Project Repository

When Does a Random-Scale Flight Path Look Lognormal?

A spherical bridge note from OpenSky trajectories

Jonathan R. Landers

June 2026

Abstract

An hour of globally sampled aircraft positions produces a tempting picture: RMS departures from endpoint great-circle arcs are positive, strongly right-skewed, and superficially lognormal. We turn that visual hypothesis into a geometric model. Each deviation is factored as $D = AQ$, where A sets route scale and Q is the RMS shape of an endpoint-pinned bridge. We prove that an independent lognormal scale suppresses the non-Gaussian cumulants of $\log Q$ at explicit rates and give a Wasserstein convergence criterion. For a Brownian bridge, the criterion yields a concrete threshold. Yet the observed log-skewness and kurtosis sharply contradict the resulting single-population prediction. The same theorem that explains a lognormal-looking center therefore diagnoses a distinct near-geodesic boundary population, and a sequential small-ball test estimates when that assignment becomes visible in time. The result is not a universal lognormal law, but a route from visual regularity to a falsifiable mixture model and a time-to-diagnosis statistic.

A visual regularity with a geometric question

The OpenSky Network makes large-scale ADS-B trajectory research possible [1]. We sampled 60 one-minute global state snapshots, retained aircraft with all 60 positions, required at least 25 km between endpoints, and removed any track with a one-minute jump above 30 km. The final cohort contains 1,780 tracks. These positions are receiver-dependent observations, not a complete census of global aviation.

For aircraft i, write $x_i(t) \in S_R^2$ for its position on the spherical Earth at normalized time $t \in [0, 1]$. The direct reference is the minor endpoint geodesic, a standard spherical construction in directional statistics [2],

$$g_i(s) = \mathrm{Exp}_{x_i(0)}\Big(s\,\mathrm{Log}_{x_i(0)}\, x_i(1)\Big), \qquad s \in [0, 1]. \tag{1}$$

Here s is dimensionless progress, and $g_i(s)$ is a position rather than a distance. If $d_{S_R^2}$ is spherical distance in kilometres, define cross-track displacement and its sampled RMS by

$$\eta_i(t) = d_{S_R^2}(x_i(t), g_i([0, 1])), \qquad D_i = \left[\frac{1}{60}\sum_{j=1}^{60} \eta_i(t_j)^2\right]^{1/2}. \tag{2}$$

The median D_i is 16.7 km and the mean is 22.7 km. Figure 1 shows why a lognormal model is initially plausible and why that impression is incomplete.

Classical lognormal modeling begins with $\log(D/d_0) \sim N(\mu, \sigma^2)$ for reference distance $d_0 = 1$ km [3]. That describes a shape but supplies no route geometry. We instead ask: *what mechanism makes a spherical path statistic look lognormal, what would falsify it, and how long must one watch before the falsifying boundary component is recognizable?*

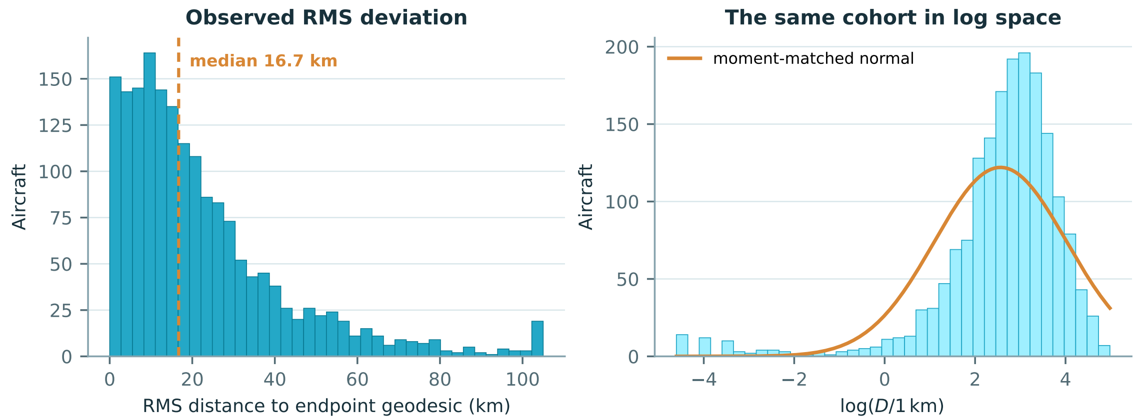

Figure 1: RMS spherical deviation for the quality-filtered cohort. Left: the application histogram, truncated visually at the 99th percentile while retaining tail observations in the end bin. Right: the 1,776 positive values in log space with a moment-matched Gaussian. The broad center looks compatible with lognormality, but the long lower tail does not.

Narrative arc.

