GENERATOR-AWARE MEASUREMENT DESIGN FOR BUDGETED DISTRIBUTION RECOVERY

Anonymous authors
Paper under double-blind review

ABSTRACT

We study high-dimensional distribution recovery under a limited measurement budget. Target individuals are sampled from an unknown target distribution, and for each sampled individual only one costly subset of variables can be measured. A pretrained generator is available for a fully observed reference population, and the target is modeled as a low-dimensional exponential tilt of the reference law. We propose *Generator-Aware Measurement Design (GAMD)*, which uses the generator to decide which variable subsets are most informative about how the target population differs from the reference. For each candidate subset, GAMD evaluates how much of the full target-shift signal can be recovered from the variables that would be observed; this quantity determines the Fisher information carried by that measurement. GAMD then combines these information values with measurement costs to allocate the budget and refines the target distribution from the acquired partial data. We prove that GAMD asymptotically matches the oracle design despite the unknown target shift and establish finite-budget guarantees. Synthetic and real-data experiments validate the proposed mechanism and demonstrate improved recovery under constrained measurements.

1 INTRODUCTION

Recovering a complex joint distribution is important in many applications, where a population model may support sampling, uncertainty quantification, and downstream analyses. Generative models provide a natural representation of such distributions, but learning or adapting them to a new population requires informative target measurements to collect data. In many settings, however, complete measurements are too costly to collect. For instance, survey respondents may answer only selected questionnaire modules (Raghunathan & Grizzle, 1995; Jang et al., 2022), targeted spatial transcriptomic assays measure only a subset of genes (Kuemmerle et al., 2024), and diagnostic systems may acquire only a subset of costly tests (Ma et al., 2018; Li & Oliva, 2021).

Motivated by these settings, we study a population-level acquisition problem in which each individual corresponds to a multivariate draw from an unknown target distribution. For each individual, we choose a subset of its variables to measure and observe only that subset. We call each allowable subset a *measurement panel*. Each panel has an associated measurement cost, and under a total measurement budget, we decide how to allocate different panels across independently sampled individuals so as to recover the complete joint distribution from the resulting partial observations.

This problem is fundamentally difficult. Restricted panels generally do not identify an unrestricted high-dimensional joint law, and their value depends on the unknown target population. We therefore need additional structure beyond the partial target measurements. In many applications, rich historical data are available from previous cohorts, survey waves, experiments, or deployments, and can be summarized by a pretrained reference generator that captures dependencies among variables. This suggests using the reference structure to decide which target measurements are most informative. Our central question is: *given a pretrained reference generator, how should we allocate a limited measurement budget in a new target population so as to recover its full distribution?*

We address this question by keeping the reference generator fixed and modeling only the reference-to-target change through a low-dimensional exponential density-ratio model. This is natural when the target largely preserves the reference support and dependence structure but reweights different regions or subpopulations. Exponential tilting provides a classical model for such population shifts

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.

(Anderson, 1979; Qin, 1998; Wu & Lu, 2016; Chen et al., 2023). This separation lets the generator retain the high-dimensional dependence learned from reference data, while the target data only estimate a small number of parameters describing how the population has changed. The acquisition problem remains challenging because this shift is unknown and the value of each panel depends on both the shift and the conditional relationships encoded by the generator.

We propose *Generator-Aware Measurement Design (GAMD)*, which uses the reference generator to decide what to measure rather than simply to fill in missing values. Because the value of a panel depends on the unknown target shift, GAMD first uses a small initialization sample to obtain a rough shift estimate. At this estimated target, the generator evaluates how much of the full target-shift signal can be recovered from each candidate panel; this quantity is exactly the panel's Fisher information and includes indirect information carried through reference dependence. GAMD then combines these panel information matrices with measurement costs to allocate the remaining budget toward the most informative measurements for full-distribution recovery. Finally, the acquired partial observations are used to refine the target-shift estimate and update the target distribution.

We provide theoretical guarantees for GAMD. As the total measurement budget B grows, GAMD achieves the optimal $1/B$ decay of the full-distribution KL risk, with a leading constant matching the local asymptotic minimax lower bound for regular information-stable prospective acquisition strategies. Thus, despite not knowing the target shift, GAMD asymptotically attains the oracle design value and its first-order recovery risk. We further establish finite-budget guarantees for the learned allocation and the resulting distribution-recovery error. Synthetic and real-data experiments validate the predicted performance and demonstrate effective recovery under constrained measurements.

Our contributions are:

- We formulate budgeted distribution recovery from partial measurements with a pretrained reference generator, and develop a generator-aware acquisition principle that converts reference conditional structure into measurement value.
- We connect partial-panel Fisher information to full-distribution KL risk, derive the resulting cost-aware acquisition criterion, and prove that GAMD matches the oracle first-order risk despite the unknown target shift, together with finite-budget guarantees.
- We validate the predicted information-design mechanism and practical effectiveness of GAMD on controlled synthetic problems and real data.

2 Related Work

Partial measurement and optimal design. Matrix-sampling and split-questionnaire designs distribute different subsets of variables across respondents to reduce measurement burden (Shoemaker, 1973; Raghunathan & Grizzle, 1995; Gonzalez & Eltinge, 2007; 2008; 2010). Subsequent work develops efficient estimators and cost-aware questionnaire designs under planned missingness (Chipperfield & Steel, 2009; 2011; Chipperfield et al., 2018), while model-based methods optimize splits using KL loss, Fisher information, or latent-variable uncertainty (Adigüzel & Wedel, 2008; Jang et al., 2022; Zhang et al., 2020), building on classical optimal-design principles (Kiefer, 1959; Chaloner & Verdinelli, 1995). These methods typically assume that the information contributed by each design choice is specified by the statistical model and optimize estimation of selected parameters, moments, or latent representations. Our setting is different: each design choice reveals only a marginal of a high-dimensional shifted distribution, and its information must be induced from the frozen reference conditional law. We use this conditional information to optimize recovery of the entire target distribution rather than a pre-specified low-dimensional estimand.

Active acquisition with generative models. Cost-sensitive and active feature acquisition selects measurements to improve prediction or decision-making under acquisition costs (Greiner et al., 2002; Ji & Carin, 2007; Saar-Tsechansky et al., 2009; Pu et al., 2017; Shim et al., 2018; Huang et al., 2018). Recent methods use latent or generative models to support sequential acquisition (Ma et al., 2018; Li & Oliva, 2021), and arbitrary-conditioning models such as VAEAC and ACFlow provide flexible conditional distributions for partially observed inputs (Ivanov et al., 2018; Li et al., 2020); related work further studies nongreedy and learned acquisition policies (Valancius et al., 2023; Guney et al., 2025). These methods are primarily *instance-wise*: they decide what to observe next for a particular

sample, often for prediction or completion. GAMD is instead *population-level and prospective*: it allocates different panels across independent target individuals to recover a shifted population law.

3 Problem Formulation

Target distribution recovery with partial measurements. Let $X = (X_1, \dots, X_d) \in \mathcal{X}$ denote a complete multivariate draw from an unknown target distribution $Q^\star$. We can sample independent target individuals $X_i \sim Q^\star$, but for each individual we observe only a selected subset of its variables. Let $\mathcal{S} \subseteq 2^{[d]}$ be a finite collection of allowable subsets, which we call *measurement panels*. Measuring panel $S \in \mathcal{S}$ reveals $X_{i,S}$ and incurs cost $c(S) > 0$. Before observing the realization of X_i, the learner chooses a panel $S_i \in \mathcal{S}$ and receives only $Y_i = X_{i,S_i}$. Panel choices may depend on previously collected observations, but not on the unobserved variables of the current individual. The total acquisition cost must satisfy $\sum_i c(S_i) \le B$, where B is the measurement budget. Thus, the budget determines both how many target individuals can be measured and which variables are observed for each of them. Our goal is to design these panel assignments so as to recover the complete target distribution $Q^\star$ as accurately as possible from the resulting partial observations.

Reference population and relative target shift. We assume access to a pretrained and frozen generative model G_0 learned from a fully observed reference population, and denote its induced joint distribution by Q_0. For every admissible panel, G_0 supports full sampling and conditional sampling of X_{S^c} given X_S, where $S^c = [d] \setminus S$. Rather than relearning a high-dimensional distribution from sparse target observations, we retain the joint structure encoded by Q_0 and model only the change from reference to target. We treat G_0 and its induced law Q_0 as fixed; Appendix G.5 analyzes errors in conditional-score computation under this law.

Before target acquisition, we fix a feature map $\phi : \mathcal{X} \to \mathbb{R}^r$ and model the target as an *exponential tilt* of the reference distribution:

$$\frac{dQ_\beta}{dQ_0}(x) = \exp\{\beta^\top \phi(x) - A(\beta)\}, \qquad A(\beta) = \log \mathbb{E}_{Q_0} e^{\beta^\top \phi(X)}, \tag{1}$$

with $Q^\star = Q_{\beta^\star}$ for an unknown $\beta^\star \in \Theta \subset \mathbb{R}^r$, where r is the number of shift features and $A(\beta) < \infty$ on Θ. Here, ϕ specifies the distributional directions along which the target may differ from the reference. It can include domain-relevant statistics, nonlinear interactions, or features learned from reference data and frozen before target acquisition. A richer ϕ enlarges the target-shift family without relearning the reference generator. The exponential-tilt model is a classical semiparametric density-ratio model with an unrestricted baseline distribution (Anderson, 1979; Qin, 1998). It is widely used to represent population reweighting and calibration (Wu & Lu, 2016; Chen et al., 2023). The model requires $Q^\star \ll Q_0$; severe support mismatch is outside our main setting.

Objective. Given the collected partial measurements and the frozen reference model, we output an estimated target distribution $Q_{\widehat{\beta}_B}$. We evaluate recovery using

$$\mathcal{R}_B = \mathbb{E}\Big[D_{\mathrm{KL}}\Big(Q_{\beta^\star} \| Q_{\widehat{\beta}_B}\Big)\Big], \qquad D_{\mathrm{KL}}(P\|Q) := \int \log\Big(\frac{dP}{dQ}\Big) dP. \tag{2}$$

Here $D_{\mathrm{KL}}(P\|Q) = +\infty$ if $P \not\ll Q$. The expectation is over target sampling and algorithmic randomness. For first-order acquisition, standard smooth f-divergences induce the same local information geometry up to scale and hence the same asymptotically optimal allocation (Appendix I.2). We use KL throughout for concreteness and analytical convenience.

4 Algorithm: Generator-Aware Measurement Design

The main difficulty in our problem is circular: the value of each panel depends on the unknown target shift $\beta^\star$, but estimating that shift requires choosing which panels to measure. Moreover, panel value is determined not only by which shift features are observed directly, but also by what unmeasured directions can be inferred through the reference dependence structure. We propose *Generator-Aware Measurement Design (GAMD)*. It resolves both issues by using a small initialization sample to localize the target shift and the frozen reference conditional law to convert this indirect dependence into panel information.

4.1 INITIALIZATION ESTIMATE

Before optimizing the main measurement allocation, GAMD needs a coarse estimate of the unknown target shift $\beta^\star$. The feature map ϕ is fixed in advance, but its mean under the target population, $\mu^\star = \mathbb{E}_{Q_{\beta^\star}}[\phi(X)]$, is unknown. The initialization stage first estimates these target feature moments from partial measurements and then converts them into an initial estimate of $\beta^\star$.

We reserve an initialization budget b_B and use a fixed, target-independent allocation $\{n_{0,S}\}_{S\in\mathcal{S}}$ satisfying $\sum_S n_{0,S}c(S) \le b_B$. The allocation must provide direct coverage of every shift feature. Let $\mathcal{D}_a \subseteq [d]$ denote the variables needed to evaluate $\phi_a(X)$, and let $n_0 = \sum_S n_{0,S}$. The fraction of initialization observations on which ϕ_a can be evaluated is $\rho_{a,B} = \sum_S (n_{0,S}/n_0)\mathbf{1}\{\mathcal{D}_a \subseteq S\}$. We require $\rho_{a,B} \ge \rho_{\min} > 0$ for every a. For example, a balanced allocation over panels covering all feature supports $\{\mathcal{D}_a\}$ satisfies this requirement whenever each feature receives a nonvanishing share of the initialization sample.

Because different features may be observed at different frequencies, we estimate their target moments by inverse-probability weighting: $\widehat{\mu}_a = n_0^{-1}\sum_{i=1}^{n_0} \mathbf{1}\{\mathcal{D}_a \subseteq S_i\}\phi_a(X_i)/\rho_{a,B}$. The factor $1/\rho_{a,B}$ corrects for the observation frequency of ϕ_a, so $\mathbb{E}[\widehat{\mu}_a] = \mathbb{E}_{Q_{\beta^\star}}[\phi_a(X)]$.

The exponential-tilt model links these moments to the shift parameter through $\mathbb{E}_{Q_\beta}[\phi(X)] = \nabla A(\beta)$. Thus $\beta^\star$ is the parameter whose model feature moments match the target moments. Using R_0 independent reference draws $X_0^{(m)} \sim Q_0$, we approximate the log-partition function by $\widehat{A}_{R_0}(\beta) = \log\Big[R_0^{-1}\sum_{m=1}^{R_0}\exp\{\beta^\top\phi(X_0^{(m)})\}\Big]$, where these reference draws do not consume the target measurement budget. We then define

$$\widetilde{\beta}_B = \arg\min_{\beta\in\Theta}\left\{\widehat{A}_{R_0}(\beta) - \beta^\top\widehat{\mu}\right\}. \tag{3}$$

With the exact log-partition function, the first-order condition is $\nabla A(\widetilde{\beta}_B) = \widehat{\mu}$, so this optimization is simply moment matching.

4.2 GENERATOR-AWARE PANEL DESIGN

Given the initialization estimate $\widetilde{\beta}_B$, GAMD uses the frozen generator to evaluate candidate panels and allocate the remaining budget. For a complete observation, the shift score is $s_\beta(X) = \phi(X) - \mu_\beta$, where $\mu_\beta = \mathbb{E}_{Q_\beta}[\phi(X)]$. If only panel S is observed, the available score is $m_{\beta,S}(X_S) = \mathbb{E}_{Q_\beta}[s_\beta(X) \mid X_S]$, and its Fisher information is $I_S(\beta) = \mathrm{Var}_{Q_\beta}[m_{\beta,S}(X_S)]$. Thus, $I_S(\beta)$ measures how much of the full target-shift signal can be recovered from panel S, including information about unmeasured directions revealed through reference dependence. This identity turns the reference conditional law into a measurement-design interface: by querying what the observed coordinates imply about the unobserved part of the full score, GAMD can value a panel even when its useful information is carried indirectly through reference dependence.

Ideally, the design would use $I_S(\beta^\star)$. Since $\beta^\star$ is unknown, GAMD instead uses the plug-in quantity $I_S(\widetilde{\beta}_B)$. This quantity is generally unavailable in closed form, so we estimate it by Monte Carlo. Because $Q_{\widetilde{\beta}_B}$ is an exponential tilt of Q_0, full and conditional draws from $Q_{\widetilde{\beta}_B}$ can be obtained by reweighting or resampling generator draws; see Appendix B. We draw N_g full samples from $Q_{\widetilde{\beta}_B}$. Their feature covariance gives $\widehat{F}$, estimating $F(\widetilde{\beta}_B) = \mathrm{Var}_{Q_{\widetilde{\beta}_B}}[\phi(X)]$, which weights how errors in different shift directions affect full-distribution recovery.

For each outer panel value $X_S = x_S$, we keep the observed coordinates fixed and draw L_g samples of the unobserved variables from $Q_{\widetilde{\beta}_B}(X_{S^c} \mid X_S = x_S)$. Averaging the resulting full-data scores gives a Monte Carlo estimate of $m_{\widetilde{\beta}_B,S}(x_S)$. We repeat this independently twice for the same x_S and use the cross-covariance of the two estimates to avoid counting Monte Carlo noise as panel information.

$$\widehat{I}_S = \Pi_{\mathbb{S}^r_+}\left\{\frac{1}{2}\left[\widehat{\mathrm{Cov}}(\widehat{m}_S^{(1)}, \widehat{m}_S^{(2)}) + \widehat{\mathrm{Cov}}(\widehat{m}_S^{(2)}, \widehat{m}_S^{(1)})\right]\right\}. \tag{4}$$

This estimator targets $I_S(\widetilde{\beta}_B)$, while the PSD projection $\Pi_{\mathbb{S}^r_+}$ ensures a valid information matrix.

Let p_S be the fraction of the remaining budget spent on panel S. Since panel S costs $c(S)$ per observation, allocation p provides estimated information per unit budget $\widehat{M}_B(p) = \sum_S p_S \widehat{I}_S / c(S)$. Its inverse measures uncertainty in estimating the target shift, while $\widehat{F}$ weights that uncertainty by its effect on the full target distribution. GAMD therefore solves

$$\widehat{p}_B \in \arg\min_{p \in \Delta_{\mathcal{S}}} \operatorname{tr}\left[\widehat{F}\, \widehat{M}_B(p)^{-1}\right], \tag{5}$$

with value $+\infty$ if $\widehat{M}_B(p)$ is not positive definite. Finally, with $B_{\text{main}} = B - b_B$, panel S receives approximately $n_{B,S} = \lfloor B_{\text{main}} \widehat{p}_{B,S} / c(S) \rfloor$ target observations.

4.3 Main Acquisition and Target Update

After obtaining $\widehat{p}_B$, GAMD spends the remaining budget according to the resulting panel counts and collects the main target observations. After the design stage, target estimation reduces to partial-data likelihood estimation. We use Fisher scoring; the generator's role is to supply the conditional scores and information matrices that are unavailable in closed form. Let $\mathcal{I}_{\text{main}}$ index these observations, with $Y_i = X_{i,S_i}$.

The main sample is used to refine the initialization estimate by maximizing the likelihood of the observed partial data. For an observation (S_i, Y_i), the score with respect to β is $m_{\beta,S_i}(Y_i)$. Hence, at a current parameter β, the main-sample score and Fisher information are $U_B(\beta) = \sum_{i \in \mathcal{I}_{\text{main}}} m_{\beta,S_i}(Y_i)$ and $J_B(\beta) = \sum_{i \in \mathcal{I}_{\text{main}}} I_{S_i}(\beta)$. The score indicates how the current target model should move, while the information matrix accounts for the curvature and relative scale of different parameter directions.

Neither quantity is generally available in closed form, because both depend on conditional expectations under Q_β. Starting from $\beta^{(0)} = \widetilde{\beta}_B$, GAMD therefore approximates them with the frozen generator at each iteration. At iteration j, we draw N_{upd} full samples from $Q_{\beta^{(j)}}$ to estimate the feature mean and the current panel information. For each observed (S_i, Y_i), we also draw L_{upd} samples of the unobserved variables from $Q_{\beta^{(j)}}(X_{S_i^c} \mid X_{S_i} = Y_i)$ and average their full-data scores to obtain $\widehat{m}_{\beta^{(j)},S_i}(Y_i)$. This gives $\widehat{U}_B^{(j)} = \sum_{i \in \mathcal{I}_{\text{main}}} \widehat{m}_{\beta^{(j)},S_i}(Y_i)$.

Using the same full samples as outer observations, we estimate $\widehat{I}_S^{(j)}$ for the acquired panels with the cross-completion construction in equation (4), and form $\widehat{J}_B^{(j)} = \sum_{i \in \mathcal{I}_{\text{main}}} \widehat{I}_{S_i}^{(j)}$. GAMD then applies the Fisher-scoring update

$$\beta^{(j+1)} = \beta^{(j)} + [\widehat{J}_B^{(j)}]^{-1} \widehat{U}_B^{(j)}, \qquad j = 0, \ldots, J-1. \tag{6}$$

After J updates, we set $\widehat{\beta}_B = \beta^{(J)}$ and return the recovered target distribution $Q_{\widehat{\beta}_B}$.

5 Theoretical Analysis

We analyze the statistical performance of GAMD by first identifying the information in a partial panel, then establishing asymptotic optimality, and finally giving finite-budget guarantees.

Assumption 1 (Regular relative-shift model)**.** *The parameter set $\Theta \subset \mathbb{R}^r$ is compact and convex, $\beta^\star$ is an interior point, and $\|\phi(X)\|_2 \le L_\phi$ Q_0-almost surely. The exponential family in equation* (1) *is minimal, meaning that there is no nonzero $v \in \mathbb{R}^r$ such that $v^\top \phi(X)$ is constant Q_0-almost surely.*

This requires a well-behaved parameter space and nonredundant bounded shift features, as satisfied when $\beta^\star$ lies away from the boundary and ϕ contains no exact linear redundancy.

Assumption 2 (Measurement design and initialization)**.** *The panel family $\mathcal{S}$ is finite, with $0 < c_{\min} \le c(S) \le c_{\max} < \infty$. There exists $p \in \Delta_{\mathcal{S}}$ such that $M_{\beta^\star}(p) \succ 0$, so the available panels collectively reveal every shift direction. The initialization budget satisfies $b_B = B^\gamma$ for some $0 < \gamma < 1$, with $n_0 \asymp b_B$ and $\inf_{a,B} \rho_{a,B} > 0$.*

The initialization sample therefore grows large enough to localize the target shift while using only a vanishing fraction of the total budget. The coverage condition is mild when the panel library covers the variables relevant to ϕ and each feature receives a fixed positive initialization share.

Algorithm 1 Generator-Aware Measurement Design (GAMD)

Require: Frozen G_0; feature map ϕ; panels $\mathcal{S}$ with costs $c(S)$; budget B; initialization budget b_B and allocation $\{n_{0,S}\}$; MC sizes $R_0, (N_g, L_g), (N_{\mathrm{upd}}, L_{\mathrm{upd}})$; updates J.
Ensure: Allocation $\widehat{p}_B$ and target estimate $Q_{\widehat{\beta}_B}$.

1: Collect initialization measurements, estimate $\widehat{\mu}$, and compute $\widetilde{\beta}_B$ via equation (3) using R_0 reference samples.
2: Draw N_g samples from $Q_{\widetilde{\beta}_B}$ and estimate $\widehat{F}$.
3: **for** $S \in \mathcal{S}$ **do**
4: Draw two L_g-completion batches for each outer X_S and compute $\widehat{I}_S$ via equation (4).
5: **end for**
6: Solve equation (5) for $\widehat{p}_B$ and allocate the remaining budget accordingly.
7: Collect the main target sample and set $\beta^{(0)} \leftarrow \widetilde{\beta}_B$.
8: **for** $j = 0, \ldots, J-1$ **do**
9: Draw N_{upd} full samples from $Q_{\beta^{(j)}}$.
10: Estimate $\widehat{U}_B^{(j)}$ using L_{upd} conditional completions per main observation, and estimate $\widehat{J}_B^{(j)}$ from the corresponding panel information.
11: Update $\beta^{(j+1)}$ via equation (6).
12: **end for**
13: **return** $\widehat{p}_B$ and $Q_{\widehat{\beta}_B}$ with $\widehat{\beta}_B = \beta^{(J)}$.

Assumption 3 (Computation accuracy)**.** *The Monte Carlo and numerical approximations used by GAMD are sufficiently accurate so that: (i) the initialization in equation* (3) *retains its statistical convergence rate; (ii) the estimates entering equation* (4) *and equation* (5)*, together with the numerical design solution, consistently recover their population counterparts; and (iii) the Monte Carlo estimates in equation* (6) *make each implemented Fisher-scoring step asymptotically equivalent to its ideal counterpart.*

This only requires computation error to remain below the relevant statistical uncertainty. It can be achieved by increasing the Monte Carlo sample sizes and tightening the optimization tolerance with B; Appendix G.1 gives explicit sufficient rates.

Recall $s_\beta(X) = \phi(X) - \mu_\beta$, where $\mu_\beta = \mathbb{E}_{Q_\beta}[\phi(X)]$. We first justify the panel information used by GAMD.

Proposition 1 (Conditional-score information)**.** *For every $S \in \mathcal{S}$, the score of the panel marginal $Q_{\beta,S}$ is $m_{\beta,S}(X_S) = \mathbb{E}_{Q_\beta}[s_\beta(X) \mid X_S]$, and hence the Fisher information contributed by panel S is $I_S(\beta) = \mathrm{Var}_{Q_\beta}[m_{\beta,S}(X_S)]$.*

A panel can therefore reveal a shift feature indirectly through dependence between its observed and unobserved variables. What is essential for panel valuation is access to this conditional structure, rather than a particular generative architecture. For a fixed feature map, a sufficiently accurate conditional-score predictor can also recover the required information; a conditional generator provides a reusable representation of the reference law that supports such queries. Appendix C formalizes this distinction and shows that covariance and other fixed-order summaries can be insufficient.

For an allocation p, define the population information per unit budget $M_\beta(p) = \sum_{S \in \mathcal{S}} p_S I_S(\beta)/c(S)$ and the corresponding design criterion $\Phi_\beta(p) = \mathrm{tr}[F(\beta) M_\beta(p)^{-1}]$, with $\Phi_\beta(p) = +\infty$ when $M_\beta(p)$ is singular. These are the population counterparts of $\widehat{M}_B(p)$ and the objective in equation (5).

Let $p^\star \in \arg\min_{p \in \Delta_{\mathcal{S}}} \Phi_{\beta^\star}(p)$ be an oracle allocation, and define $\Phi^\star_{\beta^\star} = \Phi_{\beta^\star}(p^\star)$ and $M^\star = M_{\beta^\star}(p^\star)$. The optimal allocation need not be unique, but $M^\star$ is unique.

Theorem 1 (Asymptotic optimality of GAMD)**.** *Suppose Assumptions 1–3 hold. If GAMD uses a fixed number J of Fisher-scoring updates satisfying $2^J \gamma > 1$, then*

$$\Phi_{\beta^\star}(\widehat{p}_B) \xrightarrow{p} \Phi^\star_{\beta^\star}, \qquad \sqrt{B}(\widehat{\beta}_B - \beta^\star) \Rightarrow \mathcal{N}(0, M^{\star -1}).$$

If, in addition, the sequence $\{B\|\widehat{\beta}_B - \beta^\star\|_2^2\}_{B \geq 1}$ is uniformly integrable, then

$$B\, \mathbb{E} D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\widehat{\beta}_B}) \longrightarrow \frac{1}{2} \Phi^\star_{\beta^\star}. \tag{7}$$

The constant $\frac{1}{2}\Phi^\star_{\beta^\star}$ matches the local asymptotic minimax lower bound over the regular prospective acquisition procedures defined in Appendix D.3.

The first two conclusions establish oracle design consistency and root-B efficiency. The additional uniform-integrability condition is needed only to pass from asymptotic normality to the expected-KL limit. Appendix G.3 verifies this condition for a safeguarded finite-sample implementation under the Monte Carlo rates in Appendix G.1. The matching local minimax lower bound is proved in Appendix D.3.

