

Certified Golden-Branch Maximality in a Generated Standard-Map Domain

Surya Tallavarjula^{1,2,*} and Seyed Ali Rastegar³

¹Department of Physics, University of California, Berkeley, Berkeley, CA, United States of America

²Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA, United States of America

³Department of Mathematics, University of California, Berkeley, Berkeley, CA, United States of America

*Author to whom any correspondence should be addressed.

E-mail: `surya_talla@berkeley.edu`

Keywords: standard map, invariant circles, conservative twist maps, Greene residue criterion, computer-assisted proof, validated numerics, golden mean

Abstract

We prove a computer-assisted threshold comparison theorem for invariant circles in the standard-sine conservative twist map, over a finite generated comparison domain fixed in advance. In the v1.0.1 theorem instance, the nongolden records are routed through a compact grammar consisting of two screened labels, two exact ranking records, one pruned symbolic region, one inherited-domination record, two terminal candidates with explicit exclusions, and no omitted-tail records. For this generated domain, fixed threshold branch, and fixed normalization of rotation classes, we certify that the golden rotation class has the largest threshold among the generated classes. The proof combines an a posteriori KAM certificate for the golden branch, interval-validated rational obstruction data, rational-to-irrational transport, branch identification, and finite discharge of the generated arithmetic records. The final theorem is the outward-rounded endpoint comparison

$$U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^-.$$

The result is not a proof of Greene’s conjecture over all irrational rotation numbers and is not a universality theorem for arbitrary twist maps. Its contribution is a reproducible proof architecture showing how Greene-style near-critical numerical evidence can be converted into a closed interval theorem for a specified standard-map comparison problem.

Mathematics Subject Classification (2020). 37J40; 37M21; 65G20; 65P10; 65G40.

Proof package. Release v1.0.1; archive DOI: 10.5281/zenodo.20101820.

1 Introduction

The standard map is a central model for Hamiltonian chaos, invariant-circle breakup, and the transition from ordered to stochastic motion in conservative dynamics. In the small-perturbation regime, KAM theory proves persistence of many quasiperiodic invariant objects under arithmetic and nondegeneracy hypotheses [1, 2, 3, 11, 29]. Near criticality, the question becomes sharper: for a specified rotation class, where is the breakup threshold, and how can thresholds for different rotation classes be compared with mathematically controlled error?

Greene’s residue criterion gave a powerful numerical framework for this problem by following periodic orbits whose rotation numbers approximate a target irrational [4]. For the standard map, Greene’s computations and MacKay’s renormalization picture placed the golden mean at the center of the critical-threshold story [5]. The golden class is the natural extremal candidate: its continued-fraction expansion is maximally resistant to rational approximation, and its critical behavior has

long served as a benchmark for numerical, renormalization, and rigorous-computational studies of area-preserving maps [28, 27, 21]. Aubry–Mather theory gives a complementary variational language for monotone twist maps, including ordered orbits and remnants beyond the smooth invariant-curve regime [6, 7, 8].

The challenge is that this numerical and theoretical picture does not by itself prove threshold maximality. A residue crossing suggests a breakup location for one branch, but it does not automatically provide a lower existence theorem below that value, an upper nonexistence theorem above it, a comparison across nearby irrational classes, or a discharge of the arithmetic cases under consideration. Likewise, a broad search over continued fractions is not a proof unless the generation grammar, omitted tails, branch choices, and endpoint errors are closed by explicit estimates. The problem addressed here is therefore not to produce another high-resolution computation, but to turn Greene-style near-critical evidence into a finite, replayable proof.

The specific gap addressed here is the absence of a closed comparison pipeline. Existing rigorous KAM machinery can certify persistence of a chosen invariant circle, and converse-KAM methods explain how ordered barriers obstruct persistence, but those ingredients do not by themselves produce a threshold comparison across a generated arithmetic family. The contribution of this paper is to join lower existence, branch-scoped upper obstruction, rational-to-irrational transport, branch identification, finite arithmetic discharge, and final normalization into one auditable certificate. The theorem is finite, but the dependency structure identifies exactly which estimates must be enlarged in any future continuum or larger-grammar extension.

This paper proves a bounded golden-branch maximality theorem for the standard-sine conservative twist map. We fix a map family, a threshold functional, a generated arithmetic comparison domain $\mathcal{D}_{\text{cert}}$, a branch convention, and a $\text{GL}(2, \mathbb{Z})$ normalization of rotation-class representatives. Inside that certified universe, the golden rotation class has the largest certified threshold, and equality occurs only for the normalized golden representative. The final comparison is the strict interval endpoint inequality

$$U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^-.$$

The left endpoint is the largest generated nongolden upper ceiling after upper obstruction, transport, branch identification, and generated-domain discharge. The right endpoint is the direct lower persistence anchor for the golden branch. The computation enters only after it has been translated into interval enclosures, symbolic partitions, accepted proof fields, and archived replay objects [12, 11, 13, 14, 15, 16, 21, 33, 17, 18, 19, 20].

The scope is fixed once and for all in Section ???. Within that boundary, every generated nongolden record is assigned a closed route to an upper threshold below the golden lower anchor, and the final reduction uses only the certified objects recorded in the proof ledger. The result is therefore both a theorem for a fixed standard-map comparison problem and a reproducible method for promoting near-critical numerical evidence to a structured interval proof.

Figure 1 illustrates why the proof is naturally asymmetric. Persistence below a parameter value is a lower existence statement for an invariant graph. Breakup above a parameter value is an upper obstruction statement for a specified branch. A phase portrait can motivate both claims, but it cannot join them into a threshold theorem. The joining occurs only after the lower anchor, upper ceilings, transport estimates, branch-identification witnesses, generated-domain discharge, and final endpoint comparison have all been validated.

1.1 Position within the KAM, Greene, and computer-assisted proof literature

This work sits at the intersection of four traditions. The first is classical KAM and twist-map theory: invariant curves persist under suitable smallness and arithmetic assumptions, while monotone twist




Figure 1: Phase-space orientation near the threshold regime. Panel A shows a subcritical reference at $K = 0.9716360$, with the 144/233 golden periodic surrogate drawn on the surrounding phase portrait. Panel B shows a nearby supercritical reference at $K = 0.9720370$, after visible breakup of the displayed surrogate structure. The figure gives dynamical context only: persistence, obstruction, transport, generated-domain discharge, and the final endpoint comparison are certified by the interval objects and ledgers stated later, not by visual inspection of the point clouds.

maps retain variationally ordered structures even beyond smooth persistence [1, 2, 3, 11, 29, 6, 7, 8]. The present theorem does not replace that theory; it uses it to formulate the map family, rotation class, invariant-circle problem, and threshold functional.

The second tradition is Greene’s residue criterion and the standard-map critical-threshold literature. Residue computations organize rational approximants into an effective diagnostic for criticality, and the golden mean is the canonical extremal candidate in that framework [4, 5, 28, 27]. Here, residue and crossing data are not treated as proof by themselves. They enter only after being converted into interval crossing windows, obstruction ledgers, tail/coherence checks, and endpoint inequalities.

The third tradition is computer-assisted proof for dynamical systems, including a posteriori KAM validation, interval Newton and Krawczyk methods, validated continuation, and proof-carrying numerical enclosures [14, 15, 16, 12, 11, 30, 31, 13, 21, 25, 26, 34, 33]. On the lower side, the proof uses an a posteriori existence certificate for the golden branch. On the upper side, it uses a branch-scoped obstruction principle closer in spirit to converse-KAM methods than to a floating-point residue plot [23, 22, 24, 32].

The fourth tradition is reproducible computational mathematics. The proof package uses interval semantics, routing records, manifests, checksum validation, negative controls, replay scripts, and a command-line reconstruction path so that the final small endpoint margin is attached to auditable objects rather than to informal file inspection [17, 18, 19, 20, 35]. These details matter because the theorem’s final separation is only 3.0×10^{-7} .

Throughout the paper, “golden maximality” has the bounded meaning fixed in Section ??: maximality is asserted only inside the generated comparison domain $\mathcal{D}_{\text{cert}}$ for the fixed standard-

sine map and the stated normalization conventions.

Relative to these traditions, the paper is best read as a methodology-first theorem paper. It does not introduce a new perturbative KAM theorem, a new general converse-KAM theorem, or a new numerical criterion for criticality. Instead, it assembles existing analytic ideas—a posteriori invariant-circle validation, interval Newton/Krawczyk enclosures, branch-local obstruction logic, and continued-fraction ordering—into a single finite comparison certificate. This places the work closer to validated-numerical and computer-assisted dynamics than to a claim of a new general KAM or converse-KAM theorem [25, 30, 31, 21, 32, 33].

The closest lower-side antecedent is the computer-assisted KAM framework of Figueras–Haro–Luque and the parameterization-method literature on rigorous invariant tori and invariant curves [30, 31, 13, 21]. The lower certificate in Theorem III is in that tradition: it verifies a residual, an approximate inverse/cohomology bound, nonlinear remainders, and Fourier tail estimates in a weighted analytic norm. The role of the lower certificate is to serve as one endpoint in a closed threshold comparison. The paper then joins that endpoint to branch-scoped upper obstruction, rational-to-irrational transport, threshold identification, and finite generated-domain discharge. In this sense, the contribution begins where a standalone existence certificate ends: it turns one validated near-critical existence statement into one side of an auditable comparison theorem.

This positioning also explains the role of the archive, whose purpose is to store the finite certificate data, replay the interval and symbolic checks, and make the small endpoint margin reproducible. The mathematical implications used by the theorem are the certificate-implies-theorem statements in the body of the paper; the archive records the data and computations that verify their hypotheses.

1.2 Main contribution and proof strategy

The contribution is a closed threshold-comparison pipeline rather than a refined numerical estimate. The result should be read as methodology-first: the paper proves a concrete generated-domain comparison theorem, and the proof architecture is the main object whose structure is meant to be reusable. The theorem instance supplies a direct lower a posteriori certificate for the golden branch, a branch-scoped upper-obstruction certificate for generated nongolden records, a rational-to-irrational transport estimate, a branch-identification step, a finite generated-domain discharge, and a final endpoint comparison.

The size of the generated domain is intentionally stated at the beginning rather than left to the supplementary artifact guide. In the archived release considered here, the nongolden records are routed through two screened labels, two exact ranking records, one pruned symbolic region, one inherited-domination record, two terminal candidates with explicit exclusion estimates, and no omitted-tail records inside the generated grammar. Thus, the theorem is not a continuum theorem and should not be evaluated as one. Its value is that the finite generated comparison problem is closed by explicit lower, upper, transport, identification, discharge, and normalization estimates, giving a model for what larger generated domains would require.

The proof strategy follows the final inequality backward. The lower certificate creates the golden anchor

$$K_G^- = 0.9716350 \leq \Lambda_f(\rho_G).$$

The upper-obstruction, transport, identification, and discharge stages assign every generated nongolden class $\rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}})$ a certified ceiling U_ρ satisfying

$$\Lambda_f(\rho) \leq U_\rho \leq U_{\text{ng}}^+ = 0.9716347.$$

The final reduction then uses no new numerical search; it applies the outward-rounded endpoint

comparison

$$U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^-.$$

The accompanying proof package records the theorem-facing data, replay scripts, reconstruction commands, manifests, and negative-control tests needed to audit this finite certificate [35].

1.3 What scales beyond the compact generated domain

The compact size of the present generated domain is a substantive part of the theorem statement, not an implementation accident. The release used here closes a finite arithmetic threat model: two screened labels, two exact ranking records, one pruned symbolic region, one inherited-domination record, two terminal candidates with explicit exclusion estimates, and no omitted-tail records inside the generated grammar. The natural question is therefore what a larger generated-domain theorem would require.

A useful way to state the methodology contribution is by separating the components that scale mechanically from those that would require new mathematical input. The interval Newton/Krawczyk isolation of rational branches, residue/crossing evaluation, hash-checked replay, endpoint comparison, and finite routing checks scale in the ordinary computational sense: increasing the number of generated continued-fraction cylinders or rational support denominators mainly increases the number and size of certified records. The lower a posteriori KAM certificate also scales in a standard way once a candidate branch and analytic norm have been fixed, although near-critical validation can become numerically expensive.

The non-mechanical parts are the ones that any extension to all continued-fraction coefficients bounded by a larger integer N would have to supply explicitly. First, the arithmetic generator would need a nonvacuous ranking or envelope theorem showing that every newly admitted cylinder is either dominated by an existing route or assigned to a new finite obstruction calculation. Second, the omitted-tail component would need a genuine tail ceiling rather than the vacuous complement used in this release. Third, every new near-top or terminal candidate would need branch-coherent upper obstruction data and a transport budget that survives the same endpoint comparison logic. Thus, the architecture is not merely a bookkeeping convention: it identifies exactly where a larger theorem would be routine replay, where it would be larger finite computation, and where new analytic estimates would be needed.

This is why the present theorem is framed around the proof architecture. The mathematical claim is finite, but the proof is organized so that enlarging the generated domain has a clear checklist: extend the arithmetic grammar, certify the new rational branches, prove branch-coherent upper ceilings, transport them to the irrational threshold convention, and rerun the finite discharge and endpoint reduction.

The positive statement is narrower and stronger within its boundary: every nongolden class generated by $\mathcal{D}_{\text{cert}}$ receives a certified upper threshold below the direct golden lower anchor. Enlarging the conclusion would require enlarging the generated domain and adding the corresponding lower, upper, transport, identification, and discharge estimates.

1.4 Reader’s roadmap

The manuscript can be read from the final inequality backward. The final theorem uses

$$U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^-.$$

Everything before that comparison justifies the two endpoints and proves that they apply to the same threshold functional over the generated comparison domain. Theorem III supplies the golden

lower endpoint. Theorem IV supplies the upper obstruction. Theorem V transports rational obstruction to irrational threshold bounds. The threshold-identification section aligns the branch convention. Theorem VI screens and ranks near-top challengers. Theorem VII discharges the generated arithmetic domain. Theorem VIII performs the final reduction. Supplementary Material S1 records implementation details, replay and reconstruction commands, checksum policy, validator semantics, negative controls, and additional scope clarifications.

1.5 How to read the proof objects and figures

The figures are included to orient the reader, not to replace validated inequalities. Figure 1 is purely illustrative. Figures 2–6 summarize, respectively, the dependency graph, endpoint ledgers, branch-obstruction intervals, threshold-identification seam, and support/discharge data. In every case the proof-bearing objects are the interval inequalities, symbolic routing records, and certificate lemmas stated in the text and reproduced by the archive; no visual crossing, apparent breakup, or plotted gap is used as a proof step.

2 Map family, arithmetic domain, and threshold functional

2.1 The map family

The implementation uses a lift form for a generalized standard-map forcing:

$$r_{n+1} = r_n + Kf(\theta_n), \quad \theta_{n+1} = \theta_n + r_{n+1}, \quad (1)$$

where $\theta \in \mathbb{R}$ is a lift of the circle coordinate, $r \in \mathbb{R}$ is the action-like coordinate, and K is the perturbation parameter. The associated annulus map is obtained by reducing θ modulo one. The default standard-map forcing is

$$f(\theta) = \frac{1}{2\pi} \sin(2\pi\theta). \quad (2)$$

The generalized harmonic family allows finite trigonometric forcings of the form

$$f(\theta) = \sum_{j=1}^m \frac{a_j}{2\pi\ell_j} \sin(2\pi\ell_j\theta + \varphi_j). \quad (3)$$

Definition 2.1 (Admissible forcing class). *Let $\mathfrak{F}_{\text{cert}}$ denote the set of harmonic forcings of the form (3) admitted by the proof certification path. Unless explicitly stated otherwise, the main theorem specializes to the standard forcing (2) with family label **standard-sine**.*

Definition 2.2 (Standard-sine conservative twist family). *For $f \in \mathfrak{F}_{\text{cert}}$ and K in the validated parameter regime, define*

$$F_{K,f}(\theta, r) = (\theta + r + Kf(\theta), r + Kf(\theta)).$$

The fixed family $\mathcal{F}_{\text{cert}}$ is the collection of those $F_{K,f}$ for which the proof arguments required by the final replay are available and accepted. In the archived theorem instance considered here, the proof family is the standard-sine conservative twist family specified by the archived theorem universe.

