

# The Mena Dominance Framework: From Margin and Dose to a Viability-Governed Science of Risk and Control

Arlex Orlando Murcia Mena

January 2026

## Abstract

This paper is the framework and methodological companion to Paper 1, which stated the Mena Dominance Law, and to Paper 2, which reported embodiment evidence under fixed boundary rules. It does not restate the law as a new claim and it does not add new validation data. Its purpose is to specify how the law should be instantiated, what kinds of claims different embodiments can support, how dominance trajectories should be reported, and which control and governance patterns follow when the dominance margin is treated as an operational signal.

The central object remains the time-tracked dominance margin

$$\Delta_X(t) = P_X(t) - L_X(t),$$

where  $P_X(t)$  is recoverable potential and  $L_X(t)$  is imposed load, defined in commensurate units over a declared sampling window.

This paper makes six methodological contributions. It formalizes a minimal embodiment contract that blocks arbitrary proxy instantiations while preserving cross-domain generality; states a claim ladder separating physical, meta-operational, and interpretive-only uses; clarifies boundary-faithful comparison as the governing rule for policy evaluation; extracts reusable control and governance patterns already implied by the law; specifies standard reporting metrics for fixed-boundary comparisons; and states falsifiable claim classes together with a forward research agenda.

The thesis of this paper is methodological: the scientific value of the Mena Dominance Law depends on disciplined embodiment, declared uncertainty, actuator definition where control is claimed, claim scope, and boundary-faithful comparison.

**Keywords:** Mena Dominance Framework, Mena Dominance Law, dominance margin, embodiment contract, deficit dose, Grace, Mena Horizon, boundary-faithful comparison, layered governance, recoverability, claim discipline

# Contents

<b>1</b>	<b>Role of This Paper in the Trilogy</b>	<b>4</b>
<b>2</b>	<b>Scope and Tasks of This Paper</b>	<b>4</b>
<b>3</b>	<b>Kernel Recap (Referenced, Not Re-Derived)</b>	<b>5</b>
3.1	Rate form . . . . .	5
3.2	Budget–dose form . . . . .	5
<b>4</b>	<b>Status of the Law Versus Status of an Embodiment</b>	<b>6</b>
<b>5</b>	<b>Boundary-Faithful Evaluation</b>	<b>6</b>
<b>6</b>	<b>Why Methodological Discipline Is Necessary</b>	<b>7</b>
<b>7</b>	<b>Minimal Embodiment Contract</b>	<b>7</b>
7.1	Additional requirement for meta-operational embodiments . . . . .	8
7.2	Enforcement rule . . . . .	8
<b>8</b>	<b>Embodiment Tiers and Claim Discipline</b>	<b>8</b>
8.1	Tier A: Physical embodiments . . . . .	8
8.2	Tier B: Meta-operational embodiments with tether . . . . .	9
8.3	Tier C: Interpretive-only uses . . . . .	9
<b>9</b>	<b>Core Operators for Analysis and Control</b>	<b>10</b>
9.1	Postulates and propositions . . . . .	10
9.2	Margin as headroom . . . . .	11
9.3	Dose . . . . .	11
9.4	Grace . . . . .	11
9.5	First-crossing times . . . . .	12
9.6	Mena Horizon . . . . .	12
9.7	Restoration burden . . . . .	12
<b>10</b>	<b>Design Patterns for Dominance-Governed Control</b>	<b>12</b>
10.1	Pattern 1: Boundary-respecting load shaping . . . . .	13
10.2	Pattern 2: Dual-channel routing . . . . .	13
10.3	Pattern 3: Layered admissibility gating . . . . .	13
10.4	Pattern 4: Horizon-triggered mode switching . . . . .	13
10.5	Pattern 5: Dose-aware recovery governance . . . . .	14
<b>11</b>	<b>Standard Reporting Metrics</b>	<b>14</b>
<b>12</b>	<b>Misuse Modes and Guardrails</b>	<b>15</b>
12.1	Proxy arbitrariness . . . . .	15
12.2	Goodhart and gaming . . . . .	15
12.3	Punitive ranking misuse . . . . .	15
12.4	Causal overclaim . . . . .	15
12.5	Hidden-layer depletion . . . . .	15
12.6	Boundary drift . . . . .	15

