Paper ID: FRC 100.002.001 | Version: v1.5

Accessible Phase-Space Geometry Predicts the Structure Functional in the Mixed Standard Map

Author: Hadi Servat

Affiliation: Independent Researcher

Contact: publish@fractalresonance.com

Website: fractalresonance.com

Abstract

FRC 100.002 defines the structure functional $\Sigma = S\_vM(C) - S$, equivalently the KL projection deficit from the measured phase marginal to the matched von Mises maximum-entropy member. Instrument A measures $\Sigma(\infty)$ from evolved ensembles over 12 seeds. Instrument B independently classifies the torus by finite-time Lyapunov exponent, flood-fills the accessible chaotic component, and predicts Σ from the component's phase marginal with no fitted amplitude. Over $K\_c \leq K \leq 2.0$, the original coarse sweep retains 9.7% median relative error, versus 18.9% for a fitted island-area proxy; fixed-area shuffle nulls reject random placement. A subsequently

> *frozen dense Gate 4 test selected the strongest geometry-curvature window automatically and evaluated ten unseen midpoint conditions with 12 new seeds each. There measured Σ has essentially no rank relation to regular area (ρ=0.030, one-sided permutation p=0.477), while negative entropy supplies a stronger monotonic rank baseline (ρ=0.964). The full geometry prediction has 8.7% median magnitude error but is beaten by the frozen area-only baseline at 6.6%. Estimator, seed, null, and convergence controls pass. The coarse magnitude correspondence remains a scoped result; the registered route to pillar-level fine geometric discrimination is not passed and is closed without window search.*

Paper ID: FRC 100.002.001. Answer-appendix line of FRC 100.002. The coarse-window magnitude result and geometry null remain, but the registered dense Gate 4 follow-up is negative. This version reports both at equal volume and does not promote the parent to pillar status.

1. The question this appendix decides

FRC 100.002 v2.6 §5.2 defines the **structure functional**

$$\Sigma(t) = S_{\mathrm{vM}}\big(C(t)\big) - S\big(x(t)\big) \geq 0, \tag{1}$$

with $C = |\langle e^{ix} \rangle|$ the mean resultant length (566.030), S the differential entropy of the phase marginal, and $S_{\mathrm{vM}}(C)$ the exact von Mises maximum-entropy curve of 566.030. If $q_{\mathrm{vM}}[C(p)]$ is the matched von Mises member, exponential-family projection gives

$$\Sigma(p) = D_{\mathrm{KL}}(p \,\|\, q_{\mathrm{vM}}[C(p)]). \tag{1a}$$

The inequality is therefore the standard maximum-entropy/KL projection theorem. The paper's added claim is not a new divergence; it is the zero-fitted-parameter prediction of this deficit from independently measured accessible geometry. The 100.002 pilot showed $\Sigma(\infty) > 0$ persists in the mixed regime and hypothesized (§5.4) it is anchored on the KAM hierarchy but deferred the direct comparison. This appendix asks two separable questions: does the *value* of $\Sigma(\infty)$ follow without tuning from the measured accessible geometry (magnitude); and is that value set by the *arrangement* of the regular region or merely its area (arrangement)?

2. Two operationally independent instruments

The instruments are **operationally independent in the load-bearing sense: the geometric predictor (Instrument B) never consumes the measured Σ (Instrument A).** They share the map, the drive K, the physical seed region, and the functional form (1); Instrument B contains no random numbers and reads no ensemble output.

2.1 Instrument A — $\Sigma(\infty)$, seed-averaged

Ensemble $x \sim \text{vonMises}(0, 8)$, $p \sim \mathcal{N}(0, 0.05)$, $N = 4 \times 10^4$, $T = 300$, under the Standard Map

$$p_{n+1} = p_n + K_{\text{drive}} \sin x_n, \qquad x_{n+1} = x_n + p_{n+1} \pmod{2\pi}. \tag{2}$$

$\Sigma(\infty)$ is the last-50-step mean of $S_{\text{vM}}(C) - S - (M-1)/2N$, $M = 72$ bins, with $(M-1)/2N = 8.9 \times 10^{-4}$ nats the analytic Miller–Madow bias (not fitted). Each K is run for 12 independent seeds; we report $\Sigma_{\text{meas}} \pm \text{SE}$.

