Calibrating small noise from exteriors, frontiers, cusps and density singularities

Scott Schubr
Luminary Oddities

November 2026

Abstract

We determine the sharp rates for calibrating the location and scale of a small additive noise when the latent law is known and bounded. Unknown coordinate shifts and bounded relative scales leave the sharp joint Wasserstein estimation and testing rates unchanged, and two tail counts near the strongest endpoint identify location and scale. With a known noise exterior, the scale has minimax rate $M^{-1/(r+1)}$ below $r = 1$, $(M \log M)^{-1/2}$ at $r = 1$ and $M^{-1/2}$ above, with Poisson and Gaussian local experiments. Each noise endpoint is calibrated at the fastest of three mechanisms: the strongest-window root rate, its own moving frontier, and an interior cusp in the opposite latent window, which a corrected reciprocal square-root weight exploits under unknown background. Where frontier and cusp survive at the same scale, the local experiment is a product of independent Poisson and Gaussian observations, neither of which suffices alone. Opposing latent and noise powers create a logarithmic density pole, giving rate $\log(M)/M$ with a drifted Brownian likelihood, the same law as critical bilateral gamma location.

Contents

1 Introduction **2**

2 Wasserstein recovery with unknown affine noise calibration **2**
2.1 Model and result 2
2.2 Calibrating one coordinate 3
2.2.1 Known shift and a monotone tail crossing 3
2.2.2 Moment calibration with unknown shift and scale 4
2.2.3 Unknown shift and scale from one endpoint 4
2.3 Recovery and testing on the common standardized box 5

3 Sharp scale calibration with a known noise exterior **6**
3.1 Model and results 7
3.2 Exact tip probabilities and upper bounds 7
3.3 The moment-matched comparison 8
3.4 Local experiments 9

4 Sharp affine noise calibration from frontiers and interior cusps **10**
4.1 Model and rates 10
4.2 Local observation densities 11
4.3 Upper bounds 11
4.4 Matching lower bounds 13

5 Local experiments at the frontiercusp crossover **13**
5.1 Model and statement 14
5.2 Reduction to the two windows 15
5.3 Limit of the retained experiment 16
5.4 Critical cusp information with unknown interior 17

6 Calibrating a noise endpoint from a logarithmic density singularity **18**
6.1 Model and conclusions 18
6.2 The interior logarithmic profile 19
6.3 Moment matching and Hellinger costs 20
6.4 Upper bounds 21
6.5 Local experiments 22
6.6 Direct location estimation at the variance-gamma pole 23

1 Introduction

For Gaussian noise of unknown covariance and smooth latent laws, calibration has a genuine price, determined in [3]. A bounded latent law with sharp endpoints changes this picture: the boundary layer identifies the noise. Section 2 shows that unknown affine calibration costs nothing in the recovery rates of [2]. Section 3 determines the sharp scale calibration from a known exterior. Sections 4–5 identify the frontier and cusp mechanisms and their crossover, and Section 6 treats the logarithmic singularity that links this problem to variance-gamma location [1].

2 Wasserstein recovery with unknown affine noise calibration

A known product density is observed with small additive noise whose coordinate shifts and relative scales are unknown. The sharp Wasserstein estimation and testing rates for the standardized joint noise law are unchanged by this calibration. Known standardized moments give a pilot estimate, but the known exterior of the noise density can also calibrate the scales directly. For shifts and scales together, two tail probabilities near one endpoint suffice using only the known exterior density and its positive bounds. A negative Sobolev estimate controls the moving support in the subsequent shape recovery.

2.1 Model and result

Let X have independent coordinates with known densities g_i on fixed intervals $[L_i, U_i]$. These are smooth and positive in the interior. At each endpoint, in inward depth h, assume $g_i = w_{i,\pm}(h)h^{\alpha_{i,\pm}}$, with a smooth positive $w_{i,\pm}$ up to zero and $\alpha_{i,\pm} > -1$. The data are N independent copies of

$$Y_i = X_i + b_i(c_i + a_i Z_i),$$

with Z independent of X. The $b_i > 0$ are known and sufficiently small, while a_i lies in a fixed known compact interval of positive numbers and c_i in a fixed known bounded interval.

Let $Q = [-1, 1]^d$. Fix a known reference density f_0 bounded above and below by positive constants on Q. The standardized noise density is $f = f_0 + u$, where u has mass zero, is supported in a_*Q for fixed $0 < a_* < 1$, and has torus H^s norm at most a fixed radius, with $s \geq 0$. Require uniform lower and upper density bounds on Q with strict slack around f_0. The noise coordinates may be dependent. Write $r_i = 1 + \min_\pm \alpha_{i,\pm}$, $R_* = \sum_i r_i$, and $M = N \prod_i b_i^{r_i}$.

The general model allows both c_i and a_i to be unknown, with no prescribed noise moments. We also give simpler pilots when the shifts are known or when the first two standardized moments match the known reference marginal.

Theorem 2.1. *Under the stated conditions, for $1 \le p \le 2$ the minimax expected W_p error for the standardized joint noise density has order*

$$\begin{cases} M^{-1/2}, & d + 2R_* < 2, \\ M^{-1/2}(\log M)^{1/2}, & d + 2R_* = 2, \\ M^{-(s+1)/(2s+d+2R_*)}, & d + 2R_* > 2, \end{cases}$$

as $M \to \infty$, uniformly for small b_i. The minimax separation for testing $f = f_0$ against $W_p(f, f_0) \ge \rho$, with the calibration unknown under both hypotheses and both errors at most $1/4$, has order

$$\begin{cases} M^{-1/2}, & d + 4R_* < 4, \\ M^{-1/2}(\log M)^{1/4}, & d + 4R_* = 4, \\ M^{-(s+1)/(2s+d/2+2R_*)}, & d + 4R_* > 4. \end{cases}$$

Unknown error variance can substantially change the rate in other deconvolution models, for example the unknown Gaussian variance model of Matias [7]. Delaigle and Hall [4] use phase information to treat an unknown error law. Hall and Simar [6] and Goldenshluger and Tsybakov [5] treat support endpoints under additive error. Here the convoluting density is known, and the unknown calibration acts on a compact standardized noise distribution. The known exterior or moment normalization gives direct coordinate calibration. We extend the exact window estimates and critical rates of [10, 9, 11] to this composite model. The published Gaussian deconvolution result [8] instead concerns a fixed Gaussian forward operator.

2.2 Calibrating one coordinate

Use an independent pilot portion of the data. The three methods below give, uniformly for each coordinate,

$$\mathbb{E}(|\widehat{a}_i - a_i|^2 + |\widehat{c}_i - c_i|^2) \le \frac{C}{N b_i^{r_i}}. \tag{1}$$

Under moment calibration the corresponding scale for the bound can be improved to $N\sigma_{i,2}(b_i)^2$, with $\sigma_{i,2}$ as in [11]. Constants may depend on the fixed brackets and reference. We explain the scalar argument with subscripts omitted, and reflect that coordinate if necessary so a strongest endpoint is the right one.

2.2.1 Known shift and a monotone tail crossing

Subtract the known bc from the data and fix $t_0 \in (a_*, 1)$. For a candidate scale A define

$$p_b(A) = \int_{t_0}^{1} f_{0,i}(z) \int_{0}^{bA(z-t_0)} g(U - h)\, dh\, dz.$$

At the true a, this is exactly the probability that Y after known-shift subtraction exceeds $U + bat_0$, since every contributing noise value is in the exterior where the marginal is known. The model curve $p_b(A)$ is increasing with derivative at least $c_1 b^r$ uniformly on the scale bracket, by differentiation and the endpoint taper. For example, restrict z to a fixed interval strictly inside $(t_0, 1)$, where

the noise marginal has a positive lower bound and the latent argument is at depth comparable to b. The empirical tail at $U + bAt_0$ is nonincreasing in A. Estimate a by their generalized crossing, clipped to the bracket.

At $A = a+\delta$, the true empirical-tail probability is at least $c_1 b^r \delta$ below $p_b(A)$, and the reverse gap holds at $a-\delta$. This uses monotonicity of the true tail and the derivative lower bound on the model curve. At either comparison threshold the probability is $O(b^r)$: the noise is supported on a fixed interval and exceeding the threshold requires the latent variable to be within $O(b)$ of its endpoint. The Bernoulli Bernstein inequality gives crossing error tails at most $2\exp(-c_2 N b^r \delta^2)$ within the bounded brackets. Integrating this bound proves (1); all pilot portions have size comparable to N.

2.2.2 Moment calibration with unknown shift and scale

Write $W = c + aZ$. Its support is in a fixed known interval and its density is bounded uniformly, because the scale bracket is bounded away from zero. The full-support exact moment probes of [11], after a fixed dilation to that interval, estimate $\mathbb{E}W$ and $\mathbb{E}W^2$ unbiasedly with squared risks at most $C/(N\sigma_j(b)^2)$, for $j = 1, 2$. Both σ_j^2 are at least a positive constant times b^r, and the smaller order of the two is σ_2^2. Clip the pilot moments to their fixed bounded ranges. The standardized mean μ and variance $\tau^2 > 0$ are known and $\operatorname{Var} W = a^2\tau^2$, $\mathbb{E}W = c + a\mu$. Compute a from the square root of the clipped variance, clipping it to the known positive variance range before taking the square root. Then recover c from the first moment, and clip to the shift bracket. These operations are Lipschitz on the relevant compact sets, establishing the bound.

2.2.3 Unknown shift and scale from one endpoint

We use only the known exterior and its positive upper and lower bounds. Write $s_* = c + a$ for the endpoint of W. Choose a large fixed constant $H > 2$, specified below, and then $d_0 > 0$ with $(H+1)d_0 < a_{\min}(1-a_*)$, where $a_{\min}$ is the scale bracket's lower end. A first pilot portion estimates s_* by the rescaled maximum $(\max Y - U)/b$, clipped to the known endpoint bracket. Call it $\widetilde{s}$. It never exceeds the true s_*. The event $0 \le s_* - \widetilde{s} \le d_0/10$ has probability at least $1 - \exp(-c_3 N b^r)$. Indeed, for any fixed sufficiently small $v > 0$, the probability $Y > U + b(s_* - v)$ is at least a constant times $b^r v^{r+1}$, by taking the noise within $v/(2a)$ of its endpoint and the latent variable within $bv/2$ of its endpoint. Apply this with $v = d_0/10$.