Remark 2.3. *The family is defined through both analytic form and validated numerical data because the result is a validated threshold theorem, not a perturbative KAM theorem for an unspecified open class of maps.*

2.2 Rotation classes and modular normalization

Definition 2.4 (Irrational rotation class). *An irrational rotation class is an equivalence class of irrational rotation numbers represented by a continued-fraction expansion*

$$\rho = [a_0; a_1, a_2, \dots],$$

together with the normalization conventions used by the generated arithmetic-ranking and final-reduction proof components.

Definition 2.5 (Golden class). *The golden class ρ_G is the normalized class represented by*

$$\rho_G = [1; 1, 1, 1, \dots],$$

or by its projectively equivalent representative in the final normalization domain. Its $\text{GL}(2, \mathbb{Z})$ orbit is denoted

$$\mathcal{O}_G = \text{GL}(2, \mathbb{Z}) \cdot \rho_G.$$

Definition 2.6 ($\text{GL}(2, \mathbb{Z})$ action). *For*

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}(2, \mathbb{Z}),$$

the projective action on an irrational rotation number is

$$A \cdot \rho = \frac{a\rho + b}{c\rho + d},$$

whenever the denominator is nonzero.

Definition 2.7 (Normalization domain). *A normalization domain $\mathcal{D}_{\text{norm}} \subset \mathbb{R} \setminus \mathbb{Q}$ is a selected set of representatives for the projective arithmetic classes considered by the final reduction. The final reduction proves that the golden orbit has a unique representative in $\mathcal{D}_{\text{norm}}$ and that the threshold comparison is invariant under the admissible normalization.*

Definition 2.8 (Admissible arithmetic domain). *The admissible arithmetic domain $\mathcal{D}_{\text{cert}}$ is the set of irrational classes in $\mathcal{D}_{\text{norm}}$ covered by the generated-domain discharge theorem and the final reduction. Equivalently,*

$$\mathcal{D}_{\text{cert}} = \mathcal{P}_{\text{screen}} \cup \mathcal{D}_{\text{rank}} \cup \mathcal{D}_{\text{prune}} \cup \mathcal{D}_{\text{life}} \cup \mathcal{D}_{\text{term}} \cup \mathcal{D}_{\text{omit}},$$

where the six terms are the screened finite panel, exact near-top ranking domain, theorem-pruned domain, previously dominated records, terminal candidates converted to theorem-level exclusions, and omitted-tail envelope-controlled domain.

Definition 2.9 (Arithmetic ranking coordinate). *The archived ranking argument uses a normalized continued-fraction coordinate*

$$\eta(\rho) = \rho_{\text{norm}} \in (0, 1)$$

for the canonical representative of the arithmetic class in the comparison domain. In this archived standard-forcing replay,

$$\eta(\rho_G) = 0.6180339887498949 \dots,$$

while the representative near-top nongolden classes have validated intervals

$$\eta(\text{silver}) \in [0.4142, 0.4143], \quad \eta(\text{bronze}) \in [0.3027, 0.3028].$$

This η -coordinate is part of the theorem data. Changing it changes the generated comparison problem and would require regenerating the ranking argument and downstream comparison tables. The connection with continued-fraction and Lagrange/Markoff ordering is supplied by the exact near-top ranking argument, not by an informal comparison of decimal representatives [9, 10].

2.3 Relation between the code-level ranking coordinate and arithmetic ordering

The ranking argument does not infer near-top arithmetic order from the decimal size of ρ_{norm} alone. The coordinate $\eta = \rho_{\text{norm}}$ is the stored representative coordinate used to index classes in the archived replay. The arithmetic threat ordering is proved separately by an exact continued-fraction ranking argument, which partitions the relevant continued-fraction cylinders and assigns each cylinder to a closed control mechanism. Thus intervals such as $[0.4142, 0.4143]$ and $[0.3027, 0.3028]$ should be read as identifiers/enclosures in the archived coordinate system; the near-top claim itself comes from the ranking proof, not from naming conventions.

Class	Continued-fraction pattern	η interval	Control mechanism
Golden	$[1; 1, 1, \dots]$	$0.61803\dots$	equality class / reference anchor
Silver	principal near-top cylinder	$[0.4142, 0.4143]$	screened, ranked, retired, terminal-excluded
Bronze	next nongolden near-top cylinder	$[0.3027, 0.3028]$	ranked, pruned/dominated, terminal-excluded

Definition 2.10 (Near-top window). *The near-top window is*

$$W_\eta = (\eta^*, \eta(\rho_G)].$$

In the compact release used for this manuscript, the omitted-tail part is vacuous in a precise, relative sense. The generated arithmetic domain consists of the golden representative, the displayed near-top records, the screened records, and the auxiliary records handled by pruning, routing by inherited domination estimates, and terminal candidates carrying explicit exclusion estimates. The omitted-tail set is the set-theoretic complement of those records inside the generated grammar. For release v1.0.1, that complement is empty, so $\mathcal{D}_{\text{omit}} = \emptyset$ as a component of $\mathcal{D}_{\text{cert}}$ and no numerical η^ is needed. This is not an assertion that all irrational rotation numbers outside $\mathcal{D}_{\text{cert}}$ have been covered; it is an assertion that the generated domain has no tail records left uncontrolled. In a release where the omitted-tail region is nonempty, η^* must be exported as an explicit interval and the omitted-tail part must provide a nonvacuous envelope ceiling $E(\eta^*)$.*

The grammar is deliberately modest, but it is not chosen after observing the final comparison. It is generated from a fixed arithmetic threat model: first isolate the golden reference class; next include the near-top continued-fraction cylinders predicted by the Markoff–Lagrange/continued-fraction ranking to be the most dangerous nongolden challengers; then add the finite screened records that survive the local comparison; finally route every remaining generated record by a pre-specified domination, pruning, terminal-exclusion, or omitted-tail rule. Once these rules are fixed, the proof either closes every generated record or fails. Thus, $\mathcal{D}_{\text{cert}}$ should be read as a principled finite comparison domain for testing the closed proof mechanism, not as the set of all classes that happened to be dominated.

2.4 Concrete theorem universe for the archived result

The definitions above are not meant to make $\mathcal{F}_{\text{cert}}$ or $\mathcal{D}_{\text{cert}}$ post-hoc. In the archived standard-forcing theorem, membership is determined by the generated family identifier, arithmetic partition, and support grammar before the final maximality conclusion is evaluated. The concrete replay universe used in the manuscript snapshot is:

Object	Archived replay value
Primary forcing	$f(\theta) = (2\pi)^{-1} \sin(2\pi\theta)$; family label standard-sine
Golden representative	$\rho_G = [1; 1, 1, \dots] = 0.6180339887498949\dots$ in the normalized arithmetic coordinate
Arithmetic coordinate	$\eta(\rho) = \rho_{\text{norm}}$ as exported by the archived ranking records
Golden arithmetic value	$\eta(\rho_G) = 0.6180339887498949\dots$
Near-top nongolden labels	silver, bronze
Silver η interval	[0.4142, 0.4143]
Bronze η interval	[0.3027, 0.3028]
Final comparison window	Identification overlap window [0.971634, 0.971638] in the compact downstream replay
Discharged branch witness	[0.9716352, 0.9716355]
Compressed-V target interval	[0.9716350, 0.9716370]
Theorem-VII near-top upper ceiling	0.9716347
Omitted-tail conclusion	empty omitted-tail complement inside the generated grammar
Final domain grammar	union of screened labels, exact ranking records, theorem-pruned regions, inherited-domination labels, terminal candidates with exclusion estimates, and omitted-tail records
Compact grammar census	2 screened, 2 ranked, 1 pruned, 1 inherited domination, 2 terminal-excluded, and 0 omitted-tail records

For $\mathcal{D}_{\text{cert}}$, the domain is the explicitly generated union of the continued-fraction cylinders, screened labels, pruned regions, inherited-domination labels, terminal candidates with exclusion estimates, and omitted-tail records in the archived support object. The final theorem is evaluated only after this partition has been constructed and checked.

2.5 Archived statement of the theorem universe

The public proof package mirrors the mathematical universe in an archived theorem-universe record [35]. This record is not a replacement for mathematical definitions; it is a scope-locking device. It records the map family, arithmetic coordinate, partition grammar, threshold-functional convention, branch convention, normalization convention, and replay levels. The final reduction verifies that the theorem statement agrees with the archived theorem universe.

Manuscript object	Archived ledger field	Mathematical role
$\mathcal{F}_{\text{cert}}$	map_family, forcing_label, parameter_normalization	Defines the admissible conservative family.
$\mathcal{D}_{\text{cert}}$	arithmetic_domain, partition_grammar	Defines the generated arithmetic domain.
ρ_G	golden_representative	Identifies the reference equality/top class.
η	arithmetic_coordinate	Gives the coordinate used by Theorem VII ranking.
Λ_f	threshold_functional	Records the lower/upper/transport/identification convention.
$\text{GL}(2, \mathbb{Z})$ normalization	normalization	States the equality/uniqueness convention.

Replay scope	replay_levels	Separates minimal, downstream, and full replay.
--------------	---------------	-------------------------------------------------

2.6 Pre-conclusion construction of the generated arithmetic domain

The generated arithmetic domain is constructed before the final maximality conclusion is evaluated. The generator takes as input a continued-fraction normalization grammar, the golden reference class, the near-top symbolic labels determined by exact arithmetic ranking, the records surviving the screened comparison, and the envelope/tail estimate. It then produces a finite symbolic partition together with a tail-control record. Each generated object is assigned to one of six mathematical routes: screened-panel control, exact near-top ranking, theorem-level pruning, inherited domination estimates, explicit terminal-exclusion estimates, or omitted-tail envelope control. The final theorem uses this generated partition; it does not enlarge, shrink, or relabel it after the golden comparison is known. This is the pre-success rule that gives the compact domain its mathematical meaning: the domain is small, but it is generated by arithmetic and comparison rules fixed before the final sign check.

Concretely, the generator follows a fixed finite recipe. First it fixes the continued-fraction normalization and the golden representative. Second it inserts the principal near-top nongolden cylinders selected by the exact continued-fraction/Markoff–Lagrange ranking calculation. Third it adds the finite records surviving the screened comparison stage before the final endpoint comparison is evaluated. Fourth it assigns every remaining generated record to one of the predeclared domination, pruning, terminal-exclusion, or omitted-tail routes. The generated domain is therefore not the set of examples that happened to be dominated after the fact; it is the finite set produced by this predeclared arithmetic threat model.

Domain component	Pre-conclusion generation rule	Mathematical control
$\mathcal{P}_{\text{screen}}$	finite screened labels exported by the VI finite comparison output	direct screened completion or assignment to another VII mechanism
$\mathcal{D}_{\text{rank}}$	continued-fraction near-top records generated by the ranking grammar	exact ranking plus closed near-top control
$\mathcal{D}_{\text{prune}}$	symbolic subregions whose upper envelope is already below the golden lower anchor	pruning inequality with positive margin
$\mathcal{D}_{\text{life}}$	generated records previously shown to be dominated	an explicit earlier domination estimate; not mere deletion
$\mathcal{D}_{\text{term}}$	terminal search candidates retained by the grammar	conversion to an explicit exclusion estimate
$\mathcal{D}_{\text{omit}}$	generated omitted/tail records outside the finite near-top panel	vacuous control in the archived replay, or envelope ceiling in nonvacuous releases

Proposition 2.11 (Pre-conclusion domain construction). *The arithmetic-domain generator constructs the finite symbolic partition and tail-control records before using any final golden maximality conclusion. The final replay checks the generated partition against the archived theorem universe and the VII support certificate, and the final reduction is rejected if the theorem statement refers to a family, representative grammar, or arithmetic class not present in this generated domain.*

Proof. The domain generator records labels and routing categories as fields of the VII support object and the archived theorem-universe record. The downstream validation checks those fields before the final implication and requires the failure lists for unranked, unpruned, missing, uncontrolled, and unpromoted labels to be empty. Since the final theorem uses the already generated partition and

has no operation that adds new labels, the generated domain is a pre-conclusion input to the final reduction rather than a post-hoc set of successful comparisons. \square

Published archive. The fixed theorem universe is frozen by the public release `v1.0.1`, archived as DOI 10.5281/zenodo.20101820 and cited throughout as the proof package [35]. The release records the theorem-facing upstream cache through Theorem IV and documents the verification and reconstruction procedures for the downstream layers.

Theorem-universe lock. Changing any of the family, domain, branch, threshold, or normalization conventions in the archived theorem universe requires regenerating the relevant proof components and the final replay. The manuscript theorem should not be broadened without a corresponding theorem-universe update.

2.7 Threshold functional

Let $\mathcal{I}_\rho(K, f)$ denote the assertion that the map $F_{K,f}$ admits the invariant object of rotation class ρ in the sense required by the proof program.

Definition 2.12 (Validated threshold interval). *A validated threshold interval for class ρ is an interval*

$$I_f(\rho) = [K_f^-(\rho), K_f^+(\rho)]$$

such that the lower theorem certifies persistence below or through $K_f^-(\rho)$, the upper theorem certifies obstruction above or through $K_f^+(\rho)$, and the transport-identification stack certifies that these endpoints refer to the same threshold branch.

Definition 2.13 (Threshold functional). *For $F_{K,f} \in \mathcal{F}_{\text{cert}}$ and $\rho \in \mathcal{D}_{\text{cert}}$, the threshold functional*

$$\Lambda_f(\rho)$$

is the threshold value identified by the lower-side persistence theorem, upper obstruction theorem, compressed transport contract, and threshold-identification theorem. At the archived-certificate level, $\Lambda_f(\rho)$ may be represented by a validated interval; theorem comparisons are interpreted by interval separation or by theorem-level envelope/pruning/ranking proofs.

Definition 2.14 (Strict validated comparison). *A strict validated comparison*

$$\Lambda_f(\rho) < \Lambda_f(\rho_G)$$

means that a proof component proves a positive separation between an upper ceiling for ρ and a lower anchor for ρ_G . In interval form this is typically

$$K_f^+(\rho) < K_f^-(\rho_G).$$

Definition 2.15 (Generated-domain golden maximality). *The golden class is maximal on the generated domain $\mathcal{D}_{\text{cert}}$ if*

$$\Lambda_f(\rho) \leq \Lambda_f(\rho_G) \quad \text{for all } \rho \in \mathcal{D}_{\text{cert}}.$$

It is uniquely maximal modulo the validated normalization if equality can occur only when $\rho \in \mathcal{O}_G \cap \mathcal{D}_{\text{norm}}$. The adjective “maximal” in this paper always refers to this generated-domain comparison unless a broader domain is explicitly specified.

2.8 Notation checkpoint

The symbols used below have fixed meanings throughout the paper. The theorem universe $\mathfrak{U}_{\text{cert}}$ fixes the map family, arithmetic coordinate, branch convention, threshold convention, normalization, and replay scope. The set $\mathcal{D}_{\text{cert}}$ is the generated arithmetic comparison domain, not the continuum of irrational rotation numbers. The value $K_G^- = 0.9716350$ is the direct lower persistence anchor for the golden branch, while $U_{\text{ng}}^+ = 0.9716347$ is the largest theorem-facing generated nongolden upper ceiling. The final endpoint margin $\Delta_{\text{final}} = K_G^- - U_{\text{ng}}^+ = 3.0 \times 10^{-7}$ is distinct from both the transport slack and the branch-coordinate obstruction gap recorded later.

3 Mathematical validation operators

This section spells out the mathematical operators underlying the validation language used throughout the paper. The software implementation records these operators as interval objects, Krawczyk inclusions, residue certificates, and obstruction fields; the definitions below are the corresponding continuum-level statements. The purpose is to make clear that a completion check such as `strict_krawczyk` or `analytic_incompatibility_validated` is accepted only after the associated operator inequality has been checked with outward-rounded intervals.

