

Twist-Angle Control of Driven Electromagnetic Response in van der Waals Heterostructures

Alpha-Lattice: Pre-Registered Two-Rung Verification Ladder
(Rung 1: condensed matter / nanophotonics; Rung 2: driven-state mapping)

Chris Clifton (independent)

Revision: v2.18 Date: February 17, 2026

Abstract. This document specifies a falsifiability-first experimental ladder for testing whether twist angle θ can act as a control parameter for a repeatable drive-locked non-equilibrium electromagnetic regime (DLER) in van der Waals (vdW) heterostructures. The protocol is mechanism-agnostic and pre-registered with explicit gates, thresholds, null suites, and stopping rules designed to minimize false positives. Scope is limited to coupon-scale vdW heteropackets (Rung 1) and driven-response mapping under amplitude-modulated optical pumping (Rung 2). No mechanical, inertial, or propulsion claims are made.

0.0 Response Letter (aiXiv Review Response Edition)

Weakness 1: Limited novelty / unclear distinction from known twist-angle-tunable resonances.

Action: Replace DLER framing with an operationally defined drive-locked non-equilibrium electromagnetic regime (DLER). DLER is identified by drive-referenced phase locking (phase plateau + bounded phase jitter), a finite locking range under detuning, and reproducible mode-map stationarity that exceeds amplitude-only (linear) resonance behavior.

Weakness 2: Experimental feasibility / over-integrated stack.

Action: The protocol is now strictly two-rung and coupon-scale (graphene/hBN/graphene-class heteropackets). No macroscopic integration and no mechanical correlation probe are included.

Weakness 3: Literature context for the driven regime.

Action: Motivation and references are re-centered on driven moiré systems (Floquet engineering, nonlinear transport, and driven polaritons) in vdW heterostructures. Analogies to fusion-plasma regime transitions are removed from the core argument.

Review questions addressed.

(1) Novel theoretical development: Not claimed. The work is positioned as a preregistered, falsifiable map of θ -controlled drive-locking behavior.

(2) Fabrication feasibility: coupon-scale vdW stacking and twist/strain mapping are supported by established twistrionics methods; this protocol commits only to that level.

(3) Unique signature: DLER is defined by phase-locking metrics (phase stability + locking range) plus reproducibility gates, not by amplitude peaks alone.

[v2.11 additions] Reviewer items 1–7 and thesis-ready upgrades addressed in-line: explicit primary/secondary endpoint definitions (Sec. 2.2), phase measurement specification (Sec. 4.2.1), disambiguated thermal criteria (Sec. 4.2), mode-map correlation specification (Sec. 4.2.2), increased Stage A sample size (Appendix A), enumerated mandatory null suite (Sec. 6.0), S-flag measurement uncertainty propagation (Sec. 4.3), preregistration checklist (Appendix F), data products manifest (Appendix G), and decision-flow description (Appendix H).

[v2.12 additions] Three apparatus-level specification gaps closed: observable-under-demodulation declaration (Sec. 4.2.0), modulation frequency selection procedure with response-band definition (Sec. 4.1.1), and thermometry specification with null-differential verification (Sec. 4.2.3).

[v2.13 additions] Six remaining implicit degrees of freedom eliminated: (A) PLS windowing frozen to 1 s non-overlapping blocks (Sec. 2.2), (B) f_{mod} screening context locked — dedicated calibration coupon excluded from Stage A (Sec. 4.1.1), (C) S-flag “equivalent” loophole closed with named primary/fallback proxies (Sec. 4.3), (D) secondary active-region thermal proxy added as mandatory reported check (Sec. 4.2.3), (E) locking-range detuning protocol fully frozen with step/dwell/order (Sec. 4.2.4), (F) O-flag thresholds mapped explicitly to each declared detection channel (Sec. 4.2.5). Additionally: cross-coupon mode-identity criterion added to the two-coupon recurrence requirement (Sec. 2.1), and pump wavelength/fluence specification added (Sec. 4.1).

[v2.14 additions] Stage-A gate statistic updated from residual-sum-of-squares (RSS) to a likelihood-based ΔDev under a beta-binomial model to match the bounded, discrete PLS endpoint (Sec. 2.1).

[v2.15 additions] Phase sampling and filtering requirements made internally consistent by separating (optional) raw detector sampling from the demodulated phase series $\varphi(t)$; $\varphi(t)$ sampling is now tied to the post-demod bandwidth / lock-in time constant rather than f_{mod} (Sec. 4.2.1; Appendix G).

[v2.16 additions] Mode-map stationarity made implementable by splitting full-resolution mode maps (FMM; used for the $r \geq 0.80$ reproducibility criterion) from a fast, fixed-time stationarity probe (FSP; used inside the 30 s detuning dwells). Detuning evaluation now uses FSP correlation with an explicit stabilization sub-window (Sec. 4.2.2; Sec. 4.2.4; Appendix G).

[v2.17 additions] Canonical synchronization alignment added: (A) explicit injection-locking analog (Adler/Arnold tongue) mapped to the O-flag and detuning protocol; (B) optional lock-bandwidth characterization (Δf_{lock}) with injection-strength scaling for mechanistic alignment; (C) new Cross-Project Test Library (Appendix I) mapping each Alpha-Lattice gate to its closest standard test in oscillators/photronics, phase synchrony, and open-science preregistration.

[v2.18 additions] Motivation rewritten to foreground the driven non-equilibrium gap in moiré/vdW systems and the lack of a standardized cross-material driven-response assay. References updated to include moiré overview/reproducibility context and Moiré–Floquet (driven) review sources (Appendix E).

1.A Physical Motivation and Hypothesis

Physical motivation is the driven non-equilibrium gap in moiré materials. The twistrionics/moiré field has established extraordinary geometric control (twist angle, heterostrain, reconstruction boundaries) and has largely focused on equilibrium correlated ground states. By contrast, the driven regime—what these same systems do under controlled optical or electrical pumping—is systematically under-mapped and reported through non-comparable observables across material classes. Existing tools (ultrafast pump–probe, nonlinear optical spectroscopy, driven impedance/near-field probes, Floquet engineering) demonstrate rich non-equilibrium behavior, but there is no standardized, falsifiable sweep protocol that can be applied across vdW heterostructures to test whether geometry acts as a control parameter for drive-locking dynamics.

This preregistered two-rung ladder supplies that missing assay. It operationalizes DLER as a drive-referenced phase-lock signature (phase plateau + bounded jitter), a finite locking range under detuning (Arnold tongue analog), and mode-map stationarity under controlled modulation. Because the moiré geometry sets the coupled-mode landscape, a standardized “locking fingerprint” can be archived across systems to form a response–geometry dictionary: the same metrology that detects DLER also enables comparative materials characterization and inversion (inferring plausible geometric architectures from dynamical response). Scope remains limited to Rungs 1–2: coupon-scale fabrication and θ -dependent mapping of driven-response signatures under controlled drive.

