Eormen: A Deterministic, Non-Interfering System for Observational Sensitivity in Complex Dynamics

L. A. Palmer

palmer@eormen.com

eormen.com

February 2026

1. Introduction

1.1 Motivation

Complex engineered systems rarely fail without warning. In many cases, failure is preceded by gradual, distributed structural change rather than by a single decisive threshold crossing. These early changes are often weakly expressed, non-linear, or context-dependent, and may not correspond cleanly to any one scalar indicator.

Current approaches in predictive maintenance and condition monitoring tend to prioritise forecast accuracy: estimating remaining useful life, predicting time-to-failure, or optimising intervention schedules. While valuable, such approaches face persistent challenges when the underlying system dynamics are complex, high-dimensional, or only partially observable. In these regimes, prediction models can be brittle, difficult to trust, and hard to validate under real operational constraints.

A more fundamental problem therefore precedes prediction: how to observe structural change in a system in a way that is trustworthy, reproducible, and non-interfering.

Eormen is introduced in this work to address that foundational observational problem. Rather than attempting to forecast outcomes, Eormen is designed to reveal when and how the internal structure of a system begins to deviate, under strict guarantees of determinism and non-interference. The intent is not to replace predictive systems, but to establish a reliable observational substrate upon which stronger claims may later be built, if evidence permits.

1.2 What Eormen Is

Eormen is a deterministic, non-interfering observational sensitivity system.

At its core, Eormen separates the generation of system reality from the act of observation. A reality trace is constructed once under fixed configuration and then frozen. Observers replay this frozen reality without any ability to modify it. Observational outputs are therefore guaranteed to be non-interfering by architecture, not by convention.

Within this framework, Eormen evaluates whether observational outputs respond measurably and reproducibly to structured perturbations applied under controlled conditions. Sensitivity is assessed explicitly and conservatively, using acceptance criteria that are declared in advance and preserved across iterations.

In this report, Eormen is presented as an observational system. Its demonstrated capability is the detection of measurable structural sensitivity under strict integrity constraints. It is not presented as a predictive engine, nor as an optimisation or decision-making system.

1.3 What Eormen Is Not

To avoid ambiguity and over-claim, it is important to state clearly what is not claimed in this work.

Eormen is not claimed here to be:

- a failure-time predictor,

- a remaining useful life (RUL) estimator,
- an optimisation engine,
- an autonomous decision or policy recommender,
- a machine-learning model trained to maximise forecast accuracy.

No predictive accuracy metrics are reported. No intervention policies are evaluated. No optimisation objectives are asserted.

These omissions are deliberate. The purpose of this work is to establish integrity and sensitivity at the observational level, before attempting stronger or riskier claims. Any future extension into prediction or optimisation would require separate experimental design, acceptance criteria, and evidence, and is explicitly out of scope for this report.

2. Conceptual Foundations

2.1 Observation as Perspective, Not Forecast

Eormen is grounded in the distinction between observation and prediction. Observation concerns how structure is revealed from a particular perspective; prediction concerns claims about future outcomes. While related, these are not equivalent problems.

In complex systems, early structural deviation may be observable long before any reliable forecast can be made. A system may begin to exhibit changes in coherence, constraint balance, or internal coupling without providing sufficient information to justify a time-to-failure estimate or a calibrated risk score. Treating observation and prediction as the same task can therefore lead either to overconfident forecasts or to suppressed early signals.

Eormen adopts an explicitly observational stance. It asks whether, under a fixed observational frame and strict integrity constraints, structural changes become visible in a reproducible way. It does not attempt to infer outcomes, nor to extrapolate beyond the observed window. In this sense, Eormen functions as a sensitivity instrument rather than as a forecasting oracle.

This distinction is central to the credibility of the system. Claims are limited to what the evidence can support, and observational sensitivity is treated as a valuable property in its own right.

2.2 Determinism and Trust

Trustworthy observation requires more than signal responsiveness. It requires confidence that the act of observation does not itself alter the system being observed, and that results can be reproduced exactly under identical conditions.

Eormen enforces determinism through fixed runtime profiles and locked coefficient artefacts. Given the same inputs and configuration, the system produces identical outputs. This property enables exact replay, forensic analysis, and independent verification.

Non-interference is enforced architecturally. Observers have no write-path into the reality they observe. Reality is generated once, frozen, and then replayed. This one-way flow ensures that observational outputs cannot feed back into, contaminate, or subtly alter the underlying system state.

In safety-critical or high-value operational contexts, this separation is not a convenience but a necessity. Without it, sensitivity signals may be artefacts of measurement rather than properties of the system itself.

Together, determinism and non-interference form the conceptual foundation of Eormen. They are treated as prerequisites for any later discussion of sensitivity, calibration, or utility, and they define the boundary conditions within which all results in this report should be interpreted.

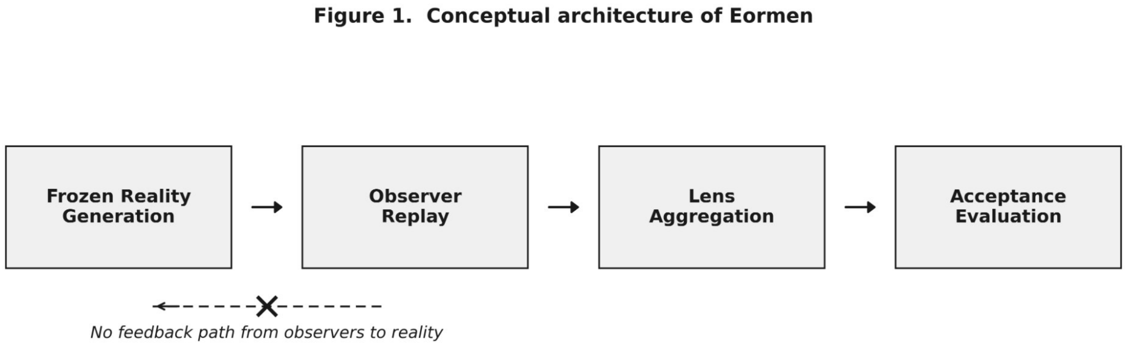

Figure 1: Conceptual architecture of Eormen. The system enforces a strict one-way flow: frozen reality is generated once, then replayed by observers, aggregated through lenses, and evaluated against acceptance criteria. No feedback path exists from observers to reality. This architectural separation guarantees non-interference by construction, not by convention.

3. Experimental Object and Data Substrate

3.1 What is being tested

This work does not test predictive accuracy, remaining useful life estimation, or fault classification performance.

The experimental object is instead the observational behaviour of Eormen under controlled conditions.