3.1 From mathematical estimates to implementation

The operators in this section are the mathematical ones used by the proof. The repository implements them in separate modules for periodic-orbit validation, interval Newton steps, lower invariant-graph validation, upper obstruction assembly, transport, branch identification, arithmetic discharge, and final reduction. The main text does not depend on a particular module name. What matters mathematically is that each numerical claim is reduced to one of the interval or symbolic checks described below. Artifact-level paths and replay commands are recorded in Supplementary Material S1.

3.2 Periodic-orbit residual and Krawczyk inclusion

Fix a rational approximant p/q to a rotation class. A lifted q -periodic orbit is represented by a vector

$$x = (x_0, \dots, x_{q-1}) \in \mathbb{R}^q, \quad x_{i+q} = x_i + p.$$

For the standard-sine forcing $f(\theta) = (2\pi)^{-1} \sin(2\pi\theta)$, the second-order form of the map gives the residual operator

$$F_{p/q,K}(x)_i = x_{i+1} - 2x_i + x_{i-1} - Kf(x_i), \quad i = 0, \dots, q-1, \quad (4)$$

where the boundary convention is $x_q = x_0 + p$ and $x_{-1} = x_{q-1} - p$. Thus a zero of $F_{p/q,K}$ is precisely a lifted p/q orbit of the standard-sine map. Its Jacobian is the cyclic tridiagonal matrix

$$(D_x F_{p/q,K}(x))_{ij} = (-2 - Kf'(x_i))\delta_{ij} + \delta_{i+1,j} + \delta_{i-1,j}, \quad (5)$$

again with the same cyclic lift convention.

Let $X = X_0 + [-r, r]^q$ be an interval box around an approximate orbit X_0 , and let C be a numerical approximate inverse of $D_x F_{p/q,K}(X_0)$. The Krawczyk operator used for the periodic-orbit validation is

$$\mathcal{K}_{p/q,K}(X; X_0, C) = X_0 - CF_{p/q,K}(X_0) + (I - CD_x F_{p/q,K}(X))(X - X_0). \quad (6)$$

All products and evaluations in (6) are interval operations. The verification accepts the periodic orbit as uniquely enclosed when

$$\mathcal{K}_{p/q,K}(X; X_0, C) \subset \text{int}(X). \quad (7)$$

By the standard interval Newton–Krawczyk theorem [14, 15, 16], (7) implies that $F_{p/q,K}$ has a unique zero in X . The repository fields called `residual_inf`, `krawczyk_margin`, `unique`, and `strict` are numerical witnesses for this inclusion; they are not used as theorem evidence unless the interval inclusion itself is accepted.

3.3 Residue, crossing windows, and interval Newton validation

Given a validated orbit enclosure X , the derivative of one standard-sine map step at phase point (θ, r) is

$$DF_K(\theta, r) = \begin{pmatrix} 1 + K f'(\theta) & 1 \\ K f'(\theta) & 1 \end{pmatrix}. \quad (8)$$

The monodromy matrix for the p/q orbit is the interval product

$$M_{p/q}(K; X) = \prod_{i=0}^{q-1} DF_K(\theta_i, r_i), \quad (9)$$

where the product is ordered along the validated orbit. Greene’s residue is then

$$R_{p/q}(K) = \frac{2 - \text{tr } M_{p/q}(K)}{4}. \quad (10)$$

The crossing certificates in Theorem IV localize a zero of

$$g_{p/q}(K) = |R_{p/q}(K)| - \frac{1}{4}. \quad (11)$$

On an interval I_K around a candidate crossing, the interval Newton step has the schematic form

$$N(K_0, I_K) = K_0 - \frac{g_{p/q}(K_0)}{g'_{p/q}(I_K)}, \quad (12)$$

where $g'_{p/q}(I_K)$ includes the derivative of the validated orbit branch with respect to K . In the implementation this derivative is obtained from the implicit equation

$$D_x F_{p/q,K}(x(K)) \frac{dx}{dK} = -\partial_K F_{p/q,K}(x(K)). \quad (13)$$

A crossing window is proof only if the interval branch enclosure, the residue interval, and the derivative/transversality checks jointly prove a unique transverse crossing of (11). The displayed crossing intervals in the obstruction figures are therefore interval-validated crossing windows, not midpoint estimates.

3.4 Lower invariant-graph validation

The lower-side validation validates an invariant graph for an irrational rotation class through a Fourier/collocation Newton scheme, in the a posteriori KAM/parameterization-method tradition [11, 30, 31, 13, 21]. In schematic form, the unknown is a periodic graph correction $h : \mathbb{T} \rightarrow \mathbb{R}$ and the invariance equation is written as

$$\Phi_{K,\rho}(h)(t) = h(t + \rho) - 2h(t) + h(t - \rho) - Kf(t + h(t)) = 0, \quad (14)$$

up to the chart convention used by the lower-side parameterization implementation. Let h_N be the finite Fourier approximation, B an approximate inverse for the linearized operator $D\Phi_{K,\rho}(h_N)$ on the resolved modes, and let $r > 0$ be a validation radius in the analytic norm used by the proof argument. The lower argument checks a radii-polynomial/Krawczyk inequality of the form

$$Y + Zr + T_N(r) < r, \quad Y = \|B\Phi_{K,\rho}(h_N)\|, \quad Z = \|I - BD\Phi_{K,\rho}(h_N)\|, \quad (15)$$

where $T_N(r)$ is an outward-rounded tail and nonlinear-remainder majorant. The small-divisor part of the proof argument supplies lower bounds for the Fourier multipliers associated with the rotation ρ . By the standard interval/Krawczyk and radii-polynomial logic [14, 15, 16], inequality (15) implies that the Newton-like graph transform maps the analytic ball $\|h - h_N\| \leq r$ strictly into itself and is contractive there; hence a true invariant graph exists in that ball. This is the mathematical content behind the lower-anchor estimates reported in Theorem III.

3.5 Upper obstruction and analytic-incompatibility operator

The upper theorem uses validated rational information to prove a branch-level obstruction, with Greene-type residue information used only after interval validation and branch/tail compatibility checks [4, 5, 14, 15, 16]. For a supported rational approximant p/q , let

$$I_{p/q}^{\text{cross}} = [K_{p/q}^-, K_{p/q}^+] \quad \text{and} \quad B_{p/q}^{\text{hyp}} = [H_{p/q}^-, H_{p/q}^+]$$

be the interval-validated residue crossing window and a validated hyperbolic band. The obstruction package verifies

$$K_{p/q}^+ < H_{p/q}^-, \quad |R_{p/q}(K)| > \frac{1}{4} \quad \text{for all } K \in B_{p/q}^{\text{hyp}}, \quad (16)$$

with branch labels, common-support identifiers, and tail-coherence fields fixed across the finite-to-tail passage. The analytic-incompatibility operator is the implication

$$\text{Core}_{p/q} \wedge \text{Tail}_{p/q} \wedge \text{HypBand}_{p/q}(B_{p/q}^{\text{hyp}}) \implies \neg \text{Persist}_\rho(K) \quad (K \geq U_\rho), \quad (17)$$

where U_ρ is the upper-front ceiling recorded by the obstruction proof component. In words: an invariant object on the identified irrational branch cannot coexist, beyond the upper ceiling, with the tail-coherent hyperbolic obstruction carried by the rational approximant ladder. The verification accepts (17) only when the finite crossing data, hyperbolic band, common support, tail coherence, and branch-identification fields all agree. A residue plot or an isolated crossing does not by itself constitute an upper theorem.

3.6 Interval arithmetic convention

All interval comparisons use outward-rounded endpoints [14, 16]. If an interval-valued quantity A is stored as $[a^-, a^+]$ and a lower anchor L is stored as $[\ell^-, \ell^+]$, a strict comparison is accepted only when

$$a^+ < \ell^-. \quad (18)$$

Displayed decimal values are therefore summaries of stored interval inequalities. The final comparison in the main theorem uses (18), not midpoint comparisons or visual separation.

4 Main theorem

The theorem is stated in two equivalent languages. The first is the invariant-circle language: the golden class has the largest certified threshold among the classes in the generated comparison domain. The second is the ledger language used by the replay: the golden branch has a lower persistence endpoint that is strictly larger than every nongolden upper ceiling produced by the discharge proof. The two languages are connected by the threshold-identification and final-reduction components.

For clarity, let

$$K_G^- := 0.9716350, \quad U_{\text{ng}}^+ := 0.9716347, \quad (19)$$

where K_G^- is the theorem-facing lower endpoint for the golden branch and U_{ng}^+ is the theorem-facing upper ceiling for the largest generated nongolden discharge record. These are not empirical fit parameters. They are outward-rounded endpoints consumed by the final reduction. The decisive sign check is

$$\Delta_{\text{final}} := K_G^- - U_{\text{ng}}^+ = 3.0 \times 10^{-7} > 0. \quad (20)$$

Main Theorem 4.1 (Validated golden threshold maximality in the generated standard-map universe). *Let $\mathfrak{U}_{\text{cert}}$ be the fixed theorem universe specified by the archived proof ledger. Let $\mathcal{F}_{\text{cert}}$ be the validated standard-sine conservative twist family, $\mathcal{D}_{\text{cert}}$ the generated arithmetic comparison domain, and Λ_f the threshold functional obtained by the lower-side persistence, upper-obstruction, transport, branch-identification, discharge, and final-reduction components.*

For the validated golden representative ρ_G ,

$$\Lambda_f(\rho_G) \geq K_G^-. \quad (21)$$

For every generated nongolden class

$$\rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}}),$$

the final discharge-and-reduction proof assigns a certified upper ceiling U_ρ satisfying

$$\Lambda_f(\rho) \leq U_\rho \leq U_{\text{ng}}^+ < K_G^-. \quad (22)$$

Consequently

$$\Lambda_f(\rho) < \Lambda_f(\rho_G) \quad \text{for every } \rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}}),$$

and therefore

$$\Lambda_f(\rho) \leq \Lambda_f(\rho_G) \quad \text{for every } \rho \in \mathcal{D}_{\text{cert}}.$$

Moreover, equality is possible only for the validated golden $\text{GL}(2, \mathbb{Z})$ -orbit representative in the normalization domain:

$$\Lambda_f(\rho) = \Lambda_f(\rho_G) \implies \rho \in \mathcal{O}_G \cap \mathcal{D}_{\text{norm}}.$$

Thus the golden class is the unique maximizing representative inside the generated, normalized theorem universe.

Equivalent reading. The proof does not say that every irrational rotation number in the continuum has been compared. It says that, after the standard-sine map, arithmetic-generation grammar, threshold branch, and $\text{GL}(2, \mathbb{Z})$ representative convention have been fixed, every generated nongolden record has a certified upper threshold ceiling below the certified golden lower endpoint. The entire final comparison is the strict interval inequality (20) together with the finite discharge theorem that assigns every generated nongolden class to a closed upper-bound route.

Logical separation. The theorem deliberately separates four claims that are easy to conflate. The lower-side statement is existence of the golden branch through the stored lower anchor. The upper-side statement is nonexistence for generated nongolden records only after the obstruction, transport, and identification fields agree. The discharge statement is discharge of the generated grammar, not of all irrational rotation numbers. The final statement is an endpoint comparison between two already certified quantities, not a new numerical experiment.

Remark 4.2 (Claim-boundary summary). *The comparison range and nonclaims are fixed in Section ???. The terms “maximal” and “unique” always refer to the generated comparison domain and the validated normalization convention in the archived theorem universe.*

Theorem 4.3 (Replay theorem). *Starting from the archived release, the minimal replay reproduces the final theorem ledger, and the downstream replay regenerates Theorem V through the final theorem from the archived heavy upstream data. Both replays return the same closed final-discharge fields used by Theorem 4.1: no active assumptions, no open hypotheses, no remaining true mathematical burden, and no failed checks. A separate construction-audit command can rebuild the theorem artifacts before rerunning the proof-facing checks and writing a run-specific manifest. The exact replay protocol, artifact identifiers, and expected acceptance fields are documented in Supplementary Material S1.*

Remark 4.4 (Mathematical theorem versus reproducibility labels). *Theorem 4.1 is the mathematical theorem, while Theorem 4.3 is the reproducibility statement. Reproducibility status labels are not proof axioms; they are accepted only when the interval inequalities, routing records, empty failure lists, and provenance checks required by the mathematical lemmas have passed.*

Remark 4.5 (Archive relativity). *Theorem 4.3 is conditional on the archived source, numerical data, and replay protocol documented in Supplementary Material S1. It is not a claim that arbitrary unarchived code proves the theorem merely by printing similar words.*

5 How the computation enters the proof

The proof is computer-assisted, but the role of the computer is limited and explicit. Each numerical stage produces a finite mathematical object: an interval enclosure, an a posteriori existence inequality, an obstruction interval, a transport bound, or a finite assignment of arithmetic cases. The paper uses these objects only through implications of the following form:

$$\text{checked finite data} \implies \text{mathematical statement.}$$

This is the usual logic of rigorous computation in dynamics and validated numerics [14, 15, 12, 11, 13, 16].

For example, a lower existence computation is not used because a plotted invariant curve looks convincing. It is used because the residual, inverse bound, small-divisor bound, nonlinear remainder, and tail estimate satisfy a radii-polynomial inequality. Similarly, a nongolden challenger is not

discarded because a finite search did not find a larger threshold. It is discarded only when the generated arithmetic record has been assigned to a closed comparison mechanism with a positive interval margin.

The repository supplies the reproducibility layer behind these finite objects: archived data, checksums, replay scripts, and negative-control tests. Those details are essential for independent checking, but they are separated from the mathematical exposition. The main text states the mathematical estimates and implications; Supplementary Material S1 records the artifact-level replay protocol and failure tests used in the public archive [35].

A strict comparison in this paper always means an interval comparison. If a challenger has an upper enclosure $U_\rho = [u^-, u^+]$ and the golden class has a lower anchor $L_G = [\ell^-, \ell^+]$ or a scalar lower endpoint K_G , then the comparison is accepted only when

$$u^+ < \ell^- \quad \text{or} \quad u^+ < K_G.$$

Midpoint comparisons and visual separations are not used as proof steps.

6 Proof architecture

The proof dependency graph is:

$$\text{I-II} \prec \{\text{III, IV}\} \prec \text{compressed V} \prec \text{identification} \prec \text{VI} \prec \text{VII} \prec \text{VIII} \prec \text{final replay}.$$

The proof separates local validation, proof finite packages, transport and identification, envelope control, discharge of the generated challenger domain, final reduction, and replay check.

Equivalently, the theorem pipeline is:

$$\boxed{\text{I-II}} \longrightarrow \boxed{\text{III, IV}} \longrightarrow \boxed{\text{V}_{\text{comp}}} \longrightarrow \boxed{\text{ID}} \longrightarrow \boxed{\text{VI}} \longrightarrow \boxed{\text{VII}} \longrightarrow \boxed{\text{VIII}} \longrightarrow \boxed{\text{Final Replay}}.$$

The Roman labels I–VIII are archival stage labels inherited from the proof package, not the numbering of theorems in this article. In the paper, setup layers I–II fix the theorem universe; the later stage labels identify the lower anchor, upper obstruction, transport, identification, screened comparison, generated-domain discharge, and final reduction components consumed by the main theorem.