Operational hypothesis (Rungs 1–2). After controlling for heterostrain, defects, and pump delivery, there exists a twist-angle interval (or discrete extrema) for which DLER can be reproducibly activated (flag $O = 1$). DLER is operationally defined by drive-referenced phase stability, a finite locking range

under small detuning, and mode-map stationarity, together with the pre-registered linewidth/Q/dissipation bounds in Sec. 4.

Null hypothesis. Once covariates are controlled, O activation and stability metrics show no statistically meaningful dependence on twist angle beyond noise and residual disorder.

1.0 Claims, Scope, and Non-Claims

Primary claim space (Rungs 1–2): twist angle θ is a control parameter for producing a repeatable drive-locked non-equilibrium electromagnetic regime (DLER) in a vdW coupon, with quantifiable spectral/impedance/mode-map signatures.

Non-claims:

- No claim that DLER is a new equilibrium phase; it is an operationally defined driven regime.
- No claim that EM LDOS modification implies inertial-mass modification.
- No reliance on HRP/SED to execute Rungs 1–2.
- No mechanical correlation, inertial, gravitational, or propulsion claim is included in this protocol.

2.0 Twist-Angle Program and Stage-Gating (Model Specified)

Small twist angles generate long-period moiré patterns; reconstruction and heterostrain can dominate and create sharp state boundaries. This motivates exploring a low-angle window without privileging any constant.

Stage A (coarse sweep): $\theta \in [0.25^\circ, 1.10^\circ]$ on a fixed grid (Appendix A).

2.1 Pre-registered Gate (“3 σ ”) [REVISED in v2.14: likelihood-based gate statistic; v2.13: mode-identity criterion added]

Let each coupon j at twist angle θ_j yield $k_j \in \{0, \dots, 60\}$ passing 1 s blocks (Sec. 2.2), with $PLS_j = k_j/60$. The Stage-A gate uses the coupon-level counts $\{k_j, \theta_j\}$ ($R(\theta)$ is still reported for visualization).

Null model (monotonic): monotone $p(\theta)$ fit by isotonic logistic regression, i.e., maximize the beta-binomial log-likelihood of k_j successes out of $n = 60$ with p constrained monotonic in θ . A single overdispersion parameter ρ is shared across θ ($\hat{\rho} = 0$ reduces to the binomial case).

Alternative model (single extremum): single-knot piecewise-linear model for logit $p(\theta)$ with one interior knot θ^* and one slope sign change, fit by maximizing the same beta-binomial likelihood over the Stage-A grid.

Test statistic: $\Delta Dev = Dev_{mono} - Dev_{knot} = 2(LL_{knot} - LL_{mono})$, where $Dev = -2 \cdot LL$ under the chosen likelihood.

Significance: permutation of θ labels across coupons within each fabrication batch (preserving each coupon’s k_j), ≥ 5000 draws (wild bootstrap permitted only if exchangeability is violated by batch constraints).

Pass condition: $p < 0.003$ and the extremum region must appear in at least two independent coupons at neighboring θ values. Otherwise Stage B is not executed.

[v2.13] Cross-coupon mode-identity criterion: The two-coupon recurrence requirement is satisfied only if the locked modes in both coupons can be identified as the same physical branch. Verification method: the spectral peak frequency (or, for microwave channel, the resonance frequency from S-parameter fit) of the locked mode in each coupon must fall within a preregistered spectral window. The window is defined as $\pm 15\%$ of the spectral peak measured in the first qualifying coupon at that

θ . If no spectral overlap exists, the modes are considered distinct branches, recurrence is not satisfied, and Stage B is not triggered for that θ . Additionally, if mode maps are available from both coupons, a cross-coupon map correlation of $r \geq 0.50$ is required as a secondary confirmation that the same spatial mode structure is present. This prevents the gate from conflating phase locking on different excitations as “recurrence.”

2.2 Primary and Secondary Endpoints [v2.11; REVISED in v2.13: PLS windowing frozen]

Primary endpoint (single metric for Stage-A gate): Phase-Lock Score $PLS(\theta)$, defined as the fraction of 1 s non-overlapping blocks within the 60 s acquisition window during which the drive-referenced phase jitter satisfies $\sigma\phi \leq 10^\circ$ rms and no 2π slips are detected within that block. $PLS = (\text{number of passing 1 s blocks}) / 60$. PLS is computed per coupon at each θ grid point; $R(\theta) \equiv$ median PLS across coupons at that grid point is reported for visualization. The Stage-A gate in Sec. 2.1 is computed on coupon-level counts (k_j) via a likelihood-based ΔDev statistic.

[v2.13 change] *The windowing that defines PLS is now frozen: 1 s non-overlapping blocks, 60 blocks per trial. This eliminates the latent degree of freedom in how $\sigma\phi$ is computed over the acquisition window and ensures PLS produces a quasi-continuous score (resolution $1/60 \approx 0.017$) rather than collapsing to a binary $\{0, 1\}$. The 1 s block length is chosen to be long enough to contain multiple cycles at all candidate f_{mod} values (≥ 100 Hz) while short enough to provide gradient information across the trial.*

Rationale: phase-lock quality directly tracks the defining signature of DLER (drive-referenced synchronization) and is less susceptible to amplitude-only artifacts than linewidth or Q proxies.

Secondary endpoints (reported but not used to trigger Stage B):

- Linewidth narrowing ratio: $\Gamma_{ON} / \Gamma_{OFF}$
- Quality factor: Q_{ON} (from ringdown or S-parameter fit)
- Near-field contrast: peak amplitude contrast in s-SNOM maps
- Impedance proxy: magnitude and phase of microwave reflection coefficient S_{11}

Mode-selection rule: At each θ , if multiple modes satisfy $O = 1$, select the mode with the highest PLS. This rule is frozen at preregistration; no post hoc mode switching is permitted. If no mode satisfies $O = 1$, $PLS = 0$ for that coupon/ θ .

3.0 Rung 1: Fabrication Roadmap (De-scoped)

3.1 Deliverable

Planar coupons comprising a twisted vdW heteropacket (e.g., graphene/hBN/graphene). No macroscopic tiles or multi-layer scaffold integration.

3.2 Mandatory Metrology

- Twist mapping: spatially resolved twist-angle map over the active area.
- Strain/heterostrain mapping (mandatory): co-registered with the twist map.
- Defect screening: optical + AFM spot checks; exclude bubble/void dominated regions.

4.0 Rung 2: Pumping, DLER Flag O, and Structural Stationarity Flag S

4.1 Pump Surrogate (fixed) [REVISED in v2.13: wavelength/fluence specification added]

Amplitude-modulated optical pumping delivered by fiber, with shutterable delivery and power metering at fiber output.

