

Auditing Record State in Constraint-Sweep Hysteresis Tests

A Loop-Scaling Diagnostic for Record-Driven versus Confound-Induced Path Dependence

Jeremy C. Jones

HoldingLight LLC — ORCID 0009-0007-2515-3774 — universalcollapse.com

Series: Universal Collapse Theory — Methods Paper • Version v1.0 • 2026-05 • CC BY 4.0

Abstract

Principle. Constraint-sweep hysteresis is observed across the biological, physical, and social sciences. The Universal Collapse Theory library treats loop area scaling with audited record state R as the S3 stabilization signature, with the single-hysteron base case $A(R) = 4\theta_0 + 4\alpha R$ proved in Jones 2026e. The S3 claim is specifically about the rate-independent component of loop area; rate-driven contributions, including record-modulated rate lag (loops whose width depends on sweep rate via R -coupled relaxation), are real but distinct phenomena and require a different diagnostic. Standard hysteresis-summary statistics describe the loop's geometry but do not by themselves test whether observed path dependence has a record-driven quasi-static component.

Result. We propose a loop-scaling diagnostic in two regimes. In discrete-state systems admitting a single-hysteron or Preisach-density description, independently audited R is compared against the model-predicted loop-area curve, with $dA/dR = 4\alpha$ in the single-hysteron base case. In continuous-state systems, the diagnostic uses a predeclared scalar response coordinate $y = f(x)$ and a sweep-rate audit to estimate the quasi-static loop area $A_{qs}(R) := \lim_{v \rightarrow 0} A(R, v)$. An S3-positive result requires $A_{qs}(R)$ to match an independently specified R -dependent prediction $A_{pred}(R)$ within declared tolerance; positive but R -independent A_{qs} indicates intrinsic static hysteresis without audited record coupling, and $A_{qs} \approx 0$ with positive finite-rate area indicates rate lag or record-modulated rate lag (a real but distinct phenomenon for which a different diagnostic, on the fitted $\tau(R)$ curve, would be appropriate). In both regimes, the diagnostic is valid only when R , the response coordinate, the formalism, and the predicted slope or curve are specified without circular use of the loop area being explained. The protocol is demonstrated on five synthetic discrete-state regimes (genuine, rate-driven, instrument hysteresis, minor-loop violation, circular- R) and three synthetic continuous-state regimes (genuine bistable, pure rate lag as negative control, combined record-coupling and rate lag). The rice transcriptome heat-stress/recovery dataset (Jones 2026h, GEO GSE74793) is specified as a step-response-loop application target for which Steps 1, 2, and 4 of the protocol are executable but Step 3 (sweep-rate audit) is not, since the dataset uses a fixed temperature-jump protocol rather than varied ramp rates.

Keywords: hysteresis; path dependence; Preisach model; record state; loop area; viscoelastic response; sweep-rate extrapolation; circular inference; double dipping; scalar response coordinate; reproducibility.

Review target

The review target is split into two acceptance levels, paralleling the structure of Auditing Independence (Jones 2026f) and Auditing Constraint Asymmetry (Jones 2026g).

Specification acceptance

Acceptance of v1.0 means the proposed audit structure is well-defined: independently specify R from non-loop channels with explicit feature-disjointness rules in omics settings; predeclare the formalism (single hysteron, Preisach density, continuous Preisach analog, bistable basin-geometry / Kramers-type model) and the scalar response coordinate $y = f(x)$ before loop area is computed; estimate the predicted slope c (or, in the continuous regime, the predicted A_{static} curve) from independent calibration, mechanistic parameter estimates, simulation truth, or a held-out training loop family disjoint from the test loop family; verify major-loop validity from pretest threshold calibration or held-out switching measurements rather than from the loop being tested; report sweep rate(s) and the rate-extrapolated quasi-static loop area where applicable; compare the audited- R loop-area curve $A_{\text{qs}}(R)$ (continuous regime) or $A(R)$ (discrete regime) against the predicted slope or curve. The structural specification can be accepted independently of whether the diagnostic outperforms standard hysteresis-summary statistics on any specific dataset.

Performance acceptance

Acceptance of the protocol's performance — a stronger claim than acceptance of the specification — requires the frozen synthetic demonstrations of §4 and a downstream real-data application that satisfies the dataset prerequisites listed in §5. v1.0 presents the synthetic demonstrations at release quality. Real-data execution against the rice transcriptome heat-stress/recovery dataset (Jones 2026h, GSE74793) is the work of a future T16 demonstration paper, with the limitation that the rate audit (Step 3) is not executable on that dataset.

Rejection criterion

The protocol should be revised or rejected if R , the formalism choice, the response coordinate y , or the predicted curve A_{pred} cannot be specified without circular use of the loop being explained; if controlled synthetic cases fail to distinguish record-driven hysteresis from the named confounds (rate-driven apparent loops, instrument hysteresis, minor-loop violation, circular R -fit, additive offset confounds, record-modulated rate lag); if the response coordinate used to compute loop area is selected from the same data used to test loop-area scaling, unless a held-out validation design is used; or if the continuous-regime formulation cannot detect record-modulated rate lag (a system with positive finite-rate loop area but $A_{\text{qs}}(R) = 0$) when sweep-rate extrapolation is available. Acceptance does not require that all confounds be disambiguated from a single sweep alone; some require multi-rate sweeps, external R proxies, or external response-coordinate calibration, and the protocol's sharpness is conditional on those external channels being available.

Stack placement

Readers outside the UCT project may skip this paragraph; the methodological claim begins in §1.

This paper is a Universal Collapse Theory Methods Paper for S3 applications. Records Across Nature, Life, and Mind (Jones 2026a) defines records as the persistence layer and states the loop-area signature S3 as a portable empirical claim. The Structuralization of Empiricism (Jones 2026b) treats S3 as the stabilization signature for path-dependent regimes in which prior resolutions persist as constraints on later updates. The Update Integrity Standard (Jones 2026c) requires an independence audit before S3 loop scaling can support a Level 3 claim. Records Amplify Hysteresis (TN-S3; Jones 2026e) supplies the formal lemma. Auditing Independence (Jones 2026f) and Auditing Constraint Asymmetry (Jones 2026g) are the parallel S1 and S2 Methods Papers this paper takes as architectural template. The present paper translates the S3 commitments into a field-deployable loop-scaling diagnostic. Formal pair: TN-S₃ (Records Amplify Hysteresis). Together they constitute the S₃ formal-protocol pair — the paired Technical Note proves the formal bound; this Methods Paper translates it into a deployable audit protocol. It is a Methods Paper, not a standard.

1. Introduction

Path-dependent response under bidirectional constraint sweeps is a recurring observable in the biological, physical, and social sciences. Soft tissue exhibits stress–recovery loops in cyclic loading (Fung 1967; Zhang, Chen & Kassab 2007); ecosystems shift between alternative stable states with different forward and reverse thresholds (Scheffer et al. 2001; Beisner, Haydon & Cuddington 2003); developmental fate commitment can be irreversible under cue reversal (Pájaro et al. 2019); institutional and policy systems lock in once thresholds are crossed (Arthur 1989; Pierson 2000); ferromagnetic and ferroelectric materials trace Preisach-form hysteresis loops under field reversal (Mayergoyz 2003); and rate-dependent viscoelastic and thixotropic systems trace loops whose area depends on sweep rate rather than on accumulated history (Wang & Ewoldt 2023). The qualitative regularity — observed path dependence under sweep reversal — is what the Universal Collapse Theory library calls the S3 stabilization signature.

Standard practice for reporting hysteresis is well established: loop area, coercive field, remanence, energy dissipation per cycle, return-point memory checks (Sethna et al. 1993), and Preisach density estimation (Mayergoyz 2003; Brokate & Sprekels 1996). These statistics describe the loop's geometry and dissipation. They are appropriate for material characterization, model description, and within-system comparison. They are not, by themselves, a test of an external record-state claim: loop area at a single condition cannot serve as evidence that accumulated records, rather than rate lag, instrument hysteresis, or another confound, produced the rate-independent component of the path dependence.

