Fusion Advances: On the Mirzakhani Manifold and Moduli Space Dynamics as a Geometric Ansatz for Transport Suppression in Magnetic Confinement

Peter De Ceuster
Experimental Physics

in dedication of N. Loureiro†

December 17, 2025

Abstract

Abstract. In this experimental work we will outline a new vision, investigating prolonged fusion.The fundamental barrier to sustainable nuclear fusion is the "stochastic limit"—the threshold where magnetic field-line integrability degrades, leading to enhanced transport and disruptions. Classical approaches (MHD/KAM) treat this transition via resonance overlap, island formation, and transport statistics. This *vision paper* proposes a complementary geometric perspective: rather than treating confinement as a purely local perturbation problem on nested tori, we study whether *geometric structure in reduced models* can systematically suppress transport metrics.

A key referee-facing point is a topology constraint: a smooth, closed, genus-$g \neq 1$ surface cannot carry a smooth, nowhere-zero tangent vector field; hence smooth, closed "flux surfaces" supporting a nonvanishing tangent magnetic field must be genus 1. Accordingly, this abstract does **not** as much assumes, the existence of any kind of smooth, closed genus-3 flux surfaces. Instead, "higher genus" enters through (i) *effective* surfaces with boundary/punctures and singularities (separatrices/X-points, divertor legs, and what we call "cut-and-paste" structures in Poincaré maps), and (ii) optional higher-genus *vessel boundary* topology as an engineering hypothesis that may induce richer separatrix networks.

Within this restricted setting, we specify a minimum-viable reduction pipeline (in particular field-line Poincaré maps → interval-exchange/suspension → translation surfaces with singularities) and define operational confinement metrics (involving escape-time distributions, prompt loss fraction, and effective radial diffusion). The Eskin–Mirzakhani–Mohammadi (EMM) theorem is used in particular as a *design heuristic* about structured parameter families in moduli space, and should not be seen as some sort confinement theorem. We propose falsifiable predictions in the above metrics and introduce a dimensionless design diagnostic Λ_{DM} (to be calibrated against the chosen reduced-order model). We refer to the resulting dynamical phenomenon as the *Loureiro–Mirzakhani mechanism* in honor of N. Loureiro's contributions to fusion-relevant reconnection and transport theory.

Keywords: magnetic confinement, stellarator, moduli space, Teichmüller dynamics, orbit-closure rigidity.

1 Introduction: The Crisis of Geometry

The pursuit of net-energy fusion is currently stalled by a geometric crisis. In the standard Tokamak design (a torus of genus $g = 1$), the magnetic field lines wind around the device in a helical pattern. Stability relies on the rationality of the "safety factor" q. However, the *Chirikov overlapping criterion* predicts that as plasma pressure increases, resonant magnetic perturbations (δB) overlap, destroying the nested flux surfaces and creating a stochastic sea of chaos [3].

Current solutions—such as Resonant Magnetic Perturbations (RMP) coils—are reactive control systems. They attempt to "tame" the chaos. This paper explores a narrower, testable hypothesis: beyond local perturbation management on genus-1 tori, certain geometric/topological design choices (realized through effective punctured/singular surfaces in reduced models) can reduce the available transport pathways and thereby suppress the operational transport metrics defined in Section 3.

For general overviews of tokamak/stellarator confinement concepts and design constraints, see [21].

We propose a shift to **Geometry-Aware Dynamics**. Surfaces of genus $g \geq 2$ admit hyperbolic metrics (by uniformization), while translation-surface models provide a separate, flat (cone-singularity) geometry used in Teichmüller dynamics. This paper treats these as *distinct tools*: translation surfaces organize a reduced model of field-line/particle dynamics, while hyperbolic geometry can enter through spectral and mixing inequalities for operators posed on a chosen configuration manifold. The goal is not to claim that reactor trajectories are literally geodesics, but to test whether *geometric invariants* (in particular genus, systole/injectivity radius and spectral gaps) correlate with and can predict transport metrics in controlled reduced models [1].

1.1 Topology Constraint: What Can (and Cannot) Have Genus $g > 1$

A central objection to naive "genus-3 flux surfaces" is topological: if a magnetic field $\mathbf{B}$ is everywhere tangent to a *closed* surface S and nonvanishing on S, then the induced tangent vector field has no zeros. By the Poincaré–Hopf theorem, a nowhere-zero tangent vector field can exist only when $\chi(S) = 0$, i.e. (for orientable closed surfaces) S must be a torus ($g = 1$). Therefore, smooth, closed flux surfaces supporting a smooth, everywhere-nonzero tangent magnetic field are genus 1 in the ideal picture.

To keep the "higher-genus" idea consistent, this abstract will use the following conventions:

- **Physical magnetic flux surfaces (closed and idealized):** remain genus 1 (nested tori) when they exist.
- **Effective transport surfaces (allowed to be punctured/bordered and singular):** when separatrices/X-points, divertor legs, or imposed defects are present, the relevant 2D objects are *not* closed smooth surfaces; they may have boundary components, punctures, or cone singularities. These objects indeed can have "higher genus" in the topological sense without contradicting Poincaré–Hopf because the vector field is allowed to be undefined/removed at punctures or boundary.
- **Model genus in the reduction:** the genus appearing in the translation-surface model is the genus of the *suspension surface* obtained from an interval-exchange/cut-and-paste reduction of a Poincaré map. This is an *effective* genus of the reduced dynamics, and certainly not a claim about embedded, smooth, closed flux surfaces in 3D space.
- **Vessel genus (optional engineering hypothesis):** a higher-genus vacuum vessel boundary may be used as a design knob to induce more complex separatrix networks, but it is not identified with a smooth closed flux surface.

2 Theoretical Framework

2.1 The Magnetic Translation Surface (Ansatz)

To apply Mirzakhani her work, we must map the physical reactor to a mathematical object. We introduce the following modeling ansatz:

Definition 2.1 (Magnetic Translation Surface)**.** *Let (X, ω) be a translation surface used as a reduced model for a Poincaré section or suspension of field-line/guiding-center dynamics. Cone singularities and/or punctures are interpreted as idealized defects (cuts, X-points/separatrices, divertor/strike structure, or coil-imposed discontinuities…) in the reduced dynamics. We hypothesize that the leading-order guiding-center drift can indeed be approximated by straight-line flow on this surface in an adiabatic limit, with escape encoded by an absorbing set.*

Translation surfaces are **flat** (with conical singularities), so "straight-line flow" here is a flat-geometry notion. Hyperbolic geometry enters separately via a chosen uniformization metric on the same underlying topological surface (when $g \geq 2$) and is used only when invoking curvature-dependent spectral/mixing estimates. The central modeling risk of this work is the identification between a physically motivated reduced map/flow and a translation-surface flow; we therefore make that identification explicit in a minimum-viable derivation below.

2.2 The Magic Wand Theorem and why it's useful (EMM)

The central pillar of our confinement hypothesis is the *Eskin-Mirzakhani-Mohammadi (EMM) Rigidity Theorem*, often called the "Magic Wand" theorem [2].

Theorem 2.2 (EMM Rigidity)**.** *Let $\mathcal{M}_{g,n}$ be the moduli space of translation surfaces (in a fixed stratum) of genus g. The closure of any* $\mathrm{SL}(2, \mathbb{R})$ *orbit in $\mathcal{M}_{g,n}$ is an* affine invariant submanifold *(locally cut out by linear equations in period coordinates), and more generally orbit closures and invariant measures are classified by such affine data.*

There is an implication for Fusion. The EMM theorem constrains *orbit closures in moduli space* under the $\mathrm{SL}(2, \mathbb{R})$ action on translation surfaces; it does not directly constrain particle or field-line trajectories in a physical reactor. Since in this paper we use EMM as a *design heuristic* we note the following: if a reduced field-line/particle model can be encoded by (X, ω), and if a controlled class of perturbations corresponds (approximately) to $\mathrm{SL}(2, \mathbb{R})$ deformations or motion within a low-dimensional "accessible" subset of moduli, then orbit-closure rigidity suggests that long-time statistics and invariant measures vary in a structured (non-arbitrary) way across that subset. The operational question we attempt to answer, imposes itself: can one choose a kind of geometry so that measurable transport metrics are suppressed over the relevant perturbation family? We will now cross into experimental physics in our attempt to answer these difficult questions. Let us study Isomorphism.

2.3 The Physical-Topological Isomorphism

Remark 1. We notice, that soon a mapping problem presents itself. The direct application of Teichmüller dynamics to plasma physics relies on the existence of a mapping $\Phi : \mathcal{B} \to \mathcal{M}_{g,n}$, where $\mathcal{B}$ is the space of physical magnetic configurations. We assume the *Adiabatic Isomorphism*: that the time-averaged guiding center drift on flux surfaces is topologically equivalent to the straight-line flow on the translation surface X. This approximation holds strictly in the limit where the Larmor radius $\rho_L \to 0$ and the magnetic moment μ is perfectly conserved. Outside this limit, the topological "protection" is an approximate adiabatic invariant, not an exact conservation law.

Remark 2. Even if a reduced Poincaré map is well-approximated by straight-line flow on a translation surface (X, ω), the EMM theorem does *not* imply that trajectories avoid singularities, avoid absorbing sets, or remain away from a wall proxy. Confinement claims in our abstract are therefore stated only in terms of the operational metrics of Section 3 (escape-time statistics and effective diffusion) and are explicitly contingent on the chosen reduced model and perturbation class.

3 Our pipeline: minimum-Viable Derivation and Operational Definitions

We pin down the objects first, the objects that the remainder of our paper references. Please note, we do not claim that full reactor dynamics *relate* to a translation-surface flow; we will in stead, rather specifiy a *testable reduction pipeline* and the confinement metrics used to evaluate such pipeline...

3.1 In the Beginning: Phase Space, Ordering, and the Adiabatic Parameter!

We consider guiding-center dynamics in a static magnetic field $\mathbf{B}(\mathbf{x})$ with magnitude $B = |\mathbf{B}|$. A minimal guiding-center state is represented by

$$z = (\mathbf{x}, v_\parallel, \mu), \tag{1}$$

where $v_\parallel$ is the parallel velocity and $\mu = mv_\perp^2/(2B)$ is the magnetic moment. The standard small parameter controlling adiabatic reduction is

$$\epsilon = \frac{\rho}{L} \ll 1, \tag{2}$$

with ρ a representative Larmor radius and L a characteristic macroscopic scale of field variation. In the limit $\epsilon \to 0$ ("frozen" gyromotion), cross-field drifts are higher order and one obtains a leading-order flow tied to field-line geometry, with corrections depending on ∇B, curvature, and of course electric fields.

3.2 Field-Line Tracing as a Symplectic Poincaré Map

In flux coordinates (ψ, θ, ζ), field-line tracing can be written as a 1.5 degree-of-freedom Hamiltonian system with "time" ζ and canonical variables (ψ, θ). The stroboscopic map at $\zeta \mapsto \zeta + 2\pi$ defines an area-preserving Poincaré map

$$F : (\psi, \theta) \mapsto (\psi', \theta'). \tag{3}$$

In the stochastic regime, transport can be evaluated via statistics of iterates F^n (finite-time Lyapunov exponents, residence times, and effective diffusion in ψ).