More precisely, the programme tests the following question:

> *Does a deterministic, non-interfering observational system produce stable, reproducible, and measurable changes in its observational outputs when the underlying system structure is perturbed, while preserving strict integrity constraints?*

The focus is therefore on observational sensitivity, not on forecasting or decision optimality.

All experiments are framed to evaluate whether Eormen's observational outputs:

- remain invariant when system structure is unchanged,
- respond measurably when structure is altered,
- do so deterministically and without interfering with the underlying system representation.

No experiment in this paper evaluates whether those observations are correct predictions of future events.

3.2 Why a public industrial dataset is required

Because the claim class is observational rather than predictive, the dataset must satisfy different criteria than those typically used in machine learning benchmarks.

The required properties are:

- structural richness rather than label fidelity,
- repeatable system trajectories rather than stochastic sampling,

- transparent provenance rather than proprietary realism,
- wide external familiarity to support audit and replication of claims.

A publicly available dataset is therefore essential. Without this, it would be impossible for external reviewers to separate system behaviour from data artefacts.

3.3 The NASA C-MAPSS dataset as a test substrate

All experiments reported in this paper use subsets of the NASA C-MAPSS turbofan engine dataset.

This dataset is widely used in the literature on degradation and maintenance, but in this work it serves a different role.

Here, C-MAPSS is used as a structural test substrate, not as a target for prediction.

Its key properties for this purpose are:

- multi-dimensional sensor streams,
- long temporal evolution per instance,
- gradual, non-abrupt structural change,
- deterministic replay from fixed initial conditions,
- absence of stochastic noise injection between runs.

These properties make it suitable for evaluating whether observational outputs change before any explicit failure label becomes decisive, which aligns with the observational framing of Eormen.

3.4 Dataset subsets and usage discipline

Different subsets of the C-MAPSS dataset were used for different experimental purposes.

- FD002, FD003, FD004 were used in the baseline integrity programme, executed across all engines, to establish large-scale determinism, non-interference, and numerical stability.
- FD001, together with selected instances from FD002 and FD004, was used in the OBS phase to test observational sensitivity under controlled windowing and probe conditions.

No data leakage occurs between phases because no learning or fitting is performed. Each run is independent and governed solely by fixed configuration artefacts.

3.5 Why this data does not invalidate the claims

It is important to state explicitly what this dataset choice does not imply.

Using C-MAPSS in this work does not mean that:

- Eormen is claimed to be tuned for turbofan engines,
- results generalise automatically to other physical systems,
- any operational thresholds are being recommended.

The dataset is used because it provides a controlled, auditable environment in which structural change exists and can be perturbed without ambiguity.

The claims in this paper are therefore conditional:

> *Given a system with gradual structural evolution and observable state variables, Eormen demonstrates deterministic observational sensitivity under strict non-interference constraints.*

That claim does not depend on the semantics of turbofan engines specifically.

3.6 Why this matters given later findings

The importance of stating the experimental object and data substrate explicitly becomes clear in light of the OBS results.

The discovery that:

- integrity could be closed at scale,
- observers could initially be dead,
- sensitivity emerged only after a theorem-legal excitation path,
- monotone calibration remained unproven,

only has meaning because the data substrate is transparent and externally recognisable.

Had a proprietary or opaque dataset been used, it would be impossible to separate:

- data artefacts,
- observer design limitations,
- acceptance semantics,
- or interpretive bias.

The dataset choice therefore directly supports the credibility of the negative findings as much as the positive ones.

3.7 Summary

This paper tests observational sensitivity, not prediction.

The NASA C-MAPSS dataset is used as a public structural substrate to make that test auditable, reproducible, and interpretable.

All claims are framed with respect to:

- what is observed,
- under what constraints,
- using what evidence,
- and what remains explicitly unproven.

This framing is essential to the scientific integrity of the results that follow.

4. Experimental Programme and Evidence Structure

4.1 Rationale for a staged experimental programme

The experimental programme was deliberately staged to avoid a common failure mode in complex system research: drawing sensitivity or usefulness conclusions before integrity has been established.

Eormen was therefore evaluated in successive phases, each designed to close a specific class of uncertainty before progressing to the next. No later phase was interpreted independently of earlier results.

The stages were:

1. Baseline integrity and stability validation.
2. Observational sensitivity testing (OBS).
3. Failure-mode diagnosis and resolution.
4. Dual-acceptance re-scoring under preserved provenance.

Each stage produced a closed set of artefacts, recorded in immutable dossier directories, and no stage required modification of results from a previous stage in order to proceed.

4.2 Baseline integrity validation (v2 consolidated programme)

The first stage addressed a single question: does the system behave deterministically and non-interferingly at operational scale?

This was evaluated using a consolidated baseline programme that executed across multiple publicly available datasets and a large number of independent system instances. The intent was not to measure sensitivity or performance, but to demonstrate:

- deterministic replay under fixed configuration,
- numerical stability,
- absence of observer interference,
- consistency across repeated executions.

This stage produced a consolidated integrity dossier that serves as the immutable reference point for all subsequent work. Only after this baseline demonstrated complete success were observational sensitivity questions considered admissible.

4.3 Introduction of the OBS phase

With integrity closed, the second stage introduced OBS: an experimental framework designed to test whether observational outputs respond measurably to structured changes in system behaviour, without altering the underlying system machinery.

The OBS phase was explicitly constrained by the following rules:

- no changes to theorem-critical components,
- no adaptive parameter fitting,
- no optimisation objectives,
- no predictive claims.

OBS was therefore not a continuation of the baseline programme, but a new layer of evaluation operating strictly on frozen system representations generated under previously validated conditions.

4.4 Window sweep experiments

The first OBS experiment class was a window sweep.

Window sweeps evaluated observational behaviour across multiple bounded temporal segments selected at different relative positions in a system's lifecycle. The purpose was to determine whether observational outputs depend on where and over what span system behaviour is examined.

The sweep design was explicitly enumerated in advance. All attempted sweep cells were recorded, and cells that were infeasible due to window placement constraints were marked as such rather than being treated as failures.

This distinction proved important later in interpreting results without conflating feasibility constraints with system integrity.

4.5 Probe experiments

The second OBS experiment class introduced probe suites.

Probe experiments applied controlled, deterministic perturbations through allowed external inputs while preserving all integrity constraints. The purpose was to test whether Eormen's observational outputs exhibit measurable sensitivity when underlying system structure is deliberately altered.

Probe experiments were:

- deterministic,
- bounded in amplitude,
- non-adaptive,
- applied under fixed configuration,
- evaluated using the same acceptance logic as non-probe runs.

Each probe family was designed to represent a qualitatively different type of structural change. Control cases were always included and treated as first-class references rather than as implicit baselines.

4.6 Initial OBS outcomes and diagnosis

The initial OBS execution produced a clear and unambiguous result: although system behaviour changed under probe perturbations, observational outputs did not.