This is more than a technicality. Distinct generating processes can produce nearly identical hysteresis loops at a single sweep condition. (i) Genuine record-coupled hysteresis: as R varies across an audited range, the rate-independent component of loop area scales with R at a rate predicted by the formalism. (ii) Confound-induced apparent record coupling: rate-dependent dynamics, instrument hysteresis, minor-loop violations, or circular R or response-coordinate construction can produce loops that mimic record-driven scaling at a single condition but fail to scale with audited R , scale tautologically, or saturate prematurely. (iii) Record-modulated rate lag: a real but distinct phenomenon in which the relaxation timescale couples to R , producing positive R -scaling at finite sweep rate but zero quasi-static residual loop area; this is not

the same claim as record-driven persistence. Single-condition loop-area summaries can be matched across these regimes; they differ in how loop area scales with audited R , how the scaling depends on sweep rate in the quasi-static limit, whether the sweep extends past both switching thresholds, and whether R and the response coordinate are constructed independently of the loop being tested.

The diagnostic proposed here is a comparison of the observed quasi-static loop-area-versus- R curve $A_{qs}(R)$ (in the continuous regime) or the rate-independent loop-area-versus- R curve $A(R)$ (in the discrete regime) against the scaling law predicted by an independently audited record-state proxy and an independently specified formalism, together with a sweep-rate audit (where executable), a major-loop validity check, and a feature-disjointness check between R and the response coordinate $y = f(x)$ used to compute loop area in high-dimensional systems.

The contribution of this paper is methodological. The single sentence the protocol is built around: a loop-area statistic becomes evidence for record-driven path dependence only when R , the response coordinate, the formalism, the predicted scaling law, the sweep bounds, and the rate audit are specified independently of the loop area being explained.

(1) A loop-scaling protocol applicable in both discrete-state and continuous-state regimes, with the same four-step structure used in Auditing Independence (Jones 2026f) for $S1$ and Auditing Constraint Asymmetry (Jones 2026g) for $S2$.

(2) A sweep-rate audit (where executable) that separates record-driven static components from rate-driven contributions to loop area via quasi-static extrapolation, with explicit handling of the case where the rate audit is not executable on a given dataset.

(3) A sharp model-conditional falsifier: comparison of the fitted $A_{qs}(R)$ or $A(R)$ curve against the model-predicted curve $A_{pred}(R)$, where the prediction is computed from sources independent of the loop being tested. Sharpness is conditional on the formalism choice, the R -source, and the response-coordinate construction; if any of these is wrong or circular, the falsifier is still informative but no longer sharp.

The sharpness comes from comparison against a prediction computed without using the loop area being explained. This breaks the circularity that affects diagnostics calibrated against the same loop they are supposed to validate — the failure mode $TN-S3$ explicitly flags as the central $S3$ vulnerability (Jones 2026e §7), and the structural problem Kriegeskorte et al. (2009) identify as "double dipping" in a neuroimaging context. The methodological problem is structurally similar to the one Bland and Altman (1986) flagged for method-comparison correlation in clinical measurement and the one Varma and Simon (2006) flagged for cross-validation under model selection: a goodness-of-fit summary computed on the same data the model was tuned to is not evidence about the data-generating process.

The paper is organized as follows. Section 2 states the formal results — discrete-state (§2.1) and continuous-state (§2.2) regimes, the circular- R and circular-projection traps (§2.3), and a diagnostic comparison (§2.4). Section 3 specifies the protocol with per-step checklists. Section 4 demonstrates it on synthetic cases in both regimes. Section 5 specifies the application to the rice transcriptome heat-stress/recovery dataset (Jones 2026h; GEO GSE74793; Wilkins et al. 2016) as a time-course application target with the rate audit marked not executable on that dataset. Sections 6–7 discuss applicable domains

and limitations. Section 8 provides a reporting template. Section 9 places the protocol in the broader UCT context. Appendix A provides reproduction code via the companion script `methods_s3_v10_figures.py`.

2. Theory: loop-area curves in two regimes

2.1 Discrete-state regime: record-shifted thresholds imply linear scaling

Setup. Let $x_t \in \{-1, +1\}$ be a binary state evolving under a continuous constraint signal $K_t \in \mathbb{R}$. Let $R \geq 0$ be a scalar summary of accumulated record state, held fixed during a single sweep; multi-sweep regimes in which R updates as records accumulate are out of scope for this protocol (Jones 2026e §3). Let $\theta_0 \geq 0$ and $\alpha \geq 0$ be parameters with $\theta_0 + \alpha R > 0$ in the nondegenerate case. The state evolves under the switching rule $x_{t+} = +1$ if $K_t \geq \theta_+(R) := \theta_0 + \alpha R$; $x_{t+} = -1$ if $K_t \leq \theta_-(R) := -\theta_0 - \alpha R$; $x_{t+} = x_t$ otherwise. The interval $[\theta_-(R), \theta_+(R)]$ is the memory zone. A closed major sweep is a quasi-static trajectory in which K increases monotonically from $K_{\min} \leq \theta_-(R)$ to $K_{\max} \geq \theta_+(R)$, then decreases monotonically back to $K_{\min}$, with initial state $x(0) = -1$ and no additional reversals. The loop area is $A(R) := |\oint x \, dK|$ under such a sweep.

Theorem 1 (Jones 2026e). Under the setup above with $\theta_0 \geq 0$, $\alpha \geq 0$, $R \geq 0$, and any closed major sweep $K(t)$ with initial state $x(0) = -1$, the absolute loop area satisfies $A(R) = 4\theta_0 + 4\alpha R$, with slope $dA/dR = 4\alpha$, strictly positive whenever $\alpha > 0$. In the limit $\theta_0 = 0$ and $\alpha R = 0$, the memory zone collapses and $A = 0$, recovering reversible response. The (K, x) trajectory traces a counterclockwise rectangle of horizontal extent $2\theta_0 + 2\alpha R$ and vertical extent 2; the absolute enclosed area is $4\theta_0 + 4\alpha R$. Full geometric and signed-integral derivation is in Jones 2026e §2.

Two identifiable objects. Two parameters organize the discrete-regime diagnostic. The intrinsic loop area $A_0 := 4\theta_0$ fixes the loop area at zero record state. The record-coupling slope $c := 4\alpha = dA/dR$ fixes the rate at which loop area grows with audited R . The discrete-state falsifier compares an empirical fit $\hat{c}$ from $A(R) \approx \hat{A}_0 + \hat{c}R$ over the audited R range against an independently determined c — from mechanistic calibration, simulation truth, or a held-out training loop family disjoint from the test loop family. Slope-independence rules are specified in §3.4. The intrinsic-area object $\hat{A}_0$ is reported alongside $\hat{c}$ but is not itself the falsifier; a fitted $\hat{A}_0$ that disagrees with an independent estimate of $4\theta_0$ is informative as a consistency check but does not by itself reject record coupling, since $\hat{A}_0$ also absorbs additive instrument-hysteresis offsets and rate-driven contributions at the nominal sweep rate (§4.1 regimes B and C).

Heterogeneous and Preisach extensions. The single-hysteron lemma is the base case for the general Preisach model (Mayergoyz 2003; Brokate & Sprekels 1996), which represents a system as a weighted superposition of independent hysterons indexed by threshold pair (θ_+, θ_-) against a Preisach density $\rho(\theta_+, \theta_-)$. Under nonnegative density and a record update that monotonically widens active threshold gaps within the swept range, aggregate loop area inherits monotone growth from the single-hysteron case, and the predicted slope is obtained by integrating against ρ . Arbitrary translations, scalings, or asymmetric reweightings of the density need not preserve monotonicity without further assumptions. In the Preisach case the diagnostic compares $\hat{c}$ against the density-integrated predicted slope, not against 4α ; the principle is identical, the predicted scalar differs.

Out-of-scope regimes. Critical-disorder and avalanche regimes — random-field Ising model near critical disorder (Dahmen & Sethna 1996), crackling-noise systems (Sethna, Dahmen & Myers 2001) — produce power-law avalanche distributions and scaling behavior in which loop area does not satisfy a linear lemma; the discrete-state diagnostic does not apply unmodified. Return-point memory (Sethna et al. 1993) is consistent with the discrete-state lemma but neither implies it nor is implied by it; the diagnostic does not require a return-point-memory check, and many continuous-state systems exhibit S3 without exhibiting return-point memory in any strict form.