Condition on this first portion and use a second pilot portion. Take thresholds $U + bt_j$ with $t_1 = \widetilde{s} - d_0$ and $t_2 = \widetilde{s} - Hd_0$. Consider the candidate rectangle $\widetilde{s} \le S \le \widetilde{s} + d_0/10$ and A in the original positive scale bracket. For either threshold the exceedance probability depends only on the known exterior: its candidate depth $D_j = S - t_j$ is less than $A(1-a_*)$. Set $F_b(t) = \int_0^{bt} g(U-h)\,dh$ for positive t, and $F_b(t) = 0$ for $t \le 0$. Then the exact model curve is

$$p_j(S, A) = p(D_j, A), \qquad p(D, A) = \int_0^{1-a_*} f_{0,i}(1-w)\, F_b(D - Aw)\, dw.$$

On the relevant positive arguments $F_b(t) \asymp b^r t^r$ and $F_b'(t) \asymp b^r t^{r-1}$. The function F_b is absolutely continuous, including at zero. Its derivative is locally integrable, so the partial derivatives of p are the integrals with F_b' and $-wF_b'$, respectively. This follows from the fundamental theorem of calculus and Fubini. The differentiated integrals are continuous even for merely bounded $f_{0,i}$, by L^1 translation and dilation continuity of the truncated derivative kernel. Using the bounds on $f_{0,i}$ and substituting $w = Dt/A$ gives

$$p(D, A) \asymp b^r D^{r+1}/A, \qquad \partial_D p(D, A) \asymp b^r D^r/A, \qquad -\partial_{\log A} p(D, A) \asymp b^r D^{r+1}/A.$$

In the last expression the integrand has the added factor Aw. All constants are uniform whenever $D/A < 1 - a_*$ and the depths lie in the fixed taper region.

Use parameter coordinates $x = (S - \widetilde{s})/d_0$ and $v = \log A$. The first log tail $\log(p_1/b^r)$ has positive derivative in x bounded away from zero, and bounded derivative in v. The second log tail $\log(p_2/b^r)$ has negative derivative in v bounded away from zero in absolute value and derivative $O(1/H)$ in x. Comparing two points one coordinate at a time gives

$$|\Delta x| \le C_1|\Delta \log p_1| + C_2|\Delta v|, \qquad |\Delta v| \le C_3|\Delta \log p_2| + (C_4/H)|\Delta x|.$$

Choose H large enough to absorb the coupled error. Thus the inverse is quantitatively stable on the full rectangle.

Use empirical tail probabilities at the two pilot thresholds. Clip each after division by b^r to fixed positive lower and upper bounds containing all model values on the rectangle. Conditional on the first portion and its good event, the resulting logarithmic statistics have squared errors $O((Nb^r)^{-1})$, by the binomial variance bound and the Lipschitz logarithm on the clipped range. A minimum-distance fit of the two model log statistics on the rectangle therefore gives the same rate for S, A. For measurability one may take the first approximate minimizer over a fixed countable dense parameter grid in (x, v), with tolerance $(Nb^r)^{-1/2}$. Recover $c = S - A$ and clip to its bracket. On the exceptional first-pilot event the clipped error is bounded and its risk contribution is exponentially small. This proves (1) without moment normalization.

2.3 Recovery and testing on the common standardized box

Let $\Gamma = (c, a)$ be the true calibration and write $G_\Gamma f$ for its pushforward under the coordinatewise affine map. These transformed densities and all calibrated candidates are supported on a fixed enlarged box, and have a uniform upper density bound. They need not have a common positive density floor on that enlarged box. On an enlarged torus containing it, however,

$$\|G_\Gamma h - G_{\Gamma'} h\|_{H^{-1}} \le C|\Gamma - \Gamma'|\|h\|_2 \tag{2}$$

for every fixed h supported on Q. Indeed vary Γ linearly between its two values. Pair with a smooth test and differentiate its composition with the affine maps. The derivative is the gradient of the test paired with the bounded affine velocity times h. Cauchy–Schwarz and the bounded Jacobian and inverse Jacobian control it by $C|\Gamma - \Gamma'|\|h\|_2$ times the H^1 test norm. Integrate along the path. Also, a common affine pushforward is an isomorphism between the supported H^{-1} norms, with uniform constants, by composing tests with the inverse affine map and taking a fixed smooth cutoff. On the standardized cube the W_2 upper bound by H^{-1} holds because the densities are bounded below there, via the Neumann transport current. Thus no lower density floor is needed on a moving support.

Use a remaining independent portion of the data for shape inference, and condition on the pilot $\widehat{\Gamma}$. The exact corner probes apply to $G_\Gamma f$ on the known enlarged box, after a fixed rescaling of the coordinate intervals. The effective sample size is still comparable to M. Use Fourier probes centered at the known trial density $G_{\widehat{\Gamma}} f_0$. Their mean coefficients are those of

$$G_\Gamma u + (G_\Gamma f_0 - G_{\widehat{\Gamma}} f_0).$$

The first component has uniformly bounded isotropic H^s norm: it is supported strictly within the enlarged box and affine scalings in a compact bracket are uniformly bounded in H^s. The second

has H^{-1} norm $O(|\Gamma - \widehat{\Gamma}|)$ by (2). Hence truncating the Fourier reconstruction at m gives squared H^{-1} error at most

$$C\left[m^{-2(s+1)} + |\Gamma - \widehat{\Gamma}|^2 + M^{-1}\sum_{|j|\le m}\langle j\rangle^{-2}\prod_i \langle j_i\rangle^{2r_i}\right]$$

conditionally on the pilot, in estimating $G_\Gamma f$ after adding back the trial density. The sum is bounded, logarithmic, or polynomial with power $d+2R_*-2$, according to the sign of that exponent, by the same Fourier count as in [10, 9]. The pilot risk in (1) is $O(M^{-1})$, since $b_i^{r_i} \ge \prod_\ell b_\ell^{r_\ell}$ for $b_\ell \le 1$.

Project the reconstruction in H^{-1} onto a countable dense set of pushforwards $G_{\widehat{\Gamma}}h$ of admissible standardized densities, within the Fourier tail tolerance. The true standardized f supplies a comparator $G_{\widehat{\Gamma}}f$, whose additional error is at most $C|\Gamma - \widehat{\Gamma}|$ by (2). Pull back the chosen estimate and use the standardized transport bound. Optimizing the cutoff gives the three estimation upper rates in Theorem 2.1.

For testing, use the usual H^{-1}-weighted Fourier probe inner-product statistic on distinct sample pairs, centered now at $G_{\widehat{\Gamma}}f_0$. The pilot is independent of the testing portion. With probability at least 15/16 uniformly in the parameters, $|\Gamma - \widehat{\Gamma}| \le CM^{-1/2}$ by the pilot risk. On this event the norm of the truncated mean under the composite null is $O(M^{-1/2})$. Under an alternative it is bounded below by

$$cW_p(f, f_0) - C(M^{-1/2} + m^{-s-1}).$$

This follows from the common normalized-box W_2 upper bound, the inverse affine norm comparison, and the two components of the centered mean. Conditional on the pilot the probe covariance bounds are unchanged by the centering. The quadratic statistic has expectation equal to the squared truncated mean, and the variance decomposition used in [9] applies. Its degenerate term has standard deviation at most

$$CM^{-1}\left(\sum_{|j|\le m}\langle j\rangle^{-4}\prod_i\langle j_i\rangle^{4r_i}\right)^{1/2},$$

and the linear term at most $CM^{-1/2}m^{(R_*-1)_+}$ times the truncated mean norm. Take the same cutoff as in the calibrated problem. The null bias from the pilot is $O(M^{-1})$, no larger than the squared separation scales. A large constant in the threshold and a larger constant in the alternative separation make the conditional errors at most 3/16 by Chebyshev's inequality. Adding the pilot failure probability gives both required error bounds.

All lower bounds follow by fixing the calibration in its known brackets. The calibrated lower alternatives of [10, 9] can preserve any fixed finite coordinate moments, so they are available in the moment-normalized class as well. This finishes Theorem 2.1.

3 Sharp scale calibration with a known noise exterior

A known compact latent density is observed with small additive noise whose interior density and relative scale are unknown. A known uniform exterior suffices for sharp scale calibration without moment normalization. The optimal rate changes at the endpoint power for logarithmic information. We obtain local Poisson and Gaussian experiments by varying the interior to cancel bulk moments.

3.1 Model and results

A known X has density g on $[L, U]$, smooth and positive inside. At endpoint $e \in \{-, +\}$ the density at inward depth h is $w_e(h)h^{r_e-1}$, with w_e smooth and positive up to zero and $r_e > 0$. We observe N independent copies of $Y = X + baZ$, with X, Z independent. The nominal scale $b > 0$ is known and tends to zero, and the relative scale a is in a known compact positive bracket with nonempty interior. A known shift, if any, is subtracted first. The noise density on $[-1, 1]$ is $f = 1/2 + u$, where u has zero mass, is supported in $[-s_*, s_*]$ for fixed $0 < s_* < 1$, and is bounded in a fixed Sobolev ball of order $s \geq 0$ and positive radius. Uniform lower and upper density bounds have strict slack around $1/2$. There are no moment constraints on f beyond its normalization. Set $r = \min_e r_e$, $E = \{e : r_e = r\}$, and $M = Nb^r$.

Theorem 3.1. *Uniformly as $M \to \infty$ and $b \downarrow 0$, the minimax expected absolute error and root mean squared error for a have order*

$$d_M = \begin{cases} M^{-1/(r+1)}, & r < 1, \\ (M \log M)^{-1/2}, & r = 1, \\ M^{-1/2}, & r > 1. \end{cases}$$

Theorem 3.2. *Let a_0 be interior to the bracket, with uniform base noise. For any finite set of t, take $a = a_0 + td_M$. There exist admissible f_t along this subexperiment such that the physical noise-scale laws aZ share any prescribed finite number of initial moments with a_0Z_0, where Z_0 is uniform. If $r < 1$ the finite experiments converge in Le Cam distance to independent Poisson processes on the strongest endpoint marks, with intensities $w_e(0)(x + t)_+^r\, dx/(2a_0r)$, $e \in E$. If $r = 1$ they converge to a Gaussian shift of scalar information $\mathcal{I} = \sum_{e \in E} w_e(0)/(4a_0)$. In this critical case the estimator below attains a centered normal limit with variance $4a/\sum_{e \in E} w_e(0)$ after normalization by $\sqrt{M \log M}$, for every admissible interior noise density and every interior scale.*

Smith [12] gives the classical transition for fixed moving-support parametric densities, at observed endpoint power one. Hall and Simar [6] and Goldenshluger and Tsybakov [5] consider endpoint estimation under additive observation error. Here the endpoint prefactor shrinks, and the unknown interior can cancel the bulk information while retaining the known moving exterior. The root-rate pilots for shape inference in Section 2 do not determine this calibration rate.

3.2 Exact tip probabilities and upper bounds

Reflect either coordinate when necessary so its endpoint is rightward. For the resulting scaled observation $v = (Y - U)/b$, the density near the tip $v = a$ is exactly

$$b^{r_e}(2a)^{-1} J_{e,b}((a - v)_+), \qquad J_{e,b}(x) = \int_0^x w_e(bz) z^{r_e-1}\, dz. \tag{3}$$

Only the known uniform noise tail contributes there. Consequently the probability of a gap at most d at a strongest endpoint has order $b^r d^{r+1}$, uniformly for all sufficiently small positive d. For $r < 1$ the clipped maximum gives a gap with tail $\exp(-cMd^{r+1})$ below a fixed gap, and exponentially small probability of a larger gap. This gives the claimed squared risk by integration.

