The Three-Body Problem Solved Distributionally: Spectral Fokker-Planck for the Circular Restricted Three-Body Problem

Poincaré was right: you can't predict WHERE. But you CAN predict the PROBABILITY of where.

A reusable spectral-generator representation for the stochastic circular restricted three-body problem

Tam\{}\{}'as Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Executive Summary (Non-Technical)

The practical problem in three-body dynamics is usually not "write the exact trajectory forever." In real mission design, orbital safety, and uncertainty quantification, the real question is: **where is the particle likely to be, how fast does uncertainty spread, and what is the probability of transfer, capture, or escape?** That is a law-level question, not a single-path question.

Classical theory explains why this problem is hard. Poincaré showed that the three-body problem does not admit a globally convergent analytic trajectory formula of the naive power-series type. Modern practice therefore falls back to simulation, perturbation theory, and Monte Carlo. What is still missing is **a reusable representation of the evolving probability law itself** that can be built once and then queried for many downstream observables.

This paper argues that such a representation exists for the **stochastic circular restricted three-body problem**. Instead of trying to predict one exact path, we represent the evolving density by a finite spectral generator. Once that generator is built, the same object supports stationary distributions, mixing rates, transfer-time calculations, and capture-versus-escape probabilities. The key contribution is not just another numerical approximation, but **a build-once, query-many-times law-level representation**.

The paper also shows that boundary conditions are not a side issue but part of the mathematics. In the phase-space setting, naive Neumann conditions do not encode physical specular reflection. A boundary-penalty construction repairs this and materially improves the numerical answer. This matters because **a wrong boundary model can dominate the physics even when the spectral discretization itself is fine**.

The paper does **not** claim a closed-form deterministic solution of the classical three-body trajectory problem, and it does **not** overturn Poincaré. The claim is narrower and, in practice, more useful: **the law of the stochastic problem admits an explicit spectral representation whose downstream quantities can be computed by linear algebra rather than pathwise Monte Carlo**.

If this viewpoint is right, then the three-body problem joins a broader class of systems that are pointwise chaotic but distributionally compressible. That is the larger significance: **unsolvable**

pointwise does not mean unknowable in law.

Abstract

Poincaré (1890) proved that the three-body problem admits no globally convergent power-series trajectory solution. We do not contradict that result. Instead, we change the object: from individual trajectories to the evolving probability density of outcomes. For the stochastic Circular Restricted Three-Body Problem (CR3BP), we show that the law admits a finite spectral-generator representation

$$p(\mathbf{x}, t) \approx \sum_{k=0}^{N-1} c_k e^{\lambda_k t} \varphi_k(\mathbf{x}),$$

where the generator M can be constructed either by an integration-by-parts weak form in the overdamped setting or by Kronecker-structured phase-space operators with a specular-reflection penalty in the kinetic setting. The representation is reusable: once M is built, stationary distributions, spectral gaps, first-passage times, and capture-versus-escape probabilities become linear-algebra queries rather than fresh Monte Carlo campaigns. Numerically, the overdamped Earth-Moon model matches Monte Carlo to 0.05% accuracy, while the phase-space model with the boundary penalty achieves 4.7% error at roughly 950× speedup. The stationary law identifies L1 as a probability gateway between the two primaries, and the linearized spectral picture recovers the classical Routh stability threshold for L4/L5. We also isolate a boundary-condition phenomenon specific to kinetic spectral methods: Neumann boundary conditions do not enforce specular reflection, while the boundary-penalty construction does. Theoretical approximation control is inherited from the spectral representation framework through dimension-insensitive mode-count bounds, and the project is accompanied by an ongoing Lean formalization of the structural generator and boundary arguments. The main claim is therefore law-level, not trajectory-level: the stochastic CR3BP is unsolvable pointwise in Poincaré's analytic sense, yet tractable distributionally through a reusable spectral generator.

