

Synchré Form Theory III: Target Inference and Certificate Construction from Data

Drew Higgins

Abstract

This paper turns the SFT decision contract into a data-to-artifact workflow. Given samples of phase vectors together with a declared model class and tolerance, the task is to infer the correct target (full torus versus coupled subtorus coset), estimate witness integers (degree or holonomy multiple) on an observed subsystem [1, 2], and attach robustness margins and convergence certificates only when justified [4, 5, 7]. The resulting output type is exhaustive inside COC: a witness stability certificate (WSC) on the correct target, or a canonical obstruction report (WOB) with evidence.

Global universality meaning

Definition 1 (Cyclic Observability Class (COC)). *The Cyclic Observability Class (COC) is the class of declared SFT contexts $\mathbf{c}$ satisfying all of the following three conditions (and it is maximal by definition, since membership is exactly these conditions):*

- *a phase observable is declared on the relevant slice/subsystem: a cycle-valued map (circle/torus/coupled target) with the minimal regularity needed to define the intended winding/holonomy notion and to test nontriviality;*
- *an admissible witness mechanism is declared (symmetry action, loop/holonomy, embedded subsystem witness, defect loop, or another explicitly listed mechanism), so the engine has a prescribed route to either a WSC or to a named obstruction code;*
- *a correct target rule is satisfied: the target type (circle/torus/CST) and any baseline statements are defined on the correct target (Haar on the target, coupled Haar on CST when coupling is present).*

Definition 2 (Universality on a class of contexts). *Fix a class of contexts $\mathcal{C}$. The SFT engine is universal on $\mathcal{C}$ if for every declared context $\mathbf{c} \in \mathcal{C}$ and every slice/subsystem where the claim is posed, the engine returns exactly one of:*

- *a Witness Stability Certificate (WSC) proving coverage on the correct target in $\mathbf{c}$ (with explicit margins), or*
- *a Witness Obstruction (WOB) belonging to a finite obstruction catalog (with evidence), certifying that the claim is blocked in $\mathbf{c}$.*

Global universality means: universal on COC. In SFT, a statement labeled *global universality* is a universality statement on COC: for every context in COC, the engine returns a WSC or a WOB on the correct target; outside COC, claims are either inapplicable (target/predicates not well-posed) or produce a boundary-type obstruction rather than a forced-coverage theorem.

COC checklist (entrywise). A declared context is treated as belonging to COOnly when all of the following are explicitly provided:

- **Phase observable (PO).** A named cycle-valued observable on the intended slice/subsystem (circle/torus/CST), together with the minimum regularity assumptions needed to define its winding/holonomy notion and to test nontriviality.
- **Admissible witness mechanism (WM).** A declared mechanism from the admissible menu (symmetry action, loop/holonomy, embedded subsystem witness, defect loop, or another explicitly listed mechanism) that specifies what the engine must attempt to certify.
- **Target rule (TR).** A declared target type (circle/torus/CST) and baseline convention on that target (Haar on the target; coupled Haar on CST when coupling is present).

When any checklist item fails, the correct engine output is a named obstruction (WOB) rather than a coverage certificate.

Well-posedness spine: what a universality claim means in SFT

SFT treats universality as a statement with an explicit boundary. The boundary is not a philosophical warning; it is part of the formal interface. Every claim is posed in a *declared context* and every answer is a finite object. This is the only sense in which SFT uses the word *meaning*: a claim is meaningful exactly when it has an auditable certificate-or-obstruction output.

Definition 3 (Declared SFT claim). *A declared SFT claim is a tuple*

$$\mathbf{q} = (\mathbf{c}, \mathbf{T}, \mathbf{S}),$$

where $\mathbf{c}$ is a declared context (with a phase observable and an admissible witness mechanism), $\mathbf{T}$ is the declared correct target (circle/torus/coupled target) together with its baseline convention (Haar on the target; coupled Haar on CST when coupling is present), and $\mathbf{S}$ is the concrete Synchré statement being tested (coverage, stability, and any declared upgrade).

Well-posedness axiom (WSC/WOB meaning). For every declared claim $\mathbf{q}$, the SFT engine returns exactly one of the following two outputs.

- **WSC (Witness Stability Certificate).** A finite certificate proving the stated claim on the declared correct target, together with explicit margins that make the claim stable in the declared perturbation model.
- **WOB (Witness Obstruction).** A finite obstruction report from a named catalog, with evidence, asserting that certification is blocked or that the claim is inapplicable in the declared regime.

The claim is *meaningful in SFT* if and only if it is posed as such an engine query and therefore admits a WSC/WOB output.

The sharp boundary can be read in one glance.

Where the claim lives	Correct output
Inside COC (phase observable + witness route + target rule)	WSC or a finite WOB (structural failure mode)
Outside COC (missing phase / missing witness route / wrong target rule)	Boundary-type WOB (inapplicable or ill-posed)

How to read any “universality” sentence in this project.

- First check whether the context is declared as belonging to COC (the checklist is part of the definition).
- Then check the target: circle, full torus, or a coupled target CST.
- Then check the object-level output: a WSC (with margins) or a WOB (with a named obstruction code and evidence).

This spine is the project’s scope wall in mathematical form. It prevents universality from becoming vague: every use of the word is anchored to the named class (COC) defined by the interface conditions and to an explicit certificate-or-obstruction output.

Target correctness axiom block (CTP/CHT). SFT treats *target correctness* as part of the meaning of any certificate-or-obstruction claim. A declared claim must state (i) the *target type* (circle / full torus / coupled target) and (ii) the *baseline measure* used for any comparison, stability, or convergence upgrade.

- **Circle or full torus targets.** When the correct target is a compact group target $\mathbb{T}_P$ or $\mathbb{T}_P$, the baseline is the (unique) Haar probability measure on that target.
- **Coupled targets (CST).** When coupling constraints restrict the image to a coupled subtorus coset,

$$\text{CST}_s := \{y \in \mathbb{T}_P : Ay = \Psi(s)\}, \quad A \in \mathbb{Z}^{r \times d},$$

then the correct target is CST_s (not the ambient torus) and the baseline is the (unique) translation-invariant probability measure on CST_s . Equivalently, for any parameterization $\eta_s : \mathbb{T}_{P_*}^m \rightarrow \text{CST}_s$, the baseline is the pushforward of Haar:

$$\mu_{H,s}^{\text{CST}} := (\eta_s)_\# \mu_H^{(m)}.$$

CTP (Coupled Target Principle). If the declared context includes coupling constraints yielding a coupled target CST_s , then all coverage statements, stability margins, and upgrades are to be interpreted *relative to* CST_s . Using the ambient torus as the target in this regime is not a weaker claim; it is a *different* claim.

CHT (Coupled Haar Theorem, as baseline convention). In the coupled regime, the baseline for comparisons is $\mu_{H,s}^{\text{CST}}$. Any convergence or mixing claim must specify this baseline when the target is CST_s .

Corollary (wrong target/baseline forces a boundary obstruction). If a declared claim uses a target/baseline pair incompatible with the declared coupling structure (e.g., it uses full-torus Haar when the correct target is CST_s), then the claim is *ill-posed in SFT*. The correct engine output in that case is the coupling-mismatch obstruction WOB-06-coupling-mismatch.

