)}$ and hence guarantees the commuting hypothesis of Proposition 4. The prior precisions are

$$\lambda_i = (n-1)\, b_i, \qquad \log b_i \sim \mathrm{Unif}\big[\log 0.1,\, \log 10\big] \text{ independently}, \qquad \Lambda^{(0)} = Q\,\mathrm{diag}(\lambda) Q^\top, \tag{45}$$

the prior mean is $\mu^{(0)} = \rho\,\hat{c}$ with $\rho = 1$, and the parameter vector is $z \sim \mathcal{N}(\mu^{(0)}, \Lambda^{(0)-1})$, so the function really is drawn from the prior the design is planned against. Observations are $y^{(t)} = m^{(t)\top} z + \sigma\varepsilon^{(t)}$ with $\varepsilon^{(t)}$ independent standard normals. The remaining parameters are $\kappa = 1$, $\|c\| = 2$, $\sigma^2 = 1/(n-1)$, $\hat{B} = \kappa + (0.3\|c\|)^2 = 1.36$, so that $\beta/\|c\| = 0.3$ and $p = 0.91$, and $B = 2n\kappa$, giving $T = \lfloor B/\kappa \rfloor = 2n$. Every method sees the same instance and the same stream of standardized noise draws, so all comparisons are paired.

Scoring. We score each run by the exact value of the max-min problem equation 4 rather than by $\mu^{(T)\top} P \mu^{(T)}$, which equals it only where (A3) holds.

Making the objective dimension-free. The objective is an *extensive* quantity: by Proposition 2 the expected gain from a design is $\mathrm{tr}(P\Lambda^{(0)-1}) - \mathrm{tr}(P\Lambda^{(T)-1})$, which aggregates the reward-weighted uncertainty removed across the $n-1$ rewarded directions. With n-independent parameters its mean therefore grows linearly in n, and curves at different n are not on a comparable footing. Two normalisations remove that growth, and both are needed. Scaling the prior precisions by $n-1$, as in equation 45, holds the total prior uncertainty $\mathrm{tr}(P\Lambda^{(0)-1}) = p\sum_{i\neq k}\lambda_i^{-1}$ at $O(1)$: it spreads a fixed amount of uncertainty over more directions, as a fixed signal variance would. On its own this makes the objective decrease like $1/n$, because each direction is then known to precision of order n and a handful of unit-precision observations barely move it. Scaling the noise variance as $\sigma^2 = 1/(n-1)$ repairs that: one observation then contributes precision $1/\sigma^2 = n-1$, a fixed multiple of the prior precisions, so the per-direction signal-to-noise ratio is dimension-free and each query removes the same fraction of a direction's uncertainty at every n. Together the two hold the mean objective within 10% over $n = 3, \ldots, 16$ with no trend (Table 3). Using $n-1$ rather than n matters, since the quantity that must be held fixed is $\sum_{i\neq k}\lambda_i^{-1}$, a sum of $n-1$ terms.

Baseline implementations. *Cost-blind A-optimality* calls the water-filling and Bendel-Mickey routines of Algorithm 1, but runs them over all n coordinates rather than the $n-1$ rewarded ones, which is what replacing P by I amounts to. Because it plans as though every query cost only κ, the design it returns has total price $\kappa T + c^\top M c > B$; we issue its queries in order and stop when the budget runs out. *Uncertainty sampling* maintains the diagonal of the posterior precision in the basis Q and repeatedly queries the affordable eigendirection of smallest precision, updating $\lambda_i \mapsto \lambda_i + 1/\sigma^2$ after each query. Since all its queries are eigendirections, the posterior precision stays diagonal in Q and this rule is implemented exactly rather than approximately. *Random* rejection-samples uniform points of the sphere until it finds an affordable one, and stops once the remaining budget falls below κ, below which nothing is affordable.

How the baselines spend their budget. Table 1 records how many queries each method actually issues and how much squared weight it places on the expensive direction, $\hat{c}^\top M \hat{c}$. Our design places none and issues every query at the base price, as Proposition 4 guarantees. The baselines place an amount that grows only weakly with n, and pay for it twice: directly, in a surcharge that buys them fewer queries, and indirectly, in information spent on a direction the reward ignores.

