Testing and inverse stability at piecewise log-affine and uniform laws under Gaussian noise

Scott Schubr
Luminary Oddities

November 2026

Abstract

Piecewise log-affine laws, which include the exponential, Laplace and uniform distributions, are the extremizers of the log-concave moment problem, and inverse problems at them have rates different from smooth references. For a fixed standardized piecewise log-affine reference with p pieces and e finite ends, and any log-concave competitor, a bounded Lipschitz discrepancy δ gives the sharp orders δ in W_1, δ or $\sqrt{\delta}$ in total variation depending on whether there is a finite end, and explicit sharp orders in W_r for $r > 1$; Hellinger distance and posterior regret have the same inverse orders. Under Gaussian noise of standard deviation σ, the sharp total variation moduli become $\max(d^2, \sigma^{3/4} d^{1/2})$ at a finite endpoint, and the same rates hold for the uniform ball and the uniform cube in every fixed dimension, where total variation is controlled by exterior mass. We also compute the exact Gaussian covariance information carried by the faces and edges of a convex body.

Contents

1 Introduction **2**

2 Sharp inverse stability at piecewise log-affine reference laws **3**
2.1 Statement of the results . . . 3
2.2 Finite measurements and comparison on affine pieces . . . 4
2.3 Total variation and the linear transport bound . . . 5
2.4 Higher Wasserstein distances . . . 5
2.4.1 The finite ends . . . 5
2.4.2 The infinite ends . . . 6
2.5 Sharpness . . . 7
2.5.1 Alternatives near a finite endpoint . . . 7
2.5.2 Whole-line alternatives for total variation and W_1 . . . 8
2.5.3 Alternatives at an infinite endpoint . . . 8
2.6 Fixed-reference tests with Gaussian errors . . . 9

3 Gaussian posterior regret and testing at piecewise log-affine laws **10**
3.1 Results . . . 10
3.2 Transfer from observed survival functions . . . 11
3.3 Transport bounds from the transferred discrepancy . . . 12
3.3.1 Hellinger bounds . . . 12
3.4 Sharp lower bounds for Hellinger distance and regret . . . 13

3.5 Tests based on observed tail flux . . . 14
3.6 Testing lower bounds . . . 15
3.6.1 Many tail scales at the quadratic threshold . . . 15

4 Vanishing Gaussian noise in tests at piecewise log-affine laws 17
4.1 Statement of the results . . . 17
4.2 Uniform measurements and the effect of one affine piece . . . 18
4.3 A measurement near the endpoint . . . 19
4.4 Inverse and testing upper bounds . . . 20
4.5 Lower bounds uniform in the noise level . . . 20
4.5.1 Smoothing a finite endpoint . . . 20
4.5.2 A regular fixed-support alternative . . . 21

5 Testing the uniform ball under Gaussian noise 22
5.1 Results . . . 22
5.2 A weighted convexity estimate . . . 23
5.3 Escape mass and bounded measurements . . . 24
5.4 A spherical boundary measurement . . . 25
5.5 Inverse and testing upper bounds . . . 26
5.6 A boundary lower bound for convex bodies . . . 27

6 Exterior mass and testing for the isotropic uniform cube 28
6.1 Results . . . 29
6.2 Rigidity at zero exterior mass . . . 29
6.3 A linear exterior mass estimate . . . 30
6.4 Coordinate tests and noisy inverse bounds . . . 31
6.4.1 A directional lower bound . . . 31

7 Gaussian covariance information at faces and edges 32
7.1 Results . . . 32
7.2 Concentration of the face contributions . . . 34
7.3 Integration of the local face profiles . . . 34
7.4 Localization at edges when face variances agree . . . 35
7.5 The integrated edge limits . . . 36
7.6 Product experiments . . . 37

1 Introduction

Piecewise log-affine laws with p pieces and e finite ends are exactly the extremizers of the log-concave signed moment problem when $2p + e$ moments are prescribed [2]. This is why moment discrepancy through order $2p + e$ is the natural measurement scale at these references, and the stability exponents below are governed by the same count. Section 2 proves the sharp inverse stability bounds at piecewise log-affine references, and Sections 3–4 transfer them to Bayes regret, Hellinger testing and vanishing Gaussian noise. Sections 5–6 treat the uniform ball and cube in fixed dimension. Section 7 computes the Hellinger and Fisher information that boundary faces and edges carry about the noise covariance. The Gaussian deconvolution background is in [1].

2 Sharp inverse stability at piecewise log-affine reference laws

We study recovery of a one-dimensional log-concave law near a fixed piecewise log-affine reference from weak discrepancies after additive noise. Finitely many bounded Lipschitz observables give a linear inverse bound in W_1. The optimal orders in total variation and higher Wasserstein distances depend on whether the reference has finite or infinite support ends. The rates hold for any fixed known additive noise law. The proof combines sign changes of differences of densities with a positive, band-limited weight. Under Gaussian noise the same observables give sharp fixed-reference tests in total variation and W_1.

2.1 Statement of the results

Let g be a fixed real log-concave density of mean zero and variance one with continuous piecewise affine log density on the interior of its support. Write p for the number of maximal affine pieces and $e \in \{0, 1, 2\}$ for the number of finite support endpoints. Put $s = 2p + e$. The alternatives f are all log-concave densities of mean zero and variance one. Values of densities at their support endpoints are irrelevant. For a fixed probability error law η, write $Kf = f * \eta$. The total variation distance is $\mathrm{TV}(f, g) = \frac{1}{2}\int |f - g|$. The notation W_r refers to the usual probability Wasserstein distance, including its rth root.

Theorem 2.1. *There are bounded Lipschitz tests $\varphi_0, \dots, \varphi_s$ depending only on η and s with the following properties. Define*

$$D(f, g) = \max_{0 \le j \le s} \left| \int \varphi_j \, d(Kf - Kg) \right|, \qquad M(f, g) = \max_{0 \le j \le s} \left| \int x^j (f - g)(x) \, dx \right|.$$

Let δ denote either $D(f, g)$ or $M(f, g)$. For sufficiently small $\delta > 0$,

$$W_1(f, g) \le C\delta, \qquad \mathrm{TV}(f, g) \le \begin{cases} C\delta, & e = 0, \\ C\sqrt{\delta}, & e = 1 \text{ or } 2. \end{cases}$$

For each fixed $r > 1$,

$$W_r(f, g) \le \begin{cases} C_r \delta^{(r+1)/(2r)}, & e = 2, \\ C_r [\delta / (\log(1/\delta))^k]^{1/r}, & e = 0 \text{ or } 1, \end{cases} \qquad k = \begin{cases} 2, & \delta = D(f, g), \\ s, & \delta = M(f, g). \end{cases}$$

Constants can depend on the fixed reference and error law, and on r for the W_r estimates. Zero discrepancy in either case implies $f = g$. All displayed orders are sharp as inverse moduli separately for either discrepancy. The moment discrepancy orders remain sharp on the fiber where moments through order $s - 1$ agree. If η has a finite first moment, all the orders for D also hold and are sharp with $W_1(Kf, Kg)$ in its place.

For example, at a standardized uniform reference the weak inverse orders in TV, W_1, W_2 are respectively $\delta^{1/2}, \delta, \delta^{3/4}$. At a standardized Laplace reference they are $\delta, \delta, \sqrt{\delta}/\log(1/\delta)$. The Archive result [3] gives logarithmic W_2 inverse stability at references with globally positive lower and finite upper bounds on the Hessian of the potential. Here the reference has affine pieces and the alternatives need only be log-concave. Kim, Guntuboyina and Samworth [4] establish direct-sample adaptation at piecewise log-affine densities by analyzing the log-concave maximum likelihood estimator. The present result is a fixed-reference inverse inequality for additive measurements.

2.2 Finite measurements and comparison on affine pieces

Choose $\epsilon > 0$ so that the characteristic function of η has no zeros in $[-2\epsilon, 2\epsilon]$. The inverse Fourier transform of a nonzero real even smooth function supported in $(-\epsilon, \epsilon)$ is a real Schwartz function w that extends to an entire function. Its real zero set is discrete and countable. Choose a outside the set of differences of two such zeros. Then $\rho(x) = w(x)^2 + w(x-a)^2 > 0$ for every $x \in \mathbb{R}$, and ρ is Schwartz with Fourier support in $[-2\epsilon, 2\epsilon]$. Put φ_j equal to the inverse Fourier transform of the transform of $x^j \rho(x)$ divided by the characteristic multiplier of convolution with reflected η. Define the quotient to be zero off the support. It is integrable and compactly supported, so φ_j is bounded and Lipschitz. Conjugate symmetry makes it real. Fourier inversion and Fubini give

$$\int \varphi_j(x+y)\, d\eta(y) = x^j \rho(x).$$

Thus the measurements give the first $s+1$ moments of $\rho(f-g)$. If η has a first moment, $D(f,g) \le CW_1(Kf, Kg)$ by the dual formula for W_1.

On each affine piece of $\log g$, the set where $f - g > 0$ is an interval by log-concavity of f. Each such interval contributes at most two sign changes across the support of g, and each exterior half-line contributes at most one. Thus $f - g$ has at most s sign changes, counting signs after discarding zero sets. Choose a cut z_i for each change and set

$$P(x) = \pm \prod_i \frac{x - z_i}{\sqrt{1 + z_i^2}},$$

with sign chosen so that $P(f-g) \ge 0$. The coefficients of P are bounded by 2^s uniformly in the cuts. Accordingly, if $\delta = D(f,g)$, then $\int \rho |P|\, |f-g| \le C\delta$. If $\delta = M(f,g)$, then $\int |P|\, |f-g| \le C\delta$. Either identity also proves uniqueness when the discrepancy is zero.

For each affine piece choose three fixed disjoint bounded closed intervals in its interior in strictly increasing order. On each of these intervals there is a subset of length uniformly bounded below on which $|P| \ge c_g > 0$. To see this, remove sufficiently small fixed neighborhoods of all cuts lying in a fixed enlargement of the intervals. The factors for the remaining cuts are uniformly bounded away from zero there, including as $|z_i|$ tends to infinity. The weight ρ has a positive minimum on these intervals. It follows that each interval contains a point at which $|f-g| \le C_g \delta$. The values of g at the three points are uniformly bounded below, so $\log(f/g) = O_g(\delta)$ at those points when δ is small. Concavity of $\log(f/g)$ on the affine piece bounds it from above to the left of the middle point by the extended right secant, and to the right by the extended left secant. These secants have uniformly bounded slopes of order δ. This proves on all of the support interior I of g that

$$f(x) \le g(x) \exp(C\delta(1 + |x|)). \tag{1}$$

In particular, writing $h = f - g$ and $h_+ = \max(h, 0)$, for each fixed $b \ge 0$ we have

$$\int_I (1 + |x|)^b h_+(x)\, dx \le C_b \delta. \tag{2}$$

This uses the fixed exponential tails of g on infinite pieces. It also uses $e^z - 1 \le z e^z$ for $z \ge 0$.

We use the uniform bound on the density of one-dimensional log-concave probability laws of variance one, and log-concavity of the distribution and survival functions [5, 4]. One way to bound the density is to start at a mode of height m. Its superlevel set at height m/e has length at most e/m. The log-concavity secants bound density beyond this interval by exponential tails of length scale $1/m$ on either side, using zero outside the support at an abrupt endpoint. Integration about the mode bounds the variance by C/m^2.

2.3 Total variation and the linear transport bound

On the whole line, both mass and the second moment of h vanish. Thus (2) controls $\int(1+x^2)|h|$ by $C\delta$. This bounds total variation and W_1, using the dual formula after subtracting the value of the test at zero.

Suppose there is a finite endpoint. Set $Q(x) = x - A$ for support (A, ∞), $Q(x) = B - x$ for support $(-\infty, B)$, and $Q(x) = (x - A)(B - x)$ for support (A, B). Let E be the set of finite endpoints. The function Q is nonnegative inside I and nonpositive outside. Its integral against h vanishes by standardization. Using (2),

$$\int_I Qh_- + \int_{\mathbb{R}\setminus I} |Q|f = \int_I Qh_+ \le C\delta, \qquad h_- = \max(-h, 0).$$

Throughout both regions $|Q(x)| \ge c_g \operatorname{dist}(x, E)$. For the quadratic this follows by comparing with the nearer endpoint, since the farther endpoint is at distance at least $(B - A)/2$. The positive inside part is handled by (2), and we obtain

$$\int \operatorname{dist}(x, E)|h(x)|\, dx \le C\delta. \tag{3}$$

The mass outside I at distance more than $\sqrt{\delta}$ is $O(\sqrt{\delta})$ by this estimate. The mass within that distance is also $O(\sqrt{\delta})$ by the uniform density bound. The total positive discrepancy is this exterior mass plus $O(\delta)$ from (2). It equals the total variation since $\int h = 0$.

Project $h(x)\, dx$ to the nearest point of E. By (3) its integrals against a one-Lipschitz function change by at most $C\delta$. The projected signed measure has zero mass and first moment of size at most $C\delta$, since the original first moment is zero. On a single point the projected measure vanishes. On two points its coefficients are opposite and of size $O(\delta)/(B - A)$. The dual formula for W_1 now gives the linear estimate.

2.4 Higher Wasserstein distances

2.4.1 The finite ends

At any finite endpoint, the reference density is uniformly bounded above and below on a fixed one-sided neighborhood. Let F_f, F_g denote the distribution functions. For x inside such a neighborhood at distance d from the finite endpoint,

$$|F_f(x) - F_g(x)| \le C \min(\sqrt{\delta}, \delta/d). \tag{4}$$

The first bound is the total variation estimate. For a single finite endpoint, use zero signed mass and (3) on the part of the support farther from that endpoint than x is. For two finite endpoints, split at the midpoint $m = (A + B)/2$. The signed mass of $(-\infty, m]$ is $O(\delta)$ from the projected measure estimate. If $A < x < m$, the signed mass in $(x, m]$ is at most $C\delta/(x - A)$ in absolute value by (3). Reflection proves the corresponding assertion at B. The same arguments bound the distribution function difference by $C\delta$ on any fixed compact subset of the interior. With no finite endpoint the latter bound follows from the whole-line total variation estimate.

Clip the random variable with density f at its finite reference endpoints. At a right finite endpoint, the log-concavity of its survival function yields

$$S_f(B + t) \le C\sqrt{\delta} \exp(-ct/\sqrt{\delta}), \qquad t \ge 0.$$

Indeed, $S_f(B) \le C\sqrt{\delta}$, while at $B - L_0\sqrt{\delta}$, for a sufficiently large fixed L_0, the positive lower bound on g and the total variation estimate give survival at least $2S_f(B)$. Use the logarithmic secant between these two points. The assertion is immediate if $S_f(B) = 0$. Reflection treats a left finite endpoint. The rth transport cost of the clipping is therefore $O(\delta^{(r+1)/2})$.

Use the increasing quantile coupling of the clipped law with g. Near a finite left endpoint, write v for the quantile probability, or for its reflection $1-v$ near a right finite endpoint. At the reference quantile the endpoint distance is comparable to v. For $v \le C_0\sqrt{\delta}$, the quantile gap is at most $C\sqrt{\delta}$, by the total variation CDF bound and the positive lower bound on g. For $C_0\sqrt{\delta} < v < v_0$, with C_0 sufficiently large and $v_0 > 0$ fixed small, bracket the clipped quantile by moving the reference quantile left and right by $C\delta/v$. These moves stay at distances comparable to v from the endpoint and the increment in the reference CDF dominates (4). The resulting rth cost over this quantile neighborhood is at most

$$C\delta^{(r+1)/2} + C\int_{C_0\sqrt{\delta}}^{v_0} (\delta/v)^r \, dv \le C_r\delta^{(r+1)/2} \qquad (r > 1).$$

On a fixed central quantile interval the compact interior CDF bound gives cost $O(\delta^r)$ by the same bracketing. This proves the asserted upper bound when both ends are finite and accounts for the finite end when there is one.