Cylinder vs torus. Equation (2) reduces only x modulo 2π; Instrument A evolves on the cylinder and lets p diffuse, as required to measure the true marginal. The Standard Map is 2π-periodic in p, so its invariant-set structure is identical on the 2π-torus; Instrument B reduces $p \bmod 2\pi$ only to build the geometric projection. Same dynamics, different projection.

2.2 Instrument B — accessible phase-space geometry (zero fitted parameters)

On a torus grid $(x, p) \in [0, 2\pi)^2$, the finite-time Lyapunov exponent from the renormalized tangent map of (2),

$$\delta p' = \delta p + K_{\text{drive}} \cos x \, \delta x, \qquad \delta x' = \delta x + \delta p', \tag{3}$$

over T_{ly} steps, classifies a cell as **regular** if $\lambda < \lambda_{\text{thr}}$; $f_{\text{island}} = 1 - \langle \mathbb{1}[\lambda \geq \lambda_{\text{thr}}] \rangle$. A periodic flood-fill (union-find across torus boundaries) isolates the chaotic component(s) touching the ensemble seed region. **The seed region uses wrapped toroidal distance:** a grid point p (represented in $[0, 2\pi)$) is in the seed if $d_{\mathbb{T}}(p, 0) = \min(p, 2\pi - p) < 0.2$, and likewise $d_{\mathbb{T}}(x, 0) < 0.5$ — implemented as `(g<0.2)|(g>2π−0.2)`, so the region straddling

$p = 0$ (where $p \sim \mathcal{N}(0, 0.05)$ actually lives) is captured. Verified in code (`accessible_component`). Assuming a uniform distribution on the accessible component, its x-marginal $\rho(x)$ gives

$$C_{\text{pred}} = \Big|\sum_x \rho(x) e^{ix}\Big|, \quad S_{\text{pred}} = -\sum_x \rho \ln \rho + \ln \tfrac{2\pi}{G}, \quad \Sigma_{\text{pred}} = S_{\text{vM}}(C_{\text{pred}}) - S_{\text{pred}}.$$

Zero fitted parameters, not choice-free. No fitted amplitude enters (4). There are declared operational choices — λ_{thr}, T_{ly}, flood-fill connectivity, seed region, grid G, and the uniform-component assumption — each fixed before scoring; convergence in the first five is in §4. **Two assumptions are probed by two distinct tests, which must not be conflated:** the *uniform-fill* assumption (that the late ensemble fills its accessible component uniformly) is tested by the sub-K_c dynamical failure (§6); whether *spatial arrangement* matters given the excluded area is tested by the shuffle null (§5). They are different questions.

3. Result: magnitude prediction

Seed-averaged Instrument A ($n = 12$), reference $G = 480$, $\lambda_{\text{thr}} = 0.03$, $T_{\text{ly}} = 800$:

K_{drive}	$\Sigma_{\text{meas}} \pm \text{SE}$	Σ_{pred}	rel. error	bias/signal
0.6	0.033233 ± 0.000179	0.067641	103.5%	3%
0.9716 (K_c)	0.010835 ± 0.000044	0.012479	15.2%	8%
1.4	0.008926 ± 0.000038	0.008057	9.7%	10%
2.0	0.008622 ± 0.000043	0.008863	2.8%	10%
2.5	0.002507 ± 0.000022	0.002145	14.5%	35%
5.0	0.000665 ± 0.000010	0.000663	0.2%	138%