- **A direct reference.** Aircraft rarely follow the shortest spherical path between observed endpoints, so the endpoint geodesic is a reference rather than an operational route.
- **Scale and shape.** Each RMS deviation is written as $D = AQ$: A sets overall magnitude, while Q captures standardized endpoint-pinned shape.
- **The smoothing prediction.** Broad variation in A should suppress the skewness and kurtosis contributed by Q, making the Brownian product nearly Gaussian in log space.
- **The boundary population.** The observed log-skewness is −2.40, with 53 tracks at or below 0.1 km RMS deviation forming a distinct lower-tail boundary.
- **Sequential diagnosis.** After the mixture is forced, prefix residuals and Wasserstein gaps ask how long one must watch before assigning a single aircraft to that boundary component.

Scale, shape, and the bridge factor

Factor the observed departure into a kilometre-valued scale and a dimensionless shape,

$$\eta_i(t) = A_i Z_i(t), \qquad D_i = A_i Q_i, \qquad Q_i = \left[\int_0^1 Z_i(t)^2\, dt\right]^{1/2}, \tag{3}$$

where $Z_i(0) = Z_i(1) = 0$. Intuitively, A_i asks *how large* the route departure is, while Z_i asks *what shape* it takes between fixed endpoints. Across a population, independent densities f_A and f_Q induce the product law

$$f_D(d) = \int_0^\infty f_A(d/q) f_Q(q)\, \frac{dq}{q}. \tag{4}$$

The Jacobian $1/q$ is the signature of multiplicative rather than additive mixing. A single lognormal is therefore not the starting mechanism; it is a possible approximation to this geometry-driven product.

When the product becomes lognormal

Set $U_\ell = \log(A_\ell/d_0)$ and $V_\ell = \log Q_\ell$ for a sequence of scale-shape populations. The index ℓ labels increasing scale dominance, not time or aircraft.

Theorem 1 (Lognormal smoothing)**.** *Suppose $U_\ell \sim N(\mu_\ell, \sigma_\ell^2)$ is independent of V_ℓ, with $m_\ell = \mathbb{E}V_\ell$ and $\tau_\ell^2 = \mathrm{Var}(V_\ell) < \infty$. Define*

$$W_\ell = \frac{\log(D_\ell/d_0) - (\mu_\ell + m_\ell)}{\sqrt{\sigma_\ell^2 + \tau_\ell^2}}.$$

If $\tau_\ell^2/(\sigma_\ell^2+\tau_\ell^2)\to 0$*, then* $W_\ell \Rightarrow N(0,1)$*. More precisely, for* $G_\ell \sim N(m_\ell, \tau_\ell^2)$*,*

$$d_{W_2}(W_\ell, N(0,1)) \le \frac{d_{W_2}(V_\ell, G_\ell)}{\sqrt{\sigma_\ell^2+\tau_\ell^2}},$$

where d_{W_2} *is quadratic Wasserstein distance [4].*

Proof. Replace V_ℓ by its Gaussian moment match G_ℓ. Then $U_\ell + G_\ell$ is exactly Gaussian with the displayed centering and variance. Convolution with the same independent U_ℓ contracts W_2. Finally, independent coupling gives $d_{W_2}(V_\ell, G_\ell) \le \sqrt{2}\tau_\ell$, so the stated variance ratio forces convergence. □

The theorem formalizes a simple picture: if fleet-to-fleet variation in route scale is broad enough, it blurs the finer distributional signature of route shape.

Proposition 1 (Exact cumulant attenuation)**.** *For one fixed population, let* $\rho = \tau_Q^2/(\sigma_A^2+\tau_Q^2)$*, where* $\sigma_A^2 = \mathrm{Var}(\log(A/d_0))$ *and* $\tau_Q^2 = \mathrm{Var}(\log Q)$*. For every* $r \ge 3$*,*

$$\kappa_r(W) = \frac{\kappa_r(\log Q)}{(\sigma_A^2+\tau_Q^2)^{r/2}}.$$

Consequently,

$$\mathrm{Skew}(W) = \mathrm{Skew}(\log Q)\rho^{3/2}, \qquad \mathrm{ExKurt}(W) = \mathrm{ExKurt}(\log Q)\rho^2.$$

Proof. Independent cumulants add, and a Gaussian has no cumulants above order two. Standardizing by the total log-variance gives the first identity; writing the third and fourth cumulants in standardized form gives the second line. □

These identities turn "approximately lognormal" into a checkable statement: shape-induced skewness decays as $\rho^{3/2}$ and excess kurtosis as ρ^2.

