Separate fits and joint compatibility in heterogeneous coverage: finite-label models with shared calibration

HUANG Man
Independent Researcher, Shanghai, China
huangman@technocentrics.com
ORCID: 0009-0003-3809-5105 Zenodo version DOI: 10.5281/zenodo.24412395

11 November 2026

hexagon:2611.00802v1

Abstract

Two preparation experiments can each fit a heterogeneous coverage model and still fail a shared-environment explanation. We characterize joint compatibility for finite-label probability representations with half-step source probes and pointwise calibration error. Absolute weaker-source void probabilities require the weighted calibration marginals to share mass at least $\kappa = (\beta_1+\beta_2-1)_+$. Combining this observed budget with exposure moments and a shared calibration band gives an exact criterion for the closure of the representation. A deterministic-anchor example requires error radius 0.10376 for separate fits and 0.33048 for joint compatibility. On an explicit three-moment family, matching bounds identify the first-order heterogeneous coefficient $\sqrt{\kappa}$; four environments preserve the anchors and give second-order stability of the minimum radius over unknown centers. A calibration-preserving moment repair and response-opening geometry yield sharp local inverse bounds with unknown, unequal group masses and first anchors. For prescribed mass and standard-deviation envelopes p_n, h_n with $h_n^2 \le bp_n$, the matching radius-estimation risk is $\min\{1, [n^{-1/2} + \min(h_n, n^{-1/4})]/p_n\}$. Uniform consistency has a corresponding two-scale necessary and sufficient criterion. Unknown envelopes introduce distinct adaptation costs. An interval honest over the full envelope union has matching expected diameter $\min\{1, n^{-1/4}/p_n\}$ even on the deterministic subclass. Under uniform protection of the deterministic regular risk, point estimators pay a sharp logarithmic cost elsewhere: one scale-free thresholded projection attains $\min\{1, [n^{-1/2} + \min(h_n, (\log(en)/n)^{1/4})]/p_n\}$, and a constrained-risk lower bound matches it. The larger mass-unrestricted moment class admits no uniformly consistent radius estimator. Simultaneous Bernoulli confidence events also give finite-sample compatibility certificates.

Keywords: truncated moment problem; heterogeneous mixtures; shared calibration; overlapping measures; partial identification; finite-sample testing.

1. Introduction

Can two preparation experiments that are separately consistent with heterogeneous coverage be explained by the same environment and a shared calibration? Separate fits constrain two response curves. Joint compatibility also constrains how their latent exposure and calibration laws can occupy one probability space. This paper identifies the information that makes those requirements differ.

Each experiment records a terminal binary void response under three source settings and two affine source probes. An environment is drawn before controlled Poisson births and held fixed during a complete repeat. The calibration can vary with that environment, subject to a pointwise error band shared by both preparations. Conditional thinning gives a mixture-of-exponentials probability representation. The weaker probe has an absolute void probability β_j in preparation j; dividing the

other probabilities by it gives the familiar normalized moment problem.

The absolute probabilities supply an additional joint restriction. Each weighted calibration marginal is a submeasure of the same environmental probability law. Their common mass must therefore be at least $\kappa = (\beta_1 + \beta_2 - 1)_+$. Theorem 1 combines this observed mass budget with the exposure-response moments and the calibration band, and proves an exact compatibility boundary for the closure of the finite-label representation. If $\beta_1 + \beta_2 \leq 1$, separate expanded profiles can be assembled on disjoint active environment blocks. Above that threshold, their required common mass can make the separate fits incompatible.

The identification difference has an explicit magnitude. The separate-fit benchmark places both preparations inside one calibration band while allowing their weighted environmental laws to be allocated independently. Take equal anchor exposure $r = 1/2$, weaker-source void probability $\beta = 3/4$, and normalized intermediate responses $(t_1, t_2) = (3/5, 4/5)$. Corollary 2.1 gives:

Requirement	Exact minimum calibration-error radius	Approximate value
Separate preparation fits within one error band	$\log(4/3)/(4\log 2)$	0.10376
Shared environment and calibration	$\log(5/2)/(4\log 2)$	0.33048

At error radius 1/8, each experiment fits its own heterogeneous coverage explanation inside the common band, while the pair has no joint representation. Theorem 2 derives the full fixed-band boundary for unequal deterministic anchors. Theorem 3 extends the exclusion to heterogeneous anchors through observable quadratic deviations. Its example passes both separate expanded profiles and the earlier determinant allowance; a fixed rational linear certificate still rejects it. These results quantify the joint identification gained by retaining absolute source-void probabilities.

Theorems 4–5 resolve the local effect of heterogeneous anchors. Centering at the observed first moment improves the response remainder to variance plus calibration width times standard deviation. For equal weak-source masses above one half and anchor mean $r \leq e^{-2}$, a common-block variance calculation identifies the exact first-order joint-boundary coefficient. Four environments attain it on the explicit three-moment family (22e); the upper bound permits every admissible third anchor. Along this family the minimum radius over unknown centers stays within $O(\sigma^2)$ of the realizing radius, whereas replacing the heterogeneous anchors by a point mass raises the required radius by an explicit positive multiple of σ. Here σ is the normalized exposure standard deviation. Corollary 5.1 converts the radius comparison into a square-root lower modulus in observed-probability distance and a local $n^{-1/4}$ estimation lower bound. Its two comparison arrays each admit a shared-environment representation.

Theorems 6–7 close the inverse problem on a specified local class. A moment-repair kernel changes the exposure moments at square-root transportation cost while preserving calibration marginals. Expanding the calibration band then opens a two-dimensional neighborhood of intermediate responses. These two constructions prove a pairwise square-root upper modulus on an observable regular-moment cone with fixed weak-source mass and first anchor. The lower-bound arrays belong to that same class, giving matching $n^{-1/4}$ estimation risk and confidence sets with a matching diameter guarantee. A mass-one counterexample shows that uniform continuity fails on the full model.

Theorems 8–9 allow both preparations' masses and first anchors to be unknown and unequal. A joint mass adjustment preserves one ambient calibration law and therefore the updated overlap budget. Affine mean adjustment and a uniform higher-moment repair then establish the same pairwise exponent over a prespecified local nuisance class. Estimation and confidence inversion vary these nuisance parameters jointly with all other raw probabilities, preserving both $n^{-1/4}$ matching conclusions.

Theorems 10–11 quantify the degeneration as masses approach one. Let p be the scale of the private masses $1-\beta_j$. In a specified triangular frontier tube, the private/common response triangle opens at order p rather than p^2. The raw-array modulus consequently has factor $1/p$, and matching pairs show minimax risk and fixed-confidence expected diameter of order $\min\{1, n^{-1/4}/p\}$. The critical condition is $np^4 \to \infty$. Corollary 11.1 uses a union of frontier tubes to attain these rates without knowing the true mass scale.

For prescribed envelope scales p, h, Theorem 12 enlarges the standard-deviation envelope from order p to h with $h^2 \le bp$, for a fixed small variance-to-mass budget b. Two-sided band expansion handles optimal bands near either endpoint or in the interior. The resulting modulus combines direct moment replacement at cost h with square-root repair, giving matching risk $\min\{1, [n^{-1/2} + \min(h, n^{-1/4})]/p\}$. The deterministic and heterogeneous lower pairs belong to this same larger class. This identifies the transition between the deterministic-anchor edge and the fourth-root regime.

Theorems 13–14 resolve unknown-envelope inference on the enlarged class. Joint projection adapts to the unknown mass scale. Honesty over the full variance-envelope union forces fourth-root expected diameter even on the deterministic subclass. Point estimation has a different optimal tradeoff: uniformly preserving the deterministic regular rate requires replacing the fourth-root term by $(\log(en)/n)^{1/4}$. A variance-thresholded pair of projections attains this cost, and a same-model constrained-risk bound proves it necessary.

The statistical analysis turns the compatibility restrictions into finite decisions. Eighteen Bernoulli frequencies from independently repeated complete experiments define one simultaneous confidence event. Theorem 15 uses it for a finite cover of possible central ratios, with all selected separating directions covered together. Its distance-based replication order calibrates inference over the full testing class. The overlap boundary supplies the model-specific identification result, and the statistical theory supplies its sampling guarantee.

Moment restrictions for exponential mixtures provide the marginal benchmark. Heckman, Robb and Walker [10, abstract] describe nonparametric Bayesian and classical tests, consistent estimation of the number of discrete support points, and method-of-moments estimation of the mixing distribution. Robb, Frydman and Robertson [11, Section 2, Proposition 1] relate finite-period survival probabilities to positive semidefinite Hankel matrices and extension to a completely monotone sequence. Their competing-risk results preserve such restrictions under stated hazard assumptions [11, Sections 3–4]. Gaussian and Lobatto extrema [13] give the sharp zero-error marginal profiles used here.

Our joint criterion asks whether two marginal fits can occupy one environmental probability space with a shared calibration band. In each preparation the normalized exposure law must satisfy its own moment restrictions. Joint feasibility also requires their weighted calibration marginals to share at least κ mass, with κ determined by the absolute void probabilities. Theorem 1 imposes both requirements simultaneously. Corollary 2.1 and Theorem 3 give explicit arrays satisfying the separate expanded profiles while violating the joint condition. These constructions establish strict additional exclusion power relative to the conjunction of separate marginal feasibility conditions in the stated finite-label model.

The testing contracts also differ. The three-period test in [11, Section 5, equations (5)–(6)] uses estimated multinomial probabilities and a statistic with an asymptotically standard normal boundary law. Section 6 here constructs simultaneous finite-sample Bernoulli confidence events and inverts joint feasibility over the common calibration band. The radius-estimation results concern the geometry of that joint identified set. Their square-root modulus and mass-margin degeneration determine the sampling rates under the observation contract in Section 6.1.

Common-submeasure geometry also appears in robust statistics. For equal masses β, the overlap condition is equivalent to total variation at most $(1-\beta)/\beta$ between normalized calibration marginals, the contamination threshold underlying Chen, Gao and Ren's modulus [15, Theorem 3.2]. Palomba [16, Sections 3–5] combines moment restrictions, population-share neighborhoods and convex duality around a specified or identified baseline law. Partial transport [18, Section 1] likewise uses couplings

with prescribed mass and dominated marginals. In the present experiment both weighted calibration marginals are unknown, each carries its own exposure-response moments, and both must lie in a shared calibration band. Their compulsory common mass is determined by observed void probabilities. Resolving these requirements together yields the explicit source-probability boundary, the exact calibration radius and the heterogeneous exclusion.

The contribution and its supporting results have distinct roles:

Role	Object resolved	Evidence
Joint identification	Observable weak-source masses, exposure moments and calibration bands constrained on one environment	Theorems 1–5; exact boundaries and matching finite-environment constructions
Local inverse geometry	Calibration-preserving repair, two-response opening and its private-mass degeneration	Lemmas 6.1–6.2, 8.1–8.2, 10.1–10.3, 12.1 and 13.1; Theorems 6, 8, 10 and 12
Marginal benchmarks	Separate moment feasibility and simpler profile or determinant restrictions	Theorems 16–18; explicit arrays comparing their exclusion power
Statistical consequences	Matching risk and honest expected diameter, the triangular critical boundary and compatibility certificates	Theorems 7, 9 and 11–15; Corollaries 5.1 and 11.1

The underlying tools include Poisson marking and thinning [1], unknown-link identification [2], exponential tilting [3], variance-sensitive concentration [4], and variance-controlled Jensen bounds [12]. Shared-frailty identifiability [7] provides a related setting in which heterogeneity and response variation can be confused. Our observation contract consists of binary terminal responses and certified source interventions.

The author's public preprints resolve different questions. The survival paper [17] establishes a moment criterion for a positive survival phase under renewed spatial catastrophes. The finite-preparation paper [9] identifies response amplitudes through recovery-front speeds and preparation histories. Its Section 8.3 supplies the nonlinear-calibration example, and Proposition 8.1 supplies the finite-cone framework evaluated in Appendix E. Theorems 1–15 here use the finite-label response representation and certified source relations; their proofs are self-contained and require no survival criterion, front-speed estimate or long-time response theorem. References [9] and [17] identify the public version-1 source records used for this comparison.

Section 2 defines the experiment. Section 3 develops the joint boundary, the calibration radius and its inference. Section 4 establishes marginal benchmarks; Section 5 gives response-data obstructions. Section 6 states the replication contract and confidence construction, Section 7 records reproducibility, and the appendices supply alternative budgets, count certificates and the cooperative-model illustration.

2. The preparation experiment and finite-label model

2.1. Probability representation

Fix an observation time $T \geq 1$ and preparation lengths $N_1 < N_2$. For proportional preparation one can take $N_j = \lfloor H_j T \rfloor$, with fixed $0 < H_1 < H_2$. Every arm begins at its preparation start with no grains. Three settings act during the old stage; the target stage is common.

An external environment E is drawn before controlled births and held fixed during a complete experiment. Conditional on E, births are Poisson and their grain marks are independent of the controlled birth realizations and unchanged across settings. The marks specify a grain's entire shape,

lifetime and motion. At the observation point, independent thinning gives

$$v_{i,j} = 1 - q_{i,j} = \mathbb{E}\Big[Z_j(E)e^{-\phi_i(E)N_j^*(E)}\Big] > 0, \qquad i = 1,2,3, \quad j = 1,2. \tag{1}$$

Here $N_j^* \geq 0$ is the finite old covering mean per unit rate, and $0 \leq Z_j \leq 1$ is the common conditional void factor from the target stage and any control-independent component. We use the star to distinguish this covering mean from the integer preparation length. The environment cannot include the future realizations of the controlled Poisson births. The calibration values ϕ_i and the environmental distribution are common across preparations. The factors Z_j and covering means may depend on preparation.

Equation (1) defines the finite-label probability model tested here. Across preparations it shares the environmental law and calibration, while Z_j and N_j^* remain preparation-dependent. Ordering the covering means, imposing a common lifetime kernel or prescribing a cross-preparation data law defines a smaller class. A specified space-time grain process may supply these additional relations.

Define the paired responses and determinants

$$\delta_{i,j} = q_{i,j} - q_{1,j} = v_{1,j} - v_{i,j}, \qquad \mathscr{D} = \delta_{3,2}\delta_{2,1} - \delta_{3,1}\delta_{2,2}, \tag{2}$$

$$\ell_{i,j} = \log(v_{i,j}/v_{1,j}), \qquad L = \ell_{3,2}\ell_{2,1} - \ell_{3,1}\ell_{2,2}. \tag{3}$$

These are products of marginal responses, rather than covariances of experimental observations.

2.2. Calibration bands and source probes

For the ordered-source experiment assume

$$0 \leq \phi_1(E) \leq \phi_2(E) \leq \phi_3(E) < \infty.$$

Set $\Delta_i = \phi_i - \phi_1$ and

$$Y_j = \Delta_3 N_j^*, \qquad dQ_j = \frac{Z_j e^{-\phi_1 N_j^*}}{v_{1,j}}\, dP, \qquad u_j = -\ell_{3,j} \geq 0. \tag{4}$$

The measure Q_j weights environments by their reference void probability and is used by the variance benchmark in Section 4.4.

Where $\Delta_3 > 0$, let $x_E = \Delta_2/\Delta_3$. Where $\Delta_3 = 0$, assign $x_E = x$. Assume that some common $x \in [0,1]$ satisfies

$$|x_E - x| \leq \epsilon \quad P\text{-almost surely}, \qquad 0 \leq \epsilon \leq 1. \tag{5}$$

The value of x can be unknown. The bound is pointwise in the environment and cannot be replaced by an unweighted average calibration error.

Two optional probes act only on the old birth source:

$$\phi_- = \phi_1 - s\Delta_3, \qquad \phi_+ = \phi_1 + s\Delta_3, \qquad s > 0, \quad \phi_- \geq 0 \quad P\text{-almost surely}. \tag{6}$$

They preserve the original Z_j, environmental distribution, preparation length and grain marks. Their void probabilities are denoted $v_{-,j}, v_{+,j}$; extend $\delta_{a,j} = v_{1,j} - v_{a,j}$ and $\ell_{a,j} = \log(v_{a,j}/v_{1,j})$ to $a = -, +$. Relation (6) is an externally certified relation between rates. Midpoints of an unknown laboratory control knob do not certify it. In particular, $\phi_1 = 0 < \phi_3$ rules out any positive weaker-probe step at that environment.

One way to enforce (6) without observing the absolute environmental rates is to build conditional Poisson source streams with nonnegative rates $c(E)$ and $d(E)$, common marks, and conditionally independent source noise. Define

$$\phi_1 = c + sd, \quad \phi_3 = c + (s+1)d, \quad \phi_- = c, \quad \phi_+ = c + 2sd.$$

Controlled superposition and thinning then certify $\Delta_3 = d$ and (6) for every environment. The construction also requires an intermediate stream at rate $c + (s + x_E)d$ with the pointwise ratio budget (5). These are source-design conditions; the terminal frequencies estimate their consequences. A device exposing only an unknown control knob supplies neither this decomposition nor the ratio budget automatically.

Sections 3 and 4.1–4.3 use half-step probes, $s = 1/2$. The variance certificate in Section 4.4 allows any feasible fixed $s > 0$; its numerical implementation accepts $10^{-6} \le s \le 1/2$.

Input or output	Experimental meaning
Observed inputs	Five-arm terminal binary frequencies in each preparation, recorded in independently repeated complete experiments
Certified design inputs	A common environmental sampling law, unchanged target and mark rules across source arms, the affine probe relation (6), and the pointwise ratio budget (5)
Population output	The joint compatibility region for the ten marginal void probabilities, and the minimum admissible calibration radius in the equal-anchor case
Finite-sample output	A rejection certificate on the simultaneous confidence event; failure to obtain a strict certificate leaves the data inconclusive

3. Joint compatibility under a shared environment

3.1. The exact common-environment boundary

Throughout Section 3 take $s = 1/2$ in (6). When the ratio varies within a band, each preparation's intermediate response depends on a joint law of exposure and calibration. A second restriction comes from the absolute weaker-probe probabilities, which normalization removes. Use $R_j = e^{-\Delta_3 N_j^*/2}$ and $m_{k,j} = v_{a_k,j}/v_{-,j}$, with $(a_1, a_2, a_3) = (1, +, 3)$. Write

$$\begin{aligned} &\beta_j = v_{-,j}, \quad U = x_E, \quad A_j = Z_j e^{-\phi_- N_j^*} \in [0,1], \\ &\mathbf{v}_j = (\beta_j, v_{1,j}, v_{+,j}, v_{3,j}, v_{2,j}), \qquad g(r,u) = (1, r, r^2, r^3, r^{1+2u}). \end{aligned} \tag{7}$$

The joint feasible set in this section is the finite-label interface (1). For a fixed calibration band $B = [l, h] \subset [0, 1]$, the half-step representation is $\mathbf{v}_j = E[A_j g(R_j, U)]$ on one probability space, with the same $U \in B$. In this section the compactified class permits $R_j = 0$ and describes the closure of the finite-covering-mean model. Positive R_j recover actual finite covering means by $N_j^* = -\log R_j$, $Z_j = A_j$, and source rates $(0, 1, 2, 3)$ with intermediate rate $1 + 2U$.

Define $\mathcal{K}_B$ as the convex hull of all ten-dimensional vectors $(a_1 g(r_1, u), a_2 g(r_2, u))$ with $a_j \in \{0, 1\}$, $r_j \in [0, 1]$ and $u \in B$. This compact convex set includes zero-mass boundaries; the overlap formulation below describes its slices with $\beta_j > 0$.

Let $\mathcal{M}_j(B)$ be the set of finite nonnegative measures ν_j on $[0, 1] \times B$ satisfying $\int g \, d\nu_j = \mathbf{v}_j$. Their masses are β_j. Denote their U marginals by λ_j. For two finite measures define their overlap as

$$\mathrm{ov}(\lambda_1, \lambda_2) = \int \min\left(\frac{d\lambda_1}{d\rho}, \frac{d\lambda_2}{d\rho}\right) d\rho, \qquad \rho = \lambda_1 + \lambda_2, \qquad \kappa = (\beta_1 + \beta_2 - 1)_+. \tag{8}$$

Theorem 1 (exact overlap boundary). Fix B and $0 < \beta_j \le 1$. The following statements are equivalent.

(a) The array belongs to the compactified common-environment model with $U \in B$.

(b) There are $\nu_j \in \mathcal{M}_j(B)$ whose calibration marginals have overlap at least κ.

(c) There are a common measure γ on $[0,1]^2 \times B$ of mass κ and residual measures σ_j on $[0,1] \times B$ of masses $\beta_j - \kappa$ such that

$$\mathbf{v}_j = \int g(r_j, u)\, d\gamma(r_1, r_2, u) + \int g(r, u)\, d\sigma_j(r, u). \tag{9}$$

Thus, if both single-preparation moment sets are nonempty, define

$$\mathcal{O}_B(\mathbf{v}) = \max_{\nu_j \in \mathcal{M}_j(B)} \mathrm{ov}(\lambda_1, \lambda_2). \tag{10}$$

The exact compatibility condition is $\mathcal{O}_B(\mathbf{v}) \geq \kappa$. The maximum is attained. A compatible compactified array has a representation with at most eleven environments and $A_j \in \{0, 1\}$. Compatibility with positive exposure coordinates gives actual finite covering means; the compactified criterion itself includes closure points.

Proof. A common environment induces $\nu_j(C) = E[A_j \mathbf{1}_{(R_j,U)\in C}]$. If p is its U marginal, then $\lambda_j \leq p$, since $A_j \leq 1$. The least measure dominating both marginals has mass $\beta_1 + \beta_2 - \mathrm{ov}(\lambda_1, \lambda_2) \leq 1$. This proves (b).

Conversely, choose a common submeasure $\eta \leq \lambda_j$ of mass exactly κ. Disintegrate ν_j conditional on U. Taking the proportion $d\eta/d\lambda_j$ of each conditional measure produces submeasures with the same U marginal η. Couple their exposure coordinates conditionally on U to form γ; the remainders are σ_j. This gives (c).

To reconstruct an environment from (c), use four disjoint blocks. Give the common block mass κ and weights $(A_1, A_2) = (1, 1)$, the two residual blocks masses $\beta_1 - \kappa, \beta_2 - \kappa$ and weights $(1, 0), (0, 1)$, and an inactive block mass $1 - \beta_1 - \beta_2 + \kappa$ with weights $(0, 0)$. Every mass is nonnegative. Unused exposure coordinates can equal one. This proves (a).

The moment constraints and domination by a common submeasure are closed under weak convergence on compact spaces. Maximizing the mass of that submeasure therefore attains (10). The same reconstruction identifies the feasible arrays with $\mathcal{K}_B$. A point in a convex hull in $\mathbb{R}^{10}$ needs at most eleven generators. Replacing zero exposure nodes by positive nodes shows that this compactified class is exactly the finite-mean model's closure. □

For equal masses $\beta_1 = \beta_2 = \beta$, write $\mu_j = \lambda_j/\beta$ and use $\mathrm{TV}(\mu_1, \mu_2) = \frac{1}{2}\int |d\mu_1 - d\mu_2|$. Then $\mathrm{ov}(\lambda_1, \lambda_2) = \beta\{1 - \mathrm{TV}(\mu_1, \mu_2)\}$. Condition (10) requires an admissible pair with $\mathrm{TV}(\mu_1, \mu_2) \leq (1 - \beta)/\beta$, automatic when $\beta \leq 1/2$. Setting the contamination share to $\eta = 1 - \beta$ gives exactly the classical threshold $\eta/(1 - \eta)$ in [15, Theorem 3.2]. The common-submeasure and coupling construction is established measure geometry, also used in partial transport [18].

The identification problem begins with the admissible pairs: each unknown marginal must lift to an exposure-calibration measure satisfying its observed anchor and fractional-response moments inside B. Theorem 1 resolves that simultaneous feasibility problem through the observed common-mass budget. Theorems 2–3 then evaluate and separate it in ways that independent moment fits miss. Retaining β_j alongside normalized moments preserves the budget; using only the normalized curves removes it.

The observed mass threshold also determines whether separate expanded profiles are sufficient; Corollary 16.1 in Section 4.1 establishes the distinction.

3.2. Closed-form boundaries and the minimum calibration radius

Suppose the anchor moments satisfy $m_{k,j} = r_j^k$ for $k = 1, 2, 3$, with $0 < r_j < 1$. The first two moments force $R_j = r_j$ under the active measure. Define

$$c_j = -2 \log r_j, \qquad t_j = \frac{v_{2,j}}{\beta_j r_j}, \qquad L_j = e^{-c_j h}, \quad H_j = e^{-c_j l}, \quad p = c_2/c_1. \tag{11}$$

These assumptions describe deterministic anchor exposure, while the calibration U and its selection weights can remain heterogeneous.

Theorem 2 (closed-form nonzero-error boundary). For a fixed nondegenerate band $B = [l, h]$, the following are necessary and sufficient for compatibility with finite N_j^* within (1).

If $\kappa = 0$, require $L_j \le t_j \le H_j$ for both groups. If $\kappa > 0$, define

$$C_j^- = \max\left(L_j, \frac{\beta_j t_j - (\beta_j - \kappa)H_j}{\kappa}\right), \qquad C_j^+ = \min\left(H_j, \frac{\beta_j t_j - (\beta_j - \kappa)L_j}{\kappa}\right). \tag{12}$$

Require $C_j^- \le C_j^+$ and, with

$$\Lambda(z) = \frac{(H_1 - z)L_1^p + (z - L_1)H_1^p}{H_1 - L_1}, \tag{13}$$

also require

$$\begin{array}{ll} p > 1: & (C_1^-)^p \le C_2^+, \quad \Lambda(C_1^+) \ge C_2^-; \\ 0 < p < 1: & \Lambda(C_1^-) \le C_2^+, \quad (C_1^+)^p \ge C_2^-; \\ p = 1: & \max(C_1^-, C_2^-) \le \min(C_1^+, C_2^+). \end{array} \tag{14}$$

For a singleton $B = \{l\}$ the condition is simply $t_j = e^{-c_j l}$.

Proof. Normalize the common measure of mass κ to a probability law Q. Its two contributions are $a_j = E_Q e^{-c_j U}$. The residual measure of mass $\beta_j - \kappa$ contributes any scalar between that mass times L_j and that mass times H_j. Solving these two inequalities for a_j gives exactly (12).

On the common block put $S = e^{-c_1 U} \in [L_1, H_1]$. Then $a_1 = ES$, $a_2 = ES^p$. For $p > 1$, the exact moment region is $a_1^p \le a_2 \le \Lambda(a_1)$: Jensen's inequality and the endpoint chord give the bounds, attained by a point mass and an endpoint mixture; their mixtures fill each vertical interval. For $p < 1$ the order reverses. Both boundaries increase with a_1, so projecting the part over $[C_1^-, C_1^+]$ gives the interval from the left lower boundary to the right upper boundary. Its intersection with $[C_2^-, C_2^+]$ is exactly (14). Theorem 1 reconstructs the environment from this common block and the two residual scalar contributions. All exposure nodes equal positive r_j, and finite mixtures attain the scalar bounds. □

For equal anchors $r_1 = r_2 = r$ and equal weaker voids $\beta_1 = \beta_2 = \beta$, with $1/2 < \beta \le 1$, Theorem 2 simplifies to

$$t_1, t_2 \in [L, H], \qquad \beta|t_1 - t_2| \le (1 - \beta)(H - L), \quad L = r^{2h},\ H = r^{2l}. \tag{15}$$

Each group's private mass $1 - \beta$ carries the difference between the two responses. This is a restriction on the joint calibration law even when every exposure moment is already known.