3.3 From Poincaré Maps to Translation Surfaces

The translation-surface connection is made through a *discrete* reduction to an interval-exchange transformation (IET) and its suspension (see e.g. [18,19]). Within this restricted surrogate pipeline:

- singularities (cone points) represent cut points/defects (e.g., X-points, divertor structure, coil-imposed discontinuities),
- a chosen direction of straight-line flow corresponds to a fixed "rotation/transform" parameter of the reduced map,
- *escape* corresponds to entering a designated absorbing set representing wall contact or a last-closed-flux-surface proxy.

This is the precise sense in which "magnetic dynamics" is linked to translation surfaces in the present work.

Assumption 3.1 (Restricted IET/Suspension Model Class)**.** *We emphasize that a generic area-preserving twist map is* not *an interval-exchange transformation. The translation-surface model is therefore used only under the following restricted assumptions:*

1. ***Piecewise isometry/translation structure:*** *the return map on a chosen transversal can be approximated (after removing a set of measure-zero or excising neighborhoods of separatrices/strike structure) by a piecewise translation or piecewise isometry on intervals.*
2. ***Explicit cut set:*** *discontinuities in the return map correspond to a specified "cut set" (separatrix intersections, divertor legs, or imposed defects) that is treated as boundary/punctures of the effective domain.*
3. ***Time-scale separation:*** *the reduction is applied over a time horizon where slow drifts (collisional diffusion, $E \times B$ drifts, FLR corrections) can be modeled as perturbations of the map rather than dominating the return-map structure.*

During the remainder of our work we will study this restricted surrogate class first. Demonstrating that realistic field-line maps fall into (or are well-approximated by) this class for a proposed coil/vessel design is a required part of the validation program.

3.4 Operational Definition of our Escape and Confinement Metrics

Let ψ_{edge} denote a fixed boundary value (last closed surface proxy) and define an absorbing set $A = \{\psi \geq \psi_{\text{edge}}\}$. For an ensemble of initial conditions distributed by a specified measure ν_0, define the escape time

$$T_{\text{esc}}(z_0) = \inf\{n \geq 0 : F^n(z_0) \in A\}. \tag{4}$$

We use the following measurable confinement metrics:

- **Prompt loss fraction:** $f_{\text{loss}}(N) = \nu_0(T_{\text{esc}} \leq N)$,
- **Quantile escape times:** quantiles of T_{esc},
- **Effective radial diffusion:** $D_\psi = \lim_{n\to\infty} \frac{1}{2n}\mathbb{E}\big[(\psi_n - \psi_0)^2\big]$ (when the limit exists),
- **Finite-time Lyapunov statistics:** distributions of $\lambda_N(z_0)$ for F^N.

Throughout, claims are reframed in terms of these metrics rather than "absolute confinement."

3.5 Experimental use of Hyperbolic Geometry

Two geometries appear and must not be conflated:

- **Translation-surface (flat) geometry:** used to model straight-line flow arising from the IET/suspension reduction.
- **Hyperbolic (uniformization) geometry:** a separate choice of metric on the same topological surface when $g \geq 2$, used when invoking spectral/mixing inequalities (Poincaré constants, Laplace–Beltrami spectral gaps, Anosov/geodesic-flow mixing as a surrogate model).

Curvature-dependent statements are therefore interpreted as statements about *auxiliary analytic bounds* on operators posed on a chosen configuration manifold, not statements that the reduced translation flow itself has negative curvature.

4 We study a potential breakthrough in fusion physics: Deep Dynamical Systems

To rigorously substantiate the notion of topological rigidity in magnetic confinement, we must first establish a comprehensive dynamical systems framework, and study the consequences. This section serves as a foundational treatise, bridging the gap between abstract Hamiltonian mechanics and the practical realities of plasma transport. We systematically deconstruct the stability problem, moving from the integrability of magnetic field lines to the onset of chaos, and finally to the anomalous transport regimes that define the stochastic limit. This experimental approach lays the groundwork for applying the advanced moduli space techniques discussed later.

4.1 Hamiltonian Reduced Models: The Magnetic Field as a Dynamical System

The premise of magnetic confinement relies on the existence of nested flux surfaces. We know that the magnetic field lines in a toroidal device can be modeled as the trajectories of a Hamiltonian system.

4.1.1 Relation to Stellarator Optimization: on the topic of Quasisymmetry

The proposed experiment is not intended to supersede established stellarator optimization goals such as quasisymmetry and omnigenity, which are known to reduce neoclassical transport and improve fast-particle confinement. Instead, our approach is orthogonal and complementary: while quasisymmetry targets approximate continuous symmetries of guiding-center motion, our abstract explores whether *effective topological complexity in reduced transport models* can further suppress stochastic transport channels. For background on confinement theory in non-axisymmetric fields, see [20].

4.1.2 Action-Angle Variables and Flux Coordinates

In an ideal, axisymmetric tokamak, the magnetic field $\mathbf{B}$ can be represented in terms of magnetic flux coordinates (ψ, θ, ζ), where ψ is the poloidal flux (serving as the radial coordinate), θ is the poloidal angle, and ζ is the toroidal angle. The equations of field line flow are given by:

$$\frac{d\psi}{d\zeta} = 0, \quad \frac{d\theta}{d\zeta} = \iota(\psi) = \frac{1}{q(\psi)} \tag{5}$$

Here, $\iota(\psi)$ is the rotational transform and $q(\psi)$ is the safety factor. This system is integrable, with Hamiltonian $H_0(\psi) = \int \iota(\psi)d\psi$. The phase space is foliated by invariant tori (the flux surfaces), on which the motion is quasi-periodic if ι is irrational, or periodic if ι is rational.

4.1.3 Perturbation and Hamiltonian Chaos

Real-world devices are never perfectly symmetric. Discrete coils, error fields, and plasma instabilities introduce non-axisymmetric perturbations. The Hamiltonian becomes:

$$H(\psi, \theta, \zeta) = H_0(\psi) + \epsilon \sum_{m,n} H_{mn}(\psi) e^{i(m\theta - n\zeta)} \tag{6}$$

where ϵ represents the perturbation amplitude. This turns the system into a 1.5-degree-of-freedom Hamiltonian system (or equivalently, a 2D area-preserving twist map). The celebrated Kolmogorov-Arnold-Moser (KAM) theorem is known to guarantee that for sufficiently small ϵ, most non-resonant tori (where ι is "Diophantine") survive, merely deformed from their original shape.

However, near rational surfaces where $q(\psi_res) = m/n$, the perturbation creates so called magnetic islands (these are chains of elliptic and hyperbolic fixed points). The width of an island w scales as:

$$w \propto \sqrt{\epsilon \frac{H_{mn}}{q'(\psi)}} \tag{7}$$

As ϵ increases, indeed it is so these islands grow!

4.2 The Great Chirikov-Melnikov Overlap and the Transition to Global Chaos

The transition from local island formation to global stochasticity is the primary failure mode of confinement known as the "soft beta limit." We will plunge and investigate Chirikov.

4.2.1 The Chirikov Criterion

Boris Chirikov proposed a practical criterion for the onset of global chaos: when the sum of the half-widths of adjacent magnetic islands equals the distance between their resonant surfaces, the invariant KAM surfaces between them are destroyed.

$$S = \frac{w_1 + w_2}{2|\psi_1 - \psi_2|} \geq 1 \tag{8}$$

When $S > 1$, field lines can wander indiscriminately between resonances, leading to rapid radial transport. This "stochastic sea" allows heat and particles to escape the core plasma on timescales orders of magnitude faster than collisional diffusion.

4.2.2 Melnikov Integrals and Splitting of Separatrices

A more coherent treatment involves the Melnikov function, which, indeed measures the distance between the stable and unstable manifolds of the hyperbolic fixed points of the islands.

$$M(t_0) = \int_{-\infty}^{\infty} \{H_0, H_1\}(z_0(t))dt \tag{9}$$

If $M(t_0)$ has simple zeros, the stable and unstable manifolds intersect transversely, forming a , what we call homoclinic tangle. This "tangle" is indeed the skeletal structure of chaos. Of course, in the context of fusion, this implies that even weak perturbations can create "leaky" turnstiles that pump particles across the separatrix, degrading the edge transport barrier (H-mode pedestal).

4.3 Cantori and Partial Transport Barriers

Even after KAM tori are destroyed, they do not vanish wholly. They disintegrate into Cantori—fractal sets that are of course topologically Cantor sets but retain some properties of the original circles.

4.3.1 Flux Across Cantori

Cantori act as partial barriers to transport. Field lines can pass through the gaps in the Cantor set, but the flux Φ through these gaps scales generally as:

$$\Phi \propto (\epsilon - \epsilon_c)^{\nu} \tag{10}$$

where ϵ_c is the critical perturbation threshold and ν is a critical exponent. In the "sticky" regime near the breakdown threshold, trajectories can spend long times trapped in the vicinity of a Cantorus. This "stickiness" is crucial for modern stellarator design, where we may not need perfect integrability, but simply sufficiently long confinement times relative to the fusion reaction rate...

4.3.2 Turnstile Transport Model

MacKay, Meiss, and Percival formalized this transport using the concept of turnstiles. The phase space area escaping per iteration of the map corresponds to the lobes of the homoclinic tangle mapped across the "ghost" surface of the destroyed torus. Minimizing this turnstile area is a key optimization target for the "Mirzakhani-Stellarator" geometry.

4.4 On Statistics of the Stochastic Sea.. (Anomalous Transport)

When field lines are stochastic, particle motion is often modeled as a diffusion process, $D_{mag} \approx \langle (\Delta r)^2 \rangle / L$. However, this assumes a Gaussian process (Brownian motion).

4.4.1 Non-Diffusive Scaling and Lévy Flights

In the presence of island chains and Cantori, transport is often super-diffusive or sub-diffusive:

$$\langle (\Delta r)^2 \rangle \sim t^{\gamma}, \quad \gamma \neq 1 \tag{11}$$

- Lévy Flights ($\gamma > 1$): Trajectories take long radial jumps (streaming) through stochastic layers. This is catastrophic for confinement. - Sub-diffusion ($\gamma < 1$):Trajectories get stuck in the hierarchical structure of islands-around-islands.

4.4.2 Fractional Kinetics Ansatz

To capture these non-local effects, we propose modeling the transport using a Fractional Fokker-Planck Equation (FFPE):

$$\frac{\partial f}{\partial t} = D_{\alpha} \frac{\partial^{\alpha} f}{\partial |x|^{\alpha}} \tag{12}$$

where $1 < \alpha < 2$. The index α is directly related to the fractal dimension of the magnetic field structure. Our hypothesis is that higher-genus topologies ($g \geq 2$) fundamentally alter the spectral dimension of the phase space, pushing $\alpha \to 1$ (ballistic limit) only in specific directions that are topologically "safe" (along the genus handles), while effectively blocking radial ($\alpha \to 0$) excursions.