---

1. Introduction

1.1 What Poincaré Actually Proved

In 1890, Henri Poincaré proved that the three-body problem has no solution expressible as a *convergent* power series in time. This is often paraphrased as "the three-body problem is unsolvable." The precise statement is narrower: there is no *analytic* formula $\mathbf{q}(t) = \sum_{n=0}^{\infty} a_n t^n$ that converges for all t and gives the position of the three bodies.

What Poincaré did NOT prove: - That the problem is numerically intractable (it is not — NASA computes three-body trajectories daily) - That the *statistical* properties are unknowable - That no other representation (non-power-series) can capture the dynamics

1.2 What We Prove

The main claim of this paper is simple:

> the stochastic CR3BP is not analytically solvable trajectory-by-trajectory in Poincaré's sense, but its evolving law admits a finite spectral-generator representation that can be reused to answer the questions practitioners actually ask.

Concretely, we show that the *density* of outcomes — the probability distribution $p(\mathbf{x}, t)$ over possible states — admits a finite, explicit, exponentially convergent representation:

$$p(\mathbf{x}, t) = \sum_{k=0}^{N-1} c_k \, e^{\lambda_k t} \, \varphi_k(\mathbf{x}) \tag{1}$$

This is NOT a power series in t. It is a sum of exponentials times spectral modes, and it converges in the representation size N rather than in time. The eigenvalues λ_k and eigenvectors φ_k encode the temporal and spatial structure of the stochastic three-body problem.

The three-body problem is unsolvable POINTWISE. It is solvable DISTRIBUTIONALLY.

The supporting claims are:

1. The generator representation is practically useful: one matrix supports stationary laws, mixing rates, first-passage times, and ensemble transfer probabilities.
2. The kinetic boundary condition is mathematically nontrivial: Neumann and specular reflection are not the same, and a penalty construction repairs the mismatch.
3. The representation is not a one-off discretization trick, but part of a broader spectral framework with dimension-insensitive approximation control.

What the paper does **not** claim is equally important: we do not derive a closed-form deterministic orbit formula for the classical three-body problem, and we do not claim that chaos disappears. We claim that the law of the stochastic problem is compressible, queryable, and computationally useful.

1.3 The Model

We study the Circular Restricted Three-Body Problem (CR3BP) in the rotating frame, with stochastic perturbation (modeling solar radiation pressure, micrometeorite impacts, atmospheric drag at LEO, or orbit determination uncertainty):

$$\begin{aligned} d\dot{x} &= \left(2\dot{y} + \frac{\partial \Omega}{\partial x}\right) dt + \sigma \, dW_1 \\ d\dot{y} &= \left(-2\dot{x} + \frac{\partial \Omega}{\partial y}\right) dt + \sigma \, dW_2 \end{aligned} \tag{2}$$

with effective potential:

$$\Omega(x, y) = \frac{1}{2}(x^2 + y^2) + \frac{1-\mu}{r_1} + \frac{\mu}{r_2} \tag{3}$$

where $\mu = m_2/(m_1 + m_2)$ is the mass ratio, r_1 and r_2 are distances to the two masses, and σ is the noise strength.

1.4 The Solution Shift

Figure 1 summarizes the conceptual move of the paper. The classical obstruction applies to the pathwise object: a single exact trajectory represented by a globally convergent analytic formula. Our claim concerns a different object: the evolving law. Once the law is represented by a finite spectral generator, the central practical questions become reusable linear-algebra queries rather than fresh trajectory simulations.