Target correctness decision flow

This page is the canonical interface for selecting the *correct target* and *baseline* in SFT. It is written to match the WSC schema fields `target_type` and `baseline`. The point is simple: if the target is chosen correctly, the engine can attempt certification; if the target is chosen incorrectly, the correct output is a named obstruction.

Decision	Action (what you declare)	If violated, output
A. Is the claim in COC?	Declare a phase observable and an admissible witness route (rotation / loop-degree / holonomy / multi-cycle) and state the target rule you are using.	If any ingredient is missing, the claim is outside COC: boundary-type obstruction (typically WOB-01).
B. Is coupling present?	If the model includes coupling constraints across phases (an integer constraint matrix χ or an equivalent coupling law), set <code>target_type=CST</code> and declare <code>baseline=coupledHaaronCS</code> T.	Using circle/torus baseline under coupling, or failing to compute/declare CST: WOB-06 .
C. Single-cycle circle observable?	If the witness route is a circle action or a single loop/defect witness producing an integer $k \in \mathbb{Z}$, set <code>target_type=circle</code> (or <code>embedded_circle</code> for subsystem-only claims) and <code>baseline=Haar</code> on $\mathbb{T}_P$.	If every admissible witness integer is 0: WOB-02 . If no loop/orbit exists in the declared domain: WOB-01 .
D. Multi-cycle torus observable?	If the witness route is a $\mathbb{T}^k$ action with response map induced by an integer matrix χ , compute the effective rank $r := \text{rank}(\chi)$. • If $r = k$, set <code>target_type=torus</code> and <code>baseline=Haar</code> on $\mathbb{T}^k$. • If $r < k$, set <code>target_type=subtorus_rank_r</code> and <code>baseline=Haar</code> on the induced r-torus (the χ-image).	Claiming full-torus forcing when the response rank is deficient (or rational locking collapses the sweep): WOB-10 .
E. Target-aware upgrades?	Any convergence or mixing upgrade must explicitly name the baseline it converges to. If the upgrade requires diffusion/mixing assumptions, record them and point to the declared evidence artifact.	Mixing/convergence claimed without an admissible hypothesis package: WOB-07 .

Allowed upgrades (summary).

Target	Baseline	Upgrades allowed (when justified)
Circle / embedded circle	Haar on $\mathbb{T}_P$	Stability margins (LCM/HMR), anti-collapse bounds, and convergence-to-Haar if a mixing model is declared.
Full torus $\mathbb{T}^k$	Haar on $\mathbb{T}^k$	Stability margins; rank-aware coverage; convergence-to-Haar if a mixing model is declared.
Subtorus rank r	Haar on induced r -torus	Same upgrades, but always relative to the induced target (never the ambient torus unless $r = k$).
CST (coupled target)	Coupled Haar on CST	All upgrades, but always relative to coupled Haar; any non-coupled baseline is ill-posed (CTP/CHT).

Universality-from-cyclic-structure schema

Theorem 1 (Universality-from-cyclic-structure (COC embedding schema)). *Fix a declared model class and let Σ denote the slice/subsystem on which a Synchré claim is posed. Assume the declaration provides:*

- **Phase observable (PO).** *A cycle-valued observable $\phi : \Sigma \rightarrow \mathbb{T}_P$ or $\Phi : \Sigma \rightarrow \mathbb{T}_P$ with the minimal regularity needed to define the intended winding/holonomy notion and to test nontriviality.*
- **Witness mechanism (WM).** *At least one admissible mechanism from the SFT menu, in one of the following structural forms:*
 - Action forcing. *An admissible continuous action $A : G \times D \rightarrow D$ on a domain $D \subseteq \Sigma$ with $G = S^1$ or $G = \mathbb{T}^k$, together with a continuous homomorphism $\chi : G \rightarrow H$ to the declared target group H such that ϕ or Φ is χ -equivariant on D in the declared tolerance (exactly or with a stated defect bound). The action-forcing mechanism is nondegenerate when the induced map on the declared target has nontrivial cyclic response (nonzero slope in the S^1 case; positive rank on the target coordinates in the $\mathbb{T}^k$ case).*
 - Embedded loop/holonomy. *An embedded loop $\gamma : S^1 \rightarrow \Sigma$ (or loop in an embedded subsystem) together with either a nonzero degree witness $\deg(\phi \circ \gamma) \neq 0$ or a nonzero holonomy multiple in a declared lift/connection model.*
- **Target rule (TR).** *A declared target type (circle/torus/coupled subtorus) and baseline convention on that target (Haar on the target, coupled Haar on CST when coupling is present).*

Then the declared context belongs to COC. Consequently, the SFT universality posture applies on this context: on the declared target, the correct engine output is either a Witness Stability Certificate (WSC) with explicit margins, or a finite Witness Obstruction (WOB) from the obstruction catalog with evidence.

Canonical cyclic families covered by the schema

The theorem is a *schema*: it reduces broad families to the same three-checklist declaration (PO/WM/TR).

The circle-action / periodic-forcing family admits a fully formal proof template in Theorem 2. Each line below is a structural criterion (not a physical promise): if the criterion holds in the adopted model class, it produces a COC embedding.

- **Periodic forcing / time-shift symmetry.** If the declared subsystem admits a continuous periodic parameter (identified with S^1) acting by reparameterization on states, and the phase observable is equivariant under that shift with nonzero slope on the declared phase circle, then the action-forcing mechanism applies.
- **Rotation / rigid cyclic response.** If the declared subsystem carries an admissible S^1 -action (or a torus action in the multi-cycle setting) and the observable is equivariant with nontrivial induced homomorphism on the declared target, then forcing applies on the corresponding orbit-image (full circle when slope is nonzero; subtorus/coset when rank is deficient and the target rule selects CST).
- **Wave-train / translation-phase reduction.** If a periodic direction induces an admissible S^1 -translation action on a domain and the phase observable is equivariant under translation with nonzero slope, the situation is again an action-forcing context.
- **Limit-cycle / periodic-orbit phase.** If the declared dynamics exhibits a periodic orbit equipped with a continuous phase coordinate (so that motion along the orbit defines an S^1 -action on that orbit) and the observable responds with nonzero winding along that cycle, then either action-forcing or an embedded loop degree witness certifies cyclic observability.
- **Defect loop / winding witness.** If the declaration provides a loop around a distinguished locus (a boundary component, defect core, or excluded set) on which the phase observable has nonzero winding (or nonzero holonomy multiple in a lift model), then the embedded witness mechanism applies.
- **Coupled drives / multi-cycle response.** If multiple cyclic inputs generate an admissible $\mathbb{T}^k$ -action and the observable is equivariant with integer-matrix response χ , then the correct target is the image subtorus (or the inferred coupled subtorus CST), and the forcing/obstruction dichotomy is posed on that target.