[v2.13] Pump optical specification: The following pump parameters are preregistered and frozen per campaign:

- **Wavelength:** Declared from the set {532 nm, 633 nm, 785 nm, 1550 nm} or a custom value with justification (e.g., resonance with a known electronic transition). The chosen wavelength is recorded and frozen before Stage-A data collection.
- **Power density range:** Declared as [P_min, P_max] in W/cm² at the sample plane. The protocol operates within this range; any run exceeding P_max is flagged as out-of-spec.
- **Spot size:** 1/e² beam diameter at the sample, measured and recorded per campaign.
- **Damage/degradation check:** Before each Stage-A campaign, a sacrificial region of the calibration coupon is exposed at P_max for 10× the trial duration (600 s). Post-exposure metrology (twist map + optical inspection) is compared to pre-exposure. If twist drift > 0.02° or visible damage is observed, P_max is reduced until the check passes. The damage-check result is recorded in Appendix G.

4.1.1 Modulation Frequency Selection [v2.12; REVISED in v2.13: screening context locked]

f_mod is selected from a pre-declared search list and frozen per Stage-A campaign. The procedure is:

1. **Candidate list:** f_mod ∈ {100 Hz, 1 kHz, 10 kHz, 100 kHz}. This spans from sub-thermal-diffusion rates to above typical phonon-polariton lifetimes in hBN-class systems ($\tau \sim 1\text{--}10$ ps implies intrinsic linewidths of order 10–100 GHz; f_mod is intentionally far below resonance frequency, probing the envelope dynamics of driven modes, not the mode frequency itself).
2. **Screening run:** Screening is performed on a **dedicated calibration coupon** fabricated from the same batch but **excluded from all Stage-A and Stage-B analysis**. The calibration coupon is tested at $\theta_{\text{ref}} = 0.60^\circ$ (mid-grid; preregistered and fixed across batches). Each candidate f_mod is tested with 3 ON/OFF cycles. The frequency yielding the lowest $\sigma\phi$ (not PLS) is selected as f_mod for the campaign. Rationale: selecting by noise floor rather than by PLS prevents the screening from optimizing directly for the primary endpoint. If two frequencies tie within measurement uncertainty, the lower frequency is chosen (conservative; slower modulation is harder to confuse with electronic pickup).
3. **Frozen value:** The selected f_mod, the screening data, and the selection rationale are recorded and deposited before Stage-A data collection begins.
4. **Nominal response band (for N2 null):** Defined as [f_mod/2, 2 × f_mod]. The N2 null control uses f_null ≥ 3 × f_mod (i.e., well outside this band). If f_mod = 1 kHz, then f_null ≥ 3 kHz.

[v2.13 change] Two modifications from v2.12: (i) screening is now performed on a dedicated calibration coupon excluded from analysis, eliminating the concern that the screening point contaminates the Stage-A dataset; (ii) frequency selection uses $\sigma\phi$ (noise floor) rather than PLS, so the screening does not optimize on the primary endpoint. θ_{ref} is fixed at 0.60° across batches for reproducibility.

4.2 DLER Flag O (binary, thresholded)

O = 1 only when all thresholds below are met; otherwise O = 0 (marginals default to O = 0).

Metric	Threshold (pre-registered)	Measurement	Notes

Linewidth narrowing	$\Gamma_{\text{ON}} \leq 0.5 \Gamma_{\text{OFF}}$; stable $\pm 10\%$ over trial	spectral FWHM	Relative to pump-OFF baseline
Quality factor	$Q_{\text{ON}} \geq 2 Q_{\text{OFF}}$ and $Q_{\text{ON}} \geq 50$	ringdown / S-parameter fit	Absolute floor prevents low-Q false pass
Drive-phase lock	$\sigma\phi \leq 10^\circ$ rms per 1 s block, no 2π slips, over 60 s trial	lock-in / heterodyne (Sec. 4.2.1)	Phase ref: pump modulation at f_{mod}
Locking range	O remains true at $\pm 0.5\%$ detuning from f_{mod} (see Sec. 4.2.4 for protocol)	frequency sweep per Sec. 4.2.4	FSP stationarity preserved at detuned points (Pearson $r \geq 0.80$ vs immediately preceding nominal; see Sec. 4.2.2 and Sec. 4.2.4).
Persistence	DLER maintained ≥ 60 s per trial	continuous monitoring	Short-dwell criterion
(i) Absolute thermal bound	$ T_{\text{ON}} - T_{\text{base}} \leq 2$ K	contact thermometry (Sec. 4.2.3)	Max excursion limit during pump-ON
(ii) ON/OFF trajectory match	$\max T_{\text{ON}}(t) - T_{\text{OFF}}(t) \leq 50$ mK over matched windows	paired ON/OFF thermometry (Sec. 4.2.3)	Differential criterion; rules out pump-heating as DLER cause
Mode reproducibility	map correlation $r \geq 0.80$ (Pearson; see Sec. 4.2.2)	re-imaging at equal timestamps	Same-state proxy; ROI mask frozen pre-run

[v2.11 change: thermal row split into (i) and (ii). v2.13 change: all numeric thresholds now appear explicitly in the threshold column; locking-range row references frozen detuning protocol in Sec. 4.2.4.]

4.2.0 Observable Under Demodulation [v2.12]

The drive-referenced phase $\phi(t)$ is extracted from one of the following pre-declared detection channels (chosen at preregistration and frozen):

(a) Near-field optical amplitude — s-SNOM tip-scattered signal demodulated at harmonic $n \geq 3$ of the tapping frequency, with the pump modulation envelope at f_{mod} serving as the slow reference. $\phi(t)$ is the phase of the f_{mod} Fourier component of the demodulated near-field amplitude time series.

(b) Microwave reflection — S_{11} magnitude and phase from a vector network analyzer or reflectometry bridge, with the pump modulation providing AM sidebands on the carrier. $\phi(t)$ is the phase of the lower sideband relative to the carrier, extracted via IQ demodulation.

(c) Photocurrent / photovoltage — Electrical contact to the heteropacket; lock-in input is the AC photocurrent at f_{mod} . $\phi(t)$ is the lock-in phase output directly.

The chosen channel, connector topology, and any intermediate amplification stages are recorded in the coupon metadata (Appendix G). If more than one channel is acquired simultaneously, only the pre-declared primary channel is used for O-flag evaluation; secondary channels are reported as exploratory.

Channel-to-threshold applicability is specified in Sec. 4.2.5.

4.2.1 Phase Measurement Specification [v2.11]

Phase reference chain: The pump modulation signal (square or sinusoidal at f_{mod}) serves as the master reference. A copy of this signal drives the reference input of the lock-in amplifier (or is used as the LO in a heterodyne detection chain). All phase measurements are therefore drive-referenced, not free-running.