<b>13 Falsifiable Claim Classes</b>	<b>16</b>
13.1 Claim A: Policy-conditional improvement . . . . .	16
13.2 Claim B: Dose and restoration burden . . . . .	16
13.3 Claim C: Layered-gating advantage . . . . .	16
13.4 Claim D: Grace exhaustion as a measurable transition . . . . .	16
13.5 Claim E: Embodiment failure condition . . . . .	16
13.6 Failure modes of the framework . . . . .	17
<b>14 Research Agenda</b>	<b>17</b>
14.1 Physical embodiments . . . . .	17
14.2 Meta-operational embodiments . . . . .	17
14.3 Methodological agenda . . . . .	17
14.4 The real job of future work . . . . .	18
<b>15 Conclusion</b>	<b>18</b>
<b>A Notation and Glossary</b>	<b>19</b>
<b>B Embodiment Card (Minimum Reporting Template)</b>	<b>19</b>
<b>C Minimum Implementation Checklist</b>	<b>19</b>

# 1 Role of This Paper in the Trilogy

The trilogy is organized as follows:

- **Paper 1:** the law paper, stating the Mena Dominance Law, its trajectory logic, Grace structure, deficit-sensitive deterioration logic where supported, and embodiment requirements.
- **Paper 2:** the evidence paper, testing whether dominance-governed control changes measurable outcomes under fixed boundary rules in executed embodiments.
- **Paper 3 (this paper):** the framework and methodological-discipline paper, stating how the law should be instantiated, reported, governed, and protected from misuse or overclaim.

Paper 1 carries the burden of the law claim. Paper 2 carries the burden of executed embodiment evidence. Paper 3 therefore neither re-proves the law nor repeats the validation work. Its task is narrower and stricter: it states the conditions under which future embodiments may make serious claims under the Mena Dominance Law.

Once a law is written in a domain-invariant form,

$$\Delta_X(t) = P_X(t) - L_X(t),$$

it becomes easy to instantiate it everywhere by analogy. This paper provides the discipline required to constrain that use. The law is general in form; embodiments are not automatic. The law is cross-domain; evidentiary standing is not flat across domains. The framework addresses that problem through embodiment discipline.

The role of this paper is therefore not to expand the law rhetorically, but to constrain its admissible use. It supplies a framework standard: a gatekeeping structure that distinguishes scientifically serious instantiations from loose analogy, distinguishes operational utility from physical-embodiment standing, and separates strong claims from exploratory ones.

## 2 Scope and Tasks of This Paper

This paper has six primary tasks.

1. It states a minimal embodiment contract required for serious use of the law.
2. It distinguishes physical embodiments, meta-operational embodiments, and interpretive-only uses.
3. It defines the standard trajectory operators used for analysis and control.
4. It extracts controller and governance patterns already implied by the law and illustrated in Paper 2.
5. It specifies reporting standards for fixed-boundary comparisons.
6. It states falsifiable claim classes and a forward research agenda.

This paper does *not* attempt four things.

1. It does not re-prove the Mena Dominance Law.
2. It does not replace Paper 1 as the carrier of the law claim.

3. It does not declare every deployment to be a physical embodiment.
4. It does not replace reliability analysis, survival analysis, control theory, or other domain tools.

The correct reading is narrower and stronger: the law supplies a governing margin trace; this paper supplies the admissibility rules for using that trace rigorously.

### 3 Kernel Recap (Referenced, Not Re-Derived)

The core object is the dominance margin

$$\Delta_X(t) = P_X(t) - L_X(t), \tag{1}$$

where  $P_X(t)$  is recoverable potential and  $L_X(t)$  is imposed load in system  $X$ , both defined in commensurate units over a declared estimation window.

The law becomes scientifically useful when the margin is treated as a trajectory rather than a snapshot. The trajectory logic from Paper 1 supplies at least five classes of outputs:

- first-crossing times to warning, critical, and irreversibility thresholds;
- deficit time and deficit area;
- deficit dose and near-boundary exposure summaries;
- Grace, understood as finite recoverability under admissible intervention;
- conditional countdowns such as the Mena Horizon in budget–dose embodiments.

Two recurrent commensurate forms appear throughout the trilogy.