This was not treated as a failure of the experimental design, but as a diagnostic outcome. Because integrity was already closed, the absence of observational response could be analysed without ambiguity.

A dedicated expressivity diagnosis was therefore performed. This diagnosis demonstrated that observer channels remained identically zero despite non-trivial changes in system behaviour.

The conclusion was that the observational pathway was structurally under-excited.

Crucially, this diagnosis did not invalidate the system; it identified a design limitation consistent with the reported evidence.

4.7 Resolution and re-execution under preserved constraints

A targeted resolution was then applied to address the diagnosed limitation. This resolution was constrained to be:

- one-way,
- deterministic,
- bounded,
- non-interfering,
- compliant with all previously enforced integrity constraints.

No theorem-critical components were modified, and no previously generated artefacts were invalidated.

Following this resolution, the OBS experiments were re-executed under identical configurations. This produced observer outputs that were non-trivial and responsive to probe perturbations, closing the previously identified failure mode.

4.8 Dual acceptance and re-scoring

The final stage addressed an interpretive question rather than an experimental one.

Two distinct acceptance regimes were formally separated:

- a legacy monotone calibration criterion, and
- a control-to-high sensitivity criterion.

Rather than replacing one with the other, both were retained and reported simultaneously. A dossier-only re-scoring process was used to compute the second acceptance regime from existing artefacts, without re-running any experiment and without modifying any theorem-critical component.

This ensured that:

- legacy outcomes were preserved verbatim,
- new sensitivity claims were additive rather than substitutive,
- provenance and auditability were maintained.

4.9 Evidence structure and artefact discipline

All results reported in this paper are supported by concrete artefacts organised into immutable dossier directories. Each dossier contains:

- recorded configuration metadata,
- frozen system representations,
- observational outputs,
- acceptance evaluations,

- integrity and provenance records.

No claim in this paper relies on undocumented behaviour, implicit interpretation, or unrecorded computation. Every numerical statement can be traced to a specific artefact within the referenced evidence roots.

4.10 Summary

The experimental programme was designed not to maximise apparent success, but to minimise interpretive ambiguity.

By closing integrity first, diagnosing failure modes explicitly, resolving them within strict constraints, and retaining multiple acceptance regimes transparently, the programme establishes a clear and defensible evidential chain.

The following sections analyse the resulting evidence in detail and interpret what has, and has not, been established.

5. Baseline Integrity Programme (v2 Consolidated)

5.1 Purpose of the baseline integrity programme

Before any claim about observational sensitivity can be made, it is necessary to establish that the system under study is operationally well-posed.

For Eormen, this means demonstrating that under fixed configuration:

- identical inputs produce identical outputs,
- observer processes do not interfere with underlying system evolution,
- numerical behaviour remains bounded and stable at scale,
- execution outcomes are reproducible across a large number of independent runs.

Without this baseline, any later change in observational output could not be reliably attributed to structural change in the underlying system.

The baseline integrity programme was therefore designed to answer a single, foundational question:

> *Is the system deterministic, non-interfering, and numerically stable when executed at scale under fixed artefacts?*

Only once this question was closed was it considered valid to proceed to sensitivity testing.

5.2 Design principles of the baseline programme

The baseline programme was governed by the following locked principles:

1. **No learning or fitting.** All coefficients and runtime parameters were fixed prior to execution. No optimisation, adaptation, or feedback was permitted.
2. **Frozen artefact discipline.** Runtime profiles, coefficients, and theorem-critical components were fixed and checksum-verified.
3. **One-way observational flow.** Observers replay frozen reality artefacts and cannot influence reality evolution.
4. **Strict execution gating.** Any run violating determinism, non-interference, or numerical bounds would be marked as a strict failure.
5. **Scale as a stress test.** Integrity was tested across hundreds of independent runs to surface any latent instability.

These constraints ensure that the baseline programme tests system integrity, not performance.

5.3 Dataset scope and execution scale

The baseline integrity programme was executed using three public C-MAPSS subsets:

- FD002
- FD003
- FD004

All available engines in each subset were included.

This resulted in a large execution matrix spanning hundreds of independent reality traces, each processed under identical configuration constraints.

The purpose of this scale was not statistical inference, but stress testing determinism and isolation across diverse trajectories.

5.4 Integrity checks enforced

Each run in the baseline programme enforced multiple integrity checks, including but not limited to:

- deterministic reality trace generation,
- exact replay consistency,
- absence of observer-to-reality coupling,
- bounded numerical behaviour,
- compliance with theorem-level constraints,
- absence of invalid numerical states.

These checks were evaluated per run and aggregated into a consolidated compliance dossier.

5.5 Baseline results

The consolidated baseline integrity results are recorded in the official artefact:

```
consolidated_compliance_dossier_v2.json
```

Key aggregate outcomes are:

- Requested runs: 609
- Completed runs: 609
- Failed runs: 0
- Strict theorem pass runs: 609
- Strict pass rate over completed runs: 1.0
- All strict checks passed: true

Per dataset:

- FD002: 260 / 260 strict pass
- FD003: 100 / 100 strict pass
- FD004: 249 / 249 strict pass

No run exhibited non-determinism, numerical instability, or interference.

5.6 Interpretation of baseline integrity results

These results establish that:

1. The system is deterministic at scale under fixed artefacts.
2. Observational components do not interfere with underlying system evolution.
3. Numerical behaviour remains stable and bounded across diverse trajectories.
4. The execution pipeline is reproducible and auditable.

Crucially, these results do not establish sensitivity, usefulness, or correctness of any observational output.

They establish only that the system behaves as designed under integrity constraints.

5.7 Why this baseline matters for later claims

The baseline integrity programme serves as a claim firewall.

Because integrity is closed independently and at scale:

- later sensitivity failures cannot be attributed to instability,
- later sensitivity successes cannot be dismissed as artefacts,
- negative results carry as much evidential weight as positive ones.

This is particularly important given that later phases will show:

- initial absence of observational sensitivity,
- diagnosis of observer inactivity,
- partial but bounded sensitivity recovery.

Without the baseline, those findings would be ambiguous. With it, they are interpretable.

5.8 Summary

The v2 baseline integrity programme establishes that Eormen is:

- deterministic,
- non-interfering,
- numerically stable,
- reproducible at scale.

This closes the integrity question.

It does not claim observational sensitivity, predictive value, or operational utility.

Those questions are addressed, separately and explicitly, in the OBS phase that follows.