We next quantify finite-budget errors. Let $\epsilon_{\mathrm{init}} = \|\widetilde{\beta}_B - \beta^\star\|$ be the initialization error, $\epsilon_{\mathrm{info}} = \|\widehat{F} - F(\widetilde{\beta}_B)\|_{\mathrm{op}} + \max_{S\in\mathcal{S}} \|\widehat{I}_S - I_S(\widetilde{\beta}_B)\|_{\mathrm{op}}$ the Monte Carlo error in the design quantities, and ϵ_{opt} the numerical suboptimality in solving equation (5). Let $\eta_{\mathrm{upd}}(\delta)$ bound the normalized Monte Carlo errors in $\widehat{U}_B^{(j)}$ and $\widehat{J}_B^{(j)}$ entering equation (6), with probability at least $1 - \delta/2$. Appendix E gives explicit bounds.

Proposition 2 (Finite-budget performance)**.** *Under Assumptions 1–2, there exist constants $C, C' > 0$ such that, for sufficiently small initialization, information-estimation, and optimization errors,*

$$\Phi_{\beta^\star}(\widehat{p}_B) - \Phi^\star_{\beta^\star} \le C(\epsilon_{\mathrm{init}} + \epsilon_{\mathrm{info}} + \epsilon_{\mathrm{opt}}). \tag{8}$$

Conditional on such initialization and design, if $\xi_B(\delta) + \eta_{\mathrm{upd}}(\delta)$ is sufficiently small, then with probability at least $1 - \delta$,

$$D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\beta^{(J)}}) \le C' \left\{ [\xi_B(\delta) + \eta_{\mathrm{upd}}(\delta)]^2 + \epsilon_{\mathrm{init}}^{2^{J+1}} \right\}, \tag{9}$$

where $\xi_B(\delta) = \widetilde{O}(\sqrt{(r + \log(1/\delta))/B_{\mathrm{main}}})$.

The first bound controls the suboptimality of the learned allocation, while the second controls the final distribution-recovery error from target-sample noise, Monte Carlo updates, and initialization error. Appendices G.4, G.5, and I.3 study fixed-completion efficiency, biased conditional scores, and target-model misspecification.

6 Synthetic Experiments

We first study GAMD in a controlled synthetic setting where the data-generating target distribution is available to the experimenter for evaluation, but is not revealed to any acquisition method.

Experimental setting. The reference distribution is constructed from a 16-dimensional four-component Gaussian mixture followed by a nonlinear coordinate transformation. We use $r = 12$ bounded unary and pairwise features to parameterize the reference-to-target shift, and generate synthetic targets by exponentially tilting the reference distribution. The candidate panels are all 120 pairs of coordinates, with equal measurement cost. Thus, each sampled target individual reveals only two of the 16 variables, while the goal is to recover the complete shifted distribution. The exact settings are given in Appendix J.2.

Baselines and evaluation. We compare GAMD with three acquisition strategies. *Uniform SQD* allocates the measurement budget uniformly across all candidate panels. *A-OSQD* uses a covariance-based A-optimal split-questionnaire design. *Discriminative Score OED* (Disc.) directly learns conditional target-shift scores from masked reference observations and constructs a design from the resulting predicted-score information. Formal definitions and implementation details are provided in Appendix J.

We report the normalized KL risk

$$R_{\mathrm{risk}} = \frac{B\, D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\widehat{\beta}_B})}{C^\star}, \qquad C^\star = \frac{1}{2}\Phi_{\beta^\star}(p^\star), \tag{10}$$

where $C^\star$ is the leading risk constant of the population oracle design. Both $\beta^\star$ and $C^\star$ are used only for generating or evaluating the synthetic experiment and are unavailable to the acquisition methods. By Theorem 1, an asymptotically optimal procedure satisfies $\mathbb{E}[R_{\mathrm{risk}}] \to 1$. We report means with one standard error.

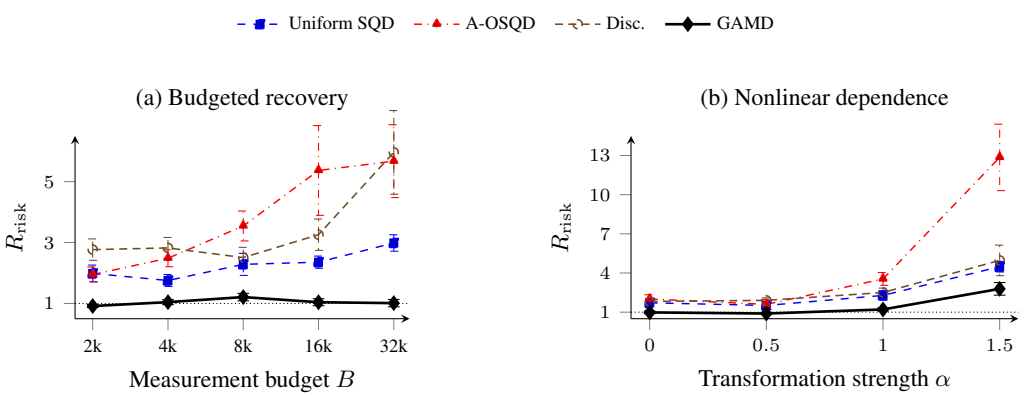

Figure 1: Synthetic distribution recovery. (a) Normalized KL risk across measurement budgets with 20 paired replications per budget. (b) Effect of increasing nonlinear dependence at $B = 8000$; the $\alpha = 1$ point reuses the corresponding 20 replications from panel (a), while the other conditions use 15 paired replications. Error bars show one standard error.

Recovery across measurement budgets. Figure 1(a) varies the total measurement budget over $B \in \{2000, 4000, 8000, 16000, 32000\}$ using 20 paired replications per budget. GAMD achieves the lowest mean normalized KL risk consistently. At the largest budget, its risk is 1.012, close to the first-order oracle reference, compared with 2.986 for Uniform SQD, 5.671 for A-OSQD, and 5.960 for Disc. Because the normalization removes the common $1/B$ scale, the remaining gap reflects differences in information collected per unit budget. The results therefore show that learning which panels are informative about the target shift can substantially improve full-distribution recovery.

Role of conditional information. We next isolate the role of the conditional information used to value panels. The synthetic reference is obtained from an underlying Gaussian mixture through the coordinate-wise transformation

$$x_j = \frac{\sinh(\alpha z_j)}{\alpha},$$

with $x_j = z_j$ at $\alpha = 0$, so α controls the strength of the additional coordinate-wise nonlinear transformation.

At fixed budget $B = 8000$, we vary $\alpha \in \{0, 0.5, 1, 1.5\}$. As α increases, covariance becomes a less complete description of the dependence relevant to panel selection, while the full conditional law remains informative. Figure 1(b) shows that GAMD has the lowest mean risk throughout. At $\alpha = 1.5$, the problem becomes harder for all methods, but A-OSQD deteriorates much more sharply, reaching 12.854 compared with 2.787 for GAMD. This supports the value of modeling conditional structure beyond covariance.

7 Real-Data Experiments

We further evaluate GAMD on the UCI Gas Sensor Array Drift dataset (Vergara et al., 2012), which contains 13,910 observations from 16 chemical sensors collected over ten temporal batches. We use batches 1–6 as the fully observed reference population, train a frozen VAEAC conditional generator (Ivanov et al., 2018), and define each candidate panel as a pair of sensors. Preprocessing, feature construction, generator structure, and other implementation details are given in Appendix J.3.

Since the real target density is unavailable, we evaluate recovery within a fixed empirical exponential-tilt family. Given evaluation features $\{\widetilde{\phi}_m\}_{m=1}^M$, define $\widehat{A}(\beta) = \log\left[\frac{1}{M}\sum_{m=1}^M e^{\beta^\top \widetilde{\phi}_m}\right]$. We report

$$D_{\widehat{A}}(\widehat{\beta}_B, \beta^\dagger) = \widehat{A}(\widehat{\beta}_B) - \widehat{A}(\beta^\dagger) - (\widehat{\beta}_B - \beta^\dagger)^\top \nabla \widehat{A}(\beta^\dagger),$$

the corresponding within-family KL divergence. Here $\beta^\dagger$ is the generating parameter for controlled shifts and the complete-data family projection for natural drift. The construction of the fixed evaluation feature pool is given in Appendix J.3.

Controlled target shift. We construct target populations by exponentially reweighting held-out reference observations, using the real sensor joint distribution as the reference while producing a controlled exponential target shift. Table 1 reports results over four measurement budgets with 30 paired replications per budget. GAMD achieves the lowest mean loss throughout. The gains are modest but consistent across budgets, showing that conditional information learned from the reference population can improve panel allocation under real high-dimensional dependence.

B	Uniform SQD	A-OSQD	Disc.	GAMD
400	0.1166 ± 0.0041	0.1185 ± 0.0040	0.1160 ± 0.0040	$\mathbf{0.1147 \pm 0.0042}$
800	0.1060 ± 0.0038	0.1085 ± 0.0035	0.1062 ± 0.0036	$\mathbf{0.1045 \pm 0.0038}$
1600	0.1066 ± 0.0020	0.1087 ± 0.0020	0.1057 ± 0.0021	$\mathbf{0.1033 \pm 0.0026}$
3200	0.1071 ± 0.0023	0.1093 ± 0.0022	0.1081 ± 0.0024	$\mathbf{0.1063 \pm 0.0028}$

Table 1: Controlled target shift on the Gas Sensor Array Drift dataset. Entries are mean $D_{\widehat{A}} \pm$ one standard error; lower is better.

Natural temporal drift. We next use later collection periods directly as target populations, taking batch 7 and the pooled batches 8–9 as two target campaigns. These targets reflect the temporal sensor drift present in the original data and need not correspond to an exact exponential tilt of the reference population. Figure 2 reports the recovery loss over three measurement budgets.

GAMD achieves the lowest mean loss for both target campaigns at every budget. For batch 7 at $B = 3200$, it reduces the loss from 0.0506 under Uniform SQD to 0.0374. The separation is larger for batches 8–9, where GAMD obtains 0.1095 at $B = 3200$, compared with 0.1288 for Uniform SQD, 0.1194 for A-OSQD, and 0.1218 for Disc.

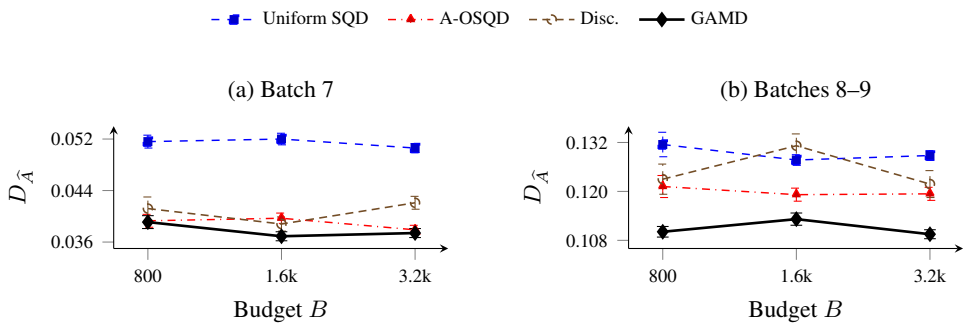

Figure 2: Recovery under natural temporal drift. Targets are taken directly from (a) batch 7 and (b) batches 8–9. Curves show mean $D_{\widehat{A}}$ with one standard error.

The gains are larger under natural drift, where the target–reference shift is more pronounced and the value of different sensor panels is less uniform. This setting leaves greater scope for generator-aware allocation to improve which measurements are collected.

8 Conclusion

We studied budgeted recovery of high-dimensional target distributions when each sampled individual can be observed only through a costly subset of variables. GAMD combines a frozen reference generator with a low-dimensional exponential-tilt model to evaluate the information carried by candidate measurement panels and allocate the acquisition budget accordingly. We established oracle design consistency, asymptotic efficiency, and the corresponding optimal KL-risk constant, together with finite-budget guarantees. Synthetic and real-data experiments show that generator-aware allocation improves distribution recovery, with particularly clear gains when useful target-shift information is carried through nonlinear or heterogeneous conditional dependence.

AI USE STATEMENT

In this work, we used generative AI tools to help develop conceptual frameworks, design and provide feedback on research methodology and experiments, and assist with method implementation. We also used generative AI tools to draft parts of the paper, summarize and analyze existing literature, and edit the manuscript to improve readability. We did not use generative AI tools for other tasks requiring disclosure. We reviewed all AI-assisted work. Implementation suggestions were tested through our experimental pipeline, and literature-based statements were verified against the cited sources. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We have aimed to make both the theoretical and empirical results reproducible. All assumptions used in the main theoretical results are stated explicitly in Section 5, with complete proofs, supporting lemmas, and sufficient Monte Carlo conditions provided in the appendices. The experimental protocols are documented in Appendix J, including the acquisition baselines, initialization and optimization procedures, synthetic data construction, numerical integration checks, and the preprocessing and generator setup for the Gas Sensor Array Drift experiments. We report the main evaluation settings, replication counts, and numerical procedures used for the reported figures and tables.

REFERENCES

Feray Adigüzel and Michel Wedel. Split questionnaire design for massive surveys. *Journal of Marketing Research*, 45(5):608–617, 2008.

JA Anderson. Multivariate logistic compounds. *Biometrika*, 66(1):17–26, 1979.

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

Rui Chen, Guanhua Chen, and Menggang Yu. Entropy balancing for causal generalization with target sample summary information. *Biometrics*, 79(4):3179–3190, 2023.

James O Chipperfield and David G Steel. Design and estimation for split questionnaire surveys. *Journal of Official Statistics*, 25(2):227, 2009.

James O Chipperfield and David G Steel. Efficiency of split questionnaire surveys. *Journal of Statistical Planning and Inference*, 141(5):1925–1932, 2011.

James O Chipperfield, Margo L Barr, and David G Steel. Split questionnaire designs: collecting only the data that you need through mcar and mar designs. *Journal of Applied Statistics*, 45(8): 1465–1475, 2018.

Jeffrey M Gonzalez and John L Eltinge. Multiple matrix sampling: A review. In *Proceedings of the Section on Survey Research Methods, American Statistical Association*, pp. 3069–3075. American Statistical Association Alexandria, VA, 2007.

Jeffrey M Gonzalez and John L Eltinge. Adaptive matrix sampling for the consumer expenditure quarterly interview survey. In *Proceedings of the Section on Survey Research Methods, American Statistical Association*, pp. 2081–2088. sn, 2008.

Jeffrey M Gonzalez and John L Eltinge. Optimal survey design: A review. *Section on Survey Research Methods–JSM*, 2010.

Russell Greiner, Adam J Grove, and Dan Roth. Learning cost-sensitive active classifiers. *Artificial Intelligence*, 139(2):137–174, 2002.

Osman Berke Guney, Ketan Suhaas Saichandran, Karim Elzokm, Ziming Zhang, and Vijaya B Kolachalama. Active feature acquisition via explainability-driven ranking. *Proceedings of machine learning research*, 267:20748, 2025.

Sheng-Jun Huang, Miao Xu, Ming-Kun Xie, Masashi Sugiyama, Gang Niu, and Songcan Chen. Active feature acquisition with supervised matrix completion. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1571–1579, 2018.

Oleg Ivanov, Michael Figurnov, and Dmitry Vetrov. Variational autoencoder with arbitrary conditioning. *arXiv preprint arXiv:1806.02382*, 2018.

Dae-Gyu Jang, Zhengyuan Zhu, and Cindy Yu. A-optimal split questionnaire designs for multivariate continuous variables. *arXiv preprint arXiv:2203.04883*, 2022.

Shihao Ji and Lawrence Carin. Cost-sensitive feature acquisition and classification. *Pattern Recognition*, 40(5):1474–1485, 2007.

Jack Kiefer. Optimum experimental designs. *Journal of the Royal Statistical Society: Series B (Methodological)*, 21(2):272–304, 1959.

Louis B Kuemmerle, Malte D Luecken, Alexandra B Firsova, Lisa Barros de Andrade e Sousa, Lena Straßer, Ilhem Isra Mekki, Francesco Campi, Lukas Heumos, Maiia Shulman, Valentina Beliaeva, et al. Probe set selection for targeted spatial transcriptomics. *Nature methods*, 21(12):2260–2270, 2024.

Lucien Le Cam and Grace Lo Yang. *Asymptotics in statistics: some basic concepts*. Springer Science & Business Media, 2012.

Yang Li and Junier Oliva. Active feature acquisition with generative surrogate models. In *International conference on machine learning*, pp. 6450–6459. PMLR, 2021.

Yang Li, Shoaib Akbar, and Junier Oliva. Acflow: Flow models for arbitrary conditional likelihoods. In *International Conference on Machine Learning*, pp. 5831–5841. PMLR, 2020.

Chao Ma, Sebastian Tschiatschek, Konstantina Palla, José Miguel Hernández-Lobato, Sebastian Nowozin, and Cheng Zhang. Eddi: Efficient dynamic discovery of high-value information with partial vae. *arXiv preprint arXiv:1809.11142*, 2018.

Yewen Pu, Leslie P Kaelbling, and Armando Solar-Lezama. Learning to acquire information. *arXiv preprint arXiv:1704.06131*, 2017.

Jing Qin. Inferences for case-control and semiparametric two-sample density ratio models. *Biometrika*, 85(3):619–630, 1998.

Trivellore E Raghunathan and James E Grizzle. A split questionnaire survey design. *Journal of the American Statistical Association*, 90(429):54–63, 1995.

Maytal Saar-Tsechansky, Prem Melville, and Foster Provost. Active feature-value acquisition. *Management Science*, 55(4):664–684, 2009.

Hajin Shim, Sung Ju Hwang, and Eunho Yang. Joint active feature acquisition and classification with variable-size set encoding. *Advances in neural information processing systems*, 31, 2018.

David M Shoemaker. *Principles and procedures of multiple matrix sampling.* Ballinger, 1973.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

Michael Valancius, Max Lennon, and Junier Oliva. Acquisition conditioned oracle for nongreedy active feature acquisition. *arXiv preprint arXiv:2302.13960*, 2023.

Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Alexander Vergara, Shankar Vembu, Tuba Ayhan, Margaret A Ryan, Margie L Homer, and Ramón Huerta. Chemical gas sensor drift compensation using classifier ensembles. *Sensors and Actuators B: Chemical*, 166:320–329, 2012.

Changbao Wu and Wilson W Lu. Calibration weighting methods for complex surveys. *International Statistical Review*, 84(1):79–98, 2016.

Chelsea Zhang, Sean J Taylor, Curtiss Cobb, and Jasjeet Sekhon. Active matrix factorization for surveys. *The Annals of Applied Statistics*, 14(3):1182–1206, 2020.

A Technical Setup and Regularity Consequences

We first establish the regularity properties used in the proofs. Under the bounded-feature exponential-tilt model, the feature mean varies smoothly with the target parameter, and the full Fisher information is positive definite. The same boundedness assumption controls derivatives of the panel likelihoods. These facts justify moment-based initialization, information estimation, and the local expansion of the final updates.

Let $\mathcal{X} = \prod_{j=1}^{d} \mathcal{X}_j$ be a standard Borel space, and write $\mathcal{X}_S = \prod_{j\in S} \mathcal{X}_j$. The dimensions d, r and finite panel library $\mathcal{S}$ are fixed throughout the asymptotic analysis. We write $\mathbb{E}_\beta = \mathbb{E}_{Q_\beta}$ for expectations of one complete observation, and $\mathbb{P}$ for probability over target data and generator simulation. A conditional law $Q_\beta(\cdot \mid S, y)$ means $Q_\beta(\cdot \mid X_S = y)$, with the observed coordinates held fixed. The notation I_r denotes the $r \times r$ identity matrix; $\|\cdot\|_*$ denotes the nuclear norm (sum of singular values), and $\lambda_{\min}$ the smallest eigenvalue of a symmetric matrix. Constants denoted by C may change between displays and depend on fixed problem parameters but not on the target budget. For $S \subseteq [d]$, write $Q_{\beta,S}$ for the marginal law of X_S under Q_β. All conditional expectations are taken in versions that are measurable in the conditioning variable. At a fixed β, $L_2(X_S)$ denotes the square-integrable functions under $Q_{\beta,S}$; unqualified expectations in the score identities are under Q_β. The symbol † on a matrix denotes its Moore–Penrose pseudoinverse. We use $\mu_\beta = \mathbb{E}_{Q_\beta}[\phi(X)]$ and $F(\beta) = \mathrm{Var}_{Q_\beta}[\phi(X)]$.

Lemma 1 (Smoothness of the exponential tilt)**.** *Under Assumption 1, A is infinitely differentiable on $\mathbb{R}^r$ and*

$$\nabla A(\beta) = \mu_\beta, \qquad \nabla^2 A(\beta) = F(\beta).$$

Moreover, on every compact parameter set all derivatives of A are uniformly bounded. Minimality implies $F(\beta) \succ 0$ for every β, and hence there exist constants $0 < f_{\min} \le f_{\max} < \infty$ such that $f_{\min} I_r \preceq F(\beta) \preceq f_{\max} I_r$ for all $\beta \in \Theta$.

Proof. Because $\|\phi(X)\|_2 \le L_\phi$, the random variable $e^{\beta^\top \phi(X)}$ is bounded above and below by finite positive constants on every compact set of β. Differentiation under the expectation is therefore justified by dominated convergence. Writing $Z(\beta) = \mathbb{E}_{Q_0} e^{\beta^\top \phi(X)}$, direct differentiation gives

$$\nabla A(\beta) = \frac{\mathbb{E}_{Q_0}[\phi(X) e^{\beta^\top \phi(X)}]}{Z(\beta)} = \mathbb{E}_{Q_\beta}[\phi(X)] = \mu_\beta.$$

Differentiating once more yields

$$\nabla^2 A(\beta) = \mathbb{E}_{Q_\beta}[\phi(X)\phi(X)^\top] - \mu_\beta \mu_\beta^\top = F(\beta).$$

Higher derivatives are joint cumulants of bounded coordinates of $\phi(X)$ and are therefore uniformly bounded on compact parameter sets.

For any nonzero $v \in \mathbb{R}^r$, $v^\top F(\beta) v = \mathrm{Var}_{Q_\beta}(v^\top \phi(X))$. The measures Q_β and Q_0 are mutually absolutely continuous because their density ratio is strictly positive. Thus, if this variance were zero, $v^\top \phi(X)$ would be constant Q_0-almost surely, contradicting minimality. Hence $F(\beta)$ is positive definite. Continuity of F and compactness of Θ then give the uniform eigenvalue bounds. In addition, bounded features give the global upper bound $F(\beta) \preceq L_\phi^2 I_r$ for every $\beta \in \mathbb{R}^r$: for any v, $v^\top F(\beta) v \le \mathbb{E}_\beta (v^\top \phi)^2 \le L_\phi^2 \|v\|_2^2$. This bound will control the KL loss even on rare events where a scoring iterate leaves the local neighborhood. □

The positive lower bound on F makes the feature mean informative about the parameter. For panel allocation, we also need a small initialization error to produce only a small change in $I_S(\beta)$. The next lemma gives a common Lipschitz bound for all candidate panels.

Lemma 2 (Uniform local Lipschitz properties)**.** *There is a finite constant L_{loc} such that, for all $\beta, \beta' \in \Theta$ and all $S \in \mathcal{S}$,*

$$\|F(\beta) - F(\beta')\|_{\mathrm{op}} + \|I_S(\beta) - I_S(\beta')\|_{\mathrm{op}} \le L_{\mathrm{loc}} \|\beta - \beta'\|_2 .$$

The panel scores, their first two parameter derivatives, and the first three derivatives of the panel log densities are uniformly bounded.

Proof. The claim for F follows from Lemma 1 and the uniform bound on the third derivative of A. The corresponding panel statements are proved in Appendix C: the panel log density is the difference of an unconditional and a conditional log-partition function. Both are cumulant-generating functions of a bounded random vector, so their first three derivatives are uniformly bounded. To also account for the parameter-dependent marginal law, differentiate $I_S(\beta) = \mathbb{E}_\beta[m_{\beta,S} m_{\beta,S}^\top]$. Its derivative contains the two bounded derivatives of $m_{\beta,S}$ and the density-derivative term $\mathbb{E}_\beta[m_{\beta,S} m_{\beta,S}^\top m_{\beta,S,k}]$, where $m_{\beta,S,k}$ is the kth score component. All are uniformly bounded, so the mean-value theorem gives the Lipschitz bound for I_S. Because $\mathcal{S}$ is finite, a common constant can be used for all panels. □

Thus a small error in the feature mean gives a small parameter error, and a small parameter error gives a small change in panel information. Individual panels may still have singular information matrices. Assumption 2 requires only that a mixture of the panels provide information in every parameter direction.

B Tilted Sampling from the Frozen Generator

The frozen generator samples from Q_0, but GAMD needs expectations under the current target model Q_β. We show how exponential reweighting gives both the full target law and its conditional distributions. A rejection sampler then supplies exact draws without evaluating the reference density; its acceptance probability describes the cost of a large reference-to-target change.

Proposition 3 (Tilted full and conditional laws)**.** *For any integrable function h,*

$$\mathbb{E}_{Q_\beta}[h(X)] = \frac{\mathbb{E}_{Q_0}[h(X) e^{\beta^\top \phi(X)}]}{\mathbb{E}_{Q_0}[e^{\beta^\top \phi(X)}]}. \tag{11}$$

For every panel S and $Q_{0,S}$-almost every x_S,

$$Q_\beta(dx_{S^c} \mid x_S) = \frac{e^{\beta^\top \phi(x_S, x_{S^c})} Q_0(dx_{S^c} \mid x_S)}{\mathbb{E}_{Q_0}[e^{\beta^\top \phi(X)} \mid X_S = x_S]}. \tag{12}$$

Proof. Equation (11) follows immediately by substituting the Radon–Nikodym derivative in equation (1):

$$\mathbb{E}_{Q_\beta}[h(X)] = e^{-A(\beta)} \mathbb{E}_{Q_0}[h(X) e^{\beta^\top \phi(X)}],$$

and $e^{A(\beta)} = \mathbb{E}_{Q_0} e^{\beta^\top \phi(X)}$.