2.4.2 The infinite ends

At every infinite end we first show that the signed density difference has a weighted tail bound

$$\int_{\text{tail}} (1 + |x|)^k |h(x)| \, dx \le C\delta, \quad k = s \text{ for } M, \qquad k = 2 \text{ for } D, \tag{5}$$

where the tail starts at a fixed sufficiently distant point. For a whole-line reference, k is even in both cases. Bound the positive part by (2) and use $|\int x^k h| \le \delta$, or the exact second-moment identity for D, to bound the negative part. Consider a reference supported on (A, ∞). For M, the integer s is odd. The polynomial $Q_s(x) = (x - A)^s$ is positive on the support and negative outside. Its integral against h is $O(\delta)$ by the moment bound. As in the proof of (3), this controls both the negative inside and the positive outside parts with weight $|x - A|^s$. For D, use the weighted sign-polynomial estimate on the fixed interval $[A - 2, A - 1]$. The same cut removal argument gives a point x_0 there with $f(x_0) \le C\delta$. One of the interior points used in (1) has density bounded below. If $f(x_0) = 0$, convexity of the support excludes any mass farther to the left. Otherwise extend the log density secant from these two points to get $f(x) \le C\delta e^{-c(x_0 - x)}$ for $x \le x_0$. This bounds the exterior second moment about A on that tail by $C\delta$. The remaining exterior part has $|x - A|^2 \le 2|x - A|$ and is controlled by (3). Apply the exact identity $\int (x - A)^2 h = 0$ and (2) to bound the negative inside second moment. This proves (5), and reflection covers support $(-\infty, B)$.

In a right infinite tail, let S_f, S_g be survival functions. For sufficiently large fixed x_0, $S_g(x) = C_g e^{-\lambda x}$ for $x \ge x_0$, with $\lambda > 0$. Equations (5) and (1) give

$$|S_f(x) - S_g(x)| \le C\delta/(1 + x)^k, \qquad S_f(x) \le S_g(x) e^{C\delta(1+x)}. \tag{6}$$

The latter follows by integrating the envelope against the exponential tail of g and absorbing the ratio of the two tail slopes into its constant. If there is a finite left endpoint, clipping there has no effect on the tails under consideration.

Write $q_f(v), q_g(v)$ for right tail quantiles with tail probability v. Put $L = \log(1/\delta)$, choose a fixed $\alpha \in (0, 1)$, and set $v_* = D_0\delta/L^k$, $v_1 = \delta^\alpha$, with D_0 a sufficiently large constant. For $v_* \le v \le v_1$,

the reference location $q_g(v) = \lambda^{-1} \log(C_g/v)$ is at least cL. Evaluate S_f at the reference quantile displaced by $\pm C\delta/(vL^k)$. These displacements are uniformly small for large D_0, and the change in S_g dominates the first error bound in (6). This brackets $q_f(v)$. The integral of the rth power of these gaps is

$$O\left((\delta/L^k)^r v_*^{1-r}\right) = O(\delta/L^k).$$

For $v_1 \le v \le v_0$ with v_0 fixed small, use a displacement $C\delta/v$ and the bound $|S_f - S_g| \le C\delta$. This costs $O(\delta^r v_1^{1-r}) = o(\delta/L^k)$.

For $v < v_*$, monotonicity and the bracket at v_* give

$$q_g(v) - q_f(v) \le C(1 + \log(v_*/v)).$$

To bound the opposite gap, invert the second inequality in (6). It gives

$$q_f(v) - q_g(v) \le C\delta(1 + \log(1/v)).$$

Indeed a fixed sufficiently large multiple of the right side added to $q_g(v)$ makes the upper survival bound at most v. Both bounds to the rth power integrate to $O(v_*)$. A left infinite tail follows by reflection. Combining the tail, finite-end and central estimates proves the upper bounds in Theorem 2.1; for $r > 1$ the finite-end contribution is lower order whenever an infinite end is present.

2.5 Sharpness

We recall the local moment chart for piecewise log-affine densities. Write the unnormalized potential as

$$V_\theta(x) = c + ax + \sum_{i=1}^{p-1} b_i (x - t_i)_+, \qquad b_i > 0,$$

with ordered knots and with finite endpoints, when present, included in the parameter θ. The s parameters vary in an open set retaining all slope and ordering conditions. At the reference the map to the first s moments in orders $0, \dots, s-1$ has invertible derivative [6]. For clarity, a row relation would give a nonzero polynomial P_0 of degree less than s. The constant and slope directions give $\int P_0 g = \int x P_0 g = 0$. For every knot the jump and location derivatives give $\int_{t_i}^B (x - t_i) P_0 g = \int_{t_i}^B P_0 g = 0$, taking $B = \infty$ if appropriate. Subtracting consecutive identities yields zero zeroth and first moments of $P_0 g$ on every open piece. There must be at least two distinct roots of P_0 in each piece, since otherwise those two identities contradict its sign. The derivative at each finite endpoint gives an additional root there. This contradicts $\deg P_0 < s = 2p + e$.

In a small chart, all polynomial moments depend smoothly on the shape parameters. The chart is locally Lipschitz in weighted L^1 for every fixed polynomial weight. Indeed, on overlapping supports the potential difference is at most $C|\theta - \theta_0|(1 + |x|)$, and the tails have uniform exponential decay. At a moving finite endpoint the symmetric difference has length $O(|\theta - \theta_0|)$ and bounded densities. We use this estimate below for normalized members of the chart to bound W_1 by the dual formula.

2.5.1 Alternatives near a finite endpoint

Assume a finite endpoint, which may be taken to be the right endpoint B by reflection. Let U be an independent centered uniform random variable on $[-1, 1]$. For small $h > 0$, convolve a chart density with the law of hU. Its moments through order $s - 1$ are smooth in the parameters and

in h^2, by the binomial formula and symmetry of U. At $h = 0$ the moment derivative in the chart is invertible. Thus for a choice $\theta(h) = \theta_0 + O(h^2)$, the construction retains exactly the moments of g through order $s - 1$. Write f_h for the resulting density. It is log-concave by preservation of log-concavity under convolution. Its sth moment discrepancy is $O(h^2)$ by the same binomial formula and smoothness.

The endpoint of the chart law changes only by $O(h^2)$, and its density near that endpoint is bounded below from inside. The convolution places mass at least ch at distance at least ch outside the original support. For example, restrict the chart input to a strip of width $h/8$ ending at its endpoint and take $U > 1/2$. It follows that

$$\mathrm{TV}(f_h, g) \geq ch, \qquad W_r(f_h, g)^r \geq c_r h^{r+1} \quad (r \geq 1).$$

For the weak discrepancies, $W_1(f_h, g) \leq Ch^2$ even before convolution with the observation error. The chart adjustment changes W_1 by $O(h^2)$. For the centered smoothing of g, write its CDF as F_g. For a displacement u subtract the linear term $-ug$ from $F_g(\cdot - u) - F_g$. The remaining L^1 norm is bounded by $\int_0^{|u|} \|g(\cdot - t) - g\|_1 \, dt$, with reflection if necessary. A log-concave density is of bounded variation, so this is $O(u^2)$. Averaging over the centered shift cancels the linear term. Convolution is a contraction under the common noise coupling, or, for the tests defining D, one can average the bounded Lipschitz functions over the noise before using the W_1 estimate. Thus $D(f_h, g) = O(h^2)$, and the same holds for $W_1(Kf_h, Kg)$ when defined. Taking h a small fixed multiple of $\sqrt{\delta}$ gives sharp total variation and W_1 rates, and sharp higher Wasserstein rates when both ends are finite.

2.5.2 Whole-line alternatives for total variation and W_1

Pick an interior point of an affine piece that is not an existing knot and add the positive hinge $h(x - t)_+$ to the potential. By the invertible chart derivative, adjust θ by $O(h)$ to fix the moments through order $s - 1$. On the original piece the first-order chart correction to the potential is affine, while the new hinge has a jump in derivative at t. Thus the first variation of the density is nonzero on a compact subinterval in this piece. Testing against a suitable bounded smooth function there proves $W_1(f_h, g) \geq ch$ and $\mathrm{TV}(f_h, g) \geq ch$. The observed W_1 and D are $O(h)$ by contraction and the weighted L^1 chart estimate, which also bounds the change in m_s. This proves sharpness with either discrepancy, including on the moment fiber.

2.5.3 Alternatives at an infinite endpoint

For a right infinite end, let R tend to infinity along that affine tail and set $m = S_g(R) \asymp e^{-\lambda R}$. For the weak discrepancy, trim the reference law at R, renormalize it, and apply an affine map to restore mean zero and variance one. Trimming and renormalization cost $O(mR)$ in W_1, by coupling the lost tail with the conditional law. The affine correction costs $O(mR^2)$, since the removed first and second moments are $O(mR)$ and $O(mR^2)$. The new support has right endpoint at most $R + o(1)$. The reference places mass comparable to m past $R + 1$, so the latent W_r cost to power r is at least $c_r m$. Contraction bounds both weak discrepancies by CmR^2 as above.

For the moment discrepancy, trim the chart density at R, and solve in the original parameters to restore the moments through order $s - 1$, including normalization. The truncated moment map differs at θ_0 from the original by $O(mR^{s-1})$. Its derivative converges uniformly to the original derivative on a neighborhood of θ_0, since the tail slopes remain bounded away from zero, with differentiated tails bounded by polynomial times exponential terms. The local inverse theorem with these uniform bounds gives a solution $\theta_R = \theta_0 + O(mR^{s-1})$. Equivalently, subtract the moment error multiplied by the inverse derivative at θ_0 at each iteration. For all large R this is a

contraction on a ball of radius a sufficiently large multiple of mR^{s-1}. The last slope changes by $O(mR^{s-1})$, so the tail of this chart density has mass comparable to m at R, since $mR^s = o(1)$. Its removed sth moment is $O(mR^s)$. The change in the untrimmed sth moment is $O(mR^{s-1})$ by smoothness. Thus $M(f_R, g) = O(mR^s)$. The law is supported to the left of R, so again $W_r^r \geq c_r m$. In either construction choose R such that mR^k is a sufficiently small fixed multiple of the prescribed discrepancy bound, with $k = 2$ or s. Then $R \asymp \log(1/\delta)$ and the lower bounds match the asserted inverse orders. Reflection gives the alternatives at a left infinite end. This proves all sharpness assertions in Theorem 2.1.

2.6 Fixed-reference tests with Gaussian errors

Corollary 2.2. *Given N independent observations from $f*\gamma_\sigma$, where γ_σ is a fixed centered Gaussian law of positive variance, consider testing the null $f = g$. For any fixed error probabilities strictly between zero and one with sum less than one, the critical separation orders are as follows. In W_1 the order is $N^{-1/2}$. In total variation the order is $N^{-1/2}$ for a whole-line reference and $N^{-1/4}$ when at least one endpoint is finite. When both endpoints are finite, the order in W_r, $r > 1$, is $N^{-(r+1)/(4r)}$.*

Proof. Estimate the $s + 1$ expectations of the bounded φ_j by their sample averages. A union bound and the variance bound for bounded variables give error $O(N^{-1/2})$, uniformly over the alternatives, with either prescribed constant error probability. The null expectations are known. Reject when the largest observed deviation exceeds a sufficiently large multiple of $N^{-1/2}$. The bounds in Theorem 2.1 give the upper orders.

For a whole-line reference the hinge construction has squared Hellinger distance $O(h^2)$ to g already before Gaussian convolution. Indeed, the difference of log densities is bounded by $Ch(1 + |x|)$, and exponential integration under g bounds the squared difference of the square roots by $O(h^2)$. Hellinger distance decreases under convolution. Take h a sufficiently small multiple of $N^{-1/2}$ in the two-point test.

For a finite endpoint consider the smoothing construction and its intermediate unsmoothed chart law $g_{\theta(h)}$. The chart satisfies $W_2(g_\theta, g) \leq C|\theta - \theta_0|$ for normalized members. To see this, the uniform CDF difference is $O(|\theta - \theta_0|)$ by the L^1 estimate. On the compact quantile ranges the densities have uniform positive lower bounds on the common interior. Near finite endpoints also use their displacements of size $O(|\theta - \theta_0|)$. Bracketing the quantiles gives gaps of the same order on these ranges. On each infinite end the quantile is affine in the logarithm of the tail probability, with coefficients smooth in the parameters on a fixed tail probability range. Its squared gap is then integrable with bound $O(|\theta - \theta_0|^2)$. Thus the Gaussian-convolved chart law and reference have Hellinger distance $O(h^2)$, by coupling their inputs and bounding the joint relative entropy of the conditional Gaussian outputs by $W_2(g_{\theta(h)}, g)^2/(2\sigma^2)$, with σ^2 the noise variance.

Adding the centered bounded jitter hU to the Gaussian noise changes its squared Hellinger distance by at most its chi-square divergence to the original noise. That divergence is

$$\mathbb{E} \exp(h^2 U U'/\sigma^2) - 1 = O(h^4),$$

where U' is an independent copy of U. The identity follows by multiplying Gaussian location likelihood ratios and integrating. The bound follows from centering and boundedness. Convolving with the common intermediate chart law can only decrease Hellinger distance. The Hellinger triangle inequality therefore gives distance $O(h^2)$ between the observed reference and smoothing alternative. Take h a sufficiently small multiple of $N^{-1/4}$. Product Hellinger bounds keep the two laws indistinguishable at the prescribed errors, and the previous section gives latent separations of the asserted orders. □

3 Gaussian posterior regret and testing at piecewise log-affine laws

For a fixed piecewise log-affine reference density, we determine the sharp inverse moduli for Hellinger distance after Gaussian noise and for the excess risk of the reference posterior mean. We also determine the fixed-reference testing rates in Wasserstein distance. A finite endpoint permits a narrow layer of displaced mass. An infinite endpoint permits discrepancies distributed across exponential tail scales. The distinction gives a critical $(\log N)^{1/4}/\sqrt{N}$ testing rate in quadratic Wasserstein distance for references supported on the whole line.

3.1 Results

Let g be a fixed log-concave density on $\mathbb{R}$ with mean zero and variance one whose log density is continuous and piecewise affine on the interior of its support, with finitely many maximal affine pieces. Write $e \in \{0, 1, 2\}$ for its number of finite support ends. Alternatives f are all one-dimensional log-concave densities with mean zero and variance one. Let $p = f * \gamma_\sigma$ and $q = g * \gamma_\sigma$, where γ_σ is a centered Gaussian density of fixed variance $\sigma^2 > 0$. Write $H^2(p, q) = \int(\sqrt{p} - \sqrt{q})^2$. For $X \sim f$ and independent Gaussian error, define $m_f(y) = \mathbb{E}_f[X \mid Y = y]$. The excess risk of using the reference posterior mean is

$$\mathcal{R}(f, g) = \mathbb{E}_f\big[(m_g(Y) - m_f(Y))^2\big] = \sigma^4 I(p\|q), \quad I(p\|q) = \int p\big[(\log p)' - (\log q)'\big]^2.$$

The last equality is the Gaussian posterior mean identity, as in the Archive inverse stability paper [3].

Theorem 3.1. *Let d be either $H(p, q)$ or $\sqrt{\mathcal{R}(f, g)}$. The sharp modulus of $W_1(f, g)$ is of order d. The sharp modulus of total variation is of order d for $e = 0$ and $d^{1/2}$ otherwise. For every fixed $r > 1$, the sharp W_r moduli are*

e	*order in* d
0	$\max\{d, d^{2/r}\}$
1	$\max\{d^{(r+1)/(2r)}, d^{2/r}\}$
2	$d^{(r+1)/(2r)}$.

Constants can depend on g, σ, r. Sharpness means matching upper and lower bounds for the supremum of the latent distance over alternatives with observed discrepancy at most a specified small radius.

Theorem 3.2. *In testing $f = g$ against standardized log-concave alternatives from N independent observations with the specified Gaussian error, the critical W_r separation for any fixed $r > 1$ and fixed error probabilities strictly between zero and one whose sum is less than one is*

reference and order	*separation*
$e = 0, \quad 1 < r < 2$	$N^{-1/2}$
$e = 0, \quad r = 2$	$N^{-1/2}(\log N)^{1/4}$
$e = 0, \quad r > 2$	$N^{-1/r}$
$e = 1$	$\max\{N^{-(r+1)/(4r)}, N^{-1/r}\}$
$e = 2$	$N^{-(r+1)/(4r)}$.

The W_1 and total variation testing orders agree with the preceding companion result Section 2.

The sharp logarithmic inverse bounds of [3] assume a reference potential with positive lower and finite upper Hessian bounds. The piecewise log-affine references here allow different rates. The weak inverse results of Section 2 use finitely many bounded measurements for general additive noise and identify a different critical role for an infinite support end. Our Hellinger and regret bounds use Gaussian tail comparisons. The critical testing rate in Theorem 3.2 arises from a sum of contributions over separated exponential tail scales. The direct-data weighted Wasserstein goodness-of-fit literature includes del Barrio, Giné and Utzet [9] and Wyłupek [10].