Canonical proof template: circle-action forcing (rotation / periodic forcing)

Theorem 2 (Circle-action forcing template). *Let Σ be a slice and let $A : S^1 \times \Sigma \rightarrow \Sigma$ be a continuous S^1 -action. Let $\phi : \Sigma \rightarrow \mathbb{T}_P$ be continuous. Assume there exists a continuous homomorphism $\chi : S^1 \rightarrow \mathbb{T}_P$ such that on a declared domain $D \subseteq \Sigma$,*

$$\phi(A(\theta, x)) = \phi(x) \oplus \chi(\theta) \quad \forall (\theta, x) \in S^1 \times D. \quad (1)$$

Write $\deg(\chi) \in \mathbb{Z}$ for the integer degree of χ .

- *If $\deg(\chi) \neq 0$, then for every $x \in D$ the orbit loop $\gamma_x : S^1 \rightarrow \Sigma$, $\gamma_x(\theta) = A(\theta, x)$ satisfies*

$$\deg(\phi \circ \gamma_x) = \deg(\chi) \neq 0,$$

and $\phi(D) = \mathbb{T}_P$. In particular, the declaration (Σ, ϕ) satisfies the COC checklist on D with: (PO) phase observable ϕ , (WM) action forcing and embedded loop witness, (TR) circle target with Haar baseline.

- If $\deg(\chi) = 0$, then ϕ is S^1 -invariant on D and factors through the orbit quotient $q : D \rightarrow D/S^1$: there exists $\bar{\phi} : D/S^1 \rightarrow \mathbb{T}_P$ with $\phi|_D = \bar{\phi} \circ q$. In this case, coverage on $\mathbb{T}_P$ can occur only through surjectivity of $\bar{\phi}$.

Proof. Every continuous homomorphism $\chi : S^1 \rightarrow S^1$ has the form [3, 1]: $\chi(e^{2\pi it}) = e^{2\pi imt}$ for a unique $m \in \mathbb{Z}$, and $\deg(\chi) = m$.

Fix $x \in D$. Define $\gamma_x(\theta) = A(\theta, x)$. By (1),

$$(\phi \circ \gamma_x)(\theta) = \phi(A(\theta, x)) = \phi(x) \oplus \chi(\theta).$$

Translation by the constant $\phi(x)$ does not change degree [1], hence $\deg(\phi \circ \gamma_x) = \deg(\chi)$.

If $\deg(\chi) \neq 0$, then χ is surjective onto $\mathbb{T}_P$. For any $y \in \mathbb{T}_P$, choose $\theta \in S^1$ with $\chi(\theta) = y \ominus \phi(x)$. Then $\phi(A(\theta, x)) = \phi(x) \oplus \chi(\theta) = y$, so $\phi(D) = \mathbb{T}_P$. The nonzero degree identity above supplies an embedded loop/holonomy witness, and the equivariance identity supplies an admissible action-forcing witness mechanism.

If $\deg(\chi) = 0$, then χ is trivial, so (1) gives $\phi(A(\theta, x)) = \phi(x)$ for all θ and all $x \in D$. Thus ϕ is constant along orbits in D and therefore factors through the orbit space [3]: there exists $\bar{\phi}$ with $\phi|_D = \bar{\phi} \circ q$. $\square$

Theorem 3 (Periodic forcing / rotation model as a COC embedding). *Let Σ be a slice in a declared model class and assume the declaration includes a continuous S^1 -action A and a continuous phase observable $\phi : \Sigma \rightarrow \mathbb{T}_P$ satisfying (1) on a domain D with $\deg(\chi) \neq 0$. Then the declared context belongs to COC on D . Consequently, on the correct circle target (with Haar baseline), the SFT engine posture applies: the correct output is either a WSC with explicit margins, or a finite WOB with evidence.*

Theorem 4 (Frame/gauge invariance (FEQ)). *Assume two declared contexts are related by a frame/gauge equivalence (FEQ) that preserves slice identification and conjugates the S^1 -action and phase observable on the declared domain D : there is a homeomorphism $F : D \rightarrow D'$ such that $F(A(\theta, x)) = A'(\theta, F(x))$ and $\phi' \circ F = \phi$. Then the degree identity $\deg(\phi \circ \gamma_x) = \deg(\chi)$ and the forcing/quotient dichotomy in Theorem 2 are invariant under the FEQ change of description.*

Canonical proof template: embedded loop / defect witnesses (degree and holonomy)

Theorem 5 (Embedded loop degree witness forces circle coverage). *Let Σ be a slice and let $i : S^1 \hookrightarrow \Sigma$ be a continuous embedding of a loop $E = i(S^1)$. Let $\phi : \Sigma \rightarrow \mathbb{T}_P$ be continuous and set $f := \phi \circ i : S^1 \rightarrow \mathbb{T}_P$. If $\deg(f) = k \neq 0$, then $f(S^1) = \mathbb{T}_P$. In particular, E is an embedded Synchré witness (ESW) for circle-type coverage with witness integer k .*

Proof. Identify $\mathbb{T}_P = \mathbb{R}/P\mathbb{Z}$ with the standard circle $\mathbb{T}_{2\pi}$ by linear rescaling; degree is invariant under this identification. Represent S^1 as $\mathbb{R}/2\pi\mathbb{Z}$ and lift f to a continuous $\hat{f} : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\hat{f}(t + 2\pi) = \hat{f}(t) + 2\pi k$. By continuity, $\hat{f}([0, 2\pi])$ is an interval whose endpoints differ by $2\pi k$. Hence its image modulo 2π meets every residue class, so f is surjective onto $\mathbb{T}_{2\pi}$, equivalently onto $\mathbb{T}_P$ [1]. $\square$

Theorem 6 (Holonomy witness forces circle coverage). *Let $i : S^1 \hookrightarrow \Sigma$ and $\phi : \Sigma \rightarrow \mathbb{T}_P$ be as above. Assume there exists a lift atlas along E with local lifts F and associated connection form α such that for the loop $\gamma := i$,*

$$\int_{\gamma} \alpha = mP \quad \text{for some integer } m \neq 0.$$

Then $\phi(E) = \mathbb{T}_P$, and the induced map $f = \phi \circ i$ has nonzero degree (indeed degree m) under the standard circle identification.

Proof. On each chart in the lift atlas, the connection form is $\alpha = dF$ for a local lift F . By additivity of integrals and the cocycle compatibility on overlaps, the loop integral of α equals the total increment of any lifted phase around one circuit of γ . Thus the lifted phase changes by mP on one circuit, i.e. the winding number is $m \neq 0$. A nonzero winding number implies nonzero degree, so surjectivity follows from Theorem 5. $\square$

Theorem 7 (Robustness of embedded witnesses under LCM/HMR). *Let $\tilde{\phi} : \Sigma \rightarrow \mathbb{T}_P$ be a perturbation of ϕ and write $\tilde{f} := \tilde{\phi} \circ i$.*

- **LDW stability (lift closeness).** *Suppose f and $\tilde{f}$ admit lifts $\hat{f}, \tilde{\hat{f}} : \mathbb{R} \rightarrow \mathbb{R}$ (after identifying $\mathbb{T}_P$ with $\mathbb{T}_{2\pi}$) such that*

$$\sup_{t \in [0, 2\pi]} |\hat{f}(t) - \tilde{\hat{f}}(t)| < \pi.$$

Then $\deg(\tilde{f}) = \deg(f)$.