2.2 Continuous-state regime: rate-extrapolated quasi-static loop scaling

Setup. Let $x(t) \in \mathbb{R}^p$ be a continuous (possibly high-dimensional) state evolving under a continuous constraint signal $K(t) \in \mathbb{R}$. Let $v > 0$ denote the sweep rate (operationally $v := |dK/dt|$ for a triangular sweep, or any operationally defined rate parameter for the chosen sweep family). Let $R \geq 0$ be a scalar summary of accumulated record state, fixed during a single sweep.

Response coordinate. For $p > 1$, the line integral $\oint x \, dK$ is vector-valued and does not admit a scalar loop-area interpretation directly. The protocol therefore requires a predeclared scalar response coordinate $y(t) := f(x(t)) \in \mathbb{R}$, with f fixed before the loop-area test is executed. The scalar loop area is then $A(R, v) := |\oint y \, dK|$ under a closed sweep at rate v . Construction rules for f are in §2.3 (admissibility) and §3.2 (declaration); two-dimensional embedding alternatives (e.g., area in a predeclared (PC₁, PC₂) plane) are noted as scope-limited extensions.

Decomposition. The protocol assumes the loop area admits a static-plus-dynamic decomposition

$$A(R, v) = A_{\text{static}}(R) + D(R, v), \quad \text{with} \quad \lim_{v \rightarrow 0} D(R, v) = 0,$$

where $A_{\text{static}}(R)$ is the rate-independent component of the loop area, identified with the record-driven contribution under the chosen formalism, and $D(R, v)$ collects all rate-driven contributions, including both record-independent rate lag (e.g., a system with constant relaxation timescale) and record-modulated rate lag (e.g., a system with R -coupled relaxation timescale $\tau(R)$). The S3 diagnostic is specifically a claim about $A_{\text{static}}(R)$, operational via the quasi-static extrapolation $A_{\text{qs}}(R) := \lim_{v \rightarrow 0} A(R, v)$. Where the decomposition holds, $A_{\text{qs}}(R) \rightarrow A_{\text{static}}(R)$ as the sweep is slowed.

Important caveat — record-modulated rate lag is a separate phenomenon. A system in which the relaxation timescale couples to record state, $\tau(R)$, can produce loop area at finite sweep rate that scales with R , but with $A_{\text{static}}(R) = 0$. This is record-modulated rate lag — a real and meaningful phenomenon, but structurally distinct from the persistent threshold structure that the S3 lemma describes. The S3 diagnostic treats record-modulated rate lag as a negative control (zero quasi-static residual) and is intentionally insensitive to it. A diagnostic for record-modulated rate lag would test the fitted $\tau(R)$ curve against an independent prediction; that is a different protocol and not the subject of this paper.

Limit caveat. The quasi-static limit may fail to exist in some systems (e.g., glassy relaxation with no rate-independent endpoint) or may contain a non-record rate-independent contribution (e.g., topological hysteresis from boundary geometry). In such cases the correct outcome is "protocol not applicable" or "record-coupling model rejected for this system," not automatic S3 failure.

Three model classes for $A_{\text{static}}(R)$. Three model classes supply candidate predicted curves $A_{\text{pred}}(R) = A_{\text{static}}(R)$ under genuine record coupling, each appropriate for a different domain. The diagnostic compares $A_{\text{qs}}(R)$ against $A_{\text{pred}}(R)$ under the class committed to before the test.

(i) Continuous-state Preisach analog. For systems with a continuous distribution of bistable elements indexed by threshold pair (θ_+, θ_-) , the aggregate quasi-static loop area is the integral of single-hysteron contributions against a Preisach density $\rho(\theta_+, \theta_-; R)$ that may itself depend on accumulated record state. Under R -coupled density shift (e.g., density translation or width modulation with R) within a fixed swept range, $A_{\text{static}}(R)$ inherits the integral structure of the discrete case and admits a model-predicted slope obtained by integrating dp/dR weighted by the per-hysteron loop contribution. The discrete-state lemma $A_{\text{static}} \propto 4\alpha$ is recovered when ρ is concentrated on a single hysteron pair.

(ii) Bistable basin-geometry model (Kramers / first-passage-time class). For systems where state transitions are escape events from metastable basins (continuous bistable gene regulation, Pájaro et al. 2019; chemical reactions in the Kramers regime; cellular fate commitment), record coupling enters as basin-geometry modification with R . In the deterministic quasi-static limit, the system tracks the stable branch until the saddle-node bifurcation, where the metastable branch vanishes; $A_{\text{static}}(R)$ follows from the R -dependence of the bifurcation thresholds $\theta_{\pm}(R)$. With branch separation normalized to 2 in symmetric cases, $A_{\text{static}}(R) = 2[\theta_+(R) - \theta_-(R)]$; for the symmetric form $\theta_{\pm}(R) = \pm(\theta_0 + \alpha R)$, this reduces to $4(\theta_0 + \alpha R)$, recovering the discrete-hysteron base case. With finite noise or finite sweep rate, Kramers / mean-first-passage-time corrections (Hänggi, Talkner & Borkovec 1990) shift the effective switching points away from the deterministic bifurcation thresholds. These corrections vanish only in the deterministic, zero-noise quasi-static limit; at finite noise, slow sweeps can still produce pre-bifurcation escape, so deterministic-quasi-static and finite-noise-quasi-static are not the same regime and the protocol's A_{qs} estimate tracks whichever is realized in the data. This class is the natural framing for the rice transcriptome target in §5.

(iii) Soft bistable continuous-state model. For systems where $x(t)$ is a continuous variable in a bistable potential $V(x; K, R)$ with R -coupled threshold structure, a closed sweep traces a smooth version of the discrete-hysteron loop. In the limit of fast intra-branch relaxation, $A_{\text{static}}(R) \rightarrow 4(\theta_0 + \alpha R)$, recovering the discrete lemma. This class is used as the genuine synthetic case in §4.2.

Confound class — pure linear viscoelastic relaxation with possibly R -coupled τ . A linear viscoelastic system with relaxation $dy/dt = -(y - K)/\tau(R)$ has zero quasi-static loop area at every R , regardless of whether τ depends on R . Such a system produces finite loop area at finite sweep rate and can show R -dependence in that finite-rate area when τ depends on R , but $A_{\text{qs}}(R) = 0$. This is record-modulated rate lag, not S3, and is used as the negative-control synthetic case in §4.2.

Falsifier. Under the decomposition above and an independently specified $A_{\text{pred}}(R)$ from one of the three model classes (or a domain-specific alternative), the falsifier is $A_{\text{qs}}(R) \neq A_{\text{pred}}(R)$ within declared tolerance. Class commitment must be made before the test; comparing a single dataset against the best-fitting class chosen post hoc is descriptive, not falsificatory. A separate negative-control falsifier is $A_{\text{qs}}(R) \approx 0$ with positive finite-rate loop area — this rejects S3 in favor of record-modulated rate lag or pure rate lag, regardless of class choice.

2.3 The circular-R trap and the circular-projection trap

Two structurally similar circularity failures threaten the diagnostic. Both are instances of the general "double dipping" problem (Kriegeskorte et al. 2009): using the same data for variable selection and selective analysis distorts descriptive statistics and invalidates inference.

Circular-R failures (inadmissible R constructions):

- (a) $R := A_{\text{observed}} / 4\alpha$ — defining R as a function of the loop area being tested produces tautological linear scaling with no falsification content.
- (b) $R :=$ coercive field, remanence, or any other geometric feature derived from the same loop — same failure mode.
- (c) $R :=$ PC₁ score fit on the same genes and samples used to compute the loop coordinate y — the omics version of the same trap. Even if the gene panel is ostensibly different, leakage through correlated samples reproduces the failure.
- (d) $R :=$ latent variable inferred from the same loop family being tested — without a held-out family, this is exploratory rather than confirmatory.

Admissible R constructions:

- (a) Prior-exposure logs, environmental history, intervention history — events that occurred before the loop being tested and were recorded outside the loop measurement system.
- (b) Predeclared molecular marker panels — gene sets, protein sets, or metabolite sets named in advance from prior literature, with feature-disjointness from the response coordinate y.
- (c) Independent phenotype measurements — physiological, morphological, or behavioral measures collected before, after, or during the sweep but not used to construct y.
- (d) Tissue-history or cell-history markers measured on disjoint samples or disjoint feature sets.
- (e) Simulation truth (in synthetic demonstrations) — the ground-truth R is known by construction and is independent of any observed quantity.
- (f) Pretest threshold calibration on a held-out sweep family — relevant for discrete-state systems where $\theta_{\pm}(R)$ is needed for major-loop verification.