Behaviour at large n. Table 2 extends the comparison to $n = 128$. For the two structured baselines the product of n and the relative shortfall is flat, confirming a $\Theta(1/n)$ decay: their waste is confined to one direction out of n, and the budget grows like $2n$. Random sampling instead approaches a nonzero limit near 8%, because its misallocation is spread over all directions rather than concentrated on one, and no amount of extra dimension dilutes it.

n	$T = 2n$	Ours	A-optimal	Uncertainty	Random
		queries issued / weight on $\hat{c}$			
3	6	6.0 / 0.00	2.9 / 0.66	4.0 / 0.51	3.0 / 0.66
6	12	12.0 / 0.00	7.5 / 0.88	8.6 / 0.85	7.4 / 1.04
10	20	20.0 / 0.00	14.8 / 1.08	15.7 / 1.08	14.2 / 1.33
16	32	32.0 / 0.00	25.7 / 1.41	26.4 / 1.40	25.4 / 1.55

Table 1: Queries issued and squared weight placed on the expensive direction, averaged over 400 instances. The budget allows $T = 2n$ queries at the base price.

		Cost-blind A-optimal		Uncertainty sampling		Random
n	instances	shortfall	$n\times$	shortfall	$n\times$	shortfall
8	600	7.23%	0.58	6.03%	0.48	18.87%
16	400	3.23%	0.52	2.73%	0.44	12.65%
32	250	1.74%	0.56	1.47%	0.47	10.06%
64	120	0.89%	0.57	0.76%	0.49	8.46%
128	60	0.44%	0.56	0.42%	0.54	7.73%

Table 2: Mean per-instance relative shortfall in expected objective against our design, extended to large n. The $n\times$ columns are n times the shortfall; their flatness is the $\Theta(1/n)$ decay. Random sampling's shortfall does not follow the same law and tends to a positive constant.

Assumption (A3) along the query path. (A3) constrains the realized posterior mean, which is random, so it can fail on individual runs; see Remark 1. It holds on 95.4% of runs under our design and on 87.0% under the baselines, the difference arising because designs that spend on $\hat{c}$ perturb $c^\top \mu^{(T)}$ away from its prior value. On runs where it fails, the score we report is the degenerate branch of equation 28, which is the true value of the max-min problem there; no method is ever credited with a value the assumptions do not support.

Full results. Table 3 gives the numbers behind Figure 2.

n	Optimal (ours)	Cost-blind A-optimal	Uncertainty sampling	Random
3	1.684 ± 0.067	1.491 ± 0.062	1.533 ± 0.063	1.364 ± 0.057
4	1.853 ± 0.067	1.697 ± 0.062	1.756 ± 0.064	1.404 ± 0.055
5	1.765 ± 0.050	1.649 ± 0.049	1.680 ± 0.049	1.372 ± 0.044
6	1.752 ± 0.044	1.643 ± 0.043	1.687 ± 0.043	1.387 ± 0.038
7	1.753 ± 0.044	1.683 ± 0.044	1.701 ± 0.043	1.435 ± 0.038
8	1.729 ± 0.038	1.641 ± 0.037	1.656 ± 0.037	1.438 ± 0.033
9	1.778 ± 0.038	1.688 ± 0.037	1.712 ± 0.038	1.454 ± 0.033
10	1.689 ± 0.033	1.611 ± 0.032	1.631 ± 0.032	1.448 ± 0.030
11	1.740 ± 0.033	1.656 ± 0.032	1.693 ± 0.032	1.483 ± 0.030
12	1.723 ± 0.030	1.675 ± 0.030	1.682 ± 0.031	1.506 ± 0.028
13	1.695 ± 0.029	1.637 ± 0.028	1.642 ± 0.028	1.480 ± 0.026
14	1.755 ± 0.029	1.693 ± 0.028	1.698 ± 0.028	1.520 ± 0.026
15	1.689 ± 0.027	1.635 ± 0.026	1.649 ± 0.027	1.468 ± 0.024
16	1.737 ± 0.026	1.701 ± 0.027	1.712 ± 0.027	1.522 ± 0.024

Table 3: Mean terminal objective $\pm$ one standard error over 2000 instances per n. The ratio of largest to smallest mean along each curve is 1.10 for our design and 1.12–1.14 for the baselines, so the normalisation leaves no trend in n.