A Proposed Temporal Difficulty Diagnostic for Hamiltonian Systems

Tamás Nagy, Ph.D.
ORCID: 0009-0004-8079-4679
tnagyphd@gmail.com
https://thalens.org

Overview

Regular and chaotic Hamiltonian motion can coexist across widely separated time scales. This note asks whether a calibrated, frequency-resolved summary of nonlinear coupling and finite-time redistribution can complement established chaos diagnostics.

The contribution is a proposed dimensionless diagnostic, its correct log-frequency derivative, and a falsifiable validation protocol. A small 36-trajectory experiment is included only as a reproducibility demonstration: one proxy has weak in-sample separation with a wide uncertainty interval, and there is no held-out test or independently established chaos label.

The construction is not a KAM theorem, an Arnold-diffusion theorem, a universal chaos criterion, or a validated classifier.

Abstract

We propose a finite-window temporal diagnostic for selected near-integrable Hamiltonian systems. It compares a normalized Fourier statistic of the perturbation Hessian along trajectories with a separately normalized finite-time redistribution estimator. We derive the log-frequency quotient derivative and state the positivity, calibration, coordinate, trajectory, and window dependencies that limit its interpretation. A deterministic 36-trajectory exploratory run compares four proxy statistics against an exploratory spectator-advantage label. After correcting the canonical modal projection and fixing the physical window, the best proxy has AUC 0.51029 (bootstrap 95% interval 0.30864–0.71193; one-sided permutation $p = 0.474$). It exceeds the input-coordinate baseline by 0.05350, but a paired score-swap test gives $p = 0.401$. The experiment therefore supplies no evidence of predictive value beyond that baseline. The proposed threshold must be fit and tested out of sample against an independent invariant-torus or chaos diagnostic.

1. Scope and Relation to Existing Diagnostics

Frequency-map analysis studies regularity and slow diffusion through numerical frequency estimates [3]. Alignment methods such as SALI distinguish ordered and chaotic trajectories through the evolution of deviation vectors [4]. The present proposal does not replace those methods. It asks whether a statistic that explicitly combines a coupling amplitude with a finite-time frequency or transport scale adds predictive information when evaluated against such an independent diagnostic.

Wavelet time-frequency analysis has already been used to extract instantaneous frequencies, identify resonance trapping and transitions, and distinguish weak from strong chaos in Hamiltonian systems [5, 6]. The present construction also developed from an earlier unpublished internal research note that proposed a temporal frequency hierarchy, a coupling-to-effective-damping quotient, a log-frequency beta function, and a candidate unit threshold [7]. The narrower contribution of this paper is the finite-window normalization, the explicit Hessian-based coupling and redistribution definitions, the derivative identity, the calibration caveats, and a falsifiable validation protocol. It does not claim to introduce time-frequency analysis itself or the underlying temporal (D, C, P) motif.

For a perturbative model with small parameter ϵ, selected regimes may display fast orbital frequencies and slower secular or resonant evolution. The scaling of those slower processes is model-dependent. No universal three-scale decomposition is assumed.

2. Proposed Temporal Diagnostic

The preceding motivation becomes operational only after the coordinate and scale choices have been fixed.

Fix dimensionless canonical coordinates $\tilde{z} = (\tilde{q}, \tilde{p})$ before calibration. Concretely, one may choose positive coordinate scales Q_i, P_i with a common action scale $Q_i P_i = A_*$, set $q_i = Q_i \tilde{q}_i$ and $p_i = P_i \tilde{p}_i$, and normalize the perturbation Hamiltonian by an energy scale E_*: $\tilde{H}_1 = H_1/E_*$. Time is measured in units A_*/E_*; below, t and ω denote the resulting dimensionless time and frequency. For a finite observation window T, choose positive dimensionless reference scales C_* and λ_* and define

$$\mathcal{D}_T(\omega) = \frac{\|C_T(\omega)\|_F/C_*}{\lambda_{\mathrm{eff},T}(\omega)/\lambda_*}, \qquad \lambda_{\mathrm{eff},T}(\omega) > 0.$$

Here $C_T(\omega)$ is a windowed temporal Fourier statistic of nonlinear coupling. The term $\lambda_{\mathrm{eff},T}$ is not physical damping. It is a proposed finite-time estimator of redistribution or decorrelation at frequency ω. Its definition, reference scales, and zero-handling rule are part of the calibration. Consequently, the numerical value 1 has no invariant meaning before calibration.