Sampling (demodulated phase $\phi(t)$, used for PLS/slip detection): $\phi(t)$ is a baseband observable after lock-in / IQ demodulation; therefore its sampling rate is set by the post-demod bandwidth, not by f_{mod} . Let B_{ϕ} be the effective baseband bandwidth of $\phi(t)$ (vendor ENBW or measured 3 dB cutoff in the calibration run). Required: sample $\phi(t)$ at $f_{s,\phi} \geq \max(10 \times B_{\phi}, 200 \text{ Hz})$. Practical substitute if only the lock-in time constant τ is known: use $B_{\phi} \approx 1/(4\tau)$ and set $f_{s,\phi} \geq \max(2.5/\tau, 200 \text{ Hz})$. Record τ (if applicable), filter order, $f_{s,\phi}$, and the B_{ϕ} estimate in the session metadata.

Optional raw-signal capture (auditability only): If the pre-demod detector waveform $v_{\text{raw}}(t)$ (or I/Q) is stored, it must be sampled at $f_{s,\text{raw}} \geq 10 \times f_{\text{mod}}$ with documented anti-alias filtering and ADC specs. Raw storage is not required for the primary endpoint; storing only $\phi(t)$ is permitted.

Filtering: The lock-in low-pass filter settings (time constant τ , filter order, and any internal decimation) are declared and frozen per campaign. τ is chosen to minimize σ_{ϕ} while preserving slip detectability; specifically, in the pre-campaign calibration run a synthetic phase-step ($\geq 180^\circ$) applied to the instrument chain must produce an excursion that satisfies the slip criterion below.

Additionally, the effective baseband bandwidth of $\phi(t)$ after filtering (vendor ENBW or measured 3 dB cutoff) must satisfy $B_{\phi} \leq f_{\text{mod}}/5$ to ensure $\phi(t)$ reflects envelope dynamics rather than modulation-cycle leakage. No additional post hoc smoothing is applied to $\phi(t)$ before jitter and slip analysis.

Phase unwrap method: Standard arctangent unwrapping with $\pm\pi$ discontinuity detection. Unwrapped phase is used for all jitter and slip computations.

Slip criterion: A 2π slip is defined as any excursion where the unwrapped phase departs from its running 1 s median by more than $\pm\pi$ and does not return within 0.5 s. Runs exhibiting ≥ 1 slip fail the drive-phase lock criterion.

Instrumental phase noise floor: Before each measurement campaign, a calibration run is performed with a stable electronic test signal (synthesizer \rightarrow detector chain, bypassing the sample). The instrumental σ_{ϕ} floor and any 60 s drift are recorded. The pre-registered threshold ($\sigma_{\phi} \leq 10^\circ$) applies to the sample measurement after quadrature subtraction of the instrumental noise floor: $\sigma_{\phi,\text{sample}} = \sqrt{(\sigma_{\phi,\text{meas}})^2 - (\sigma_{\phi,\text{cal}})^2}$. If $\sigma_{\phi,\text{cal}} > 5^\circ$, the instrument chain must be improved before data collection proceeds. The calibration run must also record τ , filter order, $\phi(t)$ output sampling rate, and the estimated B_{ϕ} (ENBW or 3 dB cutoff) used to set $f_{s,\phi}$.

Calibration log: Instrumental calibration results (noise floor, drift, slip-free verification) are saved per session and included in the data products manifest (Appendix G).

4.2.2 Mode-Map Correlation Specification [v2.11]

Map modality (frozen at preregistration): Primary = near-field amplitude (s-SNOM optical amplitude channel, s_3 or s_4 harmonic demodulation). Fallback (if s-SNOM unavailable) = spatially resolved spectral peak frequency from micro-FTIR. Two spatial products are defined: (A) Full Mode Map (FMM), used for the mode reproducibility criterion; and (B) Fast Stationarity Probe (FSP), used inside the detuning dwells (Sec. 4.2.4).

Spatial resolution and ROI mask: FMM is acquired at ≤ 50 nm pixel pitch (s-SNOM) or ≤ 1 μm pitch (micro-FTIR), over an ROI mask defined by the intersection of: (a) the $S = 1$ structural-stationarity region and (b) the area where twist-angle metrology uncertainty is $\leq 0.03^\circ$. Edge pixels within 3 pixels of the mask boundary are excluded. The mask is frozen after pre-run metrology and before pump-ON measurement. FSP is a deterministic sentinel-point subset within the same frozen ROI: $N_{\text{sen}} = 16$ points arranged as a 4×4 evenly spaced grid within the ROI bounding box; any point falling outside the mask is snapped to the nearest in-mask pixel. The sentinel set (coordinates) is saved with the run metadata.

Correlation type and threshold: Pearson correlation r is computed (i) over all unmasked pixels for FMM-to-FMM comparisons (trial-1 vs trial-2 at matched conditions: same θ , same pump state, and matched acquisition timing relative to pump-ON start), and (ii) over the 16-element sentinel vector for FSP comparisons (nominal vs detuned within the detuning sequence). Threshold: $r \geq 0.80$ for the applicable comparison.

Nuisance regressors / registration: For FMM only, if preregistered, a single linear nuisance regressor (slow spatial drift estimated from the pump-OFF baseline map sequence) may be subtracted before correlation. The regressor form and estimation method are frozen at preregistration. No additional regressors are permitted post hoc. For FSP, no drift correction, registration, or regressors are permitted; failure of r is treated as loss of stationarity.

4.2.3 Thermometry Specification [v2.12; REVISED in v2.13: secondary thermal proxy added]

Sensor type: Calibrated thin-film resistance thermometer (Cernox or equivalent) or fiber-optic fluorescence thermometer, mounted on the coupon carrier substrate within $500 \mu\text{m}$ of the active area. If non-contact thermometry (IR pyrometry) is used as the primary channel, emissivity calibration against a contact sensor is mandatory and recorded.

Absolute accuracy: ± 0.5 K (traceable calibration), sufficient for the $|T_{\text{ON}} - T_{\text{base}}| \leq 2$ K criterion (i).

Differential resolution (for criterion ii): The 50 mK ON/OFF trajectory match requires differential noise < 50 mK over the comparison window. This is verified by a pre-campaign null thermometry run: two consecutive pump-OFF segments of matched duration are recorded, and $\max |T_{\text{run1}}(t) - T_{\text{run2}}(t)|$ is computed. If this null differential exceeds 25 mK (half the threshold), the thermometry chain must be improved (e.g., longer averaging, better thermal anchoring) before data collection proceeds.

Sampling rate: ≥ 1 Hz for the $T(t)$ time series.

Time alignment: ON and OFF thermal trajectories are aligned by trigger timestamp (pump shutter TTL edge), with a maximum allowed clock skew of ± 10 ms between thermometry and pump-state logs.