Corollary 2.1 (minimum pointwise-error radius). Let $1/2 < \beta < 1$, $r \in (0, 1)$, $q = r^2$, $t_- = \min(t_1, t_2)$, $t_+ = \max(t_1, t_2)$, and

$$D = \frac{\beta(t_+ - t_-)}{1 - \beta}. \tag{16}$$

Compatibility for some central ratio and some error radius exists if and only if $q \le t_- \le t_+ \le 1$ and $D \le 1 - q$. When these conditions hold, put $z_* = \min(t_-, 1 - D)$. The exact minimum radius is

$$\epsilon_* = \frac{\log[(z_* + D)/z_*]}{-4 \log r}. \tag{17}$$

It is attained with band endpoints $-\log(z_* + D)/c$ and $-\log z_*/c$, where $c = -2 \log r$, and center their midpoint. In particular $\epsilon_* \le 1/2$. If the conditions fail, no ratio-error budget restores compatibility. For $\beta = 1$ compatibility requires $t_1 = t_2 \in [q, 1]$ and then $\epsilon_* = 0$. For $0 < \beta \le 1/2$ the minimum is instead $\log(t_+/t_-)/(-4 \log r)$, provided both responses lie in $[q, 1]$.

Proof. Write $s = r^{2h}$ and $t = r^{2l}$. The equal-anchor condition (15) is $q \le s \le t_-$, $t_+ \le t \le 1$, and $t - s \ge D$. Minimizing the band width is equivalent to minimizing $\log(t/s)$. At fixed s, choose $t = \max(t_+, s + D)$. The resulting ratio decreases with s, so the largest feasible s is z_*. Since $D \ge t_+ - t_-$, the corresponding upper endpoint is $z_* + D$. This proves (17), including $D = 0$. Any band of width $h - l$ fits a centered radius $(h - l)/2$, and every original pointwise-error band has width at most 2ϵ. Thus the minimized half-width is exactly the original model's minimum radius. With no compulsory overlap, each preparation can use its own deterministic U, giving the last formula. □

For $r = 1/2$, $\beta = 3/4$, $(t_1, t_2) = (3/5, 4/5)$, the separate profiles require only $\epsilon \ge \log(4/3)/(4\log 2) \simeq 0.10376$. The joint model requires

$$\epsilon \ge \frac{\log(5/2)}{4\log 2} \simeq 0.3304820. \tag{18}$$

Thus $\epsilon = 1/8$ passes the expanded profiles and fails the exact common-environment boundary. The observed active-mass overlap raises the minimum error radius while the exposure moments stay fixed.

Holding the normalized curves fixed exposes a second projection of the boundary. For $q \le t_- \le t_+ \le 1$, compatibility at some radius is equivalent to $\beta \le \beta_* = (1 - q)/(1 - q + t_+ - t_-)$. For $\beta > 1/2$ this rearranges $D \le 1 - q$; for $\beta \le 1/2$ it holds automatically because $t_+ - t_- \le 1 - q$. In this example $\beta_* = 15/19$: the minimum radius is $1/2$ at that mass, and any larger β excludes joint compatibility even on the full band $[0, 1]$. Every separate normalized moment calculation is identical across these mass levels.

3.3. Observable exclusions with heterogeneous anchors

The preceding boundary is exact at deterministic anchors. A necessary inequality for the full heterogeneous model extends its exclusions to open sets of arrays. For any $r \in (0, 1)$ define

$$\begin{aligned} V_j(r) &= v_{+,j} - 2rv_{1,j} + r^2\beta_j = E[A_j(R_j - r)^2], \\ m &= \min(\beta_1, \beta_2, 1 - \beta_1, 1 - \beta_2), \quad W_\epsilon(r) = 1 - r^{4\min(\epsilon, 1/2)}. \end{aligned} \tag{19}$$

Theorem 3 (stable overlap inequality). Every array in the full half-step pointwise-error model satisfies

$$|v_{2,1} - v_{2,2}| \le r\{|\beta_1 - \beta_2| + mW_\epsilon(r)\} + 3\{\sqrt{\beta_1 V_1(r)} + \sqrt{\beta_2 V_2(r)}\}. \tag{20}$$

For equal weaker voids and $R_1 = R_2 = r$, its first term is attained for every $0 \le \epsilon \le 1/2$. Thus the mass factor and the error-band width in that term are sharp. A strict violation persists under sufficiently small changes of all marginal probabilities.

Proof. Put $f(U) = r^{2U}$. Theorem 1 decomposes each calibration marginal into the same common mass and private masses $\beta_j - \kappa$. Write $H = \max_B f$, $L = \min_B f$. If $\beta_1 \ge \beta_2$, the difference of the weighted f integrals is bounded above by $(\beta_1 - \beta_2)H + (\beta_2 - \kappa)(H - L)$, and its negative has a no larger bound. Here $\beta_2 - \kappa = m$, $H \le 1$, and $H - L \le 1 - r^{4\min(\epsilon, 1/2)}$. Interchanging labels covers the other ordering.

Uniformly in $u \in [0, 1]$, the derivative of t^{1+2u} in $t \in [0, 1]$ is at most three. Consequently $|v_{2,j} - rE[A_j f(U)]| \le 3E[A_j|R_j - r|] \le 3\sqrt{\beta_j V_j(r)}$. Combining these bounds gives (20). For sharpness take $R_j = r$, put the private masses at opposite endpoints of $B = [0, 2\epsilon]$, and put the shared mass anywhere in that band. All weights and rates give an actual finite-mean representation. Continuity gives the final assertion. □

For a nondegenerate example let both active exposure laws put equal mass at $0.49, 0.51$, and let $\beta_1 = \beta_2 = 0.9$. Evaluate the intermediate response at $U_1 = 3/8$ in preparation 1 and $U_2 = 1/8$ in preparation 2. Both separate curves are complete two-point mixtures. Their common anchor

moments are $(m_1, m_2, m_3) = (0.5, 0.2501, 0.12515)$, satisfying the strict moment conditions (29). Both pass the expanded profiles with center $1/4$, error $1/8$. With $r = 1/2$,

$$V_1 = V_2 = 0.00009, \qquad |v_{2,1} - v_{2,2}| > 0.11078519, \qquad \text{right side of (20)} < 0.06864467. \tag{21}$$

The new population margin exceeds 0.04214052. For the same data the old determinant is about 0.47976040, below its variance allowance 0.48183559 at error $1/8$. Even the completely uninformative ratio budget $\epsilon = 1/2$ leaves an overlap-exclusion margin above 0.01928519. Hence interval expansion and the determinant allowance miss a common-environment obstruction retained by absolute source-void probabilities. The positive V_j and strict margins give an open neighborhood of excluded heterogeneous arrays. Rational bounds certify the population exclusion.

Inequality (20) also gives a direct finite certificate on the eighteen-frequency event. Each $\beta_j = v_{1,j} - \delta_{-,j}$ has an interval $[\beta_j^-, \beta_j^+] \subset [0, 1]$. The expression $V_j = (1 - r)^2 v_{1,j} - \delta_{+,j} - r^2 \delta_{-,j}$ has an interval by absolute-coefficient propagation. If its upper endpoint is negative, the moment model is already excluded. Otherwise use $V_j^+ = \max(0, \overline{V}_j)$, the upper endpoints for β_j, the largest interval value of $|\beta_1 - \beta_2|$, and $m^+ = \min(\beta_1^+, \beta_2^+, 1 - \beta_1^-, 1 - \beta_2^-)$ on the right. A lower endpoint for the left exceeding that bound rejects. This rule needs neither safe logarithms nor empirical variance estimates under a tilted law. Its false rejections, combined with the previous rules, remain outside the same confidence event.

Corollary 3.1 (an explicit linear certificate for arbitrary calibration). For every $r \in (0, 1)$ and $\lambda > 0$, the full half-step model satisfies

$$r\beta_1 - r^3\beta_2 + v_{2,2} - v_{2,1} - \lambda\{V_1(r) + V_2(r)\} \le r - r^3 + \frac{9}{2\lambda}. \tag{22}$$

This holds for any environmental ratio $U \in [0, 1]$, without a specified error radius. Interchanging group labels gives the reversed certificate. Its left side is a linear functional of the observed anchor and intermediate probabilities, so the radius in (26) applies directly.

Proof. For $a = 1 + 2u$ and any $s \in [0, 1]$, $\pm s^a - \lambda(s - r)^2 \le \pm r^a + 9/(4\lambda)$ by the derivative bound and completing the square. Conditional on u, maximizing the two active contributions is therefore bounded by $[r - r^a + 9/(4\lambda)]_+ + [r^a - r^3 + 9/(4\lambda)]_+$. Both terms inside the positive parts are nonnegative because $r^3 \le r^a \le r$. Their sum is the right side of (22). Averaging proves the result. □

For the heterogeneous example, take $r = 1/2$, $\lambda = 150$. The certified rational support upper bound in (22) is 0.405, while the population left side exceeds 0.42128519. This is a fixed rational separating direction for the full band $[0, 1]$, so it also certifies every central ratio at error $1/8$. Appendix D gives a finite-count certificate in this direction. It is an explicit instance of the general support construction below.

3.4. Sharp local boundaries and calibration-radius stability

The first anchor supplies a centering identity that improves the heterogeneous remainder in (20). Suppose $v_{1,j} = \beta_j r$, with $0 < r < 1$. Then $E[A_j(R_j - r)] = 0$. Put $f_u(t) = t^{1+2u}$ and $D_r(u) = \partial_t f_u(r) = (1 + 2u)r^{2u}$. Define

$$C_r = \sup_{\substack{0 \le u,t \le 1 \\ t \ne r}} \frac{|f_u(t) - f_u(r) - D_r(u)(t - r)|}{(t - r)^2}, \qquad L_r = \max_{0 \le u \le 1} |D_r'(u)|. \tag{22a}$$

Both constants are finite. For example, $C_r \le 16/r^2$: within $|t - r| < r/2$, the second derivative is bounded by $12/r$, while outside that neighborhood the numerator is at most four. Also $L_r \le 2(1 + 3|\log r|)$. Subtracting $D_r(x)E[A_j(R_j - r)] = 0$ from the linear term and applying Cauchy–Schwarz gives

$$\left| v_{2,j} - rE[A_j r^{2U}] \right| \le C_r V_j(r) + L_r \epsilon \sqrt{\beta_j V_j(r)}. \tag{22b}$$

Consequently the last term in (20) can be replaced by the sum of these two remainders whenever the observed first anchors equal $\beta_j r$. The original bound remains available for arbitrary first anchors. With $V_j = \beta_j \sigma_j^2$, the refined remainder has order $\beta_j(\sigma_j^2 + \epsilon\sigma_j)$.

Centered Taylor bounds and variance terms are established sensitivity tools [12]. Lam [19, Theorem 3.1] derives local worst-case expansions under a shrinking KL neighborhood of a baseline distribution. Here the constraints fix observed exposure moments and compulsory shared mass; the local problem determines how that mass can be allocated across two preparations.

The next result resolves the first-order coefficient of the joint boundary, including the allocation of the compulsory common mass. It uses a range of anchor means on which the derivative geometry is monotone. This yields a uniform result for an unknown central ratio.

Theorem 4 (matched local joint boundary). Fix $1/2 < \beta < 1$ and $0 < r \le e^{-2}$. Write $p = 1 - \beta$, $\kappa = 2\beta - 1$, and $\ell_r = -\log r$. For $0 \le \epsilon \le 1/2$ define

$$F_\epsilon = r(1 - r^{4\epsilon}), \qquad G_\epsilon = 1 - (1 + 4\epsilon)r^{4\epsilon}, \qquad H_\sigma(\epsilon) = pF_\epsilon + \sigma\sqrt{\kappa}\, G_\epsilon. \tag{22c}$$

There are constants $\sigma_0 > 0$ and $K < \infty$, depending only on β, r, with the following properties.

(a) Every compatible array with equal weaker voids β, common normalized first anchor r, and common normalized second anchor $r^2 + \sigma^2$, satisfies

$$|v_{2,1} - v_{2,2}| \le H_\sigma(\epsilon) + 2C_r\beta\sigma^2, \qquad 0 \le \sigma \le \sigma_0. \tag{22d}$$

The central ratio is unrestricted. The bound holds for every admissible third anchor.

(b) Set $R_+ = r + \sigma\sqrt{\kappa}$ and $R_- = r - \sigma/\sqrt{\kappa}$, and take $\sigma > 0$ small enough that these nodes lie in $(0, 1)$. Let both preparations have the same normalized anchor law

$$\mu_\sigma = \frac{1}{2\beta}\delta_{R_+} + \frac{\kappa}{2\beta}\delta_{R_-}. \tag{22e}$$

Its first three moments are r, $r^2 + \sigma^2$, and $r^3 + 3r\sigma^2 + (\kappa - 1)\sigma^3/\sqrt{\kappa}$. Among arrays matching these three anchor moments and the weaker voids, let $S_\sigma(\epsilon)$ be the supremum of $v_{2,1} - v_{2,2}$ over the compactified shared model with error radius at most ϵ. Uniformly over $\epsilon \in [0, 1/2]$,

$$H_\sigma(\epsilon) - K\sigma^2 \le S_\sigma(\epsilon) \le H_\sigma(\epsilon) + K\sigma^2. \tag{22f}$$

The lower bound is realized by a four-environment model with positive exposure nodes. Thus $\sqrt{\kappa}\, G_\epsilon$ is the exact first-order heterogeneous correction to this joint support boundary.

Proof. First fix a band $B = [l, h]$. Here h is a band endpoint, not the dispersion parameter σ. By Theorem 1, represent the two active laws using a common block of mass κ and two private blocks of mass p. Put $F = f_l(r) - f_h(r)$ and $G = D_r(l) - D_r(h)$. For the assumed r, D_r is decreasing and convex: with $a = 1 + 2u$, $D_r'(u) = 2r^{2u}(1 - a\ell_r)$ and $D_r''(u) = 4\ell_r r^{2u}(a\ell_r - 2)$. If $l < h$, then $G > 0$. Set $t(u) = (D_r(u) - D_r(h))/G \in [0, 1]$.

The zeroth-order intermediate contrast is $pF - Q$, where $Q = \int_{\text{private }1}[f_l(r) - f_U(r)] + \int_{\text{private }2}[f_U(r) - f_h(r)] \ge 0$. Define $c = \int_{\text{common}} t \in [0, \kappa]$, $e_1 = \int_{\text{private }1}(1 - t)$, $e_2 = \int_{\text{private }2} t$, and $e = e_1 + e_2$. Since $|D_r'(u)|/|\partial_u f_u(r)| \le M_r = (3\ell_r - 1)/(\ell_r r)$, we have $Ge \le M_r Q$.

For each active law, let $T_j = \int A_j t^2 - (\int A_j t)^2/\beta$ denote its unnormalized centered second moment. Its means are $p + c - e_1$ and $c + e_2$. Using $t^2 \le t$ on the common and second private blocks, and $t^2 \le 1$ on the first private block, gives

$$\begin{aligned} \beta(T_1 + T_2) &\le \beta(p + 2c + e_2) - (p + c - e_1)^2 - (c + e_2)^2 \\ &= p\kappa + 2\kappa c - 2c^2 + 2(p + c)e_1 + (\beta - 2c)e_2 - e_1^2 - e_2^2 \\ &\le p\kappa + 2\kappa c - 2c^2 + 2\beta e \le \frac{\kappa}{2} + 2\beta e. \end{aligned} \tag{22g}$$

Here $T_j = \beta \operatorname{Var}_{\nu_j/\beta}(t)$ is unnormalized. With $\bar{t}_j = \int t\, d\nu_j/\beta$, Cauchy–Schwarz reads $|\int (R_j - r)(t - \bar{t}_j)\, d\nu_j| \le \sqrt{\beta\sigma^2 T_j}$. For the linear Taylor contribution this gives the upper bound $\sigma G(\sqrt{\beta T_1} + \sqrt{\beta T_2}) \le \sigma G\sqrt{\kappa + 4\beta e} \le \sigma G\sqrt{\kappa} + (2\beta\sigma M_r/\sqrt{\kappa})Q$. Adding the two Taylor remainders yields

$$v_{2,1} - v_{2,2} \le pF + \sigma\sqrt{\kappa}G - \left(1 - \frac{2\beta\sigma M_r}{\sqrt{\kappa}}\right) Q + 2C_r\beta\sigma^2. \tag{22h}$$

Take $\sigma_0 \le \sqrt{\kappa}/(2\beta M_r)$. A singleton band follows directly from centering. Interchanging preparations proves the absolute-value bound. For any band of width at most 2ϵ, the decreasing exponential and the decreasing magnitude of D_r' imply $F \le F_\epsilon$ and $G \le G_\epsilon$. This proves (a), uniformly in the center.

For (b), use the following four atoms. Unused exposure coordinates equal one.

Environmental mass	U	(A_1, A_2)	Active exposure coordinates
p	0	$(1, 0)$	$R_1 = R_+$
$\kappa/2$	0	$(1, 1)$	$(R_1, R_2) = (R_+, R_-)$
$\kappa/2$	2ϵ	$(1, 1)$	$(R_1, R_2) = (R_-, R_+)$
p	2ϵ	$(0, 1)$	$R_2 = R_+$

The masses sum to one. Each preparation has active mass β and normalized exposure law (22e). Take the central ratio to be ϵ and use the source reconstruction following (7); all covering means are finite. The two intermediate voids are

$$y_1(\sigma, \epsilon) = \tfrac{1}{2} f_0(R_+) + \tfrac{\kappa}{2} f_{2\epsilon}(R_-), \qquad y_2(\sigma, \epsilon) = \tfrac{\kappa}{2} f_0(R_-) + \tfrac{1}{2} f_{2\epsilon}(R_+). \tag{22i}$$

Their difference is $\frac{1}{2}F(R_+) - \frac{\kappa}{2}F(R_-)$, where now $F(t) = t - t^{1+4\epsilon}$. Its constant and linear terms are precisely pF_ϵ and $\sigma\sqrt{\kappa}G_\epsilon$. Since $f_0(t) = t$, (22a) bounds the remainder by $\beta C_r \sigma^2$, uniformly over ϵ. Equation (22d) supplies the matching upper bound over every law with the same three observed anchor moments; one may take $K = 2\beta C_r$ in (22f). □

The coefficient $\sqrt{\kappa}$ records the shared allocation restriction. Separate variance bounds would allow the larger coefficient β, since $\beta^2 - \kappa = (1 - \beta)^2 > 0$. The common-block calculation (22g) and the four-atom realization identify the smaller coefficient exactly. Both the upper bound and its realization keep all three measured anchor moments; the upper bound permits any alternative representation of those moments.

Theorem 5 (minimum-radius stability and a sharp heterogeneity correction). Fix the parameters of Theorem 4 and $0 < \epsilon_0 < 1/2$. Let $\mathbf{v}^\sigma$ have weaker voids β, anchors from (22e), and intermediate voids $y_j(\sigma, \epsilon_0)$ from (22i). Let $\epsilon_{\text{joint}}(\sigma)$ be its minimum error radius in the compactified model, optimizing over the central ratio. Then, for small σ,

$$\epsilon_0 - K_0\sigma^2 \le \epsilon_{\text{joint}}(\sigma) \le \epsilon_0. \tag{22j}$$

One may take $K_0 = 3\beta C_r/[4p\ell_r r^{1+4\epsilon_0}]$. Constants and the small-dispersion range are for fixed (β, r, ϵ_0). The same bounds hold for the infimum radius over positive-exposure representations.

For comparison, replace its second and third normalized anchors by r^2, r^3, preserving β, r and the two intermediate voids, and denote the minimum radius of this deterministic-anchor array by $\epsilon_{\text{det}}(\sigma)$. It is feasible for small σ and satisfies

$$\epsilon_{\text{det}}(\sigma) - \epsilon_{\text{joint}}(\sigma) = \sigma \frac{\sqrt{\kappa}\,[1 - (1 + 4\epsilon_0)r^{4\epsilon_0}]}{4(1 - \beta)\ell_r r^{1+4\epsilon_0}} + O(\sigma^2). \tag{22k}$$

Thus the observed heterogeneous anchors change the required radius at first order, with an explicit positive coefficient, while the realizing family's joint radius remains stable to second order around ϵ_0.

Proof. The construction gives the upper bound in (22j) using positive exposure nodes. Put $d_\sigma = y_1 - y_2$. Theorem 4 gives $d_\sigma = H_\sigma(\epsilon_0) + O(\sigma^2)$ and, at any compatible radius $\epsilon \le \epsilon_0$,

$d_\sigma \le H_\sigma(\epsilon) + 2C_r\beta\sigma^2$. The construction loses at most $\beta C_r\sigma^2$ and the outer bound spends at most $2\beta C_r\sigma^2$, so $H_\sigma(\epsilon_0) - H_\sigma(\epsilon) \le 3\beta C_r\sigma^2$. Since $H'_\sigma \ge pF'_\epsilon = 4p\ell_r r^{1+4\epsilon} \ge 4p\ell_r r^{1+4\epsilon_0} > 0$ on this interval, subtraction gives the lower bound. The compactified minimum exists by compactness of the joint parameter and probability set. Every positive-exposure representation lies in that set and the displayed construction belongs to the positive class, proving the infimum assertion.

For the deterministic comparison, at $\sigma = 0$ the normalized smaller intermediate response is $t_- = (\kappa + r^{4\epsilon_0})/(2\beta) > r^{4\epsilon_0}$, while both responses lie strictly between r^2 and one. Hence the branch $z_* = 1 - D$ of Corollary 2.1 persists for small σ, with $D = d_\sigma/(pr)$. On this branch $pF_{\epsilon_{\mathrm{det}}} = d_\sigma$. Expanding the inverse of this strictly increasing function gives $\epsilon_{\mathrm{det}} = \epsilon_0 + \sigma\sqrt{\kappa}G_{\epsilon_0}/(pF'_{\epsilon_0}) + O(\sigma^2)$. Combining this identity with (22j) proves (22k). □

Corollary 5.1 (pairwise radius sensitivity and an estimation lower bound). Write $\varepsilon_*(\mathbf{v})$ for the minimum compactified compatibility radius of an array in the full-band model $\mathcal{K}_{1/2} := \mathcal{K}_{[0,1]}$, and let $a_* > 0$ be the coefficient of σ in (22k). Denote the deterministic-anchor comparison array in Theorem 5 by $\widetilde{\mathbf{v}}^\sigma$ and their common limit by $\mathbf{v}^0$. For small σ,

$$d^*_\sigma = \|\widetilde{\mathbf{v}}^\sigma - \mathbf{v}^\sigma\|_\infty = \beta\sigma^2, \qquad \varepsilon_*(\widetilde{\mathbf{v}}^\sigma) - \varepsilon_*(\mathbf{v}^\sigma) = \frac{a_*}{\sqrt{\beta}}\sqrt{d^*_\sigma} + O(d^*_\sigma). \tag{22l}$$

Every neighborhood of $\mathbf{v}^0$ therefore fails to admit a uniform pairwise Hölder bound for ε_* of any exponent greater than $1/2$ on its compatible arrays.

Under the complete-repeat sampling contract of Section 6.1, fix any open neighborhood $\mathcal{N}$ of $\mathbf{v}^0$. Let $\mathfrak{P}(\mathcal{N})$ contain all allowed sampling laws whose marginal arrays belong to $\mathcal{N} \cap \mathcal{K}_{1/2}$. With n repeats per preparation, there are $c > 0$ and n_0 such that

$$\inf_{\widehat{\varepsilon}_n} \sup_{P \in \mathfrak{P}(\mathcal{N})} E_P\,|\widehat{\varepsilon}_n - \varepsilon_*(\mathbf{v}(P))| \ge cn^{-1/4}, \qquad n \ge n_0. \tag{22m}$$

This is a local lower bound for radius estimation under the allowed coupling class.

Proof. Only the second and third anchor coordinates change between the arrays. Their absolute differences in each group are $\beta\sigma^2$ and $\beta|3r\sigma^2 + (\kappa-1)\sigma^3/\sqrt{\kappa}|$. Since $3r \le 3e^{-2} < 1$, the first is the larger for sufficiently small σ. Equation (22k) proves (22l). Both arrays converge to $\mathbf{v}^0$, so their radius difference divided by $(d^*_\sigma)^\gamma$ diverges for every $\gamma > 1/2$. This is a pairwise neighborhood statement.

For (22m), select independent Bernoulli arms and independent preparation samples for both arrays; these are allowed by Section 6.1. Their limiting void probabilities lie strictly inside $(0,1)$. The two one-repeat laws differ in four anchor probabilities by $O(\sigma^2)$, so the product-law divergence for n repeats per preparation is at most $Cn\sigma^4$. Set $\sigma = \tau n^{-1/4}$ for a sufficiently small fixed $\tau > 0$. Both arrays eventually lie in $\mathcal{N}$, and their data laws then have total variation at most $1/2$ by Pinsker's inequality. Their radius gap is at least $a_*\sigma/2$. Testing by the nearer of the two radius values shows that the larger of the two expected absolute estimation errors is at least one quarter of their radius gap times one minus this total variation. This proves the asserted lower bound. □

The two alternatives in this estimation argument each admit a shared-environment explanation at their own minimum radius. The difficulty concerns the radius required by those explanations. Equation (22l) supplies a square-root lower modulus. Section 3.5 establishes a matching upper modulus and inference on an observable regular-moment class containing these arrays. The probability-distance testing rate in Theorem 15 concerns a different target.

For $\beta = 3/4$, $r = 1/10$ and $\epsilon_0 = 1/10$, the coefficient in (22k) is approximately 3.41452. The array is given exactly by (22e) and (22i): its heterogeneous joint radius is $1/10 + O(\sigma^2)$ with upper bound $1/10$, while its deterministic-anchor comparison requires $1/10 + 3.41452\ldots\,\sigma + O(\sigma^2)$.

The centered remainder (22b) also has two distinct sharp orders across the model. To see the quadratic order at zero calibration error, take $\beta = 1$, $U = 1/4$, fixed $0 < a < \min(r, 1-r)$, and compare the symmetric law on $r \pm \sigma$ with $\nu_\sigma = (1 - \sigma^2/a^2)\delta_r + \sigma^2(\delta_{r-a} + \delta_{r+a})/(2a^2)$. Both have

moments $r, r^2+\sigma^2, r^3+3r\sigma^2$. For $f(t)=t^{3/2}$ the difference between their fractional moments is $\sigma^2\{[f(r+a)+f(r-a)-2f(r)]/(2a^2)-f''(r)/2\}+O(\sigma^4)$. The coefficient is strictly positive because $f''''>0$ on this interval. The two laws can be coupled on a common environment with constant U, giving an actual compatible array. For the mixed order, the four-atom construction at $\beta=1$ has symmetric exposure nodes and contrast $\sigma G_\epsilon+O(\epsilon\sigma^3)$; as $\epsilon\downarrow 0$, $G_\epsilon=4(\ell_r-1)\epsilon+O(\epsilon^2)$ when $r<e^{-2}$. Hence neither the variance term nor the calibration-width-times-standard-deviation term can be uniformly removed. These order examples concern (22b); Theorems 4–5 give the matched joint boundary and full-array radius result on their stated anchor family.