- **HW stability (holonomy margin).** *Suppose α and $\tilde{\alpha}$ are connection forms on E representing the declared lift/connection model for ϕ and $\tilde{\phi}$, and*

$$\left| \int_{\gamma} (\alpha - \tilde{\alpha}) \right| < P/2.$$

Then the holonomy integer multiple is unchanged.

Consequently, any ESW with nonzero witness integer remains an ESW under perturbations satisfying the corresponding margin inequality (the engine records these as $\text{LCM} < \pi$ and $\text{HMR} < P/2$ in the Synchré Stability Modulus).

Proof. For LDW, the lift-closeness bound provides a homotopy through lifted maps that never crosses a branch cut, so winding number is constant along the homotopy and degree is preserved.

For HW, the integral bound forces $\int_{\gamma} \tilde{\alpha}$ to remain within $P/2$ of the integer multiple $\int_{\gamma} \alpha = mP$, hence the nearest integer multiple is still mP and the holonomy class is unchanged. $\square$

Canonical proof template: multi-cycle forcing and correct coupled targets

Theorem 8 ($\mathbb{T}^k$ -action forcing template: image subtorus and effective target rank). *Let Σ be a slice and let $A : \mathbb{T}^k \times \Sigma \rightarrow \Sigma$ be a continuous $\mathbb{T}^k$ -action. Let $\Phi : \Sigma \rightarrow \mathbb{T}_{\mathbf{P}}$ be continuous. Assume there exists a continuous homomorphism $\chi : \mathbb{T}^k \rightarrow \mathbb{T}_{\mathbf{P}}$ such that on a declared domain $D \subseteq \Sigma$,*

$$\Phi(A(u, x)) = \Phi(x) \oplus \chi(u) \quad \forall (u, x) \in \mathbb{T}^k \times D. \quad (2)$$

Define the (possibly lower-dimensional) response subtorus

$$\text{Im}(\chi) := \chi(\mathbb{T}^k) \subseteq \mathbb{T}_{\mathbf{P}}.$$

Then for every $x \in D$,

$$\Phi(A(\mathbb{T}^k, x)) = \Phi(x) \oplus \text{Im}(\chi).$$

In particular, the action-forcing mechanism forces coverage on the coset target $\Phi(x) \oplus \text{Im}(\chi)$. A universality claim in this setting is therefore well-posed only when its declared target agrees with this coset (or an explicitly declared coupled target contained in it).

Moreover, after choosing standard identifications $\mathbb{T}^k \cong \mathbb{R}^k/2\pi\mathbb{Z}^k$ and $\mathbb{T}_{\mathbf{P}} \cong \prod_{j=1}^m \mathbb{R}/P_j\mathbb{Z}$, there exists an integer matrix $M \in \mathbb{Z}^{m \times k}$ such that χ lifts to the linear map $t \mapsto \text{diag}(P_1/2\pi, \dots, P_m/2\pi)Mt$ on universal covers. The effective target rank is

$$r := \text{rank}(M) = \dim(\text{Im}(\chi)),$$

so the response rank determines the target rank: the correct target cannot have dimension larger than r .

Proof. Fix $x \in D$. For any $u \in \mathbb{T}^k$, (2) gives $\Phi(A(u, x)) = \Phi(x) \oplus \chi(u)$. Thus

$$\Phi(A(\mathbb{T}^k, x)) = \{\Phi(x) \oplus \chi(u) : u \in \mathbb{T}^k\} = \Phi(x) \oplus \chi(\mathbb{T}^k) = \Phi(x) \oplus \text{Im}(\chi).$$

This is exactly the coset identity.

For the rank statement, observe that continuous homomorphisms between tori are induced by integer matrices on fundamental groups. Under the standard coordinate identifications, a homomorphism $\chi : \mathbb{T}^k \rightarrow \mathbb{T}_{\mathbf{P}}$ admits a lift to a linear map on universal covers whose matrix has integer entries up to the fixed diagonal period scaling. The image $\text{Im}(\chi)$ is a closed connected subgroup of $\mathbb{T}_{\mathbf{P}}$, hence a subtorus, and its dimension equals the rank of the inducing integer matrix M . $\square$

Theorem 9 (Coupling forces a coupled target (CST) and a coupled Haar baseline). *Let $\Phi : \Sigma \rightarrow \mathbb{T}_{\mathbf{P}}$ be a multi-cycle phase observable and suppose the declaration includes a coupling law on the same domain D that restricts the image to a coset of a subtorus. In the integer-matrix atlas class this has the form*

$$A\Phi(x) = \Psi(s) \quad \forall x \in D \cap \Sigma_s, \quad (3)$$

for a fixed small-integer matrix A (declared or inferred) and a slice-dependent offset $\Psi(s)$. Define the coupled target on slice s by

$$\text{CST}_s := \{y \in \mathbb{T}_{\mathbf{P}} : Ay = \Psi(s)\}.$$

Then any coverage, stability, or baseline statement is meaningful only on CST_s . In particular:

- If a claim uses the full torus as its target when (3) holds, the claim is ill-posed and the correct engine output is the coupling-mismatch obstruction (WOB-06).
- The correct uniform baseline in coupled settings is Haar on CST_s (equivalently, the pushforward of Haar on the free-coordinate parameterization of CST_s).
- In the presence of a $\mathbb{T}^k$ -action equivariance (2), the forced target is the intersection of the response coset from Theorem 8 with the coupled target; in the target-correct regime these coincide because the declaration chooses CST_s as the target.

Proof. The set CST_s defined by (3) is, by construction, a coset of the subtorus $\ker(A)$. Thus it is a legitimate torus-type target equipped with its intrinsic Haar measure. A statement phrased on the full torus while the dynamics and observable are confined to CST_s changes the meaning of “coverage” and “uniformity” and therefore is not a well-posed SFT claim.

If, in addition, (2) holds, then by Theorem 8 the action-forced image on any orbit is a translate of $\text{Im}(\chi)$. When the declaration is target-correct, the target selection rule chooses the coupled target coset that actually contains the image; this is exactly the role of CST_s . $\square$

Theorem 10 (Multi-cycle forcing yields a COC embedding on the correct target). *Assume a declared context provides: (PO) a multi-cycle phase observable Φ , (WM) an admissible $\mathbb{T}^k$ -action A with equivariance (2) on a domain D , and (TR) the target rule selecting the correct coset target (the response image coset, or a coupled target CST_s when coupling constraints are present). Then the declared context belongs to COC on D . Consequently, on the correct target (with Haar or coupled Haar baseline as appropriate), the SFT engine posture applies: the correct output is either a WSC with explicit margins, or a finite WOB with evidence.*

Theorem 11 (Frame/gauge invariance (FEQ) in the multi-cycle setting). *Assume two declared contexts are related by a frame/gauge equivalence (FEQ) that preserves slice identification and conjugates the $\mathbb{T}^k$ -action and phase observable on the declared domain D : there is a homeomorphism $F : D \rightarrow D'$ such that $F(A(u, x)) = A'(u, F(x))$ and $\Phi' \circ F = \Phi$. Then the coset identity in Theorem 8 and the coupled-target correctness statement in Theorem 9 are invariant under the FEQ change of description.*

1 Data model and declared context

A declared context specifies:

- a slice or subsystem to be observed (global slice, embedded loop, or coupled target),
- a phase observable valued in a circle or torus, with a chosen period vector $\mathbf{P}$,
- a coupling hypothesis class (none, affine diagonal, or integer-matrix constraints),
- a perturbation model and desired margins (LCM/HMR),
- whether any statistical claims are intended (RNF or NDB+MXC).