The empirical hypotheses are deliberately weak:

- calibrated values may contain information about regular versus irregular motion;
- a fitted threshold may complement established finite-time diagnostics;
- transfer failure across ensembles or models refutes universality.

None of these hypotheses identifies chaos with Arnold diffusion. Rigorous Arnold-diffusion results require near-integrability, convexity or related structure, sufficient degrees of freedom, genericity, and connecting-orbit hypotheses [1, 2].

3. Mathematical Definition

To make the proposal testable, we now specify its finite-window mathematical object before introducing the exploratory proxies.

Consider a chosen near-integrable representation in the fixed dimensionless canonical chart

$$\dot{\tilde{z}} = J\nabla_{\tilde{z}}\tilde{H}(\tilde{z}, t), \qquad \tilde{H}(\tilde{z}, t) = \tilde{H}_0(\tilde{z}) + \epsilon\tilde{H}_1(\tilde{z}, t), \qquad \tilde{z} \in \mathbb{R}^{2n}.$$

Along a specified trajectory $\tilde{z}(t)$, define

$$C_{T,jk}(\omega) = \frac{1}{T}\int_0^T \frac{\partial^2 \tilde{H}_1}{\partial \tilde{z}_j \partial \tilde{z}_k}(\tilde{z}(t), t) e^{-i\omega t}\, dt.$$

All entries of this Hessian are dimensionless, so its Frobenius norm is well-defined. The result remains trajectory-, chart-, scale-, and window-dependent: changing the canonical nondimensionalization changes the diagnostic and requires recalibration. No coordinate-invariant or ensemble-level interpretation is claimed.

Let

$$A(\omega) = \frac{\|C_T(\omega)\|_F}{C_*}, \qquad L(\omega) = \frac{\lambda_{\mathrm{eff},T}(\omega)}{\lambda_*}.$$

For $\omega > 0$, where A and L are differentiable, $C_T(\omega) \neq 0$, and $L(\omega) > 0$, the log-frequency derivative is

$$\beta_{\mathcal{D}_T}(\omega) = \frac{d\mathcal{D}_T}{d\log\omega} = \omega\left(\frac{A'(\omega)}{L(\omega)} - \frac{A(\omega)L'(\omega)}{L(\omega)^2}\right),$$

with

$$A'(\omega) = \frac{\mathrm{Re}\langle C_T(\omega), C_T'(\omega)\rangle_F}{C_*\|C_T(\omega)\|_F}.$$

Zeros of $\beta_{\mathcal{D}_T}$ on $\omega > 0$ are stationary frequencies of the calibrated diagnostic. Calling them ultraviolet or infrared fixed points is only an analogy; no stability or diffusion conclusion follows from a zero alone.

4. Exploratory 36-Trajectory Receipt

The attached machine-readable receipt is bound to the numerical source by SHA-256. In this experiment, $\rho = r_{12}$ is the initial separation coordinate of the first pair, while θ sets its initial orientation. All trajectories use 300 leapfrog steps of size $\Delta t = 5 \times 10^{-4}$, hence the same physical window $T = 0.15$. The receipt records 12 values of ρ, three angles, all 36 trajectory summaries, software versions, finite-difference scales, masses, initial-data parameters, four proxy statistics, and explicit limitations.