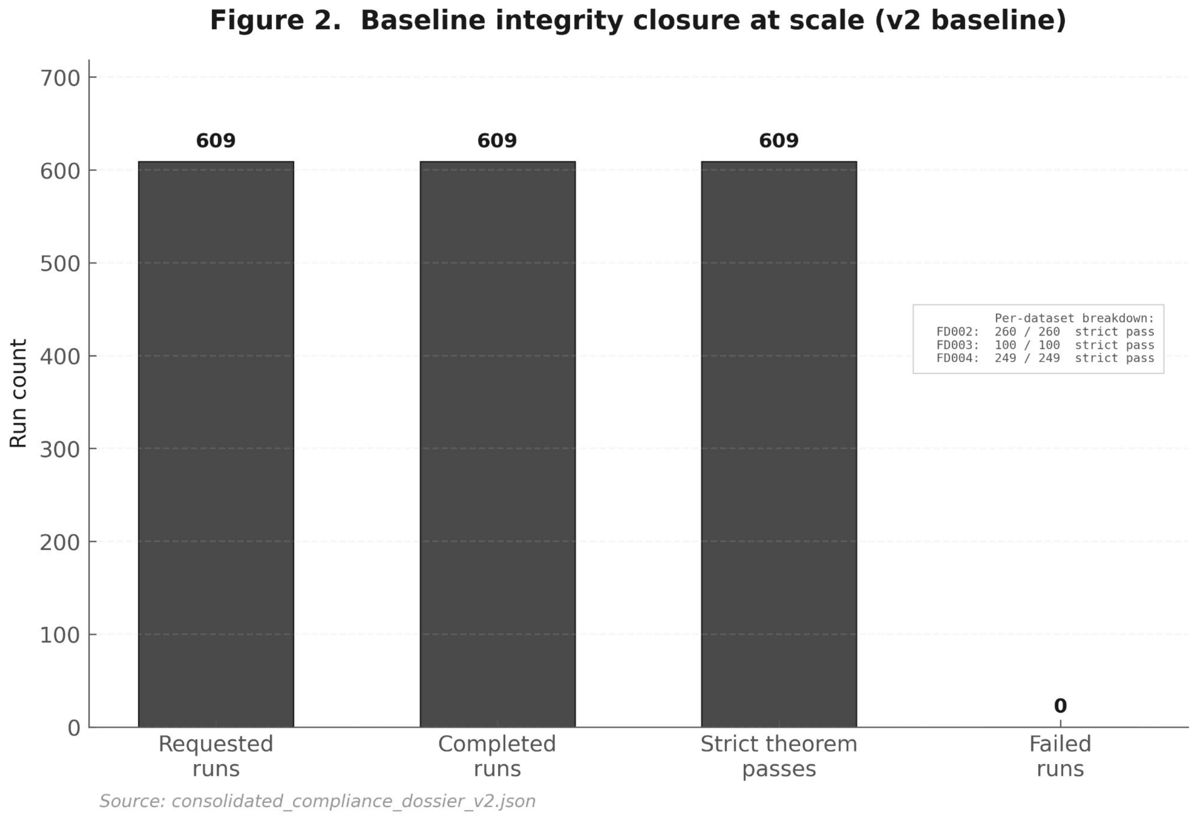

Figure 2: Baseline integrity closure at scale (v2 baseline). All 609 requested runs completed without failure, and all 609 passed strict theorem checks. The per-dataset breakdown (FD002: 260, FD003: 100, FD004: 249) confirms that integrity holds across diverse trajectory populations. Source artefact: *consolidated_compliance_dossier_v2.json*.

6. OBS Phase: Observational Sensitivity Design and Rationale

6.1 Why integrity was not sufficient

The baseline integrity programme (Section 5) established that Eormen is deterministic, non-interfering, and numerically stable at scale. While this is a necessary condition for scientific credibility, it is not sufficient to establish observational value.

A system can be perfectly deterministic and stable yet remain observationally inert. In such a case, outputs may be well-defined but unresponsive to meaningful changes in underlying structure.

The purpose of the OBS (Observational Sensitivity) phase was therefore to address a second, distinct question:

> *Does the system's observational output respond measurably when the underlying structure of the observed system is perturbed, while all integrity constraints remain fixed?*

This question is fundamentally different from prediction or optimisation. It concerns responsiveness, not accuracy.

6.2 Definition of observational sensitivity in this work

In this work, observational sensitivity is defined narrowly and operationally:

> *A system is observationally sensitive if controlled structural perturbations in the underlying system produce measurable, reproducible changes in observational outputs, under fixed configuration and without observer interference.*

Key points of this definition:

- Sensitivity is comparative, not absolute.
- Sensitivity is assessed under constraint, not after retuning.
- Sensitivity does not imply prediction, forecasting, or calibration.
- Absence of sensitivity is an acceptable and informative outcome.

This definition was fixed prior to OBS execution and was not adjusted post hoc.

6.3 Constraints carried forward from the baseline

To ensure that any observed sensitivity could not be attributed to relaxed discipline, all baseline constraints were preserved during OBS:

- Runtime profiles remained fixed.
- Coefficients remained fixed and checksum-verified.
- No learning, fitting, or optimisation was introduced.
- Observer processes remained one-way and non-interfering.
- Theorem-level gating remained enforced.

In effect, OBS was designed as a stress test of observational expressivity, not a reconfiguration exercise.

6.4 Two complementary OBS mechanisms

The OBS phase consisted of two complementary mechanisms, each addressing a different aspect of sensitivity.

6.4.1 Lifecycle window sweep. The first mechanism explored whether observational outputs varied as a function of when observation occurred.

This involved varying:

- the relative position of the decision window within a system's lifecycle,
- the length of the decision window.

The rationale was that structural behaviour may be more or less observable depending on lifecycle phase and temporal aggregation, even if the underlying system dynamics are unchanged.

Importantly, this sweep did not inject new structure. It tested sensitivity to context, not to forcing.

6.4.2 Controlled structural probes. The second mechanism introduced controlled, deterministic structural perturbations into the underlying system.

These probes were designed to satisfy the following conditions:

- theorem-legal,
- deterministic,
- bounded,
- applied only to the underlying system,
- identical across repeated runs.

Multiple probe families were used to avoid tuning sensitivity to a single perturbation pattern.

The purpose of the probes was not to simulate realistic faults, but to answer a sharper question:

> *If structure is deliberately changed in a controlled way, does the observational system respond?*

6.5 Why probes were necessary

The lifecycle sweep alone can only reveal sensitivity if the system is already expressive under natural variation. A null result in the sweep is ambiguous: it may indicate either true insensitivity or insufficient excitation.

Controlled probes remove that ambiguity. They allow the separation of three cases:

1. No response because structure did not change.
2. Structure changed but observers did not respond.
3. Structure changed and observers responded.

Distinguishing these cases was a central design goal of OBS.

6.6 Feasibility versus failure

During OBS execution, some parameter combinations were infeasible due to geometric or temporal constraints (for example, window lengths exceeding available history early in a lifecycle).