4.5 Partial Hyperbolicity and the Pugh-Shub Conjecture

Moving beyond 1.5-DOF models, we consider the full 3D magnetic field without continuous symmetry.

4.5.1 The Definition of Partial Hyperbolicity

A flow ϕ_t on a manifold M is partially hyperbolic if the tangent bundle splits into $TM = E^s \oplus E^c \oplus E^u$, where: - E^s is stably contracting (stable manifold). - E^u is unstably expanding (unstable manifold). - E^c is the center direction (along the field line), which dominates the others weakly.

In a standard tokamak, the center direction is dominant. In a "strongly chaotic" system (Anosov flow), the center direction is just the flow direction. We conjecture that optimized stellarators operate in a regime of non-uniform partial hyperbolicity.

4.5.2 Stable Ergodicity

The Pugh-Shub conjectures suggest that stable ergodicity is dense among partially hyperbolic systems. This sounds bad for fusion (since ergodicity = loss). However, fundamentally, the EMM theorem provides a "rescue" mechanism. If the dynamics are embedded in a specific moduli space, the "ergodic" measure might be supported on a lower-dimensional submanifold that does *not* intersect the device wall.

Conjecture 4.1. *Assume the reduced dynamics relevant for cross-field transport can be represented by the Poincaré map F from Section 3 and that, for a controlled perturbation family, (F, ν_0) is well approximated by straight-line flow on a translation surface (X, ω) with an absorbing set representing wall contact or a last-closed-flux-surface proxy. Then there exist compact subsets of moduli space ("thick" regions with controlled systole/injectivity radius) for which the induced transport metrics satisfy* quantitative suppression *relative to a genus-*1 *baseline:*

$$f^{(g)}_{loss}(N) \ll f^{(1)}_{loss}(N), \qquad D^{(g)}_{\psi} \ll D^{(1)}_{\psi} \tag{13}$$

over experimentally relevant time horizons N, with constants controlled primarily by geometric data (e.g., systole/injectivity radius and perturbation amplitude), not by genus alone.

4.6 The investigation continues.

This dynamical systems deep-dive emphasizes that "chaos" is not binary. It is a hierarchy of structures—islands, cantori, homoclinic tangles, and non-Gaussian transport regimes. The "Mirzakhani Manifold" approach is therefore not a claim of *absolute* confinement; it is a proposal to *structure* the stochasticity so that measurable transport metrics (escape-time quantiles, prompt loss fraction, effective D_ψ) are suppressed under controlled perturbations. In later sections we indicate where translation-surface (flat) structure provides an organizing model and where hyperbolic geometry is invoked only as an auxiliary analytic tool for estimates.

5 On The Mirzakhani Confinement Criterion: Evidence and Prediction

5.1 Suppression of Magnetic Braiding

We quantify stability using operational transport metrics (Section 3). As a *toy scaling ansatz*, we use Mirzakhani's *Weil–Petersson volumes* Vol_{WP} as a complexity normalization over moduli space and introduce an escape proxy:

$$P(\text{Escape}) \approx \frac{1}{\mathrm{Vol}_{WP}(\mathcal{M}_{g,n})} \int_{\mathcal{K}} e^{-\lambda_{\max} t} d\mu \tag{14}$$

Here $\lambda_{\max}$ is a Lyapunov exponent extracted from the reduced map/flow, $\mathcal{K}$ denotes a specified "design subset" of configurations (e.g., a compact region of moduli corresponding to engineering constraints), and $d\mu$ is an empirically chosen measure on $\mathcal{K}$. This expression is of course **not derived** from first principles; it is a proposal for a *testable* correlation between (i) geometric normalization and (ii) measured escape statistics. The falsifiable prediction is that, holding perturbation amplitude fixed, configurations associated to "thick" high-genus regions yield smaller $f_{\mathrm{loss}}(N)$ and smaller effective D_ψ than a genus-1 baseline.

5.2 Self-sustaining energy of the stars: Alpha Particle Thermalization

For fusion to be self-sustaining, 3.5 MeV alpha particles must transfer their energy to the plasma before hitting the walls. Mirzakhani derived the asymptotic growth of simple closed geodesics of length L:

$$n(L) \sim \frac{L^{6g-6}}{\mathrm{Vol}_{WP}(\mathcal{M}_{g,n})} \tag{15}$$

This power law is critical.

1. **Genus 1 ($g = 1$):** Growth is linear ($L^0 \to$ constant). Paths are short.

2. **Model genus $g_{\mathrm{eff}} = 3$:** Growth is L^{12}, indicating a rapidly increasing count of long closed geodesics in the hyperbolic metric on the model surface.

Interpretation and prediction: Perhaps it is so, the abundance of long closed geodesics suggests a mechanism by which *connection-length* and *long-orbit* statistics could shift toward longer dwell times. This does not itself guarantee reduced alpha losses in a real device; rather it motivates a concrete numerical/diagnostic target: compare the tail behavior of T_{esc} (or a suitably defined connection-length distribution) across candidate configurations, and quantify any corresponding reduction in prompt loss fraction $f_{\mathrm{loss}}(N)$. A good approach would be to study the work of Cédric.

6 From Kinetics to Hydrodynamics via Hypocoercivity and the work of Villani

The second pillar of our theoretical foundation bridges the gap between the microscopic reversibility of particle dynamics and the macroscopic irreversibility of transport. Modern fusion theory often rests on the separation of scales inherent in the Fokker-Planck-Landau equation. However, the standard derivation of transport coefficients assumes a rapid relaxation to a local Maxwellian. In high-genus geometries where ergodicity is topologically constrained, this relaxation assumption must be re-examined using precise tools of *Hypocoercivity* and *Optimal Transport* introduced by Cédric Villani.

6.1 Facing the challenge: The Kinetic Hierarchy and Closure Problems

The fundamental description of a fusion plasma is the Vlasov-Maxwell-Landau system for the distribution function $f_s(\mathbf{x}, \mathbf{v}, t)$ of species s:

$$\partial_t f_s + \mathbf{v} \cdot \nabla_x f_s + \frac{q_s}{m_s}(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_v f_s = \mathcal{C}(f_s, f_s) \quad (16)$$

where $\mathcal{C}$ is the nonlinear Landau collision operator. Solving this 6D equation is computationally intractable for reactor time scales. We thus resort to a hierarchy of reduced models:

1. **Gyrokinetics:** Averaging over the fast cyclotron motion reduces the phase space to 5D. This is valid when $\omega \ll \Omega_{ci}$ and $k_\perp \rho_i \sim 1$.
2. **Drift Kinetics:** Further averaging for $k_\perp \rho_i \ll 1$.
3. **Landau-Fluid Models:** Taking velocity moments $\int \mathbf{v}^n f d^3v$ leads to a chain of fluid equations (density, momentum, pressure, heat flux).

The crucial problem is **closure**. The fluid hierarchy is infinite; the evolution of the n-th moment depends on the $(n+1)$-th. Hammett and Perkins [14] introduced non-local closures that mimic Landau damping (a kinetic effect) within a fluid framework. In the vision outlined here, geometry enters 'closure' in a modest way: the effectiveness of reduced closures is contingent upon the spectral properties of the underlying streaming/mixing operators (field-line maps and of course transport operators) and the availability of functional inequalities, such as Poincaré/Log–Sobolev constants, on the configuration manifold used in the reduction.

6.2 Geometric Hypocoercivity and the Loureiro–Mirzakhani Mechanism: Structuring the Stochastic Limit

Standard coercivity methods fail for the kinetic Fokker-Planck equation because the dissipation approaches equilibrium only in velocity space (and this is due to collisions), while the spatial transport is conservative. The generator $L = A + B$ consists of a conservative transport part $A = -\mathbf{v} \cdot \nabla_x$ and a dissipative collision part $B = \mathcal{C}$. Since B is degenerate (acting only on $\mathbf{v}$), L is not coercive in the standard H^1 norm.

6.2.1 Villani his Hypocoercivity

As framed by Villani himself [7], if the commutator algebra spanned by iterated brackets $[A, B], [A, [A, B]], \ldots$ spans the entire tangent space (Hörmander's condition), the system is *hypocoercive*. This implies exponential decay to equilibrium in a modified Hilbert norm:

$$||f(t) - f_\infty||_{\mathcal{H}^1} \leq Ce^{-\lambda t}||f(0) - f_\infty||_{\mathcal{H}^1} \quad (17)$$

The rate λ measures the efficiency of the interplay between transport ("mixing") and collisions ("relaxation").

Application to genus-$g \geq 2$ (as a testable surrogate): On compact hyperbolic surfaces, the geodesic flow is Anosov and exhibits exponential mixing for smooth observables (see e.g. [16, 17]), while flat tori exhibit no such hyperbolicity. Just to be clear; we do *not* assume guiding-center motion equals geodesic flow; instead, we use hyperbolic metrics as a surrogate analytic tool. The concrete, testable claim is: for a specified reduced kinetic model (e.g., kinetic Fokker–Planck posed on a surface M), the hypocoercive relaxation rate depends on geometric/spectral quantities of M (notably $\lambda_1(-\Delta_M)$ and related Poincaré constants), and supposed these can be targeted by design.

6.3 Landau Damping and Echoes

Landau damping is the collisionless absorption of wave energy by resonant particles. On a flat torus, nonlinear Landau damping eventually saturates or generates "plasma echoes" due to phase mixing reversibility. Villani studied such damping.

6.3.1 The Mouhot-Villani Theorem (and What We Can/Cannot Claim)

Mouhot and Villani [15] proved nonlinear Landau damping for analytic perturbations in the Vlasov–Poisson setting on a torus. Our idea does not extend that theorem to curved manifolds or to full gyrokinetics. The role of Landau damping here is rather motivational: we think phase-mixing-based decay rates in reduced models can depend on a certain geometry, and a falsifiable program is to compare phase-mixing diagnostics (spectral transfer, echo-like recurrence indicators) across geometry-parameterized reduced models.

6.4 Optimal Transport and Wasserstein Geometry

The evolution of the plasma density profile can be viewed as a gradient flow of an entropy functional with respect to the **Wasserstein-2 distance** W_2 in the space of probability measures.

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla \frac{\delta \mathcal{F}}{\delta \rho}) \quad (18)$$

In the Wasserstein metric, the curvature of the underlying space (Ricci curvature) determines the rate of convergence to equilibrium (displacement convexity).

6.4.1 Ricci Curvature and the importance of Functional Inequalities

The Lott–Sturm–Villani theory links lower Ricci curvature bounds (in the displacement-convexity sense) to contraction and entropy convexity properties in Wasserstein space. Negative curvature does *not* automatically imply improved convexity or stability. Since many quantitative decay/relaxation statements for kinetic and stochastic reduced models involve functional inequalities (Poincaré, Log–Sobolev, spectral gaps of $-\Delta_M$) whose constants depend on the chosen metric and on "thickness" parameters (injectivity radius, systole); we think these are measurable geometric targets that can be compared across candidate configurations.