For $r > 1$ use the tail crossing of Section 2. Fix $t_0 \in (s_*, 1)$, compare the empirical tail at $v = At_0$ with $p(A) = \frac{1}{2}\int_{t_0}^1 \int_0^{bA(z-t_0)} g(U - h)\, dh\, dz$ at a strongest endpoint. They agree in expectation at the true scale, the empirical curve is nonincreasing and $p'(A) \geq cb^r$. Each trial tail probability is $O(b^r)$. Bernstein's inequality yields tail error bound $2e^{-cMd^2}$ for a clipped crossing error exceeding d.

For $r = 1$, use $N_0 = \lfloor N/\log M \rfloor$ observations for a clipped maximum $\widetilde{a} \le a$ from one strongest endpoint. Put $\Delta = M^{-1/2}(\log M)^2$. The event $0 \le a - \widetilde{a} \le \Delta/4$ fails with probability at most $\exp(-c(\log M)^3)$. On the remaining observations, at every strongest endpoint set $x = \widetilde{a} - v$ and use weight $x^{-1}\mathbf{1}_{\Delta \le x \le d_0}$, where fixed d_0 is small enough for the exact uniform tail. Sum the weights over the strongest endpoints; their windows are disjoint. The exact mean at a candidate A within the clipped interval $[\widetilde{a} - \Delta/4, \widetilde{a} + \Delta/2]$ is

$$m_{\widetilde{a}}(A) = b \sum_{e \in E} \int_{\Delta}^{d_0} \frac{J_{e,b}(x + A - \widetilde{a})}{2Ax}\, dx.$$

Since $J_{e,b}(x) = w_e(0)x + O(bx^2)$ and $J'_{e,b}(x) = w_e(0) + O(bx)$, its derivative is

$$b\,\frac{\sum_{e \in E} w_e(0)}{2A} \log(d_0/\Delta) + O(b).$$

On the good pilot event the weight variance at the true scale is $b\sum_e w_e(0)\log(d_0/\Delta)/(2a) + O(b)$. Indeed its second moment integrates $J_{e,b}(x + a - \widetilde{a})/x^2$; the additional gap term contributes only $O(1)$ to that integral, and the squared mean is $O(b^2)$. Use the clipped inverse of $m_{\widetilde{a}}$ on the empirical average. The derivative lower bound gives squared risk $O((M\log M)^{-1})$; the bad pilot has negligible bounded loss. At a true interior scale clipping is inactive with probability tending to one. The conditional Lindeberg bound follows from $\Delta^{-1}/\sqrt{M\log M} \to 0$, the pilot fraction vanishes, and $\log(d_0/\Delta) \sim \frac{1}{2}\log M$. The mean and variance formulas therefore give exactly the normal limit in Theorem 3.2.

3.3 The moment-matched comparison

Fix an interior a_0 and put $D = |a - a_0|$. Take an integer K with $2K + 2 > \max_e r_e$. Choose smooth compactly supported moment lifts p_j in $(-s_*, s_*)$ for $1 \le j \le K$, dual to the degree zero through K monomials, with zero mass. Then

$$f_a(z) = \frac{1}{2} + \sum_{j=1}^{K} \Big[(a_0/a)^j - 1\Big] \left(\int_{-1}^{1} z^j\, dz/2\right) p_j(z)$$

is admissible for small D and matches the physical moments. This follows from the moment identities; its smooth interior correction is of size $O(D)$. The two physical noise-scale densities differ in L^1 by $O(D)$, since both the moving uniform densities and their interior corrections do so.

Let q_a be the observation density along this path. We bound

$$H^2(q_a, q_{a_0}) \le C b^r \Psi_r(D), \quad \Psi_r(D) = \begin{cases} D^{r+1}, & r < 1, \\ D^2 \log(e/D), & r = 1, \\ D^2, & r > 1. \end{cases} \tag{2}$$

At latent inner depths at least Tb, for fixed large T, the Taylor expansion through K in physical noise-scale position and the matched moments give signed output size $CDb^{K+1}h^{r_e - K - 2}$ near endpoint e. The observation density is comparable to $h^{r_e - 1}$, so integrating the squared discrepancy over the density yields $O(b^{r_e}D^2)$; deep interior cost is $O(b^{2K+2}D^2)$.

In fixed scaled endpoint windows the uniform contribution to the density of v is

$$\frac{b^{r_e}}{2a}\{J_{e,b}((a - v)_+) - J_{e,b}((-a - v)_+)\}.$$

The interior correction gives $O(Db^{r_e})$ there by integrability of the one-sided kernel, and vanishes near the upper tip. At the upper tips, the $O(D)$ region covering the two tips costs $O(b^{r_e}D^{r_e+1})$ by summing the densities. At depth $h > D$ below the smaller tip, the density floor and derivative bound give a cost at most $Cb^{r_e}D^2\int_D^{d_0} h^{r_e-2}\,dh$. This is $O(b^{r_e}\Psi_{r_e}(D))$. In the remaining part of the fixed window the density floor is cb^{r_e}, since we stay a fixed distance below the tip. The correction, normalization change and upper primitive away from its tip cost $O(b^{r_e}D^2)$. The lower primitive has a moving secondary cusp: within $O(D)$ of its zero argument the difference is $O(D^{r_e})$ after division by b^{r_e}; farther away it is at most CDh^{r_e-1} on the positive side. Its squared weighted cost is therefore $O(b^{r_e}D^{2r_e+1})$ if $r_e < 1/2$, $O(b^{r_e}D^2\log(e/D))$ if $r_e = 1/2$, and $O(b^{r_e}D^2)$ otherwise. All contributions are bounded by (2), since $r_e \geq r$. The exterior beyond the windows has zero density. For D a sufficiently small multiple of d_M, the N-sample Hellinger bound is small. The two-point absolute risk lower bound establishes both claimed lower rates.

3.4 Local experiments

When $r < 1$, use on the strongest endpoints the marks $x = (a_0 - v)/d_M$. In bounded windows, (1) gives the intensity in Theorem 3.2 since $Md_M^{r+1} = 1$. Binomial Poisson approximation and bounded-window intensity convergence give convergence in total variation of the cropped point processes, independently at each endpoint. For positive marks $x > T > |t|$, the strongest-tip squared Hellinger integral begins at depth of order Td_M, and its total N-sample bound from the proof of (2) is CT^{r-1}. The far negative mark side is empty. Weaker-tip costs vanish since $r_e > r$ in $Nb^{r_e}\Psi_{r_e}(d_M)$. Secondary cusp and ordinary interior costs vanish by their higher powers of d_M. Thus the omitted data can be filled from the normalized reference exterior at that same cropped count with error tending to zero as $T \to \infty$. The omitted Poisson marks can likewise be filled from the reference intensity, whose squared-root tail discrepancy has the same T^{r-1} bound. This proves the local Poisson limit.

When $r = 1$, keep only the strongest endpoint depth windows $\Delta \leq a_0 - v \leq d_0$, with the same Δ as before. The omitted shallower tips cost at most $Cbd_M^2(1 + \log(\Delta/d_M))$ per observation by the proof of (2); all other omitted regions cost $O(bd_M^2)$. The corresponding N-sample costs vanish. Use the retained values and one common symbol for censoring all other observations. This finite experiment is equivalent asymptotically by filling the reference normalized complement. In a retained scaled depth h, the density derivative with respect to a has leading value $bw_e(0)/(2a_0)$ with error $O(bh)$, while the density is $bw_e(0)h/(2a_0) + O(b^2h^2)$. Its scalar Fisher contribution is therefore $b\sum_{e\in E} w_e(0)\log(d_0/\Delta)/(2a_0) + O(b)$. The censoring symbol has derivative $O(b)$ and mass tending to one, so contributes $O(b^2)$. The exact likelihood increments at $a_0 + td_M$ are centered under the reference and bounded by $Cd_M/\Delta = o(1)$. The second-order density Taylor remainder on the retained windows is $O(bd_M^2)$, giving total weighted squared error $O(Nbd_M^4\log(1/\Delta)) = o(1)$. The censoring symbol's Taylor remainder is also $O(bd_M^2)$. Thus the linear likelihood term converges by the triangular central limit theorem and its squared sum concentrates. The cubic log remainder vanishes by the maximum bound. The limiting information after the normalization is $\mathcal{I}$, since $\log(d_0/\Delta) \sim \frac{1}{2}\log M$. The finite likelihood vectors converge to those of the Gaussian shift. Their coordinates have mean one before and after the limit, so truncating and quantizing the likelihood vectors and regenerating the reference conditional gives convergence in Le Cam distance. This completes the local claims.

4 Sharp affine noise calibration from frontiers and interior cusps

Consider a known interval-supported density observed with small additive noise of unknown shift, scale and interior density. The known exterior of the noise creates a moving support frontier at each end of the latent interval. Each noise endpoint also creates an interior cusp at the opposite end of the latent interval. We determine the sharp coordinatewise calibration rates by comparing these two sources of information with a regular amplitude measurement. The interior cusp can give the faster rate, and its critical power gives a logarithmic improvement even with unknown smooth background.

4.1 Model and rates

Let X have known density g, supported on $[L, U]$, smooth and positive in the interior. At latent endpoint $e \in \{-, +\}$ the density at inward depth h is $w_e(h)h^{r_e-1}$, with $r_e > 0$ and w_e smooth and strictly positive up to zero. Observe N independent copies of $Y = X + b(c + aZ)$, with independent X, Z. Here $b \downarrow 0$ is known, a lies in a compact positive bracket, and c is in a compact bracket. Both brackets have nonempty interiors. The density of Z is $1/2 + u$ on $[-1, 1]$, with zero-mass u supported in $[-s_*, s_*]$ for fixed $0 < s_* < 1$ and in a fixed Sobolev ball of order $s \geq 0$ and positive radius. Uniform positive lower and finite upper bounds on the density have strict slack at $1/2$. No moments are prescribed beyond mass.

Write $S_+ = c + a$, $S_- = c - a$, $M_e = Nb^{r_e}$, and $M_* = \max(M_+, M_-)$, with $r_* = \min(r_+, r_-)$. For $m \geq 0$ set $\bar{m} = \max(2, m)$ and define

$$F_p(m) = \begin{cases} \bar{m}^{-1/(p+1)}, & 0 < p < 1, \\ (\bar{m} \log \bar{m})^{-1/2}, & p = 1, \\ 1, & p > 1, \end{cases} \qquad C_q(m) = \begin{cases} \bar{m}^{-1/(2q+1)}, & 0 < q < 1/2, \\ (\bar{m} \log \bar{m})^{-1/2}, & q = 1/2, \\ 1, & q > 1/2. \end{cases}$$

Theorem 4.1. *Uniformly as $M_* \to \infty$ and $b \to 0$, the minimax expected absolute and root mean squared errors for S_e both have order*

$$\rho_e = \min\{M_*^{-1/2}, F_{r_e}(M_e), C_{r_{-e}}(M_{-e})\}.$$

For each of a and c the corresponding order is $\max(\rho_+, \rho_-)$.

For example, suppose the opposite endpoint has power q and the own endpoint power p, with $0 < q < p < 2q < 1$. When $N \asymp b^{-\gamma}$, $\gamma > p$, the interior cusp gives the faster rate below $\gamma = (p + pq - q)/(2q - p)$ and the moving frontier gives the faster rate above it. At equality the rates have the same order. For $p \geq 2q$ and $q < 1/2$, the opposite cusp gives at least as fast a rate as either other contribution for small b.