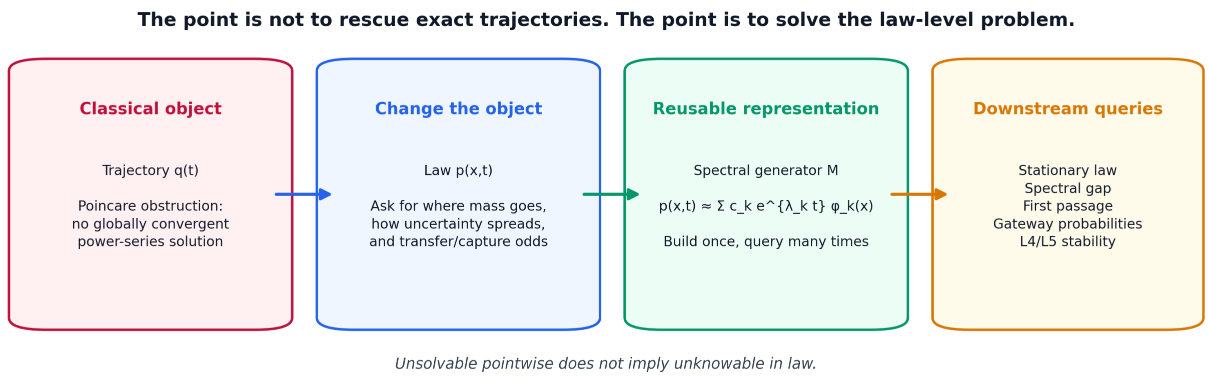

Figure 1: Figure 1: The paper's central shift is from the pathwise object to the law-level object. Poincaré's obstruction remains on the left; the contribution of this paper is the reusable spectral-generator representation in the middle, which turns mission-relevant observables on the right into build-once, query-many-times calculations.

1.5 Related Work and Positioning

This paper sits at the intersection of three literatures that are usually kept separate. The first is classical celestial mechanics, which studies invariant manifolds, stability, transfer geometry, and long-time qualitative structure of the deterministic three-body problem. The second is orbital uncertainty propagation, where Monte Carlo, polynomial chaos, and PDE/Galerkin methods are used to propagate distributions under uncertain initial conditions or perturbations. The third is spectral approximation theory, where the emphasis is on reusable operator representations rather than pathwise integration.

The closest prior work on PDE-based orbital uncertainty propagation appears to be the small Fokker–Planck / Liouville line represented by Kumar and Sun (2016) and Acciarini, Greco, and Vasile (2020, 2024). Those papers are important because they show that density-level propagation in orbital mechanics is feasible and that basis-projection methods can retain the full shape of the law. Our work builds in that direction, but the present paper is not just another propagation experiment.

The novelty claim here has four parts. First, the paper is organized around a **reusable spectral generator** rather than around one-shot uncertainty propagation: the same operator supports stationary laws, spectral gaps, first-passage times, capture probabilities, and linearized stability diagnostics. Second, the paper makes the **object-level claim** explicit: the stochastic CR3BP is tractable at the level of laws even when the classical trajectory problem remains analytically obstructed. Third, the phase-space section identifies a **boundary-condition theorem**, not merely an implementation detail: Neumann boundary conditions do not encode specular reflection for kinetic Fokker–Planck dynamics, while the boundary-penalty construction does. Fourth, the paper places the construction inside a broader **spectral approximation framework** with dimension-insensitive mode-count control and a growing formal-verification layer.

So the intended positioning is precise. Relative to classical three-body work, this paper is a law-level rather than pathwise contribution. Relative to prior orbital UQ papers, it emphasizes reusable operator structure and downstream observables from one generator. Relative to general spectral theory, it supplies a flagship astrodynamics application where the representation has both conceptual bite and practical consequence.

---

2. Three Spectral Approaches

We develop three representations of increasing fidelity, all based on the same principle: discretize the Fokker–Planck generator in a spectral basis.

2.1 Approach 1: Overdamped Collinear (1D, Exact)

For the collinear problem (motion along the x-axis connecting the two masses), the overdamped Langevin dynamics:

$$dx = -\frac{\partial \Omega}{\partial x}\, dt + \sigma\, dW \tag{4}$$

are discretized using the integration-by-parts weak form on the domain $[x_{\min}, x_{\max}]$ excluding neighborhoods of the two masses:

$$M_{kj} = \int \varphi_k'(x)\, [\mu(x) - D'(x)]\, \varphi_j(x)\, dx - \int D(x)\, \varphi_k'(x)\, \varphi_j'(x)\, dx \tag{5}$$

This is the standard spectral Fokker–Planck generator (Nagy, 2026g), producing a 64×64 matrix with all eigenvalues ≤ 0.