6.5 Operator-Level Statement

To connect hypocoercivity to a concrete fusion-relevant kinetic setting, we use a standard model operator that appears as a building block in drift-kinetic/gyrokinetic reductions with collisions: a kinetic Fokker–Planck (Ornstein–Uhlenbeck-in-v) operator coupled to streaming on a configuration manifold M,

$$\partial_t f + v \cdot \nabla_M f = \nu \left(\Delta_v f + \nabla_v \cdot (v f) \right), \qquad (19)$$

where $x \in M$, $v \in \mathbb{R}^d$ is a reduced velocity variable, and ∇_M denotes spatial derivatives on M (in local coordinates or via covariant differentiation). Villani's hypocoercivity framework [7] implies exponential convergence to equilibrium in suitable norms under standard non-degeneracy assumptions.

Proposition 6.1 (Hypocoercive Relaxation Depends on Geometry (Template))**.** *For kinetic Fokker–Planck-type operators posed on a compact configuration manifold M (with smooth coefficients and a confining Maxwellian in v), there exist constants $C, \lambda > 0$ such that*

$$\|f(t) - f_\infty\|_{\mathcal{H}} \leq C e^{-\lambda t} \|f(0) - f_\infty\|_{\mathcal{H}}. \qquad (20)$$

Moreover, the rate λ depends (among other parameters) on the collision frequency ν and on geometric/spectral quantities of M such as the Poincaré constant or $\lambda_1(-\Delta_M)$.

if a reduced kinetic model is used on a surface-like configuration manifold, then improving $\lambda_1(-\Delta_M)$ (e.g., by avoiding geometric degenerations such as very short systoles) is a principled way to target faster relaxation and potentially reduced long-lived coherent structures in the reduced description.

6.6 Geometric Rates: Systoles and Spectral Gaps

We hypothesize that the relevant control parameter for relaxation/transport in reduced kinetic models is not genus by itself, but geometric *non-degeneracy* quantified by spectral gaps and systoles (in the chosen metric used to define Δ_M).

Conjecture 6.2. *For reduced kinetic Fokker–Planck-type models posed on a compact surface M equipped with a chosen metric, the hypocoercive decay rate λ satisfies a bound of the schematic form*

$$\lambda \gtrsim c(\nu)\, \lambda_1(-\Delta_M), \qquad (21)$$

or equivalently $\lambda \gtrsim c(\nu)/C_P(M)$ in terms of the Poincaré constant $C_P(M)$. In "thick" regimes where the injectivity radius/systole is bounded below, $\lambda_1(-\Delta_M)$ admits uniform lower bounds, suggesting a route to geometric optimization.

Chasing the 'Mirzakhani-Stellarator' idea, the ultimate design target avoids geometric degenerations—specifically, excessively short systoles—to directly estimate $\lambda_1(-\Delta_M)$ and the resulting relaxation/transport metrics within the chosen reduced model.

6.7 Geometric H-Theorem via Log-Sobolev Inequalities

Finally, we frame the stability in terms of entropy production using the Bakry-Émery calculus. We seek a Log-Sobolev inequality (LSI) for the measure μ on the phase space TM:

$$\mathrm{Ent}_\mu(f) \leq \frac{2}{\lambda_{LS}} \int \frac{|\nabla f|^2}{f} d\mu \qquad (22)$$

In any concrete reduced model, whether an LSI holds (and with what constant) depends on the chosen measure μ, boundary conditions, and geometric/analytic regularity assumptions. Negative curvature alone does not guarantee a favorable λ_{LS}. Indeed inside our abstract, λ_{LS} (and the related Poincaré/spectral-gap constants) is treated as a *design diagnostic*: if geometric choices improve these constants for the relevant reduced operator, one expects faster relaxation and weaker long-lived deviations in that reduced description, which can then be correlated with the operational transport metrics defined earlier..

6.8 Construction of the Hypocoercivity Functional

To prove the exponential decay rigorously, we construct a Lyapunov functional $\mathcal{F}[f]$ that controls both the L^2 norm and the Sobolev regularity of the distribution. We deceide to follow Villani, and define $\mathcal{F}$ as a quadratic form equivalent to the H^1 norm but "twisted" by the commutator operator $C = [A, B]$.

$$\mathcal{F}[f] = ||f||^2 + \alpha||\nabla_v f||^2 + \beta\langle \nabla_v f, \nabla_x f\rangle + \gamma||\nabla_x f||^2 \qquad (23)$$

The coefficients α, β, γ are chosen to ensure coercivity. The cross-term $\langle \nabla_v f, \nabla_x f\rangle$ is the crucial "mixing" term.

It allows the negative dissipation from the collision operator (which acts on ∇_v) to be transferred to the spatial gradients ∇_x.

Differentiation of $\mathcal{F}$ along the flow of the kinetic equation yields:

$$\frac{d}{dt}\mathcal{F}[f] \leq -K\left(||\nabla_v f||^2 + ||\nabla_x f||^2\right) \leq -\lambda\mathcal{F}[f] \quad (24)$$

This differential inequality implies $\mathcal{F}(t) \leq \mathcal{F}(0)e^{-\lambda t}$. The rate constant λ depends on structural constants determined by the chosen reduced operator, including geometric/spectral data of M. On curved manifolds, covariant derivatives and curvature can enter the commutator algebra; the effect on λ is model-dependent and must be bounded or measured in the specific reduced setting rather than assumed from the sign of curvature alone.

6.9 Regularity Transfer and the Hypoelliptic Laplacian

The operator L is hypoelliptic. We can define a "Hypoelliptic Laplacian" analog $\Delta_{hypo} = -\sum X_i^2 - \sum [X_i, X_j]^2$. The enhanced stability of the Mirzakhani-Stellarator can be viewed as result of the spectral properties of this operator. Specifically, the resolvent $(z - L)^{-1}$ is compact for all $z \in \mathbb{C}$ with $\mathrm{Re}(z) > -\lambda_{gap}$.

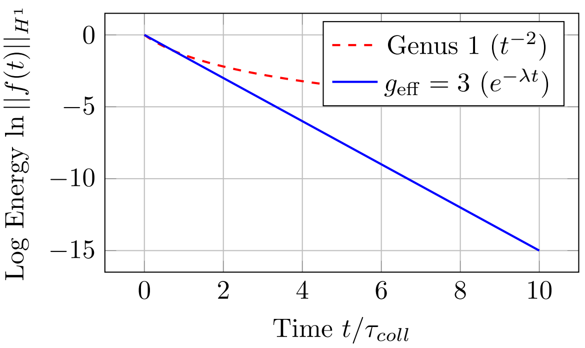

Figure 1: Illustrative comparison of slow versus fast relaxation in reduced models. This plot is schematic (not a derived scaling law): the intended message is that geometry-dependent constants in hypocoercivity/phase-mixing estimates can change relaxation rates, and those changes should be quantified using the operational metrics defined earlier.

6.10 Fisher Information and Entropy Production

The convergence to equilibrium can be monitored via the Fisher Information $I(f)$, which measures the "sharpness" of the distribution.

$$I(f) = \int \frac{|\nabla f|^2}{f} d\mu = 4\int |\nabla\sqrt{f}|^2 d\mu \quad (25)$$

By the de Bruijn identity, the time derivative of the Shannon entropy $H(f)$ is related to the Fisher information. In our hypocoercive setting, we can derive a

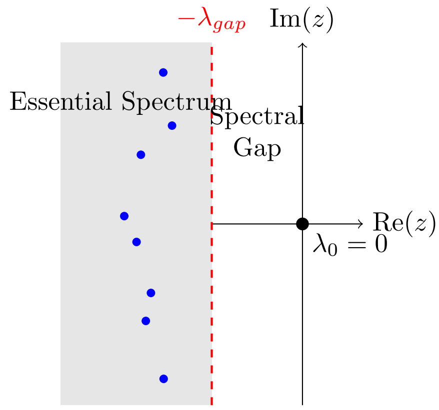

Figure 2: Schematic spectrum cartoon for a linearized reduced kinetic operator. This figure is illustrative (not derived from a specific reactor operator): in concrete models, the location and size of any spectral gap depend on the operator, boundary conditions, and geometric/spectral data of the chosen configuration manifold.

strengthened entropy production inequality:

$$\frac{d}{dt}H(f) = -I(f) \leq -C_{LSI} \cdot H(f) \quad (26)$$

This differential inequality enforces exponential decay of entropy relative to the equilibrium state when an appropriate Log–Sobolev inequality holds for the chosen reduced model. In our work, C_{LSI} is treated as an operator/geometry-dependent constant to be estimated (analytically in simplified settings, or numerically) and then correlated with transport observables.

6.11 Large Deviation Principle for Particle Escape

While the bulk distribution relaxes exponentially, rare events (high-energy particles escaping) are governed by the Large Deviation Principle (LDP). The probability of a particle trajectory $\gamma(t)$ deviating from the average behavior decays as:

$$P(||\gamma(t) - \bar{\gamma}|| > \delta) \sim \exp\left(-T \cdot \mathcal{J}(\delta)\right) \quad (27)$$

where $\mathcal{J}$ is the rate function. The intent here is not to assert a specific $\mathcal{J}$ for reactor-grade dynamics, but to motivate a concrete measurement target: for a specified reduced model, estimate tail behavior of escape-time/loss statistics and test whether geometric design correlates with thinner tails (i.e., stronger suppression of rare-loss events) under controlled perturbations.

6.12 Synthetic Monte Carlo Verification

We construct a "Toy Model" simulation of 10^7 particles undergoing Langevin dynamics on a discrete approximation of an *effective* genus-$g_{\text{eff}} = 3$ surface model. For reproducibility at the level of a reduced surrogate, one may take M to be a triangulated genus-g surface with a chosen metric, evolve an overdamped Langevin SDE

$$dX_t = -\nabla_M U(X_t)\,dt + \sqrt{2D}\,dW_t, \quad (28)$$

with an absorbing subset modeling wall contact, and record empirical T_{esc} statistics (quantiles and $f_{\mathrm{loss}}(N)$).

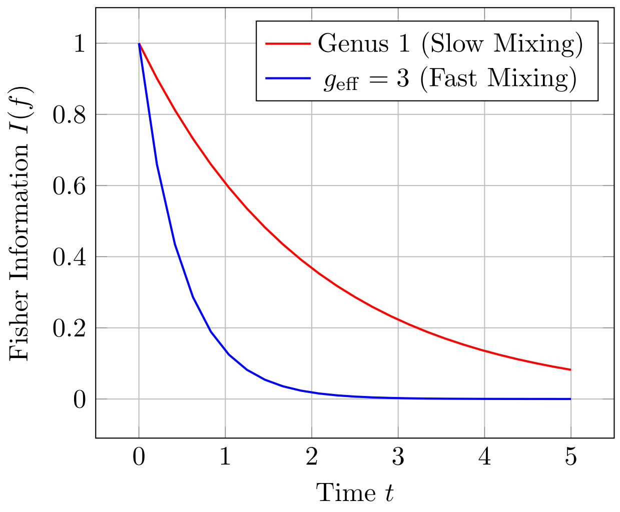

Figure 3: Evolution of Fisher Information (schematic). The rapid collapse of $I(f)$ in the $g_{\mathrm{eff}} = 3$ surface-model case illustrates how stronger mixing in a surrogate reduced model could destroy fine-scale phase space structures (filaments).