For finite data, (22b) extends directly to a centering residual: its right side gains $3|v_{1,j}-\beta_j r|$. Thus exact equality of empirical anchor ratios is unnecessary for that certificate. All quantities are linear mean functions or square-root combinations of them and can be enclosed on the event of Section 6.1. The sharp coefficient in Theorem 4 is a population statement for its specified anchor slice; noisy data must retain uncertainty in that slice when using it.

3.5. A local upper modulus and matching radius inference

Fix β, r, ϵ_0 as in Theorem 5 and let $\mathbf{v}^0$ be the common limiting array in Corollary 5.1. The upper bound uses observable restrictions on the anchor moments. Put $d_r=\min(r,1-r)/4$ and $B_r=2d_r^2$. For group j define its normalized centered moments

$$s_j^2=\frac{v_{+,j}}{\beta}-r^2, \qquad q_j=\frac{v_{3,j}}{\beta}-3r\frac{v_{+,j}}{\beta}+2r^3. \tag{22n}$$

For sufficiently small $\rho>0$, consider the compact class

$$\mathcal{D}_\rho=\left\{\mathbf{v}\in\mathcal{K}_{1/2}: v_{-,j}=\beta,\ v_{1,j}=\beta r, \quad 0\le s_j^2\le B_r/4, \quad |q_j|\le d_r s_j^2/10, \quad \|\mathbf{v}-\mathbf{v}^0\|_\infty\le\rho\right\}. \tag{22o}$$

This is a known local model class with fixed weak-source mass and first anchor. Its centered-third-moment restriction is a sufficient regularity condition for the moment repair below. The class contains both arrays of Corollary 5.1 for all sufficiently small σ. All restrictions in (22o), apart from full-band compatibility itself, are linear equalities or inequalities in the raw probabilities.

Lemma 6.1 (moment repair preserving calibration marginals). Let μ and a target moment triple have the same mean r, variances $v, v'\le B_r/4$, and centered third moments q, q' satisfying $|q|\le d_r v/10$, $|q'|\le d_r v'/10$. Suppose the target triple has a representing probability law on $[0,1]$. If $\delta=\max(|v-v'|,|q-q'|)$ is small, there is a Markov kernel changing the exposure coordinate of μ to the target moments with expected absolute change at most $C\sqrt{\delta}$. The kernel can be applied conditional on any original calibration coordinate U, leaving its marginal unchanged. Constants depend on r.

Proof. Write $d=d_r$, $B=B_r$, and assume $\delta>0$. There are two fixed laws ζ_+,ζ_- of mean r, variance B, and centered third moments $\pm dB$: use the two-point offsets $(-d,2d)$ and $(-2d,d)$ with their mean-zero weights. Their convex mixtures provide every third moment in $[-dB,dB]$ at the same mean and variance.

Put $T=20/(dB)$ and $A=4(TB+1)$. If $v\le A\delta$, then $v'\le(A+1)\delta$. Replacing the exposure independently by any target representing law costs at most $\sqrt{v}+\sqrt{v'}\le C\sqrt{\delta}$; the original U is retained.

Otherwise take $\lambda=T\delta\le 1/4$ and

$$a=\sqrt{\frac{v'-\lambda B}{(1-\lambda)v}}. \tag{22p}$$

The bounds $v>A\delta$ and $v\le B/4$ imply $0<a\le 1$ and $1-a\le 2(\lambda B+\delta)/v$. Contract R to $r+a(R-r)$ with probability $1-\lambda$, and with probability λ replace it by a law ζ from the preceding mixture. This gives the target variance. The required third moment of ζ is $[q'-(1-\lambda)a^3q]/\lambda$. It is

admissible because

$$|q' - (1-\lambda)a^3 q| \le \delta + \frac{d\lambda v}{10} + \frac{3d}{5}(\lambda B + \delta) \le \frac{5}{8} d\lambda B + (1 + 3d/5)\delta < d\lambda B. \tag{22q}$$

Here $\lambda B = 20\delta/d$ and $d \le 1/8$. The coupling cost is at most $(1-a)\sqrt{v} + \lambda(\sqrt{v} + \sqrt{B}) \le C\delta/\sqrt{v} + C\delta \le C\sqrt{\delta}$. The contraction and replacement probabilities depend only on the two moment triples, and the replacement retains the original U. Thus applying this kernel to a joint exposure-calibration measure preserves its calibration marginal. For $\delta = 0$, use the identity kernel. □

Lemma 6.2 (opening the response region). There are neighborhoods of $\mathbf{v}^0$ and ϵ_0, and constants $c, t_0 > 0$, with the following property. Take any common/private representation (9) with the fixed masses and first anchors of (22o), of an array in that neighborhood, carried by a band $[l, h]$ of width at most 2η, where η is in the radius neighborhood. Hold its anchors fixed. For $0 < t \le t_0$, the model on $[l, h + t]$ contains every intermediate-response pair within ct in ℓ^∞ distance of the original pair. The radius required for this larger band is at most $\eta + t/2$.

Proof. We first establish uniform nondegeneracy of these representations. At $\mathbf{v}^0$, every active exposure equals r. Its intermediate contrast is pF_{ϵ_0}, where $p = 1 - \beta$. If it has a representation at radius at most ϵ_0, equality in the deterministic boundary forces $l = 0$, $h = 2\epsilon_0$, the first private block to have $U = 0$, and the second to have $U = 2\epsilon_0$. The common block satisfies $\kappa^{-1} \int f_U(r)\, d\gamma = [f_0(r) + f_{2\epsilon_0}(r)]/2$. Choose $b > 0$ so that $f_b(r) > [3f_0(r) + f_{2\epsilon_0}(r)]/4$. The common mass at $U \ge b$ is then at least $\kappa/4$, by bounding $f_U(r)$ below by $f_b(r)$ below b and by $f_{2\epsilon_0}(r)$ elsewhere. It therefore has a fixed positive amount of mass with U a fixed positive distance above l. Both exposure coordinates there equal r. Each private block has positive mass p with its active exposure equal to r.

These properties persist uniformly when the array and carrying radius are sufficiently close to $\mathbf{v}^0, \epsilon_0$. To verify this, otherwise take shrinking neighborhoods and a sequence of violating representations. Weak compactness of the common and private measures, together with convergence of band endpoints, gives a limiting representation of $\mathbf{v}^0$ at radius at most ϵ_0. The preceding characterization and the open-set part of weak convergence then contradict the violation. Consequently h is bounded away from one, each private block has uniformly positive mass with exposure in a compact subinterval of $(0, 1)$, and a uniformly positive common mass has both exposures in that subinterval and $U \ge l + b$ for some $b > 0$.

Now apply three separate perturbations, each keeping all exposure coordinates and active weights fixed. Increase U by t on the first private block; increase it by t on the second private block; or decrease U to $\max(l, U - t)$ on the common block. For $t \le t_0 \le b$, all three use the enlarged band $[l, h + t] \subset [0, 1]$. The resulting intermediate-response increments have the forms $(-a_1 t, 0)$, $(0, -a_2 t)$, and $(b_1 t, b_2 t)$, respectively. The derivative $\partial_u f_u(R) = 2\log(R) R^{1+2u}$ and the mass bounds above imply that a_1, a_2, b_1, b_2 lie in a fixed compact interval of strictly positive numbers. If all four coefficients lie in $[m, M] \subset (0, \infty)$, divide the vertices by t. The edge joining $(-a_1, 0)$ and (b_1, b_2) has ℓ^∞ distance $a_1 b_2/(a_1 + b_1 + b_2) \ge m^2/(3M)$ from zero. The other distances are $a_1 a_2/(a_1 + a_2)$ and $a_2 b_1/(a_2 + b_1 + b_2)$, with the same lower bound. Thus their convex hull contains the ball with $c = m^2/(3M)$, uniformly over those coefficients. Mixing the three representations preserves their anchors and enlarged band. Finally its midpoint is an admissible center and its half-width is at most $\eta + t/2$. □

Theorem 6 (local square-root upper modulus). For sufficiently small fixed ρ, there is $C < \infty$ such that

$$|\varepsilon_*(\mathbf{v}) - \varepsilon_*(\mathbf{w})| \le C\|\mathbf{v} - \mathbf{w}\|_\infty^{1/2}, \qquad \mathbf{v}, \mathbf{w} \in \mathcal{D}_\rho. \tag{22r}$$

Together with Corollary 5.1, the exponent $1/2$ is sharp on this same class.

Proof. First, ε_* is continuous at $\mathbf{v}^0$ within the fixed-mass, fixed-first-anchor slice of compatible arrays. For the lower limit, (20) with $V_j = \beta s_j^2 \to 0$ forces $\liminf \varepsilon_* \ge \epsilon_0$. For the upper limit, fix any small $u > 0$ and choose arbitrary representing exposure laws for the anchors of a convergent

sequence of arrays. Put $M_j = ER_j^{1+4(\epsilon_0+u)}$. Centering at r gives $M_j \to r^{1+4(\epsilon_0+u)}$. Allocate their exposures independently of an endpoint calibration law on $\{0, 2(\epsilon_0+u)\}$. A required intermediate response y_j then corresponds to the fraction $\theta_j = (y_j/\beta - M_j)/(r - M_j)$ at the lower endpoint. Such an allocation exists when $0 \le \theta_j \le 1$ and $\beta|\theta_1 - \theta_2| \le p$, by Theorem 1 applied to these two-point calibration marginals. At $\mathbf{v}^0$ all inequalities are strict when $u > 0$; they persist along the sequence. Hence $\limsup \varepsilon_* \le \epsilon_0 + u$, and letting $u \downarrow 0$ proves continuity there.

Let $\Delta = \|\mathbf{v} - \mathbf{w}\|_\infty > 0$, and start with a minimizing representation for $\mathbf{v}$. Apply Lemma 6.1 to each exposure coordinate to match the anchors of $\mathbf{w}$. Normalization by fixed β gives moment changes at most a constant times Δ. For each preparation, compute one kernel from its entire normalized active exposure law and the target moments, then use this same kernel on its common and private components. On the common block apply the two preparations' kernels independently conditional on the original coordinates. The kernels match the overall moments; no blockwise moment equality is required. The active calibration marginals and their overlap remain unchanged. The new array $\mathbf{z}$ has the target anchors and a representation at radius $\varepsilon_*(\mathbf{v})$, while $|z_{2,j} - w_{2,j}| \le C\sqrt{\Delta}$, by the uniform exposure derivative bound three.

Choose ρ small enough that these repaired arrays and carrying radii lie in the neighborhoods of Lemma 6.2; this is possible since $\Delta \le 2\rho$ and the continuity just proved is at $\mathbf{v}^0$. That lemma corrects both remaining response errors using $t \le C\sqrt{\Delta}$. Thus $\varepsilon_*(\mathbf{w}) \le \varepsilon_*(\mathbf{v}) + C\sqrt{\Delta}$. Interchanging the arrays proves (22r). The pair of Corollary 5.1 belongs to $\mathcal{D}_\rho$ for all small σ, because its centered-third-moment-to-variance ratio tends to zero; it rules out every larger exponent. □

The minimax statements use the sampling experiment of Section 6.1: the raw mean array has a shared-environment representation, while every within-repeat observation coupling with those means is allowed. Independent Bernoulli arms are therefore an admissible lower-bound subexperiment. A protocol that prescribes a narrower physical replay law defines a different statistical experiment.

Theorem 7 (matching radius estimation and confidence sets). Assume the true array belongs to a specified class $\mathcal{D}_\rho$ from Theorem 6. Observe n complete repeats per preparation under Section 6.1, with arbitrary allowed within-repeat coupling. Let $\widehat{\mathbf{v}}$ be the ten empirical void probabilities, let $\widehat{\mathbf{v}}^{\mathcal{D}}$ be a measurable nearest ℓ^∞ projection onto $\mathcal{D}_\rho$, and put $\widehat{\varepsilon} = \varepsilon_*(\widehat{\mathbf{v}}^{\mathcal{D}})$. Then

$$\sup_{P:\, \mathbf{v}(P) \in \mathcal{D}_\rho} E_P|\widehat{\varepsilon} - \varepsilon_*(\mathbf{v}(P))| \le Cn^{-1/4}. \tag{22s}$$

The reverse minimax inequality holds with another positive constant for all sufficiently large n. Thus $n^{-1/4}$ is the minimax expected absolute-error order on this specified local class.

For $0 < \alpha < 1$, set $b_n = \sqrt{\log(20/\alpha)/(2n)}$, and let $\mathcal{C}_n = \mathcal{D}_\rho \cap \{\mathbf{v} : \|\mathbf{v} - \widehat{\mathbf{v}}\|_\infty \le b_n\}$. If it is nonempty, report

$$I_n = \left[\min_{\mathbf{v} \in \mathcal{C}_n} \varepsilon_*(\mathbf{v}), \max_{\mathbf{v} \in \mathcal{C}_n} \varepsilon_*(\mathbf{v})\right]. \tag{22t}$$

Otherwise report the empty confidence set, with diameter defined as zero. Uniformly over this class the confidence set covers the true minimum radius with probability at least $1 - \alpha$. On every sample path its diameter is at most $C[\log(20/\alpha)/n]^{1/4}$. For every fixed $0 < \alpha < 1/4$, the minimax worst-case expected diameter among uniformly honest, possibly empty confidence intervals on $\mathcal{D}_\rho$ has order $n^{-1/4}$.

Proof. The class is nonempty and compact; (22r) makes the radius functional continuous on it. At fixed n the empirical count vector takes only finitely many values; fix one distance-minimizing point for each possible value. This gives a measurable nearest projection onto the fixed compact set. Its distance from the true array is at most $2\|\widehat{\mathbf{v}} - \mathbf{v}\|_\infty$. Ten marginal Hoeffding inequalities give $P\{\|\widehat{\mathbf{v}} - \mathbf{v}\|_\infty > t\} \le 20e^{-2nt^2}$, without assuming independence among the ten coordinates. Integrating this tail gives $E\|\widehat{\mathbf{v}} - \mathbf{v}\|_\infty^{1/2} \le Cn^{-1/4}$, proving (22s). The lower-bound pair and Bernoulli product laws used in Corollary 5.1 lie in this same class for large n, proving the matching minimax lower bound.

The Hoeffding event with threshold b_n contains the true array in $\mathcal{C}_n$, which is then nonempty. Whenever $\mathcal{C}_n$ is nonempty, compactness and continuity give the extrema in (22t), and (22r) bounds their difference using the set's diameter $2b_n$, regardless of whether the coverage event occurs. The empty output has diameter zero. This proves coverage and the all-sample-path diameter bound, hence the expected-diameter upper bound for fixed α.

For its lower bound use the two compatible arrays in Corollary 5.1 with $\sigma = \tau n^{-1/4}$. Choose τ small enough that their data-law total variation is at most $1/4$. For any confidence set honest at level $1-\alpha$ under both laws, under the first law the probability of containing both radius values is at least $1 - 2\alpha - 1/4 > 1/4$. Its expected diameter is therefore at least a positive constant times the radius gap, which is bounded below by a constant times $n^{-1/4}$. Both arrays belong to the same $\mathcal{D}_\rho$ for large n, proving the minimax expected-diameter lower bound. □

An empty confidence set records incompatibility of the empirical confidence box with the specified local class; under that class this happens with probability at most α. This procedure fixes the local class, including β, r and the working neighborhood, before sampling. The third-moment condition is imposed on the population probabilities, with its uncertainty retained by projection or confidence-set intersection. Section 3.6 releases the fixed mass and first-anchor values. These are variational constructions over the compact moment-feasibility class; the theorem establishes their statistical properties without a general-purpose numerical solver implementation. The connection between moduli of continuity and statistical difficulty is classical in robust decision theory [15] and optimal recovery [20]; the Bernoulli risk and coverage bounds here follow from the displayed concentration arguments. The model-specific work is the moment repair preserving calibration marginals and the verified response-opening geometry. Lemma 6.2 states explicitly the geometric margin needed by this inverse problem.

Why the mass margin matters. The restriction $\beta < 1$ has a concrete role. At $\beta = 1$, take $r < e^{-2}$, fixed $0 < e < 1/2$, equal masses on $U = 0, 2e$, and the antithetic exposures $(R_1, R_2) = (r + \sigma, r - \sigma)$ and $(r - \sigma, r + \sigma)$. Both active weights are one. The contrast is $\sigma G_e + O(\sigma^3)$ with $G_e > 0$. Equation (22b), which holds at $\beta = 1$, forces every representing radius to satisfy $\sigma G_e + O(\sigma^3) \le 2C_r\sigma^2 + 2L_r\epsilon\sigma$. Thus its minimum radius has positive lower limit as $\sigma \downarrow 0$. The limiting array has deterministic exposure r and identical intermediate responses $[r + r^{1+4e}]/2$; a constant calibration ratio realizes it at radius zero. The radius functional is therefore discontinuous at this boundary array. A full-model square-root upper modulus is false. The fixed mass margin and the observable moment cone in (22o) give a local class where matching estimation is established; the cone is a verified sufficient condition, rather than a characterization of all regular arrays.

3.6. Unknown masses and first anchors

The mass and first-anchor equalities can be released without changing the local statistical exponent. Fix reference values $\beta_0 \in (1/2, 1)$, $r_0 \in (0, e^{-2})$ and $\epsilon_0 \in (0, 1/2)$, and form the reference array $\mathbf{v}^0$ as in Theorem 5 at zero dispersion. Choose known compact ranges $1/2 < \underline{\beta} < \beta_0 < \overline{\beta} < 1$ and $0 < \underline{r} < r_0 < \overline{r} < e^{-2}$. Put $d = \min(\underline{r}, 1 - \overline{r})/4$ and $B = 2d^2$. For a candidate raw array, define its own group-specific parameters $\beta_j = v_{-,j}$, $r_j = v_{1,j}/\beta_j$, $s_j^2 = v_{+,j}/\beta_j - r_j^2$ and $q_j = v_{3,j}/\beta_j - 3r_j v_{+,j}/\beta_j + 2r_j^3$. Consider

$$\begin{aligned}\mathcal{E}_\rho = \{\mathbf{v} \in \mathcal{K}_{1/2} : \ & \underline{\beta} \le \beta_j \le \overline{\beta}, \quad \underline{r} \le r_j \le \overline{r}, \\ & 0 \le s_j^2 \le B/4, \quad |q_j| \le d s_j^2/10, \\ & \|\mathbf{v} - \mathbf{v}^0\|_\infty \le \rho\}.\end{aligned} \tag{22u}$$

The ranges, reference array and neighborhood are specified before sampling; the true β_j, r_j are unknown and can differ between preparations. They vary across the candidate arrays used by the estimator and confidence procedure. All denominators stay uniformly positive. This is a compact class, generally with nonlinear moment restrictions in the raw coordinates.

Lemma 8.1 (joint adjustment of mass, mean and higher moments). Take two arrays $\mathbf{v}, \mathbf{w}$ satisfying the restrictions in (22u), and suppose $\mathbf{v}$ has a representation on a band B_0. With $\Delta = \|\mathbf{v} - \mathbf{w}\|_\infty$ small, there is an array $\mathbf{z}$ represented on that same band, with all weak-source and anchor probabilities equal to those of $\mathbf{w}$, and $|z_{2,j} - v_{2,j}| \le C\sqrt{\Delta}$. Constants depend only on the fixed ranges. The construction keeps one ambient environment and its unweighted calibration law throughout.

Proof. The functions taking raw coordinates to β_j, r_j, s_j^2, q_j are uniformly Lipschitz on the prescribed ranges. Begin with a binary-weight representation for $\mathbf{v}$, which exists by Theorem 1. We adjust each preparation on the same environment, adding conditional randomization when needed.

First change its mass from b to b'. If $b' \le b$, independently retain an active weight with probability b'/b. If $b' > b$, independently activate an originally inactive weight with probability $(b' - b)/(1 - b)$, assigning newly active exposures the current mean m. The calibration coordinate is left unchanged. Each raw response changes by at most $|b' - b|$. Uniform thinning leaves the normalized active exposure law unchanged; activation replaces it by $(b/b')\mu + (1 - b/b')\delta_m$. In either case its mean stays m and its variance and centered third moment are scaled by the same factor at most one. Thus the moment cone is preserved.

The mass step changes the active calibration marginals while preserving their domination by the same ambient calibration probability law. Consequently the adjusted marginals satisfy $\mathrm{ov}(\lambda_1', \lambda_2') \ge (\beta_1' + \beta_2' - 1)_+$, the compulsory budget for the new masses.

Next change the mean from m to m' using

$$T_{m,m'}(R) = \begin{cases} (m'/m)R, & m' \le m, \\ aR + 1 - a, \quad a = (1 - m')/(1 - m), & m' > m. \end{cases} \tag{22v}$$

This maps $[0, 1]$ into itself, has mean m', and moves every point by at most $C|m' - m|$. Its centered variance and third moment are multiplied by a^2, a^3, with the analogous $a = m'/m$ in the first branch. Since $0 < a \le 1$, the common cone $|q| \le dv/10$ and variance bound remain valid. The response change is at most $3C|m' - m|$.

Finally apply Lemma 6.1 at the new mean to match the target variance and third moment. Its proof uses only that the fixed tail nodes at mean m' with offsets $(-d, 2d)$ and $(-2d, d)$ lie in $[0, 1]$. It therefore holds uniformly for $m' \in [\underline{r}, \overline{r}]$ with this same d, B. All these means satisfy $m' - 2d \ge \underline{r}/2 > 0$ and $m' + 2d \le (1 + \overline{r})/2 < 1$. Take Δ small enough that the uniform bound on the centered-moment discrepancy also gives $\lambda = T\delta \le 1/4$ in Lemma 6.1. The source and target centered moments now differ by $O(\Delta)$, so the exposure transport cost is $O(\sqrt{\Delta})$.

For each preparation use one kernel calculated from its entire adjusted active law on all active components. These exposure kernels preserve the adjusted active calibration marginals. They may be applied independently conditional on the enlarged common environment; the two preparations still have a shared calibration coordinate on B_0. Their first three raw anchor probabilities and masses are now exactly the targets. The exposure derivative bound three controls the remaining response change. □

Lemma 8.2 (opening with varying masses and means). The conclusion of Lemma 6.2 remains valid in neighborhoods of the present reference array and radius when the two masses and first anchors are allowed to vary as in (22u). The constants are uniform over a sufficiently small such neighborhood.

Proof. Use the common/private representation of Theorem 1 with common mass $\kappa(\mathbf{v}) = \beta_1 + \beta_2 - 1$ and private masses $1 - \beta_2, 1 - \beta_1$. The prescribed ranges give $\kappa(\mathbf{v}) \ge 2\underline{\beta} - 1 > 0$ and $1 - \beta_j \ge 1 - \overline{\beta} > 0$. In the compactness contradiction of Lemma 6.2, arrays approaching $\mathbf{v}^0$ have $\beta_j \to \beta_0$, $r_j \to r_0$ and active exposure variance tending to zero. The measures and band endpoints have convergent subsequences; their limiting block masses are $2\beta_0 - 1, 1 - \beta_0, 1 - \beta_0$. At the limit the same deterministic boundary equality forces the endpoint band and the positive common mass away from its lower

endpoint. Thus the remaining margins hold uniformly: positive active mass in a fixed exposure interval $J \subset (0,1)$, positive common mass with $U \geq l + b$ for some $b > 0$, and $h \leq 1 - \tau$ for some $\tau > 0$. Together with the common/private mass bounds, these are the uniformity conditions used in Lemma 6.2. Its three perturbations still preserve all raw anchors, and its triangle proof gives the same uniform two-response opening. □

Theorem 8 (square-root modulus with unknown group-specific nuisance parameters). For sufficiently small fixed ρ, there is a finite constant C such that

$$|\varepsilon_*(\mathbf{v}) - \varepsilon_*(\mathbf{w})| \leq C\|\mathbf{v} - \mathbf{w}\|_\infty^{1/2}, \qquad \mathbf{v}, \mathbf{w} \in \mathcal{E}_\rho. \tag{22w}$$

The exponent is sharp on this same class. The values of β_j, r_j need not be known or equal across preparations.

Proof. First verify continuity of the minimum radius at $\mathbf{v}^0$ while masses and means vary. For the lower limit, use (20) at r_0; its quadratic deviations satisfy $V_j(r_0) = \beta_j[s_j^2 + (r_j - r_0)^2] \to 0$. The remaining bound converges to the strictly increasing deterministic boundary at the reference mass and mean, forcing $\liminf \varepsilon_* \geq \epsilon_0$.

For the upper limit fix a small $u > 0$ and use exposure laws representing each array's anchors. Their means tend to r_0 and variances to zero, so $M_j = ER_j^{1+4(\epsilon_0+u)} \to r_0^{1+4(\epsilon_0+u)}$. Allocating exposures independently of an endpoint calibration on $\{0, 2(\epsilon_0 + u)\}$ requires the normalized lower-endpoint fractions $\theta_j = (y_j/\beta_j - M_j)/(r_j - M_j)$. For unequal masses the exact environmental allocation condition is $\max(\beta_1\theta_1, \beta_2\theta_2) + \max(\beta_1(1-\theta_1), \beta_2(1-\theta_2)) \leq 1$, together with $\theta_j \in [0,1]$. This is the mass of the least measure dominating both two-point calibration marginals. At $\mathbf{v}^0$ and $u > 0$, the inequalities are strict, as in Theorem 6. They persist when masses and means vary near the reference. Hence $\limsup \varepsilon_* \leq \epsilon_0 + u$, proving continuity at the reference after $u \downarrow 0$.

Start from an optimal representation of $\mathbf{v}$. Lemma 8.1 produces target anchors on the same band with intermediate error $O(\sqrt{\Delta})$ relative to $\mathbf{w}$. Use Theorem 1 to put the adjusted array in common/private form (9) at its new masses, on the same band. By the continuity just proved and a sufficiently small choice of ρ, Lemma 8.2 applies uniformly to this representation. Opening its band by $O(\sqrt{\Delta})$ matches both target intermediate responses exactly. Thus $\varepsilon_*(\mathbf{w}) \leq \varepsilon_*(\mathbf{v}) + C\sqrt{\Delta}$; reverse the arrays for (22w).

The two arrays of Corollary 5.1 at the fixed interior reference values belong to $\mathcal{E}_\rho$ for all sufficiently small dispersion: their masses and means lie in the ranges, their centered-third-moment ratios tend to zero, and their raw arrays approach $\mathbf{v}^0$. They exclude every larger pairwise Hölder exponent. □

Theorem 9 (matching inference with unknown masses and means). In Theorem 7 replace $\mathcal{D}_\rho$ everywhere by $\mathcal{E}_\rho$, including the nearest projection and confidence-box intersection. This gives a measurable estimator and an honest, possibly empty confidence interval for the minimum radius. Uniformly over all allowed sampling laws with arrays in $\mathcal{E}_\rho$,

$$\sup_P E_P|\widehat{\varepsilon} - \varepsilon_*(\mathbf{v}(P))| \leq Cn^{-1/4}, \qquad \inf_{\widetilde{\varepsilon}} \sup_P E_P|\widetilde{\varepsilon} - \varepsilon_*(\mathbf{v}(P))| \geq cn^{-1/4} \tag{22x}$$

for sufficiently large n. The confidence interval has coverage at least $1 - \alpha$ and an all-sample-path diameter at most $C[\log(20/\alpha)/n]^{1/4}$, with diameter zero for an empty output. For each fixed $0 < \alpha < 1/4$, its expected-diameter order $n^{-1/4}$ is minimax among uniformly honest, possibly empty intervals on this same class.