At the initial point, let M be the relevant diagonal mass matrix and diagonalize the generalized Hessian as

$$U^\top M^{-1/2} V'' M^{-1/2} U = \mathrm{diag}(\kappa_j).$$

These spectra are indefinite: at every one of the 36 initial conditions, the raw-coordinate Hessian has two negative eigenvalues and the Levi–Civita Hessian has one. The action calculation therefore uses only the positive spectral subspace. On that subspace, $\omega_j = \sqrt{\kappa_j}$, and the implementation applies the stricter numerical retention rule $\omega_j > 0.01$. Negative, zero, and sub-floor modes are excluded; they are not interpreted as real oscillatory frequencies. For each retained mode, the canonical coordinates and action are

$$Q = U^\top M^{1/2}(q - q_0), \qquad P = U^\top M^{-1/2} p, \qquad I_j(t) = \frac{P_j(t)^2 + \omega_j^2 Q_j(t)^2}{2\omega_j}.$$

The numerical experiment uses dimensionless canonical coordinates in normalized gravitational units. The same construction is applied in the raw and Levi–Civita coordinate systems. Its regularized energy parameter is the full initial Hamiltonian

$$h = H(q_0, p_0) = V(q_0) + \sum_i \frac{p_{0,i}^2}{2m_i},$$

not the initial potential alone. Define $\sigma_{\mathrm{raw,best}}$ as the minimum temporal standard deviation of I_j over positive-frequency raw modes, and define $\sigma_{\mathrm{spectator}}$ as the temporal standard deviation of the lowest positive-frequency Levi–Civita mode. The exploratory label is then

$$W = \log\left(\frac{\sigma_{\mathrm{raw,best}}}{\sigma_{\mathrm{spectator}}}\right),$$

and the reported positive class is $W > 0$. This is a spectator-advantage label, not an independently verified chaos label. The run produced 9 positive and 27 non-positive labels. For the static coupling proxy, the quartic tensor is transformed to the same canonical modal basis and

$$\sigma_{k^*} = \max_{j<k} \frac{|A_{jk}|}{\max(|\omega_j - \omega_k|, 10^{-6})}, \qquad A_{jk} = \frac{V^{(4)}_{jjkk}}{4\omega_j\omega_k}.$$

We write η for the temporal mean of the lowest-to-highest positive frequency ratio, and $\sigma_\omega = \mathrm{sd}(\omega_{\mathrm{spectator}})/\,\mathrm{mean}(\omega_{\mathrm{spectator}})$. The in-sample results were:

Statistic	AUC	Cohen's d
σ_{k^*}	0.38683	-0.37283
η	0.49383	-0.13692
σ_ω	0.45679	-0.44193
$D_1 = \sigma\eta$	0.41564	-0.32525
$D_2 = \sigma_\omega/g$	0.46091	-0.43382
$D_3 = \sigma\sigma_\omega$	0.46914	-0.39275
$D_4 = s/g^2$	0.51029	-0.15012
$-\rho$ baseline	0.45679	0.08487

Here $\sigma = \sigma_{k^*}$ and $g = \max(1-\eta, 0.01)$, while $\eta(t)$ is the ratio of the lowest to the highest positive Levi–Civita frequency at the subsampled times. The sweep-rate proxy is

$$s = \mathrm{sd}\left(\frac{\eta(t_{i+1}) - \eta(t_i)}{t_{i+1} - t_i}\right).$$

The implementation uses $\max(g^2, 10^{-4})$ in D_4. The static quartic proxy is not numerically converged at the frozen finite-difference scales: at $\rho = 0.40$ and the aligned angle, halving both differentiation steps changes σ_{k^*} from 1.97191 to 8.21449 (a relative change of 3.16576). Accordingly, results involving σ_{k^*} are exploratory only; the headline D_4 proxy does not use it. A stratified 10,000-draw bootstrap gives a 95% AUC interval of $[0.30864, 0.71193]$ for D_4; a 10,000-permutation one-sided test gives $p = 0.474$. The AUC difference over $-\rho$ is 0.05350 with a one-sided paired score-swap $p = 0.401$. These uncorrected in-sample results do not support a predictive claim. The definitions are numerical proxies; none is proved equal to $\mathcal{D}_T$.

The code assigns descriptive zones using hard-coded boundaries at $\rho = 0.25$, 0.45, and 0.65. Therefore $\rho = 0.45$ is not a discovered transition and is not evidence for a mixing boundary. No conclusion in this paper depends on that zone assignment.