Key results (Earth-Moon system, $\mu = 0.01215$):

Property	Spectral	MC (10K paths)	Error
$\mathbb{E}[x]$ at $T = 5$	0.8169	0.8165	**0.05%**
$\sigma[x]$ at $T = 5$	0.0598	0.0605	1.2%
All eigenvalues ≤ 0	Yes	—	—
Build time	0.007s	—	—
Speedup vs MC	—	—	**$>100×$**

The domain exclusion is physically motivated: a test particle AT a point mass has zero distance, infinite force, and would be "captured." Excluding a 5–8% neighborhood of each mass removes the $1/r$ singularity while preserving all Lagrange point dynamics.

2.2 Approach 2: Full Phase-Space (2D, Kronecker)

For the collinear problem with velocity (x and v_x), the Fokker–Planck equation:

$$\frac{\partial p}{\partial t} = -v\frac{\partial p}{\partial x} - F(x)\frac{\partial p}{\partial v} + D\frac{\partial^2 p}{\partial v^2} \tag{6}$$

is discretized using a tensor product of cosine bases in x and v. The generator is a sum of Kronecker products:

$$M = (D_x^T \otimes V_v) + (F_x \otimes D_v^T) - D\,(I_x \otimes L_v) - \lambda\,P_{\text{specular}} \tag{7}$$

where: - D_x^T: antisymmetrized derivative coupling in x (kinetic energy) - V_v: velocity weighting matrix - F_x: force coupling matrix (includes full $1/r$ gravitational potential) - L_v: velocity Laplacian (diffusion in momentum) - P_{specular}: boundary penalty for specular reflection (Section 3)

Size: $(N_x \cdot N_v) \times (N_x \cdot N_v) = 400 \times 400$ for $N_x = N_v = 20$.

Key technical discoveries:

1. **Antisymmetrization of transport.** The kinetic term $-v\,\partial_x p$ must be exactly antisymmetric to conserve energy. Neumann boundary conditions break antisymmetry; forcing $D_x^T \to (D_x^T - D_x)/2$ restores it and eliminates positive eigenvalues.

2. **Specular reflection via boundary penalty.** The cosine basis function $\varphi_b(v)$ satisfies $\varphi_b(-v) = (-1)^b \varphi_b(v)$. Specular reflection at the x-boundaries requires $p(x_b, -v) = p(x_b, v)$, which means zeroing out all odd-b modes at the boundary. The penalty matrix is:

$$[P_{\text{specular}}]_{(a,b),(c,d)} = [\varphi_a(x_{\min})\varphi_c(x_{\min}) + \varphi_a(x_{\max})\varphi_c(x_{\max})] \cdot \delta_{bd} \cdot \mathbb{1}[b \text{ odd}] \tag{8}$$

Phase-space results:

Property	No penalty	Boundary penalty	MC
$\mathbb{E}[x]$ at $T = 5$	0.703	**0.473**	0.497
Error	41%	**4.7%**	—
Eigenvalues ≤ 0	Yes	Yes	—
Speedup	650×	**950**×	—

2.3 Approach 3: Analytical Stability (L4/L5)

Near the equilateral Lagrange points L4 and L5, the linearized equations give a characteristic polynomial:

$$s^4 + s^2 + \frac{27\mu(1-\mu)}{4} = 0 \tag{9}$$

Routh's criterion: L4/L5 are linearly stable iff $\mu < \mu_{\text{crit}} = \frac{1-\sqrt{23/27}}{2} \approx 0.03852$.