Proof. The class (22u) is nonempty and compact, and (22w) makes the radius map continuous on it. A measurable nearest projection therefore exists even though the class need not be convex. The ten-mean Hoeffding bound is unchanged: it covers the raw probabilities from which masses and normalized first anchors are inferred. The projection inequality and the upper modulus prove the risk upper bound; intersection with the same ten-mean box proves coverage and the diameter bound exactly as in Theorem 7. Its expectation obeys the same fixed-confidence bound.

For the two lower bounds, restrict the larger class to the compatible two-array sequence at (β_0, r_0) used in Corollary 5.1. That sequence lies in $\mathcal{E}_\rho$, so the same product-Bernoulli KL and simultaneous-coverage arguments give the displayed risk lower bound and the expected-diameter lower bound. Thus upper and lower results again use one parameter class. □

Masses and first anchors vary jointly with the radius across the raw confidence region. The projection remains defined even if an empirical weak-source frequency vanishes, since every candidate array has $\beta_j \geq \underline{\beta} > 0$. The one raw-mean event covers nuisance and radius uncertainty jointly. The ranges, moment cone and working neighborhood remain prespecified. Uniform margins keep the normalized moment map and activation probabilities controlled, while the positive common and private masses support response opening. The mass-one discontinuity example shows that the upper mass margin cannot be removed in a theorem covering the full model. The centered-moment cone and the other stated local restrictions are verified sufficient conditions; they do not characterize every possible regular nuisance configuration.

Strict interior masses without a uniform margin. Failure persists if mass one itself is excluded but masses may approach it. Fix $r \in (0, e^{-2})$ and $e \in (0, 1/2)$. In the four-environment construction take $\beta_\sigma = 1 - \sigma^2$, $\kappa_\sigma = 1 - 2\sigma^2$ and dispersion σ. Call its array $\mathbf{a}^\sigma$. Let $\mathbf{b}^\sigma$ have the same weaker voids and first anchor, deterministic exposures r, and both intermediate probabilities equal to the average of those of $\mathbf{a}^\sigma$. Both arrays have strictly positive exposure representations and masses strictly below one. The common intermediate probability of $\mathbf{b}^\sigma$, divided by β_σ, lies strictly between r^3 and r for small σ, so a constant calibration realizes it at radius zero. For $\mathbf{a}^\sigma$, (22i) gives contrast $\sigma G_e + O(\sigma^2)$. Combining the mass bound used in (20) with the centered remainder (22b) gives, at every representing radius, $\sigma G_e + O(\sigma^2) \leq r\sigma^2 + 2C_r\beta_\sigma\sigma^2 + 2L_r\epsilon\beta_\sigma\sigma$. Consequently $\liminf \varepsilon_*(\mathbf{a}^\sigma) \geq G_e/(2L_r) > 0$, whereas $\varepsilon_*(\mathbf{b}^\sigma) = 0$ and $\|\mathbf{a}^\sigma - \mathbf{b}^\sigma\|_\infty = O(\sigma)$.

Both arrays obey the same fixed centered-moment cone for small σ: the heterogeneous ratio is $|q|/s^2 = 2\sigma^3/\sqrt{1-2\sigma^2} \to 0$, and the deterministic ratio condition holds with $s^2 = q = 0$. Let $\mathcal{F}_{r,d}$ be the class of compatible arrays with equal weak-source masses $\beta \in (1/2, 1)$, common first-anchor mean r, and $0 \leq s_j^2 \leq B/4$, $|q_j| \leq ds_j^2/10$, without the upper mass bound or neighborhood box of (22u). Both sequences lie in this fixed class for small σ. No uniform modulus tending to zero is possible on $\mathcal{F}_{r,d}$, and uniform consistency of radius estimation fails there under the sampling contract of Section 6.1. Indeed, use independent Bernoulli arms for these two arrays. The weaker-probe probabilities are identical; all changed coordinates have interior limiting probabilities. Their n-repeat KL divergence is $O(n\sigma^2)$. Setting $\sigma = n^{-1}$ makes the data laws indistinguishable in total variation while their radius gap stays bounded below. The same two-point testing argument as in Corollary 5.1 then gives $\inf_{\widehat{\varepsilon}_n} \sup_{P:\, \mathbf{v}(P) \in \mathcal{F}_{r,d}} E_P|\widehat{\varepsilon}_n - \varepsilon_*(\mathbf{v}(P))| \geq c_0 > 0$ for all sufficiently large n. This failure holds even at the fixed, known mean r; it concerns the larger mass-unrestricted class, while (22u) remains a compact interior neighborhood.

Estimated moment targets already enter population-share sensitivity frameworks [16, Section 4.2], and fourth-root rates occur in overspecified mixture inverse problems [21, Section 1.1]. Wu and Yang [22, Sections 2–4 of the author version] develop moment-space projection, Gauss quadrature and moment-to-Wasserstein comparison for Gaussian mixture estimation. The target here is the minimum shared-calibration radius from finite Bernoulli marginals. The specific extension is the joint environment-preserving mass/mean repair and the uniform response geometry, which allow the raw-array modulus to remain sharp as the nuisance values vary.

3.7. The critical boundary when the mass margin shrinks

Let p denote a known scale for the unknown private masses $1 - \beta_j$. We now track constants as $p \downarrow 0$. Fix $r \in (0, e^{-2})$ and $e \in (0, 1/2)$. Choose a fixed compact interval $J \subset (0, e^{-2})$ containing r in its interior and put $d = \min(\inf J, 1 - \sup J)/4$, $B = 2d^2$. Write $s = r^{1+4e}$ and define the reference intermediate probabilities $Y_1(p) = r/2 + (1-2p)s/2$ and $Y_2(p) = (1-2p)r/2 + s/2$. For small

constants $a, p_0 > 0$, consider the triangular local classes

$$\begin{aligned}\mathcal{T}_p(a) = \{\mathbf{v} \in \mathcal{K}_{1/2} : |\beta_j - (1-p)| \le ap, \quad |r_j - r| \le ap, \quad r_j \in J,\\ 0 \le s_j^2 \le a^2p^2, \quad |q_j| \le ds_j^2/10,\\ |v_{2,j} - Y_j(p)| \le ap\}, \quad 0 < p \le p_0.\end{aligned} \tag{22y}$$

The normalized moments are defined from each candidate array as in Section 3.6. The constants a, p_0 are fixed before sampling; p may be a sequence p_n. The true masses and means remain unknown and unequal. The class localizes the exposure standard deviations and response departures to the private-mass scale.

The constant choices are uniform across mass scales. Fix r, e, J, d, B and the cross-scale enlargement factor $A = 8$. Choose the geometric constants of Lemma 10.2, including K and t_0, from the fixed inputs. Choose a maximal tube width $a_* > 0$ small enough for that geometry, the anchor repair and $e + Ka_* < 1/2$, then reduce p_0 as required. The single-scale theorem fixes any $0 < a \le a_*$; cross-scale and adaptive inference fix $0 < a \le \min\{1/16, a_*/A\}$. For each such fixed width, choose the small-distance threshold $c_a > 0$ to satisfy $c_a p_0 \le a$ and the repair/opening conditions in Theorem 10. Constants c, C may depend on these fixed choices and, for fixed-level confidence bounds, on α; they are independent of n, p, p'. The larger class $\mathcal{F}_{r,d}$ has the separate impossibility result proved in Section 3.6.

Lemma 10.1 (uniform anchor transport). The repair in Lemma 8.1 has response cost $C\sqrt{\Delta}$ with a constant independent of p for arrays in $\mathcal{T}_p(a)$, where Δ is their raw-array distance. The same holds for the relaxed tube with a replaced by $2a$.

Proof. Choose p_0 so all masses exceed $1/2$. Raw-to-normalized moments then have uniform Lipschitz constants, and the first means lie in fixed J. In the activation step the conditional probability may involve $(1-b)^{-1}$, but its total activated probability is exactly $b' - b$. Thus its raw-response cost and its change of normalized moments are $O(|b' - b|)$ uniformly, since normalization divides by $b' \ge 1/2$. Thinning has the same bound. The affine mean maps and the fixed-d, B higher-moment kernel also have constants independent of p. Validity of the activation probability follows directly from $b' \le 1$ and does not require a uniform upper mass bound. All later steps retain the adjusted calibration marginals. □

Lemma 10.2 (uniform frontier geometry). There are constants $K, c, C, t_0 > 0$, independent of p, such that, for sufficiently small fixed a, p_0:

(a) Every array in $\mathcal{T}_p(2a)$ has minimum radius in $[e - Ka, e + Ka]$.

(b) Any common/private representation of such an array carried at radius $\eta \le e + Ka$ has a band $[l, h]$ with $l \le Ca$ and $h \le 1 - \tau$ for a fixed $\tau > 0$. Each private block has at least cp mass with its active exposure in a fixed compact subinterval of $(0, 1)$. A fixed positive common mass has both exposures in that interval and $U \ge l + b$ for a fixed $b > 0$.

(c) On the enlarged band $[l, h + t]$, $0 < t \le t_0$, the same anchors admit every intermediate-response pair within cpt of the original pair.

Proof. Constants in this proof depend only on r, e, J, d; they may increase from line to line. From (22y), $V_j(r) = \beta_j[s_j^2 + (r_j - r)^2] \le Ca^2p^2$, $|\beta_1 - \beta_2| \le Cap$, and the two private masses are between $(1 - 2a)p$ and $(1 + 2a)p$. The observed intermediate contrast is $pF_e + O(ap)$, with $F_u = r(1 - r^{4u})$.

For any band $[l, h]$ of width at most 2η, replace each exposure by r in the intermediate response at a cost Cap, using the derivative bound three. The common contribution cancels. The private contributions give $v_{2,1} - v_{2,2} \le p[f_l(r) - f_h(r)] + Cap \le pF_\eta + Cap$ for $\eta \le 1/2$. Since F_u' is bounded below by $4(-\log r)r^3 > 0$ on $[0, 1/2]$, this proves the lower half of (a).

For the upper half, set $u = e + Ka < 1/2$ and choose any exposure laws representing the anchors. Uniformly in those laws, $M_j = ER_j^{1+4u} = r^{1+4u} + O(ap)$; the centering remainder is $O(a^2p^2)$ and the mean displacement is $O(ap)$. In the endpoint allocation used in Theorem 8, the lower-endpoint

calibration masses are $z_j = (v_{2,j} - \beta_j M_j)/(r_j - M_j)$. Their denominators stay bounded away from zero. At the reference array the least dominating mass equals $1 - p + pF_e/F_u$, and each normalized endpoint fraction is bounded away from zero and one. The perturbed least dominating mass is therefore at most $1 - cp(u - e) + Cap$. Choose K sufficiently large, then a sufficiently small. The endpoint allocation has mass at most one, proving feasibility at $e + Ka$ and (a).

For a representation as in (b), the sharper band inequality above implies $F_e \le r^{2l} F_{e+Ka} + Ca$. Consequently $l \le Ca$, and $h \le 2e + Ca < 1 - \tau$ after reducing a. The common mass is $\kappa = \beta_1 + \beta_2 - 1 = 1 - 2p + O(ap)$. Subtracting the first private contribution, which is $O(p)$, from $v_{2,1}$ shows that the common average of $f_U(r)$ equals $(r + s)/2 + O(p)$. Choose p_0 small so this average stays a fixed distance below r. Since $l \le Ca$, a fixed positive common mass must have $U \ge l + b$, for fixed small $b > 0$. The quadratic deviations bound the active mass of exposures outside a fixed interior interval around r by Ca^2p^2. This is smaller than half either private mass and smaller than the common calibration mass just obtained, uniformly for small a, p_0. Thus both private-interiority and joint common-interiority claims hold.

Apply the three perturbations of Lemma 6.2. Their response increments are $(-a_1 t, 0)$, $(0, -a_2 t)$, and $(b_1 t, b_2 t)$, now with $cp \le a_i \le Cp$ and $c \le b_i \le C$. Their edge distances after division by t are $a_1 a_2/(a_1 + a_2)$, $a_1 b_2/(a_1 + b_1 + b_2)$, and $a_2 b_1/(a_2 + b_1 + b_2)$. Each is bounded below by a constant times p. Hence the convex hull contains a ball of radius cpt, proving (c). Using a single lower bound for all four coefficients would lose a spurious factor of p; the separate block scales give the stated factor. □

Theorem 10 (mass-scaled modulus and matching hard pairs). There are $a_*, p_0 > 0$ such that, for each fixed $0 < a \le a_*$, constants $c, C > 0$ give, for every $0 < p \le p_0$,

$$|\varepsilon_*(\mathbf{v}) - \varepsilon_*(\mathbf{w})| \le \min\left\{\tfrac{1}{2}, \frac{C}{p}\|\mathbf{v} - \mathbf{w}\|_\infty^{1/2}\right\}, \qquad \mathbf{v}, \mathbf{w} \in \mathcal{T}_p(a). \tag{22z}$$

For each sufficiently small $0 < \sigma \le cp$, the class contains a pair at probability distance $\beta\sigma^2$, with equal weak-source masses $\beta_1 = \beta_2 = \beta = 1 - p$, whose minimum-radius gap is at least $c\sigma/p$. The constants are uniform in both p and σ.

Proof. If $\Delta = \|\mathbf{v} - \mathbf{w}\|_\infty$ is at least a sufficiently small fixed multiple of p^2, the radius range $[0, 1/2]$ gives (22z) by increasing C. Otherwise Lemma 10.1 repairs the target anchors on an optimal band of $\mathbf{v}$ with intermediate error at most $C\sqrt{\Delta}$. Restrict this case to $\Delta \le c_a p^2$, where c_a is a sufficiently small fixed constant depending on a. The repaired array then lies in $\mathcal{T}_p(2a)$. Lemma 10.2 applies to its common/private representation, using (a) for the source carrying radius; its opening corrects the remaining responses with $t \le C\sqrt{\Delta}/p \le t_0$. The band-radius increase is at most $t/2$. Reversing the arrays proves the modulus. The intermediate range $c_a p^2 < \Delta$ is again covered by the trivial bound.

For the lower pair, use (22e) and (22i) with $\beta = 1 - p$, $\kappa = 1 - 2p$, radius e and dispersion σ, and its deterministic-anchor comparison from Theorem 5. These arrays lie in $\mathcal{T}_p(a)$ when σ/p is sufficiently small: their response displacement is $O(\sigma)$, their variance is σ^2, and their centered-third-moment ratio is $2p\sigma/\sqrt{\kappa}$. Both share all intermediate responses, masses and first anchors, so their raw-array distance is exactly $\beta\sigma^2$.

The following bounds are uniform as $p \downarrow 0$ and do not use a fixed-β asymptotic remainder. The explicit construction gives $d_\sigma = pF_e + \sigma\sqrt{\kappa}G_e + O(\sigma^2)$ with remainder at most $\beta C_r \sigma^2$. For small p_0, $\kappa \ge 1/2$. In the deterministic comparison put $z = r^{4e}$. At zero dispersion the three branch margins are $t_- - z = \kappa(1 - z)/(2\beta)$, $z - r^2 > 0$, and $1 - t_+ = \kappa(1 - z)/(2\beta)$. Each has a positive lower bound independent of $p \le p_0$, since $e < 1/2$ and $\kappa \ge 1/2$. Their perturbations are $O(\sigma/p)$. Thus the branch $z_* = 1 - D$ persists uniformly when σ/p is sufficiently small. Its radius satisfies $F_{\varepsilon_{\mathrm{det}}} = d_\sigma/p$. The inverse function has uniformly bounded positive derivative in a fixed neighborhood of e. Therefore $\varepsilon_{\mathrm{det}} - e \ge c\sigma/p$. The heterogeneous array is realized at radius e, so its minimum radius is at most e, giving the claimed gap. All exposure nodes are positive for the chosen constants. □

Lemma 10.3 (comparing different mass scales). Fix $A = 8$ and $0 < a \le \min\{1/16, a_*/A\}$. For

$\mathbf{v} \in \mathcal{T}_p(a)$, $\mathbf{w} \in \mathcal{T}_{p'}(a)$, $|\varepsilon_*(\mathbf{v}) - \varepsilon_*(\mathbf{w})| \le \min\{1/2, C\sqrt{\Delta}/p\}$, where $\Delta = \|\mathbf{v} - \mathbf{w}\|_\infty$. Constants are independent of p, p'.

Proof. For $\Delta \ge c_a p^2$, the capped radius range suffices. Otherwise the masses in the two tubes imply $|p'-p| \le a(p+p')+\Delta$. Since $c_a p_0 \le a$, we have $\Delta \le ap$, hence $|p'-p| \le 4ap$ and $p' \le (1+4a)p \le 5p/4$. Moreover, $Y_1(p) = (r+s)/2 - ps$ and $Y_2(p) = (r+s)/2 - pr$, so $|Y_j(p') - Y_j(p)| \le 4ap$. The mass and response windows for $\mathbf{w}$, measured at scale p, are consequently at most $(5/4+4)ap < 8ap$; its mean window and standard deviations are at most $5ap/4$. The centered-moment cone is unchanged. Thus $\mathbf{w} \in \mathcal{T}_p(Aa)$ with the fixed $A = 8$. The source array is in this relaxed tube as well. Apply Theorem 10 there, using $Aa \le a_*$. □

Theorem 11 (critical triangular minimax boundary). Observe n complete repeats per preparation in the sampling experiment of Section 6.1. For every prescribed sequence $0 < p_n \le p_0$, without a monotonicity requirement, use the class $\mathcal{T}_{p_n}(a)$. Its scale p_n is known; the individual β_j, r_j remain unknown. The minimax expected absolute error R_n, defined in Section 6.1, satisfies, with constants independent of n, p_n,

$$c \min\left\{1, \frac{n^{-1/4}}{p_n}\right\} \le R_n(\mathcal{T}_{p_n}(a)) \le C \min\left\{1, \frac{n^{-1/4}}{p_n}\right\}. \tag{22aa}$$

Uniform consistency holds exactly when $np_n^4 \longrightarrow \infty$.

A nearest raw-mean projection followed by the minimum-radius functional attains the upper bound. Confidence-box inversion on this same class, with empty output when the intersection is empty, is honest at level $1-\alpha$ and has all-sample-path diameter at most $\min\{1/2, C[\log(20/\alpha)/n]^{1/4}/p_n\}$. For every fixed $0 < \alpha < 1/4$, its worst-case expected diameter is minimax of order $\min\{1, n^{-1/4}/p_n\}$.

Proof. For each $p > 0$ the class is nonempty and compact, and (22z) makes its radius functional continuous. The finite-count measurable projection and ten-mean concentration argument of Theorem 7 therefore applies. Taking expectations in the capped modulus gives the upper bound in (22aa), and confidence inversion gives the stated diameter bound and coverage.

For the lower bound choose the hard pair in Theorem 10 with $\sigma = c_1 \min(n^{-1/4}, p_n)$. Only four anchor probabilities differ; the changed anchor probabilities stay in a fixed compact subset of $(0,1)$ for $p_n \le p_0$. The identical weak-probe probabilities contribute zero divergence even when they approach one. Thus the n-repeat product-Bernoulli KL is at most $Cn\sigma^4 \le Cc_1^4$. Choose c_1 small to make total variation at most $1/4$. The radius gap is at least $c\sigma/p_n = cc_1 \min\{1, n^{-1/4}/p_n\}$. The two-point risk inequality proves the lower bound. The simultaneous-coverage argument of Theorem 7 proves the expected-diameter lower bound for fixed $\alpha < 1/4$. Both points belong to the same triangular class for every sufficiently large n.

The minimax risk tends to zero exactly when $p_n n^{1/4} \to \infty$. If np_n^4 fails to diverge, it is bounded along a subsequence, and the lower bound stays positive on that subsequence. This covers oscillating mass scales as well. Because the radius lies in $[0, 1/2]$, truncation to that interval makes uniform expected-error consistency and uniform probability consistency equivalent. This proves the consistency assertion. □

Corollary 11.1 (adaptation without a known mass scale). Fix the width a in the cross-scale range of Lemma 10.3 and let

$$\mathcal{U} = \bigcup_{0<p\le p_0} \mathcal{T}_p(a), \qquad \mathcal{U}_n = \bigcup_{n^{-1/2}\le p\le p_0} \mathcal{T}_p(a). \tag{22ab}$$

For sufficiently large n, project the ten empirical means onto $\mathcal{U}_n$ and report the minimum radius of a nearest point. This single estimator does not use the true p and has, for every $0 < p \le p_0$, worst-case expected error on $\mathcal{T}_p(a)$ at most $C \min\{1, n^{-1/4}/p\}$. Thus it attains the triangular minimax rate simultaneously over unknown mass scales. For the finitely many n with $n^{-1/2} > p_0$, take the estimate to be zero.

For confidence inference set $b_n = \sqrt{\log(20/\alpha)/(2n)}$ as in Theorem 7 and use $\mathcal{C}_n = \mathcal{U} \cap \{\mathbf{v} : \|\mathbf{v} - \widehat{\mathbf{v}}\|_\infty \le b_n\}$ and the closed interval from the infimum to the supremum of its radius values, with empty output for empty $\mathcal{C}_n$. The infimum and supremum are used because $\mathcal{U}$ need not be compact. This confidence set is honest over $\mathcal{U}$. On each $\mathcal{T}_p(a)$, its diameter is at most $\min\{1/2, C[\log(20/\alpha)/n]^{1/4}/p\}$ on the ten-mean coverage event. For each fixed $0 < \alpha < 1/4$, its worst-case expected diameter on every such class is at most $C_\alpha \min\{1, n^{-1/4}/p\}$, matching the lower bound of Theorem 11.

Proof. The joint set of $(p, \mathbf{v})$ defining $\mathcal{U}_n$ is compact. Lemma 10.3 makes the radius continuous on its projection $\mathcal{U}_n$; finite-count selection again gives a measurable nearest projection. If $p \ge n^{-1/2}$, projection error is at most twice empirical raw error, and Lemma 10.3, with the true scale in its denominator, gives the asserted risk bound. If $p < n^{-1/2}$, the target bound is of constant order and bounded radius loss suffices. All interval endpoints are measurable because there are only finitely many empirical count vectors at fixed n.

The true array belongs to $\mathcal{C}_n$ on the ten-mean event, proving honesty even without compactness. On that event, Lemma 10.3 compares every candidate to the true array and gives the diameter bound.

For expected diameter, also use the weaker event that each empirical weak-source probability differs from its true mean by at most $p/8$. Its complement has probability at most $4e^{-np^2/32}$. If $b_n \le p/8$, every candidate in $\mathcal{C}_n$ has both weak-source deficits between fixed positive multiples of the true p; therefore each candidate's tube scale is comparable to p. Lemma 10.3 applied between any two candidates then bounds the diameter by $C\sqrt{b_n}/p$, without requiring that the true array belongs to the confidence box. The exceptional contribution to expectation is at most $2e^{-np^2/32}$, which is bounded by a constant times $\min\{1, n^{-1/4}/p\}$. For fixed α, the regime $b_n > p/8$ is covered by the trivial bound, with an adjusted constant. This proves the expected-diameter assertion. □

The adaptive procedure uses a fixed union of frontier tubes and includes mass uncertainty inside its raw confidence region. Unlike the known-scale interval, its scale-dependent diameter guarantee is a high-probability and expected guarantee, rather than an all-sample-path guarantee. Its expected guarantee uses concentration of the weak-source frequencies in addition to the anchor-noise bound.

The three regimes are explicit. If $np_n^4 \to \infty$, the radius is uniformly estimable at order $(np_n^4)^{-1/4}$. If np_n^4 stays bounded, some radius differences remain statistically unresolved; if $np_n^4 \to 0$, the minimax error stays of constant order. For example, $p_n = n^{-\gamma}$ gives rate $n^{\gamma - 1/4}$ for $0 < \gamma < 1/4$, while $\gamma \ge 1/4$ leaves a nonvanishing minimax error. Since $(1-a)p_n \le 1 - \beta_j \le (1+a)p_n$, the same scale describes either preparation's actual mass margin, up to fixed factors.

The frontier-tube specification is fixed before sampling. Theorem 11 uses a known scale p_n; Corollary 11.1 adapts over the specified scale union without knowing its true member. Both infer the masses and first anchors jointly from the raw data. Exposure standard deviations at most ap_n and response localization at scale ap_n keep this triangular experiment near the deterministic joint frontier. These are sufficient locality conditions for the matched boundary. The earlier mass-margin counterexample has exposure standard deviation of order $\sqrt{p}$, outside this frontier tube; its full moment-cone impossibility result remains in force. The quantitative mass factor comes from the private/common response triangle, while the fourth-root factor comes from the quadratic anchor separation in the hard pair.

3.8. Jointly shrinking mass and heterogeneity scales

The standard-deviation restriction in (22y) can be enlarged to a variance budget proportional to the private mass. Keep the fixed reference r, e, J, d, B and $Y_j(p)$ from Section 3.7. For fixed sufficiently small $a, b, p_0 > 0$, and prescribed $0 \le h^2 \le bp$, define

$$\begin{aligned} \mathcal{V}_{p,h} = \{\mathbf{v} \in \mathcal{K}_{1/2} : {} & |\beta_j - (1-p)| \le ap, \quad |r_j - r| \le ap, \quad r_j \in J, \\ & 0 \le s_j^2 \le h^2, \quad |q_j| \le d s_j^2/10, \\ & |v_{2,j} - Y_j(p)| \le ap\}. \end{aligned} \tag{22ac}$$

Both masses and means vary across candidate arrays. The new scale h bounds standard deviation; it is not a band endpoint. Constants below depend on the fixed inputs and fixed a, b, but not on p, h, n. Write $\mathcal{V}_{p,h}^{(2a)}$ for the same class with the mass, first-mean and intermediate-response windows ap replaced by $2ap$; the variance bound h^2, the cone and J are unchanged. The minimax statements concern this standard-deviation envelope, including deterministic arrays and least-favorable small-dispersion pairs.

Choose K from the fixed reference inputs in a fixed neighborhood of e, then take a, b small enough for the radius bound below. Choose p_0 small enough for the private-tail and common-average margins and $bp_0 \le B/4$. Finally choose $t_0 \le \min(\tau/2, u_0/2)$ using the fixed margins constructed in the proof. These choices apply to the original and relaxed classes with one common set of constants. The class includes standard deviations much larger than p: for example $h = p^{3/4}$ satisfies the variance budget for all sufficiently small p. This range has finite positive-exposure realizations: take each normalized anchor law to put equal weights at $r - h, r + h$ and set the intermediate probabilities to $Y_j(p)$. Its third centered moment is zero; the endpoint allocation in Lemma 12.1 below realizes the two responses jointly for small p. The class retains response localization at scale p and the centered-moment cone.

Lemma 12.1 (two-sided opening under a variance budget). For sufficiently small fixed a, b, p_0, every array in $\mathcal{V}_{p,h}$ has minimum radius at most $e + K(a + b) < 1/2 - 2\tau$, for fixed $K, \tau > 0$. Any representation of such an array at a carrying radius at most that bound admits all intermediate-response pairs within *cpt* on the enlarged band $[\max(0, l - t), \min(1, H + t)]$, for $0 < t \le t_0$. Anchors are held fixed, and the required radius increases by at most t. The same K, τ, c, t_0 apply to $\mathcal{V}_{p,h}^{(2a)}$, with the same upper carrying-radius bound $e + K(a + b)$.