For the conditional statement, let $C \subseteq \mathcal{X}_{S^c}$ be measurable. Bayes' formula and equation (1) give

$$\begin{aligned} Q_\beta(X_{S^c} \in C \mid X_S = x_S) &= \frac{\int_C e^{\beta^\top \phi(x_S, z) - A(\beta)} Q_0(dz \mid x_S)}{\int e^{\beta^\top \phi(x_S, z) - A(\beta)} Q_0(dz \mid x_S)} \\ &= \frac{\int_C e^{\beta^\top \phi(x_S, z)} Q_0(dz \mid x_S)}{\mathbb{E}_{Q_0}[e^{\beta^\top \phi(X)} \mid X_S = x_S]}. \end{aligned}$$

This is exactly equation (12). □

The normalizing constant in the conditional identity depends on the observed panel, but the weight itself remains $e^{\beta^\top \phi(X)}$. Consequently, a bound on this weight yields one sampling construction for both interfaces. This is useful here because the information identities require conditional draws from a specified joint model, rather than an independently fitted collection of panel conditionals.

Proposition 4 (Exact rejection sampler)**.** *Fix β and write $E_\beta(x) = \beta^\top \phi(x)$. Suppose $\underline{E}_\beta \le E_\beta(x) \le \overline{E}_\beta$ for all x in the support of Q_0. If $X \sim Q_0$ is accepted with probability $\exp\{E_\beta(X) - \overline{E}_\beta\}$, then the accepted sample has law Q_β. The acceptance probability is at least $e^{\underline{E}_\beta - \overline{E}_\beta}$. The same construction applied to proposals from $Q_0(\cdot \mid X_S = x_S)$ produces exact draws from $Q_\beta(\cdot \mid X_S = x_S)$.*

Proof. The joint probability that a proposal falls in a measurable set C and is accepted equals

$$\int_C e^{E_\beta(x)-\overline{E}_\beta} Q_0(dx).$$

After division by the total acceptance probability, the accepted density relative to Q_0 is proportional to $e^{E_\beta(x)}$, and is therefore Q_β by equation (1). Moreover,

$$\mathbb{P}(\text{accept}) = e^{-\overline{E}_\beta}\mathbb{E}_{Q_0} e^{E_\beta(X)} \geq e^{-\overline{E}_\beta} e^{\underline{E}_\beta}.$$

Conditioning throughout on $X_S = x_S$ gives the conditional version. Assumption 1 guarantees bounded E_β uniformly on compact Θ. □

Remark 1 (Overlap diagnostic)**.** *Let $w_\beta = dQ_\beta/dQ_0$ and $\chi^2(P\|Q) = \int (dP/dQ - 1)^2 dQ$. Then $\mathbb{E}_{Q_0} w_\beta = 1$ and*

$$\mathbb{E}_{Q_0} w_\beta^2 = 1 + \chi^2(Q_\beta\|Q_0).$$

Thus the usual importance-sampling effective sample size is controlled by the χ^2 overlap between target and reference. Large weight dispersion signals poor overlap for the proposed tilt, even when the two laws have the same support. This affects the number of reference proposals needed to evaluate a panel, separately from the target-measurement cost $c(S)$.

The rejection construction provides independent exact draws. Importance resampling, sequential Monte Carlo, and architecture-specific conditional samplers provide alternative computational interfaces. Their accuracy is measured by the score and information errors appearing in the finite-budget analysis.

Exact conditional draws make each completion average unbiased for the required conditional mean. Finite-pool resampling and self-normalized importance weighting may also introduce conditional-mean bias. Cross-completion removes the additive contribution of independent simulation noise; the score perturbation bounds in Proposition 11 handle error in that conditional mean. Keeping these two effects separate will also explain why a fixed completion count is adequate for consistent panel-information estimation but can change final-estimator efficiency.

C Conditional Scores and Panel Information

This section proves Proposition 1. For an observation restricted to coordinates S, the relevant likelihood is the marginal law $Q_{\beta,S}$. Differentiating that likelihood gives the conditional score used by GAMD. We also compare the information in nested panels.

To determine what the measured coordinates reveal about β, we first integrate the full density ratio over the unobserved coordinates. For $S \in \mathcal{S}$, define

$$a_{\beta,S}(x_S) = \mathbb{E}_{Q_0}[e^{\beta^\top \phi(X)} \mid X_S = x_S].$$

The density of $Q_{\beta,S}$ relative to $Q_{0,S}$ is

$$r_{\beta,S}(x_S) = e^{-A(\beta)} a_{\beta,S}(x_S). \tag{13}$$

Lemma 3 (Observed-score identity)**.** *For every $S \in \mathcal{S}$,*

$$\nabla_\beta \log r_{\beta,S}(x_S) = \mathbb{E}_{Q_\beta}[s_\beta(X) \mid X_S = x_S] = m_{\beta,S}(x_S).$$

Furthermore,

$$-\mathbb{E}_{Q_\beta} \nabla_\beta^2 \log r_{\beta,S}(X_S) = I_S(\beta).$$

Proof. Taking logarithms in equation (13) gives $\log r_{\beta,S}(x_S) = \log a_{\beta,S}(x_S) - A(\beta)$. Because ϕ is bounded, differentiation can be moved through the conditional expectation. Thus,

$$\begin{aligned}\nabla_\beta \log a_{\beta,S}(x_S) &= \frac{\mathbb{E}_{Q_0}[\phi(X) e^{\beta^\top \phi(X)} \mid X_S = x_S]}{\mathbb{E}_{Q_0}[e^{\beta^\top \phi(X)} \mid X_S = x_S]} \\ &= \mathbb{E}_{Q_\beta}[\phi(X) \mid X_S = x_S],\end{aligned}$$

where the second equality uses Proposition 3. Subtracting $\nabla A(\beta) = \mu_\beta$ proves the score identity.

Differentiating again yields

$$\nabla_\beta^2 \log r_{\beta,S}(x_S) = \mathrm{Var}_{Q_\beta}(\phi(X) \mid X_S = x_S) - F(\beta).$$

Taking expectation and applying the law of total variance gives

$$\begin{aligned} -\mathbb{E}_{Q_\beta} \nabla_\beta^2 \log r_{\beta,S}(X_S) &= F(\beta) - \mathbb{E}_{Q_\beta} \mathrm{Var}_{Q_\beta}(\phi(X) \mid X_S) \\ &= \mathrm{Var}_{Q_\beta}(\mathbb{E}_{Q_\beta}[\phi(X) \mid X_S]) \\ &= \mathrm{Var}_{Q_\beta}(m_{\beta,S}(X_S)) = I_S(\beta). \end{aligned}$$

□

Proof of Proposition 1. Lemma 3 identifies the score of $Q_{\beta,S}$ as $m_{\beta,S}$. Its mean is zero by the tower property, so its second moment equals its covariance $I_S(\beta)$, the Fisher information of the panel marginal. □

The difference between full and panel information is the score variation left unexplained by the observed variables. The law of total variance gives this difference and compares nested panels.

Corollary 1 (Missing-information decomposition and monotonicity)**.** *For every panel S,*

$$F(\beta) = I_S(\beta) + \mathbb{E}_{Q_\beta} \mathrm{Var}_{Q_\beta}(s_\beta(X) \mid X_S). \tag{14}$$

Hence $0 \preceq I_S(\beta) \preceq F(\beta)$. If $S \subseteq T$, then

$$I_T(\beta) - I_S(\beta) = \mathbb{E}_{Q_\beta}\left[(m_{\beta,T} - m_{\beta,S})(m_{\beta,T} - m_{\beta,S})^\top\right] \succeq 0. \tag{15}$$

Proof. Equation (14) is the vector law of total variance applied to $s_\beta(X)$. For nested panels, the tower property gives $m_{\beta,S} = \mathbb{E}[m_{\beta,T} \mid X_S]$. Therefore $m_{\beta,T} - m_{\beta,S}$ is orthogonal in L_2 to every X_S-measurable vector, in particular to $m_{\beta,S}$. Expanding the covariance of $m_{\beta,T} = m_{\beta,S} + (m_{\beta,T} - m_{\beta,S})$ gives equation (15). □

In particular, enlarging a panel can only increase its Fisher information before costs are taken into account. Cost-aware design must balance this gain against the extra cost of acquiring the larger panel. The decomposition also identifies the conditional variance that will appear as simulation noise when the conditional score is approximated by completions.

For the asymptotic likelihood analysis, the score must describe not only a formal derivative but also the first-order change of the probability law. Quadratic-mean differentiability makes this statement precise: the square root of the marginal density has an L_2 expansion whose linear term is the panel score.

Lemma 4 (Quadratic-mean differentiability of panel experiments)**.** *For every $S \in \mathcal{S}$, the family $\{Q_{\beta,S} : \beta \in \Theta\}$ is differentiable in quadratic mean, with score $m_{\beta,S}$. More precisely, uniformly for β in compact subsets of the interior of Θ,*

$$\int \left[\sqrt{r_{\beta+t,S}} - \sqrt{r_{\beta,S}} - \frac{1}{2} t^\top m_{\beta,S} \sqrt{r_{\beta,S}}\right]^2 dQ_{0,S} = O(\|t\|_2^4).$$

Proof. Let $g_\beta = \sqrt{r_{\beta,S}}$. By Lemma 3, $\nabla g_\beta = \frac{1}{2} m_{\beta,S} g_\beta$. For a direction u, a second differentiation gives

$$D^2 g_\beta[u,u] = g_\beta \left\{ \frac{1}{4}(u^\top m_{\beta,S})^2 + \frac{1}{2} u^\top \nabla_\beta m_{\beta,S} u \right\}.$$

Both $m_{\beta,S}$ and its derivative are uniformly bounded by Lemma 2, so $\left\|D^2 g_\beta[u,u]\right\|_{L_2(Q_{0,S})} \leq C\|u\|_2^2$ because $\int g_\beta^2 dQ_{0,S} = 1$. The integral form of Taylor's theorem therefore implies

$$\|g_{\beta+t} - g_\beta - D g_\beta[t]\|_{L_2(Q_{0,S})} \leq C\|t\|_2^2.$$

Squaring both sides proves the claim. □

Every panel marginal therefore has a valid likelihood expansion with score $m_{\beta,S}$. We can now add these contributions over the measurements selected by GAMD and analyze the resulting estimator.

D Proof of Asymptotic Optimality

We prove Theorem 1 in three steps. First, the initialization estimate and simulated information matrices recover the oracle information matrix. Second, the main observations and a fixed number of scoring updates give the root-B Gaussian limit; the expected-KL conclusion then follows under the uniform-integrability condition stated in the theorem. Third, a local minimax lower bound shows that no regular information-stable acquisition strategy improves the resulting first-order KL constant.

D.1 Initialization Accuracy and Design Consistency

The design is evaluated at $\widetilde{\beta}_B$ rather than the unknown $\beta^\star$. We first bound this initialization error. The initialization panel counts are fixed in advance, so every feature is averaged over a predetermined set of independent target observations. The coverage weights correct for the different numbers of observations available for different features.

For initialization unit i, define the vector $Z_i \in \mathbb{R}^r$ coordinatewise by

$$(Z_i)_a = \frac{\mathbf{1}\{\mathcal{D}_a \subseteq S_i\}}{\rho_{a,B}} \phi_a(X_i).$$

Then $\widehat{\mu} = n_0^{-1} \sum_{i=1}^{n_0} Z_i$. The panel sequence is predetermined, whereas the X_i are independent draws from $Q_{\beta^\star}$.

Lemma 5 (Initialization accuracy from partial feature moments)**.** *Under Assumptions 1–2, $\mathbb{E}\widehat{\mu} = \mu_{\beta^\star}$. Moreover, there is a constant C_μ depending only on L_ϕ and $\inf_{a,B} \rho_{a,B}$ such that, with probability at least $1-\delta$,*

$$\|\widehat{\mu} - \mu_{\beta^\star}\|_2 \le C_\mu \sqrt{\frac{r + \log(1/\delta)}{n_0}}. \tag{16}$$

The same rate holds in every fixed moment order.

Let $\widetilde{\beta}_B^{\mathrm{ex}} = \arg\min_{\beta\in\Theta}\{A(\beta) - \beta^\top \widehat{\mu}\}$ denote the exact-A comparator for the implemented initialization in equation (3)*. For all sufficiently large B, with probability tending to one, $\widetilde{\beta}_B^{\mathrm{ex}}$ is an interior solution and*

$$\left\|\widetilde{\beta}_B^{\mathrm{ex}} - \beta^\star\right\|_2 \le \frac{2}{f_{\min}} \|\widehat{\mu} - \mu_{\beta^\star}\|_2 . \tag{17}$$

Consequently, $\left\|\widetilde{\beta}_B^{\mathrm{ex}} - \beta^\star\right\| = O_p(n_0^{-1/2}) = O_p(B^{-\gamma/2})$. For the implemented estimator $\widetilde{\beta}_B$ in equation (3)*, suppose $\widehat{A} = \widehat{A}_{R_0}$ satisfies $\sup_{\beta\in\Theta} \left\|\nabla\widehat{A}(\beta) - \nabla A(\beta)\right\|_2 \le \epsilon_A$ and $\nabla^2\widehat{A}(\beta) \succeq (f_{\min}/2) I_r$ for all $\beta \in \Theta$. Then*

$$\left\|\widetilde{\beta}_B - \beta^\star\right\|_2 \le \frac{4}{f_{\min}} \left(\|\widehat{\mu} - \mu_{\beta^\star}\|_2 + \epsilon_A\right). \tag{18}$$

Proof. For coordinate a,

$$\begin{aligned}\mathbb{E}\widehat{\mu}_a &= \frac{1}{n_0} \sum_{i=1}^{n_0} \frac{\mathbf{1}\{\mathcal{D}_a \subseteq S_i\}}{\rho_{a,B}} \mathbb{E}_{Q_{\beta^\star}} \phi_a(X) \\ &= \frac{\sum_S n_{0,S} \mathbf{1}\{\mathcal{D}_a \subseteq S\}}{n_0 \rho_{a,B}} \mu_{\beta^\star,a} = \mu_{\beta^\star,a}.\end{aligned}$$

Thus the vector estimator is unbiased.

Let $\rho_{\min} = \inf_{a,B} \rho_{a,B} > 0$. Since $\|\phi(X)\|_2 \le L_\phi$, we have $\|Z_i\|_2 \le L_\phi/\rho_{\min}$. The centered vectors $Z_i - \mathbb{E}Z_i$ are independent, mean zero, and uniformly bounded by $2L_\phi/\rho_{\min}$. A standard $1/2$-net argument on the Euclidean unit sphere, followed by scalar Hoeffding bounds and a union bound over a net of size at most 5^r, gives

$$\mathbb{P}\left(\|\widehat{\mu} - \mu_{\beta^\star}\|_2 > t\right) \le 2 \cdot 5^r \exp(-c n_0 t^2 / C_\mu^2)$$

for a universal $c > 0$. Solving for t proves equation (16) after changing constants. Integrating the same sub-Gaussian tail gives the fixed-order moment bound.

To convert the feature-mean error into parameter error, use the positive lower bound on the Hessian of A. Strong convexity makes the moment-matching solution stable.

Define $G_B(\beta) = A(\beta) - \beta^\top \widehat{\mu}$. By Lemma 1, G_B is $f_{\min}$-strongly convex on Θ. Since $\widetilde{\beta}_B^{\text{ex}}$ minimizes G_B, $G_B(\widetilde{\beta}_B^{\text{ex}}) \le G_B(\beta^\star)$. Strong convexity at $\beta^\star$ gives

$$G_B(\widetilde{\beta}_B^{\text{ex}}) \ge G_B(\beta^\star) + (\mu_{\beta^\star} - \widehat{\mu})^\top (\widetilde{\beta}_B^{\text{ex}} - \beta^\star) + \frac{f_{\min}}{2} \left\| \widetilde{\beta}_B^{\text{ex}} - \beta^\star \right\|_2^2.$$

Combining the two displays and applying Cauchy–Schwarz yields

$$\frac{f_{\min}}{2} \left\| \widetilde{\beta}_B^{\text{ex}} - \beta^\star \right\|_2^2 \le \left\| \widehat{\mu} - \mu_{\beta^\star} \right\|_2 \left\| \widetilde{\beta}_B^{\text{ex}} - \beta^\star \right\|_2,$$

which proves equation (17) unless the two parameters coincide, in which case it is trivial. The concentration bound just proved gives the parameter rate. Because $\beta^\star$ is interior and the estimation error tends to zero, the minimizer is interior with probability tending to one.

If A is replaced by $\widehat{A}$, the gradient at $\beta^\star$ is $\nabla \widehat{A}(\beta^\star) - \widehat{\mu}$ and has norm at most $\|\widehat{\mu} - \mu_{\beta^\star}\| + \epsilon_A$. When the strong-convexity modulus on Θ is $f_{\min}/2$, the preceding argument gives $(f_{\min}/4) \left\| \widetilde{\beta}_B - \beta^\star \right\|^2 \le (\|\widehat{\mu} - \mu_{\beta^\star}\| + \epsilon_A) \left\| \widetilde{\beta}_B - \beta^\star \right\|$, which proves equation (18). Uniform Hessian consistency gives the curvature condition with probability tending to one. If $\epsilon_A = O_p(B^{-\gamma/2})$, the implemented estimate therefore also has error $O_p(B^{-\gamma/2})$ and is interior with probability tending to one. □

The initialization is therefore consistent. Assumption 3 also makes the simulated information matrices consistent at $\widetilde{\beta}_B$; concrete Monte Carlo rates are proved in Appendix G.1. We now show that optimizing these matrices gives the oracle information in the limit. The argument needs two facts: the optimal information matrix is unique, and the objective changes continuously under small matrix errors away from singularity.

For fixed β, write $F = F(\beta)$, $I_S = I_S(\beta)$, and $A_S = I_S/c(S)$. Then $M(p) = \sum_S p_S A_S$ and $\Phi(p) = \operatorname{tr}(FM(p)^{-1})$, with value $+\infty$ for singular $M(p)$.

Lemma 6 (Existence and uniqueness of the optimal information)**.** *If there exists $p_0 \in \Delta_{\mathcal{S}}$ with $M(p_0) \succ 0$, then Φ attains its minimum on $\Delta_{\mathcal{S}}$. When $F \succ 0$, the minimizer $M^\star$ of $M \mapsto \operatorname{tr}(FM^{-1})$ over attainable information matrices is unique, although its representing allocation need not be.*

Proof. The attainable matrices form the compact convex set $\mathcal{M} = \{\sum_S p_S I_S/c(S) : p \in \Delta_{\mathcal{S}}\}$. Since $F \succ 0$, for $M \succ 0$ we have $\operatorname{tr}(FM^{-1}) \ge \lambda_{\min}(F)/\lambda_{\min}(M)$. Thus the objective tends to infinity whenever an attainable positive-definite matrix approaches a singular matrix. Extending the objective by $+\infty$ at singular matrices gives a lower-semicontinuous function on $\mathcal{M}$. The assumed feasible $M(p_0) \succ 0$ has finite objective, so a finite minimum is attained.

For a symmetric direction H, differentiating the inverse twice gives

$$D^2\{\operatorname{tr}(FM^{-1})\}[H,H] = 2\operatorname{tr}(FM^{-1}HM^{-1}HM^{-1}) = 2\operatorname{tr}(CD^2),$$

where $C = M^{-1/2}FM^{-1/2} \succ 0$ and $D = M^{-1/2}HM^{-1/2}$. If $H \neq 0$, then D^2 is nonzero and positive semidefinite, so the final trace is positive. The objective is strictly convex in M; two different minimizing information matrices would therefore have a midpoint with smaller objective. The allocation vector itself need not be unique because different vectors can represent the same matrix. □

Different panel allocations can attain the same optimum, so convergence should be stated for their information matrices rather than for the allocation vectors. Since matrix inversion amplifies errors near singularity, the following bound makes the dependence on the smallest eigenvalue explicit.

Lemma 7 (Objective perturbation on a nonsingular region)**.** *Suppose $F \succeq f_0 I_r$ for $f_0 > 0$ and $M(p) \succeq \kappa I_r$ for $\kappa > 0$ on a set $\mathcal{P} \subseteq \Delta_{\mathcal{S}}$. Let $\widehat{F}$ and $\widehat{I}_S$ satisfy $\left\|\widehat{F} - F\right\|_{\mathrm{op}} \le \delta_F$ and $\max_S \left\|\widehat{I}_S - I_S\right\|_{\mathrm{op}} \le \delta_I$. Set $\widehat{M}(p) = \sum_S p_S \widehat{I}_S / c(S)$ and $\delta_M = \delta_I / c_{\min}$, and assume $\delta_M < \kappa$. Then, for every $p \in \mathcal{P}$,*

$$\left|\mathrm{tr}(\widehat{F}\widehat{M}(p)^{-1}) - \mathrm{tr}(FM(p)^{-1})\right| \le \frac{r\delta_F}{\kappa - \delta_M} + \frac{r\,\|F\|_{\mathrm{op}}\,\delta_M}{\kappa(\kappa - \delta_M)}. \tag{19}$$

Proof. Because $\left\|\widehat{M}(p) - M(p)\right\|_{\mathrm{op}} \le \delta_M$, Weyl's inequality gives $\widehat{M}(p) \succeq (\kappa - \delta_M) I_r$. Thus both inverses exist. Add and subtract $\mathrm{tr}(F\widehat{M}^{-1})$:

$$\left|\mathrm{tr}(\widehat{F}\widehat{M}^{-1}) - \mathrm{tr}(FM^{-1})\right| \le \left|\mathrm{tr}((\widehat{F} - F)\widehat{M}^{-1})\right| + \left|\mathrm{tr}(F(\widehat{M}^{-1} - M^{-1}))\right|.$$

The first term is bounded by $\left\|\widehat{F} - F\right\|_* \left\|\widehat{M}^{-1}\right\|_{\mathrm{op}} \le r\delta_F/(\kappa - \delta_M)$. For the second, the resolvent identity gives $\widehat{M}^{-1} - M^{-1} = \widehat{M}^{-1}(M - \widehat{M})M^{-1}$. Therefore

$$\left\|\widehat{M}^{-1} - M^{-1}\right\|_{\mathrm{op}} \le \frac{\delta_M}{\kappa(\kappa - \delta_M)}.$$

Using $\|F\|_* \le r\,\|F\|_{\mathrm{op}}$ proves equation (19). □

The Lipschitz bounds imply that consistent information estimates at $\widetilde{\beta}_B$ also converge to the information at $\beta^\star$. We still need to show that the estimated optimizer avoids nearly singular information matrices. Write the population quantities at the truth as

$$\Phi_\star(p) = \Phi_{\beta^\star}(p), \qquad M_\star(p) = M_{\beta^\star}(p),$$

and the estimated objective by

$$\widehat{\Phi}_B(p) = \mathrm{tr}[\widehat{F}\widehat{M}_B(p)^{-1}], \qquad \widehat{M}_B(p) = \sum_S \frac{p_S}{c(S)} \widehat{I}_S.$$

The population objective is assigned value $+\infty$ unless $M_\star(p) \succ 0$, and the estimated objective is assigned value $+\infty$ unless $\widehat{M}_B(p) \succ 0$.

Lemma 8 (Nonsingularity and consistency of the fitted design)**.** *Under Assumptions 1–3, there is a constant $\kappa > 0$ such that*

$$\mathbb{P}\{M_\star(\widehat{p}_B) \succeq \kappa I_r\} \longrightarrow 1.$$

Moreover,

$$\Phi_\star(\widehat{p}_B) - \min_{p \in \Delta_{\mathcal{S}}} \Phi_\star(p) \xrightarrow{p} 0.$$

Furthermore, $M_\star(\widehat{p}_B) \to_p M^\star$.

Proof. By collective observability, choose a deterministic safe allocation p^{safe} with $M_\star(p^{\mathrm{safe}}) \succ 0$. By Lemma 5, $\widetilde{\beta}_B \to_p \beta^\star$. Lemma 2 and Assumption 3 then imply

$$\widehat{F} \to_p F(\beta^\star), \qquad \widehat{M}_B(p^{\mathrm{safe}}) \to_p M_\star(p^{\mathrm{safe}}).$$

Hence, with probability tending to one, $\widehat{F} \succeq f_{\min} I_r/2$, $\widehat{M}_B(p^{\mathrm{safe}})$ is positive definite, and $\widehat{\Phi}_B(p^{\mathrm{safe}}) \le C_0$ for some finite deterministic C_0.

Let $\eta_B = o_p(1)$ be the optimization gap, so $\widehat{\Phi}_B(\widehat{p}_B) \le \inf_p \widehat{\Phi}_B(p) + \eta_B$. On the same high-probability event, $\widehat{\Phi}_B(\widehat{p}_B) \le C_0 + 1$ for all sufficiently large B. For any positive-definite matrix H,

$$\mathrm{tr}(\widehat{F}H^{-1}) \ge \lambda_{\min}(\widehat{F})\,\mathrm{tr}(H^{-1}) \ge \frac{\lambda_{\min}(\widehat{F})}{\lambda_{\min}(H)}.$$

Applying this with $H = \widehat{M}_B(\widehat{p}_B)$ yields

$$\lambda_{\min}(\widehat{M}_B(\widehat{p}_B)) \geq \frac{f_{\min}}{2(C_0+1)} =: 2\kappa.$$

Finally,

$$\begin{aligned}\sup_{p\in\Delta_{\mathcal{S}}} \left\|\widehat{M}_B(p) - M_\star(p)\right\|_{\mathrm{op}} &\leq \frac{1}{c_{\min}} \max_S \left\|\widehat{I}_S - I_S(\beta^\star)\right\|_{\mathrm{op}} \\ &\leq \frac{1}{c_{\min}} \left\{ \max_S \left\|\widehat{I}_S - I_S(\widetilde{\beta}_B)\right\|_{\mathrm{op}} + L_{\mathrm{loc}} \left\|\widetilde{\beta}_B - \beta^\star\right\| \right\} = o_p(1).\end{aligned}$$

Thus $M_\star(\widehat{p}_B) \succeq \kappa I_r$ with probability tending to one.

Thus the estimated optimizer cannot approach singular information while outperforming the fixed nonsingular allocation. We can apply Lemma 7 at both the estimated and oracle designs.

Let $p^\star$ be any population minimizer. Both $M_\star(p^\star)$ and, by the eigenvalue bound above, $M_\star(\widehat{p}_B)$ have eigenvalues bounded below with probability tending to one. On the corresponding nonsingular set, Lemma 7, Lemma 2, and Assumption 3 imply

$$\left|\widehat{\Phi}_B(p) - \Phi_\star(p)\right| = o_p(1)$$

for $p = p^\star$ and $p = \widehat{p}_B$. Therefore

$$\begin{aligned}\Phi_\star(\widehat{p}_B) - \Phi_\star(p^\star) &\leq \left|\Phi_\star(\widehat{p}_B) - \widehat{\Phi}_B(\widehat{p}_B)\right| \\ &\quad + \widehat{\Phi}_B(\widehat{p}_B) - \widehat{\Phi}_B(p^\star) + \left|\widehat{\Phi}_B(p^\star) - \Phi_\star(p^\star)\right| \\ &\leq o_p(1) + \eta_B + o_p(1) = o_p(1).\end{aligned}$$

This proves value consistency.