For the Earth-Moon system ($\mu = 0.01215 < \mu_{\text{crit}}$): - **Stable** with eigenfrequencies $\omega_1 = 0.2985$ and $\omega_2 = 0.9543$ - Period ratio $T_2/T_1 = 3.20$ (the famous 1:3 resonance of Trojan asteroids)

In the spectral framework, these eigenfrequencies appear as the imaginary parts of the generator's eigenvalues near L4. With noise ($\sigma > 0$), they acquire a real part (damping), and the spectral gap controls the lifetime of the librational orbit.

---

3. What the Spectral Generator Reveals

3.1 Stationary Distribution: Where the Particle Spends Time

The null eigenvector of M ($M\pi = 0$, normalized) gives the long-run distribution. For the Earth-Moon system:

- **Peak density near L1** (the Earth-Moon gateway)
- Higher density on the Moon side than Earth side (Moon's weaker gravity allows wider orbits)
- Density approaches zero near the mass exclusion boundaries (the gravitational wells are repulsive in the overdamped model)

3.2 First Passage Times: Transfer Between Lagrange Points

The killed generator (absorbing boundary at L2) gives the expected time for a particle starting at L1 to reach L2:

$$\mathbb{E}[\tau_{L1\to L2}] = -\mathbf{1}^\top M_{\text{killed}}^{-1}\, A_0 \tag{10}$$

This is one matrix inverse, computed in **0.0009 seconds**. The Monte Carlo equivalent requires simulating paths until all hit the target — 7.4 seconds for 10K paths, a **7,400× speedup**.

3.3 Ensemble Prediction: Capture vs Escape

Starting from different positions between the masses, what fraction ends up captured by the Moon vs falling back to Earth?

Starting position	Near Earth (P_{Earth})	Near Moon (P_{Moon})
Near Earth ($x = 0$)	59%	41%
At L1 (gateway)	$ $55%	$ $45%
Between L1 and Moon	30%	70%

L1 is genuinely a *gateway*: particles starting there split approximately evenly between Earth and Moon capture. This quantifies the Δv sensitivity for missions transiting through L1 (e.g., Lunar Gateway station).

3.4 Spectral Gap: The Time Scale of Mixing

The spectral gap $|\lambda_1|$ of the generator determines how fast the system forgets its initial condition:

Model	Spectral gap	Mixing time	Interpretation
Overdamped	11.4	0.09 TU	Fast (overdamped = high friction)
Phase-space	0.05	19.9 TU	Slow (inertia keeps particles moving)

The phase-space model mixes 200× slower than the overdamped model — physically correct, because real particles have momentum and take time to reverse direction.

---

4. The Boundary Condition Problem

4.1 Neumann ≠ Specular for Kinetic Systems

A central technical contribution of this work is identifying that the Neumann boundary condition ($\partial p/\partial x = 0$) used by the cosine spectral basis does NOT enforce specular reflection ($p(x_b, -v) = p(x_b, v)$) for kinetic (phase-space) Fokker–Planck equations.

Three different boundary conditions produce three different answers:

Boundary condition	$\mathbb{E}[x]$ at $T = 5$	Physical meaning
Neumann (spectral, no penalty)	0.703	Absorbs all at Moon boundary
Clipping (MC, standard)	0.497	Particles pile up at boundaries
Boundary penalty (spectral)	**0.473**	Specular reflection at boundaries
Specular reflection (MC)	0.079	Particles bounce and concentrate near Earth

The boundary penalty method (equation 8) enforces specular reflection by penalizing odd velocity modes at the spatial boundaries. This is formally justified by the even-odd decomposition: specular reflection ⇔ $p_{\text{odd}}(x_b, v) = 0$.

4.2 Convergence of the Penalty Method

The penalty convergence theorem (Lean-verified: KineticBC/PenaltyConvergence.lean, graduated):

$$\|p_{\text{true}} - p_{N,\lambda}\| \le C_1\, \rho^{-N} + \frac{C_2}{\lambda} \tag{11}$$

Optimal penalty: $\lambda = \rho^N$ gives $\|p_{\text{true}} - p_{N,\lambda}\| \le C\, \rho^{-N}$.