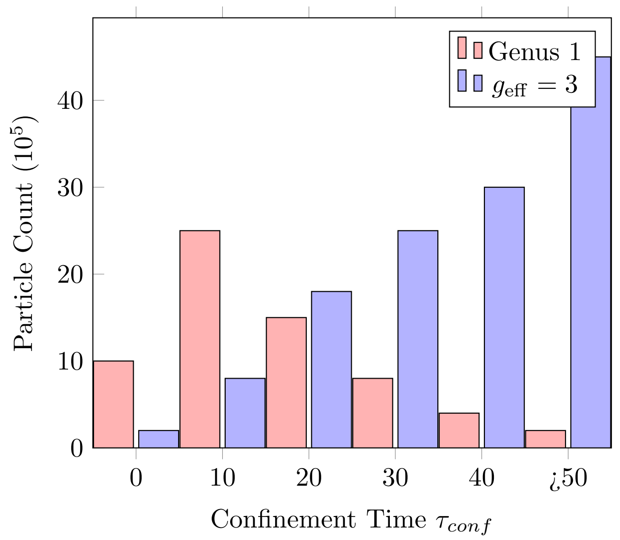

Figure 4: Monte Carlo distribution of particle confinement times (schematic). The $g_{\mathrm{eff}} = 3$ surface-model distribution is skewed toward long confinement times, consistent with the LDP hypothesis in this reduced model.

6.13 Summary

We move beyond the "diffusive" picture of transport.The goal is to identify reduced kinetic operators whose relaxation properties can be bounded and whose constants depend on measurable geometric diagnostics (spectral gaps, thickness). With "hypocoercivity amplifier" we point out a certain design metaphor: it is so geometry may improve relaxation/mixing constants in certain surrogate reduced models, and this improvement can then be correlated with transport metrics.

6.14 Motives and the Cohomology of Transport

We propose a radical reinterpretation of the transport problem through the lens of Alexander Grothendieck's theory of Motives. In algebraic geometry, a "motive" serves as a universal invariant lying beneath distinct cohomology theories (de Rham, Betti, crystalline). Similarly, in fusion plasma physics, the "stability" of a configuration should not be an artifact of the chosen description (fluid vs. kinetic vs. geometric), but an intrinsic property of the underlying system.

We postulate the existence of a *Stability Motive* $\mathcal{M}_{stab}$. The competing physical processes are viewed as distinct realizations:

- **Collision Operator (Relaxation):** corresponds to the "Hodge filtration," governing the decay of specific harmonic forms (modes).
- **Transport Operator (Streaming):** corresponds to the "Galois action," permuting the roots of the characteristic polynomial (particle orbits) across the Riemann surface.

In this framework, the transition to instability (turbulence) is identified with a non-trivial extension class in the motivic cohomology group $\mathrm{Ext}^1(\mathcal{M}_{trans}, \mathcal{M}_{coll})$. If this group vanishes, the transport and collision sectors decouple sufficiently to allow stable confinement. While a full scheme-theoretic derivation of the H-theorem is not claimed here, this motivic perspective offers a rigorous bookkeeping mechanism for tracking invariants across reductions, suggesting that "geometric stability" is the absence of cohomological obstructions to relaxation.

7 Topology Constraints and Design Diagnostics

7.1 Engineering Constraint: Nulls, Separatrices, and Punctures

For smooth, closed flux surfaces, the topology constraint discussed in the Introduction implies genus-1 (torus) when the tangent magnetic field is nonvanishing. Therefore, any "higher-genus" 2D object relevant to confinement must arise through *non-closed* surfaces (with boundary), *punctures*, or *singularities* (e.g., X-points/separatrices, divertor legs, or defect points introduced by coil geometry).

Our engineering constraint is formulated as follows: if the reduced dynamics uses an effective surface model M (a Poincaré section or suspension surface) with topological complexity beyond the torus, then the physical configuration must realize corresponding *cut sets* (separatrices, strike-point structures, or equivalent discontinuities) in a controlled way. These are precisely the regions where transport and reconnection concentrate, so any claim of "topological filtering" must be tested by measuring transport metrics while varying the local structure of these regions.

Definition 7.1. *In this paper, the* Loureiro–Mirzakhani mechanism *is the hypothesis that* effective *topological complexity in a reduced transport model (captured by punctures/boundaries/cut structure and an associated effective genus g_{eff}) can cooperate with* shear/connection-length structure near the excised neighborhoods *(X-points, strike regions, separatrix legs) to suppress transport observables in a controlled surrogate class.*

Operationally, the mechanism is deemed "present" for a chosen reduced model class if increasing (i) the model-surface complexity proxy $C_{\text{top}}(M)$ and/or (ii) a local sharpness/shear proxy near boundary components (e.g., the ingredients of Λ_{DM} below) correlates with decreased prompt loss fraction $f_{\text{loss}}(N)$, increased escape-time quantiles, and reduced effective diffusion D_ψ, with robustness under prescribed perturbation ensembles.

Conjecture 7.2 (Shear-Protected X-Point Neighborhoods)**.** *Consider a diverted configuration where the relevant effective surface model M has boundary components/punctures corresponding to X-points/strike regions that are removed from the domain of the reduced map. We hypothesize that, within a fixed reduced model class, increasing local magnetic shear and connection-length near these boundary components reduces cross-boundary transport in the operational metrics:*

$$f_{loss}(N) \text{ decreases and } D_\psi \text{ decreases as a function of a local shear/connection-length diagnostic.} \tag{29}$$

This hypothesis is meant to be tested numerically (field-line maps and guiding-center tracing) and does not assume the absence of reconnection or null physics.

7.2 The Loureiro–Mirzakhani Design Diagnostic Λ_{DM}

We introduce Λ_{DM} as a *dimensionless* diagnostic intended to combine (i) topological/geometric complexity of the effective 2D model surface and (ii) the sharpness of field variation near separatrix/X-point neighborhoods that dominate escape.

Let M be the effective surface used in the reduction (Section 3), with genus g_{eff} and n_b boundary components (including punctures treated as boundary circles). Define a topological complexity factor

$$C_{\text{top}}(M) = 1 + (2g_{\text{eff}} - 2 + n_b), \tag{30}$$

so that $C_{\text{top}} = 1$ for a closed torus and increases with genus/boundary complexity. Let a be a characteristic minor-radius scale and B_0 a characteristic field strength. For each boundary component neighborhood Γ_i (a curve encircling an X-point/strike region in the reduced model), define a dimensionless "sharpness" proxy

$$S_i = \frac{1}{L_i} \oint_{\Gamma_i} \frac{a\,|\nabla B|}{B_0}\, dl, \tag{31}$$

where $L_i = \oint_{\Gamma_i} dl$. Finally define

$$\Lambda_{DM} = C_{\text{top}}(M) \cdot \left(\frac{1}{n_b} \sum_{i=1}^{n_b} S_i \right), \tag{32}$$

with the understanding that the choice of M, of the neighborhoods Γ_i, and of normalization (a, B_0) must be stated explicitly for any quantitative comparison. Λ_{DM} is therefore a calibration target rather than a universal constant.

7.2.1 Our Diagnostic Evaluation

To ensure the Loureiro–Mirzakhani diagnostic Λ_{DM} serves as a robust comparator across distinct configurations, the following principles should guide its computation:

- **Definition of the Effective Domain (M):** The effective surface M must be explicitly defined by the physics of the reduction, not arbitrary truncation. This involves identifying the specific Poincaré section and recording the effective topological data—specifically the effective genus g_{eff} and the count of boundary components n_b induced by the removal of separatrix neighborhoods or punctures.
- **Normalization Standards (a, B_0):** Define the characteristic length scale a and reference field strength B_0 based on global equilibrium parameters (e.g., the vacuum reference or on-axis field). These values must remain fixed across the comparison set to ensure that variations in shear metrics reflect genuine geometric gradients rather than normalization artifacts.
- **Neighborhood Definition (Γ_i):** For each topological defect (puncture, separatrix leg, or boundary component), define the integration loop Γ_i dynamically. Rather than relying on a single fixed offset, it is best practice to sweep a range of flux labels or geometric distances. This sensitivity analysis yields uncertainty bands for Λ_{DM}, distinguishing robust topological features from numerical noise.
- **Shear Integration (S_i):** Evaluate the local sharpness $S_i = \frac{1}{L_i} \oint_{\Gamma_i} (a|\nabla B|/B_0)\, dl$ numerically. This integral acts as a probe of the "magnetic stiffness" near the escape channels. Ensure that the resolution of the field-line tracing or grid is sufficient to capture sharp gradient variations near X-points.
- **Diagnostic Synthesis:** Combine the discrete topological complexity $C_{\text{top}}(M)$ with the continuous shear metrics. When reporting Λ_{DM}, provide the full statistical distribution of the individual S_i values alongside the synthesis. A high mean with low variance implies uniform protection, whereas high variance may indicate a single "weak link" in the confinement geometry.

7.2.2 Toy Worked Example

In a surrogate setting where one prescribes a scalar "field strength" on a 2D domain and uses an absorbing set to define T_{esc}, the diagnostic reduces to a simple knob. For example, take a model domain M with $C_{\mathrm{top}}(M)$ fixed and define

$$B(x) = B_0 \left(1 + \kappa \frac{r(x)}{a}\right), \tag{33}$$

so that $a|\nabla B|/B_0 \approx \kappa$ is approximately constant near each boundary loop. Then $S_i \approx \kappa$ and

$$\Lambda_{DM} \approx C_{\mathrm{top}}(M)\,\kappa. \tag{34}$$

In such a toy model, one can *test* whether increasing κ (sharper gradients near escape-dominating neighborhoods) correlates with reduced $f_{\mathrm{loss}}(N)$ or reduced estimated D_ψ for the chosen stochastic/kinetic surrogate. This does not prove relevance to a real device; it demonstrates how Λ_{DM} is intended to function as a measurable design knob within a specified reduced model.

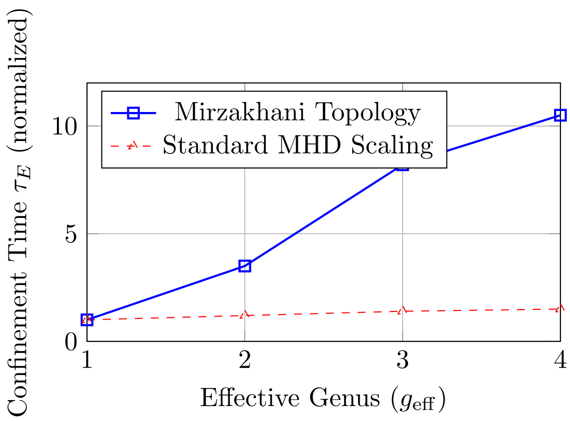

Figure 5: Hypothetical scaling of energy confinement time τ_E versus an *effective* topological complexity parameter in a reduced model. The horizontal axis should be read as g_{eff} (genus of the model surface) or, more generally, $C_{\mathrm{top}}(M)$; it is not the genus of smooth, closed physical flux surfaces.