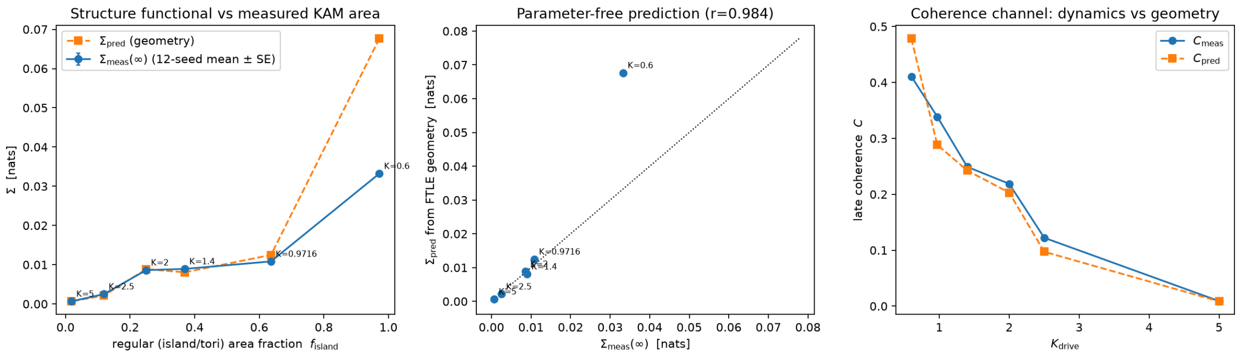

Left: seed-averaged Σ_meas(∞) (error bars within markers) and the zero-fitted-parameter Σ_pred vs the independently measured regular-area fraction. Center: Σ_pred vs Σ_meas, identity dotted; K=0.6 far above the line (sub-K_c failure). Right: coherence channel from dynamics vs geometry.

Reproduction. Instrument A reproduces the parent 100.002 pilot within $\leq 3 \times 10^{-4}$ nats; an independent re-execution regenerated all metrics bit-for-bit (§12).

The magnitude result. Over the estimator-clean window $K_c \leq K \leq 2.0$ (bias/signal $\leq 10\%$), the zero-fitted-parameter Σ_{pred} predicts the seed-averaged Σ_{meas} to **median 9.7%** (12.1% incl. the marginal $K = 2.5$). A fitted one-parameter proxy αf_{island} ($\alpha = 0.029$) achieves only 18.9% (27.8%). The geometry beats a fitted one-parameter model ~2× with one fewer degree of freedom. SE $\leq 0.5\%$ of signal in this window, so the ~10% residual is a real systematic of the uniform-component approximation, not statistical noise.

Not evidence. At $n = 6$ with all points monotone in K, Pearson r is affine-invariant (arbitrary decreasing $f(K)$ reach $r \approx 0.89$–0.97); the $r = 0.985$ of Σ_{pred} (in `kam_scores.json`) is context only, and the Spearman rank is matched by raw entropy.

4. Convergence and robustness

Grid. At $K = 1.4$: $(f_{\text{island}}, \Sigma_{\text{pred}}) = (0.3694, 0.0079)/(0.3691, 0.0081)/(0.3693, 0.0080)$ for $G = 320/480/640$. Converged.

FTLE integration time. At $K = 1.4$, $G = 320$: $\Sigma_{\text{pred}} = 0.007941/0.007911/0.008052$ for $T_{\text{ly}} = 400/800/1600$ — stable to $\sim 1.5\%$ over $4\times$.

Threshold. Σ_{pred} recomputed at $\lambda_{\text{thr}} \in \{0.01, 0.02, 0.03, 0.05, 0.08\}$ keeps median error in 6–12%. No knife-edge.

Seed region. Varying the flood-fill seed from $(0.25, 0.1)$ to $(2.0, 1.0)$ to x-wide $(3.0, 0.2)$ changes the median error by $< 1\%$.

5. Statistical geometry null: arrangement vs area

To separate *arrangement* from *area*, we compare the real Σ_{pred} against an **area-preserving shuffle null**: the same *number* of excluded cells placed uniformly at random, destroying spatial arrangement while preserving area. Because only column totals enter the x-marginal, uniform fixed-count placement induces a multivariate hypergeometric distribution over columns (each with capacity G). We draw **1000 independent shuffles per K** ($G = 320$) and report mean, standard deviation, 95% interval, and the corrected empirical one-sided tail probability $(b+1)/(1000+1)$, where b is the number of shuffles reaching the real value:

K	real Σ_{pred}	shuffle mean ± sd	shuffle 95%	p_{emp}	z
0.9716 (K_c)	0.01364	0.003773 ± 0.000294	$[0.003219, 0.004355]$	$1/1001$	33.6
1.4	0.00791	0.000916 ± 0.000076	$[0.000778, 0.001068]$	$1/1001$	92.3
2.0	0.00893	0.000523 ± 0.000042	$[0.000444, 0.000607]$	$1/1001$	198.4

The full-torus null (no exclusion) gives $\Sigma_{\text{pred}} = 0$ trivially.