The simplex is compact, and $p \mapsto M_\star(p)$ is continuous. Take any subsequence along which $M_\star(\widehat{p}_B)$ converges in distribution; by a further subsequence and a Skorokhod representation if necessary, we may reason along almost-sure convergent sample paths. Every limit point $\overline{M}$ is attainable and, by continuity of the objective on the nonsingular region, satisfies $\mathrm{tr}(F(\beta^\star)\overline{M}^{-1}) = \Phi^\star_{\beta^\star}$. Lemma 6 states that the optimal information matrix is unique, so $\overline{M} = M^\star$. All subsequential limits are the same; hence $M_\star(\widehat{p}_B) \to_p M^\star$. □

The learned allocation therefore attains the oracle objective and its unique information matrix. The lower eigenvalue bound will also keep the final scoring updates well conditioned.

Integer allocation preserves this conclusion. With $B_{\mathrm{main}} = B - b_B$ and $n_{B,S} = \lfloor B_{\mathrm{main}}\widehat{p}_{B,S}/c(S)\rfloor$, write $n_{B,S}c(S) = B_{\mathrm{main}}\widehat{p}_{B,S} - d_{B,S}$, where $0 \leq d_{B,S} < c(S)$. Then

$$\frac{J_B(\beta^\star)}{B} = \frac{B_{\mathrm{main}}}{B} M_\star(\widehat{p}_B) - \frac{1}{B}\sum_S \frac{d_{B,S}}{c(S)} I_S(\beta^\star) \xrightarrow{p} M^\star.$$

The remainder has norm $O(B^{-1})$ because the panel family is fixed. The unused main budget is less than $\sum_S c(S)$, and $b_B/B \to 0$. Thus the information limit is unchanged by either initialization expenditure or integer rounding.

D.2 Main-Sample Estimation and Asymptotic Normality

We now analyze the parameter estimate obtained from the main measurements. The allocation is random, but all its randomness is determined before collecting the main sample. Conditional on the initialization data and design simulations, the main observations are therefore independent with known panel assignments. This lets us apply likelihood expansions while retaining the data-dependent design. Let $\mathcal{G}_B$ be the sigma-field generated by all pre-main-stage randomness. Conditional on $\mathcal{G}_B$, the counts $n_{B,S}$ are fixed and the observations within each panel are independent with law $Q_{\beta^\star,S}$. Write $P^{(B)}_{\beta,\mathrm{main}}(\cdot \mid \mathcal{G}_B)$ for this conditional main-sample experiment.

Write $Y_{S,k}$ for the kth main observation assigned to panel S. For this panel, let $\ell_{\beta,S}(y) = \log r_{\beta,S}(y)$, and define the main-sample central sequence

$$\Delta_B = \frac{1}{\sqrt{B}} \sum_{S \in \mathcal{S}} \sum_{k=1}^{n_{B,S}} m_{\beta^\star,S}(Y_{S,k}). \tag{20}$$

Lemma 9 (Conditional LAN)**.** *Under Assumptions 1–3, suppose the realized counts satisfy $n_{B,S}c(S)/B - \widehat{p}_{B,S} \to_p 0$ for every S and $b_B/B \to 0$. Then, for every fixed $h \in \mathbb{R}^r$,*

$$\log \frac{dP^{(B)}_{\beta^\star + h/\sqrt{B},\mathrm{main}}(\cdot \mid \mathcal{G}_B)}{dP^{(B)}_{\beta^\star,\mathrm{main}}(\cdot \mid \mathcal{G}_B)} = h^\top \Delta_B - \frac{1}{2} h^\top M^\star h + o_{P_{\beta^\star}}(1), \tag{21}$$

and $\Delta_B \Rightarrow \mathcal{N}(0, M^\star)$. More precisely, the conditional characteristic function of Δ_B given $\mathcal{G}_B$ converges in probability to that of $\mathcal{N}(0, M^\star)$. The initialization contribution to the full GAMD likelihood is $o_p(1)$ at the $B^{-1/2}$ local scale because $b_B/B \to 0$.

Proof. Conditional on $\mathcal{G}_B$, the log-likelihood ratio is

$$\Lambda_B(h) = \sum_{S,k} \left\{ \ell_{\beta^\star + h/\sqrt{B},S}(Y_{S,k}) - \ell_{\beta^\star,S}(Y_{S,k}) \right\}.$$

Lemma 4 identifies the marginal score in quadratic mean, and Lemma 2 bounds the third derivatives of the log density. These bounds give an explicit remainder after summing over the main sample. Taylor expansion at $\beta^\star$ gives

$$\Lambda_B(h) = h^\top \Delta_B + \frac{1}{2B} \sum_{S,k} h^\top \nabla^2 \ell_{\beta^\star,S}(Y_{S,k}) h + R_B(h),$$

with $|R_B(h)| \le C \|h\|^3 N_B / B^{3/2}$, where $N_B = \sum_S n_{B,S} = |\mathcal{I}_{\mathrm{main}}|$ is the main-sample size. Because $N_B \le B/c_{\min}$, $R_B(h) = O(B^{-1/2})$.

By Lemma 3, $\mathbb{E}_{\beta^\star} \nabla^2 \ell_{\beta^\star,S}(Y) = -I_S(\beta^\star)$. The centered Hessians are independent, bounded matrices conditional on $\mathcal{G}_B$. A conditional law of large numbers therefore gives

$$\frac{1}{B} \sum_{S,k} \left\{ \nabla^2 \ell_{\beta^\star,S}(Y_{S,k}) + I_S(\beta^\star) \right\} \to_p 0.$$

The deterministic conditional mean satisfies

$$\begin{aligned} \frac{1}{B} \sum_S n_{B,S} I_S(\beta^\star) &= \sum_S \frac{n_{B,S} c(S)}{B} \frac{I_S(\beta^\star)}{c(S)} \\ &= M_{\beta^\star}(\widehat{p}_B) + o_p(1) \to_p M^\star \end{aligned}$$

by Lemma 8. This proves the quadratic term in equation (21).

The summands in equation (20) are conditionally independent, mean zero, and uniformly bounded. Their conditional covariance is $B^{-1} \sum_S n_{B,S} I_S(\beta^\star) \to_p M^\star$. The conditional Lindeberg condition is automatic from boundedness, so the conditional triangular-array central limit theorem yields $\Delta_B \Rightarrow \mathcal{N}(0, M^\star)$. The conditional covariance converges to the nonrandom matrix $M^\star$, so the same triangular-array argument makes the conditional characteristic function converge in probability to $\exp(-t^\top M^\star t/2)$ for each $t \in \mathbb{R}^r$. Averaging these bounded characteristic functions gives the unconditional Gaussian limit. Finally, the initialization sample contains $O(b_B)$ bounded score contributions. Under a local perturbation of size $B^{-1/2}$, its linear log-likelihood term is $O_p(\sqrt{b_B/B}) = o_p(1)$ and its quadratic term is $O(b_B/B) = o(1)$. Thus including the initialization data does not change the LAN limit. □

The normalized main-sample score has limiting covariance $M^\star$. The initial parameter error is larger than the main-sample statistical error, so we must show that scoring removes it. We first write the exact score and information for the main sample: Define

$$U_B(\beta) = \sum_{S \in \mathcal{S}} \sum_{k=1}^{n_{B,S}} m_{\beta,S}(Y_{S,k}), \qquad J_B(\beta) = \sum_{S \in \mathcal{S}} n_{B,S} I_S(\beta).$$

The ideal Fisher-scoring map is $T_B(\beta) = \beta + J_B(\beta)^{-1}U_B(\beta)$. The implemented map at iteration j is $\widehat{T}_{B,j}(\beta^{(j)}) = \beta^{(j)} + [\widehat{J}_B^{(j)}]^{-1}\widehat{U}_B^{(j)}$. It is evaluated only at the current iterate, with fresh simulations generated at that parameter. We analyze the ideal map first and then use Assumption 3 to control the difference made by Monte Carlo computation. This part of the proof does not require a particular finite-sample safeguard.

Lemma 10 (Local scoring recursion)**.** *There is a fixed neighborhood $\mathcal{N}$ of $\beta^\star$ such that, uniformly for random $\beta \in \mathcal{N}$,*

$$T_B(\beta) - \beta^\star = J_B(\beta^\star)^{-1}U_B(\beta^\star) + O_p\Big(B^{-1/2}\,\|\beta - \beta^\star\| + \|\beta - \beta^\star\|^2\Big). \tag{22}$$

Consequently, for an ideal update $\beta^{(j+1)} = T_B(\beta^{(j)})$ with $e_j = \beta^{(j)} - \beta^\star = o_p(1)$,

$$\|e_{j+1}\| = O_p\Big(B^{-1/2} + \|e_j\|^2\Big). \tag{23}$$

Proof. Lemma 8 and continuity of $I_S(\beta)$ imply that, on a sufficiently small fixed neighborhood $\mathcal{N}$ and with probability tending to one,

$$\lambda_{\min}\{J_B(\beta)/B\} \geq \kappa_J > 0, \qquad \sup_{\beta\in\mathcal{N}} \|J_B(\beta)/B\|_{\mathrm{op}} \leq C_J. \tag{24}$$

Thus $\big\|J_B(\beta)^{-1}\big\|_{\mathrm{op}} \leq (\kappa_J B)^{-1}$.

Let $e = \beta - \beta^\star$, and let $\dot{U}_B(\beta) = \nabla_\beta U_B(\beta)$ denote the score Jacobian. A second-order Taylor expansion of U_B gives

$$U_B(\beta) = U_B(\beta^\star) + \dot{U}_B(\beta^\star)e + R_B(e), \qquad \|R_B(e)\| \leq CB\,\|e\|^2, \tag{25}$$

where the bound follows from the uniform second derivative of the panel score. At the truth, the information identity gives $\mathbb{E}[\dot{U}_B(\beta^\star) \mid \mathcal{G}_B] = -J_B(\beta^\star)$. The difference

$$H_B = \dot{U}_B(\beta^\star) + J_B(\beta^\star)$$

is a sum of independent centered bounded matrices conditional on $\mathcal{G}_B$, so $\|H_B\|_{\mathrm{op}} = O_p(\sqrt{B})$. Also, by Lipschitz continuity of I_S,

$$\|J_B(\beta) - J_B(\beta^\star)\|_{\mathrm{op}} \leq CB\,\|e\|. \tag{26}$$

Substitute equation (25) into $T_B(\beta) - \beta^\star$:

$$T_B(\beta) - \beta^\star = e + J_B(\beta)^{-1}\{U_B(\beta^\star) - J_B(\beta^\star)e + H_Be + R_B(e)\}.$$

Add and subtract $J_B(\beta^\star)^{-1}U_B(\beta^\star)$. The inverse identity

$$J_B(\beta)^{-1} - J_B(\beta^\star)^{-1} = J_B(\beta)^{-1}\{J_B(\beta^\star) - J_B(\beta)\}J_B(\beta^\star)^{-1}$$

and equations (24)–(26) show that the inverse difference is $O(\|e\|/B)$. Because $U_B(\beta^\star) = O_p(\sqrt{B})$, the resulting term is $O_p(B^{-1/2}\,\|e\|)$. Similarly,

$$\{I_r - J_B(\beta)^{-1}J_B(\beta^\star)\}e = J_B(\beta)^{-1}\{J_B(\beta) - J_B(\beta^\star)\}e = O(\|e\|^2).$$

Finally, $J_B(\beta)^{-1}H_Be = O_p(B^{-1/2}\,\|e\|)$ and $J_B(\beta)^{-1}R_B(e) = O(\|e\|^2)$. Combining these bounds proves equation (22). Since $J_B(\beta^\star)^{-1}U_B(\beta^\star) = O_p(B^{-1/2})$, equation (23) follows. □

A scoring step cancels the linear effect of the current parameter error. What remains is the main-sample score term and a quadratic initialization remainder. The next lemma shows that the Monte Carlo errors are small enough to preserve this cancellation through the fixed updates.

Lemma 11 (Finite-step accuracy of the implemented scoring updates)**.** *Under Assumptions 1–3, let $\beta = \beta^{(j)} \in \mathcal{N}$ be a current iterate satisfying $\|\beta - \beta^\star\| = O_p(B^{-\gamma/2})$. Then, for every fixed update index j,*

$$\widehat{T}_{B,j}(\beta) - T_B(\beta) = o_p(B^{-1/2}). \tag{27}$$

If $\left\|\beta^{(0)} - \beta^\star\right\| = O_p(B^{-\gamma/2})$*, then for every fixed* $J \geq 0$*,*

$$\left\|\beta^{(J)} - \beta^\star\right\| = O_p\Big(B^{-1/2} + B^{-2^J\gamma/2}\Big). \tag{28}$$

If $2^J\gamma > 1$*, then*

$$\beta^{(J)} - \beta^\star = J_B(\beta^\star)^{-1}U_B(\beta^\star) + o_p(B^{-1/2}). \tag{29}$$

Proof. The step-equivalence clause of Assumption 3 is precisely equation (27). We first verify it from the sufficient score and information bounds in equation (48); the subsequent induction uses the step-equivalence property itself. Write $R_{U,j} = \widehat{U}_B^{(j)} - U_B(\beta)$ and $R_{J,j} = \widehat{J}_B^{(j)} - J_B(\beta)$, and set $\eta_{J,j} = B^{-1}\left\|R_{J,j}\right\|_{\mathrm{op}}$. The sufficient bounds give $\|R_{U,j}\| = o_p(\sqrt{B})$ and $\eta_{J,j} = o_p(B^{-(1-\gamma)/2})$. On the high-probability event in equation (24), Weyl's inequality and $\eta_{J,j} = o_p(1)$ imply $\left\|[\widehat{J}_B^{(j)}]^{-1}\right\|_{\mathrm{op}} = O_p(B^{-1})$. Therefore

$$\left\|[\widehat{J}_B^{(j)}]^{-1}R_{U,j}\right\| = o_p(B^{-1/2}).$$

The resolvent identity also gives

$$\left\|[\widehat{J}_B^{(j)}]^{-1} - J_B(\beta)^{-1}\right\|_{\mathrm{op}} \leq C\eta_{J,j}/B$$

with probability tending to one. The Taylor expansion equation (25), together with $\|\beta - \beta^\star\| = O_p(B^{-\gamma/2})$, yields

$$\|U_B(\beta)\| = O_p\Big(\sqrt{B} + B\left\|\beta - \beta^\star\right\|\Big).$$

Consequently,

$$\begin{aligned}&\left\|([\widehat{J}_B^{(j)}]^{-1} - J_B(\beta)^{-1})U_B(\beta)\right\| \\ &\qquad \leq O_p\Big(\eta_{J,j}B^{-1/2} + \eta_{J,j}\left\|\beta - \beta^\star\right\|\Big) = o_p(B^{-1/2}),\end{aligned}$$

because $\eta_{J,j}\left\|\beta - \beta^\star\right\| = o_p(B^{-(1-\gamma)/2}B^{-\gamma/2}) = o_p(B^{-1/2})$. Adding the two differences proves equation (27).

The information-error rate is stronger than consistency because it multiplies the current parameter error. Together with the score-error bound, it makes each simulated update differ from the exact update by $o_p(B^{-1/2})$. We now apply this conclusion successively.

The rate is proved by induction using equation (23) together with the implemented-map comparison just proved. For $J = 0$ it is the assumed initialization rate. Suppose the claim holds at step j. Then

$$\|e_j\|^2 = O_p\Big(B^{-1} + B^{-2^j\gamma}\Big) = O_p\Big(B^{-1/2} + B^{-2^{j+1}\gamma/2}\Big),$$

where $B^{-1} = O(B^{-1/2})$. Because the displayed rate is always $O_p(B^{-\gamma/2})$, the same map comparison applies at every fixed iterate. Substituting into equation (23) and adding its $o_p(B^{-1/2})$ implementation remainder proves the formula at $j + 1$.

For asymptotic linearity, apply equation (22) to the final update, whose input is e_{J-1}. By equation (28),

$$\|e_{J-1}\|^2 = O_p\Big(B^{-1} + B^{-2^{J-1}\gamma}\Big).$$

The condition $2^J\gamma > 1$ is equivalent to $2^{J-1}\gamma > 1/2$, so this squared error is $o_p(B^{-1/2})$. Also $B^{-1/2}\|e_{J-1}\| = o_p(B^{-1/2})$ because $e_{J-1} = o_p(1)$. The remainder in equation (22) is therefore $o_p(B^{-1/2})$. The map comparison shows that the final implemented update differs by another $o_p(B^{-1/2})$ term, yielding equation (29). □

The preceding expansion gives the asymptotic distribution of $\widehat{\beta}_B$. The following identity converts its parameter error into the full-distribution KL loss used in the main objective.

Lemma 12 (Local expansion of full-law KL)**.** *Uniformly for β in a neighborhood of $\beta^\star$ and $u \to 0$,*

$$D_{\mathrm{KL}}(Q_\beta \| Q_{\beta+u}) = \frac{1}{2} u^\top F(\beta) u + O(\|u\|_2^3). \tag{30}$$

Proof. Using equation (1),

$$\begin{aligned} D_{\mathrm{KL}}(Q_\beta \| Q_{\beta+u}) &= \mathbb{E}_{Q_\beta}\left[-u^\top \phi(X) + A(\beta+u) - A(\beta)\right] \\ &= A(\beta+u) - A(\beta) - u^\top \nabla A(\beta). \end{aligned}$$

Taylor expansion of A and the uniform bound on its third derivative from Lemma 1 give equation (30). □

Expected KL risk. The Gaussian limit describes typical estimation error, but an expectation can still be affected by rare runs with much larger error. Theorem 1 therefore adds the uniform-integrability condition

$$\lim_{K\to\infty} \sup_B \mathbb{E}\Big[B\|\widehat{\beta}_B - \beta^\star\|_2^2 \mathbf{1}\Big\{B\|\widehat{\beta}_B - \beta^\star\|_2^2 > K\Big\}\Big] = 0. \tag{31}$$

Under equation (31), the local KL expansion above is uniformly integrable, so the Gaussian limit can be passed through the expectation. This is a tail condition on the estimator, not an additional modeling assumption. Appendix G.3 gives a concrete sufficient implementation: finite-sample safeguards combined with explicit Monte Carlo sample-size rates.

Achievability in Theorem 1. The design conclusion follows from Lemma 8. For the estimator, Lemma 11 gives

$$\sqrt{B}(\beta^{(J)} - \beta^\star) = \left\{\frac{J_B(\beta^\star)}{B}\right\}^{-1} \Delta_B + o_p(1).$$

The normalized information converges in probability to $M^\star$, and Lemma 9 gives $\Delta_B \Rightarrow \mathcal{N}(0, M^\star)$. Slutsky's theorem yields a Gaussian limit with covariance $M^{\star-1} M^\star M^{\star-1} = M^{\star-1}$. The implemented iterates have the same asymptotic linear expansion by the update-error calculation in Lemma 11, so the Monte Carlo approximation does not change the limit under Assumption 3.

The same expansion establishes regularity under local alternatives. Lemma 9 gives joint convergence of the asymptotic linear term and the full-data log-likelihood ratio. Contiguity transfers the $o_p(1)$ remainders to $\beta_B = \beta^\star + h/\sqrt{B}$, and Le Cam's third lemma shifts the limit of Δ_B to $\mathcal{N}(M^\star h, M^\star)$. Therefore $\sqrt{B}(\widehat{\beta}_B - \beta_B)$ converges to $\mathcal{N}(0, M^{\star-1})$ for every fixed h, independently of h. The conditional likelihood expansion is uniform for bounded h, and the same boundedness arguments give locally uniform information stability. Thus GAMD belongs to the comparison class used in Appendix D.3.

For expected KL risk, assume the uniform-integrability condition equation (31) stated in Theorem 1. Let $Z_B = \sqrt{B}(\beta^{(J)} - \beta^\star)$ and $e_B = Z_B/\sqrt{B}$. The exact integral remainder for the log-partition function gives

$$B D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\beta^{(J)}}) = \int_0^1 (1-t) Z_B^\top F(\beta^\star + t e_B) Z_B \, dt.$$

By the Gaussian limit above, $Z_B \Rightarrow Z \sim \mathcal{N}(0, M^{\star-1})$ and $e_B \to_p 0$. Continuity of F implies convergence in distribution of the displayed loss to $\frac{1}{2} Z^\top F(\beta^\star) Z$. The global bound $F(\beta) \preceq L_\phi^2 I_r$ from Lemma 1 gives the domination

$$0 \le B D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\beta^{(J)}}) \le \frac{L_\phi^2}{2} \|Z_B\|_2^2.$$

The right-hand side is uniformly integrable by the estimator-level tail condition stated above. Hence expectations converge, and $\mathbb{E}[Z^\top F(\beta^\star) Z] = \mathrm{tr}(F(\beta^\star) M^{\star-1})$ gives the claim. The limit is $\frac{1}{2}\Phi^\star_{\beta^\star}$ because $M^\star$ minimizes $\mathrm{tr}(F(\beta^\star) M^{-1})$ over attainable full-budget information matrices. The next part establishes that this is the sharp local benchmark. □

D.3 Local Minimax Lower Bound

We compare GAMD with other acquisition rules that choose each panel using only past observations and have a deterministic limiting information matrix. The comparison considers target parameters $\beta^\star + h/\sqrt{B}$ near a fixed interior $\beta^\star$. This is the statistical scale at which regular estimators make errors; it does not require the target itself to be close to the reference parameter 0.

For a prospective strategy, let $P_\beta^{(B)}$ denote the joint law of the observed data and parameter-independent policy randomization. The strategy may depend on past observations but not on the unknown parameter. Its limiting information is *attainable* if it can be written as $M = \sum_S \pi_S I_S(\beta^\star)/c(S)$ with $\pi_S \geq 0$ and $\sum_S \pi_S \leq 1$; the inequality allows unused budget. A regular estimator has a limiting law for $\sqrt{B}(\widehat{\beta}_B - \beta^\star - h/\sqrt{B})$ that does not depend on fixed h under $P^{(B)}_{\beta^\star + h/\sqrt{B}}$. For a fixed acquisition strategy, its local asymptotic minimax risk is understood through the usual expanding-local-neighborhood, truncated-loss bound

$$\lim_{T\to\infty} \lim_{R\to\infty} \liminf_{B\to\infty} \inf_{\widehat{\beta}_B} \sup_{\|h\|\leq R} \mathbb{E}_{\beta^\star + h/\sqrt{B}} \min\Big\{T,\, BD_{\mathrm{KL}}(Q_{\beta^\star + h/\sqrt{B}} \| Q_{\widehat{\beta}_B})\Big\}.$$

Here the infimum is over estimators based on that strategy's observations. The supremum over local alternatives compares estimators throughout a shrinking statistical neighborhood. The bound below applies to every estimator for each information-stable acquisition strategy, and hence to the regular estimators considered in the main theorem.

Lemma 13 (Local asymptotic minimax lower bound)**.** *Consider a parameter-independent prospective acquisition strategy with an attainable limit $M \succ 0$ for which, for every fixed $R < \infty$,*

$$\frac{1}{B}\sum_i I_{S_i}(\beta^\star) \xrightarrow{p} M \quad \textit{uniformly under } \beta = \beta^\star + h/\sqrt{B},\ \|h\| \leq R, \tag{32}$$

where $M = \sum_S \pi_S I_S(\beta^\star)/c(S)$ is deterministic and positive definite for nonnegative cost shares satisfying $\sum_S \pi_S \leq 1$. Then its local asymptotic minimax full-law KL risk is at least $\frac{1}{2}\operatorname{tr}(F(\beta^\star)M^{-1})$, and hence at least $\frac{1}{2}\Phi^\star_{\beta^\star}$.

Proof. Because each selected panel is measurable with respect to the past, the policy's selection kernel is the same under every local parameter and cancels from the likelihood ratio. Under $Q_{\beta^\star}$, the panel-score increments are martingale differences and their predictable quadratic variation is the accumulated panel information. There are at most $\lfloor B/c_{\min}\rfloor$ observations; a random stopping time can be handled by zero-padding subsequent increments. The locally uniform convergence in equation (32), the bounded-increment martingale central limit theorem, and the same third-order expansion used in Lemma 9 therefore give

$$\log \frac{dP^{(B)}_{\beta^\star + h/\sqrt{B}}}{dP^{(B)}_{\beta^\star}} = h^\top \Delta_B - \frac{1}{2}h^\top M h + o_{P_{\beta^\star}}(1), \qquad \Delta_B \Rightarrow \mathcal{N}(0, M).$$

Thus the experiments are LAN with deterministic information M. Le Cam's third lemma then yields $\Delta_B \Rightarrow \mathcal{N}(Mh, M)$ under $\beta = \beta^\star + h/\sqrt{B}$, and the limit experiment can be written as

$$Z = h + M^{-1/2}\varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, I_r).$$

Lemma 12 shows that the scaled local loss converges uniformly on compact h-sets to

$$L(\widehat{h}, h) = \frac{1}{2}(\widehat{h} - h)^\top F(\beta^\star)(\widehat{h} - h).$$

The minimax risk in this Gaussian shift experiment is $\frac{1}{2}\operatorname{tr}(F(\beta^\star)M^{-1})$. To see the lower bound, place the prior $h \sim \mathcal{N}(0, \tau^2 I_r)$. The posterior covariance is $(M + \tau^{-2}I_r)^{-1}$, so the Bayes risk under L equals

$$\frac{1}{2}\operatorname{tr}\big[F(\beta^\star)(M + \tau^{-2}I_r)^{-1}\big].$$

Every minimax risk dominates every Bayes risk, and letting $\tau \to \infty$ yields the stated value. The estimator $\widehat{h} = Z$ has constant risk exactly $\frac{1}{2}\operatorname{tr}(F(\beta^\star)M^{-1})$, so the Gaussian benchmark is sharp.

The local asymptotic minimax theorem for LAN experiments (Van der Vaart, 2000; Le Cam & Yang, 2012) transfers this lower bound to the original experiment, first for truncated losses and then by monotone convergence. To compare with the full-budget oracle, write $M = \sum_S \pi_S I_S(\beta^\star)/c(S)$ with $\sum_S \pi_S \le 1$. If $\sum_S \pi_S < 1$, allocate the unused share to any admissible panel and denote the resulting full-budget matrix by M'. Then $M' \succeq M$, so $M'^{-1} \preceq M^{-1}$ and

$$\operatorname{tr}(F(\beta^\star)M^{-1}) \ge \operatorname{tr}(F(\beta^\star)M'^{-1}) \ge \Phi^\star_{\beta^\star}.$$

The same conclusion holds with $M' = M$ when the full budget is used asymptotically. This completes the proof. □

GAMD attains $\frac{1}{2}\Phi^\star_{\beta^\star}$, and every information-stable acquisition rule has local minimax risk at least this large. This proves the optimality claim in Theorem 1. Its scope is a first-order comparison as B grows, not a finite-budget lower bound on each realized loss.