Circular-projection failures (inadmissible y constructions):

When state x is high-dimensional, the response coordinate $y = f(x)$ is itself a modeling choice that can be circularly tuned. The same family of failures applies: (a) y selected to maximize the observed loop area on the test family; (b) y selected as a discriminant between high-R and low-R samples in the test family; (c) y selected by post-hoc feature inspection of the test loops. Admissible y constructions: predeclared scalar (e.g., "sum of heat-shock-protein expression"); predeclared linear projection from prior literature; PC₁ or other unsupervised summary fit on a training family disjoint from the test family; held-out validation design with explicit train/test split documented.

Feature disjointness in omics settings. When R is also derived from molecular features (e.g., a tissue-history marker panel) and y is derived from molecular features (e.g., a heat-shock response sum), the two feature sets must be disjoint. If they are not disjoint, an explicit leakage analysis must be reported showing that the shared features do not drive the observed scaling.

Why this section is load-bearing. The rest of the protocol assumes R , y , and the predicted slope are independent of the loop area being explained. Without operational rules for that independence, the protocol degrades into descriptive curve fitting and provides no falsification content. §3 implements these rules as checklist items; §5 demonstrates them in a real-data setting with explicit feature-disjointness reporting.

2.4 Diagnostic comparison

Table 1 summarizes the parallel structure across regimes. The choice of regime is determined by state-variable structure, not by framework or preference; discrete-regime statistics do not generalize directly to continuous-state data, and continuous-regime statistics do not generalize directly to discrete switching systems.

Regime	State / formalism	Statistic	Predicted scaling	Failure signature
Discrete, single hysteron	$x \in \{-1, +1\}$, $\theta_{\pm}(R) = \pm(\theta_0 + \alpha R)$	$A(R) = \oint x \, dK $	$A(R) = A_0 + cR$, $c = 4\alpha$	$\hat{c} \ll 4\alpha$; saturation; rate-driven $\hat{A}_0$
Discrete, Preisach density	Aggregate over (θ_+, θ_-) ; $\rho(\cdot; R)$	$A(R)$ integrated over ρ	model-predicted slope from $\rho(\cdot; R)$	$\hat{c} \neq \rho$ -integrated slope; ρ -fit circularity
Continuous, quasi-static	$x \in \mathbb{R}^p$; predeclared $y = f(x)$; $v \rightarrow 0$	$A_{qs}(R) = \lim_{v \rightarrow 0} \{A(R, v)\}$	$A_{pred}(R)$ from §2.2 model class	$A_{qs}(R) \approx 0$ (rate-modulated only); limit ill-defined
Continuous, fixed-rate	$x \in \mathbb{R}^p$; $y = f(x)$; single v^*	$A(R, v^*)$	$A_{static}(R)$ plus rate-driven offset	rate-driven contribution unaudited

Table 1. Diagnostic comparison across regimes. The discrete-regime falsifier is the fitted slope $\hat{c}$ versus 4α (single hysteron) or ρ -integrated slope (Preisach); the continuous-regime falsifier is $A_{qs}(R)$ versus $A_{pred}(R)$ under one of the three model classes of §2.2, with $A_{qs}(R) \approx 0$ at positive finite-rate area as a separate negative-control rejection (record-modulated rate lag). Mixing regime statistics is a category error analogous to Jones 2026f's classification-versus-continuous warning.

3. The protocol

The protocol consists of four steps, with regime-specific instantiations at Steps 1, 2, and 4. Each step is specified to be auditable independently of the loop-area data the protocol tests. Where a step is not executable on a given dataset (most commonly Step 3 on fixed-protocol time-course data), the protocol must report the step as "not executable" and specify which confound classes therefore remain unaudited.

The per-step checklists below are the minimum reporting requirements; §8 provides a consolidated reporting template.

3.1 Step 1: Independent record-state characterization

Specify R from non-loop channels. Acceptable inputs are listed in §2.3 (admissible R constructions). The audit must not use any feature of the loop being tested and must satisfy feature-disjointness from the response coordinate y in omics settings. Multi-sweep accumulation regimes (Jones 2026e §3) require a separate model and are out of scope for this protocol.

Step 1 reporting checklist. (a) R-source (admissible category from §2.3); (b) measurement timing relative to the test sweep; (c) scalar / vector / categorical structure of R; (d) independence relation to the loop statistic and to y; (e) leakage controls; (f) missingness rule; (g) confirmation that R is fixed during each test sweep.

3.2 Step 2: Response-coordinate declaration and loop-area curve construction

(2a) Declare the response coordinate $y = f(x)$. For $p = 1$, f is the identity. For $p > 1$, f must be specified before loop area is computed and must satisfy the admissibility rules in §2.3 (predeclared, prior-literature-fit, training-family-fit, or held-out validation). The declaration must be documented and timestamped before §3.2(b) is executed.

(2b) Construct the loop-area curve. Drive closed major sweeps at multiple values of audited R. Compute $A(R) := |\oint y \, dK|$ at each R. For continuous-state systems, sweep at multiple rates v where the dataset permits and report $A(R, v)$ at each rate.

Step 2 reporting checklist. (a) loop statistic (signed line integral, absolute value, or RMS); (b) scalar response coordinate f and source (predeclared / prior-literature / training-fit / held-out); (c) sweep bounds $K_{\min}$, $K_{\max}$; (d) sweep direction and closure rule (see §5 for the open-trajectory closure convention); (e) interpolation method between samples; (f) numerical integration method (trapezoid, Simpson, Green's-theorem area); (g) replicate-aggregation rule; (h) per-R uncertainty estimate (bootstrap with stated unit, parametric, or replicate-based).

3.3 Step 3: Sweep-rate audit and major-loop verification

Sweep-rate audit (where executable). Vary v at fixed R and report $A(R, v)$; extrapolate to the quasi-static limit $v \rightarrow 0$ to obtain $A_{\text{qs}}(R)$ under the §2.2 decomposition. The convergence criterion is that the relative change between the two slowest rates falls below a declared tolerance (default $\epsilon = 0.05$); if convergence is not achieved, the rate audit is reported as "partial" with the achieved tolerance stated.

Where the rate audit is not executable (single fixed protocol, time-course data with no rate variation, or a dataset where rate cannot be operationally defined), the step is reported as "not executable on this dataset" and the analysis explicitly names which confound classes therefore remain unaudited (record-modulated rate lag, record-independent rate lag, thixotropic recovery, frequency-dependent dissipation).

Major-loop verification independence. If thresholds $\theta_{\pm}(R)$ (discrete-state) or saturation extents (continuous-state) are needed for the major-loop check, they must be obtained from pretest threshold calibration, independent switching measurements, or a held-out sweep family. If they are inferred from

the same loop-area curve being tested, the verification must be reported as internal consistency rather than an independent audit.

Step 3 reporting checklist. (a) sweep-rate values $v_1 < v_2 < \dots$ (or "not executable"); (b) extrapolation model (linear in v , polynomial, exponential, model-class-specific); (c) quasi-static convergence criterion and achieved tolerance; (d) executability label per step (executable / partially-executable / not-executable) with named unaudited confound classes if not executable; (e) major-loop verification source (pretest calibration, held-out family, or internal-consistency-only flag).

3.4 Step 4: Falsifier comparison

Slope-independence rule. The predicted slope c (or, in continuous regime, the predicted curve $A_{\text{pred}}(R)$) must be obtained from independent calibration, mechanistic parameter estimates, simulation truth, or a training loop family disjoint from the test loop family. If α (single-hysteron case), the Preisach density (superposition case), or the continuous-state model class (continuous case) is fit from the same $A(R)$ curve being tested, the comparison is descriptive only and cannot serve as the sharp falsifier.

Discrete-state. Fit $A(R) \approx A_0 + cR$ over the audited R range; compare fitted slope $\hat{c}$ against predicted 4α (single hysteron) or model-predicted slope (Preisach density).

Continuous-state. Compare $A_{\text{qs}}(R)$ (where the rate audit was executable) or $A(R)$ at the slowest available sweep rate (where it was not) against $A_{\text{pred}}(R)$.