Prakasa Rao [13] obtained the classical subcritical cusp-location rate in a parametric density model. Smith [12] studied the moving-support transition, and Hall and Simar [6] studied boundaries with additive observation error. The two singularities here encode the same unknown endpoint, with different powers and possibly different effective sample sizes. The root-rate calibration used for shape inference in Section 2 and the known-shift rate in Section 3 leave this comparison open. The moment-matched lower construction uses the same principle of cancelling smoothly transmitted information as the lower construction of [14], with physical noise moments matched here to suppress contributions from the latent interior.

4.2 Local observation densities

Use reflected coordinates on the left so that either chosen latent endpoint is rightward. Let $V = (Y - U_e)/b$ in the reflected coordinates, and write $s < t$ for the two noise endpoints in those coordinates. Thus $(s,t) = (S_-, S_+)$ at the right latent endpoint and $(s,t) = (-S_+, -S_-)$ at the left. In a bounded V-window the uniform noise contribution to its density is

$$\frac{b^{r_e}}{2a}\{J_{e,b}((t-v)_+) - J_{e,b}((s-v)_+)\}, \qquad J_{e,b}(z) = \int_0^z w_e(bh)h^{r_e-1}\,dh.$$

This follows by integrating over the physical noise position in units of b. The interior correction is supported in noise positions a fixed positive distance from s, t. Hence near t the first primitive alone gives the entire density. Near s the density has the exact form

$$b^q\left\{B(v) - \frac{1}{2a}J_{e,b}((s-v)_+)\right\}, \qquad q = r_e. \tag{4}$$

Here B has uniformly bounded derivatives of every fixed order on a fixed-radius neighborhood of s. Indeed the first primitive's argument is bounded away from zero there, and differentiation of the correction kernel only uses noise positions separated from v. The density in (4) is bounded above and below by positive multiples of b^q. All neighborhoods can be chosen uniformly over the compact brackets and density class.

4.3 Upper bounds

We use observations in disjoint portions of sizes comparable to N. A root-rate pilot for (a, c) can use just the strongest latent endpoint. For completeness, first estimate its outward noise endpoint t by the clipped maximum $\widetilde{t}$. For any small fixed $d > 0$ one has $0 \le t - \widetilde{t} \le d/10$ except with probability $\exp(-c_0 M_*)$, by the first primitive and the corresponding tail mass of order $b^{r_*}d^{r_*+1}$. On a separate pilot portion, use thresholds $\widetilde{t} - d$ and $\widetilde{t} - 2d$, both in the known exterior for all candidates in $[\widetilde{t}, \widetilde{t} + d/10]$ and the positive a-bracket when d is small. The model probabilities are $b^{r_*}I_j(T)/(2A)$, where $I_j(T) = \int_0^{T-v_j} J_{e,b}(h)\,dh$ and v_j is threshold j. Uniformly on the rectangle the derivative of $\log(I_1/I_2)$ is bounded below by a positive constant, since $I_j'/I_j = (r_*+1)/(T-v_j) + O(b)$. The ratio determines T with a Lipschitz inverse, and either scaled probability then determines A with a Lipschitz inverse. Use clipped normalized empirical probabilities and an approximate best fit at tolerance M_*^{-1}. Bernstein's inequality and the fixed brackets give

$$\Pr\{|\widehat{c} - c| + |\widehat{a} - a| > z\} \le C\exp(-c_1 M_* z^2) + C\exp(-c_1 M_*) \tag{5}$$

for $z \ge C/M_*$ within the bounded error range. This gives the first term of each rate and localizes all subsequent fits. All estimates below can be clipped to fixed brackets.

Own frontier. Use reflected coordinates so the requested endpoint is the outer endpoint t, with latent exponent $p = r_e$. For $p < 1$, the gap of the maximum has tails bounded by $\exp(-c_2 M_e d^{p+1})$ up to a fixed small gap, with exponentially small probability of a larger gap. For $p = 1$ take a separate maximum pilot $\widetilde{t}$, set $\Delta = M_e^{-1/2}(\log M_e)^2$, and use the weight $x^{-1}\mathbf{1}_{\{\Delta \le x \le d_0\}}$ on $x = \widetilde{t} - V$, with fixed small d_0. The gap is at most $\Delta/4$ except with probability $\exp(-c_3(\log M_e)^4)$. Fit the exact mean

$$\frac{b^p}{2\widehat{a}}\int_\Delta^{d_0} x^{-1}J_{e,b}(x + T - \widetilde{t})\,dx$$

over $|T-\widetilde{t}|\le\Delta/2$, using an independent $\widehat{a}$ from (5). Its derivative in T is at least $cb^p\log M_e$, and replacing a by $\widehat{a}$ changes the mean by at most $Cb^p|\widehat{a}-a|$. The conditional variance of a weight is $O(b^p\log M_e)$. Inverting the clipped empirical mean gives squared error $O((M_e\log M_e)^{-1})$, because $M_*\ge M_e$. All rate arguments using a growing M_e only need to apply above a large fixed threshold, as bounded effective sizes cannot improve the root pilot as $M_*\to\infty$.

Opposite cusp below the critical power. In the opposite reflected window the requested endpoint is s up to sign, as in (4), with $0<q<1/2$ and effective size $m=M_{-e}$. Fix a bounded smooth compactly supported function ψ with integral zero and $\int\psi(u)(-u)_+^q\,du\neq0$. For a proposed center z and bandwidth h let the empirical statistic be the mean of $\psi((V-z)/h)$, divided by b^qh^{q+1}. For $|z-s|\le\epsilon h$ and ϵ small its population absolute value is at least $2\lambda>0$, uniformly at sufficiently small bandwidths, by (4) and continuity in $(z-s)/h$. For $|z-s|\ge Ah$, with A sufficiently large and the test window inside the fixed cusp neighborhood, the value is at most $\lambda/2$. To see this, the contribution of B is $O(h^{1-q})$ by cancellation, and the cusp gives either zero or $O(A^{q-1})$ by the derivative bound on $J_{e,b}$. The unscaled weight is bounded and has variance at most Cb^qh. Bernstein's inequality bounds a normalized error above $\lambda/4$ by $2\exp(-c_4mh^{2q+1})$.

Choose a sufficiently small fixed initial bandwidth, and use (5) to initialize within that bandwidth divided by two of s, except with probability $O(e^{-c_5M_*})$. Use successively halved bandwidths and a fixed lattice at each bandwidth of spacing at most ϵh. Search lattice points within distance Kh of the previous estimate, with fixed $K>2A+2$. Choose any with statistic of absolute value at least λ, or retain the previous estimate if none qualifies. Upon previous success, the search contains a candidate within ϵh of s, and all candidates belong to a deterministic set of bounded size within distance $(K+2A)h$ of s. Take the starting bandwidth small enough for these tests to be inside the fixed neighborhood. On no error larger than $\lambda/4$ over that set the new estimate is within Ah of s. This deterministic set argument permits reuse of the same data at all bandwidths. If the first failure is at h, all later moves have total size at most $2Kh$, so final error is $O(h)$. Stop at bandwidth $h_f\asymp m^{-1/(2q+1)}$. The mean squared error is bounded by

$$Ch_f^2+C\sum_{h\ge h_f}h^2e^{-c_4mh^{2q+1}}+Ce^{-c_5M_*}\le C'h_f^2,$$

where the sum is dyadic. The search is defined only by lattice points and empirical thresholds.

Opposite cusp at the critical power. Now $q=1/2$. Use (5) for $\widetilde{s},\widehat{a}$ on independent data from the fit. Use $\Delta=m^{-1/2}(\log m)^2$. The event $|\widetilde{s}-s|\le\Delta/4$ fails with probability $O(\exp(-c_6(\log m)^4))$, since $M_*\ge m$. Write $H=m^{-1/8}$ and $x=\widetilde{s}-V$. From $x^{-1/2}\mathbf{1}_{\{\Delta\le x\le H\}}$ subtract a function supported on $[H/2,H]$ to make moments of degrees $0,\dots,4$ zero. To construct the subtraction, rescale by H and solve the moment equations with polynomials times a fixed positive smooth bump on $(1/2,1)$. The fixed polynomial Gram matrix is invertible and the rescaled moments are bounded. The subtraction is thus bounded by $CH^{-1/2}$. Call the resulting weight k. Taylor's formula bounds the contribution of B to the weight mean by $O(b^qH^{11/2})$, by the five cancelled moments and uniform fifth derivative. The second moment of the weighted observation is bounded by $Cb^q\int k^2\le Cb^q\log m$.

Fit the empirical statistic using only the cusp part of its mean,

$$-\frac{b^q}{2\widehat{a}}\int_\Delta^H k(x)J_{e,b}(x+S-\widetilde{s})\,dx,\qquad|S-\widetilde{s}|\le\Delta/2.$$

The derivative has a fixed sign and absolute size comparable to $b^q\log m$, uniformly there. Indeed the uncorrected contribution to the derivative is of order $\log(H/\Delta)$, whereas the subtraction changes

the derivative integral by a bounded amount. The misspecification of the scale changes the fitted mean by at most $Cb^qH|\widehat{a}-a|$, since the cusp primitive is $O(x^{1/2})$ on the weight support. The clipped inverse then has mean squared error $O((m\log m)^{-1})$. Here the background bias divided by b^q is $O(m^{-11/16})$ and the scale estimate has squared risk $O(M_*^{-1})$. This proves the third rate in all cases where it can improve the pilot.

4.4 Matching lower bounds

Choose an interior reference (a_0, c_0) with uniform base noise. For one endpoint change its position by $D > 0$ and hold the other fixed. The two affine parameters consequently change by $D/2$ in absolute value. Match the first K moments of the physical noise position $W = c + aZ$ to the reference, for a fixed K with $2K+2 > \max(r_-, r_+)$. Explicitly, prescribe the standardized moments of Z to be those of $(W_0 - c)/a$, where $W_0 = c_0 + a_0 Z_0$. The discrepancy from the moments of Z_0 is $O(D)$. Add a linear combination of smooth lifts, dual to the monomials of degrees one through K with zero mass, supported in $(-s_*, s_*)$. Such lifts follow by inversion of the polynomial Gram matrix for a smooth positive bump there. The correction is $O(D)$ in every fixed smooth norm, so the alternative lies in the class, and the physical densities of W differ in L^1 by $O(D)$.