3.2 Transfer from observed survival functions

Denote by D the discrepancy of the finite bounded measurements in Section 2, for this reference and Gaussian noise. When D is small the comparison on affine pieces gives, on the support interior I of g,

$$f(x) \le g(x)e^{CD(1+|x|)}, \qquad \int_I (1+|x|)^b (f-g)_+(x)\,dx \le C_b D \quad (b \ge 0). \tag{7}$$

If finite endpoints exist, the projection bound from Section 2 is

$$\int \operatorname{dist}(x,E)|f(x)-g(x)|\,dx \le CD, \tag{8}$$

where E is their set. Near any finite endpoint the latent CDF discrepancy at interior distance t is bounded by $C\min(\sqrt{D}, D/t)$. On any fixed compact subinterval of I the CDF discrepancy is $O(D)$. The latter also holds on the whole line if $e = 0$, by the linear total variation bound. These statements follow from the moment and projection arguments in Section 2 using the exact standardized first two moments. Finally, in any finite outward tail the survival mass at distance t past the endpoint is bounded by $C\sqrt{D}e^{-ct/\sqrt{D}}$. This follows from log-concavity of survival, the $O(\sqrt{D})$ total variation bound and the positive density of g just inside that endpoint.

For any law a, write S_a for its right survival function. Consider a right infinite end, with $g(x) = C_0 e^{-\lambda x}$ for all sufficiently large x and some $\lambda > 0$. For some fixed M and all $x \ge M$,

$$q(x) \asymp S_q(x) \asymp g(x) \asymp e^{-\lambda x}, \qquad p(x) + S_p(x) + S_f(x) \le Ce^{-(\lambda - CD)x}. \tag{9}$$

For the upper bounds integrate (7) in the right tail, and bound the Gaussian convolution of the mass to the left of a fixed point by a Gaussian remainder. For the lower bound on q, integrate over latent inputs within a bounded interval around x in the exponential part. The same reasoning with $D = 0$ gives the upper estimate on q.

The following signed comparison is useful:

$$|S_f(x) - S_g(x)| \le 2|S_p(x) - S_q(x)| + CD(1+x)e^{-(\lambda - CD)x} \qquad (x \ge M). \tag{10}$$

Write $h = f - g$ and $K_x(t) = \mathbb{P}(t + \sigma G \ge x)$ for standard Gaussian G. If g has full support, the integral of $K_x h_+$ is at most $CD(1+x)e^{-(\lambda - CD)x}$. On a right affine piece this follows by integrating (7) against the Gaussian survival function: the convolution of $(1+t)e^{-(\lambda - CD)t}$ on a right half-line against that survival is bounded by $C(1+x)e^{-(\lambda - CD)x}$, using finite Gaussian exponential moments. On the part to the left of a fixed point use (7) and the Gaussian decay of K_x. Since $K_x \ge 1/2$ for $t \ge x$, the negative latent tail mass is bounded by twice its observed negative contribution. This contribution is bounded by the positive observed contribution plus $|S_p - S_q|$. Adding the positive latent tail mass gives (10).

If the opposite endpoint A is finite, first combine the signed part on $J = (-\infty, A+1]$. Its net mass is $O(D)$ by the compact interior CDF bound, and $\int_J |t - A||h(t)|\,dt \le CD$ by (8). The

function K_x and its derivative are bounded by Ce^{-cx^2} on J for large x. Subtracting $K_x(A)$ inside this signed block proves $|\int_J K_x h| \le CDe^{-cx^2}$. On the complement the previous bound on the observed positive contribution applies. The argument for the negative mass past x is identical after allowing the bounded signed block. This proves (10) in both cases. Reflection gives the left-end statements.

3.3 Transport bounds from the transferred discrepancy

We use the following one-dimensional transport inequality. If probability laws a, b are supported on an interval where b has positive density in the interior, then

$$W_r(a,b)^r \le r^r \int |F_a - F_b|^r b^{1-r}, \qquad r > 1, \tag{11}$$

integrating over the support interior. For continuous densities on the interior, interpolate linearly between b and a and differentiate their increasing quantile functions. At time $t < 1$, the derivative in t is the negative CDF discrepancy divided by the interpolated density, evaluated at the quantile. Its L^r norm is at most $(1-t)^{-(r-1)/r}(\int |F_a - F_b|^r b^{1-r})^{1/r}$, by the change of variable from probability to location. Integrating in time and using Minkowski gives the factor r, since the quantiles converge in L^r at $t = 1$. This convergence also follows by a mixture coupling at distance at most $(1-t)^{1/r} W_r(a,b)$. The same differentiation applies to densities continuous except at the finitely many support endpoints. If a has a clipping atom at a finite endpoint at which the density of b has a positive inner limit, spread the atom uniformly over a shrinking one-sided interval and pass to the limit. Both the transport cost of this replacement and the change in the integral tend to zero.

Apply (11) with $b = g$ and a the latent law clipped at any finite endpoints of g. The clipping cost to power r is at most $CD^{(r+1)/2}$ by the finite outward tail just described. Clipping does not change CDFs on the support interior. The central spatial region contributes $O(D^r)$. Each finite endpoint contributes at most

$$C \int_0^{c_0} \min(\sqrt{D}, D/t)^r \, dt \le C_r D^{(r+1)/2},$$

since g is bounded below on that region. Thus the only additional integral at a right infinite end is, up to constants,

$$\int_M^\infty |S_f(x) - S_g(x)|^r e^{(r-1)\lambda x}, dx.$$

3.3.1 Hellinger bounds

Set $d = H(p,q)$. Boundedness of the probes implies $D \le Cd$. Suppose $d > 0$ is small and put $L = d^{-1/2}$. On $[M, L]$, p/q is bounded by (9), so

$$\int_M^L |p-q|^2 e^{\lambda x} \, dx \le Cd^2.$$

The truncated survival integral obeys the weighted Hardy bound

$$\int_M^L \left| \int_x^L (p-q)(y) \, dy \right|^2 e^{\lambda x} \, dx \le C \int_M^L |p-q|^2 e^{\lambda y} \, dy.$$

Apply Young's inequality after multiplying by $e^{\lambda x/2}$, with the integrable kernel $e^{-\lambda(y-x)/2}\mathbf{1}_{y \ge x}$. For $x \le L/2$ the omitted survival beyond L has weighted squared integral $O(e^{-cL})$ by (9). The

remainder in (10) has squared weighted integral $O(d^2)$. Consequently the latent squared weighted flux over $[M, L/2]$ is at most Cd^2. The ratio of latent CDF discrepancy to g on this interval is bounded by (9). Use Hölder against the measure $g(x)\,dx$ if $1 < r \le 2$, or this ratio bound if $r \ge 2$, to obtain an rth weighted flux of order at most d^r for $r \le 2$ and d^2 for $r \ge 2$. The latent survival envelope alone bounds the integral past $L/2$ by $O(e^{-c_r L})$. Reflection completes the proof of the higher Wasserstein upper bounds. Total variation and W_1 use the bounded probes and Section 2 directly.

For regret, use the Poincaré inequality for the fixed log-concave density q [7]. With $u = \sqrt{p/q}$ we have $\int u^2 q = 1$, $H^2 = 2(1 - \int uq) \le 2\,\mathrm{Var}_q\, u$, and $\int |u'|^2 q = I(p\|q)/4$. Thus $H(p,q) \le C\sqrt{\mathcal{R}(f,g)}$, allowing all the same upper bounds with the square root of regret. If the discrepancy is zero, the finite measurement identification in Section 2 gives $f = g$.

3.4 Sharp lower bounds for Hellinger distance and regret

We first record a score comparison used for all three families. Let X have the reference law tilted by a density factor proportional to $e^{\theta x}$, and consider $Y = aX + b + \sqrt{v}G$, where (θ, a, b, v) is near $(0, 1, 0, \sigma^2)$. Denote the output density by $q_{\theta,a,b,v}$ and the posterior mean of X by m_y. The score is $(am_y + b - y)/v$. The conditional potential has curvature at least a^2/v, since the latent potential is convex. The Brascamp–Lieb variance bound gives a uniform bound on conditional variance, also under any additional linear tilt of the conditional distribution [8]. Integrating the resulting bound on the second derivative of the log moment generating function gives a uniform sub-Gaussian bound on the centered conditional law. In particular its centered third absolute moment is uniformly bounded. These inequalities also hold for a convex potential with hard support boundaries, by approximation or the one-dimensional strong log-concavity form of the variance bound.

At $y = 0$, m_y is uniformly bounded on the fixed parameter neighborhood by its conditional integral formula. The derivative with respect to y is $a\,\mathrm{Var}(X \mid Y = y)/v$, so $|m_y| \le C(1 + |y|)$. Differentiate the same conditional integral formula with respect to θ, a, b, v. Each derivative is the conditional covariance of X with the derivative of the unnormalized log density of the posterior. For the scale parameter a, for example, it is

$$v^{-1}\{(y - b - 2am_y)\,\mathrm{Var}(X \mid Y = y) - a\mathbb{E}[(X - m_y)^3 \mid Y = y]\}.$$

For θ it is the conditional variance. For b it is $-a/v$ times the variance. For v it is the covariance with $(y - b - aX)^2/(2v^2)$, bounded by the same centered moment estimates. Therefore the output score moves by at most

$$C(1 + |y|)(|\theta| + |a - 1| + |b| + |v - \sigma^2|) \tag{12}$$

from the score of q. Differentiation is justified by Gaussian domination of the conditional integrals for each fixed y.

Finite endpoint. Let $X \sim g$ and use the latent law $(X + hG')/\sqrt{1 + h^2}$, for independent standard Gaussian G'. It is standardized and log-concave. Its output density is $q_{0,a,0,v}$ for $a = (1 + h^2)^{-1/2}$ and $v = \sigma^2 + a^2h^2$. The parameter deviations are $O(h^2)$. By (12) and the fixed observation second moment, the regret is $O(h^4)$. A latent input strip of width proportional to h just inside a finite endpoint has probability at least ch, and with uniformly positive probability the added Gaussian pushes it at least ch outside that endpoint. Therefore the total variation is at least ch and W_r^r is at least $c_r h^{r+1}$ for every $r \ge 1$.

Infinite endpoint. Trim a right exponential tail of mass m at R. With conditional latent mean and variance after trimming, apply a positive affine map to center and standardize the law, with parameters $|a-1|+|b| = O(mR^2)$. As in Section 2, the standardized law has no mass past $R+o(1)$ and g puts mass comparable to m past $R+1$. Hence $W_r^r \geq c_r m$. We compare the output first with the untrimmed affine reference q_{0,a,b,σ^2}, denoted q_*. If $U = aX + b$ in this untrimmed reference, set $P(y) = \mathbb{P}(X \leq R \mid U + \sigma G = y)$. The trimmed output density is $p = q_* P/(1-m)$. The Gaussian differentiation identity gives

$$P'(y) = \sigma^{-2} \operatorname{Cov}(\mathbf{1}_{X \leq R}, U \mid Y = y).$$

The conditional law of U under the untrimmed affine reference is strongly log-concave with curvature at least σ^{-2}. Its variance is at most σ^2. Conditional Cauchy–Schwarz then gives

$$I(p\|q_*) = \frac{1}{1-m} \int q_* \frac{(P')^2}{P} \leq \frac{m}{\sigma^2(1-m)}.$$

By (12), the score of q_* differs from that of q in squared $L^2(p)$ norm by $O((mR^2)^2) = o(m)$. The triangle inequality in $L^2(p)$ gives $\mathcal{R} = O(m)$. Reflection treats a left infinite end.

Full support. For a whole-line reference, tilt by e^{hx} and affinely standardize. The tilt and affine parameters change by $O(h)$, so (12) gives regret $O(h^2)$. The signed log slopes at the two infinite ends differ. If the standard deviation of the tilted law is $s(h)$, the two slopes after standardization are $s(h)(\beta_L + h)$ and $s(h)(\beta_R + h)$, where $\beta_L > 0 > \beta_R$. At least one derivative at zero is nonzero. On a fixed compact interval of that affine end, the derivative of the log density is therefore a nonconstant affine function. A bounded smooth test supported there witnesses a change of size at least ch. Both W_1 and TV are at least ch, and $W_r \geq W_1$.

Taking the parameter in each example at the scale prescribed by the regret bound proves all the lower orders in Theorem 3.1 for the square root of regret. The same families give Hellinger lower bounds on the inverse moduli, since $H(p,q) \leq C\sqrt{\mathcal{R}(f,g)}$.

3.5 Tests based on observed tail flux

Use the finitely many bounded measurements defining D as a first test. Their sample deviations are $O(N^{-1/2})$ uniformly in the underlying law at any specified constant error level. Under the null, a sufficiently high threshold of this order controls the first test's type I error. Alternatives with population $D > C_0 N^{-1/2}$, for a sufficiently large C_0, are detected by this test alone. Thus for the remaining power analysis suppose $D \leq C_0 N^{-1/2}$. This is a restriction on the population, independent of the event that the measurement test rejects.

At a right infinite end take $R_N = (\log N)/\lambda$. In $[M, R_N]$, $p(x) \leq Cq(x)$ and $S_p(x) \leq CS_q(x)$ by (9). Past R_N the latent tail integral in the flux bound is $O(N^{-1})$, again by (9), since $DR_N = o(1)$. Equations (10) and (11) thus give the deterministic bound

$$W_r^r \leq C_r \left(D^r + \mathbf{1}_{e>0} D^{(r+1)/2} + \sum_{\text{infinite ends}} \left[N^{-1} + \int_M^{R_N} \frac{|S_p(x) - S_q(x)|^r}{S_q(x)^{r-1}}\, dx \right] \right). \tag{13}$$

Each left end is expressed in reflected coordinates, with its own cutoff and tail slope. The additive term involving D^r also absorbs the transferred error in (10).

Let $\widehat{S}$ be the empirical survival function. At any $x \in [M, R_N]$, Bernstein's inequality for Bernoulli counts and its integral moment consequence give

$$\mathbb{E}_p|\widehat{S}(x) - S_p(x)|^r \leq C_r\big((S_p(x)/N)^{r/2} + N^{-r}\big) \leq C_r (S_q(x)/N)^{r/2}.$$

The last step uses $S_q(x) \geq c/N$. Integrating after division by S_q^{r-1} gives the bounds $C_r N^{-r/2}$ for $1 < r < 2$, $C \log N/N$ for $r = 2$, and C_r/N for $r > 2$. This is the integral of an exponential weight: $S_q^{1-r/2} \asymp e^{(r/2-1)\lambda x}$. Use the rth root of each empirical flux $\int |\widehat{S} - S_q|^r / S_q^{r-1}$ as a tail statistic. Markov and the triangle inequality in the weighted L^r norm control both its null fluctuations and its distance to the corresponding population norm under the remaining alternatives. There are at most two tails. Combining with (13) proves the stated testing upper bounds except for the critical whole-line case $r = 2$. When there is a finite endpoint, the $r = 2$ tail cost $N^{-1} \log N$ is lower order than the finite endpoint cost $N^{-3/4}$.

For $r = 2$ with no finite endpoints, instead estimate the squared population flux directly. At one end set

$$Z_i(x) = \frac{\mathbf{1}_{Y_i \geq x} - S_q(x)}{\sqrt{S_q(x)}}, \qquad U_N = \frac{1}{N(N-1)} \sum_{i \neq j} \int_M^{R_N} Z_i(x) Z_j(x)\, dx.$$

Let $\mu = \mathbb{E}_p Z_i$, $\theta = \|\mu\|_2^2$, and let $\mathcal{K}$ be the covariance operator of Z_i on this interval. Then $\mathbb{E}_p U_N = \theta$ and

$$\operatorname{Var}_p U_N = \frac{4}{N} \langle \mu, \mathcal{K}\mu \rangle + \frac{2}{N(N-1)} \|\mathcal{K}\|_{\mathrm{HS}}^2.$$

Indeed, expand each Z_i into its mean and centered part. The linear and quadratic centered sums are uncorrelated by independence, and distinct off-diagonal pairs in the quadratic sum are uncorrelated unless they share both indices. The covariance kernel is bounded by

$$C \frac{S_q(\max(x,y))}{\sqrt{S_q(x) S_q(y)}} \leq C e^{-\lambda |x-y|/2}.$$

The first bound follows from covariance of nested survival indicators and the upper bound on S_p/S_q. The operator norm is uniformly bounded by integrating the kernel in either variable, and the squared Hilbert–Schmidt norm is $O(\log N)$. Thus $\operatorname{Var}_p U_N \leq C(\theta/N + \log N/N^2)$. Reject in either tail when its statistic exceeds a sufficiently large multiple of $\sqrt{\log N}/N$. Chebyshev gives the desired errors whenever in the alternative one flux is at least a suitably larger multiple. This detects separation of order $(\log N)^{1/4}/\sqrt{N}$ by (13).