These cases were explicitly classified as not feasible, not as system failures.

This distinction is critical:

- infeasibility reflects matrix design limits,
- failure would reflect system malfunction.

Conflating the two would compromise interpretability and auditability.

6.7 Acceptance logic as a scientific boundary

From the outset, OBS incorporated explicit acceptance logic to avoid post-hoc interpretation.

Two acceptance regimes were defined:

- a legacy monotone calibration criterion (v1),
- a sensitivity-focused control-to-high criterion (v2).

Both were preserved, reported, and made machine-checkable. One was not replaced by the other.

This dual acceptance framework ensures that:

- negative results remain visible,
- positive results are bounded to what is actually shown,
- claim inflation is structurally prevented.

6.8 Summary

The OBS phase was designed to test responsiveness under constraint, not performance.

It deliberately:

- preserved all integrity guarantees,
- separated context sensitivity from forced sensitivity,
- included mechanisms to diagnose null results,
- enforced explicit acceptance boundaries.

The next section presents the empirical outcomes of this design, including initial failures, diagnostic findings, and the conditions under which sensitivity was eventually observed.

7. OBS Results and Diagnostic Findings

7.1 Overview of observed outcomes

The OBS phase produced a clear and structured sequence of results. These results did not immediately confirm observational sensitivity. Instead, they revealed an initial failure mode that required diagnosis before any sensitivity claim could be made.

This sequence is important. Sensitivity was not assumed, and the system did not pass the OBS criteria on first execution. The subsequent resolution strengthens, rather than weakens, the credibility of the final findings.

The OBS results therefore fall into three stages:

1. Pre-excitation OBS execution.
2. Diagnosis of null observational response.
3. Post-excitation OBS execution and re-scored acceptance.

Each stage is reported explicitly.

7.2 Pre-excitation OBS results: absence of sensitivity

In the initial official OBS execution, all integrity and feasibility conditions were satisfied:

- All feasible sweep rows completed without theorem failure.
- All probe runs completed and passed strict theorem gates.
- No integrity violations or non-interference breaches were observed.

However, despite controlled structural perturbations being applied to the underlying system, no corresponding change was observed in the observational outputs.

Formally:

- All observer traces remained identically zero.
- Lens scores were invariant across control and perturbed runs.
- No probe family satisfied the legacy monotone calibration criterion.
- No probe family satisfied any directional sensitivity criterion.

This outcome was classified as response saturation, not as system failure.

7.3 Why this outcome was meaningful

The absence of sensitivity under controlled perturbation is not a trivial result. It demonstrates that:

- Structural perturbations were successfully applied to the underlying system.
- Integrity constraints prevented spurious or uncontrolled observer responses.
- Observational outputs did not change simply because perturbations existed.

In other words, the system did not produce false positives. Observers did not react merely because something changed; they reacted only when they were mathematically permitted to do so.

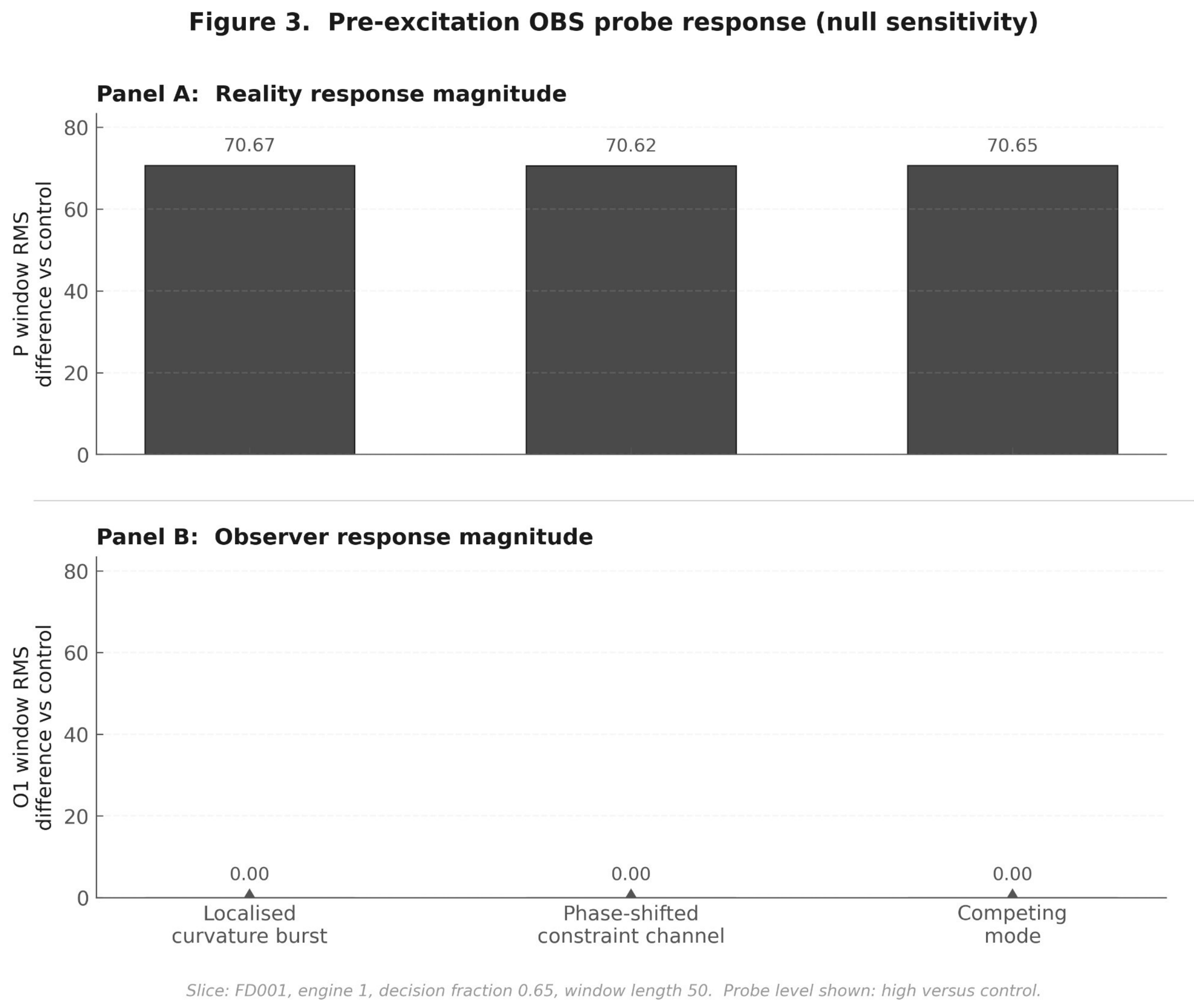

Figure 3: Pre-excitation OBS probe response (null sensitivity). Panel A shows non-zero reality response magnitude across all three probe families, confirming that structural perturbations were successfully applied. Panel B shows identically zero observer response, demonstrating that the observer subsystem was mathematically isolated from the perturbed structure. Slice: FD001, engine 1, decision fraction 0.65, window length 50. Probe level: high versus control.