Let $p = r_e$ on the own side, $q = r_{-e}$ on the opposite side, and $r_* = \min(p, q)$. With squared Hellinger distance defined by the integral of squared square-root differences, the single-draw laws P_0, P_D satisfy

$$H^2(P_0, P_D) \le C\{b^{r_*}D^2 + b^pA_{p+1}(D) + b^qA_{2q+1}(D)\}, \qquad A_\nu(D) = \begin{cases} D^\nu, & \nu < 2, \\ D^2\log(e/D), & \nu = 2, \\ D^2, & \nu > 2. \end{cases} \tag{6}$$

Away from fixed scaled endpoint windows, expand g in the shift bW to order K. At inward depth $h \ge Tb$ near an endpoint of power r, matching and the L^1 bound give pointwise discrepancy at most $CDb^{K+1}h^{r-K-2}$; the density floor is ch^{r-1}. Integrating the square over this floor gives $O(D^2b^r)$. Deep inside the bound is $O(D^2b^{2K+2})$. In each scaled window the smooth correction has size $O(Db^r)$ by kernel integrability, and is zero near the outer frontier. A fixed frontier thus contributes $O(b^rD^2)$ by pure amplitude change. A moving frontier has cost $O(b^pD^{p+1})$ in an $O(D)$ strip covering both tips, by summing the densities. At depths $h \ge D$ below the smaller tip in the frontier window the floor is cb^ph^p and the shift discrepancy is bounded by Cb^pDh^{p-1}. The added cost is $O(b^pD^2\int_D^{d_0} h^{p-2}\,dh)$. The remainder of each scaled window has density bounded below by cb^r. There the moving secondary cusp in the opposite window gives discrepancy at most Cb^qD^q on the $O(D)$ strip and at most Cb^qDh^{q-1} beyond it on the cusp side. Its cost is $O(b^q[D^{2q+1} + D^2\int_D^{d_0} h^{2q-2}\,dh])$. A fixed cusp changes only amplitude, and all remaining terms away from zeros cost $O(b^{r_*}D^2)$. This proves (6).

Putting D equal to a small constant times ρ_e makes the product Hellinger bound small, including when either effective size is bounded. For the logarithmic term use monotonicity of $D^2\log(e/D)$ near zero and the inverse order $(m\log m)^{-1/2}$. The two-point risk lower bound applies to S_e, a, and c along this submodel. Taking either endpoint for the latter two parameters proves their claimed lower rates. Their upper rates follow by taking the half-sum and half-difference of the two endpoint estimators. This completes the proof.

5 Local experiments at the frontiercusp crossover

An affine noise endpoint can move both a support frontier and an interior cusp of a small-noise convolution. At a sampling scale at which both singularities contribute, we obtain a product

experiment consisting of a moving-frontier Poisson process and a Gaussian cusp observation. At the critical cusp power the Gaussian part is a scalar shift. A moment-corrected estimator attains the critical cusp information uniformly with respect to an unknown smooth background.

5.1 Model and statement

Let X have a known density on $[L, U]$, smooth and positive inside, equal at inward depth h to $w_p(h)h^{p-1}$ near U and $w_q(h)h^{q-1}$ near L, where both weights are smooth and positive up to zero. Observe N independent copies of $X + bW$ with $W = c + aZ$ independent of X. The scale $b \downarrow 0$ is known. The parameters a, c lie in fixed compact brackets, a bounded away from zero. The density of Z is $1/2 + u$ on $[-1, 1]$, with uniform positive lower and finite upper bounds having strict slack at $1/2$. The perturbation u has zero mass, is supported in $[-s_*, s_*]$ for fixed $s_* < 1$, and lies in a fixed positive-radius Sobolev ball of nonnegative order.

Fix interior a_0, c_0 and uniform Z_0. The target is $T = c + a$ and the other endpoint is held fixed at $c_0 - a_0$. Let $T = T_0 + z\delta$, $\delta \downarrow 0$, and let z range over a fixed finite set containing zero. As in Section 4, choose smooth, compactly supported corrections u_z of size $O(\delta)$ in each fixed smooth norm, matching the first K physical noise moments to those of $W_0 = c_0 + a_0 Z_0$. Here $2K + 2 > \max(p, q)$, and K can be arbitrarily large. Write $M_p = Nb^p$, $M_q = Nb^q$,

$$A_p = \frac{w_p(0)}{2a_0 p}, \qquad A_q = \frac{w_q(0)}{2a_0 q}, \qquad B_0 = A_q(2a_0)^q.$$

The reflected opposite window has coordinate $v = (L - Y)/b$; its interior cusp moves from $s_0 = -T_0$ to $s_0 - z\delta$.

Theorem 5.1. *Suppose $N \to \infty$, $0 < p < 1$, $0 < q \le 1/2$, and $M_p\delta^{p+1} \to \tau < \infty$. For $q < 1/2$ suppose $M_q\delta^{2q+1} \to \lambda < \infty$. For $q = 1/2$ suppose $M_q\delta^2 \log(1/\delta) \to \lambda < \infty$. The finite local experiments on the moment-matched path converge in Le Cam distance to the independent product of the following two observations.*

1. *A Poisson process on the real line with intensity $\tau A_p(x + z)_+^p\, dx$.*
2. *If $q < 1/2$, a Gaussian white-noise shift on the real line with unit variance and L^2 signal*

$$\mu_z(x) = -\frac{\sqrt{\lambda}\, A_q}{\sqrt{B_0}}\left[(-z - x)_+^q - (-x)_+^q\right].$$

If $q = 1/2$, a scalar Gaussian shift with log likelihood $zG - z^2 I/2$ at the reference, where $G \sim N(0, I)$ and $I = \lambda(qA_q)^2/B_0$.

A block with zero intensity or information may be omitted.

When both limits are positive, neither singularity disappears from the local experiment. For $q < 1/2$, positive limits are possible exactly when $q < p < 2q$. Indeed the ratio of the frontier and cusp order factors is $b^{p-q}\delta^{p-2q}$. At $q = 1/2$ the ratio is $b^{p-q}\delta^{p-1}/\log(1/\delta)$; both limits can be positive when $q < p < 1$. When $p = q = 1/2$ this last ratio diverges, and when $p < q$ it diverges as well.

The conditions can be realized explicitly. For $q < 1/2$ set $\kappa = (p - q)/(2q - p)$, $\delta = b^\kappa$ and $N \sim b^{-q}\delta^{-(2q+1)}$. For $q = 1/2$ set $\kappa = (p-q)/(1-p)$, $\delta = b^\kappa(\log(1/b))^{-1/(1-p)}$ and $N \sim b^{-q}\delta^{-2}/\log(1/\delta)$. These choices give positive finite limits.

More explicitly, put

$$D_\alpha = \int_{\mathbb{R}} [(x+1)_+^\alpha - x_+^\alpha]^2 \, dx = \frac{\Gamma(1+\alpha)^2}{\Gamma(2+2\alpha)\cos(\pi\alpha)}, \qquad 0 < \alpha < 1/2.$$

For the equality use the distributional Fourier transform of x_+^α, obtained by exponential damping and the gamma integral. The difference is in L^2, and its transform away from zero has squared modulus $2\Gamma(1+\alpha)^2(1-\cos t)|t|^{-2-2\alpha}$. Plancherel and the usual cosine integral give the formula. For distinct parameters at distance $d = |z - z'|$, the affinity of the limiting experiment is

$$\exp\left\{ -\frac{\tau A_p D_{p/2}}{2} d^{p+1} - \begin{cases} \dfrac{\lambda A_q^2 D_q}{8B_0} d^{2q+1}, & q < 1/2, \\ Id^2/8, & q = 1/2. \end{cases} \right\}.$$

The exponents are the squared-root intensity distance for the Poisson block and the squared shift distance for the Gaussian block with their respective factors $1/2$ and $1/8$. If both limits are positive, the product affinity is strictly smaller than either block's affinity. Since Markov kernels can only increase affinity, neither block alone is equivalent to the full product.

The Gaussian part for $q < 1/2$ has the familiar cusp likelihood of Prakasa Rao [13] and the Gaussian signal analysis of Chernoyarov, Dachian and Kutoyants [15]. Its centered log-likelihood process is a multiple of two-sided fractional Brownian motion of Hurst index $q + 1/2$. This follows from the translation invariance and homogeneity of the L^2 norm of the signal increments. The moving frontier is the small-noise counterpart of the vanishing-support setting of Smith [12]. The point here is the joint experiment at a scale where a single parameter moves both singularities in distinct windows. The moment matching isolates these windows from the interior, as in the affine calibration lower bound of Section 4 and the cancellation principle in the Archive Gaussian inverse bound [14].

5.2 Reduction to the two windows

For an endpoint of power r, put $J_{r,b}(h) = \int_0^h w_r(bt) t^{r-1} \, dt$, $h \geq 0$. In reflected endpoint coordinates write $s < t$ for the physical noise endpoints. On bounded scaled windows the uniform contribution to the scaled observation density is exactly

$$\frac{b^r}{2a} [J_{r,b}((t-v)_+) - J_{r,b}((s-v)_+)].$$

The moment lifts contribute $O(b^r \delta)$ there and vanish near the outer frontier. To obtain the lifts, match the standardized moments to those of $(W_0 - c)/a$ using smooth functions in the allowed interior, dual to the finitely many monomials. The physical densities differ in L^1 by $O(\delta)$. Off the bounded endpoint windows, Taylor expansion of the latent density through K in bW and moment matching give single-draw squared Hellinger cost $O((b^p + b^q)\delta^2)$, as in Section 4. For example near the endpoint of power r at depth $h \geq Cb$, the discrepancy is bounded by $C\delta b^{K+1} h^{r-K-2}$, the density is bounded below by ch^{r-1}, and the exponent condition on K gives the claimed integral.

Inside the scaled windows, fixed singularities change only their amplitude, and smooth terms away from the moving zeros have weighted squared cost $O((b^p + b^q)\delta^2)$. This product cost is $o(1)$ under all hypotheses of the theorem. In the own tip window use marks $x = (T_0 - v)/\delta$. Retain $|x| \leq R$ for large fixed $R > \max_z |z| + 1$. The discarded own tip and its scaled continuation have product squared Hellinger cost at most

$$CM_p \delta^2 \int_{R\delta}^{d_0} h^{p-2} \, dh + o(1) \leq C' R^{p-1} + o(1).$$

The bound follows from the lower density of order $b^p h^p$ and the first-primitive shift bound $Cb^p\delta h^{p-1}$. The far side of this crop is empty for every z.

At the opposite cusp, when $q < 1/2$, use marks $x = (v - s_0)/\delta$ and again retain $|x| \le R$. The density is bounded below by cb^q around this cusp. Using $Cb^q\delta h^{q-1}$ for the shift discrepancy at depth $h \ge R\delta$ gives discarded cost at most

$$CM_q\delta^2 \int_{R\delta}^{d_0} h^{2q-2}\, dh + o(1) \le C' R^{2q-1} + o(1).$$

When $q = 1/2$, instead retain $\Delta \le s_0 - v \le d_0$, where $\Delta = \delta \log^2(1/\delta)$ and d_0 is fixed and small. The omitted near-cusp cost is bounded by $CM_q\delta^2[1 + \log(\Delta/\delta)] = o(1)$. This uses discrepancy $O(b^q\delta^q)$ over a strip of width $O(\delta)$ and $Cb^q\delta h^{q-1}$ beyond it.

The discarded observations can be regenerated in their original slots from the normalized reference density on the discarded region. Its mass tends to one under every z. For two measures of mass at least $1/2$, normalization increases their Hellinger distance by at most a fixed factor, by the reverse triangle inequality for the square-root norms. The independent conditional draws thus yield total variation error bounded by a constant times the square root of the discarded product Hellinger cost. The reverse kernel is the actual censoring map.