3.6 Testing lower bounds

The ordinary lower bounds follow from the alternatives for Theorem 3.1 and the product inequality $H(p^{\otimes N}, q^{\otimes N}) \leq \sqrt{N} H(p, q)$. The full-line tilt uses h of order $N^{-1/2}$. The finite endpoint uses h of order $N^{-1/4}$. An infinite end uses a trim of mass of order N^{-1}. Small constants make the product total variation smaller than one minus the sum of the specified errors. They give all the stated lower bounds other than the critical whole-line logarithm.

3.6.1 Many tail scales at the quadratic threshold

Suppose g has whole-line support. Use its right tail with $S_g(x) = c_g e^{-\lambda x}$ past a fixed point. Take $R_1, \ldots, R_J$ in this region, separated by a fixed distance $D_0 > 3$, with

$$m_j = S_g(R_j) = N^{-1/4} e^{-\lambda D_0 (j-1)} \quad \text{and} \quad m_j \geq N^{-1/2}.$$

Take all such indices, so $J \asymp \log N$. Let $a > 0$ be a small fixed constant, $b_j = a/\sqrt{N m_j}$, and for an active set A put

$$u_A(x) = \sum_{j \in A} b_j (x - R_j)_+, \qquad f_A(x) = \frac{g(x) e^{-u_A(x)}}{Z_A}, \qquad Z_A = \mathbb{E}_g e^{-u_A}.$$

These tilted densities are log-concave. Their small violations of standardization will be corrected below.

From exponential tail integration and the geometric progression of m_j,

$$1 - Z_A \le \mathbb{E}_g u_A \le CaN^{-5/8}, \qquad \mathbb{E}_g[(1+X^2)u_A(X)] \le CaN^{-5/8}(\log N)^2$$

uniformly in A. The second estimate bounds the first and second moment changes by the same order. Standardize f_A by an affine map to get $\bar{f}_A$. The coupling applying this map to the same input gives

$$W_2(f_A, \bar{f}_A) \le CaN^{-5/8}(\log N)^2 = o(N^{-1/2}).$$

For any two latent input laws, the relative entropy between their Gaussian-convolved laws is at most their squared W_2 distance divided by $2\sigma^2$. This follows by coupling the inputs and comparing their two conditional Gaussian outputs before marginalizing. Thus the product observation laws for f_A and $\bar{f}_A$ differ by $o(1)$ in total variation, uniformly in A.

For $j \in A$ and $x \in [R_j + 1, R_j + 2]$, put $v = S_g(x)$. For a small fixed $c_0 > 0$, the survival of f_A at $x - c_0 b_j$ is bounded by

$$Z_A^{-1} e^{-b_j(1-c_0 b_j)} e^{\lambda c_0 b_j} v < v$$

when N is large. Indeed $\log Z_A^{-1} = o(b_j)$ uniformly, and take $\lambda c_0 < 1/4$. The inequality follows by keeping only the jth hinge in the integrand. It brackets the difference in right quantiles below by $c_0 b_j$ for a band of probabilities of length comparable to m_j. These bands are disjoint. Consequently

$$W_2(f_A, g)^2 \ge c \sum_{j \in A} m_j b_j^2 = ca^2|A|/N.$$

The same separation, up to a constant, holds for the standardized laws for any nonempty A, by the smaller affine correction.

Choose the active indices independently with probability $\pi = J^{-1/2}$. Consider the mixture of product laws over A, first in the experiment observing the raw unstandardized densities f_A. Write $L_A = f_A/g$. For two active sets A, B,

$$\mathbb{E}_g L_A L_B = 1 + \frac{\mathrm{Cov}_g(1 - e^{-u_A}, 1 - e^{-u_B})}{Z_A Z_B} \le 1 + C\mathbb{E}_g[u_A u_B].$$

For $i \le j$, direct exponential integration gives

$$\mathbb{E}_g[(X - R_i)_+ (X - R_j)_+] = m_j \left(\frac{R_j - R_i}{\lambda} + \frac{2}{\lambda^2} \right).$$

It follows that

$$N\mathbb{E}_g[u_A u_B] \le a^2 \sum_{i \in A, j \in B} k_{i-j}, \qquad k_l = \left(\frac{D_0|l|}{\lambda} + \frac{2}{\lambda^2} \right) e^{-\lambda D_0 |l|/2},$$

a summable nonnegative kernel. The product likelihood second moment of the mixture is therefore at most

$$\mathbb{E}_{A,B} \exp\left(Ca^2 \sum_{i \in A, j \in B} k_{i-j} \right).$$

Condition on A. With $s_j = \sum_{i \in A} k_{i-j}$, we have $0 \le s_j \le \|k\|_1$ and $\sum_j s_j \le \|k\|_1 |A|$. Independence of the indices in B bounds this conditional expectation by

$$\prod_j [1 + \pi(e^{Ca^2 s_j} - 1)] \le \exp(C' a^2 \pi |A|).$$

Averaging over A by the same Bernoulli calculation gives a bound $\exp(C''a^2\pi^2 J) = \exp(C''a^2)$. Hence the mixture chi-square divergence can be made arbitrarily small by the fixed choice of a. Conditioning on $|A| \geq \sqrt{J}/2$ changes the mixture by $o(1)$ in total variation by Bernoulli concentration. Adding independent Gaussian noise cannot increase total variation, and standardization changes the noisy mixture by $o(1)$ as proved above. The resulting alternatives are indistinguishable from the noisy reference null at any prescribed pair of errors with sum less than one, while their latent separation is at least

$$c(\log N)^{1/4}/\sqrt{N}.$$

This proves the critical lower bound of Theorem 3.2.

4 Vanishing Gaussian noise in tests at piecewise log-affine laws

We test a fixed piecewise log-affine law against standardized log-concave alternatives after Gaussian noise of known standard deviation σ, allowing σ to vanish with the sample size. At a finite reference endpoint, the sharp total variation separation contributed by noise is $(\sigma^3/N)^{1/4}$. This competes with a clean-data rate N^{-1} for a single affine piece and $N^{-1/2}$ for two or more pieces. We also obtain the corresponding sharp inverse moduli for observed Hellinger distance. The proof uses a localized odd measurement at the endpoint and matches it by Gaussian smoothing of the reference before standardization.

4.1 Statement of the results

Let g be a fixed real log-concave probability density with mean zero and variance one. Its log density is continuous and piecewise affine on the interior I of its support, with p maximal affine pieces and $e \in \{0, 1, 2\}$ finite support endpoints. All alternative densities f are log-concave and have mean zero and variance one. Let $K_\sigma f = f * \gamma_\sigma$, where γ_σ is centered Gaussian with standard deviation σ, and γ_0 is the point mass at zero. Write $H^2(a, b) = \int(\sqrt{a} - \sqrt{b})^2$.

Theorem 4.1. *Uniformly for $0 \leq \sigma \leq 1$ and small $d > 0$, the supremum of $\mathrm{TV}(f, g)$ subject to $H(K_\sigma f, K_\sigma g) \leq d$ is bounded above and below by fixed positive multiples of*

$$\begin{cases} d, & e = 0, \\ \max\{d^2, \sigma^{3/4} d^{1/2}\}, & e \geq 1,\ p = 1, \\ \max\{d, \sigma^{3/4} d^{1/2}\}, & e \geq 1,\ p \geq 2. \end{cases}$$

The constants can depend on g alone.

Theorem 4.2. *Suppose N independent observations have density $K_\sigma f$. Fix type I and type II error bounds strictly between zero and one whose sum is less than one. For testing $f = g$, the critical total variation separation is bounded above and below by fixed positive multiples of the orders in Theorem 4.1 evaluated at $d = N^{-1/2}$, uniformly for $0 \leq \sigma \leq 1$. In particular, the two finite-end cases have separation*

$$\max\{N^{-1}, (\sigma^3/N)^{1/4}\} \quad \textit{for } p = 1, \qquad \max\{N^{-1/2}, (\sigma^3/N)^{1/4}\} \quad \textit{for } p \geq 2.$$

For example, the clean-data boundary scale N^{-1} holds for a standardized uniform law and a centered shifted exponential law of variance one. With two or more affine pieces the clean-data scale is $N^{-1/2}$, including references with finite support. The crossover occurs at σ of order N^{-1} in the first case and of order $N^{-1/3}$ in the second.

The Archive deconvolution result [3] studies fixed Gaussian noise with a reference potential of positive lower and finite upper curvature. The companion results Sections 2, 3 consider piecewise log-affine references and give rates at fixed positive noise. The present bounds track the noise standard deviation through the clean-data limit and distinguish a single affine piece from several pieces. Direct log-concave density estimation adapts at log-affine references [4]. The exact moment constraints on the alternatives here allow a faster fixed-reference test in the one-piece case. Related work on support endpoints under additive observation errors includes [11, 12].

4.2 Uniform measurements and the effect of one affine piece

Keep the strictly positive Schwartz weight with compact Fourier support from the proof of the bounded measurement comparison in Section 2. Dividing the Fourier transform of a fixed weighted monomial by the Gaussian characteristic function gives bounded measurement functions uniformly over $0 \le \sigma \le 1$, since the division is on a fixed compact frequency interval. Let D be the maximum absolute measurement discrepancy, with $2p + e + 1$ terms. The conclusions from the affine-piece comparison then hold with constants independent of σ. In particular, put $h_+ = (f - g)_+$ and $h_- = (g - f)_+$. For small D we have

$$f(x) \le g(x)e^{CD(1+|x|)} \quad (x \in I), \qquad \int_I (1 + |x|)^2 h_+(x)\, dx \le CD. \tag{14}$$

The same bound gives $\sup_I h_+ \le CD$ by the fixed exponential tails of g. If $e = 0$, these bounds give $\mathrm{TV}(f, g) \le CD$ by mass balance. If $e > 0$, the projection bound in Section 2 gives

$$\int \mathrm{dist}(x, E)|f(x) - g(x)|\, dx \le CD, \tag{15}$$

where E is the finite endpoint set. The compact interior CDF difference is $O(D)$. Also $\mathrm{TV}(f, g) \le C\sqrt{D}$, using (15) with the uniform density bound for variance-one log-concave laws. In the deterministic setting $D \le Cd$ follows by Cauchy–Schwarz from $H \le d$ and the uniform boundedness of the probes.

Lemma 4.3. *If $p = 1$ and $e \ge 1$, then $\mathrm{TV}(f, g) \le Cm_{\mathrm{out}}$ for every standardized log-concave density f, where $m_{\mathrm{out}} = \int_{\mathbb{R}\setminus I} f$.*

Proof. First,

$$\int_{\mathbb{R}\setminus I} (1 + x^2) f(x)\, dx \le Cm_{\mathrm{out}}.$$

For a right exterior this is immediate on bounded intervals. Choose a fixed R larger than the endpoint and larger than 3. By Chebyshev and log-concavity of survival, the logarithmic secant from -2 to R gives $S_f(R + t) \le S_f(R)e^{-ct}$, with $c > 0$ independent of f. The bound is immediate if $S_f(R) = 0$. Since $S_f(R) \le m_{\mathrm{out}}$, integrate the tail second moment. Reflect for a left exterior.

On the single affine piece the positive set of $f - g$ is an interval, so there are at most two sign changes inside I. A polynomial P with a root at each change, normalized by $\prod_i \sqrt{1 + z_i^2}$ if the roots are z_i, can be signed so that $P(f - g) \ge 0$ on I, with uniformly bounded coefficients. The exact moments of orders $0, 1, 2$ then imply

$$\int_I |P||f - g| = \left| \int_{\mathbb{R}\setminus I} Pf \right| \le Cm_{\mathrm{out}}.$$

Choose three ordered bounded subintervals inside I. Removing sufficiently small neighborhoods of the at most two roots leaves a fixed positive length in each subinterval where $|P|$ is bounded below.

Each subinterval therefore contains a point with $|f-g| \le Cm_{\text{out}}$. The density g is bounded below there. For small m_{out}, extend the secants of the concave function $\log(f/g)$ from the two adjacent ordered point pairs to get $f(x) \le g(x)e^{Cm_{\text{out}}(1+|x|)}$ on I, as in Section 2. Integrating the positive part inside gives $\int_I h_+ \le Cm_{\text{out}}$. Mass balance identifies total variation with this positive mass plus the exterior mass. For large m_{out} the inequality follows from TV ≤ 1. □

4.3 A measurement near the endpoint

Assume $e \ge 1$. Choose a finite endpoint with largest exterior mass, denoted by m, and reflect if necessary so the endpoint is a right endpoint B. When $p \ge 2$, the total variation is $O(D+m)$ by (14) and mass balance. In this case it suffices to treat $m \ge D$. When $p = 1$, total variation is at most Cm by Lemma 4.3, so no lower bound on m/D is needed. We can suppose $m > 0$ and small because the general bound TV $\le C\sqrt{D}$ holds.

In the cases just specified, the mass shortage $\int h_-$ inside any subinterval is at most TV $\le Cm$. The density of g is bounded below just inside B. Taking a sufficiently large fixed L_0 gives $S_f(B - L_0 m) \ge 2m$ for small m. The log-concavity secant for survival therefore gives

$$S_f(B+t) \le me^{-ct/m} \quad (t \ge 0), \qquad \int_B^{B+b} f(x)\,dx \ge c' \min\{m, b\} \quad (b > 0). \tag{16}$$

The second inequality subtracts survival at $B+b$ from that at B.

For $b \ge 2\sigma$, use the latent function

$$\psi_b(x) = \frac{x-B}{b} \exp\left(-\frac{(x-B)^2}{2b^2}\right).$$

It is the conditional expectation given the latent input of a bounded observed function ϕ_b. Indeed, if $a = (b^2 - \sigma^2)^{1/2}$, take

$$\phi_b(y) = \frac{b^2}{a^2}\frac{y-B}{a} \exp\left(-\frac{(y-B)^2}{2a^2}\right).$$

Direct Gaussian convolution, or differentiation of the convolution of two Gaussian functions, verifies the identity. The observed function has uniformly bounded sup norm and squared Lebesgue L^2 norm $O(b)$. The observation density is uniformly bounded, since the latent densities have variance one and are log-concave. Hence the observed variance is $O(b)$ under both the null and any alternative.

We claim for all sufficiently small b,

$$\int \phi_b\, d(K_\sigma f - K_\sigma g) \ge c\frac{\min(m,b)^2}{b} - CDb. \tag{17}$$

Beyond B, the integrand $\psi_b(f-g)$ is nonnegative. On $[B, B+b]$ the kernel is bounded below by a constant times $(x-B)/b$. If a density is bounded by M_0 and puts mass u in this interval, its first moment about B there is at least $u^2/(2M_0)$: filling adjacent to B at the density cap minimizes that moment. Using (16) proves the first term of (17). The contribution of the positive discrepancy inside the support is in absolute value at most CDb by $\sup_I h_+ \le CD$ and the L^1 norm of the kernel. The negative inside part is favorable since the kernel is nonpositive.

If a left finite endpoint A also exists, combine all the signed discrepancy over $J = (-\infty, (A+B)/2]$ before applying these interior bounds on its complement. The integral of $f-g$ on J is $O(D)$ by the interior CDF bound, and $\int_J |x-A||f-g|(x)\,dx \le CD$ by (15). On J the kernel and its derivative are bounded by Cb for all sufficiently small b, because J stays a fixed distance from B. Subtracting the constant $\psi_b(A)$ in the signed integral bounds this block in absolute value by CDb. This proves (17).

4.4 Inverse and testing upper bounds

We give one argument for the two upper bounds. In the deterministic problem put d equal to the specified Hellinger radius. In testing put $d = N^{-1/2}$. In the former setting the observation differences of the bounded probes are $O(d)$, so $D \le Cd$. Also the observed endpoint function satisfies

$$\left| \int \phi_b \, d(K_\sigma f - K_\sigma g) \right| \le Cd\sqrt{b}. \tag{18}$$

Indeed, write the density difference as the product of the difference and sum of square roots, apply Cauchy–Schwarz, and use the uniformly bounded observation densities and $\|\phi_b\|_2^2 = O(b)$.

In testing, use the sample averages of the bounded probes for D, and the average of ϕ_b for each finite endpoint, oriented outwards. Their errors relative to population means are $O(d)$ and $O(d\sqrt{b})$, respectively, with arbitrarily prescribed constant probability, by the variance bounds and Chebyshev's inequality. Reject if a bounded probe has a large absolute deviation from the null, or if an endpoint probe has a large positive deviation from the null. Under the null, sufficiently large fixed thresholds of the indicated orders control the type I error. For the alternative analysis it suffices to bound the latent distance when $D \le Cd$ and the endpoint deviations are bounded above by $Cd\sqrt{b}$, for sufficiently large constants. If a population deviation exceeds this bound, the corresponding test rejects at the desired power. Thus the needed upper bounds on population observations have the same form in the two settings.