Reading. Across $K = K_c, 1.4, 2.0$, none of the 1000 shuffled masks reaches the real prediction. With the finite-sample correction, $p_{\text{emp}} = 1/1001$ at each point; the real values sit 33.6σ, 92.3σ, and 198.4σ above their null distributions. The **spatial arrangement of the excluded regular region — its concentration into coherent blocks that shift the phase marginal off uniform — carries the predicted signal, not merely the excluded area.** This is the precise, scoped content of "the shadow of the islands" throughout the estimator-clean window. The earlier K_c area-effect reading resulted from an incorrect repeated-null sampler and is superseded by this independently regenerated fixed-area test.

6. Kept negatives, at full volume

Below K_c. At $K = 0.6$ the magnitude prediction fails by $2\times$ (0.068 vs 0.033), not a finite-time transient: extending T to 1000 and 2000 moves Σ_{meas} only to $0.0308/0.0304$. Below K_c the accessible chaotic layer is thin and sticky (cantori, slow transport), so the **uniform-fill assumption** of (4) is wrong. (This is a distinct failure from the K_c arrangement-null result of §5: here the *magnitude* is mispredicted because the ensemble does not fill uniformly; there the magnitude is right and the fixed-area null attributes the signal to arrangement, not area alone.)

Above $K \approx 2$ (estimator floor). The histogram-bias correction is 35% of signal at $K = 2.5$ and exceeds it (138%) at $K = 5.0$; these points are correction-dominated and their low relative errors are not informative. **Usable window: $K_c \le K \lesssim 2$.**

Both are retained as `results-negative-kept`.

7. What the coarse sweep establishes and does not

Establishes. For $K_c \le K \lesssim 2$, the mixed-regime $\Sigma(\infty)$ of 100.002 is quantitatively anchored on directly measured invariant-set geometry: a zero-fitted-parameter predictor reproduces its seed-averaged magnitude to $\sim 10\%$, beating a fitted one-parameter proxy $2\times$; and a 1000-shuffle fixed-area null establishes that the *arrangement* of the excluded regular region, not its area alone, carries the predicted signal (33.6–198.4σ). This converts 100.002 §5.4's "consistent with" into a measured, scoped, statistically null-controlled correspondence.

Does not establish. (i) Fine geometric discrimination beyond entropy: the registered Gate 4 test in §9 fails. (ii) Island *area/arrangement* at fixed K, not KAM *survival*. (iii) The FTLE mask marks the **excluded regular region**, which need not be exclusively KAM islands; hence the descriptive title, "shadow of the islands" as physical interpretation for the regime where arrangement dominates. (iv) No quantum gate or cross-mechanism collapse.

8. Relation to the corpus

- **566.030 / 826.829:** the max-entropy anchor; this overlays a geometric predictor on the 566.030 figure.
- **100.002 v2.6:** its answer-appendix line; the dense negative narrows the parent claim.
- **100.003:** the accessible-component geometry is the classical basin structure the collapse picture rehearses; the sub-K_c failure is a non-ergodic basin-capture instance.
- **787.787:** the FTLE/flood-fill instrument is a coherence-geometry measurement with pre-declared nulls of the kind 787 requires.

9. Registered Gate 4 dense sweep: not passed

The next step registered in v1.4 was executed without selecting a friendlier window after seeing dynamics. A geometry-only locator evaluated $K = 0.9716$ and a 0.05 grid from 1.00 to 2.00. The frozen rule selected the maximum absolute second difference of Σ_{pred}, fixing the window $1.45 \le K \le 1.55$. Confirmation then used ten unseen midpoint conditions, $K = 1.455, 1.465, \ldots, 1.545$, with 12 new seeds per condition.