5.3 Limit of the retained experiment

The own stream converges in total variation for fixed R to the cropped Poisson process. In fact its per-observation mark density multiplied by N tends in $L^1([-R, R])$ to $\tau A_p(x + z)_+^p$, and the binomial rare-event count can be replaced by a Poisson count with vanishing error.

Condition on the own marks and count k. The remaining $N - k$ observations are independent draws from the conditional law outside that crop, with the precise opposite values retained there and all other values censored to one common symbol. Ordering labels are ancillary and can be regenerated by uniform permutation. The count k is tight uniformly over the finite parameter set. The rate statements for the censored law below hold for every fixed bounded k, hence uniformly on any such set. The conditioning normalizers for the opposite densities differ from one by $O(N^{-1})$, uniformly in z. Their contribution to the squared integrals of relative increments multiplied by N vanishes. The censoring atom has mass tending to one. Its probability difference is the negative integral of the continuous part difference. By Cauchy–Schwarz its squared contribution multiplied by N vanishes whenever the continuous weighted squared integral multiplied by N is bounded, since the retained opposite probability is $o(1)$. It supplies the centering in the likelihood expansion.

For $q < 1/2$, in the bounded opposite mark window the scaled density $b^{-q}p_z(v)$ tends uniformly to B_0. The density difference is

$$p_z(s_0 + x\delta) - p_0(s_0 + x\delta) = -b^q\delta^q A_q[(-z - x)_+^q - (-x)_+^q] + o(b^q\delta^q)$$

uniformly on bounded marks. The smooth difference is $O(b^q\delta) = o(b^q\delta^q)$. Thus the relative increments tend uniformly to zero and their N-fold covariances converge to $\int_{-R}^{R} \mu_z(x)\mu_{z'}(x)\, dx$. This is the Gaussian white-noise shift on the crop.

For $q = 1/2$, write $h = s_0 - v$ on the retained region. The reference density divided by b^q differs from B_0 by $O(\sqrt{h}) + O(b)$ uniformly. The leading perturbation is $z\delta b^q q A_q h^{-1/2}$, with uniform remainder bounded by $Cb^q(\delta + b\delta h^{-1/2} + \delta^2 h^{-3/2})$. This follows from $J'_{q,b}(h) = w_q(bh)h^{-1/2}$, its derivative bound and the $O(\delta)$ change of amplitude. The weighted squared integral of the remainder, times N, is $o(1)$ by the cusp information hypothesis and $\delta/\Delta = o(1)$. Relative increments are uniformly $o(1)$. The N-fold leading covariance is $zz'\lambda(qA_q)^2/B_0$ since $\log(d_0/\Delta) \sim \log(1/\delta)$. This is the scalar Gaussian shift.

For either retained Gaussian block, relative likelihood increments are centered and tend uniformly to zero. The triangular central limit theorem gives the centered normal vector of sums, its quadratic term concentrates, and the cubic log remainder vanishes. The limit likelihood coordinates have mean one. For clarity, convergence in law of a finite likelihood vector and convergence of its coordinate means to one imply convergence of experiments here. Truncate to a large cube with small likelihood mass outside and quantize it into small boxes whose boundaries are null under the limit reference. The induced finite probabilities converge under every parameter. Under each reference, regenerating observations conditional on the quantization loses at most the mean error of the likelihood vector approximation, which tends to zero with finer boxes and a larger cube. This proves the conditional Le Cam conclusion uniformly for bounded k, and hence the joint product conclusion for the cropped experiment.

Finally let $R \to \infty$. Reference Poisson marks fill the omitted tails with squared-root intensity loss $O(R^{p-1})$. For subcritical Gaussian cusp marks, reference white noise fills the omitted tails with L^2 signal loss $O(R^{2q-1})$. These estimates follow from the same derivative bounds for the limiting power profiles. Combining with the censoring error establishes Theorem 5.1.

5.4 Critical cusp information with unknown interior

Here only assume $q = 1/2$, $m = M_q \to \infty$, and use the full model with an arbitrary fixed admissible interior. In the cusp window, write its exact density as

$$b^q \left\{ B_b(v) - (2a)^{-1} J_{q,b}((s-v)_+) \right\}, \qquad s = -T.$$

The smooth backgrounds have uniformly bounded derivatives of every fixed order on a fixed cusp neighborhood. The quantities $B_b(s)$ converge to a positive $B_\infty(s)$, by the integral formula with w_q replaced by $w_q(0)$. They are uniformly bounded above and away from zero. Let $C_0 = w_q(0)/(2a)$.

Theorem 5.2. *At an interior pair (a, c), there is an estimator using the opposite cusp with*

$$\sqrt{m \log m}\,(\widehat{T} - T) \implies N(0, 2B_\infty(s)/C_0^2).$$

The variance equals the inverse information of the critical cusp block at the same normalization. At the uniform base on cusp-only sequences in Theorem 5.1 this is the inverse information of the full local experiment along the matched path.

Use $O(N/\log m)$ pilot observations in two independent parts for the clipped maximum and two fixed-depth counts from the strongest boundary. The uniform-exterior pilot of Section 4 gives errors in $\widetilde{s}, \widehat{a}$ of squared mean $O((\log m)/m)$. Its Bernstein bound makes $|\widetilde{s} - s| \le \Delta/4$ except with probability $O(\exp(-c \log^3 m))$, with $\Delta = m^{-1/2}(\log m)^2$. On the remaining observations use $x = \widetilde{s} - v$ and the weight $x^{-1/2}\mathbf{1}_{\{\Delta \le x \le H\}}$ corrected to annihilate moments of orders zero through k. The correction is supported on $[H/2, H]$. For each fixed k the polynomial Gram construction in Section 4 gives supremum norm $O_k(H^{-1/2})$, with L^2 norm $O_k(1)$. First take $H = m^{-\epsilon}$ for fixed $0 < \epsilon < 1/2$ and k with $\epsilon(k + 3/2) > 1/2$. The residual background mean divided by b^q is $O_k(H^{k+3/2}) = o(m^{-1/2})$. Using $\widehat{a}$ in the exact cusp mean costs $O_k(b^q H|\widehat{a} - a|)$, negligible on the fluctuation scale $b^q\sqrt{\log m/m}$. The absolute derivative of the cusp mean on the pilot bracket of radius $\Delta/2$ is asymptotic to $b^q C_0 \log(H/\Delta)$ uniformly there. This follows from the derivative of J and $|S - \widetilde{s}| \le \Delta/2$, since replacing $(x + S - \widetilde{s})^{-1/2}$ by $x^{-1/2}$ under the uncorrected integral costs $O(1)$. The correction also costs $O_k(1)$ in the derivative integral. Similarly the weight variance is asymptotic to $b^q B_b(s) \log(H/\Delta)$, using continuity of the smooth background and the $O(\sqrt{x})$ cusp. The squared mean is negligible. The maximum weight is $O_k(\Delta^{-1/2})$, hence the conditional

triangular central limit theorem applies. Inverting the fitted mean, with clipping inactive with probability tending to one, yields variance $B_\infty(s)/(C_0^2(1/2-\epsilon))$ after normalization by $\sqrt{m\log m}$. Clipping is inactive because the true point is at distance at least $\Delta/4$ from the bracket edges and the estimation error is $o_{\Pr}(\Delta)$.

For the stated constant choose a deterministic sequence $\epsilon_j \downarrow 0$ and fixed k_j satisfying the strict bias condition at each stage. All bounds just used, including the background derivative bounds, hold uniformly over the brackets and standardized densities, with constants known from the model. Choose increasing deterministic sample thresholds for m so that each stage has relative errors at most $1/j$ in the derivative and variance estimates, and its normalized bias and Lindeberg bound are at most $1/j$. The weight corrections and the amplitude substitution are included in those thresholds. Also require $\log(H/\Delta)/\log m$ within $1/j$ of $1/2-\epsilon_j$. Thus a single slowly varying choice achieves the claimed normal limit. At $\delta = (m\log m)^{-1/2}$ the critical lambda of Theorem 5.1 is $1/2$, giving exactly the constant in Theorem 5.2.

6 Calibrating a noise endpoint from a logarithmic density singularity

A compact latent density and an affine noise density can have opposing power singularities whose convolution diverges logarithmically at an interior point. We locate the inducing noise endpoint at the sharp rate $(\log M)/M$ when the opposite window governs, with effective sample size M. The matched local experiment has a two-sided Brownian log likelihood with an explicit coefficient. We also determine when a moving support frontier gives a faster rate and obtain its Poisson limit. The logarithmic calculation also gives the exact rate and local experiment for direct location estimation in the critical bilateral gamma family.

6.1 Model and conclusions

Let X have known density g on $[L,U]$, smooth and positive inside. At inward depths h from its right and left endpoints, respectively, g is $w_p(h)h^{p-1}$ and $w_q(h)h^{q-1}$, where $p>0$, $0<q<1$, and the weights are smooth and positive near zero. Observe N independent copies of $Y = X + b(c+aZ)$, with independent X, Z and known $b \downarrow 0$. The parameters c, a lie in compact brackets with interior, with a bounded away from zero. The density of Z on $[-1,1]$ is f_0+u for a fixed known density f_0, smooth and positive on $(-1,1)$. Near 1 it equals $f_0(1-y) = k(y)y^{-q}$ for smooth positive k up to zero. Near -1 the reference is smooth and positive up to the endpoint. The unknown u has zero mass, is supported in a fixed strict interior interval, is bounded in a fixed Sobolev ball of nonnegative order and positive radius, and is uniformly bounded with a positive floor for f_0+u on its allowed support. The bounds have strict slack at $u=0$. Fix interior (a_0,c_0) for local analysis. The target is the right noise endpoint $T = c+a$. Let $M_p = Nb^p$, $M_q = Nb^q$, and suppose $\max(M_p, M_q) \to \infty$.

Theorem 6.1. *The minimax absolute and root mean squared errors for T have the same order, given by*

$$\begin{cases} M_p^{-1/(1+p-q)}, & p \le q, \\ (\log M_q)/M_q, & p > q. \end{cases}$$

The left noise endpoint, scale and shift may all be unknown in these upper bounds.

For the local conclusions hold the left noise endpoint fixed and set $T = T_0 + zd$ for a finite set of z containing zero, where d is the rate in Theorem 6.1. Match any fixed finite number of

physical noise moments to the reference at (a_0, c_0, f_0) by an interior perturbation, taking at least one moment and enough to have $2K_0 + 2 > \max(p, q)$.

Theorem 6.2. *Set $A_0 = a_0^{q-1} k(0) w_q(0)$. If $p \le q$, the finite local experiments converge in Le Cam distance to the Poisson process with intensity*

$$a_0^{q-1} k(0) w_p(0)\, \mathrm{B}(1-q, p)(x+z)_+^{p-q}\, dx.$$

A nonpositive exponent with a plus subscript means zero off the positive half-line. If $p > q$, the local log likelihood at reference zero converges to the Brownian experiment

$$\sqrt{\mathcal{I}}\, \mathcal{B}(z) - \frac{\mathcal{I}}{2}|z|, \qquad \mathcal{I} = \frac{\pi^2 A_0}{\sin^2(\pi q)},$$

where $\mathcal{B}$ is standard two-sided Brownian motion. The finite experiments converge in Le Cam distance.