The Brownian benchmark and its failure

A standard Brownian bridge is the simplest endpoint-pinned random shape. Its Karhunen–Loève expansion and RMS functional are

$$Z(t) = \sum_{k=1}^{\infty} \frac{\sqrt{2}\,\xi_k}{\pi k}\sin(\pi k t), \qquad Q^2 = \int_0^1 Z(t)^2 dt = \sum_{k=1}^{\infty}\frac{\xi_k^2}{\pi^2 k^2}, \tag{5}$$

with independent $\xi_k \sim N(0,1)$. This is the classical quadratic Brownian-bridge functional underlying Cramér–von Mises theory [5]; its Laplace transform is

$$\mathbb{E}e^{-\lambda Q^2} = \left(\frac{\sqrt{2\lambda}}{\sinh\sqrt{2\lambda}}\right)^{1/2}, \qquad \lambda > 0. \tag{6}$$

A two-million-draw, 500-mode numerical evaluation gives

$$\mathrm{SD}(\log Q) \approx 0.385, \qquad \mathrm{Skew}(\log Q) \approx 0.178, \qquad \mathrm{ExKurt}(\log Q) \approx -0.348.$$

Substitution in the proposition yields a concrete conclusion: $\sigma_A \gtrsim 0.50$ pushes both surviving standardized cumulants below 0.05 in magnitude. A moderately broad lognormal scale should therefore make this Brownian product look very nearly lognormal.

The observed cohort supplies a sharp diagnostic. Among the 1,776 positive D_i values,

$$\mathrm{SD}(\log(D/d_0)) = 1.465, \qquad \mathrm{Skew} = -2.40, \qquad \mathrm{ExKurt} = 8.11.$$

Variance matching under the independent Brownian model implies $\sigma_A \approx 1.413$. At that scale spread, the proposition predicts skewness about 0.003 and excess kurtosis about −0.002—essentially Gaussian in log space. The data are nowhere close. The mismatch is driven by a near-geodesic boundary: 53 tracks have recorded RMS deviation at or below 0.1 km.

From an attractive fit to a useful rejection

The mathematical arc reverses the usual modeling order. We did not choose a flexible mixture merely because a histogram fit was imperfect. Geometry first proposed $D = AQ$; the smoothing theorem then showed that the Brownian single-regime model should be *more* Gaussian in log space than the data; only then did a mixture become necessary.

A natural next model is the finite mixture form [6]

$$f_D(d) = \pi_0 f_0(d) + \sum_{c=1}^{C} \pi_c f_D(d \mid c), \qquad \pi_c \geq 0, \qquad \sum_{c=0}^{C} \pi_c = 1, \tag{7}$$

where f_0 represents the near-geodesic boundary and the continuous components represent distinct route-shape or flight-phase regimes. The current hour cannot identify all such classes, and endpoint geodesics are geometric counterfactuals rather than operationally feasible routes. Those limitations are substantive. Still, the note establishes a clean diagnostic principle: broad lognormal scaling can hide ordinary bridge-shape nonnormality, so a large residual departure is evidence for heterogeneity, dependence, or a boundary mechanism—not merely a slightly misspecified bridge.

Sequential diagnosis of the boundary population

Once the boundary component has been detected at the cohort level, the next question is individual and temporal: how long must one watch a trajectory before assigning it to that component? Let $C = 0$ denote the near-geodesic boundary population and $C = 1$ the ordinary bridge population. For the first n samples define the prefix residual

$$R_n^2 = \frac{1}{n}\sum_{j=1}^{n} \eta(t_j)^2, \qquad \mathcal{F}_n = \sigma\{\eta(t_1), \ldots, \eta(t_n)\}. \tag{8}$$

For a fixed tolerance τ, persistent small residuals are informative only when they are rare under the ordinary component.

Theorem 2 (Sequential small-ball diagnosis)**.** *Let* $\pi_0 = \mathbb{P}(C = 0)$*. Suppose that, for some* $\tau > 0$*,*

$$\mathbb{P}_1(R_n \leq \tau) \leq e^{-nI(\tau)} \qquad \text{and} \qquad \mathbb{P}_0(R_n \leq \tau) \geq 1 - \beta_n.$$

Then

$$\mathbb{P}(C = 0 \mid R_n \leq \tau) \geq \frac{\pi_0(1-\beta_n)}{\pi_0(1-\beta_n) + (1-\pi_0)e^{-nI(\tau)}}.$$

In particular, posterior confidence $1 - \alpha$ *is reached once*

$$n \gtrsim \frac{\log((1-\pi_0)/\pi_0) + \log((1-\alpha)/\alpha)}{I(\tau)}$$

up to the boundary miss term β_n*.*

Proof. Bayes' rule gives

$$\mathbb{P}(C = 0 \mid R_n \leq \tau) = \frac{\pi_0 \mathbb{P}_0(R_n \leq \tau)}{\pi_0 \mathbb{P}_0(R_n \leq \tau) + (1-\pi_0)\mathbb{P}_1(R_n \leq \tau)}.$$

The two displayed assumptions give the bound. Solving the same inequality for n gives the stated waiting-time scale. □

This is the time-series version of the mixture conclusion: not every straight-looking prefix is diagnostic, but persistent near-zero residual energy becomes decisive once ordinary bridges have too little small-ball mass left.