E Proof of the Finite-Budget Guarantee

This section proves Proposition 2 for the Monte Carlo updates actually used by GAMD. We first bound how far the learned allocation is from the oracle when initialization and information estimates are inaccurate. Conditional on such an allocation, we then bound the final KL error using the main target sample and the simulated scores. The recovery guarantee applies when the initialization and design errors are small enough for local estimation.

E.1 Finite-Sample Error in the Learned Design

Three errors affect the learned allocation: the error of $\widetilde{\beta}_B$, Monte Carlo error in the information matrices, and numerical suboptimality of the design solver. We use the main-text definitions

$$\epsilon_{\text{init}} = \|\widetilde{\beta}_B - \beta^\star\|_2, \qquad \epsilon_{\text{info}} = \|\widehat{F} - F(\widetilde{\beta}_B)\|_{\text{op}} + \max_S \|\widehat{I}_S - I_S(\widetilde{\beta}_B)\|_{\text{op}},$$

and let $\epsilon_{\text{opt}} \ge 0$ satisfy $\widehat{\Phi}_B(\widehat{p}_B) \le \inf_{p\in\Delta_{\mathcal{S}}} \widehat{\Phi}_B(p) + \epsilon_{\text{opt}}$. The perturbation bound applies to information errors at the truth, so define

$$\delta_F = \left\|\widehat{F} - F(\beta^\star)\right\|_{\text{op}}, \qquad \delta_I = \max_S \left\|\widehat{I}_S - I_S(\beta^\star)\right\|_{\text{op}}.$$

The information errors at $\beta^\star$ include both simulation error and the use of $\widetilde{\beta}_B$ in place of $\beta^\star$. By Lemma 2,

$$\delta_F + \delta_I \le \epsilon_{\text{info}} + 2L_{\text{loc}}\epsilon_{\text{init}}. \tag{33}$$

Lemma 14 (Finite perturbation of the optimized design)**.** *There exist constants $\varepsilon_0, C > 0$, depending only on a safe design, the local eigenvalue margins, $F(\beta^\star)$, and the panel costs, such that whenever $\delta_F + \delta_I + \epsilon_{\text{opt}} \le \varepsilon_0$,*

$$\Phi_{\beta^\star}(\widehat{p}_B) - \Phi^\star_{\beta^\star} \le C(\delta_F + \delta_I + \epsilon_{\text{opt}}).$$

Proof. Choose a safe allocation p^{safe} with $M_\star(p^{\text{safe}}) \succeq m_0 I_r$. For sufficiently small δ_F and δ_I, the estimated quantities satisfy $\widehat{F} \succeq f_{\min} I_r/2$ and $\widehat{M}_B(p^{\text{safe}}) \succeq m_0 I_r/2$. Consequently, $\widehat{\Phi}_B(p^{\text{safe}})$ is bounded by a finite constant C_0 that depends only on the local problem parameters. Approximate optimality gives $\widehat{\Phi}_B(\widehat{p}_B) \le C_0 + \epsilon_{\text{opt}}$. The argument in Lemma 8 then gives

$$\lambda_{\min}(\widehat{M}_B(\widehat{p}_B)) \ge \frac{f_{\min}}{2(C_0 + \epsilon_{\text{opt}})}.$$

If $\delta_I/c_{\min}$ is sufficiently small, Weyl's inequality transfers this lower bound to $M_\star(\widehat{p}_B)$. Thus both $\widehat{p}_B$ and a population optimizer $p^\star$ lie in a common nonsingular region with eigenvalue lower bound $\kappa > 0$.

Apply Lemma 7 on this region. It gives a constant C_1 such that

$$\left|\widehat{\Phi}_B(p) - \Phi_\star(p)\right| \le C_1(\delta_F + \delta_I)$$

for $p \in \{p^\star, \widehat{p}_B\}$. Therefore

$$\begin{aligned}\Phi_\star(\widehat{p}_B) - \Phi_\star(p^\star) &\le \left|\Phi_\star(\widehat{p}_B) - \widehat{\Phi}_B(\widehat{p}_B)\right| \\ &\quad + \widehat{\Phi}_B(\widehat{p}_B) - \widehat{\Phi}_B(p^\star) + \left|\widehat{\Phi}_B(p^\star) - \Phi_\star(p^\star)\right| \\ &\le 2C_1(\delta_F + \delta_I) + \epsilon_{\mathrm{opt}}.\end{aligned}$$

Absorbing constants proves the claim. □

Combining Lemma 14 with equation (33) proves equation (8). The concentration results in Appendix D.1 then make the abstract error decomposition explicit. For example, when the reference Monte Carlo error in A is negligible,

$$\epsilon_{\mathrm{init}} = O_{\mathbb{P}}\left(\sqrt{\frac{r}{n_0}}\right),$$

and Lemma 19 gives the uniform $N_g^{-1/2}$ rate for panel information up to logarithmic factors.

The bound is deterministic once the initialization and information estimates are fixed. Comparing the estimated optimum with one fixed nonsingular allocation ensures that both the estimated and oracle information matrices have eigenvalues bounded away from zero. We use this bound to control the final parameter updates.

E.2 Finite-Sample Error in Target Recovery

Even for a fixed allocation, the final estimate has two sources of error: the finite main sample and the Monte Carlo evaluation of its scores. More generator samples reduce the second error, but do not replace target measurements. For the observed main sample, the exact score and expected information are

$$U_B(\beta) = \sum_{i \in \mathcal{I}_{\mathrm{main}}} m_{\beta, S_i}(Y_i), \qquad J_B(\beta) = \sum_{i \in \mathcal{I}_{\mathrm{main}}} I_{S_i}(\beta),$$

whereas the implemented update uses $\widehat{U}_B^{(j)}$ and $\widehat{J}_B^{(j)}$ from Section 4.3. Write $B_{\mathrm{main}} = B - b_B$.

The scoring expansion uses the score at $\beta^\star$, the difference between observed and expected curvature, and the smallest eigenvalue of the information matrix. For a radius $r_0 > 0$, let $\mathcal{B}_0 = \{\beta : \|\beta - \beta^\star\| \le r_0\}$. Define the main-sample event $\mathcal{E}_B(\xi)$ by

$$\frac{1}{B_{\mathrm{main}}} \left\|U_B(\beta^\star)\right\|_2 \le \xi, \tag{34}$$

$$\frac{1}{B_{\mathrm{main}}} \left\|\dot{U}_B(\beta^\star) + J_B(\beta^\star)\right\|_{\mathrm{op}} \le \xi, \tag{35}$$

together with $J_B(\beta)/B_{\mathrm{main}} \succeq \kappa_J I_r$ for all $\beta \in \mathcal{B}_0$. For deterministic levels $\eta_U, \eta_J \ge 0$, define the update-computation event

$$\mathcal{C}_B(\eta_U, \eta_J) = \left\{ \max_{0 \le j < J} \frac{\left\|\widehat{U}_B^{(j)} - U_B(\beta^{(j)})\right\|_2}{B_{\mathrm{main}}} \le \eta_U, \right. \\ \left. \max_{0 \le j < J} \frac{\left\|\widehat{J}_B^{(j)} - J_B(\beta^{(j)})\right\|_{\mathrm{op}}}{B_{\mathrm{main}}} \le \eta_J \right\}. \tag{36}$$

This event measures the two approximations in equation (6). For total failure level $\delta \in (0,1)$, choose $\eta_U(\delta), \eta_J(\delta)$ so that this event holds with conditional probability at least $1 - \delta/2$, and set $\eta_{\mathrm{upd}}(\delta) = \eta_U(\delta) + r_0\eta_J(\delta)$. This bounds the maximum weighted sum used in the main text. We suppress the argument δ in the deterministic recursion.

Lemma 15 (Probability of the main-sample fluctuation event)**.** *Under Assumptions 1–2, there are constants $C, c > 0$ such that, conditional on any main allocation whose information is uniformly bounded below,*

$$\mathbb{P}\{\mathcal{E}_B(\xi_B(\delta)) \mid \mathcal{G}_B\} \ge 1 - \delta/2,$$

where one may take

$$\xi_B(\delta) = C\sqrt{\frac{r + \log(8r/\delta)}{B - b_B}} + C\frac{r + \log(8r/\delta)}{B - b_B}. \tag{37}$$

In particular, $\xi_B(\delta) = \widetilde{O}(\sqrt{(r + \log(1/\delta))/(B - b_B)})$.

Proof. At $\beta^\star$, the panel scores are independent conditional on the pre-main-stage sigma-field $\mathcal{G}_B$, have conditional mean zero, and are bounded in Euclidean norm by $2L_\phi$. A sphere-net Hoeffding or vector Bernstein argument gives equation (34) at the stated rate.

By the information identity, $\mathbb{E}[\dot{m}_{\beta^\star,S}(Y)] = -I_S(\beta^\star)$ when $Y \sim Q_{\beta^\star,S}$. Hence $\dot{U}_B(\beta^\star) + J_B(\beta^\star)$ is a sum of conditionally independent centered bounded matrices. Matrix Bernstein gives equation (35), with the same order up to the displayed dimension logarithm.

Finally, the normalized design information at $\beta^\star$ is bounded below on the design event established in Lemma 8. Continuity of $I_S(\beta)$ and finiteness of $\mathcal{S}$ allow r_0 to be chosen so that $J_B(\beta)/B_{\mathrm{main}} \succeq \kappa_J I_r$ throughout $\mathcal{B}_0$, apart from the $O(B_{\mathrm{main}}^{-1})$ integer-rounding error. A union bound completes the proof. □

To bound Monte Carlo error, distinguish the independent completions for different target observations from the single estimated mean $\widehat{\mu}^{(j)}$ subtracted from every score. Completion noise averages over target observations. An error in the shared mean does not, which explains the two sample-size terms below.

Lemma 16 (Finite update-computation error)**.** *Suppose fresh generator simulations are used at each of the fixed J updates. At update j, draw N_{upd} full samples from $Q_{\beta^{(j)}}$ and let $\widehat{\mu}^{(j)}$ be their feature mean. Estimate each target conditional score by the average of L_{upd} independent conditional feature completions minus $\widehat{\mu}^{(j)}$. Then, for a constant C_U,*

$$\eta_U = C_U\left\{\sqrt{\frac{r + \log(16J/\delta)}{B_{\mathrm{main}}L_{\mathrm{upd}}}} + \frac{r + \log(16J/\delta)}{B_{\mathrm{main}}} + \sqrt{\frac{r + \log(16J/\delta)}{N_{\mathrm{upd}}}} + \frac{r + \log(16J/\delta)}{N_{\mathrm{upd}}}\right\} \tag{38}$$

controls the first inequality in equation (36) *simultaneously over $j = 0, \ldots, J-1$ with probability at least $1 - \delta/4$. If the same full samples are used as outer panel samples and each distinct panel information matrix is estimated using cross-completion followed by PSD projection, then*

$$\eta_J = C'_J\left\{\sqrt{\frac{\log(16rJ|\mathcal{S}|/\delta)}{N_{\mathrm{upd}}}} + \frac{\log(16rJ|\mathcal{S}|/\delta)}{N_{\mathrm{upd}}}\right\} \tag{39}$$

controls the second inequality with probability at least $1 - \delta/4$. The constants depend only on the feature bound, panel costs, and local parameter region.

Proof. Fix an update j and condition on the main data and all randomness generated before that update. Write $\mu^{(j)} = \mu_{\beta^{(j)}}$. For target observation i, let

$$Z_{i,j} = \frac{1}{L_{\mathrm{upd}}}\sum_{\ell=1}^{L_{\mathrm{upd}}}\left\{\phi(X_{i,j}^{(\ell)}) - \mathbb{E}_{Q_{\beta^{(j)}}}[\phi(X) \mid X_{S_i} = Y_i]\right\},$$

where $X_{i,j}^{(\ell)} \sim Q_{\beta^{(j)}}(\cdot \mid S_i, Y_i)$ independently. Conditional on the main data and all pre-update randomness (so that $\beta^{(j)}$ is fixed), the vectors $Z_{i,j}$ are independent and centered, are bounded by $2L_\phi$, and have conditional second moments of order L_{upd}^{-1}. Vector Bernstein and a union bound over the fixed updates therefore control $B_{\mathrm{main}}^{-1}\left\|\sum_i Z_{i,j}\right\|_2$ by the first two terms in equation (38).

The common centering estimate contributes

$$\frac{N_B}{B_{\mathrm{main}}}\left\|\widehat{\mu}^{(j)} - \mu^{(j)}\right\|_2, \qquad N_B = |\mathcal{I}_{\mathrm{main}}|.$$

Because $N_B/B_{\text{main}} \le c_{\min}^{-1}$, vector Bernstein applied to the N_{upd} full tilted samples bounds this contribution by the last two terms in equation (38). Combining the conditional-completion and centering errors proves the score bound simultaneously over j.

For update-time information, Lemma 19 applies conditionally at each current iterate and uniformly over the finite panel family; subtracting the common $\widehat{\mu}^{(j)}$ does not alter the centered cross-covariance, and the mandatory PSD projection only changes its constant. A union bound over J updates gives the displayed uniform panel error in equation (39). Since

$$\frac{1}{B_{\text{main}}}\left\|\widehat{J}_B^{(j)} - J_B(\beta^{(j)})\right\|_{\text{op}} \le \frac{1}{B_{\text{main}}}\sum_i \left\|\widehat{I}_{S_i}^{(j)} - I_{S_i}(\beta^{(j)})\right\|_{\text{op}} \le \frac{1}{c_{\min}}\max_{S\in\mathcal{S}}\left\|\widehat{I}_S^{(j)} - I_S(\beta^{(j)})\right\|_{\text{op}},$$

the same rate controls the normalized information error. Combining the two probability bounds gives

$$\mathbb{P}\{\mathcal{C}_B(\eta_U(\delta), \eta_J(\delta)) \mid \mathcal{G}_B\} \ge 1 - \delta/2.$$

These statements may be applied up to the first exit from $\mathcal{B}_0$; the recursion below then proves that no exit occurs on the resulting event. □

The completion-noise term decreases with both the target sample size and L_{upd}. The centering term decreases with N_{upd} alone because the same estimated mean is subtracted from every target observation. Reusing the full samples as outer samples for information estimation does not require these error events to be independent: their simultaneous control follows from a union bound.

We now substitute these errors into the implemented scoring update. Score error changes the update directly, while information error also multiplies the current parameter error. Inside the radius-r_0 ball, their combined effect is bounded by $\eta_U + r_0\eta_J$, the quantity η_{upd} used in Proposition 2.

Lemma 17 (Finite recursion for the implemented updates)**.** *There are constants $a_0, a_1, a_2 > 0$ such that, on $\mathcal{E}_B(\xi) \cap \mathcal{C}_B(\eta_U, \eta_J)$, whenever $\eta_J < \kappa_J/2$ and the current iterate lies in $\mathcal{B}_0$,*

$$\left\|\beta^{(j+1)} - \beta^\star\right\| \le a_0(\xi + \eta_U) + a_1(\xi + \eta_J)\left\|\beta^{(j)} - \beta^\star\right\| + a_2\left\|\beta^{(j)} - \beta^\star\right\|^2. \tag{40}$$

Let $\zeta = \xi + \eta_U + r_0\eta_J$. If ζ and the initialization error are sufficiently small, all fixed J iterates remain in $\mathcal{B}_0$ and

$$\left\|\beta^{(J)} - \beta^\star\right\| \le C_J\left\{\zeta + \epsilon_{\text{init}}^{2^J}\right\}, \tag{41}$$

where C_J depends on the fixed number of updates but not on B.

Proof. Write $e = \beta^{(j)} - \beta^\star$, $\widehat{U} = \widehat{U}_B^{(j)}$, and $\widehat{J} = \widehat{J}_B^{(j)}$. On $\mathcal{C}_B(\eta_U, \eta_J)$,

$$\widehat{U} = U_B(\beta^{(j)}) + R_U, \qquad \widehat{J} = J_B(\beta^{(j)}) + R_J,$$

with $\|R_U\| \le B_{\text{main}}\eta_U$ and $\|R_J\|_{\text{op}} \le B_{\text{main}}\eta_J$. Since $J_B(\beta^{(j)})/B_{\text{main}} \succeq \kappa_J I_r$ on $\mathcal{E}_B(\xi)$ and $\eta_J < \kappa_J/2$, Weyl's inequality gives $\widehat{J}/B_{\text{main}} \succeq (\kappa_J/2)I_r$ and hence $\left\|\widehat{J}^{-1}\right\|_{\text{op}} \le 2/(\kappa_J B_{\text{main}})$.

A second-order Taylor expansion of the ideal score around $\beta^\star$ gives

$$U_B(\beta^{(j)}) = U_B(\beta^\star) - J_B(\beta^\star)e + H_B e + R_B(e),$$

where $H_B = \dot{U}_B(\beta^\star) + J_B(\beta^\star)$, $\|H_B\|_{\text{op}} \le B_{\text{main}}\xi$ on $\mathcal{E}_B(\xi)$, and $\|R_B(e)\| \le CB_{\text{main}}\|e\|^2$ by bounded score derivatives. For the unprojected proposal, substituting into the update yields

$$\beta^{(j+1)} - \beta^\star = \widehat{J}^{-1}U_B(\beta^\star) + \{I_r - \widehat{J}^{-1}J_B(\beta^\star)\}e + \widehat{J}^{-1}H_B e + \widehat{J}^{-1}R_B(e) + \widehat{J}^{-1}R_U.$$

The first and last terms are bounded by constants times ξ and η_U. For the second term, use

$$I_r - \widehat{J}^{-1}J_B(\beta^\star) = \widehat{J}^{-1}\{\widehat{J} - J_B(\beta^\star)\},$$

and decompose $\widehat{J} - J_B(\beta^\star) = R_J + J_B(\beta^{(j)}) - J_B(\beta^\star)$. Local Lipschitz continuity gives $\left\|J_B(\beta^{(j)}) - J_B(\beta^\star)\right\|_{\text{op}} \le CB_{\text{main}}\|e\|$, so this contribution is $O(\eta_J\|e\| + \|e\|^2)$. The Hessian-fluctuation term is $O(\xi\|e\|)$, and the Taylor remainder is $O(\|e\|^2)$. Combining the bounds proves

equation (40) for the ordinary proposal. If the optional projection in Appendix G.2 is used, convexity of Θ implies that projection cannot increase the distance to $\beta^\star$, so the same bound continues to hold.

When $\|e\| \le r_0$, the two linear products satisfy $(\xi + \eta_J)\,\|e\| \le r_0\xi + r_0\eta_J \le \max\{1, r_0\}\zeta$. After adjusting constants, equation (40) therefore implies the simpler recursion

$$\|e_{j+1}\| \le \bar{a}_0\zeta + \bar{a}_1\,\|e_j\|^2 .$$

For completeness, if $\|e_j\| \le C_j(\zeta + \|e_0\|^{2^j})$ and $\zeta \le 1$, then

$$\|e_{j+1}\| \le \bar{a}_0\zeta + 2\bar{a}_1 C_j^2(\zeta^2 + \|e_0\|^{2^{j+1}}) \le C_{j+1}(\zeta + \|e_0\|^{2^{j+1}}).$$

The base case holds with $C_0 = 1$. Since J is fixed, choose the initial error and ζ small enough that the displayed bound is below r_0 for every $j \le J$, with $\mathcal{B}_0$ contained in the interior of Θ. Induction proves that all ordinary proposals stay in this ball and equation (41) holds. Any optional safeguard is therefore inactive on this event. □

Proof of Proposition 2. The design-regret statement follows from Lemma 14 and equation (33). For the recovery statement, intersect the main-sample event from Lemma 15 with the update-computation event in equation (36). Conditional on a pre-main initialization and design satisfying the small-error conditions, each event has failure probability at most $\delta/2$, so their intersection has probability at least $1 - \delta$. Apply Lemma 17 with $\xi = \xi_B(\delta)$ and $\zeta = \xi_B(\delta) + \eta_U(\delta) + r_0\eta_J(\delta)$. On this event,

$$\left\|\beta^{(J)} - \beta^\star\right\|^2 \le 2C_J^2\left\{[\xi_B(\delta) + \eta_{\mathrm{upd}}(\delta)]^2 + \epsilon_{\mathrm{init}}^{2^{J+1}}\right\}.$$

By Lemma 1, the Hessian of A is uniformly bounded above on the local neighborhood. The exact Bregman form of KL in the proof of Lemma 12 therefore gives

$$D_{\mathrm{KL}}(Q_{\beta^\star}\|Q_{\beta^{(J)}}) \le \frac{f_{\max}}{2}\left\|\beta^{(J)} - \beta^\star\right\|^2 .$$

Absorbing constants proves equation (9). Lemma 16 supplies the stated Monte Carlo rates for a concrete implementation of Algorithm 1. □

GAMD therefore approaches the oracle design as its initialization and information estimates improve. Its final recovery error is controlled by main-sample noise, update computation, and the initialization error after J scoring steps. Appendix G.1 uses the explicit sampling bounds to choose Monte Carlo sizes for asymptotic optimality.

F Identification and the Role of Conditional Reference Information

The target law is specified by the shift parameter β, but each target observation reveals only the coordinates in an allowed panel S. We first characterize when these partial observations identify β. We then study how replacing the full reference conditional law with a linear predictor or low-dimensional marginals changes the information assigned to a panel.

F.1 Identification under Partial Panels

Even with a known reference law Q_0, different target joint distributions can give exactly the same distribution for every observed panel. The first result constructs such a pair. We then show how restricting the target to an identifiable exponential-tilt family avoids this ambiguity.

The observable function space collects every square-integrable statistic that can be computed from an admissible panel. Write

$$\mathcal{O} = \overline{\mathrm{span}}\{g(X_S) : S \in \mathcal{S},\ g \in L_2(Q_{0,S})\} \subseteq L_2(Q_0),$$

where the closure is in $L_2(Q_0)$, and let $L_2^0(Q_0)$ denote the mean-zero subspace.

Proposition 5 (Impossibility for unrestricted target shifts)**.** *Suppose there exists a bounded nonzero $h \in L_2^0(Q_0) \cap \mathcal{O}^\perp$. For every sufficiently small $\varepsilon > 0$, the probability laws*

$$\frac{dQ_\pm}{dQ_0} = 1 \pm \varepsilon h$$

are distinct but have identical marginals on every $S \in \mathcal{S}$. Moreover,

$$D_{\mathrm{KL}}(Q_\pm \| Q_0) = \frac{\varepsilon^2}{2}\mathbb{E}_{Q_0}[h^2] + O(\varepsilon^3).$$

Proof. Choose $\varepsilon < \|h\|_\infty^{-1}$, so $1 \pm \varepsilon h$ is nonnegative. Because $\mathbb{E}_{Q_0} h = 0$, both densities integrate to one. They are distinct because h is nonzero.

Fix a panel S and any bounded measurable $g(X_S)$. Orthogonality gives $\mathbb{E}_{Q_0}[h(X)g(X_S)] = 0$. Equivalently, $\mathbb{E}_{Q_0}[h(X) \mid X_S] = 0$ almost surely. The marginal density ratio of $Q_\pm$ relative to $Q_{0,S}$ is therefore

$$\mathbb{E}_{Q_0}[1 \pm \varepsilon h(X) \mid X_S] = 1.$$

Thus $(Q_+)_S = (Q_-)_S = Q_{0,S}$ for every admissible panel.

Finally, using $(1+t)\log(1+t) = t + t^2/2 + O(t^3)$ uniformly for bounded t near zero,

$$\begin{aligned} D_{\mathrm{KL}}(Q_\pm \| Q_0) &= \mathbb{E}_{Q_0}[(1 \pm \varepsilon h)\log(1 \pm \varepsilon h)] \\ &= \pm\varepsilon \mathbb{E} h + \frac{\varepsilon^2}{2}\mathbb{E} h^2 + O(\varepsilon^3 \mathbb{E}|h|^3). \end{aligned}$$

The linear term vanishes, proving the expansion. □

For example, let X_1, X_2 be independent uniform signs under Q_0 and allow only singleton panels. Taking $h(X) = X_1 X_2$ gives targets with identical uniform marginals but correlations $\mathbb{E}_{Q_\pm}[X_1 X_2] = \pm\varepsilon$. No number of singleton measurements distinguishes the two correlations. In the specified relative-shift family, direct coverage of every feature instead gives the following global identification result.

Proposition 6 (Global identification under direct feature coverage)**.** *Suppose the exponential family is minimal and, for each coordinate a, there is an observed panel S such that $\phi_a(X)$ is measurable with respect to X_S. If $Q_{\beta,S} = Q_{\beta',S}$ for every such panel, then $\beta = \beta'$.*

Proof. Equality of the relevant panel marginals implies equality of the expectation of every directly observable feature: $\mathbb{E}_{Q_\beta}\phi_a(X) = \mathbb{E}_{Q_{\beta'}}\phi_a(X)$ for all a. Hence $\nabla A(\beta) = \nabla A(\beta')$. By Lemma 1, A has positive-definite Hessian everywhere, so it is strictly convex. The gradient of a differentiable strictly convex function is injective. Therefore $\beta = \beta'$. □

Direct coverage identifies β by recovering its feature means. The information condition used by the optimality theorem is weaker: for every nonzero direction v, at least one acquired panel must have a nonzero conditional score $\mathbb{E}_\beta[v^\top s_\beta(X) \mid X_S]$. Such information can be available even when the panel does not directly measure the corresponding feature.

Proposition 7 (Local observability)**.** *For fixed β and allocation p, $M_\beta(p) \succ 0$ if and only if there is no nonzero $v \in \mathbb{R}^r$ satisfying*

$$\mathbb{E}_{Q_\beta}[v^\top s_\beta(X) \mid X_S] = 0 \quad \textit{almost surely for every } S \textit{ with } p_S > 0.$$

Proof. For any $v \in \mathbb{R}^r$,

$$\begin{aligned} v^\top M_\beta(p) v &= \sum_{S \in \mathcal{S}} \frac{p_S}{c(S)} v^\top I_S(\beta) v \\ &= \sum_{S \in \mathcal{S}} \frac{p_S}{c(S)} \mathbb{E}_{Q_\beta}\left[\mathbb{E}_{Q_\beta}[v^\top s_\beta(X) \mid X_S]^2\right]. \end{aligned}$$

Every term is nonnegative. The sum is zero exactly when every active-panel conditional expectation is zero almost surely. Thus $M_\beta(p)$ has a nontrivial null direction exactly under the stated condition. □

The initialization assumption guarantees direct estimates of the feature means. By contrast, $M_\beta(p) \succ 0$ only requires the observed panels to distinguish local parameter changes. For instance, observing a variable correlated with an unmeasured scalar shift feature can give positive information about that feature without direct coverage.