Proof. First $V_j(r) \le C(h^2 + a^2p^2)$. For any anchor-representing exposure law, centered Taylor expansion at its mean gives $ER_j^{1+4u} = r^{1+4u} + O(ap + h^2)$, uniformly for u in a fixed neighborhood of e. For the endpoint allocation put $z_j = (v_{2,j} - \beta_j M_j)/(r_j - M_j)$ and $L = \max_j z_j + \max_j(\beta_j - z_j)$. The denominators stay uniformly positive. With $L_0 = 1 - p + pF_e/F_u$ at the reference, uniform Lipschitz bounds give $|L - L_0| \le C_0(ap + h^2)$; choose C_0 to cover both the original and $2a$ windows. In the fixed neighborhood of e, $1 - L_0 = p(1 - F_e/F_u) \ge c_0 p(u - e)$ for $u \ge e$, with $c_0 > 0$. Since $ap + h^2 \le (a + b)p$, choosing $K > C_0/c_0$ and setting $u = e + K(a + b)$ makes the reference slack absorb the perturbation. Then choose a, b small enough to keep u in that neighborhood. Endpoint fractions stay strictly between zero and one. This proves the radius bound. Choose τ so the bound is less than $1/2 - 2\tau$.

Use the canonical common/private decomposition (9). Private block masses are comparable to p. Fix $J_0 = [r - \delta_0, r + \delta_0] \Subset (0, 1)$ with $\delta_0 > 0$. For each active measure, $\nu_j(R_j \notin J_0) \le C_1(h^2 + a^2p^2) \le C_1(b + a^2p_0)p$, where C_1 covers the $2a$ mean window as well. Each private block has mass at least $(1 - 2a)p$, hence retains at least $[(1 - 2a) - C_1(b + a^2p_0)]p \ge p/2$ inside J_0 after choosing a, b, p_0 small. On the common block exclude the union $\{R_1 \notin J_0\} \cup \{R_2 \notin J_0\}$; its mass is at most $2C_1(h^2 + a^2p^2)$, and both retained exposure coordinates lie in J_0.

Its normalized intermediate-response average, after replacing exposure by r, equals $(r + s)/2 + O(p + \sqrt{bp})$. For its normalized calibration law Q, write $m_Q = E_Q f_U(r)$. Reducing p_0 puts m_Q in $[r^3 + \zeta, r - \zeta]$ for a fixed $\zeta > 0$. Choose $u_0 > 0$ with $f_{u_0}(r) \ge r - \zeta/2$ and $f_{1-u_0}(r) \le r^3 + \zeta/2$. With $q_+ = Q(U \ge u_0)$ and $q_- = Q(U \le 1 - u_0)$, monotonicity gives $m_Q \ge (1 - q_+)(r - \zeta/2) + q_+ r^3$ and $m_Q \le q_- r + (1 - q_-)(r^3 + \zeta/2)$. Combining these with the interval for m_Q gives $q_+, q_- \ge \zeta/[2(r - r^3)]$. The common mass is bounded below independently of p. Excluding the joint exposure tails costs at most $C(h^2 + a^2p^2)$, so both retained common masses remain positive with both exposure coordinates interior.

Write the carrying band as $[l, H]$. Its width is at most $1 - 4\tau$. If $l \ge \tau$ and $H \le 1 - \tau$, move calibration both up and down on each private block. These four axis increments have magnitude at least *cpt* for $t \le \tau/2$, and their convex hull contains a ball of that order. If $l < \tau$, then $H < 1 - 3\tau$. Increasing calibration on each private block gives the two negative axis directions, each of size comparable to

pt. Decreasing common calibration by t, clipped at zero, gives a positive two-coordinate direction with both coefficients bounded above and below by fixed positive constants, using the common mass with $U \geq u_0$. The three edge-distance formulas in Lemma 10.2 give a cpt ball. If $H > 1 - \tau$, the width bound forces $l > 3\tau$; decrease private calibration and increase common calibration, clipped at one. The common mass with $U \leq 1 - u_0$ gives the reflected triangle and the same bound. These cases exhaust all bands. All changes preserve exposures and weights, and use the stated enlarged band. Its width increases by at most $2t$, proving the radius assertion. □

Theorem 12 (two-scale modulus and matching statistical boundary). Fix the class (22ac) before sampling and observe n complete repeats per preparation in $\mathfrak{P}_n(\mathcal{V}_{p,h})$. Put $\Gamma_h(\Delta) = \Delta + \min\{\sqrt{\Delta}, h\}$. Uniformly over $0 < p \leq p_0$ and $0 \leq h^2 \leq bp$,

$$|\varepsilon_*(\mathbf{v}) - \varepsilon_*(\mathbf{w})| \leq \min\{1/2, C\Gamma_h(\|\mathbf{v} - \mathbf{w}\|_\infty)/p\}, \qquad \mathbf{v}, \mathbf{w} \in \mathcal{V}_{p,h}. \tag{22ad}$$

The minimax absolute risk satisfies

$$R_n(\mathcal{V}_{p,h}) \asymp \min\left\{1, \frac{n^{-1/2} + \min\{h, n^{-1/4}\}}{p}\right\}. \tag{22ae}$$

A nearest raw-mean projection onto this same class attains the upper bound. Confidence-box inversion, with empty output of diameter zero, is honest and has all-sample-path diameter at most $\min\{1/2, C[b_n + \min\{\sqrt{b_n}, h\}]/p\}$. At each fixed $0 < \alpha < 1/4$ its worst-case expected diameter has the matching order (22ae). Here $b_n = \sqrt{\log(20/\alpha)/(2n)}$.

Proof. The mass and affine-mean steps cost $O(\Delta)$ and leave each source exposure variance at most h^2. At the target mean there are two repairs. Lemma 6.1 costs $C\sqrt{\Delta}$. Alternatively replace the exposure independently by any law representing the target anchors: since source and target now have the same mean and standard deviations at most h, the expected absolute change is at most $2h$. Both kernels keep the adjusted calibration marginal. Choosing the cheaper kernel yields total response error at most $C\Gamma_h(\Delta)$.

Start with an optimal representation of $\mathbf{v}$. If $\Gamma_h(\Delta) \leq c_a p$ for a sufficiently small fixed c_a, the repaired array has target anchors, variance at most h^2, and response windows at most $2ap$. Lemma 12.1 applies on the source carrying band and corrects the remaining responses using $t \leq C\Gamma_h(\Delta)/p$. Reverse the arrays to obtain (22ad). If $\Gamma_h(\Delta) > c_a p$, increase C and use the bounded radius range.

For each p, h the class is nonempty and compact, including $h = 0$. The modulus makes the radius continuous, so the finite-count projection argument applies. If η is empirical raw error, projection error is at most 2η. Ten-mean concentration gives $E\eta \leq Cn^{-1/2}$ and $E\sqrt{\eta} \leq Cn^{-1/4}$. Also $E\min(\sqrt{\eta}, h) \leq \min(E\sqrt{\eta}, h)$. The capped modulus proves the risk upper bound, and inversion of the same raw box proves coverage and the pathwise diameter bound.

Two separate pairs give the matching lower bound. For the regular pair take deterministic anchors at $\beta = 1 - p, r$, keep the second group's response $Y_2(p)$, and change the first from $Y_1(p)$ to $Y_1(p) + \delta$, where $\delta = c_1 \min(n^{-1/2}, p)$. Both arrays belong to (22ac), even when $h = 0$, for small fixed c_1. The uniform deterministic branch margins in Theorem 10 persist. Their radii satisfy $F_{\varepsilon_\delta} = F_e + \delta/p$, giving gap at least $c\delta/p$. Only one interior Bernoulli probability changes, so product-law KL is at most $Cn\delta^2$.

For the heterogeneous pair use Theorem 10's explicit construction and deterministic-anchor comparison with $\sigma = c_1 \min(h, n^{-1/4}, p)$. Both are in the same $\mathcal{V}_{p,h}$: $\sigma \leq h$, σ/p is small, and the response and moment-cone restrictions hold. Their radius gap is at least $c\sigma/p$, and KL is at most $Cn\sigma^4$. For $h = 0$ this second lower bound is zero. Choose c_1 small to bound both total variations by $1/4$. The two testing inequalities give the maximum of $\min(1, n^{-1/2}/p)$ and $\min(1, h/p, n^{-1/4}/p)$. This maximum is within fixed factors of (22ae). The simultaneous-coverage argument for each pair gives the same lower order for expected confidence diameter at fixed $\alpha < 1/4$. □

For every prescribed pair of sequences p_n, h_n satisfying $h_n^2 \leq bp_n$, the necessary and sufficient

condition for uniform consistency on these moving classes is

$$np_n^2 \longrightarrow \infty, \qquad \min\{h_n/p_n,\ n^{-1/4}/p_n\} \longrightarrow 0. \tag{22af}$$

This follows directly from (22ae), including nonmonotone sequences. If $h_n/p_n \to 0$, the heterogeneity part of the consistency criterion vanishes and the remaining requirement is $np_n^2 \to \infty$. If h_n/p_n is bounded below by a positive constant, the criterion reduces to $np_n^4 \to \infty$. Thus the earlier critical rate persists even when standard deviations substantially exceed p_n, throughout the small variance-to-private-mass regime. At $h_n = 0$, the risk is $\min\{1, n^{-1/2}/p_n\}$ and uniform consistency is equivalent to $np_n^2 \to \infty$. For example, $p_n = n^{-3/10}$ and $h_n = p_n^2$ give risk of order $n^{-1/5}$ although $np_n^4 \to 0$. In contrast, $h_n = p_n^{3/4}$ lies in the enlarged variance-budget class and retains the $np_n^4 \to \infty$ criterion.

The known inputs in Theorem 12 are the envelope scales p_n, h_n and the fixed local specification; the individual masses and means remain unknown. Corollary 11.1 provides unknown-mass-scale adaptation for its original tube family. Section 3.9 treats unknown envelopes on the larger class and identifies the additional inferential costs.

The variance budget has a geometric role: $h^2 \le bp$ prevents exposure tails from consuming an entire private block and gives slack in the endpoint allocation. The full-cone counterexample in Section 3.6 has variance p, so it lies outside this result when b is sufficiently small. Classification at unrestricted variance-to-mass ratios requires additional geometry. The extension proved here replaces the previous standard-deviation condition $O(p)$ by the substantially larger range $O(\sqrt{p})$ with a fixed small squared coefficient.

The local classes separate the available inferential guarantees as follows. Each row also carries its stated centered-moment cone and response-locality restrictions; the risks are minimax over the indicated envelopes.

Class	Prescribed inputs; unknown nuisance values	Variance envelope	Expected absolute-risk order	Scale adaptation established
$\mathcal{D}_\rho$	Fixed β, r and local reference	$s_j^2 \le B_r/4$ within the local box	$n^{-1/4}$	Fixed class
$\mathcal{E}_\rho$	Compact nuisance ranges and reference; β_j, r_j unknown	$s_j^2 \le B/4$ within the local box	$n^{-1/4}$	Joint nuisance inference within the fixed class
$\mathcal{T}_p(a)$	Envelope p; β_j, r_j unknown in ap windows	$s_j^2 \le a^2p^2$	$\min\{1, n^{-1/4}/p\}$	Corollary 11.1 adapts over its prescribed union of p scales
$\mathcal{V}_{p,h}$	Envelopes p, h; β_j, r_j unknown in ap windows	$s_j^2 \le h^2 \le bp$	$\min\{1, [n^{-1/2} + \min(h, n^{-1/4})]/p\}$	Theorems 13–14 quantify the costs of unknown envelopes

3.9. Unknown envelopes: estimation and honest confidence have different costs

Let $\mathcal{W}_p = \mathcal{V}_{p,\sqrt{bp}}$, $\mathcal{Z}_p = \mathcal{V}_{p,0}$, and

$$\mathcal{W} = \bigcup_{0<p\le p_0} \mathcal{W}_p, \quad \mathcal{W}_n = \bigcup_{n^{-1/2}\le p\le p_0} \mathcal{W}_p, \quad \mathcal{Z}_n = \bigcup_{n^{-1/2}\le p\le p_0} \mathcal{Z}_p. \tag{22ag}$$

The constants a, b, p_0 remain specified; neither the true mass scale nor an actual standard-deviation envelope is supplied to the procedures below. The full-union procedure optimizes fourth-root worst-case risk and honest expected diameter; the thresholded procedure instead preserves the deterministic regular risk and optimizes the remaining cost subject to that protection. Choose a, b small enough that the fixed enlargements $8a, 2b$ also satisfy Section 3.8. Write $\theta(\mathbf{v}) = \varepsilon_*(\mathbf{v})$.

Lemma 13.1 (cross-scale comparison and deterministic collapse). If $\mathbf{v} \in \mathcal{W}_p$ and $\mathbf{w} \in \mathcal{W}_{p'}$, then $|\theta(\mathbf{v}) - \theta(\mathbf{w})| \le \min\{1/2, C\sqrt{\Delta}/p\}$, where $\Delta = \|\mathbf{v} - \mathbf{w}\|_\infty$. For $\mathbf{v} \in \mathcal{Z}_p, \mathbf{w} \in \mathcal{Z}_{p'}$, the sharper bound is $\min\{1/2, C\Delta/p\}$.

Let $C\mathbf{v}$ replace the second and third normalized exposure moments by r_j^2, r_j^3, keeping masses, first anchors and intermediate responses. For $\mathbf{v} \in \mathcal{W}_p$, this collapse belongs to $\mathcal{Z}_p$. If $s = \max_j s_j$, then $\|\mathbf{v} - C\mathbf{v}\|_\infty \le Cs^2$ and $|\theta(\mathbf{v}) - \theta(C\mathbf{v})| \le \min\{1/2, Cs/p\}$.

Proof. For $\Delta \le c_a p$ with $c_a \le a$, the mass windows imply $|p' - p| \le 4ap$ and $p' \le 5p/4$, as in Lemma 10.3. The two arrays lie in one class at scale p with windows $8ap$ and variance budget at most $(5/4)bp$. Apply Theorem 12 with these fixed enlarged constants. Its modulus is bounded by $C\sqrt{\Delta}/p$ on the probability cube. For two deterministic arrays use its $h = 0$ linear modulus instead. When $\Delta > c_a p$, bounded radii give both claimed capped bounds after increasing C.

The collapse has positive deterministic exposure means in J. The endpoint-allocation argument of Lemma 12.1 with zero variance and the unchanged mass, mean and response windows directly constructs its full-band representation; hence it is in $\mathcal{Z}_p$. Its raw anchor changes are $\beta_j s_j^2$ and $\beta_j(3r_j s_j^2 + q_j)$, bounded by Cs^2 under the cone. Theorem 12 on $\mathcal{V}_{p,s}$ then gives the radius comparison. □

Theorem 13 (unknown-scale upper bound and sharp honest-confidence obstruction). A nearest projection onto $\mathcal{W}_n$ followed by θ has expected absolute error at most $C\min\{1, n^{-1/4}/p\}$ on every $\mathcal{V}_{p,h}$, simultaneously over unknown p, h. For the finitely many n with an empty scale range take estimate zero.

Invert the raw ten-mean box on $\mathcal{W}$, using the closed interval of infimum and supremum radius values and empty output for an empty intersection. This single confidence procedure is honest over $\mathcal{W}$. At fixed $0 < \alpha < 1/4$,

$$\inf_{I_n:\ \mathrm{coverage}_{\mathcal{W}} \ge 1-\alpha} \ \sup_{P \in \mathfrak{P}_n(\mathcal{V}_{p,h})} E_P \operatorname{diam}(I_n) \asymp \min\{1, n^{-1/4}/p\}, \qquad 0 \le h^2 \le bp. \tag{22ah}$$

The upper bound is attained by that one procedure. In particular, the same lower bound holds with the supremum restricted to $\mathcal{Z}_p$.

Proof. The sets $\mathcal{W}_n, \mathcal{Z}_n$ are compact projections of compact joint $(p, \mathbf{v})$ constraint sets. Lemma 13.1 gives continuity of the radius on them, and finite-count selection gives measurable projections. If $p \ge n^{-1/2}$, the true array lies in $\mathcal{W}_n$ and projection error is at most twice empirical raw error. The cross-scale modulus and concentration prove the risk bound. If $p < n^{-1/2}$, the capped target bound is constant and bounded loss suffices.

The full-union box covers the true array on the ten-mean event, proving honesty. On that event all candidate radii are within $C\sqrt{b_n}/p$ of the true radius. To control expected diameter, use the additional event $|\widehat{\beta}_j - \beta_j| \le p/8$ for both groups, whose complement has probability at most $4e^{-np^2/32}$. If $b_n \le p/8$, all surviving candidates have tube scales comparable to p. Lemma 13.1 applied between any two candidates bounds their diameter by $C\sqrt{b_n}/p$ on this event, even if the true array is outside the box. The exceptional expected contribution is at most $2e^{-np^2/32}$, absorbed by the capped rate. If $b_n > p/8$, bounded diameter gives the same order with a constant depending on fixed α. This is a high-probability and expected, rather than a scale-dependent all-path, diameter bound.

For the lower bound use the heterogeneous pair from Theorem 10 with $\sigma = c_1 \min(n^{-1/4}, p)$. Take p_0 small enough that $\sigma^2 \le bp$ throughout. Its deterministic comparison $\mathbf{b}^\sigma$ belongs to $\mathcal{Z}_p$, hence to every $\mathcal{V}_{p,h}$, while its heterogeneous partner belongs to $\mathcal{W}_p$. Their radius gap is at least $c\sigma/p$ and their product-law total variation is at most $1/4$. Any interval honest over $\mathcal{W}$ must cover both radii under the deterministic law with probability at least $1 - 2\alpha - 1/4 > 1/4$. This proves the lower bound at the deterministic comparison sequence itself, and thus (22ah). All interval endpoints are measurable by the finite-count argument, even though $\mathcal{W}$ need not be compact. □

Consequently a full-union honest interval can have uniformly shrinking expected diameter on the deterministic classes $\mathcal{Z}_{p_n}$ exactly when $np_n^4 \to \infty$. Knowing $h_n = 0$ instead permits the weaker

condition $np_n^2 \to \infty$ in Theorem 12. The obstruction is already present at known p: the deterministic comparisons and heterogeneous partners above share the same mass scale. The lower bound concerns a sequence of deterministic arrays inside the class, not every individual deterministic array.

Point estimation can retain the deterministic rate, but constant-factor adaptation to all the oracle rates in (22ae) is impossible. Its optimal cost under that deterministic-rate requirement is logarithmic. Put

$$q_n = \{\log(en)/n\}^{1/4}, \qquad \Psi_n(p,h) = \min\{1, [n^{-1/2} + \min(h, q_n)]/p\}. \tag{22ai}$$

Theorem 14 (sharp cost of adaptation with uniform deterministic-rate protection). There is one estimator, using neither p nor h, whose risk on every $\mathcal{V}_{p,h}$ is at most $C\Psi_n(p,h)$. In particular its risk on $\mathcal{Z}_p$ is at most $C\min\{1, n^{-1/2}/p\}$.

Conversely fix $0 < A < \infty$. Suppose there is a finite N_A, independent of p, such that the estimator sequence satisfies the following protection simultaneously for every $n \ge N_A$ and every $0 < p \le p_0$:

$$\sup_{P \in \mathfrak{P}_n(\mathcal{Z}_p)} E_P|\widehat{\theta}_n - \theta(\mathbf{v}(P))| \le A\min\{1, n^{-1/2}/p\} \tag{22aj}$$

Then there are $c_A > 0$ and $N_A' < \infty$ such that, for every $n \ge N_A'$ and every admissible p, h,

$$\sup_{P \in \mathfrak{P}_n(\mathcal{V}_{p,h})} E_P|\widehat{\theta}_n - \theta(\mathbf{v}(P))| \ge c_A\Psi_n(p,h). \tag{22ak}$$

Here N_A' may depend on N_A, A and the fixed class constants, but not on p, h; c_A depends only on A and the fixed class constants. For any protection constant at least that achieved by the displayed construction, the upper and lower bounds match within this uniformly protected family.

Proof of the upper bound. Let $\widehat{\mathbf{v}}^Z, \widehat{\mathbf{v}}^W$ be nearest raw projections onto $\mathcal{Z}_n, \mathcal{W}_n$. Define $\widehat{\beta}_j^c = \max(1/2, \widehat{v}_{-,j})$ and $\widehat{S} = \max_j[\widehat{v}_{+,j}/\widehat{\beta}_j^c - (\widehat{v}_{1,j}/\widehat{\beta}_j^c)^2]_+$. This statistic uses a variance proxy rather than an inferred envelope. It obeys $|\widehat{S} - s^2| \le C\eta$, where $s = \max_j s_j$ and $\eta = \|\widehat{\mathbf{v}} - \mathbf{v}\|_\infty$, because the clipped normalization is uniformly Lipschitz and true masses exceed $1/2$.

For a sufficiently large fixed L, set $t_n = L\sqrt{\log(en)/n}$. Report $\theta(\widehat{\mathbf{v}}^Z)$ if $\widehat{S} \le t_n$, and $\theta(\widehat{\mathbf{v}}^W)$ otherwise. For empty scale ranges report zero. If $p < n^{-1/2}$ the target Ψ_n is constant, so consider $p \ge n^{-1/2}$.

On the deterministic-projection branch, compare with $C\mathbf{v} \in \mathcal{Z}_p$. Projection error relative to it is at most $2(\eta + Cs^2)$. Lemma 13.1 and the collapse comparison bound the radius error by $\min\{1/2, C(\eta + s)/p\}$. If $h^2 \le t_n/2$, the probability of selecting the other branch is at most $20(en)^{-D}$, with D arbitrarily large by choosing L. Its bounded contribution is absorbed by $n^{-1/2}/p$. The risk is therefore at most the capped order $(n^{-1/2} + h)/p$.

If $h^2 > t_n/2$, selection of the deterministic branch implies $s \le \sqrt{t_n} + \sqrt{C\eta}$. Its expected error is at most $C(q_n + n^{-1/2})/p$. The other branch contributes at most the unconditional risk of the full-union projection, $Cn^{-1/4}/p$. Since $\sqrt{t_n}$ is a fixed multiple of q_n, in this regime $\min(h, q_n)$ is comparable to q_n. In the first regime $h \le \sqrt{L/2}q_n$, so its bound is also within a fixed factor of Ψ_n. The two regimes prove the upper bound, including the case $h = 0$.

Proof of the lower bound. Use the deterministic comparison law P_0 and heterogeneous law P_1 of Theorem 10. For their full n-repeat likelihood ratio $\Lambda = dP_1/dP_0$,

$$E_0\Lambda^2 \le \exp(Cn\sigma^4). \tag{22al}$$

Indeed a changed Bernoulli coordinate with probabilities x, y has second likelihood moment $1 + (x - y)^2/[y(1-y)]$. All changed coordinates are uniformly interior and have differences $O(\sigma^2)$; unchanged coordinates contribute exactly one. Multiply over coordinates and independent repeats.

Set $\sigma = c_1\min(h, p, q_n)$, with a sufficiently small fixed c_1, and let $d_\sigma = |\theta(\mathbf{a}^\sigma) - \theta(\mathbf{b}^\sigma)| \ge c\sigma/p$. If $\sigma \le K_A n^{-1/2}$, then either $p \le (K_A/c_1)n^{-1/2}$, so the capped regular lower bound is bounded below by an A-dependent positive constant, or $\min(h, q_n) \le (K_A/c_1)n^{-1/2}$, so Ψ_n is at most a fixed

A-dependent multiple of the regular rate. Theorem 12 therefore supplies the required lower bound in both cases. Otherwise (22aj), Markov's inequality and Cauchy–Schwarz give

$$P_1\{|\widehat{\theta}_n - \theta(\mathbf{b}^\sigma)| \geq d_\sigma/2\} \leq \exp(Cn\sigma^4/2)\sqrt{C_A n^{-1/2}/\sigma}. \tag{22am}$$

Choose K_A large and c_1 small so this is at most $1/2$ uniformly for $K_A n^{-1/2} \leq \sigma \leq c_1 q_n$. To see uniformity, the logarithm of the right side has a single interior minimum as a function of σ. At the lower endpoint it is made small by K_A; at the upper endpoint it is bounded by a constant times $n^{Cc_1^4/2-1/8}\{\log(en)\}^{-1/8}$, which tends to zero. Under P_1 the complementary event then has estimation error at least $d_\sigma/2$. This yields risk at least $c\sigma/p$. Combining with the regular lower bound gives $c_A\Psi_n$, including the capped regime. □

Uniformity of the protection start time is necessary for moving-envelope conclusions. To see this, let T_n be the protected estimator above and set $c_n = n^{-1/10}$. Use the full-union projection whenever both empirical weak-source deficits are at most $2c_n$, and T_n otherwise. At every fixed $p > 0$, the extra branch eventually has exponentially small probability, so this gated estimator satisfies regular protection from a possibly p-dependent start time. On the moving classes $p_n = c_n, h_n = q_n$, however, that branch has exponentially high probability and risk $O(n^{-1/4}/p_n)$, below Ψ_n by a diverging logarithmic factor. Thus protection only eventually at each fixed p cannot replace (22aj).

For example, take $p_n = n^{-1/10}$ and $h_n = q_n$. Oracle risk (22ae) has order $n^{-1/4}/p_n$, whereas uniform deterministic-rate protection forces order q_n/p_n, a factor $\{\log(en)\}^{1/4}$ larger. Thus simultaneous constant-factor oracle adaptation fails even when both envelopes tend to zero well inside $h_n^2 \leq bp_n$. Under the uniform protection requirement (22aj), the matching consistency criterion is $np_n^2 \to \infty$ and $\min(h_n/p_n, q_n/p_n) \to 0$.

The distinction between coverage on a union and expected length on a subclass is central to adaptive confidence theory [23, Section 2]. Constrained-risk methods quantify the cost of unusually small risk on a simpler class [24, Section 2 of the 2018 author version]. The proofs here instantiate these principles through legal finite-Bernoulli comparison pairs, their σ^2 probability separation and σ/p radius gap. Unknown mass scale can be accommodated by joint projection; unknown heterogeneity introduces the interval obstruction and the logarithmic point-estimation cost just proved.

3.10. Finite-sample certificates for the joint boundary

The overlap formulation is exact for all anchor laws, including their moment boundaries. Finite confidence certificates can test the entire resulting feasible set. Use the compact convex set $\mathcal{K}_B \subset \mathbb{R}^{10}$ defined in Section 3.1. For coefficient vectors $d_j \in \mathbb{R}^5$, define

$$s_j(u; d_j) = \max_{0 \leq r \leq 1} d_j \cdot g(r, u), \qquad h_B(d) = \max_{u \in B}\{[s_1(u; d_1)]_+ + [s_2(u; d_2)]_+\}. \tag{23}$$

The exact dual boundary is

$$\mathbf{v} \in \mathcal{K}_B \quad \Longleftrightarrow \quad d \cdot \mathbf{v} \leq h_B(d) \text{ for every } d \in \mathbb{R}^{10}. \tag{24}$$

This is finite-dimensional convex separation, within the moment-dual framework [14]. The two positive parts encode the choice of active preparation weights; omitting them or normalizing the groups separately changes the model.

For a fixed error radius write $\mathcal{K}_\epsilon = \bigcup_{x \in [0,1]} \mathcal{K}_{B_x}$, $B_x = [\max(0, x - \epsilon), \min(1, x + \epsilon)]$. This union is compact and need not be convex. Testing it by one unrestricted convex-hull separation would discard restrictions associated with the central ratio.