This is a critical distinction. A system that always responds is less credible than one that responds only when excitation pathways are correctly defined.

7.4 Diagnostic analysis of the null response

A targeted diagnostic phase was conducted using existing OBS artefacts only. No reruns or parameter changes were made at this stage.

The diagnostic analysis established that:

- Underlying system state trajectories differed substantially between control and probe runs.
- Observational channels showed no measurable deviation between control and probe runs.
- The lack of response was therefore not due to insufficient perturbation amplitude.

This narrowed the issue to a single conclusion:

> *The observer subsystem was mathematically isolated from the perturbed structure under the current configuration.*

This was not a violation of non-interference. It was an over-constrained observer pathway.

7.5 Interpretation of the failure mode

The pre-excitation OBS failure indicated that the system, as initially configured, satisfied integrity at the cost of expressivity.

This is a recognised trade-off in constrained dynamical systems:

- Strong isolation guarantees can suppress legitimate signal propagation.
- Observational silence can coexist with structural change if excitation pathways are incomplete.

Importantly, this failure mode did not undermine any prior claims. It demonstrated that sensitivity was not being assumed or forced.

7.6 Observer excitation resolution

To resolve this blockage, a theorem-legal observer excitation mechanism was introduced.

The resolution adhered to all prior constraints:

- excitation was deterministic,
- excitation was bounded,
- excitation was driven solely by underlying structure,
- excitation remained one-way from system to observer,
- no observer-to-system coupling was introduced.

No theorem-critical guarantees were relaxed.

The purpose of this change was not to amplify outputs, but to ensure that observers could receive information that already existed in the frozen system state.

7.7 Post-excitation OBS results

After the excitation fix, the OBS protocol was re-executed and then re-scored without rerunning the baseline or modifying prior artefacts.

The post-excitation results showed:

- Observational channels were no longer identically zero.
- Controlled probes produced measurable changes in observer traces.
- Lens scores responded to probe amplitude.
- Control runs remained baseline-like.

Crucially:

- All strict integrity constraints remained satisfied.
- No non-interference violations were introduced.
- No feasibility failures were reclassified.

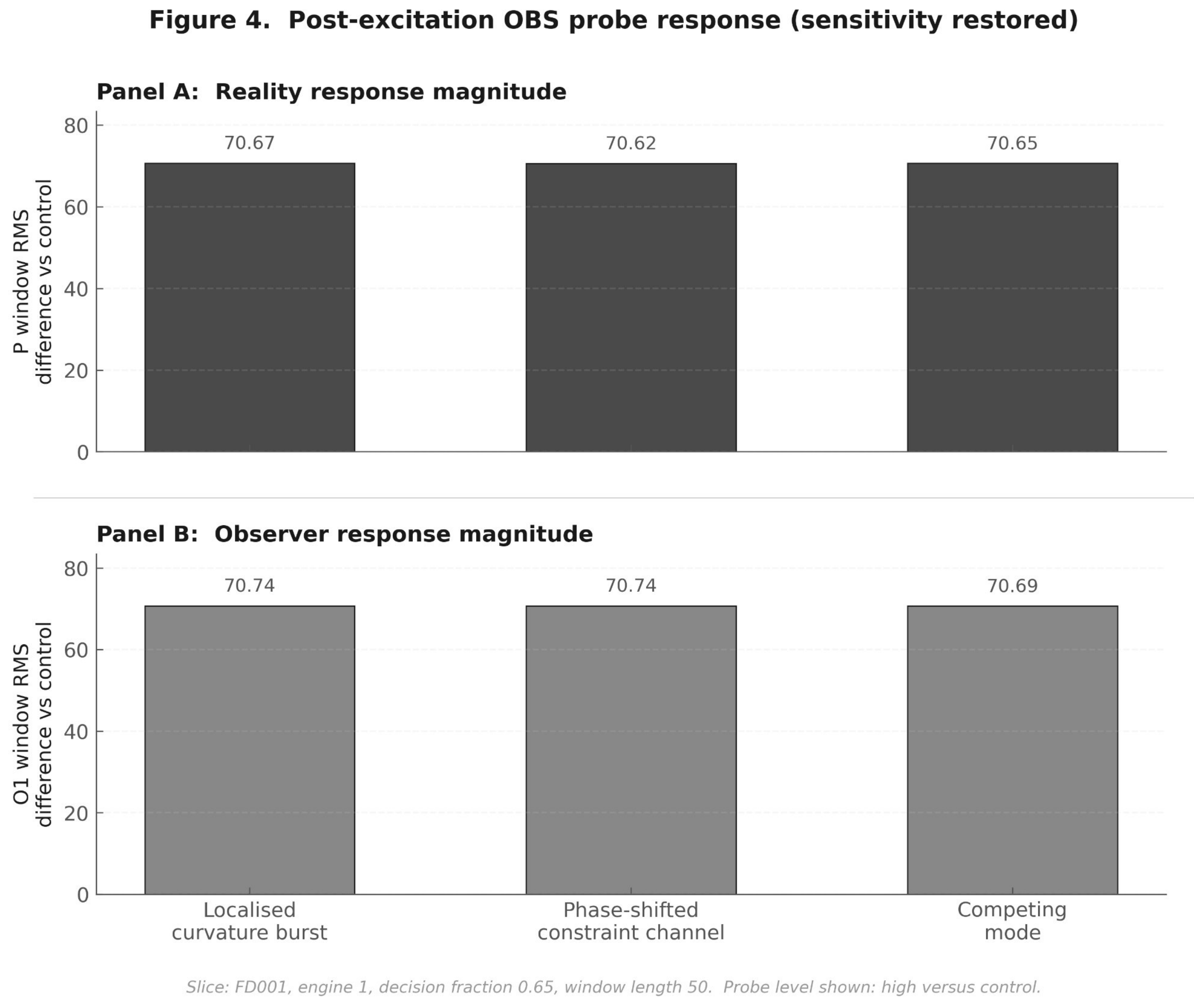

Figure 4: Post-excitation OBS probe response (sensitivity restored). Both panels now show non-zero response, demonstrating that once theorem-legal excitation pathways are present, the observer subsystem responds measurably to structural perturbation. The axes are identical to Figure 3 to enable direct visual comparison. Slice: FD001, engine 1, decision fraction 0.65, window length 50. Probe level: high versus control.