F.2 Why Full Conditional Information Matters

GAMD values a panel through the conditional mean $m_{\beta,S}(X_S) = \mathbb{E}_\beta[s_\beta(X) \mid X_S]$. A regression model restricted to a smaller class of functions may be unable to represent this mean. We first express the resulting information loss as a prediction error, then examine linear predictors and rules that use only low-dimensional reference marginals.

To quantify information lost by a restricted predictor, first consider all square-integrable functions of the observed panel. For any $v \in \mathbb{R}^r$,

$$v^\top I_S(\beta) v = \sup_{g \in L_2(X_S)} \left\{ 2\mathbb{E}[g(X_S) v^\top s_\beta(X)] - \mathbb{E}[g(X_S)^2] \right\}. \tag{42}$$

The maximum is attained by $g(X_S) = v^\top m_{\beta,S}(X_S)$, the conditional mean of the full score in direction v.

Proof of the variational identity. Fix $v \in \mathbb{R}^r$ and define $Z = v^\top s_\beta(X)$ and $h(X_S) = \mathbb{E}[Z \mid X_S] = v^\top m_{\beta,S}(X_S)$. For every square-integrable $g(X_S)$,

$$\mathbb{E}[g(X_S)Z] = \mathbb{E}[g(X_S)\mathbb{E}(Z \mid X_S)] = \mathbb{E}[g(X_S)h(X_S)].$$

Consequently,

$$\begin{aligned} 2\mathbb{E}[gZ] - \mathbb{E}[g^2] &= 2\mathbb{E}[gh] - \mathbb{E}[g^2] \\ &= \mathbb{E}[h^2] - \mathbb{E}[(g-h)^2]. \end{aligned}$$

The expression is maximized uniquely up to $Q_{\beta,S}$-null sets in $L_2(X_S)$ by $g = h$, and its maximum equals $\mathbb{E}[h^2] = v^\top I_S(\beta) v$, because $\mathbb{E}[h] = \mathbb{E}[Z] = 0$. This proves equation (42). □

Restricting the predictor class can only reduce this maximum. The next result identifies the reduction as a squared prediction error.

Proposition 8 (Predictor-class information gap)**.** *Let $\mathcal{G}_S \subseteq L_2(X_S)$ contain the zero function, and define*

$$\mathcal{I}_{S,\mathcal{G}}(v) = \sup_{g \in \mathcal{G}_S} \{ 2\mathbb{E}[g(X_S) v^\top s_\beta(X)] - \mathbb{E}[g(X_S)^2] \}.$$

Then

$$v^\top I_S(\beta) v - \mathcal{I}_{S,\mathcal{G}}(v) = \inf_{g \in \mathcal{G}_S} \mathbb{E}\left[g(X_S) - \mathbb{E}[v^\top s_\beta(X) \mid X_S] \right]^2.$$

For a square-integrable vector predictor $g(X_S) \in \mathbb{R}^r$, define

$$\mathcal{J}_S(g) = \mathbb{E}[g s_\beta^\top + s_\beta g^\top - g g^\top].$$

Then

$$I_S(\beta) - \mathcal{J}_S(g) = \mathbb{E}[(m_{\beta,S} - g)(m_{\beta,S} - g)^\top] \succeq 0. \tag{43}$$

Proof. The scalar identity follows from the square-completion calculation for equation (42), now taking the supremum only over $\mathcal{G}_S$. For the matrix identity, use the fact that g is X_S-measurable: $\mathbb{E}[g s_\beta^\top] = \mathbb{E}[g m_{\beta,S}^\top]$ and similarly for the transpose. Therefore

$$\begin{aligned} I_S - \mathcal{J}_S(g) &= \mathbb{E}[m_{\beta,S} m_{\beta,S}^\top - g m_{\beta,S}^\top - m_{\beta,S} g^\top + g g^\top] \\ &= \mathbb{E}[(m_{\beta,S} - g)(m_{\beta,S} - g)^\top]. \end{aligned}$$

The right-hand side is positive semidefinite. □

The variational value falls short of the true information by exactly the conditional-score prediction error. The matrix $\mathcal{J}_S(g)$ is a lower bound, not necessarily a covariance matrix. For the best linear predictor, orthogonality gives the more direct covariance decomposition below.

Corollary 2 (Linear and nonlinear information decomposition)**.** *Let $Z_S = z_S(X_S)$ be centered and square integrable, and let*

$$g_{\mathrm{lin}}(X_S) = \mathrm{Cov}(s_\beta, Z_S)\,\mathrm{Cov}(Z_S)^\dagger Z_S$$

be the best linear predictor of the full score from Z_S, where † denotes the Moore–Penrose pseudoinverse. Then

$$I_S(\beta) = \mathrm{Cov}(g_{\mathrm{lin}}) + \mathbb{E}[(m_{\beta,S} - g_{\mathrm{lin}})(m_{\beta,S} - g_{\mathrm{lin}})^\top]. \tag{44}$$

Thus covariance-based information is always dominated by the full conditional information, with equality exactly when the conditional score lies in the chosen linear span.

Proof. The linear predictor is the orthogonal projection of s_β onto the closed linear span generated by Z_S. Because this span is contained in the X_S-measurable subspace and $m_{\beta,S}$ is the orthogonal projection of s_β onto the full X_S-measurable subspace, the residual $m_{\beta,S} - g_{\mathrm{lin}}$ is orthogonal to g_{lin}. Therefore the covariance of $m_{\beta,S} = g_{\mathrm{lin}} + (m_{\beta,S} - g_{\mathrm{lin}})$ decomposes as in equation (44). The residual covariance vanishes if and only if the two projections coincide almost surely. □

Linear prediction is sufficient when the conditional score is linear in the chosen observed features. To see how large the allocation error can be otherwise, we construct reference laws with identical low-dimensional marginals but different optimal panels.

A design rule is called *q-marginal-based* if its output depends on a reference distribution only through all coordinate marginals involving at most q coordinates. This class includes covariance-based rules when $q = 2$, and is stronger than merely observing moments up to order q because it grants the rule the complete low-dimensional marginal distributions.

Theorem 2 (Fixed-order marginals do not determine an optimal allocation)**.** *For any integers $q \geq 1$ and $K \geq 2$ and any $\rho \in (0, 1/2)$, there exist K full-support reference distributions and K unit-cost panels such that all reference distributions have identical coordinate marginals of order at most q, but every q-marginal-based design has expected approximation ratio at least K in at least one world j:*

$$\mathbb{E}\left[\frac{\Phi_j(P)}{\Phi_j^\star}\right] \geq K. \tag{45}$$

Here P is the rule's random allocation, Φ_j is the population criterion in world j, and $\Phi_j^\star$ its minimum. The expectation allows the design rule itself to be randomized. The comparison concerns panel valuation at a fixed parameter using the specified reference summaries; the allocation receives no additional full-reference information or target data that distinguish the worlds.

Proof. Let

$$X = (H, U_{1,1:q}, \ldots, U_{K,1:q}),$$

where $U_{k,1:q} = (U_{k,1}, \ldots, U_{k,q})$ and all $U_{k,\ell}$ are independent Rademacher variables, each equal to -1 or 1 with probability $1/2$. For world $j \in \{1, \ldots, K\}$, let ξ be independent with $\mathbb{P}(\xi = -1) = \rho$ and define

$$H = \xi \prod_{\ell=1}^{q} U_{j,\ell}.$$

Because both values of ξ have positive probability, every configuration of (H, U) has positive probability; hence each reference distribution has full support.

We first verify that all marginals involving at most q coordinates are identical across worlds. A marginal not containing H consists only of independent Rademacher variables and is therefore the same in every world. Now consider a set of at most q coordinates that contains H. It contains at most $q - 1$ variables from the special block $U_{j,1:q}$, leaving at least one factor in the parity product unobserved. Conditional on all selected U variables, that unobserved factor is still an independent uniform Rademacher variable. Consequently the product $\prod_{\ell=1}^{q} U_{j,\ell}$, and hence H, remains uniform and independent of the selected coordinates. Thus every marginal of order at most q is the independent Rademacher law, regardless of j.

Use the scalar shift feature $\phi(X) = H$ and evaluate panel information at $\beta = 0$. Then $s_0(X) = H$ and $F(0) = \mathrm{Var}(H) = 1$. Let panel S_k reveal the entire block $U_{k,1:q}$ and have cost one. In world j,

$$\mathbb{E}[H \mid U_{j,1:q}] = (1-2\rho)\prod_{\ell=1}^{q} U_{j,\ell},$$

so $I_{S_j}(0) = a := (1-2\rho)^2$. For $k \neq j$, $U_{k,1:q}$ is independent of H, giving $I_{S_k}(0) = 0$. Therefore, for an allocation $p \in \Delta_K$ in world j,

$$M_j(p) = ap_j, \qquad \Phi_j(p) = \frac{1}{ap_j}, \qquad \Phi_j^\star = \frac{1}{a}.$$

The approximation ratio is $1/p_j$.

A q-marginal-based rule receives exactly the same input in all K worlds, so the distribution of its random output P is also the same. Because $\sum_j \mathbb{E}P_j = 1$, some world j satisfies $\mathbb{E}P_j \leq 1/K$. By Jensen's inequality for $x \mapsto 1/x$,

$$\mathbb{E}\left[\frac{\Phi_j(P)}{\Phi_j^\star}\right] = \mathbb{E}\left[\frac{1}{P_j}\right] \geq \frac{1}{\mathbb{E}P_j} \geq K.$$

If $P_j = 0$ with positive probability, the left-hand side is infinite and the inequality remains true. □

Knowing all marginals of a fixed order can therefore leave the best panel completely undetermined. Complete reference samples can reveal the higher-order dependence, whether it is represented by a conditional generator or a sufficiently expressive score predictor. The result distinguishes what information the acquisition rule receives, not the architecture used to process it.

A score predictor fitted for one feature map need not answer queries for new features. We now consider a reference model that can answer conditional-score queries for every function in a sufficiently rich class. Together with the unconditional feature means, these answers determine the entire conditional distribution.

Proposition 9 (Conditional-score queries determine the conditional law)**.** *Fix a panel S and suppose $\mathcal{X}_{S^c}$ is a standard Borel space. Let $\mathcal{H}$ be a countable class of bounded functions on $\mathcal{X}_{S^c}$ that is measure determining, meaning equality of their integrals identifies a probability law. At $\beta = 0$, knowledge of the conditional scores associated with scalar features $\phi_h(X) = h(X_{S^c})$ for all $h \in \mathcal{H}$, together with their unconditional means, uniquely determines $Q_0(X_{S^c} \in \cdot \mid X_S = x_S)$ for $Q_{0,S}$-almost every x_S. Conversely, the conditional law together with the same unconditional means determines all such conditional scores.*

Proof. For the scalar feature $\phi_h(X) = h(X_{S^c})$, the complete score at $\beta = 0$ is $h(X_{S^c}) - \mathbb{E}_{Q_0} h(X_{S^c})$. The observed score is therefore

$$m^h_{0,S}(x_S) = \mathbb{E}_{Q_0}[h(X_{S^c}) \mid X_S = x_S] - \mathbb{E}_{Q_0} h(X_{S^c}).$$

Knowing the score and the unconditional mean gives the conditional integral

$$\int h(z) Q_0(dz \mid x_S)$$

for every $h \in \mathcal{H}$. Because $\mathcal{H}$ is measure determining, two probability measures with the same integrals for all $h \in \mathcal{H}$ must coincide. Countability allows the exceptional null sets for individual h to be combined into one $Q_{0,S}$-null set. Hence the collection of conditional scores determines the regular conditional law almost surely. The converse follows by integrating each h against the conditional law and subtracting its given unconditional mean. □

Once the reference conditional law is known, Proposition 3 provides the conditional law for any post-hoc pair (ϕ, β). This is the precise sense in which a frozen conditional generator is reusable across different target campaigns and shift representations.

This explains the distinction between a task-specific score predictor and a reusable conditional model. Both can suffice for a fixed feature map. The conditional law also determines the answers for newly specified features and target parameters; the equivalence itself makes no claim about which representation is faster to evaluate.

G Monte Carlo Approximation and Computational Accuracy

GAMD uses generated samples to estimate the reference log-partition function, panel information, and the scores used in the final updates. We first give sample sizes sufficient for Assumption 3. We then show that the same sampling regime, together with simple finite-sample safeguards, verifies the tail condition needed for expected KL risk. Finally, we study fixed completion budgets and biased conditional-score approximations. The reference law Q_0 is fixed throughout; all errors here are computational.

G.1 Monte Carlo Sample Sizes

We state explicit sufficient conditions for Assumption 3, keeping r, $\mathcal{S}$, and J fixed as B grows. For initialization, require

$$\begin{aligned} \sup_{\beta\in\Theta} \|\nabla \widehat{A}_{R_0}(\beta) - \nabla A(\beta)\|_2 &= O_p(B^{-\gamma/2}), \\ \sup_{\beta\in\Theta} \|\nabla^2 \widehat{A}_{R_0}(\beta) - F(\beta)\|_{\mathrm{op}} &= o_p(1). \end{aligned} \tag{46}$$

Lemma 5 then preserves the target-sample initialization rate. For design, require

$$\|\widehat{F} - F(\widetilde{\beta}_B)\|_{\mathrm{op}} + \max_{S\in\mathcal{S}} \|\widehat{I}_S - I_S(\widetilde{\beta}_B)\|_{\mathrm{op}} = o_p(1), \tag{47}$$

and a numerical objective gap tending to zero in probability. These estimates are evaluated at $\widetilde{\beta}_B$; they need not converge faster than the final parameter error. Lemma 8 gives the resulting convergence of the learned design value and information matrix, without requiring a unique allocation vector.

For the final updates, let $U_B(\beta)$ and $J_B(\beta)$ be the exact score and information totals defined in Section 4.3. Sufficient bounds are

$$\begin{aligned} \max_{j<J} \|\widehat{U}_B^{(j)} - U_B(\beta^{(j)})\|_2 &= o_p(\sqrt{B}), \\ \max_{j<J} B^{-1} \|\widehat{J}_B^{(j)} - J_B(\beta^{(j)})\|_{\mathrm{op}} &= o_p(B^{-(1-\gamma)/2}). \end{aligned} \tag{48}$$

The bounds apply at the successive iterates, uniformly over the local region used in the scoring proof. In Assumption 3, asymptotic equivalence of a step means that its difference from the ideal Fisher-scoring step is $o_p(B^{-1/2})$. Lemma 11 verifies this property from equation (48). Score error changes the step directly, while information error also multiplies the current parameter error.

These rates are enough for design consistency and asymptotic normality. Expected KL risk also needs the tail condition equation (31), because convergence in probability does not control rare large simulation errors. We return to this point in Appendix G.3. We first establish the sampling bounds that make Assumption 3 achievable.

The reference sample pool is held fixed while optimizing $\widehat{A}_{R_0}$. The next lemma controls its gradient and curvature errors uniformly over the parameter set, so the bound applies at the optimizer as well as at a fixed parameter.

Lemma 18 (Monte Carlo approximation of the log-partition function)**.** *Let* $X_0^{(1)}, \dots, X_0^{(R_0)} \overset{\text{i.i.d.}}{\sim} Q_0$ *and define*

$$\widehat{A}_{R_0}(\beta) = \log\left\{ \frac{1}{R_0} \sum_{m=1}^{R_0} e^{\beta^\top \phi(X_0^{(m)})} \right\}.$$

On compact Θ*,* $\widehat{A}_{R_0}$*, its gradient, and its Hessian converge uniformly in probability to* A*,* ∇A*, and* F*, respectively. Under bounded* ϕ*, their uniform errors are* $\widetilde{O}_p(R_0^{-1/2})$ *for fixed* r *and compact* Θ*.*

Proof. The functions $e^{\beta^\top\phi(x)}$, $\phi(x)e^{\beta^\top\phi(x)}$, and $\phi(x)\phi(x)^\top e^{\beta^\top\phi(x)}$ are uniformly bounded and Lipschitz in β on compact Θ. Cover Θ by a finite Euclidean net and apply scalar or matrix Hoeffding inequalities at the net points. A union bound and the Lipschitz property extend the bounds uniformly over Θ. The denominator $R_0^{-1}\sum_m e^{\beta^\top\phi(X_0^{(m)})}$ is uniformly bounded away from zero because $\beta^\top\phi$

is bounded. To see explicitly how the three averages enter the derivatives, put $\phi_m = \phi(X_0^{(m)})$ and define

$$\widehat{Z}_k(\beta) = \frac{1}{R_0}\sum_{m=1}^{R_0} e^{\beta^\top \phi_m}\phi_m^{\otimes k}, \qquad k = 0, 1, 2,$$

where $\phi_m^{\otimes 0} = 1$, $\phi_m^{\otimes 1} = \phi_m$, and $\phi_m^{\otimes 2} = \phi_m\phi_m^\top$. Then

$$\nabla \widehat{A}_{R_0} = \frac{\widehat{Z}_1}{\widehat{Z}_0}, \qquad \nabla^2 \widehat{A}_{R_0} = \frac{\widehat{Z}_2}{\widehat{Z}_0} - \frac{\widehat{Z}_1\widehat{Z}_1^\top}{\widehat{Z}_0^2}.$$

The population quantities have the same expressions with expectations in place of averages. If $K_\Theta = \sup_{\beta\in\Theta}\|\beta\|_2 L_\phi$, both the empirical and population denominators lie in $[e^{-K_\Theta}, e^{K_\Theta}]$. Subtracting each empirical ratio from its population counterpart therefore bounds its error by a constant times the errors in $\widehat{Z}_0, \widehat{Z}_1, \widehat{Z}_2$. Taking a net mesh of polynomial order in R_0^{-1} gives the claimed uniform rate, including the logarithmic covering factor. The Hessian error eventually falls below $f_{\min}/2$ with high probability, supplying the curvature condition in equation (18). □

We next analyze the simulated panel information. Its definition, $I_S(\beta) = \mathrm{Var}_\beta(m_{\beta,S}(X_S))$, requires variation across observed panels and a conditional expectation for each panel value. Outer draws estimate the first quantity; inner completions estimate the second. Using two independent completion batches prevents their simulation variance from being counted as panel information.

Fix β and a panel S, and let $Y = X_S \sim Q_{\beta,S}$. Given Y, draw two conditionally independent batches, for $a \in \{1, 2\}$,

$$X^{(a,1)}, \ldots, X^{(a,L)} \overset{\text{i.i.d.}}{\sim} Q_\beta(\cdot \mid Y),$$

and define

$$\widehat{m}_{S,L}^{(a)} = \frac{1}{L}\sum_{\ell=1}^{L} s_\beta(X^{(a,\ell)}).$$

For $N_g \geq 2$ independent outer observations, write $a_n = \widehat{m}_{S,L,n}^{(1)}$, $b_n = \widehat{m}_{S,L,n}^{(2)}$, and let $\bar{a}, \bar{b}$ be their sample means. The sample cross-covariance in equation (4) means

$$\widehat{\mathrm{Cov}}(a, b) = \frac{1}{N_g - 1}\sum_{n=1}^{N_g}(a_n - \bar{a})(b_n - \bar{b})^\top.$$

Define the raw symmetric estimate as $\widehat{I}_S^{\mathrm{raw}} = \frac{1}{2}\{\widehat{\mathrm{Cov}}(a, b) + \widehat{\mathrm{Cov}}(b, a)\}$. Subtracting any common centering vector from all a_n, b_n leaves this centered sample covariance unchanged. Thus it is identical whether the batches use the exact μ_β or the empirical full-sample mean used by GAMD, even when that mean shares the outer samples. The population identity below is written for the exactly centered scores.

Proposition 10 (Naive variance inflation and cross-completion identity)**.** *For each a,*

$$\mathrm{Var}(\widehat{m}_{S,L}^{(a)}) = I_S(\beta) + \frac{1}{L}\{F(\beta) - I_S(\beta)\}. \tag{49}$$

For the two independent batches,

$$\mathrm{Cov}(\widehat{m}_{S,L}^{(1)}, \widehat{m}_{S,L}^{(2)}) = I_S(\beta) \qquad \textit{for every } L \geq 1. \tag{50}$$

Hence the raw symmetrized cross-covariance $\widehat{I}_S^{\mathrm{raw}}$ in equation (4) *is unbiased. GAMD then applies the PSD projection $\widehat{I}_S = \Pi_{\mathbb{S}_+^r}(\widehat{I}_S^{\mathrm{raw}})$ before optimization.*

Proof. Given Y,

$$\mathbb{E}[\widehat{m}_{S,L}^{(a)} \mid Y] = m_{\beta,S}(Y), \qquad \mathrm{Var}(\widehat{m}_{S,L}^{(a)} \mid Y) = \frac{1}{L}\mathrm{Var}(s_\beta(X) \mid Y).$$

The law of total variance gives

$$\begin{aligned}\mathrm{Var}(\widehat{m}_{S,L}^{(a)}) &= \mathrm{Var}(m_{\beta,S}(Y)) + \frac{1}{L}\mathbb{E}\,\mathrm{Var}(s_\beta(X) \mid Y) \\ &= I_S(\beta) + \frac{1}{L}\{F(\beta) - I_S(\beta)\},\end{aligned}$$

where the last equality uses equation (14).

For the cross-covariance, conditional independence of the two batches gives

$$\mathbb{E}[\widehat{m}_{S,L}^{(1)}\widehat{m}_{S,L}^{(2)\top} \mid Y] = m_{\beta,S}(Y)m_{\beta,S}(Y)^\top.$$

Also $\mathbb{E}\widehat{m}_{S,L}^{(a)} = \mathbb{E}s_\beta(X) = 0$. Taking expectations proves equation (50). The usual sample cross-covariance with denominator $N_g - 1$ is unbiased for the population cross-covariance; symmetrization preserves the expectation. □

Remark 2 (Information geometry after PSD projection)**.** *The finite-sample raw cross-covariance $\widehat{I}_S^{\mathrm{raw}}$ need not be positive semidefinite, so Algorithm 1 always uses $\widehat{I}_S = \Pi_{\mathbb{S}_+^r}(\widehat{I}_S^{\mathrm{raw}})$. If $I_S \succeq 0$ and $\left\|\widehat{I}_S^{\mathrm{raw}} - I_S\right\|_{\mathrm{op}} \le \varepsilon$, Weyl's inequality gives $\left\|\widehat{I}_S - \widehat{I}_S^{\mathrm{raw}}\right\|_{\mathrm{op}} \le \varepsilon$, and hence $\left\|\widehat{I}_S - I_S\right\|_{\mathrm{op}} \le 2\varepsilon$. Thus the projection preserves consistency and changes only constants in the concentration bounds, while ensuring that every estimated panel information matrix is positive semidefinite.*

The raw cross-covariance is unbiased, but it can have negative eigenvalues at finite sample size. PSD projection corrects these eigenvalues while changing the estimation error by at most a constant factor. The next bound controls the projected estimates for all panels simultaneously.

Lemma 19 (Uniform concentration of cross-completion estimates)**.** *Suppose $N_g \ge 2$ independent outer samples are used for each panel and the conditional draws are exact. Let $\widehat{I}_S^{\mathrm{raw}}$ be the centered raw symmetrized sample cross-covariance and $\widehat{I}_S = \Pi_{\mathbb{S}_+^r}(\widehat{I}_S^{\mathrm{raw}})$. There is a constant C_I depending on L_ϕ such that, with probability at least $1 - \delta$,*

$$\max_{S\in\mathcal{S}} \left\|\widehat{I}_S - I_S(\beta)\right\|_{\mathrm{op}} \le C_I \left\{\sqrt{\frac{\log(2r|\mathcal{S}|/\delta)}{N_g}} + \frac{\log(2r|\mathcal{S}|/\delta)}{N_g}\right\}. \tag{51}$$

The same form holds for the sample covariance estimator of $F(\beta)$. The constant can decrease with L through the variance term, although the bounded-difference term does not require $L \to \infty$.

Proof. Let $A_n = \widehat{m}_{S,L,n}^{(1)}$ and $B_n = \widehat{m}_{S,L,n}^{(2)}$. Because $\|s_\beta(X)\|_2 \le 2L_\phi$, both $\|A_n\|_2$ and $\|B_n\|_2$ are at most $2L_\phi$. Consider first the uncentered symmetrized estimator

$$\check{I}_S = \frac{1}{2N_g}\sum_{n=1}^{N_g}(A_nB_n^\top + B_nA_n^\top).$$

By Proposition 10, $\mathbb{E}\check{I}_S = I_S$. Each centered summand is a self-adjoint matrix with operator norm bounded by a constant multiple of L_ϕ^2. Its matrix variance proxy is at most a constant multiple of $N_gL_\phi^4$. Matrix Bernstein's inequality (Tropp, 2012) therefore gives, for a fixed panel,

$$\left\|\check{I}_S - I_S\right\|_{\mathrm{op}} \le CL_\phi^2\left\{\sqrt{\frac{\log(2r/\eta)}{N_g}} + \frac{\log(2r/\eta)}{N_g}\right\}$$

with probability at least $1 - \eta$. A union bound with failure probabilities of order $\delta/|\mathcal{S}|$ makes the result uniform over panels and leaves room for the centering bounds below.

Writing $\bar{A} = N_g^{-1}\sum_n A_n$ and $\bar{B} = N_g^{-1}\sum_n B_n$, the centered raw sample cross-covariance satisfies

$$\widehat{I}_S^{\mathrm{raw}} = \frac{N_g}{N_g - 1}\check{I}_S - \frac{N_g}{2(N_g-1)}(\bar{A}\bar{B}^\top + \bar{B}\bar{A}^\top).$$

The factor $N_g/(N_g-1)$ contributes only $O(N_g^{-1})$. Since the vectors are independent across outer observations and mean zero, $\mathbb{E}\|\bar{A}\|_2^2 \le 4L_\phi^2/N_g$ and similarly for $\bar{B}$. Changing one summand changes the norm of either mean by at most $4L_\phi/N_g$. A scalar bounded-differences bound for that norm, together with Jensen's inequality, therefore gives

$$\|\bar{A}\|_2 \vee \|\bar{B}\|_2 \le \frac{2L_\phi}{\sqrt{N_g}} + CL_\phi\sqrt{\frac{\log(2/\eta)}{N_g}}$$

with probability at least $1-\eta$ after a union bound. Thus the centering term is $O(\log(1/\eta)/N_g)$ and is absorbed into the second term. Finally, since $I_S \succeq 0$, the projection argument in Remark 2 gives

$$\left\|\widehat{I}_S - I_S\right\|_{\mathrm{op}} \le 2\left\|\widehat{I}_S^{\mathrm{raw}} - I_S\right\|_{\mathrm{op}},$$

which only changes the constant C_I. The proof for F is identical with $A_n = B_n = \phi(X_n) - \mu_\beta$. □

Outer observations must be independent within each panel calculation; estimates for different panels may share the same full outer sample pool. The union bound above does not require independence between panels. At fixed L_g, increasing N_g makes both $\widehat{F}$ and every PSD-projected $\widehat{I}_S$ consistent. Increasing L_g reduces completion noise but cannot replace the outer averaging, because the information matrix is the variation of a conditional mean across different observed panels.