Calibration record: Sensor model, serial number, calibration date, null-differential result, and mounting geometry are recorded per campaign in Appendix G.

[v2.13] Secondary active-region thermal proxy (mandatory reported check): Because the contact thermometer is mounted on the carrier substrate (within 500 μm but not directly on the active moiré region), a secondary thermal proxy tied to the active region is acquired and reported for every run. The proxy is one of: (i) Raman thermometry — peak-shift of the graphene G or 2D band, calibrated against the contact sensor during the null thermometry run; or (ii) spectral-shift thermometry — temperature-dependent shift of the primary polaritonic/plasmonic resonance frequency measured during pump-OFF baseline, calibrated against the contact sensor. The chosen proxy is declared at preregistration. The secondary proxy does not gate O (it is reported, not thresholded), but any discrepancy > 5 K between the contact sensor and the secondary proxy triggers a mandatory annotation in the data release and investigation of local hotspotting. If neither Raman nor spectral-shift thermometry is available, this field is recorded as “not available” with justification, and the limitation is noted in the final publication.

4.2.4 Locking-Range Detuning Protocol [NEW in v2.13]

The locking-range criterion in the O-flag table requires that O remains true under $\pm 0.5\%$ detuning of f_{mod} . The following protocol is preregistered and frozen:

5. **Detuning values:** $f_{\text{low}} = f_{\text{mod}} \times 0.995$; $f_{\text{high}} = f_{\text{mod}} \times 1.005$. If these values are within synthesizer frequency resolution, round to the nearest settable frequency and record the actual detuning. Minimum absolute detuning: 0.5 Hz (to ensure the detuning exceeds synthesizer jitter at all candidate f_{mod} values).
6. **Protocol sequence (fixed order):** (i) nominal f_{mod} for 30 s (baseline confirmation; O must be 1), (ii) step to f_{low} for 30 s, (iii) return to nominal f_{mod} for 30 s (recovery confirmation), (iv) step to f_{high} for 30 s, (v) return to nominal f_{mod} for 30 s (final recovery). Total detuning sequence: 150 s.
7. Evaluation at each detuned point: Each 30 s dwell is partitioned into (i) a 10 s stabilization sub-window immediately following the frequency step (excluded from evaluation) and (ii) a 20 s evaluation sub-window. Over the evaluation sub-window, the drive-phase lock criterion ($\sigma\phi \leq 10^\circ$, no slips) is computed from $\phi(t)$. Mode stationarity under detuning is evaluated using the Fast Stationarity Probe (FSP) defined in Sec. 4.2.2: acquire one FSP during the nominal- f_{mod} dwell immediately preceding the detuned step and a second FSP during the detuned dwell, both within their respective 20 s evaluation sub-windows, and require Pearson $r \geq 0.80$ between the two sentinel vectors. Full Mode Maps (FMM) are not required at detuned points. The linewidth, Q, persistence, and thermal criteria are not re-evaluated at detuned points (they are assessed only at nominal f_{mod} during the primary 60 s trial). Rationale: the locking-range test probes whether phase synchronization is robust to small perturbations, not whether the full O-flag profile is invariant under detuning.
8. **Pass/fail:** O’s locking-range criterion passes if and only if phase lock and mode-map stationarity are maintained at both f_{low} and f_{high} . If either detuned point fails, the locking-range criterion fails and $O = 0$ for that trial.
9. **Recovery check:** If phase lock is lost at a detuned point but recovers at the subsequent nominal step, this is recorded as a “reversible unlock” (informative but does not rescue the failed criterion).

[v2.13 rationale] This protocol freezes the detuning step size, dwell time, step order, and evaluation criteria, eliminating the reviewer concern that “O remains true under detuning” was operationally

ambiguous. The fixed sequence also provides a natural internal control: recovery at nominal f_{mod} after detuning confirms the system was not irreversibly perturbed.

4.2.4.1 Optional Characterization: Lock Bandwidth Mapping (Arnold Tongue Slice) [NEW in v2.17]

Purpose (non-gating): Align the detuning criterion with the canonical injection-locking literature by measuring the lock bandwidth Δf_{lock} at a declared theta point. This provides a quantitative locking range rather than a two-point robustness check. Δf_{lock} does not gate O unless explicitly declared at preregistration as an added secondary endpoint.

Pre-registered sweep (recommended): At a fixed theta and fixed nominal f_{mod} , sweep detuning $f = f_{\text{mod}} + \Delta f$ over a symmetric range Δf in $[-\Delta f_{\text{max}}, +\Delta f_{\text{max}}]$ with step size Δf_{step} . Default values if not otherwise justified: $\Delta f_{\text{max}} = 0.02 f_{\text{mod}}$ (2%), $\Delta f_{\text{step}} = 0.001 f_{\text{mod}}$ (0.1%) with a minimum absolute step of 0.5 Hz. At each step, dwell 12 s; discard the first 4 s for stabilization; evaluate phase lock over the remaining 8 s using the Sec. 4.2.1 criterion ($\sigma_{\phi} \leq 10$ degrees, no slips). Acquire one FSP during the nominal point and one FSP during each detuned point (within the evaluation window) and require $r \geq 0.80$ to confirm the same mode is being tracked during the sweep.

Boundary definition: Define the lower and upper detuning limits (Δf_{low}^* , Δf_{high}^*) as the most negative and most positive detuning values for which both (i) phase lock and (ii) FSP stationarity pass. Define the lock bandwidth as $\Delta f_{\text{lock}} = \Delta f_{\text{high}}^* - \Delta f_{\text{low}}^*$. Report asymmetry ($|\Delta f_{\text{high}}^*| - |\Delta f_{\text{low}}^*|$) if present. Plotting Δf_{lock} versus injection strength (below) yields a one-dimensional slice of the Arnold tongue.

4.2.4.2 Optional Scaling Check: Δf_{lock} vs Injection Strength [NEW in v2.17]

Injection-strength proxy: Declare one controllable pump parameter as the injection-strength proxy (e.g., optical modulation depth m , pump fluence F at fixed spot size, or modulation-sideband amplitude measured at the detector input). Pre-register three levels (low/nominal/high) that remain below thermal/photodamage constraints. Measure Δf_{lock} at each level using the same sweep protocol. Expected qualitative behavior in an injection-locked regime: Δf_{lock} increases monotonically with injection strength; failure to increase is a mechanistic warning flag (does not by itself negate O, but must be reported).

Reporting: For each campaign that performs this optional characterization, include Δf_{lock} tables (per theta and per injection level), raw phase traces, and FSP vectors in the Appendix G data manifest. If the optional sweep is not performed, record 'not performed' with the reason (instrument limits, time constraints, or preregistration choice).