Küchler and Tappe [17] give logarithmic growth and the difference of the one-sided finite parts for differences of gamma densities whose shape parameters sum to one. Ibragimov and Khas'minskii [16] study fixed location families with unbounded densities. Here a logarithmic pole arises at the opposite latent boundary in a shrinking convolution, with an unknown interior noise density and a simultaneously moving support frontier. The interior moment cancellation follows the affine calibration construction in Section 4, related to the Archive Gaussian inverse lower bounds [14] through their use of moment matching.

6.2 The interior logarithmic profile

Reflect the left latent endpoint, and let $v = (L - Y)/b$, $s = -T$, $h = s - v$. If y is the inward noise depth from its lower reflected endpoint, the latent kernel has argument $y + h$. Near the contact $h = 0$, the density of v divided by b^q has the form

$$P(h) = A(a)K(h) + R(h), \qquad A(a) = a^{q-1} k(0) w_q(0), \qquad K(h) = \int_0^\eta y^{-q} (y+h)_+^{q-1}\, dy,$$

where $\eta > 0$ is fixed and small, and $R(h) = R(0) + O(|h| \log(e/|h|))$, $R'(h) = O(\log(e/|h|))$ for $h \neq 0$, uniformly in the model. Indeed on $y \le \eta$ the product of the smooth weights minus its value at $(y, h) = (0, 0)$ is $O(y + |h|)$ with bounded derivatives. For $h > 0$, differentiating the residual under the integral bounds its derivative by a constant times $1 + K(h)$, since the extra factor $y + h$ absorbs one derivative. For $h < 0$ first put $t = y - |h|$ to start the kernel at $t = 0$. Then $y = t + |h|$, the weight error is $O(t + |h|)$, and one derivative on $(t + |h|)^{-q}$ is absorbed the same way. The moving upper limit contributes a bounded term. The part $y > \eta$ has bounded derivatives in h by separation from the zero of the latent kernel, including the unknown noise interior. The local residual is continuous at zero by dominated convergence on the respective half-lines in the indicated variables. Integrating the derivative proves the stated modulus.

For $h > 0$ use $t = y/(y + h)$ to write the kernel integral as $\int_0^{\eta/(\eta+h)} t^{-q}(1-t)^{-1}\, dt$. For $h < 0$ put $t = (y - |h|)/y$ to obtain $\int_0^{1-|h|/\eta} t^{q-1}(1-t)^{-1}\, dt$. The two formulas give

$$K(h) = \log(1/|h|) + j_\pm + O(|h|), \qquad K'(h) = -1/h + O(1),$$
$$j_+ = \log\eta - \gamma - \psi(1-q), \qquad j_- = \log\eta - \gamma - \psi(q),$$

where ψ is the digamma function and γ is Euler's constant. In particular $j_+ - j_- = -\pi\cot(\pi q)$. After removing an additive constant the rescaled local signal is

$$\ell(x) = \log(1/|x|) - \pi\cot(\pi q)\mathbf{1}_{\{x>0\}}.$$

For all real z, z', the differences of its translates are square integrable, and

$$\int_{\mathbb{R}} \big(\ell(x-z) - \ell(x-z')\big)^2\, dx = \frac{\pi^2}{\sin^2(\pi q)}|z-z'|. \tag{7}$$

To see this, the transform of $\log(1/|x|)$ away from the origin is $\pi/|\xi|$, and the transform of a step is $1/(i\xi)$, by their distributional derivatives. The two terms have perpendicular complex phases. The Fourier multiplier for a translation difference removes the common constant. Plancherel and $\int_{\mathbb{R}} |e^{i\xi d} - 1|^2 \xi^{-2}\, d\xi = 2\pi|d|$ give (7).

6.3 Moment matching and Hellinger costs

Translate the target by a small amount D while fixing the other endpoint. Match the first K_0 moments of the physical noise $c + aZ$ to their reference values. These equations prescribe standardized moments equal to those of $(c_0 + a_0 Z_0 - c)/a$. Their differences from the standardized reference moments are $O(|D|)$. A smooth compactly supported moment dual in the allowed interior solves them with a correction $O(|D|)$ in every fixed smooth norm. It is admissible by strict slack.

Outside the fixed scaled boundary windows, expand the latent density through K_0 in the physical noise position. Although moving the singular exterior can change the physical noise density in L^1 by more than $O(|D|)$, the Taylor remainder is smooth in the physical noise position. At physical depth $h \geq Cb$ from a latent endpoint of power r its supremum and the supremum of its derivative in the scaled noise position are $O(b^{K_0+1}h^{r-K_0-2})$. For the derivative, differentiate the Taylor polynomial and use the Taylor remainder of order $K_0 - 1$ for the derivative of the kernel. Evaluate both remainders against the known reference standardized density using $W = c + aZ$, and use the $O(|D|)$ fitted correction for the nuisance. The signed discrepancy is at most $C|D|b^{K_0+1}h^{r-K_0-2}$. The resulting Hellinger cost is $O(b^r D^2)$ after division by the floor ch^{r-1} and integration, since $2K_0 + 2 > r$. The same argument on a fixed interior region gives $O(b^{2K_0+2}D^2)$.

The right observation frontier, in $v = (Y - U)/b$, has exact density $b^p(T-v)_+^\alpha C_{b,a}((T-v)_+)$ near the tip, where $\alpha = p - q > -1$. Beta integration gives

$$C_{b,a}(0) = a^{q-1}k(0)w_p(0)\,\mathrm{B}(1-q, p).$$

The smooth positive factor and its derivatives are bounded uniformly near zero, with smooth dependence on a. This follows by rescaling the integral of $y^{-q}(T-v-y)^{p-1}a^{q-1}k(y/a)w_p(b(T-v-y))$ over $0 \leq y \leq T - v$. The translated frontier thus costs at most $Cb^p F_\alpha(|D|)$ in squared Hellinger distance, where

$$F_\alpha(d) = \begin{cases} d^{1+\alpha}, & -1 < \alpha < 1, \\ d^2\log(e/d), & \alpha = 1, \\ d^2, & \alpha > 1. \end{cases}$$

Indeed within distance $O(d)$ of the tips sum the densities, and farther inside bound the squared square-root difference by $Cb^p d^2 h^{\alpha-2}$. Outside the tip and its opposite moving logarithmic contact, stationary singularities and smooth regions cost $O((b^p + b^q)D^2)$. For the fixed frontier this uses the smooth relative change in its tip factor. Near the fixed inner cusp, split off the known noise exterior at the stationary endpoint using its inward coordinate. The exterior weight changes by

$O(|D|)$ relatively, by its smooth dependence on the scale and its positive taper. Separated noise positions can be evaluated with the standardized variable, where the latent kernel is smooth in position. Their density change is $O(|D|b^r)$ in the endpoint window of power r. The reference density is bounded below there by cb^r. On the remaining compact scaled regions away from the tip and the two contacts one can likewise split at fixed inward distances from the noise endpoints. Endpoint contributions use a smooth latent kernel or vanish, and the bounded smooth reference noise away from its endpoints can be differentiated in physical coordinates against the integrable latent kernel. The interior correction is $O(|D|b^r)$ there.

For the moving logarithmic contact, a change in parameters apart from translation can be ignored at cost $O(b^q D^2)$ in squared Hellinger distance. At a fixed aligned contact depth, the local noise weight in K changes by $O(|D|)$ relatively and the separated part changes by $O(|D|)$ in scaled density. The latter follows by using a smooth cutoff in the contact depth on the noise side and evaluating its complement in standardized noise coordinates. These changes are bounded by $C|D|$ times the reference profile, which is bounded below near the contact and has integrable logarithmic growth. It remains to compare the same profile at h and $h - D$. For $d = |D|$ and fixed bounded $x = h/d$ away from the two poles, the difference divided by $A(a_0)$ converges to $\ell(x - \operatorname{sgn} D) - \ell(x)$ and both densities are asymptotic to $A(a_0)\log(1/d)$ after division by b^q. The local remainder and its estimates are uniform in b. On bounded mark windows the squared Hellinger integrand, normalized by its prefactor $b^q d/\log(1/d)$, is dominated by a constant times the sum of one and the squared logarithms of the distances to the two poles. The denominator in the squared-root identity is bounded below by $c\log(1/d)$ there. Outside $|h| \le Rd$ but within the fixed neighborhood, the derivative bound $C/|h|$ and density floor $c(1 + \log(1/|h|))$ bound the tail by

$$Cb^q d^2 \int_{Rd}^{d_0} \frac{h^{-2}}{1 + \log(1/h)}\, dh.$$

Splitting this at $h = \sqrt{d}$ gives a normalized upper bound $C/R + o(1)$. Hence the contact alone has squared Hellinger cost

$$b^q \frac{d}{\log(1/d)} \left(\frac{\pi^2 A_0}{4\sin^2(\pi q)} + o(1) \right). \tag{8}$$

For $p \le q$, $d = M_p^{-1/(1+p-q)}$ gives a bounded $Nb^p F_\alpha(d)$, with the logarithmic contact cost and all regular costs tending to zero because $M_q \le M_p$ and $M_p d \le 1$. For $p > q$, set $d = (\log M_q)/M_q$. The frontier cost tends to zero because $M_p \le M_q$ and $\alpha > 0$, and the contact cost stays bounded. In either case taking a sufficiently small fixed multiple of d makes the product Hellinger distance small, giving the minimax lower bounds.

6.4 Upper bounds

If $p \le q$, use the clipped maximum of $(Y - U)/b$. The frontier tail mass at depth t has order $b^p t^{1+p-q}$ for small t, so the gap has tail bounded by $\exp(-cM_p t^{1+p-q})$ up to a small fixed depth, and exponentially small probability of a larger gap. Integrating proves the squared risk bound.

For $p > q$, first localize s to a fixed small interval with error probability $O(e^{-cM_q})$. The strongest latent window is at L. Its outward noise tail is known with fixed positive bounds, so the two-depth pilot of Section 2 applies verbatim. It uses the clipped outer maximum on one sample split and the two known tail counts on a separate split to fit the outward endpoint and a; the log-probability map has a uniformly Lipschitz inverse on the comparison rectangle. Bernstein's inequality at fixed tolerance gives the stated probability. This yields the inward endpoint in that window by subtracting twice the fitted scale.