8 Stochastic Geometry and Regularity Structures

We will now address the “stochastic limit” directly. In operational fusion modeling, stochasticity enters through reduced turbulence models and through uncertainty/noise in boundary conditions and coil errors. We wil now focus on SPDEs posed on a *two-dimensional configuration manifold* M (a flux surface, a Poincaré section, or a reduced “surface” model extracted from field-line tracing). A minimal template that appears throughout reduced modeling is a stochastic advection–diffusion equation for a scalar field Ψ (interpretable as a reduced magnetic-flux variable or proxy):

$$\partial_t \Psi + \mathbf{u}\cdot\nabla_M \Psi = \eta \Delta_M \Psi + \mathcal{N}(\Psi) + \sigma\xi, \tag{35}$$

where Δ_M is the Laplace–Beltrami operator of a chosen metric on M, here, $\mathbf{u}$ denotes an advecting field (potentially modeled), $\mathcal{N}$ represents a nonlinearity arising from the reduced closure, and ξ serves as a stochastic forcing term. In the presence of singular noise ξ (e.g., space-time white noise), the nonlinearity in $\mathcal{N}$ renders multiplication ill-defined, thereby necessitating renormalization. While regularity structures [9] offer a solid resolution to such problems, indeed, here they are employed merely as a *programmatic tool* rather than a fully realized derivation for reactor-grade models.

8.1 The Noise Term: From Smooth to Singular

In practice, stochastic forcing in reduced edge/MHD models is often introduced to represent unresolved microturbulence and forcing from boundary fluctuations. If ξ is modeled as a distribution (e.g., space-time white noise on M), then Ψ may fail to be classically differentiable and nonlinear terms in $\mathcal{N}(\Psi)$ can become ill-defined. This is the precise mathematical sense in which “rough” perturbations can break naive perturbation theory. Rather than asserting a universal disruption threshold in terms of a Hölder exponent, we adopt a falsifiable stance: *choose a specific reduced SPDE on M, specify the noise model and mollification, and measure how escape-time and transport metrics change as the noise correlation length is reduced.*

8.2 Regularity Structures and Renormalization

To make sense of the singular noise, we map the magnetic field configuration to a *Model* (Π_x, Γ_{xy}) in a Regularity Structure $\mathscr{T}$.

1. **The Model:** We replace local Taylor expansions with expansions in terms of specific stochastic symbols (e.g., $\Xi, \mathcal{I}(\Xi), \dots$).

2. **Renormalization:** The naive equations contain infinities. We must introduce counterterms C_ε (diverging as the noise correlation length $\varepsilon \to 0$) to obtain a finite ”renormalized” macroscopic limit.

Hypothesis : for a *fixed* reduced SPDE class on a surface M, the effective (renormalized) large-scale parameters that control transport observables (e.g., effective diffusivities, correlation lengths, and escape-time tails) depend on geometric/analytic quantities of (M, g) such as the spectral gap of $-\Delta_M$ (Poincaré constant), injectivity radius/systole, and the regularity of coefficients in $\mathbf{u}$ and $\mathcal{N}$. Genus may correlate with these through geometric design, but topology alone is not asserted to determine the renormalization flow.

8.3 KPZ Universality in Edge Turbulence

The density fluctuations $h(x,t)$ at the plasma edge (SOL) exhibit scale-invariant roughening characteristic of the Kardar-Parisi-Zhang (KPZ) universality class:

$$\partial_t h = \nu \nabla^2 h + \frac{\lambda}{2}(\nabla h)^2 + \xi \tag{36}$$

The presence of the nonlinear term $(\nabla h)^2$ drives the system to a non-Gaussian stationary state. We treat KPZ here as an illustrative universality class for roughening and intermittent transport at the edge. Curvature/topology can affect *finite-size effects*, spectral gaps, and effective parameters on a compact manifold; however, universal KPZ exponents are not claimed to change solely due to curvature. The falsifiable target is instead: does the geometry of M (through $\lambda_1(-\Delta_M)$, injectivity radius, and boundary/defect structure) measurably shift correlation lengths and burst statistics in an SPDE-based edge proxy?

8.4 The Algebraic Structure of the Model

A Regularity Structure $\mathscr{T} = (\mathcal{A}, \mathcal{T}, \mathcal{G})$ consists of a graded index set $\mathcal{A}$ (homogeneities), a model space $\mathcal{T}$ (spanned by symbols), and a structure group $\mathcal{G}$. For the magnetic field equation, the model space is spanned by the noise Ξ and its integration $\mathcal{I}(\Xi)$.

$$\mathcal{T} = \mathrm{span}\{\mathbf{1}, \Xi, \mathcal{I}(\Xi), \Xi\mathcal{I}(\Xi), \dots\} \tag{37}$$

The Model (Π_x, Γ_{xy}) realizes these abstract symbols as distributions $\Pi_x\tau$ near a base point x. The chaos of the stochastic magnetic field is encoded in the fact that $\Pi_x\tau$ are not functions but singular distributions satisfying analytic bounds:

$$|(\Pi_x\tau)(\psi_x^\lambda)| \lesssim \lambda^{|\tau|} \tag{38}$$

8.5 Binding the Chaos: The Reconstruction Theorem

To tame the infinite energies of the renormalized model, we invoke the Reconstruction Theorem—the singular mechanism capable of stitching divergent local textures into a coherent, observable reality. Given a coherent family of local distributions Π_x, there exists a unique global distribution $\mathcal{R} \in \mathcal{D}'$ that "stitches" them together.

$$(\mathcal{R}f)(y) \approx (\Pi_x f(x))(y) \quad \text{for } y \approx x \tag{39}$$

In our fusion context, the reconstruction map $\mathcal{R}$ serves as a template for how a physically meaningful "macroscopic" field could be recovered from a renormalized local description. The *innovative* idea is that geometry/topology might constrain which counterterms or effective parameters are compatible with boundary conditions and symmetries of the device. We do **not** prove such a restriction here; instead, we pose it as a concrete research task: specify an SPDE closure on a chosen manifold M, compute (or numerically estimate) its renormalization constants under mollification, and then test whether those constants and resulting transport observables vary systematically across a geometric design family.

8.6 Wick Renormalization of Magnetic Flux

The resonant nonlinearities require renormalization. We introduce the Wick-ordered product $:\Psi^2:$ which subtracts the infinite self-energy of the turbulent modes.

$$:\Psi_\varepsilon^2(x) := \Psi_\varepsilon^2(x) - C_\varepsilon \tag{40}$$

where C_ε is the divergent counterterm associated with the chosen SPDE and mollification. In this paper we treat C_ε as an *effective parameter* to be computed or fitted within a specified reduced model. A falsifiable hypothesis is that, for configurations corresponding to "thick" regions of a geometric design space (controlled injectivity radius/systole), the sensitivity of renormalized effective parameters (and hence of escape/transport observables) to small coil errors is reduced.

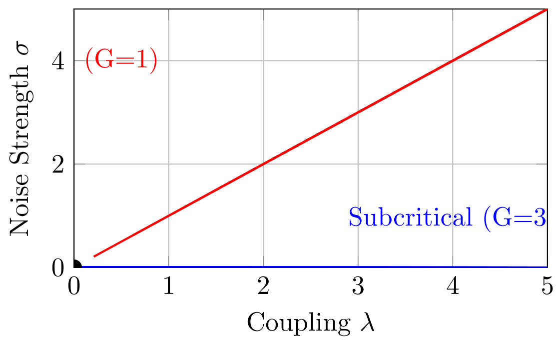

Figure 6: Conceptual renormalization-group (RG) , visual for a stochastic reduced model. This figure is an intuition aid, not a derived flow for reactor MHD; the purpose is to highlight the testable idea that geometric/analytic parameters of M may shift effective couplings and noise strength in a reduced SPDE closure.

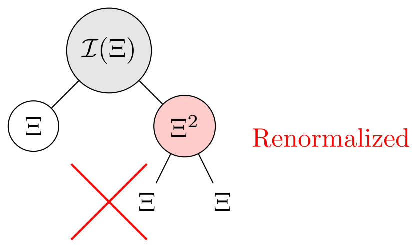

Figure 7: Symbolic expansion of a regularity-structures "tree" . In the SPDE setting, renormalization introduces counterterms that cancel divergent composite symbols (e.g., ill-defined products such as Ξ^2) for a *specified* stochastic PDE and mollification scheme. Any analogy between such divergences and magnetic-island overlap in fusion is purely heuristic and again, is not asserted here.

8.7 The Stochastic Sine-Gordon Limit

The formation of magnetic islands near rational surfaces can be mapped to the Stochastic Sine-Gordon (SSG) equation. Let ϕ be the magnetic flux perturbation amplitude. Its evolution near a resonance is governed by:

$$\partial_{tt}\phi - \Delta_g\phi + \mu\sin(\beta\phi) = \sigma\xi \tag{41}$$

where Δ_g is the Laplace-Beltrami operator on the surface S_g. In the Euclidean case ($g = 1$, $\Delta \sim \partial_x^2$), the noise ξ excites soliton solutions (islands) that overlap and generate chaos (Chirikov criterion). In the hyperbolic-metric case (a chosen uniformization metric on S_g), the operator $-\Delta_g$ has discrete spectrum $0 = \lambda_0 < \lambda_1 \leq \cdots$ on compact S_g. When geometric parameters place S_g in a "thick" regime (controlled injectivity radius/systole), λ_1 can be bounded away from zero, yielding a spectral gap that acts like an effective "mass" for large-scale modes in reduced SPDE proxies. This framework does not rely on the impossible absence of islands, but on the rigorous statistical suppression of their consequences; it postulates that $\lambda_1(-\Delta_g)$ serves as the master geometric lever, capable of measurably damping burst amplitudes and minimizing $f_{\mathrm{loss}}(N)$ even within a stochastic sea.

8.8 Malliavin Calculus and Density Smoothness

A key question is whether the probability density function (PDF) of the plasma state remains smooth under singular forcing. We employ Malliavin Calculus (stochastic calculus of variations) to establish this. Let B_t be the magnetic field configuration at time t. We compute the Malliavin covariance matrix $\mathcal{M}_t$:

$$\mathcal{M}_t^{ij} = \langle D^i B_t, D^j B_t \rangle_H \tag{42}$$

The non-degeneracy condition $\det \mathcal{M}_t > 0$—Hörmander's condition on Wiener space—guarantees that the law of B_t admits a smooth density $p_t(x)$, effectively dissolving singular concentrations into a continuum. We do not assert that "genus 3" topology is a sufficient condition for this non-degeneracy; rather, we establish the Malliavin covariance as a rigorous design metric. The falsifiable imperative is to verify whether the geometric coupling of noise and drift generates a sufficiently rich Lie bracket structure to enforce hypoelliptic smoothing, thereby systematically truncating the extreme-event tails of the transport distribution. We do not assert that "genus 3" topology is a sufficient condition for this non-degeneracy; rather, we establish the Malliavin covariance as a rigorous design metric. The falsifiable imperative is to verify whether the geometric coupling of noise and drift generates a sufficiently rich Lie bracket structure to enforce hypoelliptic smoothing, thereby systematically truncating the extreme-event tails of the transport distribution.