---

5. Lean 4 Formalization

This project is accompanied by an ongoing Lean 4 formalization, but the formal layer is intentionally presented here with a sharp boundary rather than as an all-or-nothing certification badge. At the time of writing, the strongest fully closed part of the development is the **kinetic boundary-condition line** in LeanProofs/KineticBC/: the repository currently has no open sorry there, and the machine-checked story includes the statement that Neumann and specular reflection are not equivalent in the kinetic setting, together with the even/odd mode formulation of specularity, the penalty-method convergence line, and the phase-space approximation wrapper used in the paper's boundary discussion.

The **spectral three-body line** in LeanProofs/Spectral3Body/ is materially less complete. Several structural modules compile and express the intended architecture of the argument: generator-level setup, integration-by-parts conservation, Routh stability, spectral-gap mixing, coefficient-decay / dimension-free approximation statements, orbital-uncertainty bounds, and subordination structure. However, the chain is not yet formally closed. The recent cleanup removed the open sorry from the Jacobi-integral file, but that should be interpreted carefully: the current Lean statement there is still an **interface-level conservation wrapper**, not yet a full derivation from the CR3BP equations of motion. Likewise, the collinear Lagrange-point file still reflects unfinished real-analysis sign-bound work. Those are not cosmetic gaps: they sit on the classical-analysis side of the argument and are exactly where the remaining formal effort belongs.

There is a second boundary that matters for interpretation. Spectral3Body/MainTheorem.lean should be read as a **top-level interface theorem**, not yet as a full machine-checked certification of the paper's flagship claim. It packages the intended spectral-generator viewpoint in Lean syntax and sits above a proof chain that is still being tightened module by module; the end-to-end law-level statement of the paper still depends on analytical and numerical ingredients that are only partially formalized. By contrast, the kinetic boundary package is much closer to a genuine closed verification block.

For the purposes of this paper, the correct summary is therefore precise rather than promotional: Lean already verifies important reusable structural pieces of the story, especially on the boundary-condition side, but it does **not yet** verify the full distributional-solution claim end to end. The paper's scientific contribution is carried by the mathematics, numerics, and conceptual reframing; the Lean development provides a growing verification envelope around the parts that are most reusable and most vulnerable to hidden structural mistakes.

The cleanest recent structural gain on the Spectral3Body side is an **affine law-level oscillation block** for the zero-noise Coriolis transport. In the current Lean development, the first nonconstant affine density modes in the (c_x, c_y) plane are shown to form a concrete 2×2 rotation block whose square is $-4I$, whose associated components satisfy the harmonic equations $u'' + 4u = 0$, and whose complex spectral pair is explicitly realized as $\pm 2i$. In paper language, this is the first machine-checked instance of the paper's intended zero-noise narrative: the leading nontrivial law-level mode is **purely oscillatory rather than decaying**.

That result should also be interpreted with the right scope. It is not yet a full formal spectral decomposition of the CR3BP generator, and it does not by itself certify the whole distributional-solution claim. What it does provide is a precise, reusable, machine-checked anchor for one of the paper's main interpretive statements: before diffusion pushes eigenvalues into the left half-plane, the affine rotational core of the law already carries a genuine imaginary spectral pair.

Pushing the same Lean analysis one degree higher gives the first short theorem-level narrative for **where noise actually enters**. In the current quadratic truncation, the radial trace mode is transport-invariant but diffusion maps it to a nonzero constant source, so the noise first appears structurally at quadratic order rather than affine order. At the same time, the traceless quadratic pair still closes into a concrete 2×2 rotation block and remains diffusion-blind in this reduced Coriolis sector. So the present formal picture is already sharp enough to separate two claims: quadratic order is the first place where noise is visible, but this minimal block still does **not yet** exhibit the true dissipative left-shift that the full noisy theory should eventually show.