The reference-pool and cross-completion bounds, together with the final-update bounds in Lemma 16, now give sufficient sample sizes for Assumption 3. These are generator draws and do not consume the target measurement budget.

Fix any $\nu > 0$, fixed $r, \mathcal{S}, J$, and use independent exact tilted draws with fresh simulations at each update. A sufficient choice is

$$R_0 \ge B^{\gamma+\nu}, \qquad N_g \ge B^{\nu}, \qquad L_g \ge 1 \text{ fixed}, \qquad N_{\mathrm{upd}} \ge B^{1+\nu}, \qquad L_{\mathrm{upd}} \ge B^{\nu},$$

with an optimization gap tending to zero. Integer ceilings may be taken throughout. Lemma 18 gives initialization gradient error $\widetilde{O}_p(B^{-(\gamma+\nu)/2}) = o_p(B^{-\gamma/2})$ and uniform Hessian consistency. Lemma 19 gives consistent design-stage information even with fixed L_g. For the final updates, equation (38) gives

$$\max_{j<J} \frac{\|\widehat{U}_B^{(j)} - U_B(\beta^{(j)})\|}{B} = \widetilde{O}_p(B^{-(1+\nu)/2} + B^{-1}) = o_p(B^{-1/2}),$$

using $B_{\mathrm{main}}/B \to 1$. Similarly, equation (39) gives normalized information error $\widetilde{O}_p(B^{-(1+\nu)/2}) = o_p(B^{-(1-\gamma)/2})$. These choices establish equation (46)–equation (48) and hence the three clauses of Assumption 3.

The comparatively large sufficient N_{upd} is caused by the common score-centering estimate: its error is repeated across all target observations. Independent observation-level completion noise averages across those observations, whereas common centering noise does not. With exact score centering, the information-estimation requirement alone permits smaller outer pools, subject to its stated rate. These counts measure generator computation, not additional target measurements, and therefore do not enter the target acquisition budget.

Thus design consistency allows fixed L_g, while the sharp final efficiency result requires increasingly accurate scores. Appendix G.4 quantifies the variance inflation when the final number of completions is kept fixed.

G.2 Optional Finite-Sample Safeguards

Theorem 1 does not require a particular safeguard. The following rules give one convenient finite-sample implementation and are used in Appendix G.3 to verify the tail condition equation (31).

The matrix inverses in GAMD require positive-definite estimated information. On an exceptional Monte Carlo run where this fails, we use the following fallback rules. If $\widehat{F}$ is not positive definite or no estimated panel mixture is positive definite, use the initialization cost allocation $p_S^{\mathrm{init}} =$

$n_{0,S}c(S)/\sum_T n_{0,T}c(T)$ for the main sample. If $\widehat{J}_B^{(j)}$ is singular, retain the current parameter at that step. Otherwise compute the proposal in equation (6) and project it onto the compact convex parameter set:

$$\beta^{(j+1)} = \Pi_\Theta\Big(\beta^{(j)} + [\widehat{J}_B^{(j)}]^{-1}\widehat{U}_B^{(j)}\Big). \tag{52}$$

Here Π_Θ is Euclidean projection, distinct from the matrix PSD projection in equation (4). It ensures $\|\beta^{(j)} - \beta^\star\|_2 \le \mathrm{diam}(\Theta)$ even on a run with inaccurate information estimates.

The fallback allocation is feasible and has positive-definite population information. To verify the latter claim, suppose $v^\top m_{\beta,S}(X_S) = 0$ for every panel receiving initialization observations. For each feature ϕ_a, a covering panel makes $\phi_a(X)$ measurable with respect to X_S, so

$$\mathbb{E}_\beta[(v^\top s_\beta(X))\phi_a(X)] = \mathbb{E}_\beta[(v^\top m_{\beta,S}(X_S))\phi_a(X)] = 0.$$

The left-hand side is the ath component of $F(\beta)v$. Hence $F(\beta)v = 0$, and minimality gives $v = 0$. Direct coverage and bounded positive costs therefore provide a valid fallback without knowing $\beta^\star$.

Under Assumption 3, the estimated information matrices are nonsingular near $\beta^\star$ with probability tending to one, and the fixed scoring iterates converge to the interior point $\beta^\star$. Hence these safeguards are asymptotically inactive and do not change the Gaussian limit. On the local event used in Proposition 2, the ordinary proposals also remain in the interior, so the finite-budget bound is unchanged.

For expected risk, asymptotic inactivity alone is not enough: rare failures must have sufficiently small contribution to the expectation. Appendix G.3 combines the fallback rules and compact projection with exponentially small failure probabilities to establish the required uniform integrability.

G.3 Verifying the Expected-Risk Tail Condition

The expected-KL conclusion of Theorem 1 assumes equation (31). We now verify it for the safeguarded implementation above. On ordinary runs, the local scoring bounds give root-B error control. On rare ill-conditioned runs, the fallback rules and compact projection keep the estimator bounded. Exponentially small failure probabilities then make the contribution of those runs negligible.

Lemma 20 (Moment bound for the Monte Carlo estimator)**.** *Suppose Assumptions 1–2 hold and $2^J\gamma > 1$. Use the safeguards in Appendix G.2, exact independent full and conditional tilted draws, fresh simulations at each update, and a deterministic design-optimization tolerance tending to zero. For some fixed $\nu > 0$, suppose*

$$\begin{gathered} R_0 \ge B^{\gamma+\nu}, \qquad N_g \ge B^\nu, \qquad L_g \ge 1 \text{ fixed}, \\ N_{\mathrm{upd}} \ge B^{1+\nu}, \qquad L_{\mathrm{upd}} \ge B^\nu. \end{gathered} \tag{53}$$

Then, for every fixed $q > 2$ and all sufficiently large B_0,

$$\sup_{B \ge B_0} \mathbb{E}\Big\|\sqrt{B}(\widehat{\beta}_B - \beta^\star)\Big\|_2^q < \infty. \tag{54}$$

Consequently, the uniform-integrability condition equation (31) *holds.*

Proof. Write $D_\Theta = \sup_{\beta,\beta'\in\Theta}\|\beta - \beta'\|_2$. Compact projection gives $\|\beta^{(j)} - \beta^\star\|_2 \le D_\Theta$ for every simulated outcome. The projection and singular-matrix fallback require no knowledge of $\beta^\star$. We will show that they are inactive on an event with exponentially small failure probability, and use the deterministic bound D_Θ elsewhere.

Initialization. We first bound the initialization error in arbitrary fixed moment order. Bounded features and the fixed coverage lower bound imply, by integrating the concentration inequality in Lemma 5,

$$\mathbb{E}\|\widehat{\mu} - \mu_{\beta^\star}\|_2^k \le C_k n_0^{-k/2} \qquad \text{for every fixed } k \ge 2.$$

The reference-pool gradient at the fixed point $\beta^\star$ has the same kind of bound. Indeed, for $w_m = e^{\beta^{\star\top}\phi(X_0^{(m)})}$,

$$\nabla\widehat{A}_{R_0}(\beta^\star) - \mu_{\beta^\star} = \frac{R_0^{-1}\sum_m w_m\{\phi(X_0^{(m)}) - \mu_{\beta^\star}\}}{R_0^{-1}\sum_m w_m}.$$

The numerator is an average of independent, centered, bounded vectors. Its summands are centered because exponential reweighting changes Q_0 to $Q_{\beta^\star}$. The denominator is bounded below by a positive constant depending only on Θ and L_ϕ. Consequently its kth moment is at most $C_k R_0^{-k/2}$.

The empirical Hessian satisfies $\inf_{\beta\in\Theta} \lambda_{\min}(\nabla^2 \widehat{A}_{R_0}(\beta)) \geq f_{\min}/2$ except on an event of probability at most Ce^{-cR_0}. To verify this fixed-tolerance statement, cover compact Θ by a finite net fine enough that the uniformly bounded third derivatives contribute at most $f_{\min}/4$ between a point and its nearest net point. Bounded-summand concentration for the empirical zeroth, first, and second weighted moments gives the remaining $f_{\min}/4$ bound at every net point; the denominators in their ratios are uniformly positive. The net has a fixed size, so the union bound retains an exponential rate. This is the fixed-tolerance version of Lemma 18.

On the Hessian event, apply the strong-convexity argument in Lemma 5 using the gradient at $\beta^\star$. On its complement use D_Θ. With $x_0 = \|\widetilde{\beta}_B - \beta^\star\|_2$, we obtain

$$\mathbb{E}x_0^k \leq C_k\{n_0^{-k/2} + R_0^{-k/2} + e^{-cR_0}\} \leq C_k B^{-k\gamma/2}. \tag{55}$$

The same argument at a fixed small threshold $t > 0$ gives $\mathbb{P}(x_0 > t) \leq C_t(e^{-c_t n_0} + e^{-c_t R_0})$.

Main-sample and Monte Carlo errors. Next consider the main sample and update computation. Put $N_B = |\mathcal{I}_{\text{main}}|$ and define the nonnegative random error levels

$$\begin{aligned} a_B &= B_{\text{main}}^{-1}\big(\|U_B(\beta^\star)\|_2 + \|\dot{U}_B(\beta^\star) + J_B(\beta^\star)\|_{\text{op}}\big), \\ u_B &= \max_{j<J} B_{\text{main}}^{-1}\|\widehat{U}_B^{(j)} - U_B(\beta^{(j)})\|_2, \\ v_B &= \max_{j<J} B_{\text{main}}^{-1}\|\widehat{J}_B^{(j)} - J_B(\beta^{(j)})\|_{\text{op}}. \end{aligned}$$

Conditional on the initialization and learned allocation, the terms forming a_B are sums of independent centered bounded vectors or matrices. Since $N_B \leq B_{\text{main}}/c_{\min}$ and r is fixed, integrating their concentration bounds gives $\mathbb{E}a_B^k \leq C_k B^{-k/2}$. These constants do not depend on the realized allocation.

For u_B, condition additionally on all randomness before each update, including the main observations. The current parameter is then fixed. The score error is the sum of independent centered conditional-feature completions, minus N_B times the error of the common full-sample feature mean. Integrating bounded-summand concentration for these two averages gives

$$\mathbb{E}u_B^k \leq C_k\{(B_{\text{main}}L_{\text{upd}})^{-k/2} + N_{\text{upd}}^{-k/2}\}. \tag{56}$$

The first term uses independence of all conditional completions across target observations and their repeated draws. The second uses $N_B/B_{\text{main}} \leq c_{\min}^{-1}$. The same bounds apply at adaptive iterates because they hold conditionally with constants uniform on Θ; a maximum over fixed J only changes C_k.

For v_B, the cross-completion concentration in Lemma 19 gives kth moment $C_k N_{\text{upd}}^{-k/2}$ for each PSD-projected panel estimate. The common centering estimate cancels from the centered sample cross-covariance. Moreover,

$$v_B \leq c_{\min}^{-1} \max_{j<J, S\in\mathcal{S}} \|\widehat{I}_S^{(j)} - I_S(\beta^{(j)})\|_{\text{op}}.$$

Taking the maximum over the fixed panel collection and fixed updates therefore yields

$$\mathbb{E}v_B^k \leq C_k N_{\text{upd}}^{-k/2}. \tag{57}$$

No independence between the score-centering and information errors is needed for these bounds.

A high-probability regular event. We now construct the event on which local scoring is well conditioned. Choose a ball $\mathcal{B}_0$ of radius r_0 around $\beta^\star$ contained in the interior of Θ. By the comparison with a fixed nonsingular allocation in Lemma 8, sufficiently small initialization and design-information errors imply $J_B(\beta)/B_{\text{main}} \succeq \kappa_J I_r$ throughout a sufficiently small such ball. This uses continuity in β and the $O(B_{\text{main}}^{-1})$ rounding error. At fixed error thresholds, initialization concentration

and conditional design-sample concentration show that this event fails with probability at most $C(e^{-cn_0}+e^{-cR_0}+e^{-cN_g})$. The deterministic optimization tolerance is below the required threshold for all sufficiently large B.

Intersect this event with sufficiently small fixed bounds on x_0, a_B, u_B, v_B. Bounded-summand concentration at the truth and conditional concentration at each of the fixed updates show that the resulting event, denoted $\mathcal{H}_B$, satisfies

$$\mathbb{P}(\mathcal{H}_B^{\mathrm{c}}) \le Ce^{-cB^a}, \qquad a = \min\{\gamma, \nu\} > 0. \tag{58}$$

For the update errors one can use the conservative fixed-threshold bounds with exponents B_{main} and N_{upd}; increasing L_{upd} only improves them. Conditional union bounds over $j < J$ remain valid even though later iterates depend on earlier simulations.

Moment bound on and off the regular event. On $\mathcal{H}_B$, $v_B < \kappa_J/2$ ensures $\widehat{J}_B^{(j)}/B_{\mathrm{main}} \succeq (\kappa_J/2)I_r$ as long as the current iterate is in $\mathcal{B}_0$. The Taylor calculation in equation (25), now retaining the numerical errors rather than using O_p notation, bounds the unprojected proposal by

$$x_{j+1} \le C\{a_B + u_B + (a_B + v_B)x_j + x_j^2\}, \qquad x_j = \|\beta^{(j)} - \beta^\star\|_2.$$

Projection onto convex Θ cannot increase distance to $\beta^\star$, so the inequality also holds for the projected update. While $x_j \le r_0$, put $d_B = a_B + u_B + r_0 v_B$ and absorb the linear products to obtain $x_{j+1} \le C(d_B + x_j^2)$. Induction over fixed J gives

$$x_J \le C_J\{d_B + x_0^{2^J}\} \quad \text{on } \mathcal{H}_B. \tag{59}$$

Choose the fixed thresholds defining $\mathcal{H}_B$ small enough that this bound, at every $j \le J$, is below r_0. The induction then also proves that every unprojected proposal stays in $\mathcal{B}_0$: no singular fallback or parameter projection is invoked on $\mathcal{H}_B$.

The sample sizes in equation (53) and the moment estimates above imply $\mathbb{E}d_B^q \le C_q B^{-q/2}$. Applying equation (55) at the larger but fixed order $q2^J$, equation (59) yields

$$\begin{aligned}\mathbb{E}[(\sqrt{B}x_J)^q \mathbf{1}_{\mathcal{H}_B}] &\le C_q B^{q/2}\{\mathbb{E}d_B^q + \mathbb{E}x_0^{q2^J}\} \\ &\le C_q\{1 + B^{q(1-2^J\gamma)/2}\}.\end{aligned}$$

This is uniformly bounded because $2^J\gamma > 1$. On the complement, compact projection and equation (58) give

$$\mathbb{E}[(\sqrt{B}x_J)^q \mathbf{1}_{\mathcal{H}_B^{\mathrm{c}}}] \le B^{q/2} D_\Theta^q \mathbb{P}(\mathcal{H}_B^{\mathrm{c}}) \le C_q B^{q/2} e^{-cB^a} \longrightarrow 0.$$

Adding the two parts proves equation (54). Finally, with $Z_B = \sqrt{B}(\widehat{\beta}_B - \beta^\star)$,

$$\mathbb{E}[\|Z_B\|_2^2 \mathbf{1}\{\|Z_B\|_2^2 > K\}] \le K^{-(q-2)/2}\mathbb{E}\|Z_B\|_2^q.$$

The bound is uniform in B and tends to zero as $K \to \infty$, proving uniform integrability. □

The polynomial sample sizes above are sufficient rather than necessary. The key distinction is that Assumption 3 controls the typical root-B behavior, whereas Lemma 20 also controls rare outcomes strongly enough to justify convergence of the expected KL risk.

G.4 Fixed-Completion Efficiency

Suppose each target observation uses a fixed number of conditional completions in the final score. The simulated score is still unbiased, but its variance exceeds the exact panel information. We derive the resulting covariance for estimators that admit a simulated-score asymptotic linearization. Proposition 2 separately controls the finite number of updates used by GAMD.

The centering by μ_β below is exact, or has error negligible on the root-B scale. A shared finite-pool mean instead contributes the common-error term controlled by equation (38).

For an observation $Y = X_S$, let

$$\psi_{\beta,S,L}(Y,\xi) = \frac{1}{L}\sum_{\ell=1}^{L} s_\beta(X^{(\ell)}), \qquad X^{(\ell)} \overset{\text{i.i.d.}}{\sim} Q_\beta(\cdot \mid Y),$$

where ξ denotes the independent simulation randomness. Then $\mathbb{E}_\xi[\psi_{\beta,S,L}(Y,\xi) \mid Y] = m_{\beta,S}(Y)$.

The variance contains both target information and conditional simulation noise. At the truth, the law of total variance and equation (14) give

$$\begin{aligned}\operatorname{Var}(\psi_{\beta^\star,S,L}) &= \operatorname{Var}(\mathbb{E}[\psi_{\beta^\star,S,L} \mid X_S]) + \mathbb{E}\operatorname{Var}(\psi_{\beta^\star,S,L} \mid X_S) \\ &= I_S(\beta^\star) + \frac{1}{L}\{F(\beta^\star) - I_S(\beta^\star)\}.\end{aligned} \tag{60}$$

The added covariance is largest in score directions that the observed panel leaves uncertain. It vanishes for a score fully determined by that panel.

Fix $p \in \Delta_{\mathcal{S}}$ with $M_{\beta^\star}(p) \succ 0$ and counts satisfying $n_{B,S}c(S)/B \to p_S$. Write $M(p) = M_{\beta^\star}(p)$ and $L = (L_S)_{S\in\mathcal{S}}$, where $L_S \geq 1$ is the number of completions for panel S, and define

$$V_L(p) = \sum_S \frac{p_S}{c(S)}\left[I_S(\beta^\star) + \frac{1}{L_S}\{F(\beta^\star) - I_S(\beta^\star)\}\right]. \tag{61}$$

Assume the estimator admits the simulated-score expansion

$$\sqrt{B}(\widehat{\beta}_L - \beta^\star) = M(p)^{-1}\frac{1}{\sqrt{B}}\sum_i \psi_{\beta^\star,S_i,L_{S_i}} + o_p(1). \tag{62}$$

This is the regularity requirement on the estimating algorithm in the following result; conditional unbiasedness alone does not establish it.

Theorem 3 (Fixed-completion sandwich covariance)**.** *Consider an estimator $\widehat{\beta}_L$ satisfying equation* (62)*, with conditionally unbiased scores ψ_{β,S,L_S} and independent simulation randomness across observations. Assume differentiation may be interchanged with conditional simulation expectation and the normalized Jacobian converges to $-M_{\beta^\star}(p)$. Then*

$$\sqrt{B}(\widehat{\beta}_L - \beta^\star) \Rightarrow \mathcal{N}\big(0, M_{\beta^\star}(p)^{-1}V_L(p)M_{\beta^\star}(p)^{-1}\big). \tag{63}$$

If $\{B\|\widehat{\beta}_L - \beta^\star\|_2^2\}$ is uniformly integrable, the corresponding expected full-KL risk has leading term

$$\frac{1}{2B}\operatorname{tr}\big[F(\beta^\star)M(p)^{-1}V_L(p)M(p)^{-1}\big]. \tag{64}$$

Proof. For each target observation, the simulated score is conditionally unbiased for the exact panel score. Therefore the negative population score derivative, after budget normalization, is $M_{\beta^\star}(p)$, as for the exact observed-score equation. By equation (60), the budget-normalized variance of the summed simulated scores converges to $V_L(p)$. The Lindeberg condition follows from bounded ϕ for fixed L_S. Using the assumed simulated-score asymptotic linearization gives

$$\sqrt{B}(\widehat{\beta}_L - \beta^\star) = M(p)^{-1}\frac{1}{\sqrt{B}}\sum_i \psi_{\beta^\star,S_i,L_{S_i}} + o_p(1),$$

which implies equation (63) by the triangular-array central limit theorem. Applying the local KL expansion in Lemma 12 and the stated uniform-integrability condition yields equation (64). □

Since $F - I_S \succeq 0$, $V_L(p) \succeq M(p)$, so finite simulation cannot improve efficiency. If $\min_{S:p_S>0} L_S \to \infty$, then $V_L(p) \to M(p)$ and the ideal covariance $M(p)^{-1}$ is recovered. If L_S is fixed, the practically optimal allocation may instead minimize the sandwich criterion in equation (64); GAMD uses the ideal criterion when conditional simulation is sufficiently accurate or inexpensive relative to target measurement.

The inverse information $M(p)^{-1}$ describes sensitivity of the estimating equation, and $V_L(p)$ is the covariance of its simulated scores. Fixed unbiased completion noise therefore increases the first-order variance. A systematic error in the conditional-score mean can instead shift the limiting parameter, which we study next.

G.5 Approximate Conditional Scores

Conditional-score approximations need not be unbiased, for example when a learned predictor replaces exact conditional sampling. We compare such approximations with the exact score under the same fixed law Q_0. A constant error cancels from a covariance estimate but not from a likelihood score, so its effects on design and on final estimation must be treated separately.

Fix β and S, and let $m_S = m_{\beta,S}(X_S)$. Suppose an approximate procedure returns $\widetilde{m}_S(X_S)$. Because information is a covariance, define the centered approximation error

$$e_S = \{\widetilde{m}_S - \mathbb{E}\widetilde{m}_S\} - m_S, \qquad \mathbb{E}\,\|e_S\|_2^2 \le \varepsilon_S^2.$$

Let $\widetilde{I}_S = \mathrm{Cov}(\widetilde{m}_S)$. Here expectations are under Q_β, so the comparison holds at a common parameter and common outer law. Centering isolates the part relevant to a covariance: adding a constant vector to every predicted score cannot change the panel information.

Proposition 11 (Conditional-score error implies information error)**.** *Under the preceding definition,*

$$\left\|\widetilde{I}_S - I_S\right\|_{\mathrm{op}} \le 2\sqrt{\|I_S\|_{\mathrm{op}}}\,\varepsilon_S + \varepsilon_S^2. \tag{65}$$

Proof. Because m_S is centered,

$$\widetilde{I}_S = \mathbb{E}[(m_S + e_S)(m_S + e_S)^\top].$$

Hence

$$\widetilde{I}_S - I_S = \mathbb{E}[m_S e_S^\top + e_S m_S^\top + e_S e_S^\top].$$

For unit vectors u, v,

$$\left|u^\top \mathbb{E}[m_S e_S^\top] v\right| \le \sqrt{\mathbb{E}(u^\top m_S)^2}\sqrt{\mathbb{E}(v^\top e_S)^2} \le \sqrt{\|I_S\|_{\mathrm{op}}}\,\varepsilon_S.$$

Thus each cross term has operator norm at most $\sqrt{\|I_S\|_{\mathrm{op}}}\varepsilon_S$. The final term has norm at most $\mathbb{E}\,\|e_S\|^2 \le \varepsilon_S^2$. The triangle inequality proves equation (65). □

The information error is linear in small conditional-score error, with a quadratic correction. Combining it with the inverse-information perturbation bound gives the corresponding allocation guarantee.

Corollary 3 (Design stability under approximate conditional scores)**.** *At fixed β, abbreviate $F = F(\beta)$ and $M(p) = M_\beta(p)$. Suppose the oracle and approximate optimizers lie in a common region where $M(p) \succeq \kappa I_r$, and the information perturbation is small enough to be below $\kappa/2$. If $\left\|\widetilde{F} - F\right\|_{\mathrm{op}} \le \varepsilon_F$ and $\max_S \mathbb{E}\,\|e_S\|^2 \le \varepsilon^2$, then an ϵ_{opt}-optimal design for the approximate objective satisfies*

$$\Phi(\widetilde{p}) - \Phi(p^\star) \le C\left(\varepsilon_F + \varepsilon + \varepsilon^2 + \epsilon_{\mathrm{opt}}\right)$$

for a constant C depending on κ, F, the costs, and the maximal panel information.

Proof. Proposition 11 converts the L_2 score error into a uniform operator-norm information error of order $\varepsilon + \varepsilon^2$. Lemma 7 then bounds the uniform objective error on the nonsingular region. The standard two-sided comparison between the approximate optimizer and the exact optimizer, used in Lemma 14, adds the optimization gap and proves the result. □

A conditional sampler and a direct score predictor with the same error receive the same design guarantee. The additional benefit of a conditional-law representation is its ability to answer queries for new shift features, as established in Proposition 9.