For $e = 0$, total variation is $O(d)$ from the uniform global probes alone. If $e \ge 1$ and σ is above a sufficiently small fixed positive constant, use $\mathrm{TV} \le C\sqrt{D}$; this is at most a fixed multiple of $\sigma^{3/4} d^{1/2}$. In the remaining noise range choose $b = \max(2\sigma, A_0 d^2)$, with fixed A_0 large and with d small. This lies in the small-bandwidth regime of (17). In the cases in which m must be bounded, (17) and the population observation bounds give

$$\frac{\min(m, b)^2}{b} \le C(d\sqrt{b} + Db) \le C' d\sqrt{b}.$$

If $m \ge b$, this implies $b \le C' d\sqrt{b}$, contradicting $b \ge A_0 d^2$ for sufficiently large A_0. Hence

$$m \le C d^{1/2} b^{3/4} \le C' \max(d^2, \sigma^{3/4} d^{1/2}).$$

Lemma 4.3 yields the bound for $p = 1$. For $p \ge 2$, combine this with $\mathrm{TV} \le C(D + m)$, including the case $m < D$. These arguments prove the upper bounds of Theorems 4.1 and 4.2.

4.5 Lower bounds uniform in the noise level

4.5.1 Smoothing a finite endpoint

Let $X \sim g$ and G' be independent standard Gaussian. Use the latent alternative

$$X_h = \frac{X + hG'}{\sqrt{1 + h^2}}.$$

It is log-concave and standardized. A strip of width proportional to h just inside a finite endpoint has mass at least ch. With a uniformly positive probability the centered Gaussian shift pushes that mass outside the original support. The affine scale change of the endpoint is only $O(h^2)$. Thus $\mathrm{TV}(f_h, g) \ge ch$, where f_h is the density of X_h. The clean total variation is also $O(h)$, by bounded variation of g for translation and the $O(h^2)$ affine correction. For the latter estimate, on overlapping supports the log density change is at most $Ch^2(1+|x|)$, and the moving finite endpoints

cost intervals of length $O(h^2)$. The tails on the overlapping supports are uniformly exponential. Since Hellinger distance contracts under convolution and $H^2 \le 2\mathrm{TV}$, we obtain

$$H^2(K_\sigma f_h, K_\sigma g) \le Ch. \tag{19}$$

To measure the advantage of smoothing by the noise, put $q_v = g * \gamma_{\sqrt{v}}$ for $v > 0$. We claim

$$\int \frac{(q_v'')^2}{q_v} \le Cv^{-3/2} \qquad (0 < v \le 2). \tag{20}$$

In the sense of distributions, $Dg = \ell g\,dx + \sum_B c_B \delta_B$, where $\ell = (\log g)'$ on the interior is bounded and the sum runs over finite endpoints. The density is continuous at all interior knots, so only the endpoints contribute atoms. Write $q_v'' = (Dg) * \gamma'_{\sqrt{v}}$. For the continuous part, Jensen with the conditional law of the input given the Gaussian observation bounds the weighted squared norm by

$$\mathbb{E}\left[\ell(X)^2 \frac{(Y-X)^2}{v^2}\right] \le C/v, \qquad Y = X + \sqrt{v}G.$$

For a right endpoint term let $s = \sqrt{v}$, $z = (y - B)/s$, and let φ denote standard Gaussian density. On a fixed interior interval adjacent to B, g is bounded below. Therefore

$$q_v(B + sz) \ge c \int_0^{c_0} \varphi(z + t)\,dt$$

with a fixed $c_0 > 0$ for all $s \le \sqrt{2}$. Integrate only up to a fixed multiple of $(1 + |z|)^{-1}$ to bound the last integral below by $c'\varphi(z)/(1 + |z|)$. The weighted squared norm of $\gamma_s'(\cdot - B)$ is at most

$$Cs^{-3} \int \frac{(\varphi'(z))^2(1 + |z|)}{\varphi(z)}\,dz \le C's^{-3}.$$

Reflection treats a left endpoint. The squared norm of the finite sum is bounded by the sum of the squared norms up to a fixed factor. This proves (20).

The heat equation gives $\partial_v \sqrt{q_v} = q_v''/(4\sqrt{q_v})$ with L^2 norm at most $Cv^{-3/4}$. For $\sigma > 0$, set $a = (1+h^2)^{-1/2}$ and $v = \sigma^2 + a^2h^2$. Integrating the heat derivative from σ^2 to v bounds the Hellinger difference between the unscaled reference at these two variances by $Ch^2\sigma^{-3/2}$. At the larger variance v, the change in Gaussian observation from latent X to latent aX has relative entropy at most $(1-a)^2\mathbb{E}X^2/(2v)$ by comparing the conditional Gaussian location laws. Its Hellinger distance is at most Ch^2/σ. The observed law from X_h is the law of $aX + \sqrt{v}G$, so the triangle inequality proves

$$H^2(K_\sigma f_h, K_\sigma g) \le Ch^4/\sigma^3 \qquad (0 < \sigma \le 1). \tag{21}$$

For radius d, take h a small fixed multiple of d^2 when $\sigma \le d^2$, and of $\sigma^{3/4}d^{1/2}$ when $\sigma > d^2$. Equations (19) and (21) respectively verify the radius. This gives the two endpoint lower scales in either experiment.

4.5.2 A regular fixed-support alternative

Suppose $p \ge 2$. Keep all endpoints of the support fixed and vary the constant, linear slope, and each of the $p - 1$ positive jumps and knot locations in the convex potential

$$V(x) = c + ax + \sum_{i=1}^{p-1} b_i(x - t_i)_+.$$

On this open set of $2p$ parameters, the derivative of the map to the first three moments in orders $0, 1, 2$ has rank three. If the derivative had a nonzero row relation, the corresponding polynomial P would have degree at most two. Varying c, a gives $\int Pg = \int xPg = 0$, while varying the jump and knot gives $\int_{t_i}^{B}(x - t_i)Pg = \int_{t_i}^{B} Pg = 0$ for each knot, with B the right support endpoint possibly infinite. These conditions subtract to zero zeroth and first moments of Pg on every affine piece. There must be two sign changes of P on each of the p open pieces, which is incompatible with degree at most two.

Thus the first three moment level set is a smooth manifold of positive dimension near the reference. Choose a smooth parameter curve tangent to a nonzero vector of its tangent space. The derivative of the potential along the curve is nonzero as a function up to null sets. Indeed, vanishing would first force each knot displacement derivative to vanish, from its step at that knot with nonzero jump coefficient. The remaining affine and hinge terms are linearly independent on the support. A bounded compactly supported smooth test therefore has a nonzero first variation against the density. In particular this standardized fixed-support family satisfies $\mathrm{TV}(f_h, g) \geq ch$.

On the common support, $|\log(f_h/g)| \leq Ch(1 + |x|)$, and the reference has exponential tails. Consequently $H(f_h, g) \leq Ch$, by squaring the change in the square-root density and integrating. Gaussian convolution only decreases the distance. This proves a lower total variation bound of order d at observed Hellinger radius d, uniformly in the noise level. A full-line reference necessarily has $p \geq 2$, so this covers the remaining Hellinger lower cases.

For testing, all lower families can be taken at observed radius $cN^{-1/2}$ with the fixed factor $c > 0$ arbitrarily small. The product Hellinger distance is then at most c. Product total variation is therefore smaller than one minus the sum of the specified error bounds for sufficiently small c. Two-point testing gives the lower bounds of Theorem 4.2. This also applies at $\sigma = 0$, even when the alternative has nonzero mass outside the reference support.

5 Testing the uniform ball under Gaussian noise

Fix a Euclidean ball scaled to have covariance identity. Among isotropic log-concave laws, total variation from its uniform law is controlled by the mass outside the ball. With known Gaussian error of standard deviation σ, the sharp testing separation is $\max\{N^{-1}, (\sigma^3/N)^{1/4}\}$, uniformly down to zero noise in each fixed dimension. The upper test applies to all isotropic log-concave alternatives. It uses only observed radii when alternatives are centrally symmetric. The matching lower bound uses a Gaussian perturbation at the boundary and a heat information estimate for uniform laws on convex bodies.

5.1 Results

Fix an integer $d \geq 2$, put $R = \sqrt{d+2}$, and let $B = \{x \in \mathbb{R}^d : |x| \leq R\}$. Write $g = \mathbf{1}_B/|B|$, where $|B|$ is its volume. For an alternative density f, write $m(f) = \int_{B^c} f$. The alternatives are log-concave probability densities on $\mathbb{R}^d$ with mean zero and covariance I_d.

For $\sigma > 0$ let γ_σ be the Gaussian density of covariance $\sigma^2 I_d$ and write $K_\sigma f = f * \gamma_\sigma$, with $K_0 f = f$. For densities a, b use $H^2(a, b) = \int(\sqrt{a} - \sqrt{b})^2$ and $\mathrm{TV}(a, b) = \frac{1}{2}\int |a - b|$.

Theorem 5.1. *There is a constant $C_d < \infty$ such that $\mathrm{TV}(f, g) \leq C_d m(f)$ for every alternative. Uniformly over $0 \leq \sigma \leq 1$ and sufficiently small $\epsilon > 0$, the sharp order of the supremum of $\mathrm{TV}(f, g)$ subject to $H(K_\sigma f, K_\sigma g) \leq \epsilon$ is*

$$\max\{\epsilon^2, \sigma^{3/4}\epsilon^{1/2}\}.$$

For any fixed positive test error bounds whose sum is less than one, the sharp total variation separation from g with N observations from $K_\sigma f$ is

$$\max\{N^{-1}, (\sigma^3/N)^{1/4}\}.$$

For centrally symmetric alternatives the test depends only on the observed radii. All constants may depend on d and, for testing, the fixed errors.

This is a multivariate boundary case of the piecewise log-affine testing phenomenon in Section 4. The Archive Gaussian inverse theorem [3] uses globally positive reference curvature and observes a different inverse order. Here the uniform reference has a hard boundary. Existing noisy-support results such as [12] concern estimating unknown support. The test here takes the ball to be known and measures total variation of the latent law. Our alternatives can have arbitrary angular variation consistent with isotropy and log-concavity.

5.2 A weighted convexity estimate

Throughout the proof, dimensional constants may change between occurrences. We use the uniform fixed-dimensional density and exponential tail bounds for the isotropic log-concave class, and its uniform lower density bound in a small ball around the origin [13, Theorem 5]. The same bounds give weak compactness, with covariance preserved under weak limits by the exponential tails.

We first prove an estimate for a nonnegative weight w that is bounded and Lipschitz on $\mathbb{R}^d$, is positive on the open ball B°, and has log-concave restriction $w\mathbf{1}_B$. We will only use weights for which w and ∇w are either compactly supported or rapidly decreasing. For a smooth positive isotropic log-concave density $f = e^{-V}$, set

$$Z = \int wf, \qquad \nu(dx) = Z^{-1}w(x)f(x)\,dx, \qquad \mu = \int x\,d\nu.$$

Suppose Z is bounded below, $Z \le |B|^{-1}\int_B w + C\eta$, and

$$\left|\int f\nabla w\right| + \left|\int f(dw + x\cdot\nabla w)\right| \le C\eta. \tag{22}$$

Assume that on a ball of fixed positive radius about μ, contained in B°, both f and w are bounded below by fixed positive constants. Then

$$f(x) \le |B|^{-1} + C'\eta \qquad (x \in B). \tag{23}$$

To prove the estimate, integration by parts gives

$$a = \mathbb{E}_\nu \nabla V = Z^{-1}\int f\nabla w, \qquad b = \mathbb{E}_\nu X\cdot\nabla V = Z^{-1}\int f(dw + x\cdot\nabla w).$$

These identities are immediate for the smooth positive densities produced by Gaussian regularization of an isotropic law and scalar standardization. Their score has at most affine growth, by the Gaussian conditional variance bound or the Lipschitz bound on the negative log density Hessian. Cutoffs then justify integration by parts for either weight. We work with these regularized densities and pass to the limit below.

Let $u = \nabla V(\mu)$ and $U(x) = V(x) - V(\mu) - u\cdot(x-\mu)$. Then $U \ge 0$. Since $\mu = \mathbb{E}_\nu X$, convexity yields

$$\mathbb{E}_\nu U \le \mathbb{E}_\nu (X-\mu)\cdot\nabla U = b - \mu\cdot a \le C\eta.$$

The lower bounds on f, w on the core imply that the Lebesgue integral of U there is $O(\eta)$. Convexity and the average over centered balls show that $U \le C\eta$ on a concentric ball of half the radius. For any x with $\nabla U(x) \ne 0$, choose y on this smaller ball in the direction of $\nabla U(x)$ from μ, at a fixed radius $r_0 > 0$. The supporting inequality at x gives

$$(x-\mu)\cdot\nabla U(x) \ge U(x) - U(y) + (y-\mu)\cdot\nabla U(x) \ge r_0|\nabla U(x)| - C\eta.$$

Thus $\mathbb{E}_\nu|\nabla U| \le C\eta$, and $|u| = |a - \mathbb{E}_\nu \nabla U| \le C\eta$. In particular,

$$\int w|\nabla f| = Z\mathbb{E}_\nu|\nabla V| \le C\eta.$$

Apply the L^1 Poincaré inequality for the fixed log-concave probability measure proportional to $w\mathbf{1}_B$. It follows from the Cheeger inequality for a log-concave probability [14], by applying it to the two parts about a median and using coarea. Replacing the median by the mean costs a factor at most two. Hence

$$\int_B w|f - c_B| \le C\eta, \qquad c_B = \frac{\int_B wf}{\int_B w} \le |B|^{-1} + C\eta.$$

On the smaller core we have $U = O(\eta)$ and $|u| = O(\eta)$, so $f(x) = f(\mu) + O(\eta)$ there by the uniform density bound. The core has weighted mass bounded below. It follows that $f(\mu) \le |B|^{-1} + C\eta$. The tangent inequality for V on all of B then gives $f(x) \le f(\mu)e^{-u\cdot(x-\mu)} \le |B|^{-1} + C'\eta$, as claimed.

5.3 Escape mass and bounded measurements

Let $m = \int_{B^c} f$. We have

$$\int_{B^c} (1 + |x|^2) f(x)\, dx \le C_d m. \tag{24}$$

For a short proof, fix $R_1 > \max(R, 3)$. Each standardized coordinate marginal is log-concave. Chebyshev's inequality and a fixed logarithmic survival secant from -2 to R_1 give that its right survival is at most Cme^{-ct} at $R_1 + t$. Use the sign-reflected marginal for the left survival. Indeed the tail at R_1 or $-R_1$ already has mass at most m. A union bound for the coordinates gives the same exponential tail, with dimensional constants, for $|X|$ beyond $\sqrt{d}R_1$. Integrating the resulting second moment proves (24).

Use $w_0(x) = (R^2 - |x|^2)_+$ in the weighted estimate. The exact mass and trace constraints give $Z = R^2 - d + O(m) = 2 + O(m)$, with $Z \ge 2$, and

$$\int f\nabla w_0 = -2\int_B xf(x)\, dx = O(m),$$
$$\int f(dw_0 + x\cdot\nabla w_0) = \int_B [dR^2 - (d+2)|x|^2] f(x)\, dx = O(m).$$

The last integrals on the full space would be zero by centering and $R^2 = d + 2$, and the exterior errors are bounded by (24). The probability density $w_0 f/Z$ is log-concave, supported in B and uniformly bounded. Each of its one-dimensional marginals is bounded above by a dimensional constant because the cross-sections of B have bounded volume. Thus its covariance is bounded below by a positive multiple of the identity: a density bounded by L puts mass at most $1/2$ within distance $1/(4L)$ of its mean. The covariance is also bounded above since the support is bounded. After whitening about its own barycenter μ, apply the lower density bound on the isotropic core from [13]. Transforming back shows that $w_0 f/Z$ is bounded below on a fixed-radius ball about μ, contained in B°. The individual lower bounds for w_0, f follow from their uniform upper bounds. So (23) with $\eta = m$ yields $f \le |B|^{-1} + Cm$ on B. The positive discrepancy inside B is $O(m)$, and outside it equals the exterior mass. By mass balance this proves $\mathrm{TV}(f, g) \le Cm$.