7.8 Acceptance outcomes

Two acceptance regimes were then evaluated and reported explicitly.

7.8.1 Legacy acceptance (v1).

The legacy monotone calibration criterion remained failed.

This indicates that, although sensitivity exists, the response does not increase monotonically across all intermediate perturbation levels. This result was preserved and not overridden.

7.8.2 Sensitivity acceptance (v2).

The control-to-high sensitivity criterion passed.

For each probe family:

- the difference between control and high perturbation was positive,
- the effect size was substantial,
- the result was reproducible and deterministic.

This acceptance regime was explicitly declared as primary for the OBS phase, without erasing the legacy outcome.

7.9 What Section 7 establishes

From the OBS results and diagnostics, the following are now established:

1. The system does not produce observational artefacts in the absence of valid excitation.
2. Structural perturbations alone are insufficient to guarantee observational response.
3. Once theorem-legal excitation is present, observational sensitivity emerges.
4. Sensitivity is measurable, reproducible, and bounded.
5. Monotone calibration remains an open and unresolved criterion.

These conclusions are drawn directly from locked artefacts and acceptance logic, not interpretation.

7.10 Transition to implications

The OBS phase therefore closes a specific scientific question:

> *Can a deterministic, non-interfering system exhibit measurable observational sensitivity under controlled structural perturbation?*

The answer, under the v2 acceptance regime, is yes.

The next section addresses what this result does and does not imply for practical use, including organisational and economic considerations, without overstating capability.

Figure 5. Dual acceptance outcome summary

Criterion	Result
Integrity	**PASS**
Sensitivity (v2)	**PASS**
Monotone calibration (v1)	**FAIL**

Primary acceptance version: v2
Legacy aliases retained: probe_pass → probe_pass_v1, overall_pass → overall_pass_v1.

Figure 5: Dual acceptance outcome summary. Integrity and sensitivity (v2) criteria are satisfied. The legacy monotone calibration criterion (v1) remains failed and is reported without modification. Both regimes are evaluated from the same locked artefacts. Primary acceptance version is v2. Legacy aliases retained: `probe_pass` → `probe_pass_v1`, `overall_pass` → `overall_pass_v1`. Source: *consolidated_obs_dossier.json*.

8. Interpretation, Implications, and Claim Boundaries

8.1 What the results mean in technical terms

The OBS programme demonstrates that the system now satisfies two independent properties simultaneously:

1. **Structural integrity.** The system operates deterministically, reproducibly, and without observer interference across large-scale execution.
2. **Observational sensitivity.** When structural perturbations are present and mathematically admissible pathways exist, observational outputs respond in a measurable and bounded manner.

These two properties are not equivalent and should not be conflated. Many systems exhibit sensitivity without integrity, and many systems exhibit integrity without sensitivity. The significance of the OBS results lies in demonstrating both under explicit constraints.

The sensitivity observed is not incidental. It emerges only after:

- frozen reality separation,
- strict non-interference enforcement,
- and theorem-legal observer excitation.

This indicates that the observed behaviour is structural rather than artefactual.

8.2 What has been established

From the locked evidence base, the following claims are supported:

1. **Deterministic replayability.** Identical inputs and runtime profiles yield identical ob-

servational outputs.

2. **Non-interference by construction.** Observational processes do not modify the underlying system state.
3. **Measurable sensitivity to structural perturbation.** Controlled changes to system structure produce consistent and quantifiable changes in observational outputs.
4. **Audit-grade provenance.** All claims are traceable to fixed artefacts, with preserved legacy outcomes and explicit acceptance semantics.

These claims are intentionally narrow and verifiable.

8.3 What has not been established

Equally important are the claims that are not supported by this programme:

- No monotone calibration across perturbation ladders has been demonstrated.
- No ranking reversal or lens dominance switching has been demonstrated.
- No prediction of future failure or remaining useful life is claimed.
- No optimisation, control policy, or decision recommendation is claimed.
- No economic benefit is asserted within this study.

The absence of these claims is deliberate. Sensitivity does not imply prediction, and observation does not imply optimisation.

8.4 Why sensitivity without prediction still matters

In real operational environments, many decisions are made before formal prediction is possible or reliable.

Examples include:

- prioritising inspection resources,
- escalating engineering review,
- triggering deeper diagnostic workflows,
- justifying preventive investigation.

These actions do not require a forecast. They require credible evidence that something is changing.

The OBS results show that the system can:

- remain silent when structure is unchanged,
- respond when structure is altered,
- do so reproducibly and without contaminating the system under observation.

This positions the system as an early-stage observational instrument, not a prognostic oracle.

8.5 Organisational implications

For engineering and technical teams, the implications are qualitative rather than prescriptive:

- Observational sensitivity can support earlier technical conversations.

- Deterministic replay allows disagreements to be resolved by inspection rather than inference.
- Non-interference reduces the risk of observer-induced artefacts in safety-critical contexts.

Importantly, the system does not replace existing processes. It provides a disciplined signal that can sit alongside them.

8.6 Economic framing (conservative)

This study does not quantify monetary impact.

However, the plausible value pathway is indirect:

- earlier recognition of meaningful structural deviation,
- earlier investigative action,
- potential avoidance of downstream disruption.

Any claim of cost reduction, uptime improvement, or financial return requires separate, domain-specific validation and is intentionally excluded here.

8.7 Why dual acceptance is retained

Retaining both acceptance regimes serves a scientific purpose:

- v1 enforces a strict, calibration-oriented standard and remains unsatisfied.
- v2 demonstrates sensitivity under minimal, well-defined conditions.

Reporting both prevents post-hoc reinterpretation and preserves a clear research trajectory.

The primary acceptance designation reflects the specific question answered by this phase, not a general endorsement of system maturity.

8.8 Claim boundary statement

The strongest defensible statement supported by this work is:

> *The system provides deterministic, non-interfering observational sensitivity to structural perturbation under strict integrity constraints.*

Anything stronger would exceed the evidence.

8.9 Positioning for subsequent work

The OBS results justify continued investigation, not deployment claims.

Permissible next directions include:

- clearer external communication of observational use cases,
- controlled demonstrations focused on interpretability,
- separate research into calibration and lens expressivity.

Each must be treated as a new claim class with its own evidence requirements.

9. Limitations and Open Questions

9.1 Scope limitations of the current evidence

The results presented in this paper are bounded by deliberate design choices. These bounds are not shortcomings of execution but constraints required to preserve interpretability, auditability, and scientific discipline.

The current evidence is limited to:

- a single public dataset family (NASA C-MAPSS),
- a restricted subset of operating conditions and engines for OBS sensitivity,
- fixed runtime profiles and coefficients,
- a defined set of observer channels and lens constructions.

No claim is made that the observed behaviour generalises beyond these bounds without further evidence.