4.2.5 O-Flag Threshold Applicability by Detection Channel [NEW in v2.13]

The O-flag criteria in Sec. 4.2 apply to all declared channels, but not every metric is evaluable for every channel. The following mapping is preregistered:

O-flag Criterion	(a) s-SNOM	(b) Microwave S_{11}	(c) Photocurrent
Linewidth narrowing	Yes: spectral FWHM from s-SNOM spectra	Yes: 3 dB bandwidth from S_{11} fit	N/A unless spectral data acquired independently
Quality factor	Yes: from spectral peak fit	Yes: from S-parameter fit (primary method)	N/A unless spectral data acquired independently
Drive-phase lock	Yes: $\phi(t)$ from f_{mod} component of demodulated	Yes: $\phi(t)$ from IQ sideband extraction	Yes: $\phi(t)$ from lock-in phase output

	amplitude		
Locking range	Yes (per Sec. 4.2.4)	Yes (per Sec. 4.2.4)	Yes (per Sec. 4.2.4)
Persistence	Yes	Yes	Yes
Thermal (i) + (ii)	Yes (independent of channel)	Yes (independent of channel)	Yes (independent of channel)
Mode reproducibility	Yes: near-field amplitude maps	Conditional: requires spatially resolved S_{11} scan	N/A: no spatial map from photocurrent

Scoring rule: $O = 1$ requires that all **applicable** criteria pass. Criteria marked “N/A” for the declared channel are omitted from the O evaluation for that channel. However, if a channel produces N/A for ≥ 3 criteria (i.e., linewidth, Q , and mode reproducibility are all unevaluable), that channel is considered **insufficient for a standalone DLER claim**; the O flag is downgraded to $O = 0$ with annotation “insufficient channel coverage.” This prevents a minimal-instrumentation setup from trivially passing O by having most criteria be inapplicable.

[v2.13 addition: this table and scoring rule explicitly tie each O -flag metric to each detection channel, closing the concern that a lab using channel (c) could not evaluate half the criteria.]

4.3 Structural Stationarity Flag S (binary) [REVISED in v2.13: proxy loophole closed]

Goal: prevent pump-induced reconstruction/domain-wall motion from masquerading as a state effect.

$S = 1$ only if both:

(a) Pre/post stability: post-run twist/strain/moiré proxy maps correlate with pre-run maps at ≥ 0.90 and mean inferred twist drift $\leq 0.02^\circ$ over active area.

(b) In-run stability: a fast moiré proxy observable shows no monotonic drift exceeding a preregistered bound (default: $< 5\%$ in the proxy metric baseline value across a full run).

[v2.13] Named in-run proxies (replacing former “or equivalent” language): The in-run stability proxy is one of the following, declared at preregistration and frozen:

Primary proxy: Optical moiré imaging (wide-field reflection contrast at a wavelength sensitive to moiré periodicity). The 5% bound applies to the spatial-mean moiré contrast over the ROI.

Fallback proxy: Ellipsometric contrast (Δ/Ψ ratio from spectroscopic ellipsometry at a fixed wavelength). The 5% bound applies to the spatial-mean ellipsometric contrast over the ROI.

If neither is available: S is not evaluable for criterion (b), and the run is marked $S = 0$ (invalid for DLER claims) with annotation “in-run proxy unavailable.” The run may still be analyzed for structural dynamics. This hard rule prevents the open-ended “or equivalent” from reintroducing researcher choice.

If either criterion fails, set $S = 0$ and the run is invalid for any DLER claim (may be analyzed as structural dynamics).

[v2.11 — Reviewer Item 7] Measurement uncertainty propagation for S criterion: The 0.02° twist-drift bound applies to the drift estimate after subtracting measurement noise variance. Specifically, if the single-measurement twist uncertainty is $\sigma_{\theta, \text{meas}}$ (determined from repeated metrology on a reference coupon), then the drift significance is evaluated as: $\text{drift}_{\text{net}} = |\text{mean}(\theta_{\text{post}} - \theta_{\text{pre}})|$ and the bound is applied to $\text{drift}_{\text{net}}$ only when $\text{drift}_{\text{net}} > 2\sigma_{\theta, \text{meas}}$ (i.e., the drift must exceed twice the measurement uncertainty to be considered real). If $\text{drift}_{\text{net}} \leq 2\sigma_{\theta, \text{meas}}$, the drift is treated as

within noise and S is not failed on this criterion alone. The value of $\sigma_{\theta, \text{meas}}$ is recorded per fabrication batch in the data products manifest.

5.0 Stopping Rules, Blinding, Replication

5.1 Stopping Rules

- Stage B only if Sec. 2.1 gate passes.
- Rung 2 futility stop (preregistered): if Stage A fails the Sec. 2.1 gate across two independent fabrication batches, terminate and publish the full null suite and metrology.
- If repeated runs fail $S = 1$ (structural stationarity), terminate and publish Rungs 1–2 plus a structural-dynamics note.
- **Publish-null commitment:** Rungs 1–2 will be published regardless of outcome.

5.2 Blinding and Preregistration

Pump-state labels are randomized by a key-holder not involved in analysis. Scripts are frozen and time-stamped before unblinding (OSF). Any peeking is logged; results become exploratory.

5.3 Inter-lab Reproducibility

Any positive θ -extremum or DLER activation claim requires at least one independent replication attempt prior to strong interpretation.

6.0 Mandatory Null Suite [v2.11]

The following null controls are mandatory for every Stage-A grid point. Each null run follows the identical acquisition, storage, and analysis pipeline as the primary DLER run. Results for all nulls are reported regardless of primary outcome.

#	Null Control	Protocol	Rationale / What It Rules Out
N1	Pump unmodulated (CW)	Same average optical power, but modulation depth set to zero (CW illumination).	Isolates modulation-specific effects from steady-state heating or photodoping. If $O = 1$ under CW, the DLER claim is invalid.
N2	Wrong modulation frequency	f_{mod} shifted to $f_{\text{null}} \geq 3 \times f_{\text{mod}}$, outside the $[f_{\text{mod}}/2, 2 \times f_{\text{mod}}]$ response band (Sec. 4.1.1). Same average power and modulation depth.	Tests whether phase locking is frequency-specific rather than a broadband artifact.
N3	No-twist control coupon	θ well outside the ROI ($\geq 5^\circ$ or AB-stacked bilayer, i.e., $\theta = 0^\circ$ or 60°).	Establishes that any DLER signature requires moiré-scale modulation, not just a bilayer.
N4	Dead-stack control	Replace one active layer with an inert analog (e.g., replace graphene with amorphous carbon of matched thickness).	Tests whether the heterostructure's electronic coupling is necessary, not just optical absorption.