Layer	Checker invariant	Mathematical implication
III	the direct-anchor residual, small-divisor, cohomology-inverse, frame, nonlinear, tail, branch/chart, and graph-consumption checks pass at the direct anchor	analytic lower-anchor persistence for the golden branch
IV	interval obstruction, common-support, tail-coherence, and analytic-incompatibility checks hold	upper nonexistence/front ceiling
V	transport majorants are smaller than the preserved comparison gap	rational finite data transport to the irrational branch
ID	branch labels, chart identifiers, overlap window, and native-tail witness agree	lower/upper/transport objects identify one threshold branch
VI	screened top-gap and envelope estimate fields are valid	local/screened golden comparison estimate

VII	all six generated-domain-discharge failure fields are empty	all nongolden challengers are controlled
VIII	final implication, normalization, and universe-matching checks pass	the validated comparison is exactly the main theorem

Architecture Principle 6.1 (Separation of theorem roles). *Each layer has one primary theorem role. Lower persistence, upper obstruction, transport, identification, envelope control, challenger discharge, and final reduction are not interchangeable. This prevents local numerical evidence from being mistaken for generated-domain arithmetic discharge.*

Proposition 6.2 (Acyclic dependency structure). *The proof dependency graph is acyclic. No theorem layer uses a downstream theorem object as an input to its own proof.*

Proof. Each theorem object is used only by later layers in the order listed above. The only subtle relationship is between VI and VII. VI exports a finite screened comparison output, while VII proves the discharge theorem using that finite output. VI does not use the conclusion of VII to prove its local screened statement. Hence there is no cycle. \square

Final assembly in ledger form

The final implication can be read as a three-line ledger. Theorem III gives the golden lower endpoint

$$K_G^- = 0.9716350 \leq \Lambda_f(\rho_G). \quad (23)$$

Theorem VII, after the upper-obstruction, transport, identification, and screened-envelope inputs have been consumed, gives for every generated nongolden class a ceiling

$$\Lambda_f(\rho) \leq U_\rho \leq U_{\text{ng}}^+ = 0.9716347. \quad (24)$$

The final reduction uses the outward-rounded endpoint comparison

$$U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^-. \quad (25)$$

Thus every generated nongolden class is strictly below the golden lower certificate before the equality and normalization clauses are applied. This final endpoint comparison is separate from the transport-gap inequality and from the branch-coordinate obstruction gap in Theorem IV.

Dependency audit table

Layer	Burden discharged	Primary dependency	Main failure mode guarded against
I–II	theorem universe, normalization, and shared workstream conventions	certified-universe and setup objects	later theorem refers to a different family or convention
III	lower golden persistence at $K_G^- = 0.9716350$	direct-anchor a posteriori certificate	confusing a local diagnostic row with the near-critical lower anchor
IV	branch-scoped upper obstruction and ceiling data	obstruction atlas, support, tail, and incompatibility ledgers	treating a residue plot as an upper theorem

V	rational-to-irrational transport with positive margin	compressed transport contract	spending more gap than the comparison has available
Identification	common threshold branch for lower, upper, and transported objects	overlap window, branch labels, and native-tail witness	comparing incompatible branches
VI	screened local comparison and envelope estimate	screened challenger outputs and top-gap ledger	missing a near-top nongolden challenger
VII	generated-domain discharge	six-route grammar and empty failure fields	replacing generated-domain closure by a hand-picked finite list
VIII	final theorem reduction and normalization	universe match, endpoint comparison, and equality convention	proving a statement broader than the certified universe
Replay/recon.	audit of accepted proof objects and regenerated artifacts	archived proof ledger, replay protocol, reconstruction procedure, and negative controls	accepting status labels without checking proof fields

7 Numerical anchors and reproducibility snapshot

This section records the theorem-facing numerical ledgers. Its purpose is to prevent three different quantities from being conflated: the final endpoint gap, the transport-preservation budget, and the Theorem-IV branch-coordinate obstruction gap. All three are positive inequalities, but they enter the proof at different layers.

Endpoint ledger: the final comparison

The final theorem uses the lower golden endpoint

$$K_G^- = 0.9716350 \tag{26}$$

and the largest generated nongolden discharge ceiling

$$U_{\text{ng}}^+ = 0.9716347. \tag{27}$$

The final reduction consumes the strict interval sign check

$$\Delta_{\text{final}} = K_G^- - U_{\text{ng}}^+ = 3.0 \times 10^{-7} > 0. \tag{28}$$

This is the narrowest theorem-facing comparison. It is therefore made at the level of stored interval endpoints and certified support fields, not by midpoint comparison, graphical separation, or informal decimal rounding.

Endpoint field	Displayed value	Logical role
Golden lower endpoint K_G^-	0.9716350	lower persistence certificate consumed by the final theorem
Generated nongolden upper ceiling U_{ng}^+	0.9716347	largest upper ceiling after generated-domain discharge
Final endpoint gap $K_G^- - U_{\text{ng}}^+$	3.0×10^{-7}	final sign check proving strict nongolden domination

Transport and identification ledger

The transport step has its own positive inequality. It preserves a finite comparison through rational-to-irrational passage, branch alignment, tail control, and rounding charges. In the compact theorem-facing ledger, the charged transport majorant is

$$\delta_{\text{tot}} = 1.800001 \times 10^{-6},$$

while the available comparison budget at that layer is

$$\Delta_{\text{transport}} = 1.0 \times 10^{-5}.$$

Thus

$$1.0 \times 10^{-5} - 1.800001 \times 10^{-6} = 8.199999 \times 10^{-6} > 0. \quad (29)$$

This inequality is not the final theorem margin. It says that the transport and branch-identification procedure does not destroy the separation that the downstream layers require. The two numbers often compared by eye, 1.800001×10^{-6} and 3.0×10^{-7} , therefore bound different objects: the former is a pre-final transport charge inside a 10^{-5} -scale budget, while the latter is the final comparison between two already certified theorem-facing endpoints.

The identification ledger also records a native-tail witness interval

$$[0.9716352, 0.9716355]$$

and an overlap window

$$[0.971634, 0.971638].$$

The witness has width 3.0×10^{-7} and sits inside the overlap window used to align the lower, upper, and transported branches. The local screened shell also records a current challenger upper value 0.9716348; compared to the witness lower endpoint 0.9716352, this gives a separate witness-level margin of approximately 4.0×10^{-7} . The theorem-facing final endpoint comparison, however, remains (28).

Quantity	Displayed value	Role
Transport charge δ_{tot}	1.800001×10^{-6}	rational/branch/tail/rounding charge in Theorem V
Available transport budget	1.0×10^{-5}	gap scale reserved for transport and branch effects
Unused transport budget	8.199999×10^{-6}	positive slack in the transport-preservation inequality
Native-tail witness interval	$[0.9716352, 0.9716355]$	branch-identification witness
Witness width	3.0×10^{-7}	geometric width of the branch witness, not the transport error
Witness-width/top-gap slack	9.7×10^{-6}	difference between 10^{-5} and the witness width
Local screened challenger upper	0.9716348	intermediate screened-shell comparison value

Upper-obstruction ledger

Theorem IV contains a larger positive branch-coordinate incompatibility gap. This number lives in the obstruction chart; it is not the final endpoint margin. In the selected strict bridge profile, the theorem-facing obstruction ledger records:

Theorem-IV obstruction field	Displayed value	Role
Selected bridge profile	strict	theorem-facing upper-obstruction profile
Certified upper interval	[0.970272906, 0.9717820875]	upper branch interval in the obstruction chart
Certified barrier interval	[1.0366258098, 1.0367050854]	hyperbolic barrier interval in the obstruction chart
Certified incompatibility gap	0.0648437222	local branch-coordinate separation
Gap/localization ratio	42.9661518	robustness ratio for the obstruction localization
Certified tail begins	$q = 144$	start of certified upper-tail regime
Tail denominator support	[144, 233]	rational support range in the theorem-facing upper package
Supporting entry count	5	number of supporting ledger entries in the selected obstruction profile
Missing hypotheses	empty	no unresolved obstruction-side hypotheses consumed downstream

The obstruction gap is deliberately much larger than the final endpoint gap because it measures a different object. The value 0.0648437222 is a local order-separation margin inside a validated branch chart; it proves that a branch cannot persist beyond the certified upper front. The final endpoint gap 3.0×10^{-7} compares the largest generated nongolden upper front with the golden lower persistence anchor. The obstruction margin is therefore not “spent down” into the final endpoint gap. Rather, it certifies the validity of an upper ceiling, and the closeness of that upper ceiling to the golden lower anchor reflects the near-critical nature of the comparison. The final theorem is obtained only after the obstruction result is transported, branch-identified, discharged over the generated domain, and reduced to the endpoint inequality (28).

Generated-domain census

The discharge step is finite after the comparison domain has been generated. In the compact archived instance, the generated nongolden grammar is routed through six closure mechanisms. The theorem is accepted only when no generated record remains uncontrolled.

Record type	Count	Meaning	Open?
Screened labels	2	silver and bronze labels exported by the screened comparison	no
Exact ranking records	2	silver and bronze near-top continued-fraction records	no
Pruned symbolic regions	1	dominated cylinder-level region with upper ceiling below the golden anchor	no

Inherited-domination records	1	deferred route with closed domination provenance	no
Terminal candidates with exclusion estimates	2	terminal candidates promoted to explicit proof routes	no
Omitted-tail records	0	empty complement inside the generated grammar	no

The zero omitted-tail count means that the complement inside the generated grammar is empty. It does not mean that all possible irrational rotation numbers outside the generated comparison domain have been compared.

The public archive was also tested against deliberately corrupted proof paths: replacing the compressed transport object by a diagnostic shell, making a discharge failure list nonempty, changing branch labels, eroding the comparison margin, or changing the theorem universe without regenerating the dependent data. These tests are reproducibility checks rather than additional mathematical assumptions; their role is to show that the replay is fail-closed [17, 18, 19, 20, 35].

8 Why the finite checks imply the theorem

The proof uses numerical data only after they have been converted into finite mathematical checks. This section records the logical implications used later in the theorem chain.

8.1 Interval enclosure soundness

All threshold comparisons are interval comparisons. If a challenger has upper enclosure $U_\rho = [u^-, u^+]$ and the golden class has lower anchor K_G , then strict domination means

$$u^+ < K_G. \quad (30)$$

The reported margin is

$$\Delta_{\rho,G} = K_G - u^+ > 0. \quad (31)$$

This rule is used for direct challenger comparison, pruning, screened-panel completion, omitted-tail control, and the final reduction. Midpoints and plotted separations never replace (30).

8.2 Branch-identification soundness

The lower existence result and the upper obstruction result must refer to the same threshold branch. The identification step checks compatibility of branch labels, local charts, overlap windows, and the native-tail witness interval. Only after these checks pass are the lower anchor, the upper front, and the transport object interpreted as statements about one threshold functional Λ_f .

8.3 Transport and gap preservation

The rational-to-irrational transport step carries finite rational data to the limiting irrational branch. Its error budget has the schematic form

$$\Delta_{\text{available}} > \delta_{\text{rat}} + \delta_{\text{branch}} + \delta_{\text{tail}} + \delta_{\text{round}}. \quad (32)$$

Here δ_{rat} controls the finite rational approximant error, δ_{branch} controls chart and branch alignment, δ_{tail} controls the high-denominator tail, and δ_{round} is the outward-rounding allowance. When (32) holds, the strict separation visible at the finite level remains strict after all four effects are charged against the available gap.




Figure 2: Acyclic theorem dependency graph and replay boundary. The graph records the formal dependency order: setup layers I–II fix the universe, Theorem III supplies the golden lower anchor, Theorem IV supplies the upper-obstruction front, and the downstream layers transport, identify, discharge, normalize, and replay the final theorem object. The downstream replay starts from archived heavy upstream certificates; the construction command rebuilds the theorem artifacts before rerunning the proof-facing checks. The proof is the sequence of finite implications and accepted proof objects represented by the arrows, not the diagram itself.

$$\begin{array}{c}
 U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^- \\
 \overbrace{\Delta_{\text{final}} = 3.0 \times 10^{-7} > 0} \\
 \begin{array}{ll}
 \text{transport inequality} : 1.0 \times 10^{-5} - 1.800001 \times 10^{-6} = 8.199999 \times 10^{-6} > 0 & \\
 \text{identification witness} : [0.9716352, 0.9716355] & \text{branch-local witness, not the final endpoint gap} \\
 \text{Theorem-IV obstruction gap} : 0.0648437222 & \text{branch-coordinate incompatibility, not the final endpoint gap}
 \end{array}
 \end{array}$$

Figure 3: Ledger geometry for the final comparison. The top line is the theorem-level endpoint inequality: the largest generated nongolden upper ceiling lies below the golden lower anchor by 3.0×10^{-7} . The lower lines show two different positive checks used earlier: the transport-preservation slack and the branch-coordinate obstruction gap. The figure keeps these numerical scales distinct; the proof uses the stored outward-rounded fields behind the displayed values.

8.4 Finite discharge of the generated arithmetic domain

Theorem VII is not a list of examples. It is a finite assignment map over the generated continued-fraction records. Each nongolden record is routed to at least one closed control mechanism: exact near-top ranking, pruning, screened-panel completion, routing by inherited domination estimates, terminal candidates carrying explicit exclusion estimates, or omitted-tail control. The discharge implication is valid only when no generated record remains unassigned. Thus search termination is never used as evidence by itself; it is used only after every terminal candidate has been routed to a mathematical exclusion mechanism.

8.5 Final reduction

The final reduction combines three statements. First, the identified threshold functional gives a golden lower endpoint K_G^- and a nongolden upper ceiling U_ρ for every generated nongolden class. Second, the generated-domain discharge theorem proves that the worst such ceiling is bounded by U_{ng}^+ . Third, the $\text{GL}(2, \mathbb{Z})$ normalization identifies the only equality class. The final theorem is then a direct consequence of the strict endpoint comparison

$$U_\rho \leq U_{\text{ng}}^+ = 0.9716347 < 0.9716350 = K_G^- \leq \Lambda_f(\rho_G). \quad (33)$$

Lemma 8.1 (Endpoint assembly). *Assume that Theorem III supplies $\Lambda_f(\rho_G) \geq K_G^-$, Theorem VII supplies $\Lambda_f(\rho) \leq U_\rho \leq U_{\text{ng}}^+$ for every $\rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}})$, and $U_{\text{ng}}^+ < K_G^-$. Then every generated nongolden class is strictly dominated by the golden class.*

Proof. For each generated nongolden ρ ,

$$\Lambda_f(\rho) \leq U_\rho \leq U_{\text{ng}}^+ < K_G^- \leq \Lambda_f(\rho_G),$$

so $\Lambda_f(\rho) < \Lambda_f(\rho_G)$. The proof uses only endpoint inequalities and does not require estimating the exact value of $\Lambda_f(\rho_G)$. \square

Theorem 8.2 (Finite-check soundness). *Assume that the lower existence estimate, upper obstruction estimate, transport budget, branch-identification checks, arithmetic-discharge assignment, and final normalization checks all hold with the displayed interval margins and empty unresolved case lists. Then the final theorem statement follows over the generated theorem universe $\mathfrak{U}_{\text{cert}}$.*

Proof. The proof follows the acyclic dependency graph in Section 6. Setup layers I–II fix the map family, normalization, and workstream data. Theorem III gives the golden lower anchor K_G^- . Theorem IV gives upper obstruction fronts. Theorem V transports rational information to the irrational threshold setting while preserving the strict gap required by the downstream layers. The threshold-identification step ensures that the lower, upper, and transported data refer to one threshold branch. Theorem VI supplies the screened top-gap estimate. Theorem VII discharges the generated nongolden domain and gives the upper ceiling U_{ng}^+ . Theorem 8.1 gives strict nongolden domination. Theorem VIII applies the $\text{GL}(2, \mathbb{Z})$ normalization and universe-matching checks. Since no layer uses a downstream conclusion to prove an upstream input, the chain is non-circular and proves the stated theorem. \square

9 Setup layers I–II: workstream package

Setup layers I–II provide the foundational workstream package. They fix the family, normalization conventions, critical-surface scaffolding, and renormalization-side objects used downstream.

Exported field	Primary consumers	Mathematical role
family_ label=standard-sine	III, IV, VIII	fixes the map family throughout the theorem chain
normalization convention	VII, VIII	fixes the arithmetic representative convention and equality case
critical-surface object	III, IV, ID	supplies the branch scaffold used by lower/upper threshold objects
renormalization/workstream package	V, VI, final replay	supplies the theorem-universe/workstream context used downstream
empty open-hypothesis fields	all downstream layers	certifies that the final proof path does not call unresolved workstream assumptions

Proposition 9.1 (Setup layers I–II: validated workstream package). *There exists a workstream proof component*

$$C_{\text{I--II}}$$

accepted by the final replay such that:

- (i) *the map family and normalization conventions are fixed;*
- (ii) *the renormalization/workstream and critical-surface objects needed downstream are exported;*
- (iii) *the proof argument is compatible with the lower-anchor persistence, upper obstruction, transport, identification, envelope, discharge, and final reduction layers;*

(iv) the final replay uses $C_{\text{I}--\text{II}}$ with no active assumptions and no open hypotheses relevant to the main theorem.