The numerical source hash is 1e213a089bcc54acdcbbb18ae3eb2494a9a6784812226f816958a9ef552ab76e, and the receipt hash is db11fbba9d927a7b5907b5dd04e9827b9d9f4ec290662b4be46f1ab192eb7acd.

5. Verification Boundary

A supplementary manifest lists ten elementary implication checks concerning nonnegativity, monotonicity, and conditional quotient or AUC statements. It does not calculate or certify the reported AUC values. These checks do not establish Hamiltonian regularity, KAM-torus persistence, Arnold diffusion, universality, or the empirical performance of the proxy. This note therefore claims no formal verification of its scientific conclusions. The machine-readable receipt records the numerical-source hash, software versions, the frozen numerical parameters used by the public source, and every trajectory summary used in the table; it is a numerical reproducibility artifact only.

6. Validation Protocol

A meaningful test requires the following preregistered sequence:

1. specify the canonical nondimensionalization, trajectory ensemble, window, C_*, λ_*, and the estimator $\lambda_{\mathrm{eff},T}$;
2. freeze proxy choice and regularization constants before evaluation;
3. fit any threshold on a training ensemble only;
4. compare with independent frequency-map, SALI, Lyapunov, or invariant-torus diagnostics on held-out trajectories;
5. report uncertainty, permutation baselines, and sensitivity to window and coordinates;
6. test the frozen protocol on another Hamiltonian model.

The primary falsifier is lack of out-of-sample improvement over established diagnostics or lack of transfer across models.

7. Limitations and Nonclaims

1. $\lambda_{\mathrm{eff},T}$ is not yet operationally fixed.
2. The 36-trajectory run is small and in-sample; its wide uncertainty interval includes chance performance.
3. The exploratory label $W > 0$ is not a chaos ground truth.
4. The proxy depends on coordinates, trajectory, window, and regularization; the tested static quartic proxy is materially finite-difference sensitive.
5. No KAM theorem, Arnold-diffusion theorem, universal threshold, or coordinate-invariant classifier is claimed.
6. No formal theorem in the historical artifact proves the physical proposal.

8. Conclusion

The temporal-difficulty construction is a testable research proposal, not an established law. The corrected exploratory run does not show predictive value beyond the input-coordinate baseline. Its value will therefore be determined only by a frozen out-of-sample study against independent chaos diagnostics. The present numerical receipt makes this negative result reproducible while preventing it from being mistaken for validation.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used AI-based tools for manuscript drafting, literature search, symbolic computation verification, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

[1] Cheng, Chong-Qing (2012). Arnold diffusion in nearly integrable Hamiltonian systems. *arXiv*. DOI: 10.48550/arXiv.1207.4016

[2] Cheng, Chong-Qing and Xue, Jinxin (2015). Arnold diffusion in nearly integrable Hamiltonian systems of arbitrary degrees of freedom. *arXiv*. DOI: 10.48550/arXiv.1503.04153

[3] Laskar, Jacques (1999). Introduction to Frequency Map Analysis. *Hamiltonian Systems with Three or More Degrees of Freedom, NATO ASI Series 533*. DOI: 10.1007/978-94-011-4673-9_13

[4] Skokos, Ch. and Antonopoulos, Ch. and Bountis, T. C. and Vrahatis, M. N. (2002). Smaller alignment index (SALI): Determining the ordered or chaotic nature of orbits in conservative dynamical systems. *arXiv*. DOI: 10.48550/arXiv.nlin/0210053

[5] Chandre, C., Wiggins, S. R., and Uzer, T. (2003). Time-frequency analysis of chaotic systems. *Physica D: Nonlinear Phenomena* 181(3–4), 171–196. DOI: 10.1016/S0167-2789(03)00117-9

[6] Vela-Arevalo, Luz Vianey (2002). *Time-Frequency Analysis Based on Wavelets for Hamiltonian Systems*. Ph.D. dissertation, California Institute of Technology. DOI: 10.7907/8MBB-3Z60

[7] Nagy, Tamás (2026). *The Three-Body Problem Through the (D, C, P) Lens: Four New Tool Concepts*. Unpublished internal research note, 27 April 2026.