Data consist of samples of phase values, either along a loop parameter (for winding) or over time (for drive increments and coupling inference).

2 Obstruction exhaustiveness by predicate exhaustion

Data are allowed to produce certificates only when the declared regime validators hold. When a validator fails, the correct output is a finite obstruction code with auditable evidence. The predicate exhaustion lemma suite fixes this boundary formally.

Predicate exhaustion: minimal completeness skeleton

SFT treats *validator predicates* as the formal boundary between certificate-claims and obstruction-claims.

The audit-facing specification of the predicate suite and its canonical mapping appears in `VALIDATOR.PREDICATE.SPEC.tex` (and in the runnable contract files under `engine/contracts`). Inside the declared regime, a witness stability certificate (WSC) is permitted only when the listed predicates hold on the declared target. When a predicate fails, the correct output is a witness obstruction (WOB) with a canonical code and a finite evidence checklist.

Validator predicates and canonical WOB codes

Write V_{XX} for the validator predicate associated to obstruction code **WOB-XX**. The intended reading is:

- if V_{XX} holds, then **WOB-XX** is *not* triggered by that failure mode,
- if V_{XX} fails, then **WOB-XX** is the correct obstruction code to report (with evidence).

The predicates are:

- **V₀₁ (subsystem/observable present)**. There exists an admissible subsystem (slice, orbit-domain, or embedded loop/defect domain) and a phase observable to a declared target $\mathbb{T}_P$ or $\mathbb{T}_P$. If this fails, report **WOB-01-no-subsystem**.
- **V₀₂ (nontrivial winding/holonomy)**. In an embedded mechanism, some certified loop carries a nonzero integer witness (degree or holonomy multiple). If this fails, report **WOB-02-trivial-winding**.
- **V₀₃ (no pinch on the declared target)**. On the declared target, the observable image does not miss a positive-length arc neighborhood required for SCL. If this fails, report **WOB-03-pinch**.
- **V₀₄ (no tear: definability/regularity)**. The hypotheses needed to define degree/holonomy in the adopted model class hold on the evidence domain (liftability/atlas consistency/continuity at the declared tolerance). If this fails, report **WOB-04-tear**.
- **V₀₅ (holonomy consistency)**. In holonomy-based mechanisms, the lift/connection interface is consistent at the declared tolerance so the holonomy integer is well-defined. If this fails, report **WOB-05-holonomy-defect**.
- **V₀₆ (coupled target correctness)**. In coupled settings, the selected target is the coupled subtorus coset CST_s and the baseline is coupled Haar on CST_s . If this fails, report **WOB-06-coupling-mismatch**.
- **V₀₇ (mixing claim justified)**. Any statistical convergence claim is accompanied by an admissible certificate hypothesis (RNF or NDB+MXC, with declared parameters). If this fails, report **WOB-07-mixing-not-justified**.
- **V₀₈ (orbit realized on the evidence domain)**. In action-equivariant mechanisms, the evidence domain contains a χ -active circle-type orbit segment needed for the forcing template. If this fails, report **WOB-08-orbit-collapse**.
- **V₀₉ (equivariance within tolerance)**. The equivariance defect is within the declared admissible tolerance on the evidence domain. If this fails, report **WOB-09-equivariance-defect**.
- **V₁₀ (no resonant locking for the intended target)**. The response rank and the inferred constraint structure do not collapse the intended target (full torus versus subtorus coset). If this fails, report **WOB-10-resonant-locking**.
- **V₁₁ (connected slice/component availability)**. The declared measurement domain intersects the orbit-bearing/loop-bearing component required by the chosen mechanism. If this fails, report **WOB-11-disconnected-slice**.

- **V₁₂ (slope-zero quotient surjectivity when needed).** When the induced homomorphism χ is trivial, the factor map through the quotient is surjective onto the declared target. If this fails, report WOB-12-slope-zero-quotient-nonsurjective.
- **V₁₃ (constraint hypothesis uniqueness).** In constraint/target inference, the declared data and tolerance determine a unique small-integer constraint hypothesis within the declared finite search. If this fails, report extttWOB-13-ambiguity.
- **V₁₄ (sufficient excitation for inference).** The sample count and excitation are sufficient to support stable inference of constraint rank and target rank at the declared tolerance. If this fails, report extttWOB-14-insufficient-excitation.
- **V₁₅ (stability margin away from resonance).** The best constraint candidates are separated from the declared tolerance by a positive stability margin; near-relations do not make the decision unstable. If this fails, report extttWOB-15-resonance-masking.

Lemma suite (each predicate failure forces its WOB code)

Lemma 1 (WOB-01 is forced by $\neg V_{01}$). *If no admissible subsystem/observable exists in the declared regime, then no WSC exists in that regime.*

Proof. Every WSC is defined on a declared evidence domain and records a target and observable. If such data do not exist in the declared regime, then no WSC object satisfying the regime validators can exist. $\square$

Lemma 2 (WOB-02 is forced by $\neg V_{02}$). *If every admissible loop mechanism in the evidence domain has zero degree and zero holonomy multiple, then no embedded-loop WSC with nonzero witness integer exists.*

Proof. Embedded-loop WSCs require an integer witness invariant that persists under the LCM/HMR margins. If all admissible loops have witness integer 0, then the mechanism cannot certify forced coverage by winding/holonomy and the obstruction is the trivial-winding class. $\square$

Lemma 3 (WOB-03 is forced by $\neg V_{03}$). *If the observable image misses a required arc neighborhood on the declared target, then SCL on that target cannot hold, hence no target-correct WSC exists.*

Proof. SCL asserts surjectivity onto the declared target. Missing an arc neighborhood contradicts surjectivity, so no WSC proving SCL on that target can exist. $\square$

Lemma 4 (WOB-04 is forced by $\neg V_{04}$). *If the adopted model class does not provide the regularity needed to define the witness invariant (degree/holonomy) robustly at the declared tolerance, then no WSC with certified margins exists.*

Proof. The WSC format requires a well-defined integer witness together with margins certifying its stability. If the witness cannot be defined in the adopted model class on the evidence domain, the certificate interface cannot be satisfied. $\square$

Lemma 5 (WOB-05 is forced by $\neg V_{05}$). *If holonomy consistency fails at the declared tolerance, then holonomy-based certification is ill-posed in that regime and no holonomy WSC exists.*

Proof. Holonomy WSCs require a connection/lift interface whose loop integral is well-defined modulo the period. If the interface is inconsistent, the holonomy integer is not stable or not defined, and the holonomy-defect obstruction is the correct output. $\square$