#### 3.1 Rate form

When potential and load are naturally expressed as matched rates,

$$\Delta_R(t) = P_R(t) - L_R(t). \tag{2}$$

Examples include power margin, throughput margin, reserve margin, metabolic-power margin, or response-capacity margin.

#### 3.2 Budget–dose form

When a stored budget is depleted by a rate of use,

$$\Delta_E(t) = E(t) - \int_{t_0}^t \dot{E}_{\text{load}}(\tau) d\tau, \tag{3}$$

where  $\dot{E}_{\text{load}}$  is defined so that its integral shares units with  $E(t)$ . In such embodiments, margin can be converted into a conditional time-to-boundary under declared load assumptions.

The scientific point is not the algebraic identity alone. The point is that once a bounded system admits a declared viability boundary and a commensurate governing channel, the resulting margin trace supplies a compact trajectory grammar for warning, deficit, recoverability, and control.

## 4 Status of the Law Versus Status of an Embodiment

A central methodological distinction must be maintained:

The law and the standing of a particular embodiment are not the same thing.

Paper 1 states the Mena Dominance Law in general form. That generality does *not* mean every proposed instantiation immediately inherits the same evidentiary strength. A weak embodiment does not weaken the law; it weakens the embodiment. A strong embodiment does not make every analogy scientific; it shows that the law is scientifically productive when the contract is satisfied.

This distinction avoids two opposite errors:

1. **Collapse by overextension:** treating every proxy construction as if it were a physical embodiment.
2. **Collapse by dilution:** assuming that because some deployments are operational or proxy-based, the law itself is merely interpretive.

Neither is correct. The law is one thing; claim strength is embodiment-dependent.

The framework therefore adopts a doctrine of split status:

1. **Status of the law:** fixed by the statement and scope of Paper 1.
2. **Status of the embodiment:** earned by measurement discipline, reproducible construction, declared uncertainty, and falsifiability.
3. **Status of the claim:** limited by the tier and evidence quality of the embodiment.

This split is what allows the Mena Dominance Law to remain universal in operational form without pretending that every use of the law has the same scientific standing.

## 5 Boundary-Faithful Evaluation

The governing comparative rule of the framework is boundary-faithful evaluation: policy comparisons are meaningful only when violation rules, thresholds, cadence, and endpoint definitions are held fixed.

A central lesson of Paper 2 is that policy comparisons are scientifically interpretable only when the boundary is held fixed. If the policy changes and the boundary also changes, the comparison is confounded.

This paper therefore adopts the following rule:

When comparing policies under the Mena Dominance Law, boundary definitions, violation rules, logging cadence, and endpoint definitions must remain fixed unless the explicit object of study is a boundary redesign.

This rule preserves interpretability because it separates the effect of control from the effect of changed standards. It also makes cross-policy comparison portable across domains. A battery cutoff, a reserve floor, a thermal limit, an error budget, or a governance threshold can all be handled in the same grammar once the boundary is fixed.

Boundary-faithful evaluation is not a stylistic preference. It is an admissibility condition. Comparisons that change the boundary while claiming policy advantage are outside MDL-grade evidence unless the redesign of the boundary is itself the declared object of study.

## 6 Why Methodological Discipline Is Necessary

A domain-invariant law faces two opposite failure modes.

The first is underspecification. Users may treat the law as an inspiring metaphor, define arbitrary proxies for potential and load, and claim predictive status without stating a pipeline, uncertainty band, or falsification condition. That turns a law into rhetoric.

The second is overclaim. Users may observe that the law is physically meaningful in some embodiments and then assume that all deployments inherit the same evidentiary standing. They do not.

This paper therefore enforces three principles.

1. **Law status is preserved.** The law claim remains the one stated in Paper 1; Paper 3 governs how embodiments may legitimately use it.
2. **Embodiment status is earned.** A claimed instantiation must specify how  $P_X(t)$ ,  $L_X(t)$ , and  $\Delta_X(t)$  are constructed and tested.
3. **Claim scope is tiered.** Strong claims attach only where the embodiment warrants them.

Methodological discipline is therefore not secondary to the framework. It is the mechanism by which a general law avoids collapse into uncontrolled analogy, un-auditable proxy use, or borrowed credibility.

## 7 Minimal Embodiment Contract

For an embodiment to support scientific or operational claims under the Mena Dominance Law, it must satisfy the following contract.