Theorem 15 (complete finite covering certificates). Cover $[0, 1]$ by finitely many center intervals $J_g = [a_g, b_g]$. Put $\widetilde{B}_g = [\max(0, a_g - \epsilon), \min(1, b_g + \epsilon)]$. For each g choose d_g and a certified upper bound on $h_{\widetilde{B}_g}(d_g)$. Reject when a lower confidence bound for $d_g \cdot \mathbf{v}$ exceeds that upper bound for every g.

(a) On the eighteen-frequency event this rule never rejects $\mathcal{K}_\epsilon$. It can be combined with the other certificates in this paper, and coefficients and intervals can be selected from the same fixed-size data if all support bounds are valid.

(b) Every array outside $\mathcal{K}_\epsilon$ has such a finite strictly separating certificate. More quantitatively, if $\operatorname{dist}_\infty(\mathbf{v}, \mathcal{K}_\epsilon) \geq \delta > 0$, a center cover with $O(1 + \delta^{-1})$ intervals and coefficient vectors of total ℓ^1 norm one suffices. There is an absolute C such that

$$n \geq C\delta^{-2} \log(36/\alpha) \tag{25}$$

complete repeats per preparation give rejection on the simultaneous event, hence power at least $1 - \alpha$. This is a guarantee for valid support certificates; it does not assume an unverified numerical optimizer returns them.

(c) Over the full marginal model and unrestricted within-repeat coupling contract, the $\delta^{-2} \log(1/\alpha)$ order is necessary for a test with size α and power at least $3/4$ uniformly over alternatives at distance at least δ, for sufficiently small δ and $\alpha \leq 1/8$. The lower-bound alternatives are ordinary shape violations in the full class.

Proof. At one environment each active weight can be maximized over its endpoints zero and one. The exposure coordinates maximize independently conditional on u, giving (23). The separating-hyperplane characterization of a compact convex set gives (24).

For a linear functional the raw means satisfy

$$d_j \cdot \mathbf{v}_j = (\textstyle\sum_{k=0}^{4} d_{k,j}) v_{1,j} - d_{0,j}\delta_{-,j} - d_{2,j}\delta_{+,j} - d_{3,j}\delta_{3,j} - d_{4,j}\delta_{2,j}. \tag{26}$$

The corresponding absolute-coefficient radius is valid for every coefficient vector on the same event. A null central ratio lies in some J_g and its calibration band is contained in $\widetilde{B}_g$. It cannot be separated there. This proves (a) without a further multiple-testing allocation.

For (b), the function r^{1+2u} has u derivative bounded by $2/e$ on the whole compact domain, with the continuous zero-node extension. Moving a calibration value by h therefore moves each coordinate of a generated array by at most $2h/e$. The same bound holds after mixing. If every center interval has length at most $e\delta/8$, the outer-band class for that interval lies within $\delta/4$ of the class at any chosen center in it. It is still at distance at least $3\delta/4$ from the alternative. Convex separation in the ℓ^∞ metric supplies a coefficient vector of ℓ^1 norm one with at least that margin. This yields the stated finite cover. A concrete fixed-separation procedure scans a finite rational net of the coefficient unit sphere, with mesh a sufficiently small constant times δ, and computes support upper bounds to the same accuracy. The support functions and data functionals are Lipschitz in those coefficients on the bounded observation domain, so this finite search retains a positive gap. Its computational cost is separate from the number of replicates.

Every empirical primitive mean has radius at most $C_0(\sqrt{t/n} + t/n)$, since its frequency is at most one. A raw arm void has radius at most three times this bound. Thus (26), with total coefficient norm one, has radius at most $3C_0(\sqrt{t/n} + t/n)$ uniformly over all selected coefficients. Under (25), its estimate and lower endpoint together spend less than half the separation margin. All intervals then reject on the event. Support bounds can themselves be obtained to any positive tolerance by finite grids: the r derivative of $d_j \cdot g$ is bounded by $3\|d_j\|_1$, and the u derivative by $2|d_{4,j}|/e$. Grid maxima with these remainder bounds give certified upper approximations. Rational nodes and exponents admit outward root enclosures as in Appendix C. Strict gaps also permit rational approximation of separating coefficients. Thus no exact nonlinear optimization oracle is required for strict certificates.

For (c), use the null with $R_j = 1/2$, $A_j = 3/4$, $U = 1$; all five marginal void probabilities are strictly inside $(0, 1)$, and $v_{2,j} = v_{3,j} = 3/32$. Choose independent Bernoulli arms and independent groups, which are allowed couplings. Change only the first group's intermediate void to $3/32 - 2\delta$. Every null array has $v_{2,1} \geq v_{3,1}$, so this alternative is at ℓ^∞ distance at least δ from the full null union. The

product-law Kullback–Leibler divergence for n repeats is at most $C_1 n\delta^2$, with an absolute constant for small δ. Any rejection event with alternative probability at least $3/4$ and null probability at most α has binary divergence at least $(3/4)\log(1/\alpha) - \log 2$. The data-processing inequality gives the claimed order. □

Increasing the number of center cells or searching more separating directions spends no additional confidence probability on the same event. Each reported support bound still requires certification. Theorem 15 gives a finite search construction; Section 7 specifies the implemented closed-form and scalar certificates.

4. Marginal benchmarks and simpler certificates

4.1. Exact marginal profiles at zero calibration error

The half-step design $s = 1/2$ admits an exact description of the shared-position information. We first impose $\epsilon = 0$, so $x_E = x$ is an unknown constant common to the preparations. Normalize at the weaker probe and define

$$R_j = e^{-\Delta_3 N_j^*/2}, \qquad d\mu_j = \frac{Z_j e^{-\phi_- N_j^*}}{v_{-,j}}\, dP, \qquad m_{k,j} = \frac{v_{a_k,j}}{v_{-,j}}, \quad (a_1,a_2,a_3) = (1,+,3). \tag{27}$$

Then $0 < R_j \le 1$, $m_{k,j} = E_{\mu_j} R_j^k$, and

$$w_j := \frac{v_{2,j}}{v_{-,j}} = E_{\mu_j} R_j^{1+2x}. \tag{28}$$

Thus four anchor probabilities identify the first three moments of a probability measure on the unit interval. The unknown intermediate source position becomes an unknown fractional exponent shared by two moment problems.

For one preparation suppress j. We give explicit formulas for the interior moment class

$$H_R = \begin{pmatrix} m_1 & m_2 \\ m_2 & m_3 \end{pmatrix} \succ 0, \qquad H_{1-R} = \begin{pmatrix} 1-m_1 & m_1-m_2 \\ m_1-m_2 & m_2-m_3 \end{pmatrix} \succ 0. \tag{29}$$

Put $d = m_1 - m_2$, $e = m_2 - m_3$ and define a three-point measure μ_L on $\{0, z, 1\}$ by

$$z = e/d, \qquad \lambda_z = \frac{d}{z(1-z)}, \quad \lambda_1 = \frac{m_1 m_3 - m_2^2}{d-e}, \quad \lambda_0 = \frac{(1-m_1)e - d^2}{e}. \tag{30}$$

Define a two-point measure μ_G as follows:

$$S = \frac{m_3 - m_1 m_2}{m_2 - m_1^2}, \qquad p = S m_1 - m_2, \qquad a, b = \frac{S \mp \sqrt{S^2 - 4p}}{2}, \qquad g_a = \frac{b - m_1}{b - a}, \quad g_b = 1 - g_a. \tag{31}$$

Under (29), all weights are positive and sum to one, $0 < z < 1$, and $0 < a < b < 1$. Both measures match $m_0 = 1, m_1, m_2, m_3$. Set

$$\begin{aligned} G_r &= g_a a^r + g_b b^r, \qquad J_r = \lambda_z z^r + \lambda_1, \\ F_-(x) &= \min(G_{1+2x}, J_{1+2x}), \quad F_+(x) = \max(G_{1+2x}, J_{1+2x}). \end{aligned} \tag{32}$$

The notation J_r distinguishes the Lobatto sum from the determinant L.

Theorem 16 (sharp shared-calibration profile). Suppose $0 < v_{-,j} \le 1$ and (29) holds for both preparations.

(a) Among all measures on $[0,1]$ with the given first three moments, the exact range of ER^{1+2x} is

$$[F_-(x), F_+(x)]. \tag{33}$$

For $1 < r < 2$ the lower and upper endpoints are G_r, J_r; for $2 < r < 3$ their order reverses. At $r = 1, 2, 3$ both equal m_r.

(b) Each set

$$I_j = \{x \in [0,1] : F_{-,j}(x) \le w_j \le F_{+,j}(x)\} \tag{34}$$

is a possibly empty closed interval. The given marginal probability array is in the closure of the finite-covering-mean representation (1), with half-step probes and a common constant ratio, if and only if

$$I_1 \cap I_2 \neq \varnothing. \tag{35}$$

(c) Under the pointwise error contract (5), let $I_j = [A_j, B_j]$ be the profiles computed from the observed anchors and intermediate probabilities. A necessary condition is

$$[\max(0, A_1 - \epsilon), \min(1, B_1 + \epsilon)] \cap [\max(0, A_2 - \epsilon), \min(1, B_2 + \epsilon)] \neq \varnothing. \tag{36}$$

An empty I_j already excludes its preparation. In particular, a profile gap $A_1 - B_2 > 2\epsilon$ excludes the shared-ratio contract. Part (c) is a necessary condition; the necessary-and-sufficient assertion in (b) uses $\epsilon = 0$.

Proof. We provide the low-order argument explicitly. The two matrices in (29) are the integrals of $R(1,R)^{\mathsf{T}}(1,R)$ and $(1-R)(1,R)^{\mathsf{T}}(1,R)$ when a representing measure exists. Conversely, their strict positivity gives $d > e > 0$, the positive weights in (30), and direct substitution verifies all four moments of μ_L. This also proves existence of a representing measure from (29).

The polynomial $q(t) = t^2 - St + p$ is orthogonal to $1, t$ under μ_L. It has two distinct zeros in $(0,1)$. Indeed, if it had fewer than two interior sign changes, multiplying it by the polynomial of degree at most one with those sign-change zeros would have constant sign on $[0,1]$ and would not vanish on all three support points; its integral would contradict orthogonality. Polynomial division by q shows that evaluation at its two roots with the weights in (31) integrates every polynomial of degree at most three correctly. The weights are positive: the integral of the square of either degree-one cardinal polynomial equals the corresponding weight. This proves the claims about μ_G.

For $f(t) = t^r$, $1 < r < 2$, the fourth derivative is positive on $(0,1]$. Its cubic Hermite interpolant at the two double nodes a, b is a pointwise minorant, since its remainder has the sign of $(t-a)^2(t-b)^2$. The cubic interpolant at the single nodes $0, 1$ and double node z is a pointwise majorant, since its remainder has the sign of $t(t-z)^2(t-1)$. To cover the nonsmooth zero endpoint, apply the remainder formula first to $(t+\eta)^r$, $\eta > 0$, and pass to its uniform limit as $\eta \downarrow 0$. Integrating the cubic polynomials uses only the four matched moments and yields $G_r \le ER^r \le J_r$. For $2 < r < 3$, the fourth derivative is negative and the inequalities reverse. Integer exponents give the known moments. Convex mixtures of μ_G and μ_L attain every value between the two bounds. These are the classical Gaussian and Lobatto extremal constructions; their quadrature exactness and derivative-sign principle are reviewed in [13, Sections 2.1 and 2.2.2]. Their role here is to eliminate the preparation-specific mixing laws while retaining the unknown common exponent.

For a finite covering mean, $R = 0$ is excluded. The Gaussian extremum is attained; the noninteger Lobatto extremum is a closure point. To preserve the first three moments while removing its zero atom, replace its nodes by $\delta, z_\delta, 1$, where

$$z_\delta = \frac{e - \delta d}{d - \delta(1 - m_1)}. \tag{37}$$

Solve the first three linear moment equations for the weights. The identity $E[(R-\delta)(R-z_\delta)(R-1)] = 0$ then gives the fourth equation. For sufficiently small $\delta > 0$ the nodes and weights remain positive and converge to (30). Mixing these measures with μ_G attains every strict intermediate value with positive nodes. Thus the finite-positive-R ranges are $[G_r, J_r)$ for $1 < r < 2$ and $(J_r, G_r]$ for $2 < r < 3$, with the stated closed ranges as their closures.

Both functions in (32) are continuous and strictly decreasing on $[0,1]$, because each quadrature sum has a positive-weight node in $(0,1)$. Therefore (34) is a closed interval. It is empty if $w \notin [m_3, m_1]$; $w = m_1, m_2, m_3$ respectively gives the singleton $0, 1/2, 1$.

For sufficiency in (b), choose a common x in the intersection and measures realizing (33). Use positive-node approximations when needed. On the common product environment $E = (R_1, R_2)$ set

$$Z_j = v_{-,j}, \quad N_j^* = -\log R_j, \quad (\phi_-, \phi_1, \phi_+, \phi_3) = (0, 1, 2, 3), \quad \phi_2 = 1 + 2x.$$

This reconstructs all the marginal probabilities, exactly or in the required limit. Necessity follows from the change of measure. For a convergent sequence of compatible arrays, take a subsequence of their common ratios converging in $[0, 1]$ and of their measures converging weakly on $[0, 1]$. Joint continuity of t^r on $[0, 1] \times [1, 3]$ transmits both moments and intermediate responses to the limit.

For (c), write $x_\pm = \min(1, \max(0, x \pm \epsilon))$. Pointwise monotonicity gives $F_-(x_+) \le w \le F_+(x_-)$. Inverting the two decreasing boundaries gives $A - \epsilon \le x \le B + \epsilon$, intersected with $[0, 1]$. Applying this to both preparations proves (36). □

Theorem 16 characterizes the zero-error marginal compatibility boundary for this design and interior moment class. Sections 3.1–3.10 retain the absolute weak-source masses to resolve the nonzero-error common-environment problem. The variance inequality in Theorem 18 is a simpler necessary outer constraint. Its constant-step generality and explicit error allowance remain useful when half-step probes are unavailable. On half-step data, the sharp profile can contain substantially more information than that outer inequality.

Corollary 16.1 (when expanded profiles are sufficient). In the interior moment class (29), if $\beta_1 + \beta_2 \le 1$, the expanded-profile intersection (36) is necessary and sufficient for compatibility in the closure under the pointwise error contract. If $\beta_1 + \beta_2 > 1$, it is in general only necessary.

Proof. For fixed $B = [l, h]$, one preparation's exact range is $F_{-,j}(h) \le w_j \le F_{+,j}(l)$. Monotonicity gives necessity. Constant $U = h$ and $U = l$ with the extremal measures of Theorem 16 attain the two endpoints; mixtures of these joint (R, U) measures fill the interval while preserving the anchor moments. This condition is equivalent to $I_j \cap B \neq \varnothing$. When $\kappa = 0$, Theorem 1 places the two groups on disjoint active blocks, so the separate conditions are sufficient. Existence of a common center gives (36).

For failure above the threshold, take the same deterministic $R = r \in (0, 1)$ and the same $\beta > 1/2$ in both groups, with normalized intermediate responses $w_1 = r^{1+2l}$ and $w_2 = r^{1+2h}$, where $l < h$. Each separately fits the band, but strict monotonicity forces its entire calibration marginal to opposite endpoints. Their overlap is zero, contradicting $\kappa = 2\beta - 1 > 0$. Theorem 2 gives the full boundary for this anchor class. Section 3.3 supplies a counterexample with both moment matrices strictly positive, so the insufficiency also holds inside (29). □

4.2. Polynomial certificates from separate profiles

The same eighteen frequencies used in Theorem 18 certify a profile separation through cubic polynomial bounds. Their use follows the moment-dual principle [14]; the shared unknown source position supplies the two incompatible inequalities.

Choose a threshold x_0 and an error budget satisfying $0 \le \epsilon < \min(x_0, 1/2 - x_0)$. Put $r_h = 1 + 2(x_0 + \epsilon)$ and $r_l = 1 + 2(x_0 - \epsilon)$, so both lie in $(1, 2)$. Let $p(t) = \sum_{k=0}^{3} p_k t^k$ be the cubic Hermite interpolant of t^{r_h} at any two distinct double nodes in $(0, 1)$. Let $q(t) = \sum_{k=0}^{3} q_k t^k$ interpolate t^{r_l} at the single nodes $0, 1$ and any double node $z \in (0, 1)$. The preceding remainder argument gives the global bounds $p(t) \le t^{r_h}$ and $q(t) \ge t^{r_l}$ on $[0, 1]$.

For any cubic coefficient vector c, define the observable linear functional

$$\mathcal{A}_j(c) = c_0 v_{-,j} + c_1 v_{1,j} + c_2 v_{+,j} + c_3 v_{3,j}, \qquad H_1 = \mathcal{A}_1(p) - v_{2,1}, \quad H_2 = v_{2,2} - \mathcal{A}_2(q). \tag{38}$$

Theorem 17 (simultaneous linear separation). Under (1), (5) and half-step probes, H_1 and H_2 cannot both be positive. On the simultaneous event of Section 6.2, lower confidence endpoints for these two linear functionals therefore give a type-I-error-α rejection when both endpoints are positive. This rejection can be combined with both legs of Theorem 18 on that same event.

Proof. If the common central ratio satisfies $x \le x_0$, then $x_E \le x_0 + \epsilon$, and

$$v_{2,1}/v_{-,1} = E_{\mu_1} R_1^{1+2x_E} \ge E_{\mu_1} R_1^{r_h} \ge E_{\mu_1} p(R_1).$$

Thus $H_1 \le 0$. If $x \ge x_0$, the reversed monotonicity and the majorant give $H_2 \le 0$. Every x satisfies at least one of these alternatives.

For completeness, the signed differences in (2) give

$$\mathcal{A}_j(c) - v_{2,j} = (\textstyle\sum_k c_k - 1)v_{1,j} + \delta_{2,j} - c_0\delta_{-,j} - c_2\delta_{+,j} - c_3\delta_{3,j}. \tag{39}$$

Use the baseline estimate and signed-difference estimates from Section 6.2. A valid radius is

$$\mathcal{S}_j(c) = |\textstyle\sum_k c_k - 1| b_j + r_{2,j} + |c_0| r_{-,j} + |c_2| r_{+,j} + |c_3| r_{3,j}. \tag{40}$$

Subtract this radius from the estimate of H_1 or H_2, with the appropriate sign. Every linear functional is covered on the same event. Hence every combined false rejection lies in the complement of that event. □

The polynomial certificate needs no safe logarithms and remains defined when log rectangles are unsafe. Nodes and thresholds can be fixed by design. Selecting them from the same data also preserves the mathematical coverage argument if each selected polynomial has the proved global bound: the event covers the entire finite mean vector, and (39)–(40) then hold for every coefficient vector. The numerical implementation uses fixed nodes for the example below; it does not substitute rounded polynomial coefficients for an exact global-bound proof.

4.3. A shared-position exclusion beyond separate curves

Consider the probability measure assigning equal mass to $R = 1/10, 1/2, 9/10$. In both preparations take $v_- = 4/5$ and the anchor moments

$$(m_1, m_2, m_3) = (1/2, 107/300, 57/200). \tag{41}$$

For the intermediate arm use the same full mixture curve, evaluated at $x_1 = 3/8$ in preparation 1 and $x_2 = 1/8$ in preparation 2. Each preparation separately is generated by the constant-rate Poisson mixture $v_j(\phi) = (4/5)ER^\phi$, with anchor rates $0, 1, 2, 3$ and intermediate rate $1 + 2x_j$. Thus every single-preparation Laplace shape and moment restriction holds at its own intermediate position. Its conditional exposure has three well-separated values; both probe dispersions are strictly positive.

For (41), the Gaussian nodes are $a, b = 1/2 \mp 2\sqrt{6}/15$ with equal weights, and the Lobatto weights at $0, 1/2, 1$ are $16/75, 43/75, 16/75$. The profiles are

Preparation	Intermediate moment w_j	Common-ratio profile I_j, approximately
1	0.3822345169	$[0.3720447100, 0.3815973607]$
2	0.4510952373	$[0.1219377785, 0.1349282925]$

Their gap is approximately 0.2371164175. Equation (36) excludes every pointwise ratio-error budget below approximately 0.1185582087. The displayed profile endpoints are numerical summaries; the following algebraic witness certifies exclusion at the fixed nonzero budget $\epsilon = 1/10$ without relying on rounded profile roots.

Take $x_0 = 1/4$, use the Gaussian nodes for the minorant of $t^{17/10}$ and the Lobatto internal node $1/2$ for the majorant of $t^{13/10}$. Since the four moments match the corresponding quadrature rules, the exact population margins are

$$H_1 = \tfrac{4}{5}(G_{17/10} - w_1) > 0.00393195, \qquad H_2 = \tfrac{4}{5}(w_2 - J_{13/10}) > 0.00393297. \tag{42}$$

Outward rational root enclosures certify these inequalities. The two preparations consequently reject the shared calibration contract even with a pointwise ratio error of 0.1.

For the same array,

$$\begin{aligned} &u = -\log(57/100), \quad D_s = \log(107/75), \quad \overline{K} = 128/75, \\ &L = u\log(w_2/w_1) \simeq 0.0931117091, \quad B_0 = \tfrac{64}{171}u \simeq 0.2103836887. \end{aligned} \tag{43}$$

The original determinant allowance fails to reject even at zero ratio error. Its failure follows directly from $w_2/w_1 < 6/5$ and $\log(6/5) < 1/5 < 64/171$. The profile and linear certificate therefore strictly strengthen that allowance on the same anchor and intermediate information. Positive definiteness and both strict margins persist in a neighborhood of this array. The separation covers an open set of heterogeneous curves, beyond perturbations of a deterministic-exposure example.

A deterministic synthetic five-bit fixture with 10^{12} complete repeats per preparation, obtained by rounding a nested coupling of these marginals, yields lower endpoints above 0.0039293 and 0.0039310 at $\alpha = 0.05$ for (38). Appendix C records the counts and their rational certificates. The population construction diagnoses a failure of the shared-calibration contract.

4.4. Variance certificates for general probe steps

For $u, K \geq 0$ define

$$M(u,K) = u + \frac{1}{2}e^u K,$$

$$B_\epsilon(u,K) = \frac{1}{8}u_2 e^{u_1} K_1 + \epsilon\{u_2 M(u_1,K_1) + u_1 M(u_2,K_2)\}. \tag{44}$$

This function is nondecreasing in every component of u and K.

Theorem 18 (variance allowance and observable certification).
Suppose (1), (4) and (5) hold.

(a) If $\mathrm{Var}_{Q_j}(Y_j) \leq K_j < \infty$, then

$$L \leq B_\epsilon(u,K). \tag{45}$$

No positive exponential moment of Y_j is needed for this part.

(b) If the probes satisfy (6), let

$$D_{s,j} = \log \frac{v_{-,j}v_{+,j}}{v_{1,j}^2}.$$

Then

$$D_{s,j} \geq 0, \qquad \mathrm{Var}_{Q_j}(Y_j) \leq \overline{K}_j(s) := \frac{e^{D_{s,j}} - 1}{s^2}. \tag{46}$$

Thus (45) holds with $K_j = \overline{K}_j(s)$. The lifetime kernels can be unknown, and the original environment distribution can have infinite rate moments.

(c) Under the fixed-design replication contract in Section 6, eighteen Bernoulli frequencies provide simultaneous intervals at confidence $1-\alpha$. Write $\overline{D}_{s,j}$ for a safe upper endpoint for $D_{s,j}$. Reject if any such endpoint is negative. Otherwise, when all required log intervals are safe, put

$$K_j^+ = \frac{e^{\overline{D}_{s,j}} - 1}{s^2}, \qquad U_j = \max(0, -\underline{\ell}_{3,j})$$

and reject if

$$\underline{L} > B_\epsilon(U, K^+). \tag{47}$$

The union of these two rejection rules has type-I error at most α for every law satisfying the full contract. If a required interval is unsafe, that leg supplies no certificate. A safe probe group can still reject by curvature when the main determinant intervals are unsafe.

The first rejection uses classical log-Laplace convexity. The second calibrates a residual-mixture allowance using measured probe responses. Their use of a shared confidence event avoids an independence requirement between $\underline{L}$ and K^+.

4.4.1. Proof of the variance allowance

For a single preparation suppress the index j and write

$$h(t) = \log E_Q e^{-tY}, \quad 0 \le t \le 1, \quad u = -h(1), \quad \mu = E_Q Y.$$

Finite variance here includes a finite second moment. Differentiating under the expectation gives

$$h'(t) = -E_{Q_t} Y, \qquad h''(t) = \mathrm{Var}_{Q_t}(Y), \qquad \frac{dQ_t}{dQ} = \frac{e^{-tY}}{E_Q e^{-tY}}.$$

These identities extend to $t = 0$ by dominated convergence. Since $Y \ge 0$,

$$\mathrm{Var}_{Q_t}(Y) \le E_{Q_t}(Y - \mu)^2 \le e^u K. \tag{48}$$

Taylor's integral formula yields

$$u \le \mu \le u + \tfrac{1}{2} e^u K.$$

Moreover $E_{Q_t} Y$ decreases with t, because its derivative is $-\mathrm{Var}_{Q_t}(Y)$. Therefore

$$\sup_{0 \le t \le 1} |h'(t)| \le \mu \le M(u, K). \tag{49}$$

Convexity and (48) bound the chord gap:

$$0 \le d := xh(1) - h(x) \le \tfrac{1}{2} x(1-x) e^u K \le e^u K/8. \tag{50}$$

For completeness, subtract $\frac{1}{2} e^u K\, t(1-t)$ from $th(1) - h(t)$. The result has zero endpoints and nonnegative second derivative, hence is nonpositive.

By (5), $e^{-x_E Y}$ is between the exponentials with parameters $\max(0, x - \epsilon)$ and $\min(1, x + \epsilon)$. Taking expectations and logs, and using (49), gives

$$r := \log E_Q e^{-x_E Y} - h(x), \qquad |r| \le \epsilon M(u, K). \tag{51}$$

No independence between x_E and Y is used. In each preparation,

$$\ell_{2,j} = -xu_j - d_j + r_j, \qquad \ell_{3,j} = -u_j.$$

The common x term cancels:

$$L = u_2 d_1 - u_1 d_2 - u_2 r_1 + u_1 r_2. \tag{52}$$

Using $d_2 \ge 0$, (50) and (51) proves (45). Interchanging preparations gives the companion lower bound

$$L \ge -\tfrac{1}{8} u_1 e^{u_2} K_2 - \epsilon\{u_2 M(u_1, K_1) + u_1 M(u_2, K_2)\}. \tag{53}$$

The exponential-family covariance identity is standard [3, Section 3.5.1]; the one-sided allowance in (45) is its consequence for this preparation determinant.

4.4.2. Proof of probe certification

The weaker source rate gives

$$E_{Q_j} e^{sY_j} = v_{-,j}/v_{1,j} < \infty, \qquad E_{Q_j} e^{-sY_j} = v_{+,j}/v_{1,j}.$$

Thus all moments needed below exist. If Y' is an independent copy under Q_j, exchangeability gives

$$e^{D_{s,j}} = E_{Q_j} e^{sY_j} E_{Q_j} e^{-sY_j} = E \cosh[s(Y_j - Y')] \ge 1 + s^2 \mathrm{Var}_{Q_j}(Y_j). \tag{54}$$

This proves (46). The independent copy is used in the proof and does not impose independence between experimental arms.