Fixed-rate caveat. When Step 3 is not executable, comparison of $A(R)$ at the slowest or only available sweep rate is conditional or exploratory unless independent evidence rules out rate-driven confounds (record-modulated rate lag, record-independent rate lag, thixotropic-analog dynamics, frequency-dependent dissipation). Such a result may be reported as S3-consistent under stated caveats — naming the specific unaudited confound classes — but not as a full rate-audited S3-positive result. The protocol does not permit a confirmatory S3-positive claim from a fixed-rate dataset alone.

Tolerance band — synthetic demonstrations. When the predicted slope c_{pred} is simulation truth, the default decision rule is $|\hat{c} - c_{\text{pred}}| < Z \cdot \text{SE}(\hat{c})$ for declared Z ($Z = 1.96$ for two-sided 95% CI), or analogous in the continuous regime. This is the rule used in §4.

Tolerance band — empirical applications. When the predicted slope or curve is obtained from independent calibration, mechanistic estimates, or a held-out training family, the SE-only rule is generally insufficient: a large standard error can let almost any $\hat{c}$ pass, and a small standard error can flag substantively trivial discrepancies as failures. The empirical rule must combine statistical uncertainty with a pre-declared substantive equivalence tolerance δ :

$$|\hat{c} - c_{\text{pred}}| < \max(\delta, Z \cdot \text{SE}(\hat{c})),$$

or, more conservatively, $|\hat{c} - c_{\text{pred}}| < \delta$ with the confidence interval reported separately. Pre-registration of δ , Z , and SE construction (bootstrap with stated resampling unit, profile-likelihood, or analytic) is required for the decision to count as confirmatory. The choice of δ is domain-specific and must be justified in terms of what discrepancy from c_{pred} would be substantively meaningful for the claim being tested; it is not a free parameter to be tuned to the observed $\hat{c}$.

Step 4 reporting checklist. (a) predicted curve $A_{\text{pred}}(R)$ (or scalar c) with derivation source; (b) fitted curve and confidence interval (bootstrap with stated resampling unit, profile-likelihood, or analytic); (c) tolerance band (synthetic-demo SE rule, or empirical $\max(\delta, Z \cdot \text{SE})$ rule, with declared values of δ and Z); (d) decision rule (accept / reject / exploratory); (e) named confound signature where deviation is observed; (f) sensitivity analysis across alternative admissible model classes from §2.2 (continuous regime) or alternative Preisach-density specifications (discrete-Preisach regime).

4. Synthetic demonstrations

The protocol is demonstrated on synthetic cases in both regimes. Frozen simulator code is archived as the companion file `methods_s3_v10_figures.py`; numerical outputs and figures are reproduced exactly from that script with random seed 0 (discrete) and seed 1 (continuous bootstrap).

Statistical definitions used throughout §4. (i) Bootstrap resampling unit: R -levels are resampled with replacement ($n = 2,000$ resamples); for each resample the slope $\hat{c}$ is refit and the 95% CI is the $[2.5, 97.5]$ percentile of the bootstrap slope distribution. The predicted slope c_{pred} is treated as known by construction in these synthetic cases. (ii) rate-dep statistic: for each R , compute $(\max_v A - \min_v A) / \text{mean}_v A$ across the audited rate range; report the mean of this normalized range over R . A rate-dep value near zero indicates rate-independent loop area; values near one or larger indicate strongly rate-driven loops. (iii) Tolerance band (synthetic-demonstration rule only): a regime is labeled PASS if $|\hat{c} - c_{\text{pred}}| < 1.96 \cdot \text{SE}(\hat{c})$ in the discrete case, or $|c_{\text{qs}} - c_{\text{pred}}| < 1.96 \cdot \text{SE}(c_{\text{qs}})$ in the continuous case (where SE is the bootstrap standard deviation). REJECT decisions use the same threshold. This synthetic-only rule is appropriate here because c_{pred} is simulation truth; empirical applications must use the equivalence rule of §3.4 with a pre-declared substantive tolerance δ .

4.1 Discrete-state regime: single hysteron with five confound structures

Five regimes are simulated with $\theta_0 = 0.4$, $\alpha = 1.0$, $K_{\text{max}} = 2.0$ (for non-minor-loop cases), $R \in [0, 1]$ in nine increments, sweep rates $v \in \{0.5, 1.0, 2.0, 4.0\}$, and 2,000 bootstrap resamples for slope CIs. (A) Genuine: $\theta_{\pm}(R) = \pm(\theta_0 + \alpha R)$, rate-independent. (B) Rate-driven: $\theta_{\pm}(R) = \pm(\theta_0 + \gamma v)$ with $\gamma = 0.20$ — the rate-driven increment vanishes as $v \rightarrow 0$, leaving the intrinsic baseline $A = 4\theta_0$ at quasi-static, with no R -coupling at any rate. (C) Instrument hysteresis: $\theta_{\pm} = \pm\theta_0$ fixed, no R -coupling, no rate dependence. (D) Minor-loop violation: same as (A) but $K_{\text{max}} = 0.7 < \theta_{+}(R)$ for $R > 0.3$. (E) Circular- R : same dynamics as (A), but R is operationally defined as $A_{\text{observed}} / 4$.

Regime	Predicted $A(R)$	$\hat{c} \pm 95\% \text{ CI}$	$\hat{A}_0$	rate-dep.	Decision
(A) Genuine	linear, slope $4\alpha = 4.000$	4.002 (+3.997, +4.007)	1.599	0.000	PASS
(B) Rate-driven	slope ≈ 0 ; v -dependent area	0.000 (-0.000, +0.000)	2.401	0.903	REJECT (rate)
(C) Instrument hyst.	fixed offset; slope ≈ 0	0.000 (-0.000, +0.000)	1.601	0.000	REJECT (instr.)

Regime	Predicted A(R)	$\hat{c} \pm 95\% \text{ CI}$	$\hat{A}_0$	rate-dep.	Decision
(D) Minor-loop	saturation/truncation at high R	-2.387 (-4.160, +0.000)	1.893	0.000	INVALID (major-loop)
(E) Circular R	tautological by construction	4.000 (+4.000, +4.000)	-0.000	0.000	INADM.

Table 2. §4.1 discrete-regime results. Predicted slope is $4\alpha = 4.000$; predicted intrinsic area is $4\theta_0 = 1.600$. Regime (A) recovers both within tight CI. Regime (B) shows zero R-slope with high rate-dep statistic; the rate-driven increment grows with ν and vanishes as $\nu \rightarrow 0$, leaving the intrinsic baseline $4\theta_0$ at quasi-static. $\hat{A}_0$ for regime (B) is $4(\theta_0 + \gamma\nu)$ at $\nu = 1$, reported as a consistency check; the predicted $\hat{A}_0 = 4\theta_0$ at $\nu \rightarrow 0$ is recovered by extrapolation in the continuous regime but not directly tested here. Regime (C) shows zero R-slope with no rate dependence and $\hat{A}_0$ matching $4\theta_0$, the diagnostic signature of instrument hysteresis with no record coupling. Regime (D) shows saturation at high R: the fit slope is dominated by the truncated high-R points where the sweep cannot reach $\theta_-(R)$, invalidating the major-loop assumption rather than the lemma. Regime (E) gives a perfect slope by construction but the test is tautological since R was derived from A; INADM. = inadmissible by §2.3.




Figure 1. §4.1 discrete-regime demonstration. Left: $A(R)$ at nominal sweep rate $\nu = 1$, with predicted line from Theorem 1. Regime A tracks the prediction; regimes B and C are flat in R (no R-coupling); regime D rises until $K_{\max} < \theta_-(R)$, then collapses. Right: rate-dependence at fixed R range. Regime A is rate-independent (record-driven); regime B shows growing loop area with ν (rate lag, vanishing as $\nu \rightarrow 0$); regimes C and D are rate-independent but flat in R. Generated by `methods_s3_v10_figures.py`.