Lemma 6 (WOB-06 is forced by $\neg V_{06}$). *In a coupled setting, if the target is not declared as CST_s with coupled Haar baseline, then any claimed certificate is target-incorrect and no target-correct WSC exists without repairing the target.*

Proof. Coupling changes the maximal invariant set and the correct baseline measure. A certificate that uses the wrong target/baseline is not a WSC in the declared target-correct regime. $\square$

Lemma 7 (WOB-07 is forced by $\neg V_{07}$). *If a statistical convergence claim is asserted without an admissible hypothesis certificate (RNF or NDB+MXC), then the claim is outside the permitted regime and the correct output is the mixing-not-justified obstruction.*

Proof. SFT treats convergence as an additional artifact class whose hypotheses must be declared and auditable. Without an admissible hypothesis certificate, no convergence certificate can be issued, so the regime boundary is triggered. $\square$

Lemma 8 (WOB-08 is forced by $\neg V_{08}$). *If the χ -active orbit needed by the equivariant forcing template is absent from the evidence domain, then action-equivariant forcing cannot be invoked there and no forcing-based WSC exists on that domain.*

Proof. The forcing template certifies coverage by transporting phases along a realized orbit. If no such orbit occurs on the evidence domain, the mechanism is unavailable and the orbit-collapse class is the correct obstruction. $\square$

Lemma 9 (WOB-09 is forced by $\neg V_{09}$). *If equivariance defect exceeds the declared tolerance, then equivariance is not valid in the adopted model class at that tolerance and no equivariance-based WSC exists.*

Proof. Equivariance-based WSCs require the defect to be within tolerance so that the induced homomorphism controls orbit images. If the defect is too large, the required mechanism hypotheses fail. $\square$

Lemma 10 (WOB-10 is forced by $\neg V_{10}$). *If rank deficiency or resonant locking collapses the intended target, then no certificate can truthfully assert coverage of the larger target; the correct output is the resonant-locking obstruction (and the target must be revised).*

Proof. The correct target is determined by the response rank and by any valid integer constraints. If the intended target is larger than the reachable coset/subtorus, coverage of the larger target is impossible. $\square$

Lemma 11 (WOB-11 is forced by $\neg V_{11}$). *If the declared measurement domain does not meet the required orbit-bearing/loop-bearing component, then no certificate for that mechanism can be produced from that domain.*

Proof. All certificate mechanisms require evidence on a domain supporting the relevant orbit or loop. If the domain misses that component, the mechanism is inapplicable there. $\square$

Lemma 12 (WOB-12 is forced by $\neg V_{12}$). *If χ is trivial and the quotient factor map is not surjective onto the declared target, then SCL on that target is impossible and no target-correct WSC exists.*

Proof. When χ is trivial, equivariance forces the observable to factor through the quotient. If the induced factor is not surjective, the declared target cannot be covered. $\square$

Lemma 13 (extttWOB-13 is forced by egV_{13}). *If multiple incompatible small-integer constraint hypotheses fit within the declared tolerance and search, then the target is not uniquely determined in the declared regime and no target-correct certificate can be issued without additional excitation or a refined target declaration.*

Proof. A target-correct certificate requires a declared target type and, when constraints are part of the target selection, a determinate constraint hypothesis. If the declared data/tolerance admit more than one incompatible constraint hypothesis within the search, any choice is non-auditable without additional information, so the declared regime treats this as an explicit inference obstruction. $\square$

Lemma 14 (extttWOB-14 is forced by egV_{14}). *If the sample count or excitation is insufficient to support stable constraint inference at the declared tolerance, then inference is ill-posed in the declared regime and the correct output is extttWOB-14-insufficient-excitation.*

Proof. Stable inference requires enough samples and enough motion to distinguish candidate constraint hypotheses. When these are absent, any inferred constraints can change under small perturbations of the data, so no target-correct certificate can be issued in the declared regime. $\square$

Lemma 15 (extttWOB-15 is forced by egV_{15}). *If the best candidate constraints are within the declared tolerance by too small a margin, or if a near-relation lies just above tolerance, then the constraint decision is unstable at the declared tolerance and the correct output is extttWOB-15-resonance-masking.*

Proof. A positive stability margin is needed so that the inferred constraints persist under admissible perturbations. When the margin is too small or a near-relation is present, small perturbations can change the inferred target rank or constraint set, so the declared regime treats the situation as a resonance-masking obstruction. $\square$

Validator predicate specification (audit interface)

The completeness claim in the declared regime is enforced by a finite suite of *validator predicates*. Each predicate V_{XX} corresponds to one canonical obstruction code WOB-XX. For an issued WSC, the applicable predicates must hold; otherwise the correct output is a WOB with the associated code.

For audit-level synchronization between the paper and the runnable artifacts, the normative mapping is also recorded in:

- engine/contracts/VALIDATOR_PREDICATE_SPEC.md
- engine/contracts/validator_predicates.json
- engine/contracts/wob_codebook.json

Pred.	Canonical WOB code	Applicability (minimal)	Audit hook (artifact-level)
V01	WOB-01-no-subsystem	always	audit_wsc.py: requires a nonempty target and a predicate ledger entry for V01
V02	WOB-02-trivial-winding	mechanisms LDW/HW	audit_wsc.py: requires witness integer $k \neq 0$ and a V02 ledger entry

Pred.	Canonical WOB code	Applicability (minimal)	Audit hook (artifact-level)
V03	WOB-03-pinch	when arc-neighborhood coverage is asserted	recorded as an evidence reference (no recomputation in v0.1)
V04	WOB-04-tear	mechanisms LDW/HW	<code>audit_wsc.py</code> : requires declared perturbation model + margins and a V04 ledger entry
V05	WOB-05-holonomy-defect	when holonomy multiple is invoked	<code>audit_wsc.py</code> : requires <code>margins.HMR.P</code> and a V05 ledger entry
V06	WOB-06-coupling-mismatch	target is CST (or mechanism CMT)	CTP/CHT (coupling target correctness) enforced: requires <code>target.type=cst</code> and baseline declaring coupled Haar on CST
V07	WOB-07-mixing-not-justified	mechanism MXC	<code>audit_wsc.py</code> : requires a convergence certificate (RNF or NDB+MXC) and a V07 ledger entry
V08	WOB-08-orbit-collapse	mechanism AWT	<code>audit_wsc.py</code> : requires a V08 ledger entry with evidence reference
V09	WOB-09-equivariance-defect	mechanism AWT	<code>audit_wsc.py</code> : requires a V09 ledger entry with evidence reference
V10	WOB-10-resonant-locking	when full-torus target is asserted	recorded as an evidence reference (rank/constraint inference expands in SFT III)
V11	WOB-11-disconnected-slice	when component restriction is used	recorded as an evidence reference (connectedness tests expand with payloads)
V12	WOB-12-slope-zero-quotient-nonsurjective	when χ is trivial and quotient is invoked	recorded as an evidence reference (quotient surjectivity expands with payloads)
V13	WOB-13-ambiguity	when CST/constraint inference is invoked	recorded as a benchmark/tool evidence reference (no recomputation in v0.1)
V14	WOB-14-insufficient-excitation	when CST/constraint inference is invoked	recorded as a benchmark/tool evidence reference (no recomputation in v0.1)

Pred.	Canonical WOB code	Applicability (minimal)	Audit hook (artifact-level)
V15	WOB-15-resonance-masking	when CST/constraint inference is invoked	recorded as a benchmark/tool evidence reference (no recomputation in v0.1)

The WSC artifact format includes a lightweight predicate ledger field `validator.predicates` that records the applicable predicates as “holds” with short evidence references. The audit layer checks the parts that are readable from the artifact interface and rejects target-incorrect coupled claims.