The coefficient one in (46) is sharp. For an equally weighted two-point law $Y = m \pm w/2$, with $m \geq w/2$,

$$\frac{e^{D_s} - 1}{s^2 \operatorname{Var} Y} = \left[\frac{\sinh(sw/2)}{sw/2}\right]^2 \longrightarrow 1.$$

For deterministic Y, the certified variance is exactly zero. The improvement over a second-moment bound is its cancellation of a common mean. Feasibility of the weaker probe already implies a coarse finite second-moment bound through exposure domination, as shown in Appendix A; the measured probes provide the finer, mean-cancelled budget.

4.4.3. Resolution of the determinant allowance

Proposition 1 (resolution of the determinant allowance). For any ordered main void probabilities $v_{1,j} \geq v_{2,j} \geq v_{3,j} > 0$, regardless of whether the coverage null holds,

$$L \leq u_1 u_2, \qquad B_\epsilon(u, K) \geq 2\epsilon u_1 u_2. \tag{55}$$

Consequently the population margin $L - B_\epsilon(u, K)$ is nonpositive whenever $\epsilon \geq 1/2$ and $K_j \geq 0$. The curvature leg remains a separate restriction.

Proof. When $u_j > 0$, put $t_j = -\ell_{2,j}/u_j \in [0, 1]$. Then $L = u_1 u_2(t_1 - t_2) \leq u_1 u_2$. If either $u_j = 0$, ordering makes both relevant log responses zero and the same inequality holds. Since $M(u, K) \geq u$, (44) gives the second inequality. □

Every $x_E \in [0, 1]$ satisfies (5) with $x = 1/2, \epsilon = 1/2$. This uninformative envelope therefore cannot resolve the ordered-response alternatives through the determinant allowance. Useful residual-calibration information must be tighter. More specifically, for an alternative family with $u_j \asymp P$, $L \asymp P^3$ and $K_j = O(P^2)$, the variance term in (44) is $O(P^3)$ while the ratio-error term is at least of order ϵP^2. A positive cubic-scale margin requires $\epsilon = O(P)$; a sufficiently small constant and the variance allowance must still leave a strictly positive margin. The cooperative example in Appendix E is detected directly by curvature and does not by itself establish power for this determinant leg.

There is a concrete cross-preparation restriction beyond fitting a separate exponential curve to each preparation. Take $s = 1/4$, a common target factor $Z = 9/10$, reference rate $\phi_1 = 1$, endpoint rate $\phi_3 = 3$, and deterministic covering means $N_1^* = 1/10$, $N_2^* = 1/5$. Let the intermediate rates be 2 in the first preparation and 3/2 in the second, with probes at rates 1/2 and 3/2. The relative log probabilities are

Preparation	ℓ_2	ℓ_3	ℓ_-	ℓ_+	Required intermediate ratio
1	$-1/10$	$-1/5$	$1/20$	$-1/20$	$1/2$
2	$-1/10$	$-2/5$	$1/10$	$-1/10$	$1/4$

The baseline voids are $v_{1,j} = (9/10)e^{-N_j^*}$. Within each preparation every entry is on the exact exponential curve $v_j(\phi) = Ze^{-\phi N_j^*}$. Thus all its single-preparation Laplace shape and moment constraints hold. Both local probe dispersions equal zero, and the measured budgets in (46) are zero. Across preparations, however,

$$L = \frac{1}{50}, \qquad L - B_\epsilon(u, 0) = \frac{1 - 8\epsilon}{50}. \tag{56}$$

The determinant certificate separates this array for $\epsilon < 1/8$. The array isolates a preparation-dependent calibration change. Each separate curve fits its own intermediate position, while a common-position requirement excludes their joint fit. The boundary in Section 3 retains this restriction, and (44)–(47) supply a simpler variance allowance with a finite confidence construction.

5. What main-response data alone identify

5.1. Unknown calibration and the rank constraint

For linear old rates $\rho, 2\rho, 3\rho$, conditional Poisson coverage gives

$$q_{3,j} - 2q_{2,j} + q_{1,j} = -E[Z_j e^{-\rho N_j^*}(1 - e^{-\rho N_j^*})^2] \le 0. \tag{57}$$

This sign is preserved by environmental averaging. It is not preserved by arbitrary common nonlinear calibration. For example, take $\phi(a) = a^2$, $Z_j = z_j > 0$ and $N_j^* = c_j A(T)$, where $A(T) \downarrow 0$. Then

$$\frac{q_{3,j} - 2q_{2,j} + q_{1,j}}{A(T)} \longrightarrow 2z_j c_j \rho^2 > 0. \tag{58}$$

A second preparation provides a calibration-free leading restriction under a sparse-residue contract.

Proposition 2 (common-calibration response rank).
In (1), let the three ϕ_i be deterministic finite nonnegative numbers, common across preparations, without a monotonicity assumption. If

$$E[Z_j N_j^*] = O(A), \qquad E[Z_j (N_j^*)^2] = o(A), \qquad A \downarrow 0, \tag{59}$$

then $\mathscr{D}/A^2 \to 0$.

Proof. The bound $|e^{-z} - 1 + z| \le z^2/2$ for $z \ge 0$ gives

$$\delta_{i,j} = (\phi_i - \phi_1)m_j + r_{i,j}, \qquad |r_{i,j}| \le \tfrac{1}{2}(\phi_i^2 + \phi_1^2)E[Z_j (N_j^*)^2] = o(A).$$

The leading matrix is an outer product; its determinant vanishes. Each remaining term is $O(A)o(A)$ or $o(A)^2$. □

For a fixed, repeatably conditioned environment e, no sparsity is needed:

$$\ell_{i,j|e} = -(\phi_i(e) - \phi_1(e))N_j^*(e), \qquad L_e = 0. \tag{60}$$

Both statements require a calibration shared across preparations. Equation (60) takes conditional probabilities before logarithms. It does not assert that the determinant formed from unobserved environment averages is zero.

The experimental restriction in (59) is separability of a common control vector and a preparation vector. Equation (60) is its exact conditional counterpart. These are direct algebraic consequences of the probability model, with a different role from rank ordering in unknown-link regression [2].

5.2. Finite mixture saturation

Proposition 3 (finite response compatibility).
Given any finite collection of triples $0 \le q_{1,\lambda} < q_{2,\lambda} < q_{3,\lambda} < 1$, there is a common strictly increasing three-point calibration $(0, 1, R)$ and, for each label λ, a two-point environment satisfying (1) and matching all three probabilities exactly.

Proof. Put $v_i = 1 - q_i$, $u_\lambda = v_2/v_1$, $w_\lambda = v_3/v_1$ and choose a common integer

$$R > \max_\lambda \frac{\log w_\lambda}{\log u_\lambda}.$$

Then $R \ge 2$ and $u_\lambda^R < w_\lambda < u_\lambda < 1$. The increasing polynomial $g_R(z) = (1 - z^R)/(1 - z)$ has a unique root $z_\lambda \in (0, u_\lambda)$ of

$$g_R(z_\lambda) = \frac{1 - w_\lambda}{1 - u_\lambda}.$$

Set

$$Z_\lambda = v_1, \quad \theta_\lambda = \frac{1-u_\lambda}{1-z_\lambda}, \qquad N^*_\lambda = \begin{cases} 0 & \text{with probability } 1-\theta_\lambda, \\ -\log z_\lambda & \text{with probability } \theta_\lambda. \end{cases}$$

Both probabilities lie strictly between zero and one. Direct substitution gives $E[Z_\lambda e^{-\phi_i N^*_\lambda}] = v_i$ for $\phi_i = 0, 1, R$. □

The two-point support in Proposition 3 is labelwise. All labels can be placed on a common product environment, whose total support generally contains more than two points. A shared mother Poisson process and a common target Bernoulli factor within each label realize nested triples $X_1 \le X_2 \le X_3$. A nested binary triple has only the outcomes $000, 001, 011, 111$, whose probabilities are determined by its three marginals. Thus the construction also matches its complete triple law. For independent experimental blocks at finitely many labels it matches the product data law.

If labels include finitely many observation times, set $N^*_j(T) = 0$ after the largest selected time. Conditions (59) then hold eventually. Consequently the asymptotic contract alone supplies no nontrivial uniformly valid finite test for those independent blocks: a matching null law has exactly the same data distribution. This finite-label construction does not enforce a common lifetime kernel or a single space-time process. The next construction enforces both.

5.3. A fixed Cox grain process with false positive averaged determinants

At each integer site and birth row, let singleton grains have conditionally independent Poisson births and a fixed lifetime $h = 12k$, where $k \ge 1$ is an integer specifying one law. Draw a global environment once:

$$P(E = 1) = \varepsilon, \qquad \beta_0 = \varepsilon/(100k), \quad \beta_1 = 2/k. \tag{61}$$

An independent stream of rate $\eta > 0$ has lifetime survival

$$S(a) = (a+1)^{-2}, \qquad a = 0, 1, \ldots, \qquad B = \sum_{a \ge 0} S(a) < 2. \tag{62}$$

Every grain remains at its birth site and is present at age a when its lifetime exceeds a. This gives a spatially translation-invariant frozen-grain mechanism. At constant forcing its birth rule is also time-translation invariant. Old controls multiply both streams by $i = 1, 2, 3$; the target multiplier is one.

For any T, H, define

$$n_H = \min\{\lfloor HT \rfloor, (12k - T)_+\}, \quad m = \min(T, 12k),$$

$$J_H = \sum_{a=T}^{T+\lfloor HT \rfloor - 1} S(a), \quad J_0 = \sum_{a=0}^{T-1} S(a).$$

The same law for all times gives

$$v_{i,H}(T) = E \exp[-\beta_E(m + i n_H) - \eta(J_0 + i J_H)]. \tag{63}$$

This is a genuine conditional Poisson birth-and-lifetime model with a common target factor and nested preparation windows.

At $T = k$, $H_1 = 1, H_2 = 10$, first set $\eta = 0$. Then

$$v^0_{i,j} = (1-\varepsilon)e^{-c\varepsilon(1+iH_j)} + \varepsilon e^{-2(1+iH_j)}, \qquad c = 1/100. \tag{64}$$

As $\varepsilon \downarrow 0$, both $\mathscr{D}^0/\varepsilon^2$ and L^0/ε^2 converge to

$$\begin{aligned} K = c[&10e^{-4}(1-e^{-2})^2 - e^{-22}(1-e^{-20})^2] \\ &- e^{-26}(1-e^{-2})(1-e^{-20})(e^{-2} - e^{-20}) > 10^{-3}. \end{aligned} \tag{65}$$

To verify the expansion, write $\delta_{2,j} = \varepsilon[cH_j + a_j(1-b_j)] + O(\varepsilon^2)$ and $\delta_{3,j} = \varepsilon[2cH_j + a_j(1-b_j^2)] + O(\varepsilon^2)$, where $a_j = e^{-2(1+H_j)}$, $b_j = e^{-2H_j}$. Multiply the two pairs in (2). Since $v_{1,j} \to 1$, (3) has the same coefficient. The elementary bounds $.135 < e^{-2} < .136$, $e^{-4} > .018$, $e^{-22} < 10^{-9}$ and $e^{-26} < 10^{-11}$ prove the last inequality.

Take $\eta(\varepsilon) = \varepsilon^2/10000$. Since $J_0, J_H \le B < 2$, adding this nonzero tail changes each probability by $O(\varepsilon^2)$ and each determinant by $O(\varepsilon^3)$. For all sufficiently small ε, both averaged determinants are positive and all hole probabilities are $O(\varepsilon)$. Conditional on either environment, (60) still gives $L_E = 0$, while (57) remains nonpositive before and after averaging.

There is also an explicit finite certificate. For $\varepsilon = 10^{-3}$ and $\eta = 10^{-10}$, rational exponential and logarithm bounds give

$$\mathscr{D}^0 > 10^{-9}, \qquad L^0 > 10^{-9}, \qquad \max q^0_{i,j} < 131/100000.$$

Here $J_0(k) < 2$ and $J_H(k) \le 1$ for every $k \ge 1$. Each void changes by at most 5η, so each response changes by at most 10η. Product expansion bounds the determinant change by $40(131/100000)\eta + 200\eta^2 < 6 \times 10^{-12}$. The voids exceed $1/2$, so the corresponding bound for L is $160(131/100000)\eta + 800\eta^2 < 3 \times 10^{-11}$. Thus, uniformly over k,

$$\mathscr{D}(k) > 9 \times 10^{-10}, \qquad L(k) > 9 \times 10^{-10}, \qquad \max q_{i,j}(k) < .002. \tag{66}$$

The rational enclosure procedure is specified in Section 7.

For each fixed k, when $T \ge 12k$ the compact old grains disappear. Formula (63) becomes $\overline{Z}_T e^{-i\eta J_H}$, hence $L = 0$ exactly. Also

$$J_H(T) \sim \frac{H}{1+H}\frac{1}{T},$$

by comparison with the integral of $(a+1)^{-2}$; the weighted first old moment is $O(T^{-1})$ and the second is $O(T^{-2})$. Thus this law satisfies the pointwise sparse-residue contract. Its stationary void curve at normalized rate $a \ge 0$ is

$$v_{\mathrm{stat}}(a) = e^{-a\eta B}[(1-\varepsilon)e^{-12c\varepsilon a} + \varepsilon e^{-24a}], \tag{67}$$

independent of k. Its leading old-tail response is independent of k as well. Yet (66) can occur after any prescribed time by selecting a larger k. Stationary calibration and that leading tail coefficient do not give a uniform sign onset for this family. A single member does not remain positive at all late times.

6. Complete replication and simultaneous inference

6.1. The joint frequency event

For preparation j, a complete repeat gives

$$(X_{1,j,r}, X_{2,j,r}, X_{3,j,r}, X_{-,j,r}, X_{+,j,r}) \in \{0,1\}^5,$$

where $X = 1$ denotes covered or damaged and $X = 0$ denotes void or open. Across $r = 1, \dots, n_j$, these vectors are independent and identically distributed. Their internal dependence is unrestricted. The two preparation samples may also be dependent, provided each has the stated independent-repeat property.

For a class $\mathcal{A}$ of compatible marginal arrays, define $\mathfrak{P}_n(\mathcal{A})$ to consist of all joint laws of two size-n preparation samples whose group marginal laws are $Q_j^{\otimes n}$, where Q_j is a law on $\{0,1\}^5$ with void means $1 - E_{Q_j} X_{i,j} = v_{i,j}$ for one array $\mathbf{v} \in \mathcal{A}$. The two group samples may have any coupling. The shared-environment representation is imposed on $\mathbf{v}$; the admissible Q_j include every five-bit law with those means. In particular, independent Bernoulli arms and independent preparation samples belong

to this experiment. The minimax radius risk used above is $R_n(\mathcal{A}) = \inf_{\widehat{\varepsilon}} \sup_{P \in \mathfrak{P}_n(\mathcal{A})} E_P |\widehat{\varepsilon} - \varepsilon_*(\mathbf{v}(P))|$. A specified Poisson replay protocol gives a more structured observation law within the marginal framework.

For a mixed-law probability the environment is redrawn in each complete repeat. For a conditional-environment probability the same environment is retained while the controlled noise is independently redrawn. These estimate different probabilities. Separate spatial sites from a long-range field are not automatically independent repeats.

For each preparation measure the baseline void frequency $\widehat{v}_j$ and, for $a \in \{2, 3, -, +\}$, the two frequencies

$$\widehat{p}^{\pm}_{a,j} = \frac{1}{n_j} \sum_r \mathbf{1}\{X_{a,j,r} - X_{1,j,r} = \pm 1\}.$$

There are $2(1 + 4 \cdot 2) = 18$ Bernoulli means. Let

$$t = \log(36/\alpha), \qquad z_j = t/n_j, \qquad R(y, z) = \min\left\{1, \frac{5}{3}z + \sqrt{2yz + \frac{7}{3}z^2}\right\}. \tag{68}$$

Lemma 1 (simultaneous mean coverage).
With probability at least $1 - \alpha$, all eighteen true means differ from their empirical frequencies by at most their radius R.

Proof. If B is Bernoulli with mean μ, then for integers $k \geq 2$, $E|B - \mu|^k \leq \mu$ and $k! \geq 2 \cdot 3^{k-2}$. Hence, for $|\lambda| < 3$,

$$\log E e^{\lambda(B-\mu)} \leq \frac{\mu\lambda^2}{2(1 - |\lambda|/3)}.$$

Chernoff's argument for the empirical mean gives

$$P\{|\widehat{\mu} - \mu| > \sqrt{2\mu z} + 2z/3\} \leq 2e^{-t}.$$

Write $d = |\widehat{\mu} - \mu|$. If $d \geq 2z/3$, use $\mu \leq \widehat{\mu} + d$, square the inequality, and solve the quadratic to obtain $d \leq 5z/3 + \sqrt{2\widehat{\mu}z + 7z^2/3}$. If $d < 2z/3$ this bound also holds. Truncation at one is valid. A union bound over eighteen means proves the result. □

The radius follows directly from the Bernoulli Bernstein argument above; [4] supplies the broader variance-sensitive inference context.

6.2. Safe logarithms and the determinant interval

On the event in Lemma 1,

$$v_{1,j} \in V_j = [\max(0, \widehat{v}_j - b_j), \min(1, \widehat{v}_j + b_j)], \qquad b_j = R(\widehat{v}_j, z_j),$$

$$\delta_{a,j} \in I_{a,j} = [\widehat{\delta}_{a,j} - r_{a,j}, \widehat{\delta}_{a,j} + r_{a,j}],$$

where $\widehat{\delta} = \widehat{p}^{+} - \widehat{p}^{-}$ and $r_{a,j} = R(\widehat{p}^{+}_{a,j}, z_j) + R(\widehat{p}^{-}_{a,j}, z_j)$.

For $f(v, d) = \log(1 - d/v)$, call the rectangle $V_j \times I_{a,j}$ safe when

$$\inf V_j > 0, \qquad \inf V_j - \sup I_{a,j} > 0. \tag{69}$$

On a safe rectangle, $f_d = -1/(v - d) < 0$ and $f_v = d/[v(v - d)]$ has constant sign along any horizontal edge. The minimum and maximum therefore occur at the four corners, even for negative responses. Their extrema give $[\underline{\ell}_{a,j}, \overline{\ell}_{a,j}]$.

For the four main log responses, let $\widehat{\ell}_{i,j} = f(\widehat{v}_j, \widehat{\delta}_{i,j})$ and let $\eta_{i,j}$ be the maximum distance from this value to its interval endpoints. Then a valid determinant half-width is

$$\begin{aligned}\mathcal{R} = {} & |\widehat{\ell}_{3,2}|\eta_{2,1} + |\widehat{\ell}_{2,1}|\eta_{3,2} + |\widehat{\ell}_{3,1}|\eta_{2,2} + |\widehat{\ell}_{2,2}|\eta_{3,1} \\ & + \eta_{3,2}\eta_{2,1} + \eta_{3,1}\eta_{2,2}.\end{aligned} \tag{70}$$

Thus $\underline{L} = \widehat{L} - \mathcal{R}$ is a lower confidence endpoint. It follows by expanding both products, without assuming that their estimates are independent or unbiased.

Each safe probe pair gives

$$\overline{D}_{s,j} = \overline{\ell}_{-,j} + \overline{\ell}_{+,j}.$$

Under the null it is at least $D_{s,j} \geq 0$ on the same event. If some $\overline{D}_{s,j} < 0$, the curvature rejection is valid. Otherwise, when all rectangles needed for the determinant leg are safe, $K_j \leq K_j^+$ and $u_j \leq U_j$. Monotonicity of (44) gives

$$L \leq B_\epsilon(u, K) \leq B_\epsilon(U, K^+),$$

which excludes (47) on the coverage event. Both false rejections are therefore contained in its complement. This proves Theorem 18(c).

For the exact conditional null (60), both preparations must use the same fixed or repeatably conditioned environment e, with its calibration shared across preparations. Three-arm repeats then need only ten frequencies: two baselines and eight signed difference frequencies. Taking $t = \log(20/\alpha)$ and using (69)–(70), $\underline{L} > 0$ rejects that conditional class. Applying this zero allowance to an unknown environmental average is invalid.

6.3. What rare pairing pays for

Suppose the void probabilities are bounded below by a fixed positive constant and the actual paired differences satisfy

$$P(X_{a,j} \neq X_{1,j}) \leq P, \qquad a = 2, 3, -, +. \tag{71}$$

The arms can have much larger individual damage probabilities. Under monotone replay using a mother Poisson process, (71) is paid for by the response probability itself; mere closeness of marginal means does not imply it for an arbitrary coupling.

Let $n_1 = n_2 = n$ and $z = t/n$. Lemma 1 and the corner derivatives give log-endpoint errors

$$\eta = O\Big(\sqrt{Pz} + z + P\sqrt{z}\Big) = O\Big(\sqrt{Pz} + z\Big), \tag{72}$$

where the final bound uses $P \leq 1$ and sufficiently small z. The determinant error is $O(P\eta + \eta^2)$.

Proposition 4 (sufficient replication orders). Fix a void lower bound $v_* > 0$, a feasible step $s > 0$ and a finite variance-envelope constant C_K. There are constants $C < \infty$ and $P_0 > 0$, depending only on these fixed quantities, with the following property. Consider any fixed-design experiment satisfying Section 6.1, (71), $0 < P \leq P_0$, and $v_{a,j} \geq v_*$ for every arm. Let $t = \log(36/\alpha)$ and take n complete repeats per preparation.

(a) Suppose the main responses are ordered, $D_{s,j} \geq 0$, $\overline{K}_j(s) \leq C_K P^2$, $0 \leq \epsilon \leq 1$, and $\Delta = L - B_\epsilon(u, \overline{K}(s)) > 0$. If

$$n \geq C\frac{P^3}{\Delta^2}\log(36/\alpha), \tag{73}$$

then, on the simultaneous event of Section 6.2, the determinant certificate rejects unless the curvature leg has already rejected. In particular, its rejection probability is at least $1 - \alpha$.

(b) If $D_{s,j} = -g < 0$ for one preparation, then

$$n \geq C\frac{P}{g^2}\log(36/\alpha) \tag{74}$$

is sufficient for curvature rejection on that event. For this part the constant depends only on v_*, and no variance envelope is needed.

Proof. On the coverage event, a Bernoulli frequency whose true mean is at most P has empirical radius $O(\sqrt{Pz} + z)$: substitute $\widehat{p} \leq P + R(\widehat{p}, z)$ into (68) and solve its square-root bound. The baseline radius is $O(\sqrt{z} + z)$. The rare-pair contract gives $|\delta_{a,j}| \leq P$. Together with the void bound

this gives $u_j = O(P)$ and $|D_{s,j}| = O(P)$. Thus $\Delta = O(P^2)$ in part (a) and $g = O(P)$ in part (b). Either sample condition implies $z = O(P)$. The difference rectangles consequently have $d = O(P)$ throughout. For sufficiently small P_0, the common void bound makes every required log rectangle safe; the baseline-coordinate derivative is $O(P)$ and the difference-coordinate derivative is uniformly bounded. This proves (72) on the event with uniform constants.

For (a), the determinant error is $O(P\eta + \eta^2)$. The measured variance budget changes by $O(\eta/s^2)$, and the upper log endpoints change by $O(\eta)$. On $\overline{K}_j(s) \le C_K P^2$ the resulting allowance error is again $O(P\eta + \eta^2)$, uniformly over $0 \le \epsilon \le 1$. Since $0 < \Delta \le L = O(P^2)$, (73) with a sufficiently large constant pays both the square-root and linear terms in (72), and makes their total contribution smaller than Δ. It also ensures safe rectangles by taking P_0 small enough. For (b), $g = O(P)$ and (74) makes the two probe-log errors smaller than g. □

These are uniform sufficient orders under specified envelopes. In the balanced regime $u_j \asymp P$, a cubic determinant margin additionally requires the calibration precision described in Proposition 1. Substitution of an estimated margin alone does not establish these sample conditions. The linear certificate in Theorem 17 has its own explicit finite radii (40); Proposition 4 concerns the two variance-test legs. Moving the probe step toward zero, adaptive stopping, or using correlated spatial observations as repeats changes the sampling contract.

7. Reproducible analytic and synthetic checks

The package contains deterministic analytic calculations, fixed synthetic count fixtures and implementations of the confidence formulas. Section 6 proves the probability guarantees for complete repeats.

The scalar sign calculations use rational interval endpoints. For $x \ge 0$, a finite positive Taylor sum for e^x has remainder at most the next term divided by $1 - x/(m+2)$ when $x < m + 2$. Taking reciprocals encloses e^{-x}. For $0 < v \le 1$, the series

$$-\log v = \sum_{r=1}^{m} \frac{(1-v)^r}{r} + \mathcal{E}, \qquad 0 \le \mathcal{E} \le \frac{(1-v)^{m+1}}{(m+1)v}$$

gives log enclosures. Each rational operation is rounded outward to a fixed rational grid. These bounds certify (66) and (86).

The general-step count calculator implements (68)–(70) and the two variance-test rejection legs. It accepts two sets of five-bit outcome counts, α, a fixed probe step and a valid ratio-error bound. Bits are ordered $(1, 2, 3, -, +)$. The calculator uses 80-digit Decimal arithmetic with numerical guards. Generic decisions near a numerical boundary require separate interval certification; the rational scripts certify the specified scalar and count decisions.

The local-boundary module reconstructs the four environments, checks their three anchor moments, evaluates the radius comparison at decreasing dispersion, checks the common-block variance inequality on a fixed-seed set of one thousand configurations, and evaluates the probability distance and Bernoulli divergence of the radius-comparison pair. These are deterministic construction diagnostics using 80-digit Decimal arithmetic; the analytic bounds in Section 3.4 establish the uniform claims.

The radius-upper module checks both moment-repair branches on 160 fixed-seed pairs, target moments and exposure support, the calibration-marginal kernel identity, the response-opening triangle, and the positive radius lower bound in the mass-one discontinuity example.

The unknown-mass/mean module checks twelve finite shared-environment fixtures with unequal mass and mean changes. It verifies thinning and activation, the ambient calibration law and updated overlap budget, affine support/cone preservation, and both higher-moment repair branches against the target moments.

The critical-mass module checks four shrinking-mass realizations and hard pairs in the subcritical, critical and consistent regimes. It verifies exact raw distances, uniform Taylor remainders,

deterministic-branch margins, the response triangle's linear private-mass scale and the product-Bernoulli divergence. These are deterministic finite-construction checks; Section 3.7 proves the uniform modulus and matching risk bounds. The projection estimators and confidence inversions remain variational constructions.

The two-scale module verifies positive-exposure symmetric anchor laws with dispersion exceeding the private-mass scale, their joint endpoint allocations, and both lower-bound pairs including the zero-heterogeneity edge. It uses Decimal construction diagnostics. The uniform two-sided opening, modulus and minimax conclusions are proved in Section 3.8; projection and confidence inversion are variational procedures.

The half-step modules compute the population profiles, the polynomial certificate, the deterministic-anchor nonzero-error boundary and minimum radius, and the stable overlap count certificate. Three hundred constructed common environments check the closed-form boundary across unequal masses and exposure ratios, with compatible common-law controls. Theorem 15 supplies the general center-cover search as a finite construction; the implemented programs cover the listed closed-form and scalar certificates.