4.2 Continuous-state regime: bistable system with three confound structures

Three regimes are simulated with $\theta_0 = 0.4$, $\alpha = 1.0$, $K_{\max} = 2.0$, $R \in [0, 1]$ in seven increments, and sweep rates $\nu \in \{0.05, 0.1, 0.2, 0.5, 1.0\}$. $A_{\text{qs}}(R)$ is estimated by linear extrapolation in ν over the three slowest rates. (A) Genuine bistable + small fixed lag: $y(t)$ relaxes with timescale $\tau_A = 0.02$ toward a hysteron-target switching at $\theta_{\pm}(R) = \pm(\theta_0 + \alpha R)$. The static loop area is $4(\theta_0 + \alpha R)$ by §2.2 model class (iii), recovered as $\nu \rightarrow 0$. (B) Pure rate lag (negative control): no hysteron; $dy/dt = (K - y)/\tau(R)$ with $\tau(R)$

$= \tau_0 + \lambda R$ ($\tau_0 = 0.05$, $\lambda = 0.25$). Quasi-static loop area is zero at every R . (C) Combined: hysteron-target as in (A), but with R -coupled relaxation timescale $\tau_C(R) = \tau_A + \delta R$ ($\delta = 0.10$). Both $A_{\text{static}}(R) > 0$ and an additional R -coupled rate-lag contribution are present.

Regime	Predicted $A_{\text{static}}(R)$	$A_{\text{qs}}(0)$	$A_{\text{qs}}(1)$	$\hat{e}_{\text{qs}} \pm 95\% \text{ CI}$	rate-dep.	Decision
(A) Genuine	$4(\theta_0 + \alpha R)$, slope 4.000	1.600	5.601	4.000 (+3.999, +4.002)	0.025	PASS
(B) Pure rate lag	0 at all R	0.000	0.003	0.003 (+0.002, +0.004)	2.506	REJECT (rate-mod.)
(C) Combined	$4(\theta_0 + \alpha R)$, slope 4.000	1.600	5.601	4.000 (+3.999, +4.002)	0.068	PASS + rate flag

Table 3. §4.2 continuous-regime results. Quasi-static loop area $A_{\text{qs}}(R)$ is obtained by linear $v \rightarrow 0$ extrapolation from the three slowest rates. Regime (A) recovers $A_{\text{static}}(R) = 4(\theta_0 + \alpha R)$ within tight CI on the slope and at both endpoints, with very small rate-dep statistic — a clean PASS for genuine record-driven static residual. Regime (B), the negative control, shows $A_{\text{qs}} \approx 0$ at all R despite finite-rate loop area at higher v (rate-dep ≈ 2.5 , the dominant feature of this regime); the diagnostic correctly rejects S3 in favor of record-modulated rate lag. Regime (C) recovers the same A_{static} slope as (A) — record coupling is intact — but the rate-dep statistic is approximately three times larger, indicating a residual R -coupled rate-lag component on top of the static record structure. The Step-3 sweep-rate audit is what disambiguates (A) from (B) and (A) from (C).




Figure 2. §4.2 continuous-regime demonstration. Left: $A(R, v)$ curves at five sweep rates per regime (faded lines), with the $v \rightarrow 0$ extrapolated $A_{\text{qs}}(R)$ shown as bold curves with markers. Genuine (A) and Combined (C) sit on the predicted line $A_{\text{static}} = 4(\theta_0 + \alpha R)$; Pure rate lag (B) sits at zero across all R despite having visible finite-rate loop area at faster sweeps. Right: sweep-rate audit at $R \approx 0.5$. The genuine and combined cases have nearly rate-independent loop area near $4(\theta_0 + 0.5\alpha) = 3.6$ across all rates with star (★) extrapolation hitting the predicted value. The pure-rate-lag case grows from near zero at slow v to ~ 1.3 at $v = 1$, with the $v \rightarrow 0$ extrapolation correctly returning to zero. Generated by `methods_s3_v10_figures.py`.

5. Application: rice transcriptome heat-stress/recovery (continuous-state regime, step-response application target)

The rice transcriptome heat-stress/recovery dataset (Jones 2026h; Wilkins et al. 2016; GEO GSE74793) provides a path-dependent biological system in which transcriptome state response to environmental constraint (temperature) can be measured along a time course of heat exposure and recovery. The dataset comprises rice leaf transcriptome measurements sampled every 15 minutes during a near-instantaneous 30 → 40 °C heat shock and the subsequent recovery period, across biological replicates and conditions. Critically, the dataset uses a fixed temperature-jump protocol rather than varied temperature ramp rates; the sweep-rate audit (§3.3) is therefore not executable on this dataset. Steps 1, 2, and 4 are executable subject to the constraints below.

Step 1 (R characterization). Candidate R variables must be separated into record proxies and stratification covariates. Prior heat-shock exposure history (recorded experimentally, not inferred from the test loop) is a record-proxy candidate. Developmental stage, cultivar, tissue identity, batch, and circadian timing are stratification covariates or blocking variables — they index biological context rather than accumulated record state, and should be modeled as covariates rather than folded into R unless an explicit accumulated-record interpretation is justified. v1.0 will commit to a primary R and a covariate set, with feature-disjointness from §3.2's response coordinate verified explicitly.

Step 2 (response coordinate and loop-area curve). The response coordinate $y = f(x)$ must be predeclared. Three candidate forms, in increasing order of leakage risk: (i) a literature-defined heat-shock-response gene panel scored as a weighted sum (lowest leakage if R is also literature-defined and the two panels are disjoint); (ii) PC_1 from a held-out training set of biological replicates not used in the test loop family; (iii) an unsupervised summary fit on the test family with explicit leakage analysis. v1.0 will commit to one. Loop area is computed as $|\oint y \, dK|$ under the temperature trajectory.

Step-response geometry (not a smooth sweep). Because GSE74793 uses a step-change protocol — temperature transferred near-instantaneously from 30 °C to 40 °C, held, then transferred back to 30 °C — rather than a continuous up/down ramp, $|\oint y \, dK|$ is interpreted here as a step-response loop statistic, not as a smooth sweep-area statistic. The integral receives contributions only at the temperature transitions; during the fixed-temperature plateaus $dK = 0$ and the integrand contributes nothing, even though $y(t)$ continues to evolve. Recovery dynamics during plateaus are represented implicitly through the response values y reached just before each return transition, and explicitly through the separately reported endpoint gap (closure convention paragraph below). A continuous-ramp S3 test — one in which the sweep-rate audit of Step 3 is executable — would require a different dataset with controlled ramp-rate variation; GSE74793 cannot supply that audit, and any S3-consistent result on this dataset is therefore a step-response-geometry result under stated rate-confound caveats, not a full rate-audited S3-positive result.

Step-transition convention for the loop integral. Because the temperature trajectory $K(t)$ is piecewise-constant with discontinuous jumps, the integral $\oint y \, dK$ is not strictly defined without a convention pairing each K -jump with a y value. The default convention for this protocol is left-continuous in y : each K -jump is paired with the response value $y(t^-)$ measured immediately before the transition. Right-continuous ($y(t^+)$) and midpoint ($(y(t^-) + y(t^+))/2$) conventions are admissible alternatives but must be predeclared with an independent justification (e.g., a calibrated instantaneous-response correction). Whichever convention is chosen, it must be applied uniformly across all R levels and reported in the integration-

method field of the Step 2 reporting checklist (§3.2). For 15-minute-sampled time-series data such as GSE74793, the practical effect of convention choice is small but nonzero, and the convention should be reported regardless.

Closure convention for open trajectories. When $K(T) = K(0)$ but $y(T) \neq y(0)$ — i.e., temperature returns to baseline but transcriptome state does not — the (K, y) trajectory is not closed. The protocol imposes closure by appending a vertical segment at fixed $K = K(0)$ connecting $y(T)$ back to $y(0)$. Because this vertical segment contributes zero to $\int y \, dK$, the loop-area statistic is unaffected by the choice. The endpoint gap $\Delta y := |y(T) - y(0)|$ is reported separately as a non-closure diagnostic; large Δy indicates incomplete recovery and may itself carry record-state information orthogonal to the loop area. This convention applies to any time-course application of the protocol where the response coordinate does not return to its starting value within the measurement window.

Step 3 (rate audit). Not executable on GSE74793. The dataset uses a single fixed temperature-jump protocol; ramp rate is not varied. Confound classes that therefore remain unaudited: record-modulated rate lag (transcriptional relaxation timescale coupling to R), record-independent rate lag, thixotropic-analog recovery dynamics, and frequency-dependent dissipation in transcriptome-state trajectory. Any S3-consistent result on this dataset is therefore conditional on rate-driven confounds being ruled out by other means (e.g., independent literature, auxiliary single-cell time-course data, or biological-mechanism arguments). Per the §3.4 fixed-rate caveat, the result is reported as S3-consistent under stated caveats, not as a confirmatory S3-positive result.