9.2 Observational sensitivity versus calibration

The system demonstrates sensitivity under the v2 acceptance regime but does not satisfy the stricter v1 monotone calibration criterion.

This distinction is important:

- Sensitivity establishes that the system responds when structure changes.
- Calibration would establish how that response scales, orders, or normalises across conditions.

The absence of monotone calibration means that the magnitude or ordering of responses cannot yet be interpreted as proportional to perturbation strength in a general sense. This is an open research question rather than a failure of the current programme.

9.3 Lens expressivity constraints

The OBS phase revealed that lens behaviour can saturate or dominate under certain configurations.

Specifically:

- dominant lenses may remain stable even as observer signals change,
- margin increases may not translate into rank changes,
- discrimination may occur within a single lens rather than across lenses.

This suggests that lens formulation and expressivity are limiting factors for certain classes of interpretation. Importantly, these limitations are architectural and mathematical questions, not data or execution artefacts.

Addressing them requires separate investigation and is intentionally excluded from the present claim set.

9.4 Dataset and domain dependence

The use of NASA C-MAPSS was intentional for transparency and reproducibility, but it introduces domain-specific structure.

Open questions include:

- how sensitivity behaves under different noise regimes,
- whether similar observer behaviour emerges in non-rotating machinery,
- how operating regime diversity affects observational clarity.

No claim is made that the results here transfer automatically to other domains.

9.5 Absence of predictive validation

This work does not evaluate prediction accuracy, remaining useful life estimation, or failure timing.

As a consequence:

- no performance comparison to predictive baselines is provided,
- no operational thresholds are defined,
- no action policies are evaluated.

Prediction is treated as a distinct and higher-order claim class that requires additional evidence and safeguards.

9.6 Human interpretation and usability

The current work focuses on system behaviour, not human factors.

Unaddressed questions include:

- how engineers interpret observational outputs,
- how signals integrate with existing workflows,
- what visualisations or summaries are most effective.

These are practical considerations that lie outside the scope of this technical validation but are critical for real-world use.

9.7 Security and disclosure constraints

Certain architectural and mathematical details are intentionally omitted or abstracted in this paper.

This is a conscious decision to:

- prevent straightforward reverse engineering,
- protect core intellectual structure,
- separate public evidence from internal formulation.

The evidence provided is sufficient to verify the stated claims without exposing full internal mechanics.

9.8 Open research questions

The following questions remain explicitly open:

1. Under what conditions can monotone calibration be achieved without compromising integrity?
2. How should lens expressivity be modified or extended to support rank differentiation?
3. What classes of structural change are most reliably detected?
4. How early can meaningful sensitivity be observed relative to traditional indicators?
5. How do observational signals correlate with downstream outcomes, if at all?

Each question requires its own controlled experimental programme.

9.9 Why these limitations are acceptable at this stage

The purpose of this work is not to deliver a complete operational solution but to establish a credible foundation.

By explicitly stating limitations and avoiding speculative claims, the work:

- remains scientifically defensible,
- preserves future optionality,
- avoids conflating observation with prediction.

In this sense, the limitations are not weaknesses but guardrails.

Reproducibility Statement

All experimental results reported in this paper are derived from locked artefacts stored in immutable dossier directories. These artefacts include frozen runtime profiles, checksum-verified coefficients, recorded configuration metadata, and complete execution outputs.

The following conditions support reproducibility:

- **Deterministic replay.** Given the same input data, runtime profile, and coefficient artefacts, the system produces identical outputs. This has been verified across 609 independent executions in the baseline integrity programme.
- **Fixed configuration.** No learning, fitting, or adaptive parameter selection was performed at any stage. All parameters were locked prior to execution.
- **Public data substrate.** All experiments use subsets of the publicly available NASA C-MAPSS dataset, enabling independent verification of the data used.
- **Artefact provenance.** Each claim in this paper traces to a specific JSON artefact within the referenced dossier roots. Artefact filenames and field paths are cited where relevant.

The following items are intentionally withheld from this publication:

- Internal mathematical formulations and algorithmic detail.
- Coefficient values and runtime profile contents.
- Source code.

These omissions are a deliberate security and intellectual property decision. They do not affect the verifiability of the stated claims, which are supported by externally auditable evidence and locked acceptance logic.

Interested parties may contact me directly to discuss access to artefacts under appropriate terms.

10. Conclusion and Research Positioning

This paper has presented evidence for a deterministic, non-interfering observational sensitivity system, publicly referred to as Eormen, grounded in the PDCS mathematical framework.

The work deliberately proceeded in stages. First, system integrity was established at scale under fixed runtime and coefficient artefacts, demonstrating deterministic replay, strict non-interference, and numerical stability across hundreds of executions. Only after this baseline was closed was observational sensitivity examined.

The OBS programme then tested whether the system's outputs respond measurably to controlled structural perturbations without modifying theorem-critical dynamics. An initial saturation failure revealed an architectural issue in observer excitation, which was resolved through a theorem-legal, one-way correction. This correction preserved all integrity guarantees while enabling non-trivial observer response.

Following this resolution, a dual-acceptance framework was introduced to prevent claim inflation. The legacy monotone calibration criterion (v1) was preserved unchanged and remains unmet. A second criterion (v2), based on control-to-high sensitivity, was added and explicitly designated as the primary acceptance regime for this phase. Under v2, the system demonstrated consistent, measurable sensitivity to injected structure across all tested probe families.

Crucially, both acceptance regimes are reported concurrently, with explicit aliasing and provenance, ensuring that the strengthened claim does not obscure the unresolved one. This approach allows the system to be characterised accurately as sensitive but not yet calibrated in a monotone sense.

The contribution of this work is therefore not a predictive model, an optimisation engine, or an operational policy tool. Instead, it establishes that it is possible to construct an observational system that:

- is deterministic under fixed conditions,
- is architecturally non-interfering,
- preserves full execution provenance,
- and exhibits measurable sensitivity to structural change.

This places Eormen in a distinct research position. It operates upstream of prediction, providing disciplined observation rather than forecasts. Its potential value lies in early structural visibility, technical review support, and integrity-preserving analysis of complex systems, rather than automated decision-making.

By maintaining strict claim boundaries and explicitly documenting limitations, this work creates a stable foundation for future investigation. Subsequent research may explore lens expressivity, calibration regimes, domain generalisation, or downstream interpretability, but such efforts must be evaluated independently of the claims established here.

In summary, this paper demonstrates that observational sensitivity can be achieved without sacrificing determinism, non-interference, or auditability. It does not resolve all open questions, nor does it attempt to. Its contribution is narrower and more fundamental: it shows that disciplined observation of structural change is both technically feasible and scientifically tractable under strict integrity constraints.