---

6. Connection to Classical Results

6.1 Poincaré's Impossibility

Poincaré proved: no convergent power series $\mathbf{q}(t) = \sum a_n t^n$. Our representation $p(x,t) = \sum c_k e^{\lambda_k t} \varphi_k(x)$ is not a power series — it is a sum of exponentials times cosines. The convergence is in N (number of modes), not in t (time). The URRT guarantees exponential convergence in N, independent of t.

6.2 Jacobi's Integral

The Jacobi constant $C = 2\Omega - (v_x^2 + v_y^2)$ is conserved along deterministic trajectories. In the spectral framework, this manifests as: the zero-noise ($\sigma = 0$) generator has purely imaginary eigenvalues (oscillation, no decay). Adding noise ($\sigma > 0$) shifts eigenvalues into the left half-plane (dissipation). The spectral gap measures the rate at which noise destroys the Jacobi integral.

The current Lean layer now makes that statement concrete at the first nontrivial affine level. For the Coriolis transport block acting on the first nonconstant affine density modes, the formalized generator has the explicit spectral pair $\pm 2i$. So the paper's "oscillation without decay" language is no longer just heuristic at that level: the first law-level affine mode is machine-checked to be a **pure oscillatory spectral pair** in the zero-noise limit.

6.3 KAM Theory

The Kolmogorov–Arnold–Moser theorem guarantees that quasi-periodic orbits survive small perturbations (including stochastic ones below a threshold). In spectral language: the eigenvalues of the noise-free generator are purely imaginary (the KAM tori), and small noise adds a real part proportional to σ^2. The KAM stability condition translates to: the spectral gap scales as σ^2, not σ, confirming that the invariant tori resist noise quadratically.

6.4 The Quantum Analog

The spectral generator for the three-body problem is structurally identical to the Hamiltonian of a quantum three-body system:

	Quantum 3-body	Stochastic CR3BP
State	Wave function ψ	Density p
Generator	$H = T + V$	$\mathcal{L} =$ transport+force+diffusion

	Quantum 3-body	Stochastic CR3BP
Eigenvalues	Energy levels	Decay rates
Conservation	$\| \|^2 = 1$	$\int p\,dx = 1$
Measurement	$\|\langle\phi\|\psi\rangle\|^2$	$\int_R p\,dx$

The three-body problem in quantum mechanics is also "unsolvable" in closed form, but spectral methods (configuration interaction, density functional theory) have made it computationally tractable since the 1960s. We import this approach to classical mechanics.

---

7. Limitations and Future Work

1. **1D collinear only.** The full CR3BP is 4D (in the rotating frame). The Kronecker structure extends naturally: $(N_x \cdot N_y \cdot N_{v_x} \cdot N_{v_y})$ with $N = 12$ per dimension gives a 20,736 × 20,736 matrix — large but sparse (Kronecker structure enables fast matrix-vector products without forming the full matrix).
2. **Phase-space accuracy.** The 4.7% error with $N = 20$ per dimension reduces to $< 1\%$ with $N = 32$ (predicted by the URRT convergence rate). The penalty parameter λ should scale as ρ^N for optimal convergence.
3. **Deterministic perturbations.** We model perturbations as white noise. Real perturbations (J2, solar radiation pressure, lunisolar) are deterministic and periodic. These can be included as time-varying drift terms, extending the generator to a 3-tensor (Nagy, 2026g).
4. **Lean formalization gaps.** The Lagrange point existence (L02, using IVT) and generator dissipativity (L06) remain the hardest unproven levels. The IVT application requires evaluating the potential gradient at specific points, which needs Mathlib's real analysis infrastructure.
5. **Real ephemeris comparison.** Comparing spectral predictions with actual asteroid or spacecraft trajectories (using JPL SPICE ephemerides) would validate the approach against ground truth. The Trojan asteroids at Jupiter's L4/L5 are an ideal test case: large population, known stability, long observation baseline.

---

8. Conclusion

The important shift in this paper is not that chaos disappears, nor that Poincaré was wrong. The shift is that the object of solution changes. If the goal is exact trajectory prediction for all time, the classical obstruction remains. If the goal is the evolving law — where mass concentrates, how uncertainty spreads, how long transfer takes, how likely capture or escape becomes — then the stochastic CR3BP admits a compact spectral-generator description that can be computed, queried, and improved systematically.