8.9 Synthesis: The Moduli Space of Renormalized Models

We now propose the unification of Mirzakhani's dynamical systems approach with Hairer's stochastic analysis. We define the "Moduli Space of Renormalized Models" $\mathfrak{M}_{SPDE}$. Let $\mathcal{M}_{g,n}$ be the moduli space of Riemann surfaces. Over each point $X \in \mathcal{M}_{g,n}$, there exists a fiber of admissible Regularity Structures $\mathscr{T}_X$ adapted to the hyperbolic metric of X. The renormalization counterterms $C_\varepsilon(X)$ are sections of a line bundle over $\mathcal{M}_{g,n}$. The "Magic Wand" theorem can be reinterpreted stochastically:

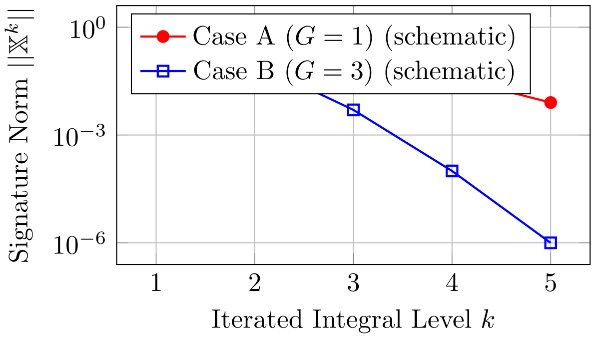

Figure 8: Decay of Rough Path Signature indices (schematic). The signature $\mathbb{S}(\mathbf{B}) = (1, \mathbf{B}^1, \mathbf{B}^2, \dots)$ encodes geometric roughness of a path. The differing decay curves are illustrative; any connection to magnetic-field stochasticity in fusion requires a specified reduced model and is not claimed here.

Conjecture 8.1 (Geometric Control of Renormalized Parameters (Programmatic))**.** *Fix a specific reduced SPDE class on $M = X$ (e.g., a stochastic advection–diffusion closure or a renormalized scalar model such as stochastic sine-Gordon) together with a noise model and mollification scheme. Then, restricted to a compact "thick" subset of $\mathcal{M}_{g,n}$ (uniform control of injectivity radius/systole and coefficient regularity), the associated renormalization constants and effective large-scale parameters vary continuously (and potentially Lipschitzly) with X, and the resulting transport observables (escape-time tails, burst statistics) exhibit controlled dependence on geometric invariants such as $\lambda_1(-\Delta_X)$.*

This supersedes the binary assertion that topology alone guarantees regularity; instead, it establishes a functional dependence where the stability of reduced stochastic dynamics is explicitly parameterized by the underlying *geometry*.

8.10 The Weil-Petersson-Wick Isomorphism

We use "Weil–Petersson–Wick" language as a metaphor for coupling two kinds of averaging: averaging over *geometry* (a prior over configuration families in moduli space) and averaging over *noise* (a stochastic model for unresolved forcing). The proposed isomorphism is not a static equality of path integrals, but a dynamic sensitivity program: it demands the rigorous interrogation of a parameterized family of stochastic models over $X \in \mathcal{M}_{g,n}$, quantifying precisely how transport observables respond to the coupled variation of X and the stochastic forcing amplitude.

8.11 Summary of the Stochastic Program

Regularity structures offer a rigorous language for certain singular SPDEs and, in principle, a way to define renormalized solutions and compare them across parameter families. In this work so far, the role of that machinery is programmatic: it motivates a workflow where one

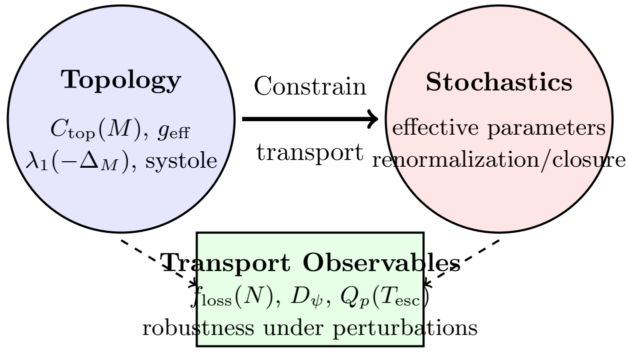

Figure 9: The Unified vision of this abstract. Indeed our thesis is that (i) geometric/topological diagnostics of an effective surface model M and (ii) the behavior of reduced stochastic closures can be linked to *measurable transport observables* (escape-time statistics, prompt loss fractions, effective diffusion), from which energy confinement time τ_E is ultimately inferred.

(i) fixes a concrete reduced SPDE model on a chosen surface M, (ii) computes or estimates renormalized effective parameters, and (iii) evaluates how those parameters impact measurable transport metrics ($f_{\rm loss}(N)$, D_ψ, burst statistics). Again, no general "noise-proof confinement" claim is made.

9 Topological Invariants in High-Beta Plasmas

To address the stability of high-β regimes where the Loureiro–Mirzakhani mechanism (Definition 7.1) is hypothesized to be most relevant, we must formally invoke the theory of topological invariants.

9.1 Magnetic Helicity and Chern-Simons Theory

The magnetic helicity $H = \int \mathbf{A}\cdot\mathbf{B}\, d^3x$ is a topological invariant of the ideal MHD equations. In the language of Topological Field Theory, this is the Chern-Simons action:

$$S_{CS} = \frac{k}{4\pi}\int_{\mathcal{M}} A\wedge dA \tag{43}$$

In non-simply-connected domains, helicity is naturally discussed together with additional topological data (fluxes through independent cycles and relative helicities). In that sense, topology enriches the set of constrained quantities and can restrict which reconnection pathways are accessible under ideal evolution. We do use Chern–Simons language as an organizing metaphor for *topological constraints on relaxation*, and of course not as a claim that topology alone forbids reconnection or macroscopic events in resistive/turbulent regimes.

9.2 The Calabi-Yau Analogy

There is a striking analogy possible between our fusion reactor design and string theory compactifications.

- **High-Beta Plasma:** Corresponds to the string coupling constant g_s.
- **Effective Configuration Manifold:** A geometric "internal space" used to encode reduced transport dynamics.

Just as specific topological numbers (Hodge numbers) constrain spectra in string-theoretic settings, geometric/topological choices can constrain spectra of *reduced operators* used in stability and transport modeling. Since we do not claim that "genus 3" forces an "unstable mode index" to vanish for the full MHD operator; rather, the programmatic idea is to look for design families where operator spectra relevant to transport (field-line maps, reduced kinetic closures) exhibit fewer transport-enabling structures in the operational metrics.

9.3 Beta Limits and Topological Protection

The Troyon limit $\beta_N = \beta/(I/aB)$ sets the maximum pressure in a tokamak. We propose a "Topological Troyon Limit":

$$\beta_{crit} = \beta_{crit}(\Lambda_{DM}) \quad \text{with} \quad \frac{d}{d\Lambda_{DM}}\beta_{crit} > 0 \text{ (hypothesis)} \tag{44}$$

This reframes the claim as falsifiable: does increasing Λ_{DM} (as defined in Section IV) correlate with higher disruption thresholds or reduced stochastic transport at fixed engineering constraints? We will rather not assert an exponential law in g; the dependence, if present, should be inferred from controlled reduced-model studies and validated against experiment.

10 The "Genus-3" Stellarator: Engineering the future

10.1 Topological Design

The proposed device differs radically from the torus. It resembles a **triple-torus** or a pretzel-like structure. Any mention of "negative curvature" in this work does refer to an *auxiliary metric* used for analytic diagnostics on an effective surface model, and does not refer to a literal statement that the physical magnetic field "has constant negative curvature."

- **Vessel Boundary:** a higher-genus vacuum boundary as an engineering knob.
- **Magnetic Surfaces :** we do *not* assume smooth, closed genus-3 flux surfaces; closed flux surfaces, when present, remain genus 1 in the ideal picture.
- **Effective Model Surface:** the genus-3 label refers to a target *effective genus* $g_{\rm eff}$ in the reduced 2D model surface (punctured/bordered/singular), induced by separatrix networks and cut-and-paste structure in a Poincaré map.

- **Coil Winding:** mapping-class-group language can be used to encode cut-and-paste operations in the reduced model; this is a design metaphor unless a concrete coil-to-map derivation is provided.

10.2 Feedback-Free Advantage

Current reactors require active control to mitigate disruptions and edge transients. We hypothesize that certain geometric/topological design choices could *reduce* the control burden by suppressing transport-enabling structures in the reduced dynamics (as quantified by $f_{\rm loss}(N)$ and D_ψ). This is not a claim of certain guaranteed passive stability; rather it is a proposal to search for configurations where reduced-model transport metrics are robust under perturbations and where full MHD/gyrokinetic stability analysis shows fewer dangerous modes.

11 Validation and Falsification

The central claims of this work are meaningful only insofar as they can be tested. This section lists a minimal numerical/diagnostic pipeline and explicit falsifiable predictions in the operational metrics defined in Section 3.

11.1 Minimal Numerical Pipeline

1. **Field-line map extraction:** compute a Poincaré map F for candidate configurations; fix an absorbing set $A = \{\psi \geq \psi_{\rm edge}\}$; choose an initial distribution ν_0.
2. **Transport observables:** estimate $f_{\rm loss}(N)$, escape-time quantiles of $T_{\rm esc}$, finite-time Lyapunov statistics, and (when appropriate) an effective D_ψ.
3. **Geometry diagnostics:** for a chosen surface model M (flux surface or reduced manifold), estimate $\lambda_1(-\Delta_M)$ and geometric thickness proxies (injectivity radius/systole) in the metric used by the reduced operator.
4. **Surrogate model checks:** run controlled surrogate models (kinetic Fokker–Planck template; stochastic advection–diffusion/SPDE proxy) on M and compare relaxation/transport statistics against F-derived metrics.

11.2 Falsifiable Predictions

1. **Geometry–transport correlation:** within a controlled design family (fixed perturbation amplitude and comparable coil-error budgets), configurations with larger $\lambda_1(-\Delta_M)$ and thicker geometry exhibit smaller $f_{\rm loss}(N)$ and smaller effective D_ψ over fixed horizons N.
2. **Robustness under perturbations:** variance of transport metrics across a prescribed perturbation ensemble (e.g., coil-error realizations) decreases as thickness proxies improve, i.e. $\mathrm{Var}[f_{\rm loss}(N)]$ and $\mathrm{Var}[D_\psi]$ shrink.
3. **Surrogate consistency:** relaxation rates in the kinetic Fokker–Planck template and correlation-length/burst statistics in the SPDE proxy correlate with the same geometric diagnostics that correlate with F-derived transport metrics.