For measurements after noise use the different weight

$$w_1(x) = (R^2 - |x|^2)^2 e^{-\alpha|x|^2}, \qquad 0 < \alpha < 1/2$$

fixed. It is nonnegative on the whole space, vanishes on the sphere, and has log-concave restriction to B. For each component of

$$w_1, \qquad x w_1, \qquad \nabla w_1, \qquad d w_1 + x \cdot \nabla w_1$$

there is a uniformly bounded observed function whose conditional expectation given the latent input is that component, for $0 \le \sigma \le 1$. Indeed, the Fourier transform is a polynomial times a Gaussian with exponent $-|\xi|^2/(4\alpha)$; divide by $e^{-\sigma^2|\xi|^2/2}$. The remaining polynomial times Gaussian is integrable uniformly by $\alpha < 1/2$, so its inverse is bounded. It preserves radiality for scalar radial components and odd parity for the vector components.

Let D be the largest absolute discrepancy of these decoded expectations from their values under g. For g, the two derivative integrals in (22) vanish by the divergence theorem, because $w_1 = 0$ on the sphere. The mean integral $\int x w_1 g$ vanishes by symmetry. For f with small D, we therefore have $Z = |B|^{-1} \int_B w_1 + O(D)$, $\mu = O(D)$, and (22) with $\eta = D$. The standard isotropic core around zero provides the needed bounds near μ for this weight. It follows that

$$(f - g)_+ \le CD \quad \text{on } B. \tag{25}$$

When alternatives are centrally symmetric, their two vector expectations vanish exactly, so the two scalar radial observations suffice.

To pass both conclusions to an arbitrary log-concave density, smooth by an independent centered Gaussian of variance tending to zero and restore isotropy by scalar standardization. This preserves centering, log-concavity and, when imposed, symmetry. The resulting densities converge in L^1 by continuity of translation, convolution approximation and dilation in L^1. Their observed measurement differences, exterior masses and total variation converge. The positive-part bound (25) passes almost everywhere along a subsequence. This completes both comparisons.

5.4 A spherical boundary measurement

For $b > 0$ and $r \ge 0$ let

$$Q_b(r) = \mathbb{P}(|r e_1 + bG| \ge R), \qquad J_b(r) = \int_{\partial B} \gamma_b(r e_1 - z)\, dS(z).$$

Here G is standard Gaussian in $\mathbb{R}^d$ and dS is surface measure. The second function is the density at radius R of $|r e_1 + bG|$. Set $c_b = Q_b(R)/J_b(R)$ and $\psi_b(r) = Q_b(r) - c_b J_b(r)$. We show, uniformly over $0 < b \le 2$, that

$$\begin{aligned} c_b &\asymp b, \qquad \psi_b(r) \le 0 \quad (r \le R), \\ \psi_b(r) &\ge c \min\{(r - R)/b, 1\} \quad (r \ge R). \end{aligned} \tag{26}$$

The Gaussian radius density at $s > 0$ with input radius r is

$$k_b(s \mid r) = (2\pi b^2)^{-d/2} s^{d-1} e^{-(r^2+s^2)/(2b^2)} M(rs/b^2), \qquad M(t) = \int_{S^{d-1}} e^{t\omega_1}\, dS(\omega).$$

The function $A = M'/M$ is nonnegative and increasing for $t \ge 0$; its derivative is the variance of ω_1 under the corresponding exponential tilt. It is positive for $t > 0$ and less than one. For $s > R$,

$$\partial_r \log \frac{k_b(s \mid r)}{k_b(R \mid r)} = b^{-2}[sA(rs/b^2) - RA(rR/b^2)] \ge b^{-2}(s - R)A(rR/b^2).$$

Hence $Q_b/J_b = \int_R^\infty k_b(s \mid r)/k_b(R \mid r)\, ds$ is increasing and ψ_b has the correct sign.

The elementary spherical integral estimate $M(t) \asymp e^t(1+t)^{-(d-1)/2}$ for $t \geq 0$ follows by integrating near $\omega_1 = 1$: in polar angle θ, $1 - \cos\theta \asymp \theta^2$ and the surface factor is $\sin^{d-2}\theta$. The same bounds by integration over a cap and its complement give uniform constants. In particular $k_b(s \mid r) \asymp 1/b$ for $r, s \in [R, R+b]$ and $b \leq 2$. Also $Q_b(R) \geq 1/2$, since a positive outward normal component already exits the ball. This proves $c_b \asymp b$. For $r \in [R, R+b]$, the derivative of Q_b/J_b is bounded below by a positive constant, by integrating the displayed logarithmic derivative over $s \in [R, R+b]$ and using $A(rR/b^2) \geq c_d > 0$. Thus $\psi_b(r) \geq c(r-R)/b$ on that interval. On $[R, \infty)$, J_b is decreasing since its logarithmic derivative is $(-r + RA(rR/b^2))/b^2 < 0$. The survival Q_b is increasing: condition on the transverse Gaussian coordinates and use that $\mathbb{P}(|r+bG_1| \geq s)$ increases with $r \geq 0$ for any $s \geq 0$. Therefore ψ_b remains bounded below beyond $R+b$, proving (26).

The cap estimate also gives $\|J_a\|_\infty \leq C/a$ for $0 < a \leq 2$. For input radius comparable to R, use the estimate for M in the formula for $k_a(R \mid r)$. Away from that range the Gaussian factor $e^{-(r-R)^2/(2a^2)}$ absorbs the prefactor a^{-d}. Moreover $\int_{\mathbb{R}^d} J_a(|x|)\,dx = |\partial B|$.

For $b \geq 2\sigma$, let $a = (b^2 - \sigma^2)^{1/2}$ and define the observed function

$$\phi_b(y) = Q_a(|y|) - c_b J_a(|y|).$$

The Gaussian semigroup identity gives conditional expectation $\psi_b(|X|)$ for latent input X. The observed function is uniformly bounded for $0 < b \leq 2$, since $a \asymp b$. Under an alternative with escape mass m,

$$\mathbb{E}_f \phi_b(Y)^2 \leq C(b+m). \tag{27}$$

For the surface term, its Lebesgue squared L^2 norm, multiplied by c_b^2, is $O(b)$ by the uniform and integral bounds on J_a. The observed density is uniformly bounded. For the probability term, use $Q_a^2 \leq Q_a$ and the semigroup to get expectation at most $\mathbb{P}_f(|X + bG| \geq R) \leq m + Cb$. For the last step, the chance of crossing from inside is at most $\mathbb{P}(b|G| \geq R - |X|)$. The latent density cap and integration in spherical shells bound the interior contribution by $C\int_0^R \mathbb{P}(b|G| \geq t)\,dt \leq Cb$. In particular (27) is $O(b)$ under the null.

By (25), the adverse contribution of positive interior discrepancy to the mean difference is at most CDb. Indeed $\int_B |\psi_b(|x|)|\,dx \leq Cb$ by the same escape integral and the surface measure integral. The negative interior discrepancy is favorable because $\psi_b \leq 0$ there. The uniform density cap gives $\int_{R<|x|<R+t} f \leq Ct$ for $0 < t \leq 1$. Thus at least $m/2$ of the exterior mass lies at distance at least cm from the sphere, with a fixed sufficiently small c. Equation (26) bounds the exterior contribution from below, and we conclude

$$\mathbb{E}_f \phi_b(Y) - \mathbb{E}_g \phi_b(Y) \geq c \min\{m^2/b, m\} - CDb. \tag{28}$$

5.5 Inverse and testing upper bounds

For observed Hellinger radius ϵ, the bounded global measurements give $D \leq C\epsilon$. The absolute difference of the spherical observation means is at most $C\epsilon\sqrt{b+m}$ by Cauchy–Schwarz with the difference of square roots and (27). For testing, set $\epsilon = N^{-1/2}$. Use a preliminary test of the bounded global measurements, with thresholds of order ϵ about their null values. It controls the first type I error budget by Chebyshev and rejects with the required power when $D > C\epsilon$. For the spherical test reject if the observed sample mean exceeds its reference value by a sufficiently large constant times $\epsilon\sqrt{b}$. Under the null this controls the second type I budget by (27). Under an alternative with population mean difference greater than a sufficiently large constant times $\epsilon\sqrt{b+m}$, the same variance bound gives the required power. Thus for both upper bounds it suffices to consider $D \leq C\epsilon$ and a population boundary difference at most $C\epsilon\sqrt{b+m}$.

Choose $b = \max(2\sigma, A_0\epsilon^2)$ for a large fixed A_0, and take ϵ small enough that $b \le 2$. Equation (28) shows

$$\min\{m^2/b, m\} \le C\epsilon\sqrt{b+m} + CDb.$$

If $m \ge b$, then $CDb \le C'\epsilon m$ and the inequality implies $m \le C''\epsilon^2$ after absorption. This contradicts $m \ge b \ge A_0\epsilon^2$ when A_0 is sufficiently large. In the other case $m^2/b \le C\epsilon\sqrt{b}$, using $D \le C\epsilon$ and $b \le 2$. It follows that

$$m \le C\epsilon^{1/2}b^{3/4} \le C' \max\{\epsilon^2, \sigma^{3/4}\epsilon^{1/2}\}.$$

The escape mass bound proves the asserted total variation estimates.

5.6 A boundary lower bound for convex bodies

The lower bounds hold for uniform laws on arbitrary fixed isotropic convex bodies. We first establish the heat information bound used in the construction. Let $K \subset \mathbb{R}^d$ be a compact convex body with nonempty interior and with the origin in its interior. Write $q_b = |K|^{-1}\mathbf{1}_K * \gamma_b$, where b is a Gaussian standard deviation. For $0 < b \le \sqrt{2}$, we claim

$$\int_{\mathbb{R}^d} \frac{(\Delta q_b)^2}{q_b} \le C_K b^{-3}. \tag{29}$$

Choose a fixed inner ball $B(0, r_0) \subset K$. For each $y \in K$, the convex combination of y with this ball using weight proportional to b contains a ball of radius comparable to b within distance $O(b)$ of y, contained in K. It follows that q_b is bounded below uniformly on K. Its inner strip of width t has volume $O_K(t)$: the body $(1 - t/r_0)K$ is at distance at least t from the boundary when $0 < t \le r_0$, by the same convex combination. For large t use bounded volume.

For $y \in K$ use $\Delta q_b(y) = -|K|^{-1}\int_{K^c} \Delta\gamma_b(y - z)\, dz$ because the Gaussian Laplacian integrates to zero. Its absolute value is at most

$$C_K b^{-2}\mathbb{E}\big[(|G|^2 + d)\mathbf{1}_{|G| \ge \mathrm{dist}(y,\partial K)/b}\big].$$

The lower bound on q_b, the strip volume and the Gaussian tail bound control the integral over K by $C_K b^{-3}$. For example, sum over strips with distance to the boundary in $[jb, (j+1)b]$: their volumes are at most $C_K(j+1)b$ and the expectation decays as Ce^{-cj^2}.

For the exterior apply conditional Jensen to $\Delta q_b/q_b$, written as the conditional expectation of $b^{-2}(|G|^2 - d)$ given $X + bG$, with X uniform on K. The contribution to (29) is at most

$$b^{-4}\mathbb{E}\left[(|G|^2 - d)^2\mathbf{1}_{X+bG\notin K}\right].$$

Conditionally on G, an exit requires that X lie within distance $b|G|$ of the boundary. This event has probability at most $C_K b|G|$ by the strip estimate. This again gives $C_K b^{-3}$, proving (29).

For the ball the leading constant of this heat information has a normal-profile formula:

$$\lim_{b\downarrow 0} b^3 \int \frac{(\Delta q_b)^2}{q_b} = \frac{|\partial B|}{|B|}\int_{\mathbb{R}} \frac{t^2\varphi(t)^2}{\Phi(-t)}\, dt, \tag{30}$$

where φ, Φ are standard Gaussian density and CDF. The two estimates just given show that the scaled integral outside the tube $\mathrm{dist}(y, \partial B) \le Tb$ is bounded by a quantity tending to zero as $T \to \infty$, uniformly in small b. For the inside estimates keep only strips with $j \ge T - 1$. For the outside estimates, any crossing from $X \in B$ to that exterior region has $|G| > T$, and apply the same conditional strip bound with this extra restriction. In the tube write $y = (R + bt)u$, with $|u| = 1$.

The event $y - bG \in B$ converges for almost every G to $G \cdot u \geq t$. Therefore $q_b(y) \to |B|^{-1}\Phi(-t)$ and

$$b^2 \Delta q_b(y) \longrightarrow |B|^{-1}\mathbb{E}[(|G|^2 - d)\mathbf{1}_{G\cdot u \geq t}] = |B|^{-1} t\varphi(t).$$

On fixed bounded t intervals the denominator is uniformly positive for small b, since radial ball probabilities decrease with the distance of the Gaussian center from the origin. The numerator is uniformly bounded after scaling by the corresponding unconditioned Gaussian moment. The volume element in this parametrization is $b(1 + bt/R)^{d-1} dt\, dS$, with surface measure on ∂B. Dominated convergence on bounded tubes followed by the uniform tube estimate proves (30).

Now let X be uniform on a fixed mean-zero convex body K whose uniform covariance is I_d, and use

$$X_h = a(X + hG'), \qquad a = (1 + h^2)^{-1/2},$$

with independent standard Gaussian G'. This is isotropic and log-concave. The Gaussian jitter puts at least order h mass outside K. To see this, take inputs in the homothetic shell $K \setminus (1-h)K$, which has volume of order h. Write any such point as $(1 - t)z$ for a boundary point z and $0 \leq t \leq h$. A unit supporting normal n at z gives a fixed positive probability of exit, since crossing its supporting hyperplane requires only

$$G' \cdot n > [(a^{-1} - 1) + t]\langle z, n\rangle / h \leq C_K.$$

It follows that $\mathrm{TV}(f_h, g) \geq c_K h$ for the reference density g uniform on K. In the clean experiment $\mathrm{TV}(f_h, g) \leq C_K h$: translation of the indicator of a convex body changes its L^1 norm by at most perimeter times the translation length, and shrinking by a costs $O(h^2)$ by the homothetic volume formula. Thus the observed squared Hellinger distance is at most $C_K h$, by contraction.

For $\sigma > 0$ the alternative observation law is that of $aX + \sqrt{v}G$, with $v = \sigma^2 + a^2h^2$. The bound (29) gives

$$\left\| \partial_v \sqrt{q_{\sqrt{v}}} \right\|_2 \leq C_K v^{-3/4}$$

by the heat equation. Integration from σ^2 to v bounds the Hellinger change of the unscaled reference by $C_K h^2 \sigma^{-3/2}$. The affine correction at variance v costs at most $C_d h^2/\sigma$ in Hellinger distance by comparing the conditional Gaussian means given the same X in the two experiments and using relative entropy. Hence

$$H^2(K_\sigma f_h, K_\sigma g) \leq C_K \min\{h, h^4/\sigma^3\} \quad (0 < \sigma \leq 1).$$

If $\sigma \leq \epsilon^2$, take h a sufficiently small multiple of ϵ^2, also allowing $\sigma = 0$ by the clean bound. Otherwise take h a small multiple of $\sigma^{3/4}\epsilon^{1/2}$. This proves the Hellinger lower bounds in Theorem 5.1 at the ball, and the same endpoint lower for every isotropic convex body. At the ball the alternatives used here are radial. For testing take ϵ a sufficiently small multiple of $N^{-1/2}$. The product Hellinger bound and the two-point testing inequality give the claimed lower separation at any prescribed positive error bounds whose sum is less than one.

6 Exterior mass and testing for the isotropic uniform cube

For any isotropic log-concave law in a fixed dimension, its total variation from the isotropic uniform cube is bounded by a constant times its mass outside the cube. The proof uses equality in the Prékopa–Leindler inequality and a compactness argument for normalized differences of log densities. Coordinatewise testing yields sharp total variation separation after Gaussian observation errors, uniformly down to the clean experiment. The same comparison gives the sharp Hellinger inverse modulus with arbitrary known Gaussian covariance.

6.1 Results

Fix $d \geq 2$ and put $I = [-\sqrt{3}, \sqrt{3}]$, $K = I^d$. The density $g = \mathbf{1}_K/|K|$ has mean zero and covariance I_d. We call a probability density isotropic when its mean is zero and its covariance is I_d. Let $g_1 = \mathbf{1}_I/|I|$ be the common coordinate marginal of g, and let f_i denote the ith coordinate marginal of f. The total variation distance is $\mathrm{TV}(f, g) = \frac{1}{2}\int |f - g|$.

Theorem 6.1. *There is $C_d < \infty$ such that every isotropic log-concave density f on $\mathbb{R}^d$ satisfies*

$$\mathrm{TV}(f, g) \leq C_d \int_{K^c} f \leq C_d \sum_{i=1}^{d} \mathrm{TV}(f_i, g_1).$$

In particular, the uniform law is the unique isotropic log-concave law supported on K.