Finite obstruction list is complete in the declared regime

Inside the declared regime, the obstruction interface is not an open-ended list of failure stories. It is a finite catalog whose members are forced by validator predicate failures. This is the sense in which the catalog is *complete*: every certificate failure inside the regime is recognized as exactly one canonical obstruction code with auditable evidence.

Canonical obstruction catalog (WOB)

WOB code	Meaning (declared regime)
WOB-01-no-subsystem	No admissible cyclic subsystem/observable is available in the measured model.
WOB-02-trivial-winding	All candidate loops have degree/holonomy multiple 0.
WOB-03-pinch	Image misses an arc neighborhood on the declared target (non-surjective).
WOB-04-tear	Regularity/definability needed for degree/holonomy fails at the declared tolerance.
WOB-05-holonomy-defect	Lift/connection consistency needed for holonomy certification fails.
WOB-06-coupling-mismatch	Coupling target/baseline is wrong (CST not computed or baseline not defined).
WOB-07-mixing-not-justified	Statistical convergence claim made without an admissible RNF or NDB+MXC hypothesis.
WOB-08-orbit-collapse	No χ -active circle-type orbit exists on the declared evidence domain.
WOB-09-equivariance-defect	Equivariance defect exceeds the declared admissible tolerance.
WOB-10-resonant-locking	Rank deficiency/rational dependence collapses the intended torus target.
WOB-11-disconnected-slice	Measurement domain misses the orbit/loop component required by the mechanism.
WOB-12-slope-zero-quotient-nonsurjective	χ is trivial and the induced quotient factor map is not surjective.

Theorem (catalog completeness within the declared regime)

Fix the declared regime described above (declared target type, admissible mechanism, admissible FEQ class, declared tolerances and margins, and in coupled settings the target-correctness rule selecting CST_s with coupled Haar baseline). Let WSC mean: *there exists a witness stability certificate satisfying the regime validators on the declared correct target*. Let WOB mean: *there exists a witness obstruction whose code is one of the fifteen canonical codes and whose evidence satisfies its checklist*. Then, inside the declared regime,

$$\neg \text{WSC} \iff \text{WOB}.$$

More concretely: a validator predicate failure $\neg V_{XX}$ forces the obstruction code WOB-XX , and if all validator predicates hold then the regime permits a WSC on the correct target.

Necessity (each obstruction code is realized by an audited atlas artifact)

Each canonical WOB code occurs in the shipped witness atlas as a boundary artifact with concrete evidence. The following obstruction artifacts are included under `engine/artifacts/` and are checked by the audit suite:

WOB code	Realizing atlas obstruction artifact
WOB-01-no-subsystem	WOB-ATLAS-0013.json
WOB-02-trivial-winding	WOB-ATLAS-0014.json
WOB-03-pinch	WOB-ATLAS-0015.json
WOB-04-tear	WOB-ATLAS-0016.json
WOB-05-holonomy-defect	WOB-ATLAS-0031.json
WOB-06-coupling-mismatch	WOB-ATLAS-0011.json
WOB-07-mixing-not-justified	WOB-ATLAS-0020.json
WOB-08-orbit-collapse	WOB-ATLAS-0019.json
WOB-09-equivariance-defect	WOB-ATLAS-0012.json
WOB-10-resonant-locking	WOB-ATLAS-0009.json
WOB-11-disconnected-slice	WOB-ATLAS-0017.json
WOB-12-slope-zero-quotient-nonsu rjective	WOB-ATLAS-0018.json

This realizability statement is what makes the catalog *necessary*: no code can be removed without deleting a genuine, audited boundary mechanism.

3 Target inference: full torus versus coupled target

CST inference theorem (small-integer constraints with stability)

The practical bottleneck for multi-cycle systems is deciding whether the correct target is the full torus or a coupled subtorus coset. SFT treats this as a bounded integer-relation problem with a stability margin and an obstruction fallback [9, 15].

Theorem 12 (CST inference with stability or obstruction). *Fix a sampling model and a tolerance corresponding to the declared perturbation scale. There is a conservative procedure that, given sampled phase vectors in $\mathbb{T}_{\mathbf{P}}$, outputs exactly one of:*

- **CST output.** A candidate small-integer constraint description $Ay = \Psi$ together with a quantitative residual bound, constraint rank $r := \text{rank}(A)$, and an explicit stability margin within the declared tolerance.
- **Full-torus output.** A declaration that no constraint was found in the declared bounded search, together with an explicit quantitative gap to the best relation in that search.
- **Obstruction output.** A canonical WOB code with evidence indicating why target inference is ill-posed at the declared tolerance (typically ambiguity, insufficient excitation, or resonance masking).

3.1 Identifiability and stability of inferred constraints

SFT treats target inference as a certificate-producing step. To prevent over-reading, the inference tool reports two quantitative objects:

- an *excitation summary* computed from the centered increment matrix $\Delta y_i := y_{i+1} - y_i$ after lift-unwrapping, including a numerical rank and a smallest nonzero singular value (a conservative persistent-excitation proxy),
- a *stability margin* m that measures how far the inferred constraint residuals sit below the declared tolerance.

The identifiability claim below is deliberately narrow: it is about *stability within the declared finite search space* (bounded coefficient size, bounded number of constraints), and it is conditional on a declared observation-noise budget.

Theorem 13 (Conservative identifiability in the declared small-integer regime (tool-matched)). *Fix a period $P > 0$ and a dimension $k \geq 1$. Fix declared inference parameters*

- coefficient bound $M \in \mathbb{N}$ (the tool enumerates $a \in [-M, M]^k \cap \mathbb{Z}^k \setminus \{0\}$),
- relation cap $L \in \mathbb{N}$ (the tool selects up to L non-multiple relation rows),
- tolerance fraction $\alpha \in (0, 1)$ and absolute tolerance $\tau := \alpha P$.

Assume the tool's basic regime checks are satisfied: the sample size satisfies $n \geq \max(25, 8k)$ and not all coordinates are weak at threshold 4τ (otherwise the tool returns WOB-14 or WOB-13).

Let $y_0, \dots, y_{n-1} \in \mathbb{T}_P^k$ be sampled phase vectors and suppose a lift choice produces unwrapped representatives $\tilde{y}_i \in \mathbb{R}^k$. Assume a bounded observation model

$$\|\tilde{y}_i - \tilde{y}_i^{\text{true}}\|_\infty \leq \varepsilon \quad \text{for all } i,$$

for some declared noise budget $\varepsilon \geq 0$.