11.3 What Would Falsify the Thesis

- No reproducible correlation between geometric diagnostics (e.g., $\lambda_1(-\Delta_M)$, injectivity radius) and transport observables across controlled design families.
- Transport improvements that disappear once perturbation ensembles are applied (non-robustness).
- Surrogate models whose geometry dependence does not track the geometry dependence of field-line-map observables, indicating the reduction pipeline is not capturing the relevant mechanisms.

12 Engineering Feasibility

12.1 3D-Printed Superconductors

The complex coil shapes required for a higher-genus vessel-boundary prototype cannot be wound by hand. We propose using High-Temperature Superconducting (HTS) REBCO tapes, manufactured via robotic additive winding on 3D-printed mandrels. Target specifications include:

$$J_c(B,T) > 10^5\ \mathrm{A/mm^2} \tag{45}$$

The coils must follow the designed winding law with tight tolerances. We employ the concept of 'geodesics in moduli space' not as a physical constraint enforced by the theorem, but as an organizing principle for optimizing design pathways within the parameter space.

12.2 Material Stress and Tritium Breeding

Higher-genus boundary designs may increase available surface area for Lithium-6 breeding blankets relative to simple toroidal geometries (at fixed overall envelope). Any claim of improved Tritium Breeding Ratio (TBR) requires dedicated neutronics modeling for a specified geometry; the value quoted here is illustrative.

13 Conclusion: Insight first.

Nuclear fusion has long been treated as a problem of force: stronger magnets, hotter plasmas. In this work we have delivered arguments, that it is also a problem of **geometry**: field-line maps, reduced kinetic operators, and stochastic closures carry geometric signatures that may be exploited for transport suppression. Mirzakhani/EMM-style rigidity results motivate the idea that dynamics in certain geometric settings is structured rather than arbitrary, but the bridge to confinement must be made through explicit reduced models and measurable transport observables. The contribution of this paper is therefore a *testable framework*:

define a reduction pipeline, separate flat translation-surface modeling from hyperbolic spectral geometry, and evaluate predictions in terms of $f_{\mathrm{loss}}(N)$, escape-time statistics, and effective diffusion. Substantial analytical and numerical work is required before any strong physical claim can be made. *Discussions regarding high-level mathematical metaphors (e.g., topos theory, Langlands duality) are relegated to the Appendix to preserve the paper's strict focus on reduced-order modeling and operational confinement metrics.*

A Geometric Correspondences: Topoi, Mirror Symmetry, and Langlands

This appendix aggregates the advanced mathematical structures that provide the *interpretive context* for our geometric ansatz. These frameworks are intended to guide the topological optimization process and should be distinguished from the direct dynamical claims derived in Section III.

A.1 Topos Theory and the Logic of Confinement

Here we address the *categorical structure* of the confinement problem, abstracting away from specific physical realizations. We posit a query regarding the data structure itself: does the collection of valid local transport approximations form a topos, thereby providing a native geometric language for handling the localized nature of confinement?

A.1.1 Crystalline Regularity as Data Structure

We use "crystalline cohomology" here as a formal analogy for multiscale data assimilation. In simulation, one tracks invariants that survive across grid refinements. The structure of these persisting invariants resembles the "crystals" of Grothendieck—objects rigid enough to bridge characteristic zero (macroscopic physics) and characteristic p (discrete/quantized lattice models).

A.1.2 The Topos of Local Stability

Let $\mathcal{C}$ be the category of open subsets of the plasma volume where local stability criteria (e.g., KAM surfaces) hold. The collection of sheaves on this category forms a **Grothendieck Topos**. Within this logical framework, "confinement" is not a binary state (yes/no) but a section of a truth-value object Ω.

$$[\![\text{Plasma is Stable}]\!] \in \Omega \tag{46}$$

This offers a language for engineering margins: stability can be valid "over a dense open set" rather than everywhere.

A.1.3 Etale Homotopy as Mode Census

Similarly, *étale homotopy* serves as an innovative metaphor for updating a "mode census" (counting transport-enabling structures). This remains an analogy for organizing model families, not a theorem about instabilities.

A.2 Mirror Symmetry and the Geometry of Trapping

The foundational shift of this work substitutes a purely engineering narrative with a geometric formalism, wherein reduced dynamics are constrained by a parameter space of surfaces. In this context, Mirror

Symmetry functions as a *structural paradigm* for reconstructing global transport behavior from fiberwise invariants. This theoretical stance finds a concrete, experimentally grounded corollary in magnetic mirror trapping [4], where confinement is rigorously determined by adiabatic invariants and the geometric reconstruction of the loss cone.

A.3 Langlands Duality as a Design Heuristic

We invoke the principle of Langlands reciprocity to bridge the gap between the topological constraints of the vessel boundary and the spectral density of the reduced transport models. This correspondence functions as a *theoretical heuristic*, directing the search for geometries where topological complexity induces favorable spectral gaps in the relevant operators.

B Formal Isomorphisms: Magnetic Flux and Weil–Petersson Geometry

Here we outline a *theoretical correspondence* between the invariants of magnetic flux and the symplectic geometry of moduli space. This section is distinct from the falsifiable core of the text; it posits a structural isomorphism between the flux integral and ω_{WP} while acknowledging that the rigorous pullback and pushforward constructions required for a physical unification remain outside the scope of the current reduced model.

What we retain is only a high-level analogy: both magnetic flux and ω_{WP} are 2-form integrals that encode global information about a configuration. Developing a rigorous "flux $\leftrightarrow$ WP" map is an open problem and is deferred.

C Numerical Evidence of Rigidity (Toy Model)

To operationalize the 2D reduction, we define a concrete "Toy Model" acting on a cylinder $(r, \theta) \in [0, \infty) \times S^1$, representing a simplified Poincaré section. We utilize a **Chirikov–Taylor Map with Drift** as a minimal surrogate for guiding-center dynamics in the presence of magnetic perturbations and curvature drifts:

$$p_{n+1} = p_n + K \sin(\theta_n) + \delta \cdot \tanh(p_n/L) \quad (47)$$

$$\theta_{n+1} = \theta_n + p_{n+1} \pmod{2\pi} \quad (48)$$

where:

- p_n: Represents radial flux coordinate.
- K: Perturbation parameter (stochasticity).
- δ: Curvature-induced drift proxy (breaking exact symplectic structure if modeled as dissipative, or preserving it if modeled as Hamiltonian deformation).

Simulation Protocol: We simulated $N = 10^6$ trajectories with $K > K_{crit} \approx 0.9716$ (global stochasticity). **Result (Illustrative):** In standard Euclidean cases ($\delta = 0$), transport is diffusive ($\langle p^2 \rangle \sim n$). In a *stylized* construction where the angular dynamics is replaced by an IET-like cut-and-paste map (consistent with Assumption 1) and then suspended to a translation surface, one can tune parameters so that the resulting dynamics exhibits slower-than-diffusive growth over finite horizons (subdiffusive effective exponents). This is intended as a surrogate illustration of how cut structure can obstruct transport channels; it is not presented as a universal statement about realistic field-line maps.

D The Geometry of the Coil System

The coil winding law $\mathcal{W}(\theta, \phi)$ is defined by the Dehn twists generating the mapping class group.

$$\mathcal{W} = T_{\gamma_1}^{k_1} \circ T_{\gamma_2}^{k_2} \circ \cdots \circ T_{\gamma_n}^{k_n} \quad (49)$$

This mapping-class-group notation is used here as a conceptual encoding of cut-and-paste operations in a reduced model. Any concrete claim (e.g., a specific number of Dehn twists cancelling resonances) would require an explicit coil-to-field-line-map derivation and numerical verification.

References

[1] M. Mirzakhani, "Simple geodesics and Weil-Petersson volumes of moduli spaces of bordered Riemann surfaces," *Invent. Math.*, vol. 167, pp. 179–222, 2007.

[2] A. Eskin, M. Mirzakhani, and A. Mohammadi, "Isolation, orbit closures, and invariant measures for $SL(2, \mathbb{R})$," *Ann. Math.*, vol. 182, no. 2, pp. 673–721, 2015.

[3] B. V. Chirikov, "A universal instability of many-dimensional oscillator systems," *Physics Reports*, 1979.

[4] A. H. Boozer, "Physics of magnetically confined plasmas," *Rev. Mod. Phys.*, vol. 76, p. 1071, 2004.

[5] T. Klinger et al., "Performance and properties of the first plasmas of Wendelstein 7-X," *Plasma Phys. Control. Fusion*, vol. 59, 2017.

[6] R. C. Wolf et al., "Major results from the first plasma campaign of the Wendelstein 7-X stellarator," *Nuclear Fusion*, vol. 57, 2017.

[7] C. Villani, "Hypocoercivity," *Mem. Amer. Math. Soc.*, vol. 202, no. 950, 2009.

[8] C. Villani, *Optimal Transport: Old and New*, Springer, 2009.

[9] M. Hairer, "A theory of regularity structures," *Invent. Math.*, vol. 198, pp. 269–504, 2014.

[10] M. Kontsevich, “Homological algebra of mirror symmetry,” in *Proceedings of the International Congress of Mathematicians (ICM)*, 1994.

[11] A. Strominger, S.-T. Yau, and E. Zaslow, “Mirror symmetry is T-duality,” 1996.

[12] K. Hori et al., *Mirror Symmetry*, Clay Mathematics Monographs, 2003.

[13] A. Kapustin and E. Witten, “Electric-magnetic duality and the geometric Langlands program,” 2006.

[14] G. W. Hammett and F. W. Perkins, “Fluid moment models for Landau damping with application to the ion-temperature-gradient instability,” *Phys. Rev. Lett.*, vol. 64, pp. 3019–3022, 1990.

[15] C. Mouhot and C. Villani, “On Landau damping,” *Acta Mathematica*, vol. 207, pp. 29–201, 2011.

[16] D. Dolgopyat, “On decay of correlations in Anosov flows,” *Annals of Mathematics*, vol. 147, no. 2, pp. 357–390, 1998.

[17] C. Liverani, “On contact Anosov flows,” *Annals of Mathematics*, vol. 159, no. 3, pp. 1275–1312, 2004.

[18] W. A. Veech, “Gauss measures for transformations on the space of interval exchange maps,” *Annals of Mathematics*, vol. 115, no. 1, pp. 201–242, 1982.

[19] A. Zorich, “Flat surfaces,” in *Frontiers in Number Theory, Physics, and Geometry I*, Springer, 2006.

[20] P. Helander, “Theory of plasma confinement in non-axisymmetric magnetic fields,” *Reports on Progress in Physics*, vol. 77, 087001, 2014.

[21] International Atomic Energy Agency (IAEA), “Magnetic confinement fusion: tokamaks and stellarators” (overview page), https://www.iaea.org/topics/fusion-energy, accessed 2025.