The adaptation module checks the product-Bernoulli likelihood second-moment identity, its fourth-order scaling, logarithmic alternative budgets and the absolute-risk transfer bound on 24 finite comparison fixtures. These Decimal diagnostics accompany the analytic proofs in Section 3.9; they neither optimize the variational estimators nor simulate coverage.

The profile module computes the linear certificate in Theorem 17. For the heterogeneous array (41), exact rational root comparisons enclose the Gaussian nodes, fractional powers and Hermite coefficients. Interval Gaussian elimination retains the coefficient uncertainty in the linear functions, and outward rational square-root bounds propagate the confidence radii. The rational choice $t = 6.58 > \log(720)$ is certified by a positive Taylor partial sum for $e^{6.58}$. It supplies the displayed finite-sample lower endpoints at $\alpha = 0.05$. The profile roots in the table are high-precision numerical summaries; the positive population and count margins are certified independently of those roots.

Synthetic checks cover a deterministic-coverage null; the cross-preparation array (56), including its separate exact exponential curves and loss of a positive allowance margin at $\epsilon = 1/8$; negative signed responses; unsafe logarithms; a safe curvature rejection when the main determinant is unsafe; the heterogeneous shared-calibration separation and a compatible common-ratio control; the fixed Cox counterexample; finite-mixture reconstruction; and exact enumeration of the cooperative finite cone. Fixed seeds and all constructions are included in the check script. The replication unit remains a complete repeat, independently of how many coordinates are recorded within it.

8. Discussion

Separate heterogeneous coverage fits do not guarantee a shared explanation. The absolute weaker-source void probabilities retain a common environmental mass budget that disappears under separate normalization. Combining that budget with exposure moments and the calibration band gives the exact joint boundary for the closure of the finite-label probability representation.

The local theory also quantifies what exposure heterogeneity can change. On the anchor family of Theorems 4–5, matching upper bounds and four-environment realizations determine the first-order support correction $\sigma\sqrt{\kappa}G_\epsilon$. The full array has a minimum calibration radius stable to second order around its realizing radius. The deterministic-anchor comparison differs at first order with the explicit coefficient in (22k). These results retain all three anchor moments and optimize over the unknown central ratio.

The minimum calibration radius makes the gain concrete. With $r = 1/2$, $\beta = 3/4$ and normalized intermediate responses $3/5, 4/5$, separate fits within one calibration band need radius about 0.10376, while joint compatibility needs about 0.33048. The stable overlap inequality extends the distinction to heterogeneous anchors. Its strict example remains outside the joint model even when separate

profiles and the variance-adjusted determinant are feasible.

Finite-sample certificates preserve these restrictions on one eighteen-frequency event. The $\delta^{-2}\log(1/\alpha)$ replication order and its matching lower bound concern the entire testing class; the lower-bound construction uses ordinary shape violations. They calibrate the sampling cost of compatibility testing. Corollary 5.1 addresses radius estimation: a quadratic change in anchor probabilities can produce a linear change in the required radius, yielding a local $n^{-1/4}$ risk lower bound. This lower bound uses two compatible arrays rather than a shape violation. Theorems 6–7 establish the matching upper rate on the observable local moment class (22o), using a calibration-preserving moment repair and verified response-opening margin. The estimator and confidence interval are defined by compact variational problems. Theorems 8–9 release the fixed nuisance values: the same square-root exponent and matching statistical orders hold when each preparation's mass and first anchor vary over the compact ranges in (22u). The mass and mean adjustments preserve a common environment before the higher-moment repair.

Theorems 10–11 quantify the shrinking-margin case on the stated frontier tube: the factor $1/p_n$ in the sharp modulus leads to risk $\min\{1, n^{-1/4}/p_n\}$ and the critical quantity np_n^4. A scale-adaptive construction attains the same estimation and fixed-confidence expected-diameter orders without knowing p_n. The mass-one counterexample identifies a boundary where the radius map loses continuity.

Theorem 12 separates the mass scale from the exposure standard-deviation envelope. The condition $h^2 \le bp$ permits dispersions much larger than p, while preserving private-block interior mass. Its matching risk contains a regular $n^{-1/2}/p$ term and a capped heterogeneity term $\min(h, n^{-1/4})/p$. Uniform consistency therefore requires both $np^2 \to \infty$ and $\min(h/p, n^{-1/4}/p) \to 0$. The former np^4 boundary is recovered whenever h/p stays bounded below; smaller heterogeneity can permit consistency beyond that boundary.

Unknown-envelope inference changes these oracle guarantees. Theorem 13 proves that full-union honesty forces fourth-root expected diameter even on deterministic arrays. Theorem 14 shows that point estimation can uniformly retain the deterministic regular rate, with an optimal logarithmic price over the heterogeneity envelopes. Its thresholded projection uses only the observed raw frequencies; both lower bounds use compatible pairs in the same finite-label model.

The particle-process examples explain why these observation contracts matter. A frozen Cox environment can change averaged determinant signs, and the cooperative recovery model in Appendix E yields a source-probe curvature exclusion. A specified common space-time grain process adds restrictions beyond the finite-label representation. A rejection of the representation excludes every grain law satisfying its observation contract; a feasible array must also satisfy the additional process constraints to admit that physical interpretation.

For an application, the essential inputs are the absolute source-void probabilities, the certified affine probe relation and the pointwise calibration-error budget. Independently repeated complete experiments estimate the required binary means. These inputs connect the joint feasibility question to moment problems, partial identification, heterogeneous mixtures and finite-sample tests.

Appendix A. Alternative budgets and sharpness

A.1. Bounded residual support

If Y_j has essential support width at most w_j, every exponential tilt has variance at most $w_j^2/4$. To see this, write its support as $[m, m+w]$ and use $E[(Y-m)(m+w-Y)] \ge 0$ to obtain $\operatorname{Var} Y \le (EY-m)(m+w-EY) \le w^2/4$. Repeating the chord-gap proof gives

$$0 \le xh_j(1) - h_j(x) \le x(1-x)w_j^2/8 \le w_j^2/32.$$

Since $m \le u_j \le m + w_j$, $\sup Y_j \le u_j + w_j$, so (51) can instead be bounded by $\epsilon(u_j + w_j)$. Thus

$$L \le \frac{u_2 w_1^2}{32} + \epsilon(2u_1u_2 + u_2w_1 + u_1w_2). \tag{75}$$

A support-width bound is not an empirical ordinary variance estimate.

Both 1/32 in (75) and 1/8 in (45) are asymptotically sharp within actual frozen grain processes. For $w > 0$, take a low environment with lifetime-five grains at rate $w/4$. In a high environment use lifetime-two grains at rate $2w/3$ and lifetime-five grains at rate $w/12$. Use $T = 1$, preparation lengths one and four, old multipliers $1, 2, 3$, and target multiplier one. Add an environment-independent stream of rate $\eta = w^2$ with survival (62). Choose

$$P(E = \text{low}) = (1 + e^w)^{-1}, \quad P(E = \text{high}) = e^w/(1 + e^w).$$

In the short preparation, total reference void exponents differ by w; the reference-void weights are exactly equal. Hence Y_1 under Q_1 has two equally likely values of width w, and $K_1 = w^2/4$. In the long preparation $Y_2 = b = 2w + 2\eta J_2$ is constant, where $J_2 = \sum_{a=1}^{4} S(a)$. With $x = 1/2$,

$$L = b\left[\tfrac{1}{2}\log\cosh(w/2) - \log\cosh(w/4)\right], \qquad u_2 = b.$$

Expansion at zero gives $L/(u_2w^2/32) \to 1$ and $L/(u_2e^{u_1}K_1/8) \to 1$. The same example realizes the two-point sharpness in (54), and its nonzero lifetime tail persists for every $w > 0$.

A.2. Exposure domination with known lifetime kernels

Let $C_j = -\log Z_j + \phi_1 N_j^*$ be the total reference covering mean; put $C_j = \infty$ when $Z_j = 0$. If a valid geometric bound gives $Y_j \le r_j C_j$ wherever $C_j < \infty$, then

$$E_{Q_j} Y_j^2 = \frac{E[Y_j^2 e^{-C_j}]}{v_{1,j}} \le \frac{4e^{-2} r_j^2}{v_{1,j}}, \tag{76}$$

using $\sup_{c \ge 0} c^2 e^{-c} = 4e^{-2}$. Zero-weight environments contribute zero. No original-environment rate moment is required.

For finitely many fixed known lifetime survival kernels $S_k(a)$ and arbitrary correlated, finite almost surely environmental rates β_k, set

$$J_{0,k} = \sum_{a=0}^{T-1} S_k(a), \qquad J_{j,k} = \sum_{a=T}^{T+N_j-1} S_k(a).$$

For old multipliers $1, 2, 3$ and target one,

$$C_j = \sum_k \beta_k (J_{0,k} + J_{j,k}), \quad Y_j = 2\sum_k \beta_k J_{j,k}, \quad r_j = \max_k \frac{2J_{j,k}}{J_{0,k} + J_{j,k}}. \tag{77}$$

Assume $S_k(0) = 1$, so the denominators are positive for $T \ge 1$. A safe lower confidence bound for $v_{1,j}$ can be used in (76) on the same coverage event. Unknown kernels require an independent geometric certification.

If the stationary reference void is uniformly at least $v_0 > 0$, then finite-reference voids are at least v_0, and Cauchy–Schwarz gives $u_j \le E_{Q_j} Y_j \le 2e^{-1} r_j/\sqrt{v_0}$. At $\epsilon = 0$, (45) implies the uniform bound

$$L \le \frac{e^{-3} r_1^2 r_2}{v_0^{3/2}} \exp\left(\frac{2e^{-1} r_1}{\sqrt{v_0}}\right). \tag{78}$$

For fixed kernels with $r_j = O(A(T))$ this is $O(A(T)^3)$. This stronger kernel contract excludes the moving compact-lifetime family used for nonuniformity in Section 5.3.

For probes satisfying (6), $\phi_1 \ge s\Delta_3$ implies $Y_j \le C_j/s$. Thus (76) already gives a coarse finite budget for that subclass. The two measured probes retain the finer variance information and remove the squared mean.

Appendix B. A one-sided observation obstruction

The wider representation experiment in this appendix observes arbitrary nonnegative finite Y under a distribution Q, with no tail budget. Each environment draw permits finitely many binary void queries with conditional means e^{-tY}, $t \geq 0$. Queries can be paired or sequentially chosen. The total number of independently redrawn environments is at most a fixed m. No weaker-forcing query corresponding to a negative Laplace parameter is available.

Proposition 5. If a possibly randomized upper confidence bound $U \in [0, \infty]$ satisfies

$$P_Q\{\mathrm{Var}_Q\, Y \leq U\} \geq 1 - \alpha$$

for every finite-variance law in this experiment, then at $Q_0 = \delta_0$,

$$P_{Q_0}\{U = \infty\} \geq 1 - \alpha. \tag{79}$$

Proof. Suppose $P_{Q_0}(U < \infty) > \alpha$. Choose a finite c with $P_{Q_0}(U \leq c) > \alpha$. Set $Q_M = (1-\delta)\delta_0 + \delta\delta_M$. With probability at least $(1-\delta)^m$ all environment draws equal zero and the complete query data have exactly their Q_0 distribution. Choose $\delta > 0$ small enough that $(1-\delta)^m P_{Q_0}(U \leq c) > \alpha$ and then choose M so large that $\delta(1-\delta)M^2 > c$. On that identical-data event coverage fails, contradicting the assumed bound. □

This is an elementary rare-contamination obstruction for the stated query experiment, within the general nonparametric impossibility setting discussed in [8]. Known-rate exposure control in Appendix A and the weaker-rate intervention in (6) provide the additional information supporting finite budgets.

Appendix C. A self-contained finite-count witness

For the array in Section 4.3, the exact minorant p is defined by value and derivative interpolation of $t^{17/10}$ at $a, b = 1/2 \mp 2\sqrt{6}/15$. The exact majorant q interpolates $t^{13/10}$ at $0, 1$ and twice at $1/2$. The following displayed coefficient centers enclose the corresponding exact coefficients with error at most 5×10^{-13} per entry; $q_0 = 0$ exactly.

Coefficient	Minorant p	Majorant q
0	−0.006914452375	0
1	0.162611267906	0.512648562186
2	1.013345008781	0.711063898866
3	−0.170765293415	−0.223712461052

Each root enclosure is obtained by integer comparisons: for rational $y \geq 0$, choose consecutive grid points $k/10^{70}, (k+1)/10^{70}$ bracketing the positive root through their exact integer powers. Applying this to the Gaussian square root and fractional powers, then solving the four Hermite equations by outward interval elimination, gives coefficient enclosures contained in the displayed cells. Bounds on the exact polynomials follow from interpolation; the rounded centers need not retain their contact properties.

The complete five-bit integer counts, with bits ordered $(1, 2, 3, -, +)$, are:

Outcome	Preparation 1	Preparation 2
11111	200000000000	200000000000
11101	400000000000	400000000000
01101	94212386486	39123810152
00101	20454280180	75542856514

Outcome	Preparation 1	Preparation 2
00100	57333333333	57333333333
00000	228000000001	228000000001

All unlisted outcomes have zero count, and each column sums to 10^{12}. Use $\alpha = 1/20$, $s = 1/2$, $x_0 = 1/4$, $\epsilon = 1/10$ and the conservative rational value $t = 329/50 > \log(720)$. The last inequality follows by summing the positive Taylor series of $e^{329/50}$ through degree 39; its partial sum already exceeds 720. Compute the radii in (68) by outward rational square-root bounds, and propagate the displayed coefficient cells through (39)–(40). The resulting exact rational lower endpoints exceed 0.0039293 and 0.0039310. This calculation includes the coefficient uncertainty and establishes the stated finite-count rejection without a floating-point boundary decision.

Appendix D. Finite counts for the nonzero-error overlap exclusion

For the heterogeneous example in Section 3.3, the five-bit counts below use the same order $(1, 2, 3, -, +)$ and 10^{12} complete repeats per preparation.

Outcome	Preparation 1	Preparation 2
11111	100000000000	100000000000
11101	450000000000	450000000000
01101	182358160847	71572961889
00101	42551839152	153337038110
00100	112455000000	112455000000
00000	112635000001	112635000001

All other counts vanish. Use $r = 1/2$, $s = 1/2$, $\epsilon = 1/8$, $\alpha = 1/20$ and the certified rational $t = 6.58$ from Appendix C. Propagate the eighteen-frequency radii into the weak-void intervals, quadratic-deviation upper endpoints and absolute intermediate-difference lower endpoint as specified after (21). Fraction endpoints and outward square-root bounds give a finite overlap-margin lower bound above 0.04125. The same counts do not reject through the older variance-adjusted determinant.

A second, entirely linear certificate uses (22) at $r = 1/2$, $\lambda = 150$. In the five-coordinate order of (7) its coefficient vectors are $d_1 = (-37, 150, -150, 0, -1)$ and $d_2 = (-301/8, 150, -150, 0, 1)$. The certified rational support upper bound is $81/200$. The rational radii from (26) give a count-margin lower bound above 0.01540. This single direction is valid over the whole band $[0, 1]$, so it is also a finite covering certificate for any allowed central ratio at error $1/8$.

The population intermediate probabilities are algebraic: $\beta ER^{7/4}$ and $\beta ER^{5/4}$ for equal weights on $49/100, 51/100$. Integer fourth-root enclosures certify the population margins in (21); rational logarithmic remainders separately certify the old allowance's positive slack. The counts are a specified deterministic fixture. The certificate verifies the finite decision under the stated replication contract.

Appendix E. Cooperative recovery as an application

Random-grain models with long-range dependence and nonlinear readouts provide a related probabilistic setting [6]. Here the cooperative interval process illustrates the source-probe statistics.

We evaluate the new statistics in the interval-disaster model shared with [9]. Its finite-cone reduction [9, Proposition 8.1] applies to this protocol. The calculation below is a self-contained short-protocol illustration of the inference rules. Let the hole set evolve on $\mathbb{Z}$ by

$$H_{n+1} = (H_n \cap (H_n + 1)) \cup D_{n+1}, \qquad H_0 = \varnothing. \tag{80}$$

At each integer origin, a Poisson event of length $r \geq 2$ and rate $a\nu_r$ places r consecutive holes in the current row. Rows and origin-length streams are independent. Assume

$$C = \sum_{r\geq 2} \nu_r > 0, \qquad A = \sum_{r\geq 2} r\nu_r < \infty. \tag{81}$$

All intervals have length at least two, so $A \geq 2C$. This is a noisy one-dimensional AND eroder, related to the classical dynamics considered in [5]. Its interval births may be unbounded and correlated across sites.

Prepare for $j = 1$ or $j = 2$ old updates at rate a, then make one target update at any fixed rate $b \geq 0$. Let $v_{a,j}(b)$ denote the terminal open probability.

Proposition 6 (exact short-protocol evaluation).
For the model (80)–(81),

$$\begin{aligned} v_{a,1}(b) &= e^{-bA} g_1(a), \quad g_1(a) = e^{-aA}(2 - e^{-aC}), \\ v_{a,2}(b) &= e^{-bA} g_2(a), \quad g_2(a) = e^{-2aA}(4 - 3e^{-aC}). \end{aligned} \tag{82}$$

For old rates $a = \rho, 2\rho, 3\rho$, define L by (3). Then, as $\rho \downarrow 0$,

$$L = 2C^2(4A - 3C)\rho^3 + O(\rho^4) > 0 \tag{83}$$

for all sufficiently small $\rho > 0$. For preparation one and every feasible $0 < s \leq 1/2$, use $a_\pm = (1 \pm 2s)\rho$. Its source-affine probe curvature satisfies

$$D^{\mathrm{coop}}_{s,1} = \log \frac{g_1((1-2s)\rho) g_1((1+2s)\rho)}{g_1(\rho)^2} < 0. \tag{84}$$

These conclusions include every nonzero law $\nu_r \sim r^{-p}\mathcal{L}(r)$ with $p > 3$ and eventually positive slowly varying $\mathcal{L}$. No long-time limit is used.

Proof. Write O_x for the event that site x is open in one disaster row. Poisson avoidance gives

$$P(O_x) = e^{-aA}, \quad P(O_x \cap O_{x+1}) = e^{-a(A+C)}, \quad P(O_{x-1} \cap O_x \cap O_{x+1}) = e^{-a(A+2C)}. \tag{85}$$

The number of origins of a length-r interval hitting k consecutive sites is $r + k - 1$.

After one old row the hole set is exactly that row. At the next deterministic erosion step the terminal site is open if either of its two ancestors is open. The independent target row is open at that site with probability e^{-bA}. Inclusion-exclusion proves the first formula in (82).

For two old rows, write A_x, B_x for the hole indicators in the first and second rows. At an intermediate site,

$$H_{2,x} = B_x \vee (A_x A_{x-1}).$$

Let U_x be its open event. Independence across rows gives

$$P(U_x) = 2e^{-2aA} - e^{-a(2A+C)}.$$

Both U_x and U_{x-1} require the second row open at two adjacent sites. The first-row requirement is $O_{x-1} \cup (O_x \cap O_{x-2})$. For intervals of length at least two, the incidence measures for the pair $\{x-2, x\}$ and the triple $\{x-2, x-1, x\}$ coincide: any interval hitting the middle site hits an endpoint. Hence the first-row union has probability

$$e^{-aA} + e^{-a(A+2C)} - e^{-a(A+2C)} = e^{-aA}.$$

It follows that $P(U_x \cap U_{x-1}) = e^{-a(2A+C)}$. The final erosion step is open with probability $2P(U_x) - P(U_x \cap U_{x-1})$, giving g_2.

Let $f_j = \log g_j$. Taylor expansion of these explicit analytic functions at zero gives

$$f_1(a) = -(A - C)a - C^2a^2 + O(a^3), \qquad f_2(a) = -(2A - 3C)a - 6C^2a^2 + O(a^3).$$

The target factor cancels from every log ratio. Substituting $\ell_{i,j} = f_j(i\rho) - f_j(\rho)$ into (3) gives (83); its coefficient is positive because $A \geq 2C > 0$.

Finally,

$$f_1''(a) = -\frac{2C^2e^{-aC}}{(2 - e^{-aC})^2} < 0, \qquad a \geq 0.$$

Strict concavity proves (84), including $a_- = 0$ at $s = 1/2$. The zero-rate case is an ordinary finite update, so it needs no zero-intensity stationary phase or limiting argument. □

For a concrete law take $\nu_r = \mathbf{1}_{\{r=2\}} + 10^{-9}r^{-4}$, $r \geq 2$. This has a genuine nonzero unbounded tail. At $\rho = .1$,

$$2 \leq A \leq 2 + 10^{-9}/4, \qquad 1 \leq C \leq 1 + 5 \cdot 10^{-9}/48.$$

The bounds follow from decreasing-series integral comparisons after the first term. Substitution in (82), with outward rational exponential and log bounds, gives

$$.002675298 < L < .002675300. \tag{86}$$

For the pure dimer law the tighter enclosure is $.002675298677 < L < .002675298678$. Equation (84) supplies a strict negative probe curvature for the heavy-tailed law itself, for every fixed permitted step.

The positive L rejects the exactly conditioned coverage class (60). The two-moment reduction here is the $L \leq 3$ case of the finite-cone framework in [9]; the new evaluations concern the present determinant and probe statistic. By Section 5.3 it does not alone reject unrestricted Cox mixtures. Equation (84) rejects the affine-probe contract even under arbitrary allowed mixing. For a three-arm residual-mixture test one must compare L with a separately certified allowance (45), or add the two probes and apply Theorem 18.

Statements and declarations

Competing interests. The author is employed by a public fund management institution in China. This research was conducted independently and does not represent the views of the author's employer, any fund, or any client. The study did not use non-public employer or client information and does not constitute investment advice.

Data and code availability. All results are mathematical. The accompanying reproducibility directory contains deterministic analytic certificates, synthetic count fixtures, the finite count calculator, and the half-step profile and polynomial-certificate modules. The proofs and explicit probability arrays are contained in this manuscript. No empirical dataset or original-field simulation is used.

Declaration of generative AI and AI-assisted technologies in the writing process. During the preparation of this work, the author used OpenAI GPT-6.1 Sol and xAI Grok 4.7 to assist with literature synthesis, development of argumentative and theoretical structure, expansion of the author's propositions into manuscript drafts, mathematical and code checks, and language revision. The author takes responsibility for the content. No AI system is listed as an author. This declaration records the tools used and does not assert that a human has independently certified every proof or line of code.

References

[1] Last, G. and Penrose, M. (2017). *Lectures on the Poisson Process*. Cambridge University Press. https://doi.org/10.1017/9781316104477. Author manuscript, 21 August 2017: https://stoch.math.kit.edu/img/Last/lastpenrose2017.pdf.

[2] Balabdaoui, F., Durot, C. and Jankowski, H. (2019). Least squares estimation in the monotone single index model. *Bernoulli* **25**(4B), 3276–3310. https://doi.org/10.3150/18-BEJ1090. Author version 2.

[3] Wainwright, M. J. and Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* **1**(1–2), 1–305. https://doi.org/10.1561/2200000001.

[4] Maurer, A. and Pontil, M. (2009). Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory*. Official conference paper. Author version 1.

[5] Rakovszky, T., Gopalakrishnan, S. and von Keyserlingk, C. (2024). Defining stable phases of open quantum systems. *Physical Review X* **14**(4), 041031. https://doi.org/10.1103/PhysRevX.14.041031.

[6] Surgailis, D. (2024). Scaling limits of nonlinear functions of random grain model, with application to Burgers' equation. *Stochastic Processes and their Applications* **174**, 104390. https://doi.org/10.1016/j.spa.2024.104390. Author version 2.

[7] Balan, T. A. and Putter, H. (2019). Nonproportional hazards and unobserved heterogeneity in clustered survival data: When can we tell the difference? *Statistics in Medicine* **38**(18), 3405–3420. https://doi.org/10.1002/sim.8171.

[8] Bertanha, M. and Moreira, M. J. (2020). Impossible inference in econometrics: Theory and applications. *Journal of Econometrics* **218**(2), 247–270. https://doi.org/10.1016/j.jeconom.2020.04.016. Author version 5.

[9] Huang, M. (2026). *Finite-preparation response under heavy-tailed catastrophes: exact amplitudes and separation from Poisson coverage*. Zenodo preprint, version 1, published 3 November 2026. Version DOI.

[10] Heckman, J. J., Robb, R. and Walker, J. R. (1990). Testing the mixture of exponentials hypothesis and estimating the mixing distribution by the method of moments. *Journal of the American Statistical Association* **85**(410), 582–589. DOI. The historical summary in Section 1 uses the publisher abstract, checked 11 November 2026. The full article was unavailable for this revision; the detailed moment-condition and test-statistic comparison uses the open text of [11].

[11] Robb, R., Frydman, H. and Robertson, A. (2017). Heterogeneity versus duration dependence with competing risks: an application to the labor market. *Applied Stochastic Models in Business and Industry* **33**(5), 465–475. DOI and open text.

[12] Liao, J. G. and Berg, A. (2019). Sharpening Jensen's inequality. *The American Statistician* **73**(3), 278–281. DOI. Author-hosted text.

[13] Gautschi, W. (2002). The interplay between classical analysis and (numerical) linear algebra—a tribute to Gene H. Golub. *Electronic Transactions on Numerical Analysis* **13**, 119–147. Journal text. Author text.

[14] Bertsimas, D. and Popescu, I. (2005). Optimal inequalities in probability theory: a convex optimization approach. *SIAM Journal on Optimization* **15**(3), 780–804. DOI. Author-hosted text.

[15] Chen, M., Gao, C. and Ren, Z. (2016). A general decision theory for Huber's ϵ-contamination model. *Electronic Journal of Statistics* **10**(2), 3752–3774. DOI. Author version 2, revised 16 January 2017.

[16] Palomba, F. (2026). *Sensitivity Analysis in Population Shares.* Working paper, version 3 November 2026, accessed 4 November 2026. Author manuscript.

[17] Huang, M. (2026). *An exact moment criterion for survival under renewed spatial catastrophes.* Zenodo preprint, version 1, published 3 November 2026. Version DOI.

[18] Kitagawa, J. and Pass, B. (2015). The multi-marginal optimal partial transport problem. *Forum of Mathematics, Sigma* **3**, e17. DOI and open text.

[19] Lam, H. (2016). Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research* **41**(4), 1248–1275. DOI. Author manuscript.

[20] Donoho, D. L. (1994). Statistical estimation and optimal recovery. *The Annals of Statistics* **22**(1), 238–270. DOI. Author-hosted scan.

[21] Le, D., Pham, T., Rinaldo, A. and Ho, N. (2026). *Weak limits of parameter estimation in strongly identifiable mixture models.* arXiv preprint 2610.25857, version 1, 19 October 2026. Versioned author text. Preprint DOI.

[22] Wu, Y. and Yang, P. (2020). Optimal estimation of Gaussian mixtures via denoised method of moments. *The Annals of Statistics* **48**(4), 1981–2007. DOI. Author version 2, 13 April 2019.

[23] Cai, T. T. and Low, M. G. (2004). An adaptation theory for nonparametric confidence intervals. *The Annals of Statistics* **32**(5), 1805–1840. DOI. Author reprint, arXiv version 1.

[24] Duchi, J. C. and Ruan, F. (2021). A constrained risk inequality for general losses. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, *Proceedings of Machine Learning Research* **130**, 802–810. Published version. The comparison here uses the author preprint arXiv:1804.08116, version 1, 22 April 2018. Versioned author text. Preprint DOI.