Frozen Gate 4 quantity	Result

$\rho(f_{\text{island}}, \Sigma_{\text{meas}})$	0.030
one-sided permutation p	0.477
$\rho(f_{\text{island}}, -S_{\text{meas}})$	0.964
condition-resampling stability interval for $\Delta\rho$	$[-1.765, -0.128]$
geometry median relative magnitude error	8.7%
frozen area-only median relative magnitude error	6.6%

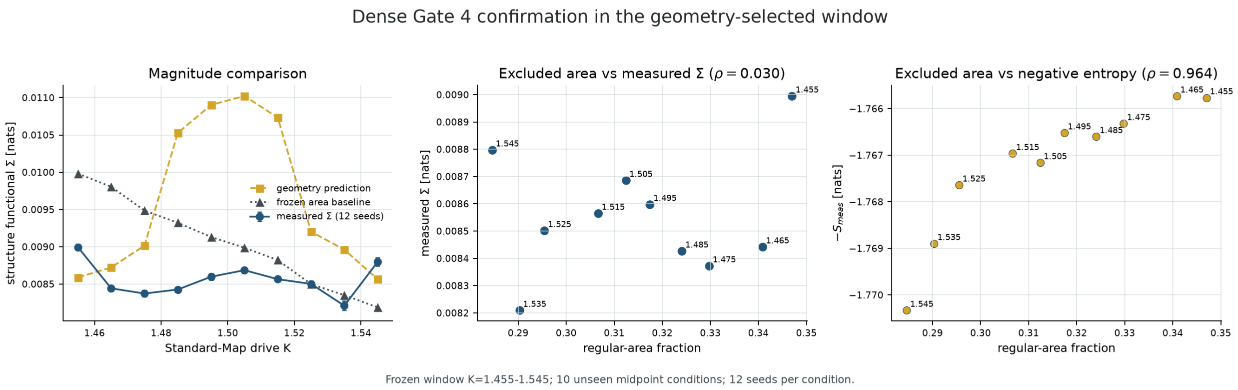

Dense Gate 4 confirmation. The geometry prediction peaks near K=1.50 while measured structure remains comparatively flat; negative entropy supplies the stronger monotonic rank baseline across the same K window.

All estimator-floor, seed-precision, fixed-area-null, and three-sentinel grid/time convergence checks pass. The failure is therefore not assigned to those controls. The stability interval resamples adjacent fixed K conditions and is not a population confidence interval; the entropy comparison is a monotonic rank baseline, not a causal attribution after removing K.

Decision. Gate 4 is not passed. The dense run does not support fine geometric ranking beyond entropy, and the full accessible-component predictor does not beat the frozen area-only magnitude baseline in the selected window. The original coarse magnitude correspondence remains reported, but this registered route to pillar status is closed negatively. Searching another window would constitute a new registered claim, not a continuation of this confirmation.

10. Kill conditions and promotion status

The coarse-window claim fails if FTLE is not converged in G or T_{ly}; the zero-fitted-parameter prediction for $K_c \leq K \lesssim 2$ exceeds $\sim 30\%$ under reasonable λ_{thr}; the geometry fails to beat the fitted one-parameter proxy; the repeated fixed-area shuffle null is not rejected; the sub-K_c failure is only transient; or Σ_{meas} fails to reproduce the parent pilot. Those coarse checks pass.

Pillar promotion additionally required the dense Gate 4 ranking and baseline conditions. Those conditions fail in §9. The paper therefore retains the scoped coarse result while recording non-promotion at the stronger gate.

11. Reproducibility appendix

Fixed seeds throughout.

File	Role
`FRC-100-002-001-v1.5.md`	this paper
`kam_sigma_probe.py`	both instruments + sweep driver
`seeds_and_nulls.py`	multi-seed uncertainty, T_{ly} convergence, single and repeated geometry nulls
`score_and_figure.py`	scores, fitted-proxy baseline, context correlations, figure
`test_score_consistency.py`	regression test binding the score artifact to the archived 12-seed headline
`robustness.py`	Σ_{pred}-vs-threshold sweep
`kam_metrics.json`, `seeds_nulls.json`, `kam_scores.json`	outputs
`fig_kam_sigma.png`	figure
`gate4/`	frozen protocol, dense data, controls, scores, report, figure, and review notes
`CHECKSUMS.txt`	SHA-256 manifest

Rerun: `python3 kam_sigma_probe.py`; `python3 seeds_and_nulls.py seeds <K...>` / `nulls <K...>` / `tly <K>` / `repnull <K...>`; `python3 score_and_figure.py`; `python3 test_score_consistency.py`; `python3 robustness.py`.