Let the observation law be $K_\Sigma f$, the convolution of f and an independent Gaussian with known positive semidefinite covariance Σ. Write $\sigma_i = \sqrt{\Sigma_{ii}} \in [0, 1]$, allowing exact clean coordinates. Put $s = \max_i \sigma_i$ and $H^2(a, b) = \int(\sqrt{a} - \sqrt{b})^2$.

Theorem 6.2. *The sharp order of the supremum of $\mathrm{TV}(f, g)$ subject to $H(K_\Sigma f, K_\Sigma g) \leq \epsilon$ is $\max\{\epsilon^2, s^{3/4}\epsilon^{1/2}\}$, uniformly in Σ for small ϵ. With N independent observations, the fixed-reference testing separation in latent total variation is*

$$\max\{N^{-1}, (s^3/N)^{1/4}\},$$

with constants depending on dimension and prescribed positive error bounds whose sum is less than one.

Unlike the uniform Euclidean ball comparison in Section 5, the cube theorem uses the exact one-dimensional marginals at a zero-escape limit. A normalized difference in the interior is concave. Its coordinate marginals exclude both a signed boundary deficit and any nonzero concave interior difference. This supplies a linear bound without a quantitative version of the Prékopa–Leindler equality theorem. The testing results then follow from the scalar boundary experiment Section 4. The geometry is different from Gaussian deconvolution at positively curved reference potentials [3].

6.2 Rigidity at zero exterior mass

Let f be isotropic, log-concave, and supported on K. Each coordinate marginal is log-concave by Prékopa's theorem, is supported on I, and has mean zero and variance one. The one-piece exterior mass inequality of Section 4 shows that it is exactly g_1.

For clarity, suppose the first coordinate marginal is g_1. On the open interval I°, the slices $z \mapsto f(t, z)$ have the same positive integral. The marginal is constant there, using its log-concave continuous representative. For any $t_0, t_1 \in I^\circ$ and $0 < \lambda < 1$, log-concavity compares the two slices at t_0, t_1 with the slice at $(1-\lambda)t_0 + \lambda t_1$ by the Prékopa–Leindler inequality. Since their integrals are equal, it is an equality case. Dubuc's equality description, as formulated in [15, Theorem 1.1], states here that the normalized slices are translates of the same density, with the intermediate translation equal to the affine combination of the other two. Consequently the conditional barycenter $M(t) = \mathbb{E}[X_{-1} \mid X_1 = t]$ is affine in t. The slope equals the cross covariance between X_{-1} and X_1, since $\mathrm{Var}(X_1) = 1$, so it is zero. The normalized slices therefore coincide, and the first coordinate is independent of the remaining coordinates. Repeating the same argument on the remaining isotropic log-concave marginal gives the product distribution, which is g.

6.3 A linear exterior mass estimate

Suppose the first inequality in Theorem 6.1 fails. By the preceding rigidity this yields a sequence of isotropic log-concave densities f_j such that

$$a_j = \mathrm{TV}(f_j, g) > 0, \qquad m_j = \int_{K^c} f_j = o(a_j).$$

In particular $m_j \to 0$. The isotropic log-concave class has uniform exponential tails in fixed dimension [13], so it is weakly compact with preserved first and second moments. Any weak limit of a subsequence is isotropic, log-concave and supported on K, so it is g. Along a subsequence we therefore have L^1 convergence by Proposition 2 of [16]. The convex-function argument in their proof also gives locally uniform convergence on K°. In particular $a_j \to 0$.

On fixed compact subsets of K° put $L_j = \log(f_j/g)$, defined there for all large j. The function is concave, and its L^1 norm on each such subset is $O(a_j)$, since f_j is bounded below there and $\int |f_j - g| = 2a_j$. We use the following elementary consequence. On a smaller compact subset, $a_j^{-1} L_j$ is uniformly bounded and equicontinuous. For the convex function $-L_j$, a bound above follows from Jensen over small centered balls contained in a larger compact subset. For a lower bound, if its value at a point is $-M$, convexity bounds it by $-M/2 + O(a_j)$ on the points halfway from that point to a fixed small ball around it, using the upper bound just proved. The L^1 norm then forces $M = O(a_j)$ on the smaller subset. Secant slopes inside a still larger compact subset give the local Lipschitz bound. By diagonal subsequence selection, L_j/a_j converges locally uniformly to a concave function L on K°.

A supergradient of L_j at the origin has norm $O(a_j)$ by the local slope bound. Its supporting affine function bounds L_j above on all of K°, with the convention $\log 0 = -\infty$. Since K is bounded, this yields the global bound

$$\frac{(f_j - g)_+}{a_j} \le C_d \quad \text{on } K.$$

On K° the normalized density difference converges locally uniformly to gL, since $f_j = g e^{L_j}$ there. Also the normalized total positive mass inside K is $1 - m_j/a_j \to 1$. Dominated convergence therefore gives

$$\int_K g L_+ = 1. \tag{31}$$

Consider the signed measures $(f_j - g)\,dx/a_j$, of total variation norm two. Their mass outside K tends to zero in absolute value. By passing to a further subsequence the restrictions to K converge weakly to a signed finite measure ζ. Its restriction to K° is $gL\,dx$ by local uniform convergence. There is no positive boundary part, by the uniform bound on the normalized positive density inside K. Thus

$$\zeta = gL\,dx - \nu,$$

where ν is a nonnegative measure supported on ∂K and L is integrable on K.

For each coordinate, the scalar one-piece bound gives

$$\mathrm{TV}(f_{j,i}, g_1) \le C \int_{I^c} f_{j,i} \le C m_j.$$

Therefore the coordinate projections of ζ are all zero signed measures. The projections of its absolutely continuous part have no atoms at the endpoints of I. It follows that ν assigns zero mass to both faces of K perpendicular to every coordinate axis. These faces cover ∂K, so $\nu = 0$. In particular the function L has zero marginal integral along each coordinate.

We finish by showing that an integrable concave function with all those zero marginals must vanish. For each i, the distributional second derivative $\partial_{ii}L$ is a nonpositive measure on K°. Test it against $\varphi(x_i)\chi_n(x_{-i})$, where φ is nonnegative, smooth and compactly supported in I°, and $0 \le \chi_n \le 1$ are smooth compactly supported functions increasing to one in the other coordinates. The value tends to zero by integrability of L and the zero marginal integral:

$$\int L(x)\varphi''(x_i)\chi_n(x_{-i})\,dx \longrightarrow \int L(x)\varphi''(x_i)\,dx = 0.$$

By monotone convergence for the nonnegative measure $-\partial_{ii}L$, this forces $\partial_{ii}L = 0$. Since the distributional Hessian of a concave function is negative semidefinite, all its entries now vanish. For instance apply negativity of the directional second derivative in the directions $e_i \pm e_k$ to eliminate the off-diagonal entries. Thus L is affine. Each of its coordinate marginals is zero, so the affine coefficients and constant are zero. This contradicts (31). The first inequality of Theorem 6.1 follows. The second follows from the union bound on coordinate escape events and the fact that the reference coordinate law has no exterior mass.

6.4 Coordinate tests and noisy inverse bounds

The Hellinger distance cannot increase under coordinate projection. Hence an observed discrepancy at most ϵ implies the same bound in each scalar experiment $f_i * \gamma_{\sigma_i}$ versus $g_1 * \gamma_{\sigma_i}$. The one-piece result of Section 4 gives

$$\mathrm{TV}(f_i, g_1) \le C \max\{\epsilon^2, \sigma_i^{3/4}\epsilon^{1/2}\}.$$

Sum and use Theorem 6.1 to prove the upper inverse bound.

For testing, apply the same scalar tests coordinate by coordinate. Allocate the type I budget among the fixed d tests and reject if any of them rejects. The observations within each scalar sample are independent across the N observations. The union bound for type I error does not require independence between coordinates. If the joint latent total variation is above a sufficiently large multiple of the stated separation, Theorem 6.1 gives a coordinate above its scalar threshold. Its scalar test controls the type II error at the required level.

6.4.1 A directional lower bound

For the lower bounds, let v be a unit eigenvector for the largest eigenvalue τ^2 of Σ. If $\Sigma = 0$, take any unit v. We have $s \le \tau \le \sqrt{d}\,s$. For X uniform on K and an independent scalar standard Gaussian G, consider

$$X_h = T_h(X + hGv), \quad T_h = I_d + (a-1)vv^\top, \quad a = (1+h^2)^{-1/2}.$$

The alternative is log-concave and exactly isotropic.

The following estimates hold uniformly over unit v for every fixed isotropic convex body K. Disintegrate the uniform law over $z \in v^\perp$. At a fixed z, its conditional distribution along v is uniform on a chord of length $\ell(z)$, and the marginal density of z is $\ell(z)/|K|$. For z in a fixed small ball in $v^\perp$ inside the projection of an inner ball of K, the chords have lengths bounded below by a fixed positive constant. A strip of width proportional to h at the right endpoint of each such chord exits under X_h with uniformly positive probability, since the correction $a - 1 = O(h^2)$ moves that coordinate by only $O(h^2)$. The transverse coordinates are unchanged. The strips have total mass of order h, so $\mathrm{TV}(f_h, g) \ge ch$. The clean distance is also at most Ch, by the bounded variation translation estimate for the uniform density and an $O(h^2)$ change in total variation under

T_h. For the latter, T_hK lies between $(1-Ch^2)K$ and $(1+Ch^2)K$ because K contains a fixed inner ball and $T_h - I_d = O(h^2)$. Use the volume and determinant changes to bound the L^1 difference. Consequently the clean squared Hellinger distance is $O(h)$, and this bound holds under any noise by contraction.

For $\tau > 0$, first retain only the noise of variance τ^2 along v. The remaining Gaussian covariance $\Sigma - \tau^2 vv^\top$ is positive semidefinite, so adding that noise only decreases Hellinger distance. The projected marginal in $v^\perp$ is unchanged in the rank-one noise experiment. We use the one-dimensional heat bound uniformly over its chords. For a uniform chord of length $\ell > 0$ let q_b be its convolution with the Gaussian of standard deviation $b > 0$. If $\ell \geq b$, then

$$\int (q_b'')^2/q_b \leq C\ell^{-1}b^{-3}.$$

Indeed q_b'' is the signed difference of the two endpoint Gaussian derivatives divided by ℓ. Near either endpoint use the denominator bound obtained by integrating the Gaussian over an inner segment of length b adjacent to that endpoint. In scaled coordinate t the denominator is bounded below by $c\ell^{-1}\varphi(t)/(1+|t|)$, as in the scalar heat estimate of Section 4. Each endpoint then contributes at most $C\ell^{-1}b^{-3}$. If $\ell < b$, conditional Jensen applied to the Gaussian second-derivative polynomial gives $\int (q_b'')^2/q_b \leq 2b^{-4}$. In either case multiplying by ℓ gives an upper bound Cb^{-3}. Integrating against the transverse marginal, the projection area is bounded in fixed dimension. Thus the squared norm of the variance derivative of the square-root density in the rank-one experiment is at most Cb^{-3}.

The reference variance along v is τ^2, while the alternative is a latent scale change by a followed by Gaussian variance $\tau^2 + a^2h^2$ along v. Integrating the variance derivative bounds the unscaled Hellinger difference by $Ch^2\tau^{-3/2}$. The affine correction at the larger variance has Hellinger distance at most Ch^2/τ, by coupling the latent chord input with itself and comparing the two conditional Gaussian means. This combines to

$$H^2(K_\Sigma f_h, K_\Sigma g) \leq C\min\{h, h^4/\tau^3\} \quad (\tau > 0),$$

since $\tau \leq \sqrt{d}$, with the clean bound available if $\tau = 0$. Taking h a small multiple of $\max\{\epsilon^2, s^{3/4}\epsilon^{1/2}\}$ makes the Hellinger distance at most ϵ. It proves the inverse lower bound. At ϵ a small multiple of $N^{-1/2}$, the product Hellinger bound and two-point testing give the testing lower bound as well. This completes Theorem 6.2.

7 Gaussian covariance information at faces and edges

We compare two small Gaussian smoothings of a density with a convex support and a positive boundary trace. The leading squared Hellinger distance is a boundary integral of one-dimensional Gaussian profile discrepancies. The leading covariance Fisher information depends on the normal quadratic form of the covariance direction. For uniform boxes, covariance changes preserving all coordinate variances have no face contribution. Their first nonzero information comes from pairs of faces meeting at codimension-two edges. We give exact profile integrals at both orders.

7.1 Results

Let $K \subset \mathbb{R}^d$, $d \geq 2$, be a compact convex body with nonempty interior. Let g be a probability density supported on K, Lipschitz and bounded above and below by positive constants on K from the inside. For a positive semidefinite matrix A and $b > 0$, write $q_{b,A}$ for the density of $X + bA^{1/2}G$,

where $X \sim g$ and G is independent standard Gaussian. For $A = 0$ use $q_{b,0} = g$. Let $n(x)$ be the outward unit normal at almost every boundary point, and write $a_A(x) = \sqrt{n(x)^\top A n(x)}$. Use $H^2(p,q) = \int (\sqrt{p} - \sqrt{q})^2$.

For $a \geq 0$ define

$$P_a(t) = \Phi(-t/a) \quad (a > 0), \qquad P_0(t) = \mathbf{1}_{t<0},$$

where Φ is the standard Gaussian CDF. Write $\varphi = \Phi'$. For positive definite A and symmetric E define the Fisher information for the scalar parameter in $A + uE$ at $u = 0$ by

$$I_b(A; E) = \int \frac{(\partial_u q_{b,A+uE}|_{u=0})^2}{q_{b,A}} = \frac{b^4}{4} \int \frac{(E : \nabla^2 q_{b,A})^2}{q_{b,A}}.$$

Here $E : M = \operatorname{tr}(E^\top M)$.

Theorem 7.1. *For positive semidefinite A, C,*

$$\lim_{b \downarrow 0} \frac{H^2(q_{b,A}, q_{b,C})}{b} = \int_{\partial K} g(x) \int_{\mathbb{R}} (\sqrt{P_{a_A(x)}(t)} - \sqrt{P_{a_C(x)}(t)})^2 \, dt \, dS(x).$$

For positive definite A and symmetric E,

$$\lim_{b \downarrow 0} \frac{I_b(A; E)}{b} = \frac{J_1}{4} \int_{\partial K} g(x) \frac{(n(x)^\top E n(x))^2}{a_A(x)^3} \, dS(x), \qquad J_1 = \int_{\mathbb{R}} \frac{t^2 \varphi(t)^2}{\Phi(-t)} \, dt.$$

For the edge result take an axis-aligned box K and $g = \mathbf{1}_K / |K|$. A codimension-two face F is specified by distinct coordinates $i < j$ and their outward signs $s_i, s_j \in \{-1, 1\}$. Let A_F be the covariance of $(s_i Z_i, s_j Z_j)$ when Z has covariance A. For a positive definite two-by-two covariance M, write φ_M for its Gaussian density and

$$S_M(t_1, t_2) = \mathbb{P}(Z_1 \geq t_1, Z_2 \geq t_2), \qquad Z \sim N(0, M).$$

Use $|F|$ for its $(d-2)$-dimensional area, with counting measure when $d = 2$.

Theorem 7.2. *Suppose A, C are positive definite and have the same diagonal. Then*

$$\lim_{b \downarrow 0} \frac{H^2(q_{b,A}, q_{b,C})}{b^2} = \sum_F \frac{|F|}{|K|} \int_{\mathbb{R}^2} (\sqrt{S_{A_F}(t)} - \sqrt{S_{C_F}(t)})^2 \, dt.$$

The coefficient is positive exactly when $A \neq C$. For positive definite A and a symmetric covariance direction E with zero diagonal,

$$\lim_{b \downarrow 0} \frac{I_b(A; E)}{b^2} = \sum_F \frac{|F|}{|K|} E_{ij}^2 \int_{\mathbb{R}^2} \frac{\varphi_{A_F}(t)^2}{S_{A_F}(t)} \, dt.$$

In particular the leading Fisher form is diagonal in the off-diagonal covariance coordinates.

Brosse and Dalalyan [17] obtain local Gaussian density and derivative profiles near boundary and corner strata for isotropic smoothing. The results here concern integrated discrepancies between covariance shapes. They require negligible contributions away from the appropriate boundary strata, including regions where the smoothed density appears in a denominator. Classical results on heat content and perimeter include [18, 19]. At a uniform ball the scalar Fisher limit reduces to the heat information constant in Section 5. The role of a hard boundary differs from Gaussian inversion at uniformly curved log-concave references [3].