Choose a smooth bounded compactly supported function φ satisfying

$$\int \varphi(u)\,du = 0, \qquad \int \varphi(u)\ell(-u)\,du \neq 0.$$

Such a function exists since the profile is nonconstant and locally integrable. At a candidate center z, estimate the mean of $\varphi((v-z)/h)$, divided by $b^q h$. For $|z-s| \le \epsilon h$ its absolute expectation is at least $2\lambda_0 > 0$ for a fixed λ_0 when ϵ and h are small. This follows from the profile expansion, cancellation of the constants and local translation continuity of the logarithm in L^1. For a center at distance at least Ah from s and a probe supported in the fixed contact region, its expectation is at most $\lambda_0/2$ in absolute value for large fixed A, by cancellation and the $C/|v-s|$ derivative bound. The variance of the unscaled weight is at most $Cb^q h \log(e/h)$, since the integral of the logarithmic pole over any interval of length $O(h)$ is at most $Ch\log(e/h)$. Bernstein bounds the probability of a fixed normalized error by

$$2\exp\{-cM_q h/\log(e/h)\}.$$

Starting at a small fixed bandwidth h_0, halve the bandwidth at each stage and search a lattice of spacing ϵh within distance $K_1 h$ of the previous estimate, where $K_1 > 2A+2$. Initialize from the pilot within $h_0/2$ of the true contact. Take a candidate whose absolute statistic exceeds λ_0, or retain the previous estimate if none does. On success through the preceding stage, all candidates therefore lie in a fixed deterministic set of bounded size within $O(h)$ of the true contact, and one lies within ϵh. Thus, even with the same data at all stages, the unconditional probability of a first failure at h is at most $C\exp(-cM_q h/\log(e/h))$. After the first failure at h the final error is $O(h)$ by the geometric bound on the size of all later moves. Stop at $h_f \asymp (\log M_q)/M_q$. The sum of the squared failure costs is $O(\sum_{h \ge h_f} h^2 \exp[-cM_q h/\log(e/h)]) = O(h_f^2)$, since $M_q h/\log(e/h) \gtrsim h/h_f$ for the dyadic range. The pilot failure contributes exponentially small bounded loss after clipping. This gives the upper rate.

6.5 Local experiments

For $p \le q$ retain the right frontier marks $x = (T_0 - v)/d$ in $[-R, R]$, with $R > \max_z |z| + 1$. The scaled intensity converges in L^1 by the tip formula and integrability of the translated power on a bounded interval. The count is binomial with probability $O(1/N)$, giving the stated Poisson limit on the crop. The omitted tip squared Hellinger cost multiplied by N is at most $CR^{\alpha-1}$, by integrating $h^{\alpha-2}$ from Rd in the tip. All other costs vanish as above. One may fill the omitted observations from the normalized reference exterior conditional on their count. Normalization costs a fixed factor in Hellinger distance since its mass tends to one. Fill the omitted Poisson marks from the reference intensity with the same square-root tail estimate, and then send R to infinity.

For $p > q$ retain opposite contact marks $x = (s_0 - v)/d$ in $[-R, R]$, excluding radius ϵ around zero and all candidate pole locations z. Censor the rest to one symbol. Along this path $s = s_0 - zd$, and on the retained mark set, the single-observation scaled density in v divided by b^q is $A_0 \log(1/d)(1+o(1))$ uniformly under each parameter. Its change from the reference divided by b^q is $A_0[\ell(x-z) - \ell(x)] + o(1)$ uniformly there. The relative increments therefore vanish uniformly. The censoring symbol restores zero mean and makes a vanishing contribution to the N-fold covariance. Indeed its probability change is the negative integral of the density change over the retained region, whose mass tends to zero, so Cauchy–Schwarz bounds the contribution by a vanishing multiple of the continuous covariance bound. Since $M_q d/\log(1/d) \to 1$, the log likelihood expansion and triangular central limit theorem give unit white noise on the crop with signal $\sqrt{A_0}[\ell(x-z) - \ell(x)]$. The quadratic term concentrates and the cubic term is negligible by the uniform relative increment bound. The finite likelihood vectors converge with mean-one limit coordinates. Truncating and

quantizing them under the reference gives convergence in Le Cam distance, as in Section 5. The omitted Hellinger costs vanish as $R \to \infty$ and $\epsilon \to 0$, by the tail and squared-log domination in the proof of (8). The same bounds fill the omitted Gaussian white-noise regions. The variogram (7) and polarization identify the full centered Gaussian log likelihood as the Brownian motion in Theorem 6.2, with drift equal to minus half the variance.

6.6 Direct location estimation at the variance-gamma pole

Let $0 < q < 1$ be fixed. Observe N copies of $Y = \mu + G_+ - G_-$, where independent gamma variables G_+, G_- have shapes $1-q, q$ and rates λ_+, λ_-. The location μ lies in a fixed compact interval with interior and the rates lie in fixed compact positive brackets. The case $q = 1/2$ is the critical univariate variance-gamma family. Nitithumbundit and Chan [18, Sections 3.2 and 5.3] simulate this case. The power-pole theorem they use below the critical shape leaves the logarithmic case outside its hypotheses. The following result concerns the experiment and the optimal estimation rate. Their leave-one-out estimator is a separate procedure.

Theorem 6.3. *The minimax absolute and root mean squared errors for μ have order $(\log N)/N$, including with unknown gamma rates. At fixed rates, the finite local experiments for $\mu = \mu_0 + z(\log N)/N$ converge to the Brownian likelihood in Theorem 6.2 with*

$$\mathcal{I} = \frac{\pi^2}{\sin^2(\pi q)} \frac{\lambda_+^{1-q} \lambda_-^q}{\Gamma(1-q)\Gamma(q)}.$$

In the symmetric standard case $q = 1/2$ and $\lambda_+ = \lambda_- = 1$, the coefficient is π.

Write $x = Y - \mu$. For $x > 0$, the density is

$$A e^{-\lambda_+ x} \int_0^\infty (y+x)^{-q} y^{q-1} e^{-(\lambda_+ + \lambda_-)y}\, dy, \qquad A = \frac{\lambda_+^{1-q} \lambda_-^q}{\Gamma(1-q)\Gamma(q)}.$$

For $x < 0$ exchange the positive and negative gamma variables. The local analysis in Section 2 applies after the change of orientation $h = -x$: splitting both integrals at a small fixed depth leaves the logarithmic coefficient A, the same log-plus-step profile up to reflection, and a continuous residual whose derivative is $O(\log(e/|x|))$. For example, a common cutoff on the positive gamma coordinate gives the two integrals in $K(h)$. On each side the exponential factors are smooth with value one at the contact and the tail of the integral and its derivatives converge exponentially. These statements are uniform on the rate brackets. Away from zero the logarithmic derivative of the density is bounded on each outer ray, by differentiation of the displayed integral and the bound $(y + |x|)^{-1} \le |x|^{-1}$. Hence translation outside a fixed contact window has squared Hellinger cost $O(D^2)$. Now (8) applies with $b^q = 1$ and $A_0 = A$. It proves the two-point lower bound, and the contact-only likelihood proof with $M_q = N$ gives the Brownian limit.

For the upper bound with unknown rates, a fixed-accuracy location pilot comes from the first three raw moments. Put $u = 1/\lambda_+$, $v = 1/\lambda_-$. The mean, variance V, and third cumulant C_3 are

$$\mu + (1-q)u - qv, \qquad V = (1-q)u^2 + qv^2, \qquad C_3 = 2[(1-q)u^3 - qv^3].$$

The last two determine u, v with a uniformly Lipschitz inverse. For any two parameter pairs put $\Delta_+ = (1-q)(u^2 - u'^2)$, $\Delta_- = q(v^2 - v'^2)$. The changes of $V, C_3/2$ are $\Delta_+ + \Delta_-$ and $B_+\Delta_+ - B_-\Delta_-$, where each $B_\pm$ is a secant slope of $t^{3/2}$ on a compact positive bracket of squared scales. Both slopes are positive and bounded above and away from zero, so the inverse is bounded uniformly.

The mean then determines μ. To obtain exponential fixed-tolerance concentration, truncate the raw moment integrands at a sufficiently large fixed radius and choose an approximate fit to their empirical values using the untruncated model moments. Exponential tails on the brackets make the deterministic truncation bias uniformly as small as needed, and Hoeffding's inequality controls the bounded empirical integrands. The resulting pilot is within any fixed small tolerance of μ, except with probability $O(e^{-cN})$. On independent data the logarithmic multiscale estimator applies with $b^q = 1$, using the known profile ℓ up to reflection. Its bounds only require a uniform positive lower bound on A, and therefore give the claimed rate with unknown rates.

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, *Location estimation at variance-gamma logarithmic poles*, Luminary Research Archive, 2026.

[2] S. Schubr, *Recovering small additive noise from boundaries, vertices and corners*, Luminary Research Archive, 2026.

[3] S. Schubr, Sharp transport identification with unknown Gaussian covariance, Luminary Research Archive, edition 3 (2026). `https://research.luminaryoddities.org/sharp-transport-identification-unknown-gaussian-covariance`.

[4] A. Delaigle and P. Hall, Methodology for non-parametric deconvolution when the error distribution is unknown, J. Roy. Statist. Soc. B 78 (2016), 231–252. DOI.

[5] 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.

[6] P. Hall and L. Simar, Estimating a changepoint, boundary, or frontier in the presence of observation error, J. Amer. Statist. Assoc. 97 (2002), 523–534. DOI.

[7] C. Matias, Semiparametric deconvolution with unknown noise variance, ESAIM Probab. Stat. 6 (2002), 271–292. Article.

[8] S. Schubr, *Sharp Wasserstein Stability for Gaussian Deconvolution of Log-Concave Laws*, Luminary Research Archive, edition 4. Archive edition.

[9] S. Schubr, *Recovering small additive noise from boundaries, vertices and corners*, the part titled *Recovering joint noise from corners of product densities*, Luminary Research Archive, 2026.

[10] S. Schubr, *Recovering small additive noise from boundaries, vertices and corners*, the part titled *Critical Wasserstein rates for one-dimensional endpoint noise*, Luminary Research Archive, 2026.

[11] S. Schubr, *Recovering small additive noise from boundaries, vertices and corners*, the part titled *Information scales for mixed noise moments*, Luminary Research Archive, 2026.

[12] R. L. Smith, Maximum likelihood estimation in a class of nonregular cases, Biometrika 72 (1985), 67–90. DOI.

[13] B. L. S. Prakasa Rao, Estimation of the location of the cusp of a continuous density, Ann. Math. Statist. 39 (1968), 76–87. DOI.

[14] S. Schubr, *Sharp Wasserstein Stability for Gaussian Deconvolution of Log-Concave Laws*, Luminary Research Archive, edition 4 (2026). Archive edition.

[15] O. V. Chernoyarov, S. Dachian, and Yu. A. Kutoyants, On parameter estimation for cusp-type signals, Ann. Inst. Stat. Math. 70 (2018), 39–62. DOI.

[16] I. A. Ibragimov and R. Z. Khas'minskii, Asymptotic behavior of statistical estimates of the shift parameter for samples with unbounded density, J. Soviet Math. 16 (1981), 1035–1041. DOI.

[17] U. Küchler and S. Tappe, On the shapes of bilateral Gamma densities, Statist. Probab. Lett. 78 (2008), 2478–2484. Author manuscript.

[18] T. Nitithumbundit and J. S. K. Chan, Maximum leave-one-out likelihood method for the location parameter of variance gamma distribution with unbounded density, J. Stat. Comput. Simul. 93 (2023), 2642–2671. DOI.