Step 4 (falsifier comparison). Compare fitted $A(R)$ against the predicted curve $A_{\text{pred}}(R)$ under the chosen formalism — the bistable basin-geometry / Kramers class of §2.2(ii) is the natural framing for transcriptional commitment dynamics. The predicted curve must be obtained independently, from prior literature on heat-shock response or mechanistic models of transcriptional commitment, not fit from the test family.

v1.0 will name the primary R , the response coordinate, the formalism, and the predicted-slope source explicitly. Numerical execution against the existing T16 analysis is the subject of a companion empirical paper.

6. Discussion

6.1 Domains

Direct applications: stress–recovery in soft tissue (Fung 1967; Zhang et al. 2007); ecosystem regime shifts (Scheffer et al. 2001; Beisner et al. 2003); developmental fate commitment under cue reversal (Pájaro et al. 2019); institutional and policy lock-in (Arthur 1989; Pierson 2000); ferromagnetic and ferroelectric Preisach systems (Mayergoyz 2003); rate-dependent viscoelastic and thixotropic systems (Wang & Ewoldt 2023, useful as a comparison case where the rate audit is the primary diagnostic and the protocol's negative-control regime is the relevant outcome); memory and belief systems where prior commitments alter response to subsequent evidence (Lord, Ross & Lepper 1979). Choice of regime — discrete or continuous — is determined by state-variable structure, not framework preference.

6.2 The independence audit as the load-bearing step

Neither the single-hysteron lemma nor the static-plus-dynamic decomposition is novel in itself. What this protocol contributes is the pairing of these textbook predictions with an independently audited R proxy, an independently declared response coordinate, and a sweep-rate audit (where executable) that distinguishes static record-driven structure from rate-driven contributions including record-modulated rate lag. The audit converts a descriptive loop-area statistic into a falsifiable record-state claim. Without independent R, independent y, and independent slope, the protocol can verify the form of loop scaling (linear, sublinear) but not whether the scaling is record-driven.

6.3 Comparison to alternative diagnostics

Standard hysteresis-summary statistics — loop area, coercive field, remanence, energy dissipation per cycle (Mayergoyz 2003) — are single-condition descriptive summaries. They share the limitation, structurally similar to the one Bland and Altman (1986) flagged for method-comparison correlation in clinical measurement, that genuine record-coupling and rate-driven apparent hysteresis can produce indistinguishable single-condition values under plausible parameter substitutions. The agreement-curve approach of Auditing Independence (Jones 2026f) and the latency-curve approach of Auditing Constraint Asymmetry (Jones 2026g) addressed the same structural problem — comparing observed scaling against an independently specified prediction — for S1 and S2 respectively; the present protocol is the S3 analog.

Preisach density estimation (Mayergoyz 2003; Brokate & Sprekels 1996) is more diagnostic in principle but is typically calibrated against the same loop family being audited, reproducing the circular-fit failure mode (Varma & Simon 2006). The protocol of §3 differs in requiring the predicted slope or curve to come from a non-test source. Where a Preisach density is genuinely available from independent calibration (mechanistic micromagnetics, separate validation experiments), the protocol incorporates it as the predicted slope in §3.4.

Return-point memory tests (Sethna et al. 1993) are complementary but address a different signature; the present protocol does not require a return-point-memory check, since many continuous-state systems exhibit S3 without exhibiting return-point memory in any strict form. Avalanche-scaling diagnostics (Dahmen & Sethna 1996; Sethna, Dahmen & Myers 2001) apply to a different regime — critical disorder — where the linear lemma does not hold and a separate scaling-form diagnostic is appropriate.

Rheology-aware diagnostics for thixotropic and viscoelastic systems (Wang & Ewoldt 2023) overlap most directly with the rate-audit component of this protocol and are complementary in continuous-state physical systems. The present protocol's contribution is the pairing of a rate audit with an independently audited R-scaling check, allowing record-coupling claims to be separated from record-modulated rate lag in systems that may exhibit both.

Where the protocol is most useful: settings with multiple R levels, an independently constructible record-state proxy, and either a multi-rate sweep design or external grounds for ruling out rate-driven contributions. Where it is least useful: single-condition material characterization, systems where record state is intrinsically inferable only from loop geometry, and critical-disorder regimes where the linear lemma does not apply.

7. Limitations

Eight limitations are stated explicitly:

- (1) Loop scaling does not prove metaphysical record persistence. The protocol tests whether observed scaling is consistent with a stated record-coupling model; it cannot establish record persistence as a property of the world.
- (2) The slope c (or curve A_{pred}) is model-relative. Single-hysteron 4α , Preisach-density predictions, and continuous-state model class predictions yield different predicted scalings; sensitivity analysis across plausible formalisms is essential.
- (3) Continuous-state systems require sweep-rate extrapolation, not just R -variation. Mixing classification-style discrete-regime statistics with continuous-state data is a category error analogous to the warning in Jones 2026f.
- (4) Multi-sweep R accumulation (Jones 2026e §3) requires a separate model and is not implied by the linear lemma. The protocol assumes R fixed during each sweep.
- (5) Shared environmental history across audited R values can dominate apparent R -scaling. The audit must include the full exposure chain, not only the immediate pretest history.
- (6) A fitted slope mismatch can reject a record-coupling model but cannot by itself identify the exact confound. Identification requires further audit (independent residual analysis, simulation injection tests, mechanistic ablations).
- (7) Projection dependence. In high-dimensional systems, loop area depends on the chosen response coordinate $y = f(x)$. Different f choices can yield different loop-area curves on the same dataset; the protocol's sharpness is conditional on f being predeclared and admissible per §2.3. Sensitivity analysis across alternative admissible f choices is recommended.
- (8) Dataset eligibility. Some datasets exhibit path dependence but cannot support the full audit because they lack independent R , multiple R levels, valid major-loop coverage, or sweep-rate variation. The rice transcriptome target (§5) is one such case: Steps 1, 2, and 4 are executable but Step 3 is not. Honest reporting of which steps are executable on a given dataset is part of the protocol.

8. Reporting template

The following template specifies the minimum information that should accompany any use of the protocol. The format is modeled on the reporting templates of Auditing Independence (Jones 2026f) and Auditing Constraint Asymmetry (Jones 2026g), consolidated with the per-step checklists of §3.

Field	Specification
Latent record state	What accumulated history is hypothesized to drive path dependence?
Channel / sample units	What counts as a sweep? (specimen, condition, replicate, run)

Field	Specification
Regime	Discrete-state (single hysteron / Preisach) or continuous-state
Loop statistic	$A = \oint y \, dK $, signed, RMS, or other
Response coordinate y and source	Identity ($p = 1$) or predeclared $f(x)$; source per §3.2(2a)
Formalism	Single hysteron, Preisach density, continuous Preisach, bistable basin / Kramers, soft bistable continuous (§2.2)
Predicted scaling $A_{\text{pred}}(R)$ or c	Functional form and source (independent of test loop)
R-source	Record proxy vs. stratification covariate; admissibility class (§2.3)
Feature disjointness (omics)	$R \text{ features} \cap y \text{ features} = \emptyset$, or leakage analysis reported
Sweep bounds	K_{min} , K_{max} ; major-loop verification source
Closure convention	Closed by construction; or vertical-segment closure with reported endpoint gap (§5)
Sweep rate(s)	$v_1, \dots, v_k$ or "single fixed rate"
Rate-audit executability	Executable / partial / not-executable; named unaudited confounds
Quasi-static convergence	Achieved tolerance ε between two slowest rates (continuous regime)
Replicate count	Independent-units estimate per R level
Reference baseline	Simulation truth, leave-one-out, model expectation
Falsifier statistic	$\hat{c}$ vs. c_{pred} , or $A_{\text{qs}}(R)$ vs. $A_{\text{pred}}(R)$
Confidence interval	Bootstrap (resampling unit specified), profile-likelihood, or analytic
Tolerance band	Synthetic-demo SE rule, or empirical $\max(\delta, Z \cdot \text{SE})$ rule with δ pre-declared
Decision	Pass / reject / exploratory / inadmissible / not-applicable
Confound signature	Named confound class consistent with observed deviation
Sensitivity analysis	Across alternative f , formalism, p , or noise model
Claim level	Method demo / Level 3 empirical (UIS terminology)

Field	Specification
UIS ledger link	Independence-audit ledger entry

Table 4. Reporting template for the loop-scaling protocol. Every entry should be filled before the diagnostic is treated as informative.