Verification. The certified workstream package constructs a proof component from the validated renormalization package, critical-surface object, and family-normalization data. The final replay checks that later theorem components accept these fields and that no unresolved workstream hypotheses are called. Therefore the package is closed for the final theorem. \square

10 Stage III (archived Theorem III): lower existence at the golden anchor

Theorem III supplies the lower side of the final comparison. The statement needed downstream is simple: the golden invariant circle exists at the parameter value

$$K_{\text{III}} = 0.9716350.$$

This is the lower anchor against which the nongolden upper ceilings are compared. The theorem does not require, and we do not claim here, that this computation alone proves a full interval of persistence around the anchor. Its role is to provide the exact lower endpoint needed by the strict comparison

$$0.9716347 < 0.9716350.$$

Theorem 10.1 (Theorem III: lower golden anchor). *For the standard-sine map in the theorem universe fixed above, there is an invariant graph with golden rotation number at*

$$K_{\text{III}} = 0.9716350.$$

Consequently, the lower threshold endpoint for the golden class satisfies

$$K_f^-(\rho_{\text{G}}) \geq K_{\text{III}}.$$

Proof. The proof is an application of the a posteriori parameterization theorem for analytic invariant circles [11, 28, 30, 31, 13, 21]. We use the weighted Fourier algebra

$$\|u\|_{\nu} = \sum_{k \in \mathbb{Z}} |u_k| \nu^{|k|}, \quad \nu = 1.001,$$

which corresponds to an analytic strip of width $(2\pi)^{-1} \log \nu$ in the angle variable. The use of narrow complex strips and high-resolution Fourier enclosures near the golden standard-map threshold follows a line of rigorous and semi-rigorous work on small-divisor estimates and golden invariant curves [28, 27, 21]. The finite calculation encloses the Fourier modes $|k| \leq M/2$ with $M = 8192$ and bounds the omitted tail by the same weighted norm. One seeks an embedding $W(\theta)$ satisfying

$$F_{K_{\text{III}}}(W(\theta)) = W(\theta + \rho_{\text{G}}), \tag{34}$$

where F_K is the standard-sine map and ρ_{G} is the golden rotation number. The computation constructs a high-resolution Fourier approximation W_M and then applies an a posteriori theorem: if the residual, the inverse/cohomology bounds, the nondegeneracy estimates, the nonlinear remainder, and the Fourier tail are all sufficiently small in the chosen analytic norm, then an exact invariant graph exists near W_M .

In the archived computation the validation radius is $r = 3 \times 10^{-5}$. Let Y denote the interval-enclosed residual/Newton defect, Z the defect in the approximate inverse and reducibility frame, and Q the quadratic nonlinear majorant after the finite Fourier enclosure and tail estimate have been included. Equivalently, the tail majorant $T_M(r)$ is absorbed into the displayed Qr^2 bound after outward rounding. The nondegeneracy/twist condition used by the a posteriori theorem is checked in the same weighted norm through the resolved cohomology inverse and branch-normalization estimates. The checked inequality is the radii-polynomial condition

$$Y + Zr + Qr^2 < r. \quad (35)$$

The conservative displayed bounds are

$$\begin{aligned} Y &\leq 7.06239340228929 \times 10^{-8}, \\ Z &\leq 0.22895403677540946, \\ Q &\leq 551.0006605725524, \\ Y + Zr + Qr^2 &\leq 7.435155631800474 \times 10^{-6}, \\ r - (Y + Zr + Qr^2) &\geq 2.2564834368199524 \times 10^{-5}. \end{aligned}$$

The last line gives a positive existence margin. These numbers are not fitted midpoints: the finite Fourier residual, approximate-inverse defect, nonlinear majorant, and tail contribution are evaluated as outward-rounded interval quantities and then inserted into the scalar radii-polynomial inequality (35). The artifact replay recomputes the displayed interval bounds from the stored Fourier/collocation data, while the text uses only the resulting inequalities. The validation radius $r = 3 \times 10^{-5}$ is a radius in the analytic function norm around the approximate parameterization, not an uncertainty interval in the parameter K . The value $K_G^- = 0.9716350$ is the largest lower anchor used by the final theorem in the archived certificate; lowering it would weaken the final endpoint comparison, while increasing it would require a separate a posteriori validation at the larger parameter.

The resolved small-divisor estimate for the golden rotation is

$$\min_{1 \leq |k| \leq 4095} |e^{2\pi i k \rho_G} - 1| \geq 0.001087432533666174,$$

and the corresponding resolved cohomology-inverse bound is

$$\|\mathcal{L}_{\rho_G}^{-1}\|_{\text{resolved}} \leq 919.5972798685693.$$

Together with the tail and nonlinear estimates, these inequalities put the computation in the hypotheses of the a posteriori invariant-circle theorem. Hence a true golden invariant graph exists at K_{III} . The interval-arithmetic and Krawczyk/radii-polynomial background is standard in validated numerics [14, 15, 16, 25, 21]; Supplementary Material S1 records the corresponding artifact objects and replay checks. \square

The important point for the final theorem is that the lower-side computation proves exactly what the reduction needs. A potential ambiguity would be to read a local validation far from the critical regime as the near-critical lower anchor. The final proof does not do this: it uses the direct lower anchor above and nothing stronger. The parameter-window and artifact-object details are recorded in Supplementary Material S1.

11 Stage IV (archived Theorem IV): upper nonexistence from rational obstructions

Theorem IV supplies the upper side of the threshold comparison. Its purpose is to convert rigorously enclosed rational-approximant data into an upper obstruction for the limiting branch used later by transport and threshold identification. The theorem is branch-scoped in the sense fixed in Section ??: it proves the finite obstruction statement needed for the standard-sine family, the selected branch charts, and the generated theorem universe used here [4, 22, 23, 24, 32].

The point of this section is to make the upper argument readable without opening the repository. The computation contributes finitely many mathematical objects: rational branch enclosures, residue/crossing certificates, hyperbolic barrier intervals, branch-chart compatibility checks, support-tail checks, and outward-rounded separation inequalities. The theorem uses these objects only through the obstruction implication stated below. A plotted crossing, a large residue, or an isolated hyperbolic periodic orbit is never accepted by itself as a nonexistence theorem.

11.1 What exactly is the obstruction?

For the upper theorem, persistence means existence of an invariant circle represented in the same branch chart and function space as the threshold functional. In that chart an invariant circle is represented by a parameterization h satisfying

$$\Phi_{K,\rho}(h) = 0, \quad (36)$$

with the analyticity, monotonicity, orientation, and normalization conventions fixed by the lower and identification steps. The upper obstruction is therefore not the qualitative statement that a nearby rational residue has crossed a visible threshold. The proof needs a contradiction inside the same branch chart as (36). This is the reason for carrying branch labels, common-support fields, and tail-coherence fields through the Theorem-IV artifacts.

Let p/q be a supported rational approximant to the limiting class ρ . In the branch chart $\mathcal{T}_{p/q}$, let $X_{p/q}(K)$ denote the interval tube enclosing the rational branch, and let $\mathcal{H}_{p/q}(K)$ denote the validated hyperbolic/barrier strip associated to that branch. The two objects are compared only after they have been expressed over the same angular base in the same transverse coordinate z . If the chart orientation is chosen so that the barrier lies above the candidate branch, and if

$$X_{p/q}(K)(\vartheta) = [x_{p/q}^-(\vartheta), x_{p/q}^+(\vartheta)], \quad \mathcal{H}_{p/q}(K)(\vartheta) = [h_{p/q}^-(\vartheta), h_{p/q}^+(\vartheta)],$$

then the signed order separation is

$$\text{sep}(X_{p/q}(K), \mathcal{H}_{p/q}(K)) = \inf_{\vartheta} (h_{p/q}^-(\vartheta) - x_{p/q}^+(\vartheta)), \quad (37)$$

with the sign reversed in the orientation-reversed chart. Thus $\text{sep} > 0$ means positive order separation throughout the common chart. It is not a Euclidean distance in the annulus and it is not, by itself, the final threshold margin. It is the branch-coordinate quantity that allows an interval obstruction to be promoted to a nonexistence statement.

The corrected obstruction slack is obtained by subtracting all allowances that could weaken the branch-coordinate separation:

$$\mathfrak{D}_{\rho,K}(h) := \inf_{p/q \in S_\rho} \left(\text{sep}(X_{p/q}(K), \mathcal{H}_{p/q}(K)) - E_{p/q}^{\text{chart}}(h) - E_{p/q}^{\text{tail}} - E_{p/q}^{\text{round}} \right). \quad (38)$$

Here \mathcal{S}_ρ is the finite support ladder for the branch. The term E^{chart} accounts for the width of the graph-to-branch chart comparison, E^{tail} accounts for passing from the supported rational ladder to the limiting irrational branch, and E^{round} is the outward-rounding allowance. A positive lower bound for (38) means that, after every recorded uncertainty has been charged, the barrier still lies strictly on the obstructing side of the branch tube.

For any invariant graph h satisfying (36) on the chosen monotone branch, the ordering convention gives the necessary persistence inequality

$$\mathfrak{D}_{\rho,K}(h) \leq 0. \quad (39)$$

The Theorem-IV computation proves the opposite inequality on the same branch tube:

$$\mathfrak{D}_{\rho,K}(h) \geq \gamma_\rho > 0, \quad K \geq U_\rho. \quad (40)$$

The contradiction between (39) and (40) is the upper nonexistence conclusion. The role of residues and hyperbolic periodic branches is to construct the barrier strip and prove the positive separation; the nonexistence theorem comes only after branch coherence and tail coherence place that barrier on the same branch as the candidate invariant graph.

11.2 Three local implications used by Theorem IV

The upper proof is easiest to read as three finite implications: branch-order separation, rational obstruction, and tail/branch coherence.

Lemma 11.1 (Branch-order separation). *Fix a monotone branch chart and suppose that the rational branch tube $X_{p/q}(K)$ and the barrier strip $\mathcal{H}_{p/q}(K)$ are expressed in that chart. If the corrected separation*

$$\text{sep}(X_{p/q}(K), \mathcal{H}_{p/q}(K)) - E_{p/q}^{\text{chart}} - E_{p/q}^{\text{tail}} - E_{p/q}^{\text{round}} =: \gamma_{p/q}(K)$$

is strictly positive, then no invariant graph satisfying (36) can remain in the corresponding branch tube at that parameter.

Proof. The point that needs justification is the order constraint imposed by persistence. We use two standard facts for exact monotone twist maps. First, the Birkhoff graph theorem and Aubry–Mather ordering imply that essential invariant circles and ordered periodic/minimizing branches are graphs with a consistent cyclic order over the angular coordinate [6, 7, 8]. Second, the converse-KAM obstruction framework of MacKay–Percival, and its rigorous implementations for standard-map-type systems, uses such ordered barriers to rule out invariant circles when a validated barrier separates the admissible branch in the same order coordinate [23, 22, 24]. The branch chart used here is chosen precisely to put the candidate invariant graph, the rational branch tube, and the obstruction strip into one transverse order coordinate. In this coordinate, persistence of the invariant graph in the identified branch tube implies that the obstruction strip cannot lie strictly on the forbidden side of the branch after the chart-width, tail, and rounding allowances have been charged; otherwise the graph would have to cross the ordered barrier. This is the necessary inequality (39). The strict positivity condition $\gamma_{p/q}(K) > 0$ is the interval-certified negation of that inequality in the same chart. The contradiction is therefore not a consequence of a residue value alone; it is the standard twist-map order obstruction, localized to the validated branch chart and made quantitative by the corrected separation estimate. \square

Lemma 11.2 (Rational obstruction ledger). *For the theorem-facing upper branch, the archived interval calculations produce a coherent rational obstruction ledger with no missing hypotheses. In the selected strict bridge profile, the ledger records the following outward-rounded quantities:*

<i>Ledger field</i>	<i>Certified value</i>
<i>Selected bridge profile</i>	<i>strict</i>
<i>Rotation class recorded in the upper object</i>	$\rho = 0.6180339887498949$
<i>Certified upper localization interval</i>	$[0.9702729060, 0.9717820875]$
<i>Certified hyperbolic/barrier interval</i>	$[1.0366258098, 1.0367050854]$
<i>Upper localization width</i>	$1.5091815186 \times 10^{-3}$
<i>Barrier width</i>	$7.9275611016 \times 10^{-5}$
<i>Certified branch-coordinate incompatibility gap</i>	$6.4843722239 \times 10^{-2}$
<i>Gap/localization-width ratio</i>	42.9661518127
<i>Supporting entry count</i>	5
<i>Support fraction floor</i>	1.0
<i>Entry coverage floor</i>	1.0
<i>Certified tail start</i>	$q = 144$
<i>Certified tail denominators</i>	$\{144, 233\}$
<i>Missing upper hypotheses</i>	\emptyset

Consequently the finite rational obstruction data give a positive corrected separation for the selected upper branch after the recorded localization, support, tail, and rounding allowances are charged.

Proof. The periodic-orbit component first isolates the rational branches by interval Newton/Krawczyk enclosures and evaluates the residue/crossing quantities with outward rounding. The obstruction assembly then records a common upper localization interval and a common hyperbolic/barrier interval over the same branch chart. The lower endpoint of the barrier interval lies strictly above the upper endpoint of the localization interval after the supported-neighborhood allowances have been applied; the recorded branch-coordinate gap is

$$1.0366258098 - 0.9717820875 \geq 6.4843722239 \times 10^{-2}.$$

The support fraction floor and entry coverage floor are both 1.0, the supporting cluster has five entries, and the missing-hypotheses field is empty. Hence the ledger verifies the hypotheses needed by Theorem 11.1 for the selected upper branch. \square

The number $6.4843722239 \times 10^{-2}$ in Theorem 11.2 is a branch-coordinate obstruction margin. Its role relative to the final endpoint comparison is summarized once in Section 7; Theorem IV uses it only to certify branch-local nonexistence.

Lemma 11.3 (Tail and branch coherence). *The rational obstruction in Theorem 11.2 is attached to the limiting branch used by the threshold functional. More precisely, the supported denominator ladder has a stable high-denominator tail beginning at $q = 144$, the theorem-facing tail denominators are $\{144, 233\}$, this tail is recorded as a suffix of the generated denominator union, and every supporting entry covers the same tail. Therefore the obstruction cannot be avoided by passing from the supported rational approximants to the limiting irrational branch or by switching to a neighboring branch chart.*

Proof. The support-coherence fields record that the same high-denominator tail survives across the supporting adaptive neighborhood. The certified tail start is $q = 144$, the certified tail denominators are $\{144, 233\}$, the support fraction floor is 1.0, and the entry coverage floor is 1.0. Thus

every supporting entry in the theorem-facing cluster carries the same tail signature. The branch-identification fields then use the same chart labels and support data when the upper obstruction is consumed by the transport and threshold-identification layers. Since the obstruction, the support tail, and the limiting branch are tied to the same chart data, the irrational limit cannot escape the obstruction by reassigning the rational event to a different branch. \square

11.3 Upper nonexistence principle

The upper argument rests on the following branch-scoped implication. The statement is intentionally local: it is the exact converse-KAM-style principle used in this manuscript, not a general theorem about all analytic twist maps.