The same idea can be expressed in the transport language already used in the smoothing theorem. Let $Y_n = \log(R_n/d_0)$, and let $P_{0,n}$ and $P_{1,n}$ be the laws of Y_n under the boundary and ordinary components. Define

$$\Delta_n = W_2(P_{0,n}, P_{1,n}).$$

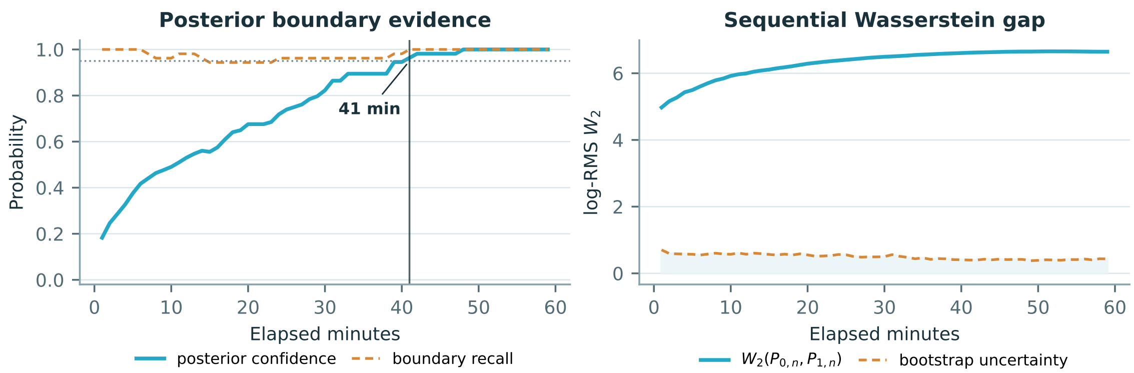

Figure 2: Sequential boundary diagnostics using the final endpoint geodesic as the reference. Left: empirical posterior confidence and boundary recall for the event $R_n \leq 0.1$ km. The prior boundary mass is $53/1780$, so early small residuals are not enough; confidence becomes high only after persistence. Right: the empirical Wasserstein gap between the log-prefix RMS laws of the boundary and ordinary components, plotted against a bootstrap uncertainty floor.

A conservative critical time is

$$T_W = \inf\{n : \Delta_n > r_{0,n}(\alpha) + r_{1,n}(\alpha) + m\},$$

where $r_{c,n}(\alpha)$ are sampling radii and m is a chosen margin. For one observed aircraft with value y_n, the point-to-law score

$$S_n(y_n) = W_2^2(\delta_{y_n}, P_{1,n}) - W_2^2(\delta_{y_n}, P_{0,n})$$

is positive when the observation is closer to the boundary law than to the ordinary law. Thus the practical stopping rule combines individual posterior evidence with population separation:

$$T_* = \max\{T_{\text{post}}, T_W\}, \qquad T_{\text{post}} = \inf\{n : \mathbb{P}(C = 0 \mid \mathcal{F}_n) \geq 1 - \alpha\}.$$

In the present hour, with $\tau = 0.1$ km and $\pi_0 = 53/1780$, the final-endpoint diagnostic crosses high posterior confidence after roughly forty-one elapsed minutes. The online problem without the final endpoint is harder, because many ordinary trajectories are locally straight; that distinction is exactly what the waiting-time theorem is meant to expose.

References

[1] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, "Bringing up OpenSky: A large-scale ADS-B sensor network for research," *Proc. 13th ACM/IEEE IPSN*, pp. 83–94, 2014. doi:10.1109/IPSN.2014.6846743.

[2] K. V. Mardia and P. E. Jupp, *Directional Statistics*. Wiley, 2000. doi:10.1002/9780470316979.

[3] J. Aitchison and J. A. C. Brown, *The Lognormal Distribution*. Cambridge University Press, 1957.

[4] C. Villani, *Optimal Transport: Old and New*. Springer, 2009. doi:10.1007/978-3-540-71050-9.

[5] T. W. Anderson and D. A. Darling, "Asymptotic theory of certain goodness of fit criteria based on stochastic processes," *Ann. Math. Statist.*, vol. 23, no. 2, pp. 193–212, 1952. doi:10.1214/aoms/1177729437.

[6] G. J. McLachlan and D. Peel, *Finite Mixture Models*. Wiley, 2000. doi:10.1002/0471721182.