7.2 Concentration of the face contributions

Write F for a bounded Lipschitz extension of the values of g on K to the whole space. Such an extension follows from the usual infimum formula for a Lipschitz extension, followed by clipping. Here and below the actual density under convolution is $F\mathbf{1}_K$. The volume of the inner strip of K of width t is at most $C_K t$. To see this after translating an interior point to zero, choose a fixed inner ball of radius r_0; the homothetic body $(1-t/r_0)K$ lies at distance at least t from the boundary for $t \le r_0$, by convexity. For larger t use bounded volume.

We first localize the Hellinger integral. If $W = A^{1/2}G$, then for $y \in K$ write

$$q_{b,A}(y) - g(y) = \mathbb{E}[F(y - bW) - F(y)] - \mathbb{E}[F(y - bW)\mathbf{1}_{y-bW \notin K}].$$

The first term is $o(b)$ in $L^1(K)$. At almost every y the derivative of F exists, and its linear contribution has Gaussian expectation zero. The difference quotient is dominated by the Lipschitz constant times $\mathbb{E}|W|$, so dominated convergence gives the assertion. For the second term on inner points at distance more than Tb from the boundary, a nonzero indicator requires $|W| > T$. The strip volume bound, after integrating in y at each fixed W, bounds the integral by $Cb\mathbb{E}[|W|\mathbf{1}_{|W|>T}]$. This estimate holds even for a degenerate A. The same proof applies to C. Use $(\sqrt{p} - \sqrt{q})^2 \le |p - q|$ to bound the Hellinger integrand on these inner points. Outside K at distance more than Tb, simply bound by $q_{b,A} + q_{b,C}$. Conditionally on W, a displacement bW to that exterior region requires $|W| > T$ and a latent point within $b|W|$ of the boundary. The same strip estimate applies. In consequence the integral divided by b outside the Tb boundary tube has upper limit tending to zero as T tends to infinity.

For Fisher information assume A positive definite. Let

$$P_E(G) = G^\top A^{-1/2} E A^{-1/2} G - \operatorname{tr}(A^{-1}E).$$

Differentiation of the Gaussian density gives

$$b^2 E : \nabla^2 q_{b,A}(y) = \mathbb{E}[P_E(G) F(y - bA^{1/2}G)\mathbf{1}_{y-bA^{1/2}G \in K}].$$

The polynomial has mean zero. For $y \in K$, subtracting $F(y)$ isolates an $O(b)$ Lipschitz error and a missed Gaussian tail bounded by $C\exp(-c\operatorname{dist}(y, \partial K)^2/b^2)$. The smoothed density is bounded below uniformly on K: a convex combination of y with a fixed inner ball contains a ball of radius proportional to b, within $O(b)$ of y, on which the latent density is bounded below. The squared derivative integral divided by the density, after scaling by b^{-1} for the squared numerator just displayed, thus has interior contribution bounded by $O(b) + o_T(1)$ off the tube. This follows by summing over strips of integer multiples of b, using their volume bound and the Gaussian estimate.

On the exterior, write the derivative ratio $b^2(E : \nabla^2 q_{b,A})/q_{b,A}$ as the conditional expectation of $P_E(G)$ given $Y = X + bA^{1/2}G$. Conditional Jensen bounds its squared integral against $q_{b,A}$ off the exterior tube by

$$\mathbb{E}[P_E(G)^2 \mathbf{1}_{\operatorname{dist}(Y,K) > Tb}].$$

This is at most $Cb\mathbb{E}[P_E(G)^2 |W| \mathbf{1}_{|W|>T}]$ by the conditional strip argument. This proves localization for $I_b(A; E)/b$ as well.

7.3 Integration of the local face profiles

A compact convex body with interior has a finite cover of its boundary by Lipschitz graph charts. After choosing the vertical unit vector e for one chart pointing outward, the body is locally below a graph $x_d = \eta(x')$. Use a continuous partition of unity supported in smaller chart cylinders and equal

to one in sum near the boundary. It suffices by localization to integrate each chart contribution with the vertical coordinate $x_d = \eta(x') + bt$, $|t| \le T$, then send T to infinity. On a smaller cylinder, vertical distance and distance to the boundary are comparable near the graph, by its Lipschitz bound, so the omitted contribution tends uniformly to zero after scaling. The integration element in these coordinates is $b\,dx'\,dt$.

At a differentiability point $x = (x', \eta(x'))$ of the graph, put $s = t(n \cdot e)$. The condition $x + bte - bA^{1/2}G \in K$ converges to $n \cdot A^{1/2}G \ge s$ almost surely whenever the latter boundary has zero probability. This follows directly by the first-order graph expansion. If the normal variance vanishes, the statement holds for $t \neq 0$, which suffices in the integral. The Lipschitz assumption gives the inside density trace $g(x)$, and the uniform density bound therefore gives

$$q_{b,A}(x + bte) \longrightarrow g(x) P_{a_A(x)}(s).$$

The same holds for C.

For positive definite A, the polynomial formula and the same graph limit yield

$$b^2 E : \nabla^2 q_{b,A}(x + bte) \longrightarrow g(x)(n^\top E n) \frac{s}{a_A(x)^3} \varphi(s/a_A(x)).$$

To compute the coefficient, condition $P_E(G)$ on the standard normal $n \cdot A^{1/2}G/a_A(x)$. Its conditional value at z is $(n^\top E n)(z^2 - 1)/a_A(x)^2$. Integrating $(z^2 - 1)\varphi(z)$ above $s/a_A(x)$ gives $(s/a_A(x))\varphi(s/a_A(x))$.

For bounded t, the Gaussian polynomial numerator is uniformly bounded. In the positive definite case, the denominator is bounded below by a positive constant depending on T, using the inner ball argument also for points within $O(Tb)$ of K. Dominated convergence applies in each chart. The surface element is $dx'/(n \cdot e)$, and $ds = (n \cdot e)\,dt$, so the resulting element is $dS\,ds$. Then sum the charts and let T tend to infinity using the localization estimates. The Hellinger expression is the first limit in Theorem 7.1. For the Fisher expression, division by $g(x)P_{a_A(x)}(s)$ and the substitution $s = a_A(x)t$ give the stated factor $J_1/a_A(x)^3$ and the covariance derivative contributes the factor one quarter. This proves both limits.

7.4 Localization at edges when face variances agree

Write the box as a product of compact intervals of positive length. For positive definite matrices A, C with the same diagonal, put $A_u = (1 - u)A + uC$, $0 \le u \le 1$, and $E = C - A$. The path derivative is

$$\dot{q}_{b,A_u} = b^2 \sum_{i<j} E_{ij} \partial_{ij} q_{b,A_u}.$$

This follows by differentiating the Gaussian Fourier multiplier, using that the diagonal entries of E vanish. Let $\mathcal{E}$ be the union of the codimension-two faces. We prove, uniformly over the covariance path,

$$\int_{\operatorname{dist}(y,\mathcal{E}) > Tb} \frac{(\partial_{ij} q_{b,A_u}(y))^2}{q_{b,A_u}(y)}\,dy \le C b^{-2} \delta(T), \qquad \delta(T) \longrightarrow 0. \tag{32}$$

The same bound holds on the whole space with $\delta(T)$ replaced by a fixed constant. These estimates also apply to any single positive definite covariance.

Indeed, $\partial_{ij} \mathbf{1}_K$ is the signed sum of surface measures on the four codimension-two faces with coordinate indices i, j. The sign at a pair of outward signs s_i, s_j is $s_i s_j$, by differentiating the two

interval indicators. Write γ_{b,A_u} for the Gaussian density of covariance $b^2 A_u$. For one such face F and $z \in F$ define

$$K_z(y) = \int_0^{c_0 b} \int_0^{c_0 b} \gamma_{b,A_u}(y - z + r s_i e_i + s s_j e_j)\, dr\, ds,$$

where $c_0 > 0$ is fixed. The corresponding inner collar is contained in K for all sufficiently small b, and its volume element is $dr\, ds\, dS(z)$. Thus $q_{b,A_u}(y) \geq |K|^{-1} \int_F K_z(y)\, dS(z)$. With $w = (y - z)/b$, integrate only to $c_0 b/(1 + |w|)$ in each inward coordinate. Uniform ellipticity of A_u makes the Gaussian density on these displacements comparable from below to its value at $y - z$. Hence

$$K_z(y) \geq c b^2 \gamma_{b,A_u}(y - z)/(1 + |w|)^2.$$

Cauchy–Schwarz in the surface integral now bounds the squared Gaussian convolution of that face divided by q_{b,A_u} by

$$C \int_F \frac{\gamma_{b,A_u}(y - z)^2}{K_z(y)}\, dS(z) \leq C b^{-2} \int_F (1 + |(y - z)/b|)^2 \gamma_{b,A_u}(y - z)\, dS(z).$$

On the region in (32) the scaled distance from every point of every face exceeds T. Integrating in y and using Gaussian tails proves (32), including the four-face sum.

By the derivative formula and (32), $\|\partial_u \sqrt{q_{b,A_u}}\|_2^2$ is $O(b^2)$, and is $O(b^2 \delta(T))$ when restricted outside the edge tube. Apply Minkowski's inequality to the integral of the square-root derivative over $u \in [0, 1]$. The Hellinger integral outside the tube is $O(b^2 \delta(T))$. For the Fisher integral with a zero-diagonal direction E, its exact directional formula and (32) give the same localization.

7.5 The integrated edge limits

Fix $T > 0$. At a face F of codimension two, use coordinates

$$y = z + b(s_i t_i e_i + s_j t_j e_j), \qquad z \in F, \quad (t_i, t_j) \in [-T, T]^2.$$

The volume element is $b^2\, dt_i\, dt_j\, dS(z)$. The union of these rectangular tubes covers the Tb neighborhood of $\mathcal{E}$ except for parts in tangential spill regions of volume $O_T(b^3)$, near intersections with further faces. Overlaps also have volume $O_T(b^3)$: distinct non-opposite face pairs can meet only where at least three box coordinates are within $O_T(b)$ of an endpoint. In dimension two the vertex tubes are disjoint for small b. The Hellinger integrand is bounded by a constant on these sets. For Fisher information the covariance derivative is uniformly bounded by the Gaussian polynomial formula, and the density is bounded below by $c_T > 0$ throughout these sets since they lie within $O_T(b)$ of K. These sets therefore contribute $o(b^2)$ to the relevant integrals.

For a point z in the relative interior of F, the event $y - bA^{1/2}G \in K$ converges to the two outward signed constraints $s_i Z_i \geq t_i$ and $s_j Z_j \geq t_j$, with $Z = A^{1/2}G$. All other box constraints disappear in this scaling. Hence

$$q_{b,A}(y) \longrightarrow |K|^{-1} S_{A_F}(t), \qquad q_{b,C}(y) \longrightarrow |K|^{-1} S_{C_F}(t).$$

For the covariance derivative in a zero-diagonal direction E, only the signed face convolution for the pair i, j and this edge survives. Rescale the tangential Gaussian integration to all of $\mathbb{R}^{d-2}$. The resulting density is the two-coordinate Gaussian marginal. Thus

$$\partial_u q_{b,A+uE}(y)\big|_{u=0} \longrightarrow |K|^{-1} s_i s_j E_{ij} \varphi_{A_F}(t).$$

All other edge measures are at a positive distance from each fixed relative interior point, so their scaled Gaussian convolutions vanish. The formula also applies for a vertex when $d = 2$, where no tangential integration is needed.

Dominated convergence for each fixed T now gives the profile integrals over $[-T, T]^2$. Let $T \to \infty$ by the localization estimates of the preceding section. They also show finiteness of the integrals in the resulting limits. If $A \neq C$ with the same diagonal, some pair covariance differs. The Gaussian quadrant survival functions for that pair cannot coincide everywhere, since they determine the bivariate law. Their continuity makes the corresponding squared profile integral positive. This proves Theorem 7.2.

7.6 Product experiments

For two distinct covariance shapes let $C_* > 0$ denote their limiting squared Hellinger coefficient. It is the coefficient of b in Theorem 7.1 when their normal variances differ on a set of positive boundary area. For a uniform box with equal coordinate variances and different covariance matrices, it is the coefficient of b^2 in Theorem 7.2. Write $k = 1$ or $k = 2$, respectively. Then, if $Nb^k \to \lambda \in [0, \infty)$,

$$H^2(q_{b,A}^{\otimes N}, q_{b,C}^{\otimes N}) \longrightarrow 2(1 - e^{-\lambda C_*/2}).$$

This follows by raising the single-observation affinity $1 - H^2/2$ to the Nth power. If $Nb^k \to 0$, total variation tends to zero. If $Nb^k \to \infty$, affinity tends to zero and total variation tends to one. In particular, the number of observations needed to distinguish fixed changes in visible normal variances is of order b^{-1}, while distinguishing fixed changes in off-diagonal covariance at a box with fixed coordinate variances requires order b^{-2}.

Statements and Declarations

Funding. This research was self-funded.

Competing interests. The author owns Luminary Oddities and has no other relevant competing interests.

AI assistance. OpenAI Codex assisted with derivations, proof development, source checks and independent internal review of each part. Anthropic Claude edited, consolidated and cross-referenced the parts and checked computations. Scott Schubr directed the research and is responsible for the content.

References

[1] S. Schubr, Sharp Wasserstein stability for Gaussian deconvolution of log-concave laws, Luminary Research Archive, edition 4 (2026). `https://research.luminaryoddities.org/a-sharp-logarithmic-inverse-bound-for-gaussian-estimation`.

[2] S. Schubr, Entropy curvature under Gaussian smoothing: restoration and moment-matched obstructions, Luminary Research Archive, edition 14 (2026), section "The signed moment problem for log-concave laws". `https://research.luminaryoddities.org/local-and-eventual-entropy-concavity-under-gaussian-smoothing`.

[3] S. Schubr, *Sharp Wasserstein stability for Gaussian deconvolution of log-concave laws*, Luminary Oddities Research Archive, edition 4 (2026). Edition 4.

[4] A. K. H. Kim, A. Guntuboyina and R. J. Samworth, *Adaptation in log-concave density estimation*, Ann. Statist. 46 (2018), 2279–2306. DOI.

[5] A. Saumard and J. A. Wellner, *Log-concavity and strong log-concavity: a review*, Stat. Surv. 8 (2014), 45–114. arXiv:1404.5886.

[6] S. Schubr, *entropy edition 14*, the part titled *The signed moment problem for log-concave laws*, Luminary Research Archive, 2026.

[7] S. G. Bobkov, *Isoperimetric and analytic inequalities for log-concave probability measures*, Ann. Probab. 27 (1999), 1903–1921. Author copy.

[8] H. J. Brascamp and E. H. Lieb, *On extensions of the Brunn–Minkowski and Prékopa–Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation*, J. Funct. Anal. 22 (1976), 366–389. DOI.

[9] E. del Barrio, E. Giné and F. Utzet, *Asymptotics for L_2 functionals of the empirical quantile process, with applications to tests of fit based on weighted Wasserstein distances*, Bernoulli 11 (2005), 131–189. DOI.

[10] G. Wyłupek, *Verifying the validity of exponentiality*, ALEA 21 (2024), 1807–1822. DOI.

[11] A. Goldenshluger and A. B. Tsybakov, *Estimating the endpoint of a distribution in the presence of additive observation errors*, Statist. Probab. Lett. 68 (2004), 39–49. DOI.

[12] J. Capitao-Miniconi, E. Gassiat and L. Lehéricy, *Support and distribution inference from noisy data*. arXiv:2304.09452v2.

[13] A. K. H. Kim and R. J. Samworth, *Global rates of convergence in log-concave density estimation*, Ann. Statist. 44 (2016), 2756–2779. arXiv:1404.2298.

[14] S. G. Bobkov, *Isoperimetric and analytic inequalities for log-concave probability measures*, Ann. Probab. 27 (1999), 1903–1921. DOI.

[15] K. J. Böröczky, A. Figalli and J. P. G. Ramos, *A quantitative stability result for the Prékopa–Leindler inequality for arbitrary measurable functions*, Ann. Inst. H. Poincaré Anal. Non Linéaire 41 (2024), 565–614. DOI.

[16] M. Cule and R. J. Samworth, *Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density*, Electron. J. Stat. 4 (2010), 254–270. arXiv:0908.4400.

[17] N. Brosse and A. Dalalyan, *Boundary-layer asymptotics for Gaussian-smoothed singular measures*. arXiv:2607.04514v1.

[18] M. Miranda Jr., D. Pallara, F. Paronetto and M. Preunkert, *Short-time heat flow and functions of bounded variation in $\mathbb{R}^N$*, Ann. Fac. Sci. Toulouse Math. 16 (2007), 125–145. DOI.

[19] M. van den Berg and K. Gittins, *Uniform bounds for the heat content of open sets in Euclidean space*. arXiv:1407.6139.