Method note. Computation executed and verified in-session by AI assistants under the author's direction, in the corpus's reproduce-from-seeds discipline.

12. Adversarial-review chain (v1.0 → v1.5)

Three independent reviews within one session.

Review 1 (independent agent, re-ran code). Reproduced all metrics bit-for-bit; confirmed instrument independence; showed the v1.0 headline Pearson r is affine-invariant/non-evidential at $n = 6$. → v1.1 promoted the zero-fitted-parameter magnitude over correlation.

Review 2 (independent model). Seven corrections, all adopted → v1.2: full-precision tables; **12-seed uncertainty** (honest median 9.7–12.1%, not the single-seed 8.5%); "zero-fitted-parameter" wording with choices listed; explicit T_{ly} convergence; **single geometry nulls**; cylinder/torus clarification; descriptive title.

Review 3 (independent model). Three corrections, all adopted → v1.3:

1. **Repeated fixed-area shuffle null** — 1000 realizations per K with mean, sd, 95% interval, and corrected empirical tail probability. Independent regeneration corrected an erroneous repeated-null table in the supplied draft: arrangement is rejected against random fixed-area placement at all three clean-window points (zero exceedances; $p_{\mathrm{emp}} = 1/1001$; $z = 33.6, 92.3, 198.4$).
2. **Separated the two assumptions** the nulls test: the shuffle probes *arrangement given area*; the sub-K_c failure probes *uniform fill*. §2.2, §5, §6 now keep them distinct.
3. **Wrapped toroidal seed region** stated explicitly and verified in code (already correctly implemented as $d_{\mathbb{T}} <$ threshold).

Plus wording: "the excluded regular region" rather than "the islands," since the FTLE mask does not uniquely isolate KAM islands.

No review required retraction; each correction narrowed and strengthened the result. The clearest example of the process working: the single-shuffle "arrangement at K_c" claim did not survive a proper statistical null, and the paper now says so.

Review 4 and registered confirmation → v1.5. An independent audit found the v1.3 release score had consumed single-seed values; v1.4 repaired the artifact and bound it to the archived 12-seed means. The predeclared dense Gate 4 test then selected its window using geometry only and failed on unseen dynamical conditions. Independent recomputation confirmed the score and required narrower interpretation of the adjacent-condition stability interval, entropy baseline, fixed-area null, and local checksum chronology. All limits are carried into §9 and `gate4/REVIEW-NOTES.md`.

13. Version note

Version 1.5 retains the v1.4 artifact-consistency correction and adds the completed registered dense Gate 4 result. The new result is negative: fine ranking is absent, entropy supplies the stronger monotonic rank baseline, and the frozen area-only predictor has lower magnitude error in the selected window. No coarse point is removed and no replacement window is searched.

Connections

- [[FRC-100-002]] — parent basin statement
- [[FRC-566-030]] — exact von Mises / $Q(C)$ curve
- [[FRC-826-829]] — von Mises exchange-rate identity
- [[FRC-100-003]] — basin capture and collapse
- [[FRC-787-787]] — coherence-geometry measurement discipline

References

1. B. V. Chirikov, "A universal instability of many-dimensional oscillator systems," *Physics Reports* 52, 263–379 (1979).
2. J. D. Meiss, "Symplectic maps, variational principles, and transport," *Rev. Mod. Phys.* 64, 795 (1992).
3. G. Benettin, L. Galgani, A. Giorgilli, J.-M. Strelcyn, "Lyapunov characteristic exponents for smooth dynamical systems," *Meccanica* 15, 9 (1980).
4. G. A. Miller, "Note on the bias of information estimates," in *Information Theory in Psychology* (1955).
5. E. T. Jaynes, "Information theory and statistical mechanics," *Phys. Rev.* 106, 620 (1957).
6. K. V. Mardia, P. E. Jupp, *Directional Statistics*, Wiley (1999).
7. H. Servat, *FRC 100.002 v2.6 — Coherence in Chaos* (preprint, 2026-07-12).
8. H. Servat, *FRC 566.030 v1.3 — Reciprocity in Action* (preprint, 2026-07-09).
9. H. Servat, *FRC 826.829 v1.5 — Mathematical Foundations* (preprint, 2026-07-09).