Suppose the inference tool returns a CST candidate (A^, Ψ^*) with:*

- constraint rank $r = \text{rank}(A^*)$ and induced target rank $d := k - r$,
- joint residual bound $\rho := \max_i \|A^* \tilde{y}_i - \Psi^*\|_\infty$,
- stability margin $m := \tau - \rho > 0$,
- excitation flag `excitation_ok=true`, meaning the reported increment-rank proxy is at least d .

Let the tool's best single-relation residual be ρ_{best} with first non-multiple runner-up residual ρ_{ru} (both reported), and define the ranking gap

$$\Delta := \rho_{\text{ru}} - \rho_{\text{best}} \geq 0.$$

If $\varepsilon \leq \min(m/2, \Delta/2)$, then the inference outcome is stable within the declared finite search:

- the CST feasibility decision cannot flip under perturbations of size ε (constraints remain feasible at tolerance τ), and
- the tool's best-relation ranking used in row selection is stable, hence the returned (A^*, Ψ^*) does not change under such perturbations.

Equivalently, the tool's reported envelopes $\varepsilon \leq m/2$ and $\varepsilon \leq \Delta/2$ provide an auditable safety region for trusting the inferred constraint in the declared small-integer regime.

Interpretation. The theorem does not claim that all couplings are identifiable from all data. It says: within the declared finite coefficient regime, the tool reports two concrete safety envelopes: $\varepsilon \leq m/2$ for feasibility stability, and $\varepsilon \leq \Delta/2$ for stability of the best-relation ranking used in row selection. When excitation is weak or the margin is small, SFT treats the correct output as an inference obstruction (typically WOB-14, WOB-13, or WOB-15).

Theorem 14 (Null-model upper bound for spurious small-integer constraints (tool gate)). *Fix a period $P > 0$, a tolerance $\tau \in (0, P/2]$, and a dimension $k \geq 1$. Let $y_1, \dots, y_n$ be independent Haar samples on $\mathbb{T}_P^k$. Fix a nonzero integer vector $a \in \mathbb{Z}^k \setminus \{0\}$ and set*

$$z_i := a \cdot y_i \pmod{P} \in \mathbb{T}_P.$$

Then $(z_i)_{i=1}^n$ are independent Haar samples on $\mathbb{T}_P$. Moreover, the event

$$\exists \Psi \in \mathbb{T}_P \text{ such that } \max_i \text{dist}(z_i, \Psi) \leq \tau$$

is equivalent to all z_i lying in some arc of length 2τ , and satisfies the bound

$$\mathbb{P}\left(\exists \Psi : \max_i \text{dist}(z_i, \Psi) \leq \tau\right) \leq n \left(\frac{2\tau}{P}\right)^{n-1}.$$

If the inference search enumerates at most N distinct relation directions a (after collapsing integer multiples), then by a union bound the probability of finding any such relation within tolerance obeys

$$p_{\text{upper}} \leq N n \left(\frac{2\tau}{P}\right)^{n-1}.$$

If, in addition, the selected relation is required to satisfy the same tolerance on an independent holdout set of size n_{ho} , then the combined null-model bound multiplies:

$$p_{\text{upper,combined}} \leq \left(N n \left(\frac{2\tau}{P}\right)^{n-1}\right) \cdot \left(N n_{\text{ho}} \left(\frac{2\tau}{P}\right)^{n_{\text{ho}}-1}\right).$$

Interpretation. The bound is conservative and uses only Haar geometry and a union bound on the declared finite search. It supplies an auditable gate: a candidate constraint is accepted only when it is holdout-stable at tolerance τ and the reported $p_{\text{upper,combined}}$ is below a declared threshold.

Definition 4 (Target inference problem). *Given sampled phase vectors $y(t) \in \mathbb{T}_{\mathbf{P}}$, determine whether the image behaves as:*

- full-rank on $\mathbb{T}_{\mathbf{P}}$, or
- restricted to a coupled subtorus coset CST described by integer constraints $Ay = \Psi$.

4 Witness integer inference from samples

Loop samples (degree)

If a loop is sampled by phases $p_0, p_1, \dots, p_N \in [0, P)$, lift-unwrapping provides a lifted sequence $\tilde{p}_i$ with small jumps, and the degree estimate is

$$k \approx \frac{\tilde{p}_N - \tilde{p}_0}{P} \in \mathbb{Z}.$$

This is stable under the LCM margin: if two lifts remain within π , the integer cannot change [1].

Holonomy samples

If the observable is realized via a local lift atlas or a connection, holonomy inference uses integrated loop data to estimate the multiple m such that the loop integral equals mP . Stability is controlled by the HMR margin.

5 Certificate construction

A data-driven WSC records:

- the target (full torus or CST) and baseline rule,
- the witness integer k (degree) or m (holonomy multiple),
- stability margins (LCM/HMR),
- pointers to the data used and the inference method used to compute the fields.

A data-driven WOB is the correct output when inference detects a finite obstruction: no subsystem, trivial winding, coupling mismatch, resonance masking, and so on.

6 Convergence certificates from data

SFT treats statistical claims as artifacts:

- **RNF certificate.** Provide evidence that increments are nonresonant (up to declared tolerance) and state Haar convergence on the correct target [4, 5, 6].
- **NDB+MXC certificate.** Declare a noise-to-diffusion bridge and record a verified mixing inequality with a rate from the thresholds ledger [7, 8, 13].

References

- [1] Hatcher, Allen. *Algebraic Topology*. Cambridge University Press, 2002.
- [2] Munkres, James R.. *Topology*. Prentice Hall, 2000.
- [3] Bredon, Glen E.. *Introduction to Compact Transformation Groups*. Academic Press, 1972.
- [4] Walters, Peter. *An Introduction to Ergodic Theory*. Springer, 1982.

- [5] Einsiedler, Manfred and Ward, Thomas. *Ergodic Theory with a View Towards Number Theory*. Springer, 2011.
- [6] Kuipers, L. and Niederreiter, Harald. *Uniform Distribution of Sequences*. Wiley, 1974.
- [7] Pavliotis, Grigorios A.. *Stochastic Processes and Applications*. Springer, 2014.
- [8] Evans, Lawrence C.. *Partial Differential Equations*. American Mathematical Society, 2010.
- [9] Lenstra, Arjen K., Lenstra, Hendrik W., and Lovász, László. *Factoring polynomials with rational coefficients*. Mathematische Annalen, 261, (4), 515–534, 1982.
- [10] Hardy, G. H. and Wright, E. M.. *An Introduction to Diophantine Approximation*. Oxford University Press, 2008.
- [11] Katok, Anatole and Hasselblatt, Boris. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1995.
- [12] Ferguson, Helaman R. P. and Bailey, David H. *A polynomial time, numerically stable integer relation algorithm*. Technical report (PSLQ), 1999.
- [13] Meyn, Sean P. and Tweedie, Richard L. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2009.
- [14] Doeblin, Wolfgang. *Exposé de la théorie des chaînes simples constantes de Markov à un nombre fini d'états*. Revue Mathématique de l'Union Interbalkanique, 2, 77–105, 1937.
- [15] Ferguson, Helaman R. P. and Bailey, David H. *A Polynomial Time, Numerically Stable Integer Relation Algorithm*. RNR Technical Report RNR-91-032, July 14, 1992.