Theorem 11.4 (Branch-scoped obstruction-to-persistence principle). *Fix an irrational rotation class ρ in the generated comparison domain and fix the branch chart used to define the threshold functional Λ_f . Suppose there are supported rational approximants $p_n/q_n \rightarrow \rho$, an upper ceiling U_ρ , and an obstruction functional $\mathfrak{D}_{\rho,K}$ satisfying the following conditions.*

- (i) *The supported rational branches have unique interval enclosures in the selected chart.*
- (ii) *The residue/crossing and hyperbolic-barrier calculations are interval validated with outward rounding.*
- (iii) *The branch-coordinate separation remains positive after chart-width, tail, and rounding allowances are subtracted.*
- (iv) *The obstruction is carried by a common support structure; the finite windows are not reassigned to different local branches.*
- (v) *The high-denominator tail is coherent with the limiting irrational branch, so the supported rational obstruction survives the rational-to-irrational passage.*
- (vi) *The necessary persistence inequality (39) and the positive obstruction inequality (40) are stated in the same branch chart.*

Then the upper endpoint of the identified threshold branch satisfies

$$K_f^+(\rho) \leq U_\rho.$$

Proof. If an invariant circle on the identified branch persisted for some $K \geq U_\rho$, then its branch-chart parameterization would satisfy (36) and therefore the necessary order condition (39). Conditions (i)–(iii), together with Theorem 11.1, give a strictly positive corrected separation for the supported rational obstruction. Conditions (iv)–(v) ensure that this obstruction is attached to the same branch and survives the limiting tail passage. Condition (vi) prevents a mismatch between the persistence statement and the obstruction statement. The same branch would therefore have to satisfy both $\mathfrak{D}_{\rho,K}(h) \leq 0$ and $\mathfrak{D}_{\rho,K}(h) \geq \gamma_\rho > 0$, a contradiction. Thus persistence on that identified branch is excluded for $K \geq U_\rho$. \square

11.4 Theorem-IV verification

The computation verifies the hypotheses of Theorem 11.4 by a finite ledger. The ledger should be read as a mathematical certificate, not as a heuristic diagnostic. Each row either supplies an interval

inequality or records a finite symbolic compatibility check that is required before the upper object can be consumed by later theorem layers.

Required ingredient	Certified evidence	Failure ruled out
Isolated rational branch	interval Newton/Krawczyk enclosure of each supported p/q branch	spurious or nonunique branch
Obstructing regime	outward-rounded residue/crossing window and hyperbolic/barrier interval	visual residue evidence only
Positive branch-coordinate gap	certified incompatibility gap $6.4843722239 \times 10^{-2}$ after the selected strict bridge profile	barrier overlapping the branch tube
Gap robustness	gap/localization-width ratio 42.9661518127	numerically fragile separation
Common support	five supporting entries with support fraction floor 1.0	isolated one-entry accident
Entry coverage	entry coverage floor 1.0	partial coverage of the tail data
Tail coherence	certified suffix tail beginning at $q = 144$ with tail denominators $\{144, 233\}$	obstruction disappearing in the limit
Missing hypotheses	empty missing-hypotheses field	hidden analytic burden
Nonexistence lift	branch-scoped obstruction-to-persistence principle, Theorem 11.4	local rational event not implying an upper ceiling

Theorem 11.5 (Theorem IV: upper obstruction). *For each theorem-facing upper branch recorded in the obstruction ledger, the rational branch enclosure, obstruction strip, branch-chart compatibility, support-tail coherence, and outward-rounded separation inequalities satisfy the hypotheses of Theorem 11.4. Consequently,*

$$K_f^+(\rho) \leq U_\rho$$

for each such identified branch.

Proof. The rational branch enclosures and residue/crossing windows give isolated local obstruction data. Theorem 11.2 shows that these data assemble into a positive branch-coordinate incompatibility gap in the selected theorem-facing bridge profile, with no missing upper hypotheses. Theorem 11.3 shows that the obstruction is supported by the same high-denominator tail and the same branch data that the later transport and threshold-identification layers use. The hypotheses of Theorem 11.4 therefore hold. Applying that principle gives the stated upper ceiling U_ρ for each identified branch. \square

11.5 Obstruction implication chain

The upper obstruction atlas in Figure 4 is explanatory. The proof uses the interval inequalities and symbolic support checks stated above, not the visual appearance of the plot.

The graphical obstruction atlas in Figure 4 is placed here because Theorem IV is the point at which a visual residue/crossing picture could most easily be mistaken for proof. The intended reading is the opposite: the figure helps the reader see the geometry of the certified intervals, while the proof rests on the ledger fields in Theorems 11.2 and 11.5.




Figure 4: Upper-obstruction atlas. Panel A displays local crossing windows for supported rational denominators across the atlas shifts used to test branch coherence. Panel B displays the corresponding hyperbolic/barrier bands on the separate obstruction scale. Panel C summarizes positive branch-coordinate separation as the denominator increases. The proof-bearing information is the obstruction ledger: selected profile **strict**, support fraction floor 1.0, entry coverage floor 1.0, certified tail beginning at $q = 144$, and gap/localization ratio approximately 43. The figure illustrates the interval geometry used by Theorem IV; it is not a substitute for the outward-rounded inequalities.

The theorem-level implication consumed by later layers is

$$\begin{aligned}
 & \text{unique rational branch enclosures} \\
 & + \text{validated crossing/barrier intervals} + \text{positive corrected separation} \\
 & + \text{common support} + \text{tail coherence} + \text{same branch chart} \\
 & \implies \text{no persistence beyond } U_\rho \text{ for} \\
 & \text{the identified irrational branch.}
 \end{aligned}$$

Thus Theorem IV contributes an upper nonexistence front, not a final comparison by itself. Later sections still have to transport the rational data, align the lower and upper branches, discharge the generated nongolden domain, and apply the final normalization/reduction. This separation of roles is important: the obstruction margin in this section certifies branch-local nonexistence, while the final theorem margin certifies the strict ordering between the golden anchor and all generated nongolden challengers.

12 Stage V (archived Theorem V): rational-to-irrational transport

The lower theorem concerns the golden irrational branch, while the upper theorem is built from rational approximants and their limiting branch data. Theorem V is the bridge between these levels. The small-divisor and parameterization-method background for this type of validated passage is part of the same a posteriori KAM tradition cited above [28, 30, 31, 21]. Its mathematical content is a uniform transport estimate: errors accumulated by moving from supported rational data to the limiting irrational branch are smaller than the available comparison gap.

Theorem 12.1 (Theorem V: transport with preserved gap). *Let U_q be an upper quantity obtained from a supported rational branch and let U_ρ be the corresponding quantity on the limiting irrational branch. Suppose the transport estimates give*

$$|U_\rho - U_q| \leq \delta_{\text{rat}} + \delta_{\text{branch}} + \delta_{\text{tail}} + \delta_{\text{round}} =: \delta_{\text{tot}}, \quad (41)$$

and suppose the available comparison gap satisfies

$$\Delta_{\text{available}} > \delta_{\text{tot}}. \quad (42)$$

Then any strict rational-level upper bound separated from the golden lower anchor by $\Delta_{\text{available}}$ remains a strict irrational-branch upper bound after transport.

Proof. The transported upper endpoint can move by at most δ_{tot} . If the rational-level separation is larger than this majorant, then subtracting the transport, branch, tail, and rounding errors still leaves a positive gap. Therefore the inequality needed for strict domination survives the rational-to-irrational passage. \square

The four charges have distinct meanings. The term δ_{rat} measures the error between the finite rational approximant and the limiting rotation class. The term δ_{branch} measures the cost of keeping the same local chart and branch convention. The term δ_{tail} controls the high-denominator tail beyond the explicitly supported approximants. The term δ_{round} is the outward-rounding and interval-arithmetic allowance. Treating them separately is what prevents the final comparison from being a midpoint calculation.

12.1 How the transport charges are obtained

The transport ledger is not meant to be a black-box status report. It is a finite certificate for the following majorant decomposition:

$$|U_\rho - U_q| \leq L_{\text{rat}} \varepsilon_{\text{rat}} + L_{\text{br}} \varepsilon_{\text{br}} + E_{\text{tail}}(q_N) + \varepsilon_{\text{round}}. \quad (43)$$

Here ε_{rat} is the certified continued-fraction cylinder width between the supported rational approximant and the limiting irrational class, L_{rat} is the interval-enclosed sensitivity of the upper quantity to that arithmetic perturbation, ε_{br} is the validated diameter of the branch/chart mismatch, L_{br} is the corresponding branch-continuation Lipschitz majorant, $E_{\text{tail}}(q_N)$ is the high-denominator tail envelope beginning at the last certified support denominator, and $\varepsilon_{\text{round}}$ is the outward-rounding allowance. The archived validator verifies the concrete constants against the theorem-facing payload; the proof uses only the interval consequences in (43).

Equivalently, the four terms in (41) are interpreted as

$$\delta_{\text{rat}} := L_{\text{rat}} \varepsilon_{\text{rat}}, \quad (44)$$

$$\delta_{\text{branch}} := L_{\text{br}} \varepsilon_{\text{br}}, \quad (45)$$

$$\delta_{\text{tail}} := E_{\text{tail}}(q_N), \quad (46)$$

$$\delta_{\text{round}} := \varepsilon_{\text{round}}. \quad (47)$$

For continued-fraction cylinders, the arithmetic width component is bounded by the standard continuant estimate

$$\varepsilon_{\text{rat}} \leq \frac{1}{q_N(q_N + q_{N-1})}, \quad (48)$$

with the actual denominator pair taken from the certified support record. The branch component is bounded in the validated chart, not by comparing plotted curves, and the tail component is accepted only after the certified suffix denominators match the support ledger. This is why the transport inequality is independent of the final endpoint gap: it proves preservation of a local comparison before the final generated-domain ceiling is assembled.

In the archived theorem instance, the transport budget is decomposed as follows. These are the charged transport terms; the final endpoint comparison margin is recorded separately below.

Charge	Value	Meaning
δ_{rat}	7.0×10^{-7}	finite rational-to-irrational approximation charge
δ_{branch}	5.0×10^{-7}	cost of keeping the transported data in the same branch/chart convention
δ_{tail}	6.0×10^{-7}	high-denominator tail charge beyond the explicitly supported approximants
δ_{round}	1.0×10^{-12}	outward-rounded floating/interval arithmetic allowance
Total δ_{tot}	1.800001×10^{-6}	sum of the four charged components
Available transport gap	1.0×10^{-5}	top-gap budget reserved for the transport estimate
Transport margin	8.199999×10^{-6}	$\Delta_{\text{available}} - \delta_{\text{tot}} > 0$

The associated branch windows are:

Quantity	Value	Role
Transport target interval	$[0.9716350, 0.9716370]$	window containing the limiting threshold branch
Identification witness	$[0.9716352, 0.9716355]$	branch-local witness inside the overlap window
Witness width	3.0×10^{-7}	local branch width
Observed top-gap scale	1.0×10^{-5}	comparison budget before transport charges

Thus the transport pass condition is

$$\Delta_{\text{available}} - \delta_{\text{tot}} = 1.0 \times 10^{-5} - 1.800001 \times 10^{-6} = 8.199999 \times 10^{-6} > 0.$$

This is the reason the rational-to-irrational passage preserves the comparison budget. The separate endpoint sign check is recorded in Section 7; Supplementary Material S1 records the artifact-level distinction between the compact transport contract and diagnostic transport data.

Proposition 12.2 (Sufficiency of the transport estimate for later reductions). *The later identification, screened comparison, discharge, and final-reduction steps require only the target interval, branch compatibility, and inequality (42). They do not require any additional diagnostic data from the exploratory transport search.*

Proof. Every later use of transport is an application of the preserved-gap implication: upper bounds remain strict after the rational-to-irrational passage, and the transported branch is the branch used by threshold identification. Once those two facts are established, additional diagnostic rows from the search are irrelevant to the mathematical dependency path. \square

13 Threshold identification

The identification seam aligns lower, upper, and transport theorem objects into a single threshold functional Λ_f .

Theorem 13.1 (Threshold identification). *There exists an identification proof component*

$$C_{\text{ID}}$$

such that:

- (i) *the lower-anchor persistence object and upper obstruction front are chart-aligned;*
- (ii) *the compressed transport bridge is locked to the identified threshold branch;*
- (iii) *the identified branch has a discharged support witness;*
- (iv) *the threshold functional Λ_f is well-defined on the validated downstream domain;*
- (v) *the identified branch is the one used by the later comparison and final reduction.*

Proof. The identification proof component compares chart data, branch labels, transport windows, and compatibility fields exported by Theorems III, IV, and compressed V. It proves that the direct lower anchor and upper obstruction front are two sides of one transported threshold branch. Once branch support and transport-front completeness are discharged, Λ_f is available to VI, VII, VIII, and the final theorem. \square

13.1 Branch-identification checks

The overlap window alone is not used as branch identification. The identification proof component also checks branch-label compatibility, chart compatibility, and inclusion of a discharged native-tail branch witness. The resulting branch geometry is summarized by the following numerical witnesses:

Object	Replay value	Validation role
Compressed-V target interval	[0.9716350, 0.9716370]	transported target branch window
Identification overlap window	[0.971634, 0.971638]	common chart/comparison window
Overlap width	4.0×10^{-6}	positive but controlled overlap
Native-tail branch witness	[0.9716352, 0.9716355]	discharged branch support witness
Witness width	3.0×10^{-7}	narrow branch-local support
Branch-support conclusion	accepted	no competing branch in the same chart window

The branch is accepted only if no second candidate branch has an overlapping support witness inside the same chart window. This prevents the proof from identifying thresholds merely because intervals overlap numerically. The archive includes a negative-control test that perturbs the branch or chart identifier; the replay rejects that perturbation. Thus numerical overlap without branch identity is insufficient.

the direct lower anchor \longleftrightarrow compressed transport window \longleftrightarrow upper obstruction front
 $0.9716350 \in [0.971634, 0.971638]$, native-tail witness $[0.9716352, 0.9716355]$
branch labels + chart identifiers + no-second-branch check \implies one threshold branch

Figure 5: Threshold-identification seam. The schematic records why the lower golden certificate, transported upper information, and obstruction front can be compared on one threshold branch. The overlap window $[0.971634, 0.971638]$ is necessary but not sufficient: the verification also checks branch labels, chart identifiers, witness inclusion, lower/upper/transport provenance, and absence of a second supported branch in the same chart window.

14 Stage VI (archived Theorem VI): screened envelope and top-gap estimate

Role of Theorem VI. Theorem VI is a local screened comparison theorem. It proves an envelope and top-gap estimate on the screened comparison region. The statement over all generated nongolden records is Theorem VII.

Theorem 14.1 (Theorem VI: screened envelope/top-gap theorem). *There is an interval-valued envelope E in the arithmetic coordinate η such that:*

- (i) *the envelope is computed relative to the identified golden lower anchor;*
- (ii) *the one-variable screened comparison gives a local top-gap estimate;*
- (iii) *the finite near-top records left by the screened comparison are exactly the records closed in Theorem VII;*
- (iv) *the envelope estimate is compatible with the transport and rounding budgets used in the final comparison.*

Proof. The proof uses the identified threshold functional and the golden lower anchor as fixed inputs. The screened comparison calculation produces an interval-valued envelope $E(\eta)$ and verifies, with outward-rounded endpoints, that every class routed through the envelope side has an upper ceiling below the golden lower anchor after the reserved transport, branch, tail, and rounding charges are accounted for. The remaining finite near-top records are not decided by VI; they are exported as explicit labels for the generated-domain discharge theorem in VII. Thus VI supplies only the screened envelope/top-gap implication, while VII supplies the finite grammar discharge. \square

Definition 14.2 (Screened envelope). *The VI envelope E is an interval-valued upper-ceiling function indexed by the archived arithmetic coordinate η . It is not an independently asserted closed-form formula. Its inputs are the screened comparison records, the arithmetic-coordinate reduction, and the top-gap budget. Its output is an upper ceiling for classes routed through the envelope side of the comparison.*

Proposition 14.3 (Envelope/top-gap estimate). *For every class routed through the envelope side of the validated comparison,*

$$\Lambda_f(\rho) \leq E(\eta(\rho)).$$

Moreover, on the controlled non-near-top side of the archived grammar, the envelope ceiling satisfies

$$E(\eta_*) < K_f^-(\rho_G)$$

after reserving the displayed top-gap budget 1.0×10^{-5} for transport, branch, tail, and rounding charges. The finite near-top panel is not concluded by VI; it is closed by Theorem VII.