Approximation during the final update is more delicate because a systematic score error can bias the estimator, rather than merely change the allocation. Let $\widetilde{m}_{\beta,S} = m_{\beta,S} + a_{\beta,S}$ and define the normalized population bias under allocation p by

$$b_p = \sum_S \frac{p_S}{c(S)} \mathbb{E}_{Q_{\beta^\star,S}}[a_{\beta^\star,S}(X_S)].$$

Proposition 12 (Effect of biased final conditional scores)**.** *Suppose* $M_{\beta^\star}(p) \succ 0$ *and* $\max_S \|\mathbb{E}a_{\beta^\star,S}(X_S)\| \le \varepsilon_G$*. Assume the approximate score has bounded local second moments and locally Lipschitz derivatives, and its empirical score and derivative fluctuations have the usual* $B^{-1/2}$ *order. Let* $\delta_G = o(1)$ *bound the discrepancy between its normalized population Jacobian and the ideal population-score Jacobian on a local neighborhood; the latter equals* $-M_{\beta^\star}(p)$ *at the truth. Suppose the panel counts equal* $Bp_S/c(S) + O(1)$ *and the approximate estimating equation has a consistent local root* $\widehat{\beta}$*. Then*

$$\widehat{\beta} - \beta^\star = M_{\beta^\star}(p)^{-1} b_p + O_p(B^{-1/2} + \varepsilon_G^2 + \delta_G \varepsilon_G), \tag{66}$$

and therefore $D_{\mathrm{KL}}(Q_{\beta^\star} \| Q_{\widehat{\beta}}) = O_p(B^{-1} + \varepsilon_G^2)$*. For the leading bias not to affect root-*B *centering, the aggregate bias must be negligible on that scale;* $\varepsilon_G = o(B^{-1/2})$ *is a sufficient, not necessary, condition because panel biases can cancel. Matching the ideal efficiency constant additionally requires negligible centered approximation noise.*

Proof. The normalized approximate estimating equation has the expansion

$$0 = W_B + b_p - M_{\beta^\star}(p)(\widehat{\beta} - \beta^\star) + R_B,$$

where $W_B = B^{-1} \sum_i \{\widetilde{m}_{\beta^\star,S_i}(Y_i) - \mathbb{E}\widetilde{m}_{\beta^\star,S_i}(Y_i)\} = O_p(B^{-1/2})$ contains both the exact-score and approximation fluctuations, and, under the stated local smoothness and derivative controls, $\|R_B\| = O_p(\left\|\widehat{\beta} - \beta^\star\right\|^2) + O_p(\delta_G \left\|\widehat{\beta} - \beta^\star\right\|)$. Local invertibility and a standard contraction argument first imply $\left\|\widehat{\beta} - \beta^\star\right\| = O_p(B^{-1/2} + \varepsilon_G)$. Substituting this preliminary rate into the remainder and solving the linearized equation gives equation (66); products involving $\delta_G B^{-1/2}$ are absorbed into the $O_p(B^{-1/2})$ term. The KL claim follows from the quadratic upper bound in Lemma 12. □

More target observations reduce sampling noise but do not remove a persistent mean score error b_p. Consequently, accurate information matrices can be sufficient for good allocation even when the corresponding biased scores are unsuitable for final estimation. This differs from changing the reference law or misspecifying the target family, considered in Appendix I.3.

H Properties of the Optimal Allocation

GAMD minimizes a convex function of the panel cost shares. This section derives the objective derivatives and a Frank–Wolfe stopping bound for its numerical solver. It also proves that an optimal design can use a small number of panels and that the population objective is invariant to invertible feature transformations. For fixed β, abbreviate $F = F(\beta)$, $I_S = I_S(\beta)$, $M(p) = M_\beta(p)$, and $\Phi(p) = \Phi_\beta(p)$, and write $A_S = I_S/c(S)$. Throughout this section $F \succ 0$.

Proposition 13 (Convexity and directional derivatives)**.** *The extended-valued function* $\Phi(p) = \mathrm{tr}(FM(p)^{-1})$*, with value* $+\infty$ *for singular* $M(p)$*, is convex on* $\Delta_{\mathcal{S}}$*. At every* p *with* $M(p) \succ 0$*,*

$$\frac{\partial \Phi(p)}{\partial p_S} = -\gamma_S(p), \qquad \gamma_S(p) = \frac{1}{c(S)} \mathrm{tr}\left[M(p)^{-1} F M(p)^{-1} I_S\right].$$

Moreover, $\sum_S p_S \gamma_S(p) = \Phi(p)$*.*

Proof. For a positive-definite matrix M and symmetric direction H, the derivative of the inverse is $D(M^{-1})[H] = -M^{-1}HM^{-1}$. Therefore

$$D\{\mathrm{tr}(FM^{-1})\}[H] = -\mathrm{tr}(M^{-1}FM^{-1}H).$$

A second differentiation in the same direction gives

$$D^2\{\mathrm{tr}(FM^{-1})\}[H,H] = 2\,\mathrm{tr}(FM^{-1}HM^{-1}HM^{-1}). \tag{67}$$

Let $C = M^{-1/2}FM^{-1/2} \succeq 0$ and $D = M^{-1/2}HM^{-1/2}$. The right-hand side equals $2\,\mathrm{tr}(CD^2) \ge 0$. Thus $M \mapsto \mathrm{tr}(FM^{-1})$ is convex on the positive-definite cone. Composition with the affine map

$p \mapsto M(p)$ preserves convexity, and the lower-semicontinuous extension by $+\infty$ is convex on the whole simplex.

Taking $H = A_S$ gives the coordinate derivative and hence the formula for γ_S. Finally,

$$\begin{aligned}\sum_S p_S \gamma_S(p) &= \operatorname{tr}\left[M(p)^{-1} F M(p)^{-1} \sum_S p_S A_S\right] \\ &= \operatorname{tr}[M(p)^{-1} F] = \Phi(p).\end{aligned}$$

□

The sensitivity $\gamma_S(p)$ is defined in Proposition 13. The following result turns it into a computable upper bound on the remaining design suboptimality.

Corollary 4 (Equivalence theorem and Frank–Wolfe gap)**.** *A feasible $p^\star$ with $M(p^\star) \succ 0$ is optimal if and only if $\gamma_S(p^\star) \le \Phi(p^\star)$ for every S, with equality for every S such that $p^\star_S > 0$. For any p with $M(p) \succ 0$, the Frank–Wolfe gap*

$$g_{\mathrm{FW}}(p) = \max_{S \in \mathcal{S}} \gamma_S(p) - \Phi(p)$$

satisfies $0 \le \Phi(p) - \Phi(p^\star) \le g_{\mathrm{FW}}(p)$.

Proof. The directional derivative from $p^\star$ toward the simplex vertex e_S that assigns all mass to panel S is

$$\nabla\Phi(p^\star)^\top (e_S - p^\star) = -\gamma_S(p^\star) + \sum_T p^\star_T \gamma_T(p^\star) = -\gamma_S(p^\star) + \Phi(p^\star).$$

A convex differentiable function is minimized over the simplex exactly when all feasible directional derivatives are nonnegative, giving the inequality. Because the weighted average of the γ_S over active panels equals Φ, every active panel must attain equality.

For the gap bound, convexity gives

$$\Phi(p^\star) \ge \Phi(p) + \nabla\Phi(p)^\top (p^\star - p).$$

Rearranging and maximizing the right-hand side over all simplex points yields

$$\begin{aligned}\Phi(p) - \Phi(p^\star) &\le \max_{q \in \Delta_{\mathcal{S}}} \nabla\Phi(p)^\top (p - q) \\ &= -\Phi(p) + \max_S \gamma_S(p) = g_{\mathrm{FW}}(p).\end{aligned}$$

□

The Frank–Wolfe gap gives the optimization tolerance used in the performance bounds. Two additional properties describe the optimum itself: some optimal allocation uses only finitely many informative panels, and an invertible reparameterization of the same shift family does not change the population objective.

Proposition 14 (Sparse support and basis invariance)**.** *Let $d_r = r(r+1)/2$. There exists an optimal allocation supported on at most $d_r + 1$ panels. If $\widetilde{\phi} = C\phi + b$ for an invertible matrix C, the corresponding design objective is unchanged after the natural reparameterization.*

Proof. The matrices $M(p)$ form the convex hull of the finite set $\{A_S : S \in \mathcal{S}\}$ in the vector space of symmetric $r \times r$ matrices, whose dimension is d_r. By Carathéodory's theorem, every attainable information matrix, in particular an optimal one, can be represented as a convex combination of at most $d_r + 1$ atoms.

For the basis transformation, the centered feature and score transform as $\widetilde{s} = Cs$. At corresponding parameters $\widetilde{\beta} = C^{-\top}\beta$, the additive feature constant is absorbed by the log-partition function. Hence $\widetilde{F} = CFC^\top$, $\widetilde{I}_S = CI_SC^\top$, and $\widetilde{M} = CMC^\top$. Therefore

$$\operatorname{tr}(\widetilde{F}\widetilde{M}^{-1}) = \operatorname{tr}(CFC^\top C^{-\top} M^{-1} C^{-1}) = \operatorname{tr}(CFM^{-1}C^{-1}) = \operatorname{tr}(FM^{-1}).$$

□

I Extensions Beyond the Main Setting

We conclude with three examples that clarify the scope of the model and loss. Gaussian reference laws give a closed-form panel information matrix. Replacing KL with another smooth f-divergence preserves the first-order optimal allocation. When the target is outside the assumed shift family, a KL projection separates approximation error from error within the family.

I.1 Gaussian Models

If the reference law is Gaussian and the target differs only in its mean, the conditional score is linear. We calculate its information directly to recover the Gaussian covariance-design special case. Here $\phi(X) = X$ is unbounded; the calculation does not invoke or extend the bounded-feature concentration results.

Proposition 15 (Classical Gaussian optimal design as a special case)**.** *Suppose $Q_0 = \mathcal{N}(0, \Sigma)$ with $\Sigma \succ 0$ and $\phi(X) = X$. Write $\Sigma_{S,S}$ for the covariance submatrix indexed by S and $\Sigma_{:,S}$ for its selected columns. Then $Q_\beta = \mathcal{N}(\Sigma\beta, \Sigma)$, $F(\beta) = \Sigma$, and*

$$I_S(\beta) = \Sigma_{:,S}\Sigma_{S,S}^{-1}\Sigma_{S,:}. \tag{68}$$

Thus GAMD reduces exactly to covariance-based Gaussian information design.

Proof. The Gaussian moment-generating function gives $A(\beta) = \frac{1}{2}\beta^\top\Sigma\beta$. Tilting the Gaussian density by $e^{\beta^\top x}$ completes the square and shifts the mean to $\Sigma\beta$ without changing the covariance. Hence $s_\beta(X) = X - \Sigma\beta$ and $F(\beta) = \Sigma$. For a jointly Gaussian vector,

$$\mathbb{E}[X - \Sigma\beta \mid X_S] = \Sigma_{:,S}\Sigma_{S,S}^{-1}(X_S - (\Sigma\beta)_S).$$

Taking the covariance and using $\mathrm{Var}(X_S) = \Sigma_{S,S}$ gives

$$I_S = \Sigma_{:,S}\Sigma_{S,S}^{-1}\Sigma_{S,S}\Sigma_{S,S}^{-1}\Sigma_{S,:},$$

which simplifies to equation (68). □

I.2 Other Distributional Losses

KL is one choice of full-distribution loss. We now show that other smooth f-divergences assign the same local quadratic weight to parameter error, up to a positive constant. They therefore give the same first-order optimal allocation within the specified shift family.

Proposition 16 (Common local geometry of smooth f-divergences)**.** *Let $D_f(P\|Q) = \int f(dP/dQ)dQ$, where $f : (0,\infty) \to \mathbb{R}$ is convex and three times differentiable near one, $f(1) = f'(1) = 0$, and $f''(1) > 0$. Then, uniformly for β in compact sets,*

$$D_f(Q_\beta\|Q_{\beta+u}) = \frac{f''(1)}{2}u^\top F(\beta)u + o(\|u\|^2). \tag{69}$$

Consequently, every such divergence induces the same first-order optimal allocation as KL.

Proof. Let q_β denote a density relative to Q_0 and set $R_u = q_\beta/q_{\beta+u}$. A Taylor expansion of the log density ratio gives

$$\log R_u = -u^\top s_\beta(X) + O(\|u\|^2)$$

in $L_3(Q_\beta)$, uniformly on compact parameter sets. Therefore $R_u = 1 - u^\top s_\beta(X) + O(\|u\|^2)$. Since $D_f(Q_\beta\|Q_{\beta+u}) = \mathbb{E}_{Q_{\beta+u}} f(R_u)$ and $f(1+t) = \frac{1}{2}f''(1)t^2 + o(t^2)$, changing the expectation from $Q_{\beta+u}$ to Q_β affects only higher-order terms, while

$$\mathbb{E}_{Q_\beta}(u^\top s_\beta(X))^2 = u^\top F(\beta)u.$$

This proves equation (69). Multiplication of the loss by the positive scalar $f''(1)$ does not change the minimizing allocation. □

More general losses remain compatible with the same panel information $M_\beta(p)$ but may replace $F(\beta)$ by a different local loss matrix. Thus the conditional-information component of GAMD is loss agnostic, while the outer matrix in the design criterion reflects the chosen full-law objective.

I.3 Misspecified Target Shifts

The main theorem assumes that the true target belongs to the exponential-tilt family. For a target $P^\star$ outside this family, define its best full-law approximation by minimizing $D_{\mathrm{KL}}(P^\star \| Q_\beta)$. We first decompose the resulting KL loss, then compare this full-law projection with the parameter targeted by partial-data likelihood.

Proposition 17 (KL projection decomposition)**.** *Let $P^\star \ll Q_0$ with $D_{\mathrm{KL}}(P^\star \| Q_0) < \infty$, and suppose*

$$\beta^\dagger \in \arg\min_{\beta \in \Theta} D_{\mathrm{KL}}(P^\star \| Q_\beta)$$

is an interior minimizer. Then $\mathbb{E}_{P^\star}\phi(X) = \mathbb{E}_{Q_{\beta^\dagger}}\phi(X)$, and for every $\beta \in \Theta$,

$$D_{\mathrm{KL}}(P^\star \| Q_\beta) = D_{\mathrm{KL}}(P^\star \| Q_{\beta^\dagger}) + D_{\mathrm{KL}}(Q_{\beta^\dagger} \| Q_\beta). \tag{70}$$

Proof. Using equation (1),

$$D_{\mathrm{KL}}(P^\star \| Q_\beta) = D_{\mathrm{KL}}(P^\star \| Q_0) - \beta^\top \mathbb{E}_{P^\star}\phi(X) + A(\beta).$$

The first-order condition at the interior minimizer is therefore $\nabla A(\beta^\dagger) = \mathbb{E}_{P^\star}\phi(X)$, proving moment matching. Subtracting the objective at $\beta^\dagger$ from that at β gives

$$A(\beta) - A(\beta^\dagger) - (\beta - \beta^\dagger)^\top \nabla A(\beta^\dagger).$$

On the other hand, direct calculation under the exponential family gives

$$D_{\mathrm{KL}}(Q_{\beta^\dagger} \| Q_\beta) = A(\beta) - A(\beta^\dagger) - (\beta - \beta^\dagger)^\top \nabla A(\beta^\dagger).$$

The two expressions coincide, establishing equation (70). □

The first term in equation (70) is the approximation error of the shift family, while the second measures estimation error relative to its best full-law approximation. This decomposition defines a useful evaluation loss under misspecification.

Partial-data likelihood can target a different parameter. For a fixed cost allocation p, its population objective per unit budget is

$$\mathcal{L}_p(\beta) = \sum_S \frac{p_S}{c(S)} \mathbb{E}_{P^\star_S}[\log r_{\beta,S}(X_S)],$$

where $P^\star_S$ is the true target marginal on panel S. Differentiating under the expectation, justified by bounded panel scores, gives

$$\nabla \mathcal{L}_p(\beta) = \sum_S \frac{p_S}{c(S)} \mathbb{E}_{P^\star_S}[m_{\beta,S}(X_S)].$$

Thus an interior partial-likelihood optimum solves a weighted sum of panel-score equations. The full-law projection instead solves $\mu_\beta = \mathbb{E}_{P^\star}\phi(X)$. These equations agree at $\beta^\star$ under the well-specified model, but need not agree under misspecification: the conditional expectation defining $m_{\beta,S}$ is taken under Q_β, not $P^\star$. Consequently the limiting partial-likelihood estimate can depend on the allocation even though the best full-law projection does not.

J Experimental Details and Additional Results

This appendix provides the experimental protocols, baseline implementations, numerical settings, and additional numerical validation for Sections 6 and 7.

J.1 Common Evaluation Protocol

Each target individual is measured through a single panel. All panels used in our experiments have unit cost, so the total budget B equals the number of acquired target individuals, including those used for initialization. Within each replication, all methods share the same underlying target draws and initialization observations. Only the variables selected by a method are revealed to its acquisition

procedure; complete target observations are retained solely for evaluation. Main-stage cost shares are converted to integer panel counts by largest-remainder rounding, so that the full remaining budget is allocated.

The initialization budget follows

$$b_B = \min\left\{ \left\lceil a\sqrt{B/b} \right\rceil, \lfloor 0.2B \rfloor \right\},$$

with $(a, b) = (400, 8000)$ for the synthetic experiments and $(160, 1600)$ for real-data experiments. This gives initialization sizes $200, 283, 400, 566, 800$ for the synthetic budget sweep at $B = 2000, 4000, 8000, 16000, 32000$, respectively, and 400 for the nonlinearity experiments. For real-data experiments, the corresponding sizes are $80, 114, 160, 227$ at $B = 400, 800, 1600, 3200$.

Acquisition baselines. All baselines allocate the main-stage budget over the same panel set $\mathcal{S}$. Since all experimental panels have unit cost, a main-stage allocation $p \in \Delta_{\mathcal{S}}$ assigns approximately $B_{\text{main}} p_S$ measurements to panel S, where $B_{\text{main}} = B - b_B$.

Uniform SQD. The uniform design does not use reference or target information and assigns $p_S^{\text{unif}} = \frac{1}{|\mathcal{S}|}, S \in \mathcal{S}$. It therefore provides a design-agnostic benchmark for the value of adaptive panel allocation.

A-OSQD. A-OSQD is a covariance-based A-optimal split-questionnaire design. Let $Z \in \mathbb{R}^{16}$ denote the reference representation used by this baseline and let $\Sigma_{\text{ref}} = \text{Cov}(Z)$. For a panel S, let P_S be the corresponding coordinate-selection matrix. Under the Gaussian mean-shift proxy, the panel information is

$$J_S^{\text{A}} = P_S^\top (\Sigma_{\text{ref},SS})^{-1} P_S.$$

The resulting information under allocation p is

$$M_{\text{A}}(p) = \sum_{S \in \mathcal{S}} \frac{p_S}{c(S)} J_S^{\text{A}},$$

and A-OSQD solves

$$p^{\text{A}} \in \arg\min_{p \in \Delta_{\mathcal{S}}} \text{tr}\left[M_{\text{A}}(p)^{-1} \right].$$

Thus, A-OSQD uses only second-order reference dependence and an identity loss weight; unlike GAMD, it does not use the estimated target shift or the nonlinear conditional score.

Discriminative Score OED. Disc. provides a stronger alternative that attempts to learn the target-relevant panel information directly. Recall that GAMD values panel S through the conditional score $m_{\beta,S}(x_S) = \mathbb{E}_{Q_\beta}[s_\beta(X) \mid X_S = x_S]$. Instead of computing this conditional expectation through the reference generator, Disc. trains a predictor $\widehat{m}_\theta(S, x_S)$ to approximate $m_{\widetilde{\beta}_B, S}(x_S)$ from masked reference observations. Its estimated panel information is $\widehat{I}_S^{\text{disc}} = \widehat{\text{Cov}}_{\widetilde{\beta}_B}[\widehat{m}_\theta(S, X_S)]$, where the covariance is computed on an independent reference-validation set using tilt weights at $\widetilde{\beta}_B$. Disc. then solves the same full-distribution design problem as GAMD,

$$\widehat{p}^{\text{disc}} \in \arg\min_{p \in \Delta_{\mathcal{S}}} \text{tr}\left[\widehat{F}_{\text{disc}} \left(\sum_{S \in \mathcal{S}} p_S \widehat{I}_S^{\text{disc}} \right)^{-1} \right],$$

with $\widehat{F}_{\text{disc}}$ estimated from the same validation pool. Thus, GAMD and Disc. differ primarily in how they estimate the conditional panel information: GAMD queries the frozen conditional generator, whereas Disc. learns a task-specific conditional-score predictor.

The Disc. predictor takes a masked observation and its binary mask as input, with two width-64 $\tanh$ hidden layers and an r-dimensional linear output. It is trained at $\widetilde{\beta}_B$ using tilt-weighted squared error, an ℓ_2 penalty of 10^{-4}, and 100 gradient steps with an initial learning rate of 10^{-2} and linear decay. The synthetic validation pool contains 4000 observations. Because the predictor is task-specific, a new network is fitted for each target campaign.

GAMD and Disc. optimize their respective full-law objectives using Frank–Wolfe initialized from the uniform allocation, with at most 300 iterations. We use a dual-gap tolerance of 10^{-6} for synthetic GAMD and 10^{-4} for Disc. and the Gas experiments. Within each comparison, all methods use the same final estimator and parameter constraints.

Uncertainty quantification. For experiments with n paired replications, we report the mean loss together with its standard error $s/\sqrt{n}$.

J.2 Synthetic Experiments

Synthetic construction. We construct a 16-dimensional reference distribution from an equally weighted four-component Gaussian mixture, followed by the coordinate-wise transformation $x_j = \sinh(\alpha z_j)/\alpha$, with $x_j = z_j$ at $\alpha = 0$. Letting $u_j = x_j/s_j$, where s_j is the reference standard deviation of coordinate j, we use the 12-dimensional feature map

$$\phi(x) = \big(\tanh(u_1), \ldots, \tanh(u_6), \tanh(u_1u_7), \ldots, \tanh(u_6u_{12})\big)^\top.$$

All 120 coordinate pairs are admissible measurement panels. The target distribution is obtained by exponentially tilting the reference distribution, with the tilt magnitude calibrated to give an importance-weight effective-sample-size fraction of approximately 0.5. The true target parameter is used only to generate target observations and to compute oracle evaluation quantities; it is never available to the acquisition methods.

Initialization balances the six panels $(a, a+6)$, $a = 1, \ldots, 6$, so that every shift feature is directly observed on a nonvanishing fraction of the initialization sample. Unless otherwise stated, GAMD performs $J = 4$ final scoring updates.

Numerical accuracy of conditional-information estimation. Finally, we verify that the QMC approximation used for conditional-score and panel-information calculations is sufficiently accurate at the numerical settings used in the experiments. We compare the default QMC estimator with a higher-accuracy adaptive reference on conditional means, aggregate panel information, and end-to-end normalized risk.

Table 2 summarizes the resulting discrepancies. Conditional-mean errors are below 10^{-3} at the median and approximately 2×10^{-3} at the 95th percentile. The aggregate-information operator error is below 2×10^{-3}, and replacing the default scoring calculation by the adaptive reference changes normalized risk by only 0.038 on average. These checks indicate that the reported recovery results are not materially driven by numerical integration error.

Table 2: Numerical validation of the conditional-information computation.

Statistic	Observed error
Median relative conditional-mean error	9.488×10^{-4}
95th-percentile relative conditional-mean error	2.162×10^{-3}
Maximum absolute coordinate error	2.236×10^{-3}
Mean aggregate-information operator error	1.421×10^{-3}
Maximum aggregate-information operator error	1.977×10^{-3}
Mean absolute normalized-risk difference	0.038105

Implementation details. Reference moments and oracle quantities are computed from fixed pools of $20{,}000$ observations. Conditional expectations are evaluated using scrambled Sobol QMC. The budget sweep uses order-14 conditional-score integration and order-12 full-law moment estimation; panel information is estimated from 128 outer observations with two independent conditional estimates. The nonlinearity experiment uses the same panel-information resolution with cached order-14 conditional scores. Further numerical settings, including optimization tolerances, follow the common protocol in Appendix J.1.

J.3 Gas Sensor Experiments

Dataset and preprocessing. We use the UCI Gas Sensor Array Drift dataset, which contains $13{,}910$ observations collected over ten temporal batches from 16 chemical sensors. Each sensor provides eight raw measurements, giving a 128-dimensional observation. Batches 1–6 define the fully observed reference population; we use 4746 observations for generator training and hold out 1187 for validation.

Preprocessing is fitted using only the reference-training split. For each sensor, its eight raw measurements are standardized coordinate-wise using the corresponding reference-training mean and standard deviation. We then fit a one-dimensional principal component separately to each standardized sensor block. Let z_a denote the resulting principal-component score for sensor a, $a = 1, \dots, 16$. The fitted standardization parameters and principal-component directions are kept fixed and applied unchanged to the validation and target populations.

The 16 sensor-level scores are used only to define the low-dimensional reference-to-target shift. Specifically, we take

$$\phi(x) = \big(\tanh(z_1), \dots, \tanh(z_8), \tanh(z_1 z_9), \dots, \tanh(z_8 z_{16})\big)^\top,$$

giving $r = 16$ bounded shift features. The first eight features capture marginal changes in sensors $1, \dots, 8$, while the remaining eight capture selected cross-sensor changes between sensor pairs $(a, a+8)$. These transformations are fixed for all target campaigns.

The measurement space remains the original sensor space. A panel consists of two complete sensor blocks, so each acquisition reveals all eight raw measurements from each of two sensors, for a total of 16 observed variables. All $\binom{16}{2} = 120$ sensor pairs are admissible and have equal cost. Thus, the PCA representation is used only to parameterize the population shift and does not reduce what is measured from an acquired sensor.

Frozen conditional model. We train a VAEAC conditional generator on the reference-training split and freeze it before target acquisition. The model has input dimension 128, hidden width 512, and latent dimension 64. All generator and calibration quantities are fixed before the target replications.

Evaluation metric. Because the real target density is not available in closed form, we evaluate recovery within the empirical exponential-tilt family. For an empirical log-partition function $\widehat{A}$, we report

$$D_{\widehat{A}}(\widehat{\beta}_B, \beta^\dagger) = \widehat{A}(\widehat{\beta}_B) - \widehat{A}(\beta^\dagger) - (\widehat{\beta}_B - \beta^\dagger)^\top \nabla \widehat{A}(\beta^\dagger). \tag{71}$$

For the controlled experiment, $\widehat{A}$ is the empirical log-partition function of the reference-validation population, and $\beta^\dagger = \beta^\star$. For natural drift, we first draw $8,000$ complete samples from the frozen reference generator G_0 and compute their shift features. We then apply an affine calibration so that the resulting feature vectors match the reference-training feature mean and covariance. The empirical log-partition function of these calibrated feature vectors is used as $\widehat{A}$, and $\beta^\dagger$ is the projection of the complete target population onto this empirical tilt family. Complete target observations are used only to compute $\beta^\dagger$ for evaluation and are never available to the acquisition methods.

Controlled target shift. We construct controlled targets by exponentially reweighting held-out reference observations with probabilities proportional to $\exp\{\beta^{\star\top}\phi(x)\}$, with the tilt magnitude calibrated to an effective-sample-size fraction of approximately 0.5. This produces an exact exponential tilt of the empirical reference-validation distribution. We evaluate $B \in \{400, 800, 1600, 3200\}$ using 30 paired replications per budget.

Natural temporal drift. We next use later temporal batches directly as target populations. Batch 7 defines the first target campaign, while batches 8–9 are pooled for the second. No artificial reweighting is applied, so these targets reflect the temporal drift present in the original sensor measurements.

We evaluate $B \in \{800, 1600, 3200\}$ using 20 paired replications per campaign and budget.

Implementation details. Initialization balances the eight sensor pairs $(a, a+8)$, $a = 1, \dots, 8$. Within each experiment, all acquisition methods use the same target draws, initialization observations, final estimator, conditional-completion procedure, and parameter constraints; only the measurement allocation differs. The controlled and natural-drift experiments use at most 10 and 8 scoring updates, respectively, with 32 and 16 conditional completions per acquired observation.