9. Broader context

This protocol was developed within a structural framework on records and constraint-guided collapse (Jones 2026a, 2026b, 2026c, 2026e) as a falsifier for the loop-area S3 signature. The methods contribution stands independently: an audit for record-coupled hysteresis claims in any domain where loop area can be measured against an externally audited record-state proxy. The internal framework is one motivating context; the protocol's correctness and applicability do not depend on adopting it.

Appendix A. Reproduction code

The companion script `methods_s3_v10_figures.py` reproduces the results of §4 and Figures 1 and 2 exactly, with random seed 0 (discrete) and seed 1 (continuous bootstrap), and dependencies NumPy and Matplotlib only. The script is archived alongside this Methods Paper. The minimal core of the discrete-state simulator is reproduced below for inspection; the continuous-state simulator and figure code are in the companion file.

```
import numpy as np

# NumPy 1.x / 2.x compatibility: np.trapezoid was added in NumPy 2.0,
# replacing the deprecated np.trapz. The companion script uses this shim.
try:
    integrate_trapezoid = np.trapezoid
except AttributeError:
    integrate_trapezoid = np.trapz

THETA0, ALPHA = 0.4, 1.0

def hysteron_sweep(K, theta_p, theta_m, x_init=-1):
    x = np.empty(len(K)); state = x_init
    for i, k in enumerate(K):
        if k >= theta_p: state = 1
        elif k <= theta_m: state = -1
        x[i] = state
    return x

def loop_area(K, x):
    return abs(integrate_trapezoid(x, K))

def major_sweep(K_max, n=4000):
    up = np.linspace(-K_max, K_max, n // 2)
    dn = np.linspace(K_max, -K_max, n // 2)
    return np.concatenate([up, dn])
```

```
# Regime A (genuine):  $A(R) = 4\theta_0 + 4\alpha R$ 
for R in np.linspace(0, 1, 9):
    K = major_sweep(K_max=2.0)
    tp = THETA0 + ALPHA * R
    A = loop_area(K, hysteron_sweep(K, tp, -tp))
    print(f'R={R:.2f}  A={A:.4f}  predicted={4*THETA0 + 4*ALPHA*R:.4f}')
```

References

- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131. <https://doi.org/10.2307/2234208>
- Beisner, B. E., Haydon, D. T., & Cuddington, K. (2003). Alternative stable states in ecology. *Frontiers in Ecology and the Environment*, 1(7), 376–382. [https://doi.org/10.1890/1540-9295\(2003\)001\[0376:ASSIE\]2.0.CO;2](https://doi.org/10.1890/1540-9295(2003)001[0376:ASSIE]2.0.CO;2)
- Bland, J. M., & Altman, D. G. (1986). Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327(8476), 307–310. [https://doi.org/10.1016/S0140-6736\(86\)90837-8](https://doi.org/10.1016/S0140-6736(86)90837-8)
- Brokate, M., & Sprekels, J. (1996). *Hysteresis and Phase Transitions*. Springer. <https://doi.org/10.1007/978-1-4612-4048-8>
- Dahmen, K., & Sethna, J. P. (1996). Hysteresis, avalanches, and disorder-induced critical scaling: A renormalization-group approach. *Physical Review B*, 53(22), 14872–14905. <https://doi.org/10.1103/PhysRevB.53.14872>
- Fung, Y. C. (1967). Elasticity of soft tissues in simple elongation. *American Journal of Physiology*, 213(6), 1532–1544. <https://doi.org/10.1152/ajplegacy.1967.213.6.1532>
- Hänggi, P., Talkner, P., & Borkovec, M. (1990). Reaction-rate theory: fifty years after Kramers. *Reviews of Modern Physics*, 62(2), 251–341. <https://doi.org/10.1103/RevModPhys.62.251>
- Jones, J. C. (2026a). *Records Across Nature, Life, and Mind*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7H6DY>
- Jones, J. C. (2026b). *The Structuralization of Empiricism*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/J4GZ9>
- Jones, J. C. (2026c). *Update Integrity Standard (UIS v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/DWM29>
- Jones, J. C. (2026d). *Objectivity from Records: An Exponential Consensus Bound (UCT Technical Note v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/6M7N3>
- Jones, J. C. (2026e). *Records Amplify Hysteresis: A Loop-Area Lemma (UCT Technical Note v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/QJMSZ>
- Jones, J. C. (2026f). *Auditing Independence in Multi-Channel Measurement: An Agreement-Curve Diagnostic for Genuine versus Correlated Redundancy (UCT Methods Paper v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7U8SK>
- Jones, J. C. (2026g). *Auditing Constraint Asymmetry in Latency-Based Resolution Tests: A Latency-Curve Diagnostic for Neutrality versus Confound-Induced Delay (UCT Methods Paper v1.0)*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/HRKWT>

- Jones, J. C. (2026h). Constraint-Sweep Hysteresis in Rice Transcriptome State During Heat Stress and Recovery: A Reproducible, Theory-Neutral Test on GEO GSE74793. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/KZ8TP>
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature Neuroscience*, 12, 535–540. <https://doi.org/10.1038/nn.2303>
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11), 2098–2109. <https://doi.org/10.1037/0022-3514.37.11.2098>
- Mayergoyz, I. D. (2003). *Mathematical Models of Hysteresis and Their Applications* (2nd ed.). Academic Press/Elsevier. <https://doi.org/10.1016/B978-0-12-480873-7.X5000-2>
- NCBI GEO. Series GSE74793. Rice gene expression in heat stress and dehydration stress — time series. <https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE74793>
- Pájaro, M., Otero-Muras, I., Vázquez, C., & Alonso, A. A. (2019). Transient hysteresis and inherent stochasticity in gene regulatory networks. *Nature Communications*, 10, 4581. <https://doi.org/10.1038/s41467-019-12344-w>
- Pierson, P. (2000). Increasing returns, path dependence, and the study of politics. *American Political Science Review*, 94(2), 251–267. <https://doi.org/10.2307/2586011>
- Scheffer, M., Carpenter, S., Foley, J. A., Folke, C., & Walker, B. (2001). Catastrophic shifts in ecosystems. *Nature*, 413(6856), 591–596. <https://doi.org/10.1038/35098000>
- Sethna, J. P., Dahmen, K., Kartha, S., Krumhansl, J. A., Roberts, B. W., & Shore, J. D. (1993). Hysteresis and hierarchies: Dynamics of disorder-driven first-order phase transformations. *Physical Review Letters*, 70(21), 3347–3350. <https://doi.org/10.1103/PhysRevLett.70.3347>
- Sethna, J. P., Dahmen, K. A., & Myers, C. R. (2001). Crackling noise. *Nature*, 410(6825), 242–250. <https://doi.org/10.1038/35065675>
- Varma, S., & Simon, R. (2006). Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7, 91. <https://doi.org/10.1186/1471-2105-7-91>
- Wang, Y., & Ewoldt, R. H. (2023). Thixotropy, antithixotropy, and viscoelasticity in hysteresis. *Journal of Rheology*, 67(6), 1199–1219. <https://doi.org/10.1122/8.0000620>
- Wilkins, O., Hafemeister, C., Plessis, A., et al. (2016). EGRINs (Environmental Gene Regulatory Influence Networks) in rice that function in the response to water deficit, high temperature, and agricultural environments. *The Plant Cell*, 28(10), 2365–2384. <https://doi.org/10.1105/tpc.16.00158>
- Zhang, W., Chen, H. Y., & Kassab, G. S. (2007). A rate-insensitive linear viscoelastic model for soft tissues. *Biomaterials*, 28(24), 3579–3586. <https://doi.org/10.1016/j.biomaterials.2007.04.040>

This paper is part of the Universal Collapse Theory library. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure

AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author's own, and the author takes full responsibility for the content.

Citation

Jones, J. C. (2026). Auditing Record State in Constraint-Sweep Hysteresis Tests: A Loop-Scaling Diagnostic for Record-Driven versus Confound-Induced Path Dependence (UCT Methods Paper v1.0). HoldingLight LLC.