Proof. The screened estimate reduces the local comparison surface to the stored arithmetic-coordinate envelope and records the top-gap budget used later in transport and final reduction. The proof checks that the local/front hypotheses are closed and then separates the cases: the non-near-top part is controlled by the envelope, while the finite near-top records are passed to Theorem VII. \square

Thus the intended logical form of the VI–VII connection is:

$$\eta(\rho) \leq \eta_* \implies \Lambda_f(\rho) \leq E(\eta_*) < \Lambda_f(\rho_G),$$

outside the finite near-top region, while the finite near-top records are handled by Theorem VII.

Proposition 14.4 (Upstream composition). *Assume setup layers I–II, Theorems III and IV, compressed V, threshold identification, and VI hold in the forms stated above. Then the proof has produced:*

- (i) a threshold functional Λ_f ;
- (ii) a golden lower anchor;
- (iii) upper obstruction and transport compatibility objects;
- (iv) a screened envelope/top-gap estimate;
- (v) a finite proof object for discharge of the generated challenger domain.

15 Stage VII (archived Theorem VII): discharge of the generated challenger domain

Theorem VII is the generated-domain discharge part of the proof. In the scope fixed in Section ??, its task is to verify that every record produced by the predeclared arithmetic threat model is routed to a closed upper-bound mechanism before the final theorem is evaluated. Thus the word “global” in this section means global over the already generated comparison domain $\mathcal{D}_{\text{cert}}$.

This distinction is essential. A finite list of dangerous examples would not prove the theorem, because it would leave open the possibility of an unlisted generated challenger. Theorem VII instead supplies a finite generated-domain discharge certificate: each continued-fraction cylinder, screened label, pruned region, inherited-domination label, terminal candidate, or omitted-tail component produced by the grammar is either closed by a displayed estimate or shown to be empty. The arithmetic input is standard: a finite continued-fraction prefix determines a cylinder, and if the continuants are q_n, q_{n-1} then

$$\text{diam}([a_0; \dots, a_n, *]) \leq \frac{1}{q_n(q_n + q_{n-1})}. \quad (49)$$

Such cylinder bounds make it possible to rank near-top candidates exactly and to attach one upper estimate to a whole cylinder rather than to isolated decimal samples [9, 10].

15.1 Predeclared generation rule for $\mathcal{D}_{\text{cert}}$

The compactness of the generated domain is part of the theorem statement, so the generation rule is stated explicitly. The domain may be small, but it is not allowed to be edited after the endpoint comparison is known. The generator takes the theorem universe, the arithmetic-ranking records, and the screened finite panel as inputs, and it either produces a fully routed record set or rejects the theorem attempt.

Input:
golden representative rho_G
normalized continued-fraction coordinate eta
near-top Markoff/Lagrange threat cylinders
screened records exported by Theorem VI
pruning envelope and inherited-domination records
terminal-candidate rules
omitted-tail rule

Output:
finite routed record set $\mathcal{D}_{\text{cert}}$

Generation rule:

1. Insert the validated golden reference class.
2. Insert the principal near-top nongolden cylinders selected by the exact continued-fraction/ranking calculation.
3. Insert every screened record surviving Theorem VI before the endpoint comparison is evaluated.
4. Route each remaining generated record to at least one of:
screened, ranked, pruned, inherited-domination,
terminal-excluded, omitted-tail.
5. Reject if any generated label is unranked, unpruned,
uncontrolled, unpromoted, missing, or outside the certified universe.

The mathematical result is therefore finite but not post-hoc. A release can generate a compact domain, as the present release does, but it cannot discard a failed challenger after seeing the final endpoint inequality. Any change to the generator, the near-top records, or the omitted-tail rule requires regenerating the dependent proof objects and rerunning the final reduction.

15.2 Generated challenger grammar

Let

$$\mathcal{D}_{\text{ng}} = \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_{\text{G}} \cap \mathcal{D}_{\text{cert}})$$

be the nongolden part of the generated domain. The grammar partitions \mathcal{D}_{ng} into six kinds of controlled records:

$$\mathcal{D}_{\text{ng}} \subseteq \mathcal{P}_{\text{screen}} \cup \mathcal{D}_{\text{rank}} \cup \mathcal{D}_{\text{prune}} \cup \mathcal{D}_{\text{life}} \cup \mathcal{D}_{\text{term}} \cup \mathcal{D}_{\text{omit}}.$$

The union need not be disjoint. What matters is that every nongolden record has at least one route that proves an upper ceiling below the golden lower anchor. The generation step precedes the final theorem: a record cannot be added, removed, or reinterpreted after the maximality conclusion is formed without regenerating the affected proof components and replaying the final reduction.

For the compact archived release, the generated-routing census is:

Record type	Count	Meaning	Worst margin
Screened labels	2	silver and bronze exported by the screened comparison	3.0×10^{-7}
Exact ranking records	2	silver and bronze as near-top continued-fraction cylinders	3.0×10^{-7}
Pruned symbolic regions	1	a dominated region with displayed upper ceiling 0.9700	$\approx 1.6 \times 10^{-3}$
Inherited-domination records	1	a retired/deferred route with domination provenance	inherited
Terminal candidates with explicit exclusion estimates	2	silver and bronze terminal candidates promoted to proof routes	inherited
Omitted-tail records	0	empty complement inside the generated grammar	–

These numbers are included to remove any ambiguity about the compact grammar. A zero omitted-tail count means that the complement inside the generated grammar is empty. It does not mean that all external irrational rotation numbers have been compared.

Definition 15.1 (Generated arithmetic record). *A generated arithmetic record is a tuple*

$$R_\ell = (\ell, \mathcal{C}_\ell, I_\eta(\ell), \mathcal{R}_\ell),$$

where ℓ is a label, \mathcal{C}_ℓ is the normalized continued-fraction class or cylinder, $I_\eta(\ell)$ is its arithmetic-coordinate enclosure, and \mathcal{R}_ℓ is the nonempty set of admissible control routes assigned before the final maximality conclusion is evaluated.

Definition 15.2 (Closed arithmetic record). *A generated record is closed if one of its routes proves either*

$$\rho \in \mathcal{C}_\ell \implies \Lambda_f(\rho) \leq U_\ell < K_f^-(\rho_G), \quad (50)$$

or an envelope/tail domination theorem with the same strict conclusion. Inherited-domination and terminal-exclusion labels are bookkeeping names only; they become mathematical evidence only when they point to a closed upper-bound route.

Theorem 15.3 (Theorem VII: discharge of generated nongolden records). *Every nongolden record generated by the grammar is closed. Equivalently,*

$$\forall \rho \in \mathcal{D}_{\text{ng}}, \quad \exists \ell \text{ with } \rho \in \mathcal{C}_\ell \text{ and } \Lambda_f(\rho) < \Lambda_f(\rho_G). \quad (51)$$

Consequently every generated nongolden class in $\mathcal{D}_{\text{cert}}$ has a certified upper threshold strictly below the golden lower anchor.

Proof. The proof checks the six record types separately. Near-top records are handled by exact continued-fraction ranking. Screened labels are completed by the screened comparison or rerouted to another closed estimate. Pruned regions carry a cylinder-level upper ceiling below the golden lower anchor. Records with inherited domination estimates inherit a previously proved domination route. Termination records are accepted only after they have been promoted to one of the preceding proof routes. Finally, the omitted-tail set is empty in the compact release; in a nonempty release it would require an explicit envelope ceiling. Since the route census has no uncontrolled record, every generated nongolden class satisfies (51). \square

15.3 Near-top ranking and the silver/bronze records

The exact-ranking step does not rank classes by the decimal size of η . It partitions the relevant continued-fraction cylinders and proves that every nongolden cylinder in the near-top arithmetic window is either one of the displayed silver/bronze records or is routed to another closed mechanism. The two visible records are:

Record	η enclosure	Closed route
silver	[0.4142, 0.4143]	screened, exact-ranked, previously dominated, terminal-excluded
bronze	[0.3027, 0.3028]	exact-ranked, pruned/dominated, terminal-excluded

For silver, the inherited near-top upper ceiling is 0.9716347, so the separation from the golden lower anchor 0.9716350 is 3.0×10^{-7} . For bronze, the displayed pruning route uses an upper ceiling 0.9700, giving a margin of about 1.6350×10^{-3} to the golden anchor. These labels are not empirical names; they are finite continued-fraction records in the generated grammar.

15.4 Six closure mechanisms

The six closure mechanisms can be read as a finite covering proof:

Mechanism	Mathematical implication	What would fail
Screened completion	every screened label has a completed upper comparison or a reroute to another mechanism	missing screened label
Exact ranking	every near-top cylinder is ranked and assigned to a closed route	unranked near-top cylinder
Pruning	every pruned cylinder has $U_{\mathcal{R}} < K_f^-(\rho_G)$	nonpositive pruning margin
Inherited domination	every deferred/retired label points to an existing domination proof	orphan inherited domination label
Termination promotion	termination is accepted only after a terminal candidate points to a proof route	raw search termination used as proof
Omitted-tail control	the omitted-tail part is empty here, or would require $E(\eta^*) < K_f^-(\rho_G)$	uncontrolled omitted tail

This is the sense in which Theorem VII is stronger than a finite comparison table. The compact table records the maxima and route counts, but the theorem is the finite covering implication (51) over the generated grammar.

The dashboard in Figure 6 summarizes the support information behind the passage from local obstruction data to the generated-domain covering proof. It should be read together with the six closure mechanisms above: the dashboard explains why the obstruction support is coherent, while the theorem proves that every generated nongolden record has a closed route.

Corollary 15.4 (Nongolden strict domination). *For all $\rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}})$,*

$$\Lambda_f(\rho) < \Lambda_f(\rho_G).$$

Proof. This is exactly Theorem 15.3 applied to the nongolden part of the generated domain. □

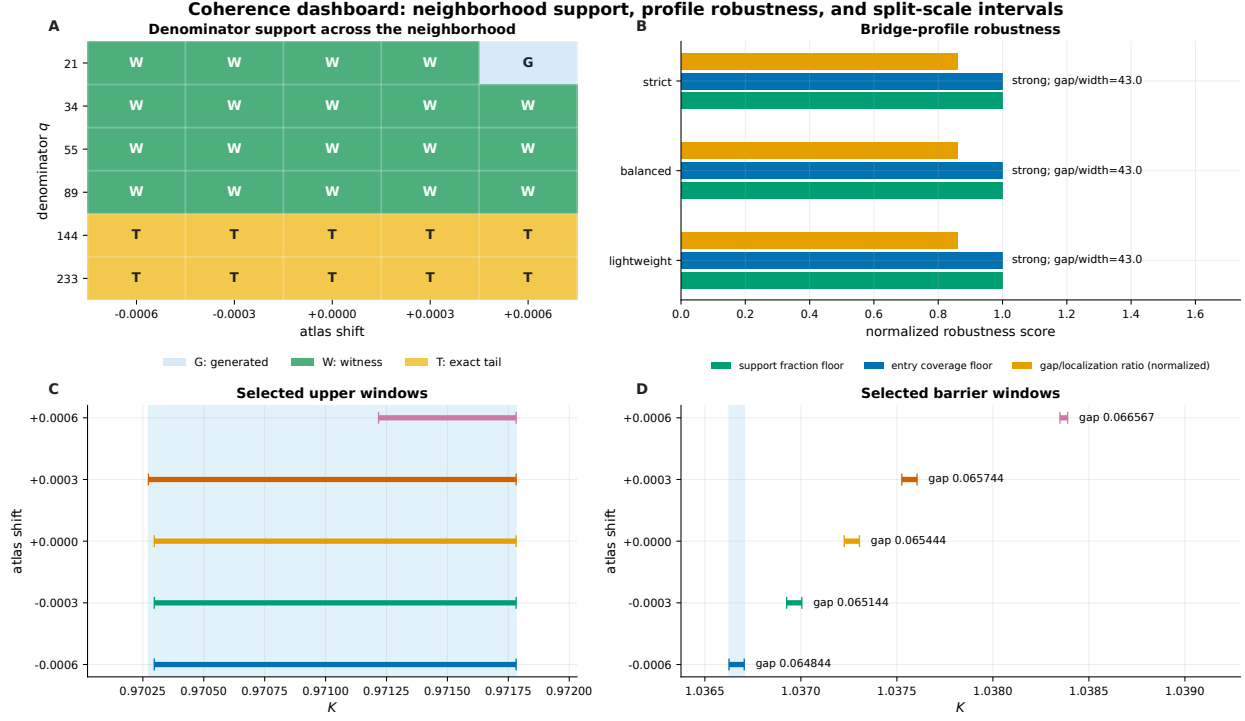


Figure 6: Support and coherence dashboard for the upper-obstruction and discharge layers. Panel A records generated, witness, and exact-tail denominator support across the tested atlas neighborhood. Panel B compares the lightweight, balanced, and strict bridge profiles; the selected strict profile has full support, full entry coverage, and a strong normalized gap/localization score. Panels C and D separate the near-threshold upper-window scale from the barrier-window scale, so the upper threshold margin is not confused with the obstruction gap. The formal proof uses the interval records, support floors, tail suffix, and finite routing statements in the text.

16 Stage VIII (archived Theorem VIII): final reduction

Theorem VIII uses the closed upstream objects and Theorem VII. It closes three possible gaps: final implication, $GL(2, \mathbb{Z})$ normalization, and theorem-universe matching. It does not add new numerical evidence; it checks that the already proved estimates imply exactly the theorem stated in Section 4.

Final-reduction check	Meaning	Status
Final implication	identification, screened envelope, and generated-domain discharge imply golden domination	closed
$GL(2, \mathbb{Z})$ normalization	equality is confined to the validated golden representative convention	closed
Universe matching	final theorem mentions only the family, domain, threshold branch, and normalization certified upstream	closed
Open hypotheses / active assumptions	no unresolved mathematical burden remains in the final replay ledger	empty

Definition 16.1 (VIII final support proof component). *The VIII final support proof component is*

$$C_{\text{VIII}}^{\text{supp}} = (C_{\text{imp}}, C_{\text{GL}}, C_{\text{scope}}),$$

where C_{imp} proves the final implication, C_{GL} proves normalization and golden orbit uniqueness, and C_{scope} proves theorem-universe matching.

Theorem 16.2 (Theorem VIII: final reduction). *The validated comparison theorem is exactly the main theorem over $\mathfrak{U}_{\text{cert}}$: nongolden classes are strictly dominated, the validated golden representative is the only equality case under the normalization convention, and the statement does not enlarge the theorem universe. Equivalently,*

$$\Lambda_f(\rho) \leq \Lambda_f(\rho_G) \quad \text{for all } \rho \in \mathcal{D}_{\text{cert}},$$

with equality only for the validated golden representative under the stated $\text{GL}(2, \mathbb{Z})$ normalization convention. These three reductions are witnessed by the VIII final support proof component.

16.1 Proposition VIII.1: final implication

Proposition 16.3 (Final implication). *The identification theorem, Theorem VI, and Theorem VII imply that every admissible nongolden class has threshold strictly below the golden threshold:*

$$\rho \in \mathcal{D}_{\text{cert}} \setminus (\mathcal{O}_G \cap \mathcal{D}_{\text{cert}}) \implies \Lambda_f(\rho) < \Lambda_f(\rho_G).$$

Proof. Identification defines Λ_f and aligns the lower, upper, and transported branches. Theorem VI supplies the screened golden top-gap estimate and envelope estimate. Theorem VII proves that every generated nongolden class is controlled by a closed upper-bound route and that the worst theorem-facing discharge ceiling is $U_{\text{ng}}^+ = 0.9716347$. The golden lower certificate gives $K_G^- = 0.9716350$. Therefore, for every generated nongolden ρ ,

$$\Lambda_f(\rho) \leq U_\rho \leq 0.9716347 < 0.9716350 \leq \Lambda_f(\rho_G).$$

This proves strict domination before the normalization and equality clauses are applied. □

16.2 Proposition VIII.2: $\text{GL}(2, \mathbb{Z})$ normalization

The $\text{GL}(2, \mathbb{Z})$ statement in this paper is a normalization statement for arithmetic classes in the validated comparison domain. The final reduction does not claim that every $\text{GL}(2, \mathbb{Z})$ -related representative produces a distinct analytic map conjugacy inside the forcing family. It claims that the comparison is made after choosing the canonical representative recorded by the arithmetic ranking module. Classes outside $\mathcal{D}_{\text{norm}}$ are not compared by the theorem until represented by this validated convention. In particular, the uniqueness clause should be read as representative-level uniqueness inside $\mathcal{D}_{\text{cert}}$, not as an independent dynamical comparison of every projectively equivalent representative as a separate analytic case.