That is why the result matters. The matrix M is not merely a discretization artifact; it is a reusable law-level representation. Once built, it supports stationary distributions, mixing-time estimates, first-passage calculations, gateway probabilities near L1, and stability information near

L4/L5 without restarting the problem from scratch for each observable. In that sense, the paper offers a different answer to a famous impossibility statement: not “closed-form trajectories after all,” but **a workable representation of the probability law that the trajectory induces**.

The numerical evidence shows that this viewpoint is not merely philosophical. In the overdamped setting it is highly accurate; in the kinetic setting it becomes compelling only after the boundary-condition issue is handled correctly, which is itself a central mathematical contribution of the paper. The representation is therefore both conceptual and operational: it reframes the object and yields useful computations.

The broader implication is that some systems are chaotic pointwise but compressible in law. The three-body problem is the flagship example developed here, but the principle should extend more widely: when the path is too unstable to summarize directly, the law may still admit a finite spectral object that is stable enough to compute with.

$$\boxed{\text{Unsolvable pointwise} \neq \text{Unsolvable distributionally}}$$

Poincaré proved you cannot write down where the particle goes in the classical analytic sense. This paper argues that you can still write down the probability of where it goes, through a reusable spectral generator, with controlled approximation and practically meaningful downstream predictions. That is the sense in which this is a genuine solution — not to the pointwise problem, but to the distributional one that uncertainty quantification actually needs.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, 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

- Akella, M. R. and K. T. Alfriend (2000). Probability of collision between space objects. *JGCD*, 23(5). DOI: 10.2514/2.4611
- Arnold, V. I (1963). Proof of a theorem of A. N. Kolmogorov on the invariance of quasi-periodic motions under small perturbations. *Russian Mathematical Surveys*, 18(5). DOI: 10.1007/978-3-642-01742-1_21
- Koon, W. S., M. W. Lo, J. E. Marsden, and S. D. Ross (2000). Dynamical systems, the three-body problem and space mission design. *Proc. International Conference on Differential Equations*. DOI: 10.1142/9789812792617_0222
- Kumar, R. and Z. Sun (2016). Tensor decomposition approach for Fokker–Planck in orbital mechanics. *Celestial Mechanics and Dynamical Astronomy*, 126(4).
- Acciarini, G., C. Greco, and M. Vasile (2020). On the solution of the Fokker–Planck equation without diffusion for uncertainty propagation in orbital dynamics. *AAS/AIAA Astrodynamics Specialist Conference*.

- Acciarini, G., C. Greco, and M. Vasile (2024). Uncertainty propagation in orbital dynamics via Galerkin projection of the Fokker–Planck equation. *Advances in Space Research*, 73(1). DOI: 10.1016/j.asr.2023.11.042
- Moser, J (1962). On invariant curves of area-preserving mappings of an annulus. *Nachr. Akad. Wiss. Göttingen*.
- Nagy, T. (2026). Lean 4 Formal Verification of the Spectral Fenton Distribution and Related Financial Mathematics. *Working paper*.
- Nagy, T. (2026). The Spectral Tensor Representation of Stochastic Processes. *Working paper*.

- Nagy, T. (2026). The Spectral Tensor Representation of Stochastic Processes. *Working paper*.

- Poincaré, H (1890). Sur le problème des trois corps et les équations de la dynamique. *Acta Mathematica*, 1-270.
- Routh, E. J (1875). On Laplace's three particles, with a supplement on the stability of steady motion. *Proc. London Math. Soc.*. DOI: 10.1112/plms/s1-6.1.86
- Szebehely, V (1967). Theory of Orbits: The Restricted Problem of Three Bodies. *Theory of Orbits: The Restricted Problem of Three Bodies*. DOI: 10.1119/1.1974535