N5	Randomized pump-state labels	Blind analysis with shuffled ON/OFF labels (already specified in Sec. 5.2).	Direct false-positive rate calibration for the analysis pipeline. Explicitly tied to PLS and all secondary endpoints.
----	------------------------------	-----------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------

All null results are included in the final data release regardless of primary outcome. Any null that yields $O = 1$ triggers a mandatory investigation and is reported as an anomaly.

Appendix A: Sweep Schedule [v2.11]

Stage A: $\theta = \{0.25, 0.35, 0.45, 0.60, 0.80, 1.10\}^\circ$, with $n = 3\text{--}5$ coupons per grid point (minimum $n = 3$; target $n = 5$ where fabrication yield permits). Plus one dedicated calibration coupon at $\theta_{\text{ref}} = 0.60^\circ$ for f_{mod} screening (Sec. 4.1.1) and damage check (Sec. 4.1); this coupon is excluded from all Stage-A and Stage-B analysis.

Hierarchical aggregation rule: If coupon yield forces $n < 3$ at any grid point, the Stage-A gate is still computed on coupon-level k_j via the likelihood-based isotonic-vs.-knot comparison in Sec. 2.1; unequal coupon counts across θ automatically downweight sparse angles through the binomial trial counts ($n = 60$ per coupon). No linear-Gaussian mixed-effects model is used for the primary gate.

Stage B (conditional): smallest region containing extremum; default ROI $\theta \in [0.35^\circ, 0.55^\circ]$ with step 0.02° , $n \geq 7$ coupons.

Appendix E: References (selected, load-bearing)

- Bistritzer, R. & MacDonald, A. H. Moiré bands in twisted double-layer graphene. Proc. Natl Acad. Sci. USA 108, 12233–12237 (2011).
- Cao, Y. et al. Correlated insulator behaviour in magic-angle graphene superlattices. Nature 556, 80–84 (2018).
- Yoo, H. et al. Atomic and electronic reconstruction at the vdW interface in twisted bilayer graphene. Nat. Mater. 18, 448–453 (2019).
- Carr, S. et al. Electronic-structure methods for twisted moiré layers. Nat. Rev. Mater. 5, 748–763 (2020).
- Andrei, E. Y. et al. The marvels of moiré materials. Nat. Rev. Mater. 6, 201–206 (2021).
- Lau, C. N. et al. Reproducibility in the fabrication and physics of moiré materials. Nature 602, 41–50 (2022).
- Rodriguez-Vega, M. et al. Low-frequency and Moiré–Floquet engineering: A review. Ann. Phys. 435, 168434 (2021).
- Oka, T. & Kitamura, S. Floquet engineering of quantum materials. Annu. Rev. Condens. Matter Phys. 10, 387–408 (2019).
- Bukov, M., D’Alessio, L. & Polkovnikov, A. Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering. Adv. Phys. 64(2), 139–226 (2015).
- Basov, D. N. et al. Polaritons in van der Waals materials. Science 354, aag1992 (2016).
- Woessner, A. et al. Highly confined low-loss plasmons in graphene–BN heterostructures. Nat. Mater. 14, 421–425 (2015).
- Nosek, B. A. et al. The preregistration revolution. PNAS 115, 2600–2606 (2018).
- Adler, R. A study of locking phenomena in oscillators. Proc. IRE 34, 351–357 (1946).
- Pikovsky, A., Rosenblum, M. & Kurths, J. Synchronization: A Universal Concept in Nonlinear Sciences. Cambridge Nonlinear Science Series 12, Cambridge University Press (2001).

Appendix F: Preregistration Checklist [v2.11; updated v2.12, v2.13]

Item	Specification / Reference
Primary endpoint	Phase-Lock Score PLS(θ); 1 s non-overlapping blocks — Sec. 2.2
Secondary endpoints	Linewidth ratio, Q, near-field contrast, S_{11} — Sec. 2.2
Stage-A gate	Isotonic (monotone) vs. single-knot (one extremum) using likelihood-based ΔDev + permutation, $p < 0.003$, 2-coupon recurrence + mode-identity — Sec. 2.1
Stage-A sample size	$n = 3\text{--}5$ coupons per θ ; hierarchical fallback if $n < 3$ — App. A
DLER flag O thresholds	Linewidth, Q, phase lock, locking range, persistence, thermal (i)+(ii), mode corr. — Sec. 4.2
S-flag thresholds	Pre/post corr. ≥ 0.90 , drift $\leq 0.02^\circ$, in-run drift $< 5\%$ — Sec. 4.3
Null suite	N1–N5: CW, wrong f_{mod} , no-twist, dead-stack, label shuffle — Sec. 6.0
Stopping rules	Gate fail $\times 2$ batches \rightarrow terminate; S-fail \rightarrow terminate — Sec. 5.1
Blinding	Pump-state labels randomized; scripts frozen pre-unblinding — Sec. 5.2
Publish-null commitment	Full Rungs 1–2 published regardless of outcome — Sec. 5.1
Data release	All raw data, calibration logs, analysis scripts deposited on OSF — App. G
Mode-selection rule	Highest PLS mode; frozen at preregistration — Sec. 2.2
Phase measurement spec	Reference chain, sampling, filtering, unwrap, slip criterion, cal — Sec. 4.2.1
Mode-map correlation spec	FMM vs. FSP definitions, Pearson r , ROI mask, sentinel set, nuisance regressor limits — Sec. 4.2.2
Observable declaration [v2.12]	Detection channel (a)/(b)/(c) declared and frozen — Sec. 4.2.0
f_{mod} selection [v2.12]	Candidate list, screening on cal coupon ($\sigma\phi$), frozen value, response band — Sec. 4.1.1
Thermometry spec [v2.12]	Sensor type, accuracy, null-differential run, sampling, alignment — Sec. 4.2.3
PLS windowing [v2.13]	1 s non-overlapping blocks, 60 blocks per trial — Sec. 2.2

Locking-range protocol [v2.13]	Fixed sequence, 30 s dwells partitioned into 10 s stabilization + 20 s evaluation; phase-lock + FSP stationarity eval at detuned points — Sec. 4.2.4
Channel-threshold map [v2.13]	Per-channel O-flag applicability + insufficient-coverage rule — Sec. 4.2.5
Mode-identity criterion [v2.13]	$\pm 15\%$ spectral window + cross-coupon $r \geq 0.50$ — Sec. 2.1
S-flag named proxies [v2.13]	Primary: optical moiré imaging; fallback: ellipsometric contrast — Sec. 4.3
Pump optical spec [v2.13]	Wavelength, power density, spot size, damage check — Sec. 4.1
Secondary thermal proxy [v2.13]	Raman or spectral-shift thermometry; reported check, not gated — Sec. 4.2.3

Appendix G: Data Products Manifest [v2.11; updated v2.12, v2.13]

The following files are saved per run and deposited in the OSF repository:

Data Product	Format	Notes
Raw phase time series	Binary / HDF5	Demodulated $\varphi(t)$ sampled at $f_{s,\varphi} \geq \max(10 \times B_\varphi, 200 \text{ Hz})$ (or $\geq \max(2.5/\tau, 200 \text{ Hz})$ if only τ is known); record τ/B_φ per session.
Raw spectral data	HDF5	Full spectra per spatial point (s-SNOM or micro-FTIR)
Near-field / mode maps	TIFF + metadata JSON	Amplitude and phase maps, per trial, with pixel coordinates
Twist / strain maps	HDF5	Pre-run and post-run, with uncertainty estimates
Thermometry logs	CSV	$T(t)$ for ON and OFF runs, matched timestamps
Pump power log	CSV	Fiber-output power vs. time, including modulation waveform
Instrumental calibration	CSV + notes	Phase noise floor, drift, slip-free verification per session
Coupon metadata	JSON	Fabrication batch, θ target, metrology summary, S-flag result

Analysis scripts	Python / Julia	Frozen, time-stamped, version-controlled (Git hash recorded)
Blinding key	Encrypted file	Released only after analysis scripts are frozen
f_mod screening data [v2.12]	CSV + JSON	All candidate frequencies tested, $\sigma\phi$ per frequency, selection rationale
Null thermometry run [v2.12]	CSV	Paired pump-OFF thermal trajectories, null-differential result
Observable declaration [v2.12]	JSON	Detection channel choice, connector topology, amplification chain
Damage-check result [v2.13]	JSON + images	Pre/post metrology at P_max, pass/fail, any P_max reduction
Secondary thermal proxy [v2.13]	CSV + calibration	Raman or spectral-shift T(t), calibration against contact sensor
Detuning sequence log [v2.13]	CSV	f_low/f_high actual values, $\sigma\phi$ and FSP-corr at each step, pass/fail
Fast Stationarity Probe (FSP) values [v2.16]	CSV + metadata JSON	16-point sentinel vector per dwell (nominal/detuned), coordinates, acquisition timing; used for detuning stationarity r
Lock-bandwidth sweep data (Delta f_lock) [v2.17 optional]	CSV + HDF5	Detuning sweep points, dwell windows, pass/fail per step, delta f_low*, delta f_high*, Delta f_lock; injection-strength level metadata

Appendix H: Decision Flow [v2.11; updated v2.12, v2.13]

The following decision flow governs the progression through the protocol:

Step 0 — Campaign Setup [v2.12/v2.13]. Declare and freeze: detection channel (Sec. 4.2.0), pump wavelength and power density range (Sec. 4.1), in-run S-flag proxy choice (Sec. 4.3), secondary thermal proxy choice (Sec. 4.2.3). Perform f_mod screening on calibration coupon (Sec. 4.1.1). Perform null thermometry verification (Sec. 4.2.3). Perform pump damage check (Sec. 4.1). Deposit all pre-campaign declarations on OSF.

Step 1 — Rung 1 Fabrication + Metrology. Fabricate coupon batch for Stage-A θ grid. Perform mandatory twist, strain, and defect metrology. Coupons failing defect screening are excluded.

Step 2 — Rung 2 Stage A: Coarse Sweep. For each θ grid point ($n = 3-5$ coupons): run primary DLER protocol (60 s trial + 150 s detuning sequence per Sec. 4.2.4) and full null suite (N1–N5). Compute O flag and S flag per run. Record PLS(θ) as primary endpoint. If $S = 0$, run is excluded from DLER analysis (retained for structural-dynamics analysis).

Step 3 — Stage-A Gate Evaluation. Apply preregistered isotonic-vs.-knot test (Sec. 2.1) to the coupon-level counts $\{k_j\}$ (k_j successes out of 60; equivalently $PLS_j = k_j/60$), using ΔDev and the permutation scheme. Verify cross-coupon mode identity at candidate extremum

points (Sec. 2.1). If $p < 0.003$ AND 2-coupon recurrence with mode identity is satisfied: PASS → proceed to Stage B. If FAIL: repeat with second independent fabrication batch. If FAIL $\times 2$ batches: TERMINATE → publish full null suite, metrology, and null result.

Step 4 — Stage B: ROI Refinement. Fine-grid sweep within extremum ROI ($\Delta\theta = 0.02^\circ$, $n \geq 7$ coupons). Full null suite at each point. Characterize DLER boundaries and reproducibility.

Step 5 — Replication. Any positive result requires at least one independent replication attempt (Sec. 5.3) before strong interpretation.

Step 6 — Publication. Publish Rungs 1–2 results, all null data, metrology, calibration logs, and analysis scripts regardless of outcome (publish-null commitment).

Appendix I: Cross-Project Test Library (Canonical Analogs) [NEW in v2.17]

This appendix maps each Alpha-Lattice gating construct to its closest established analog in adjacent fields (nonlinear oscillators / photonics injection locking, phase synchrony metrics, and open-science falsifiability controls). The mapping is informational and intended to (i) reduce reviewer ambiguity, (ii) import known failure modes, and (iii) clarify which measurements are gating versus supportive.

Alpha-Lattice element	Closest canonical analog	Primary observable / expected behavior	What it rules out (artifact class)
O-flag phase lock criterion (σ_ϕ , slips)	Injection locking / Adler phase dynamics	Drive-referenced phase stability; no phase slips while driven	PLL/demodulation artifacts; drift masquerading as coherence
+0.5% detuning survival (Sec. 4.2.4)	Locking range robustness check	Lock persists under small detuning; recovers at nominal	Single-frequency coincidences; ultranarrow resonant tuning artifacts
Optional Δf_{lock} sweep (Sec. 4.2.4.1)	Arnold tongue boundary mapping	Quantified lock bandwidth; boundary where phase lock fails	Ambiguous locking claims; post hoc detuning choices
PLS = $k/60$ windowed pass fraction	Phase synchrony metrics (PLV/PLI-style windowing)	High fraction of fixed windows pass a fixed lock criterion	Cherry-picked stable intervals; post hoc windowing
Mode-map correlation (FMM) and FSP $r \geq 0.80$	Mode identity / stationarity checks	Same spatial/spectral mode tracked across time and detuning	Mode hopping; scan drift; re-pointing artifacts
Stage-A gate statistic (isotonic vs knot; permutation)	Pre-registered extremum / change-point detection	Evidence for a single θ ROI extremum beyond monotone null	Look-elsewhere effects; p-hacking across θ grids
Mandatory null suite (CW, wrong f, dead stack, etc.)	Negative controls / sham conditions	O collapses under nulls while holding under true condition	Instrument cross-talk; analysis bias; unintended feedback loops
Preregistration + publish-null rule	Registered reports / OSF preregistration discipline	Frozen analysis; all outcomes disclosed	Outcome switching; selective reporting; threshold tuning