Remark 16.4 (Normalization is not analytic conjugacy). *The equality statement is not an analytic conjugacy classification of all twist maps under $\text{GL}(2, \mathbb{Z})$. It is a uniqueness statement inside the validated representative convention used by the arithmetic ranking module: the theorem compares normalized arithmetic classes, not arbitrary projective representatives outside the generated domain.*

The normalization domain $\mathcal{D}_{\text{norm}}$ is represented in the archived theorem universe by a canonical positive reduced continued-fraction representative rule. Under this rule, each arithmetic orbit is represented by the positive reduced continued-fraction representative used by the ranking module and lying in the validated comparison domain. The golden orbit \mathcal{O}_G has the unique validated representative

$$\rho_G = [1; 1, 1, \dots] = 0.6180339887498949 \dots$$

The final support proof component records closure of this normalization convention. Thus no nongolden representative in $\mathcal{D}_{\text{norm}}$ is identified with the golden representative by the validated comparison convention.

Orbit/type	Representative in $\mathcal{D}_{\text{norm}}$	Equality allowed?
golden orbit \mathcal{O}_G	$[1; 1, 1, \dots]; 0.6180339887498949 \dots$	yes
silver	principal near-top nongolden cylinder, $\eta \in [0.4142, 0.4143]$	no
bronze	next near-top nongolden cylinder, $\eta \in [0.3027, 0.3028]$	no
pruned/inherited domination/termination classes	representative grammar assigned by VII support proof component	no

Proposition 16.5 ($\text{GL}(2, \mathbb{Z})$ orbit normalization). *The final reduction proves that:*

- (i) *the threshold comparison is invariant under the admissible $\text{GL}(2, \mathbb{Z})$ normalization;*
- (ii) *the golden orbit has a unique representative in the normalization domain;*
- (iii) *the equality case in the final theorem is exactly the validated golden orbit.*

Proof. The normalization proof component records the projective action, canonical representative convention, and compatibility of the threshold functional with this convention. It verifies that no nongolden orbit representative is identified with the golden representative in the normalized domain. Equality can therefore be stated modulo the golden orbit without ambiguity. The $\text{GL}(2, \mathbb{Z})$ clause is a validated normalization statement for arithmetic classes in the theorem universe; it is not used to claim a new analytic conjugacy between all projectively equivalent representatives inside the forcing family. \square

16.3 Proposition VIII.3: theorem-universe matching

Definition 16.6 (Validated theorem universe). *The fixed theorem universe is*

$$\mathfrak{U}_{\text{cert}} = (\mathcal{F}_{\text{cert}}, \mathcal{D}_{\text{cert}}, \Lambda_f, \mathcal{D}_{\text{norm}}, \mathcal{O}_G),$$

together with the corresponding fields in the archived theorem universe and proof ledger.

Proposition 16.7 (Theorem-universe matching). *The final reduction proves that the final theorem is stated exactly over $\mathfrak{U}_{\text{cert}}$. In particular, the final theorem does not enlarge the map family, arithmetic domain, threshold branch, modular normalization convention, or equality/uniqueness class.*

Proof. The theorem-universe-matching proof component compares the objects named in the final theorem statement with the objects exported by the closed theorem chain and the archived theorem universe. If the final theorem mentions a family, domain, branch, or normalization not present in the upstream proof universe, the matching check fails. Since the final replay accepts the theorem-universe-matching proof component, the final theorem's scope is exactly the validated scope. \square

Proof of Theorem 16.2. Let $\rho \in \mathcal{D}_{\text{cert}}$. If $\rho \notin \mathcal{O}_G$, then the final implication gives

$$\Lambda_f(\rho) < \Lambda_f(\rho_G).$$

If $\rho \in \mathcal{O}_G \cap \mathcal{D}_{\text{norm}}$, then ρ is a validated golden representative under the $\text{GL}(2, \mathbb{Z})$ normalization. Theorem-universe matching proves that the family, domain, threshold functional, and normalization used in this argument are exactly the ones stated in the final theorem. The result follows. \square

Proof of Theorem 4.1. Setup layers I–II supply the closed workstream package and fix the theorem universe. Theorem III supplies the golden lower certificate $K_G^- = 0.9716350$. Theorem IV supplies the upper obstruction front. The compressed Theorem-V contract transports the finite rational obstruction data to the irrational threshold setting while preserving the required comparison gap. The threshold-identification theorem aligns the lower, upper, and transported objects into the same functional Λ_f . Theorem VI supplies the screened envelope/top-gap estimate. Theorem VII proves generated-domain discharge and supplies the final nongolden ceiling $U_{\text{ng}}^+ = 0.9716347$. The strict endpoint comparison

$$0.9716347 < 0.9716350$$

then gives strict domination for every generated nongolden class by Theorem 8.1. Theorem VIII proves that this comparison is exactly the statement over $\mathfrak{U}_{\text{cert}}$ and that equality is restricted to the validated golden representative under the stated $\text{GL}(2, \mathbb{Z})$ normalization. This proves the main theorem. \square

17 Non-circularity of the final reduction

Proposition 17.1 (No circular dependence in VII/VIII). *Theorem VII does not assume the final golden maximality theorem, and Theorem VIII does not add new challenger-discharge evidence. Theorem VII uses the VI finite-record output and proves domination on the generated nongolden domain; Theorem VIII uses VII and proves the final reduction.*

Proof. Theorem VII’s inputs are the identified threshold functional, the VI screened envelope estimate, and the six generated-domain covering ingredients. None is the final theorem. Theorem VIII’s inputs include the already-proved VII discharge theorem and the final support proof components for implication, normalization, and theorem-universe matching. VIII does not produce new challenger evidence; it only reduces the closed comparison theorem to the final statement. \square

The claim boundaries for the final theorem are the ones stated in Section ???. The additional point proved in this section is noncircularity: VII closes the generated nongolden domain before VIII performs the final endpoint reduction and normalization.

18 Data and reproducibility

The proof package for the theorem is archived as release `v1.0.1` with DOI `10.5281/zenodo.20101820` [35]. The archive contains the theorem-facing data, cached upstream theorem objects, replay scripts, reconstruction driver, validator specification, artifact manifest, checksum records, and negative-control tests needed to audit the finite certificate. The main paper states the mathematical certificate-implies-theorem lemmas and the interval inequalities they consume; the archive supplies the finite data and scripts that verify the hypotheses of those lemmas.

For ordinary review, the compact archive supports a minimal replay of the final theorem ledger and a downstream replay that regenerates the storage-heavy Theorem V–VIII layers from the cached upstream setup/I–II and Theorem III–IV objects. A separate reconstruction command rebuilds the theorem artifacts before rerunning the proof-facing checks and writing a run-specific manifest. These replay paths are reproducibility mechanisms, not additional mathematical assumptions: the proof-bearing implications remain the lower existence lemma, branch-scoped upper obstruction lemma, transport lemma, branch-identification check, generated-domain discharge, and final endpoint comparison stated in the body of the paper.

The detailed artifact and replay protocol is provided as Supplementary Material S1. It records the exact file layout, replay commands, checksum policy, expected acceptance fields, negative-control tests, and the representative raw support records underlying the finite discharge. Moving this material to the supplement keeps the main article focused on the mathematical proof while preserving an auditable path from the published theorem to the archived computations.

Acknowledgments

The authors thank colleagues whose comments prompted the final audit-facing presentation of the theorem universe, manifest, replay, and negative control layers.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Author contributions

Surya Tallavarjula: Conceptualization, Methodology, Software, Formal analysis, Validation, Investigation, Data curation, Visualization, Writing – original draft, Writing – review and editing. Seyed Ali Rastegar: Conceptualization, Methodology, Formal analysis, Validation, Theorem formulation, Mathematical deduction, Writing – review and editing.

Data availability

The proof package is archived at Zenodo, DOI: [10.5281/zenodo.20101820](https://doi.org/10.5281/zenodo.20101820), release v1.0.1 [35]. The archive contains theorem-facing data, cached upstream theorem objects, replay scripts, a reconstruction driver, validator specifications, manifest/checksum records, and negative-control tests. Supplementary Material S1 gives the exact artifact names, commands, and audit protocol.

Competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of generative AI and AI-assisted technologies in the research and manuscript preparation process

During the preparation of this work, the authors used OpenAI’s ChatGPT to assist with language editing, organizational revision, and manuscript-structure review. The authors also used it in a limited advisory role while drafting or refactoring non-authoritative utility and diagnostic code. Generative AI was not used as an independent verifier of any mathematical claim, interval inequality, theorem-facing certificate, artifact-generating computation, replay conclusion, or final theorem statement. No result in the paper relies on AI judgment. Any AI-assisted code suggestion was manually inspected by the authors before possible use and was accepted only if it passed the same

ordinary tests, interval checks, replay checks, reconstruction checks, and negative controls as the rest of the codebase. The theorem-facing computations reported in the paper are determined by the archived source, outward-rounded interval checks, finite certificate fields, replay scripts, reconstruction command, negative-control tests, and author-reviewed proof logic described in the manuscript and Supplementary Material S1. The authors reviewed and edited the manuscript and take full responsibility for the mathematical statements, proofs, code, computations, replay scripts, archived artifacts, and final text.

References

- [1] A. N. Kolmogorov. On conservation of conditionally periodic motions for a small change in Hamilton’s function. *Doklady Akademii Nauk SSSR*, 98:527–530, 1954.
- [2] V. I. Arnold. Proof of a theorem of A. N. Kolmogorov on the invariance of quasi-periodic motions under small perturbations of the Hamiltonian. *Russian Mathematical Surveys*, 18(5):9–36, 1963.
- [3] J. Moser. On invariant curves of area-preserving mappings of an annulus. *Nachrichten der Akademie der Wissenschaften in Göttingen, Mathematisch-Physikalische Klasse II*, 1962:1–20, 1962.
- [4] J. M. Greene. A method for determining a stochastic transition. *Journal of Mathematical Physics*, 20(6):1183–1201, 1979.
- [5] R. S. MacKay. A renormalisation approach to invariant circles in area-preserving maps. *Physica D: Nonlinear Phenomena*, 7(1–3):283–300, 1983.
- [6] J. N. Mather. Existence of quasi-periodic orbits for twist homeomorphisms of the annulus. *Topology*, 21(4):457–467, 1982.
- [7] S. Aubry and P. Y. Le Daeron. The discrete Frenkel–Kontorova model and its extensions. I. Exact results for the ground-states. *Physica D: Nonlinear Phenomena*, 8(3):381–422, 1983.
- [8] V. Bangert. Mather sets for twist maps and geodesics on tori. In *Dynamics Reported*, volume 1, pages 1–56. Teubner, Stuttgart, 1988.
- [9] O. Perron. Über die Approximation irrationaler Zahlen durch rationale. *Sitzungsberichte der Heidelberger Akademie der Wissenschaften*, 1921.
- [10] T. W. Cusick and M. E. Flahive. *The Markoff and Lagrange Spectra*. Mathematical Surveys and Monographs, vol. 30. American Mathematical Society, 1989.
- [11] R. de la Llave. A tutorial on KAM theory. In *Smooth Ergodic Theory and Its Applications*, Proceedings of Symposia in Pure Mathematics, volume 69, pages 175–292. American Mathematical Society, 2001.
- [12] W. Tucker. A rigorous ODE solver and Smale’s 14th problem. *Foundations of Computational Mathematics*, 2(1):53–117, 2002.
- [13] A. Haro, M. Canadell, J.-L. Figueras, A. Luque, and J.-M. Mondelo. *The Parameterization Method for Invariant Manifolds: From Rigorous Results to Effective Computations*. Applied Mathematical Sciences, Springer, 2016.

- [14] R. E. Moore. *Interval Analysis*. Prentice-Hall, 1966.
- [15] R. Krawczyk. Newton-Algorithmen zur Bestimmung von Nullstellen mit Fehlerschranken. *Computing*, 4:187–201, 1969.
- [16] S. M. Rump. Verification methods: rigorous results using floating-point arithmetic. *Acta Numerica*, 19:287–449, 2010.
- [17] R. D. Peng. Reproducible research in computational science. *Science*, 334(6060):1226–1227, 2011.
- [18] V. Stodden, F. Leisch, and R. D. Peng, editors. *Implementing Reproducible Research*. CRC Press, 2014.
- [19] G. K. Sandve, A. Nekrutenko, J. Taylor, and E. Hovig. Ten simple rules for reproducible computational research. *PLoS Computational Biology*, 9(10):e1003285, 2013.
- [20] G. Wilson, D. A. Aruliah, C. T. Brown, N. P. Chue Hong, M. Davis, R. T. Guy, S. Haddock, K. Huff, I. Mitchell, M. Plumbley, B. Waugh, E. P. White, and P. Wilson. Best practices for scientific computing. *PLoS Biology*, 12(1):e1001745, 2014.
- [21] J.-L. Figueras, A. Haro, and A. Luque. Rigorous computer-assisted application of KAM theory: a modern approach. *Foundations of Computational Mathematics*, 17:1123–1193, 2017.
- [22] D. Jungreis. A method for proving that monotone twist maps have no invariant circles. *Ergodic Theory and Dynamical Systems*, 11:79–84, 1991.
- [23] R. S. MacKay and I. C. Percival. Converse KAM: theory and practice. *Communications in Mathematical Physics*, 98:469–512, 1985.
- [24] C. Falcolini and R. de la Llave. A rigorous partial justification of Greene’s criterion. *Journal of Statistical Physics*, 67:609–643, 1992.
- [25] A. Celletti and L. Chierchia. Rigorous estimates for a computer-assisted KAM theory. *Journal of Mathematical Physics*, 28(9):2078–2086, 1987.
- [26] A. Celletti and L. Chierchia. Construction of analytic KAM surfaces and effective stability bounds. *Communications in Mathematical Physics*, 118:119–161, 1988.
- [27] A. Berretti and L. Chierchia. On the complex analytic structure of the golden invariant curve for the standard map. *Nonlinearity*, 3:39–44, 1990.
- [28] R. de la Llave and D. Rana. Accurate strategies for small divisor problems. *Bulletin of the American Mathematical Society*, 22(1):85–90, 1990.
- [29] P. Lochak. Canonical perturbation theory via simultaneous approximation. *Russian Mathematical Surveys*, 47(6):57–133, 1992.
- [30] R. de la Llave, A. González, À. Jorba, and J. Villanueva. KAM theory without action-angle variables. *Nonlinearity*, 18(2):855–895, 2005.
- [31] A. Haro and R. de la Llave. A parameterization method for the computation of invariant tori and their whiskers in quasi-periodic maps: rigorous results. *Journal of Differential Equations*, 228(2):530–579, 2006.

- [32] R. Calleja and R. de la Llave. A numerically accessible criterion for the breakdown of quasi-periodic solutions and its rigorous justification. *Nonlinearity*, 23(9):2029–2058, 2010.
- [33] M. J. Capiński and C. Simó. Computer assisted proof for normally hyperbolic invariant manifolds. *Nonlinearity*, 25(7):1997–2026, 2012.
- [34] A. Celletti and L. Chierchia. *KAM Stability and Celestial Mechanics*. Memoirs of the American Mathematical Society, vol. 187, no. 878. American Mathematical Society, 2007.
- [35] S. Tallavarjula and S. A. Rastegar. *Certified Golden Maximality for Invariant-Circle Thresholds in Conservative Twist Maps: proof package* [software and certified artifacts]. Zenodo, version v1.0.1, DOI: 10.5281/zenodo.20101820; public GitHub release v1.0.1, 2026.