- (E1) **Boundary declaration.** State the viability boundary, the operational meaning of violation, and the detection rule.
- (E2) **Channel definition.** Define  $P_X(t)$  and  $L_X(t)$  in commensurate units over the same declared window.
- (E3) **Observables.** List the measured signals, data sources, and sampling cadence.
- (E4) **Estimator or pipeline.** Specify how observables are transformed into  $P_X(t)$ ,  $L_X(t)$ , and  $\Delta_X(t)$ , including windows, filters, hold assumptions, and any re-anchoring logic.
- (E5) **Uncertainty.** Declare uncertainty or resolution  $\epsilon_X$  for  $\Delta_X(t)$ . Near-boundary claims are indeterminate when  $|\Delta_X(t)| \leq \epsilon_X$ .
- (E6) **Dose operator.** If dose is used, define the operator and justify the choice of  $\phi$  in

$$D_X(t) = \int_{t_0}^t \phi(\max(0, -\Delta_X(\tau))) d\tau.$$

- (E7) **Intervention frame.** If control, governance, or recoverability claims are made, identify at least one admissible actuator  $u(t)$  or intervention-relevant evaluation frame and state the admissibility constraints.
- (E8) **Outcomes and evaluation protocol.** Predeclare outcomes, comparison rules, and control-cost reporting under fixed boundary definitions.
- (E9) **Falsifiability hook.** State an explicit observation that would break the embodiment mapping.

## 7.1 Additional requirement for meta-operational embodiments

If  $P_X(t)$  and  $L_X(t)$  are operational proxies rather than direct physical measures, the mapping must be tethered to measurable consequences such as restoration time, backlog growth, resource depletion, incident rates, error accumulation, loss events, or comparable observable burdens. Without such a tether, the use is interpretive only.

## 7.2 Enforcement rule

Failure of any required element blocks higher-tier claims for that embodiment.

That enforcement rule is what prevents arbitrary proxy constructions from borrowing the credibility of stronger embodiments.

# 8 Embodiment Tiers and Claim Discipline

Not all embodiments support the same claims. The law is one thing; the evidentiary standing of a given instantiation is another.

Table 1: Claim ladder for embodiments of the Mena Dominance Law.

Tier	Embodiment type	Minimum condition	Allowable claim strength
Tier A	Physical embodiment	Commensurate physical units, declared measurement pipeline, uncertainty, and falsification hook	Direct physical-domain MDL application claims, trajectory claims, and boundary claims within the declared embodiment
Tier B	Meta-operational embodiment with tether	Reproducible proxy pipeline tied to measurable consequences and declared thresholds	Operational routing, classification, countdown, burden, and policy-conditional claims within the tethered operational frame
Tier C	Interpretive-only use	Conceptual or exploratory mapping without serious tether, uncertainty discipline, or falsification	Explanatory or heuristic use only; no strong predictive or physical-embodiment claims

## 8.1 Tier A: Physical embodiments

In Tier A,  $P_X(t)$  and  $L_X(t)$  are defined in commensurate physical units with a declared measurement or estimation pipeline and stated uncertainty.

### Allowable claims.

- direct physical-domain MDL application claims within the declared embodiment;
- time-to-boundary outputs;

- physical deficit time, area, and dose comparisons;
- Grace and response-based recoverability statements when an admissible intervention frame is declared;
- deficit-sensitive restoration-burden or deterioration claims when explicitly tested.

**Control-claim requirement.** Control, governance, or response-based recoverability claims within Tier A additionally require an admissible actuator or intervention-relevant evaluation frame.

**Not automatically allowed.** Tier A does not authorize a universal mechanism claim beyond the declared embodiment. It authorizes MDL-based physical-domain trajectory and control claims within the declared embodiment, not a universal mechanism claim beyond it.

## 8.2 Tier B: Meta-operational embodiments with tether

In Tier B,  $P_X(t)$  and  $L_X(t)$  are operational proxies, thresholds are governance-defined, and the mapping is tethered to measurable outcomes.

**Allowable claims.**

- routing logic and mandate assignment;
- classification stability and robustness;
- policy-conditional outcome differences under fixed rules;
- operational countdowns and recoverability statements when the tether is explicit.

**Not allowable without additional evidence.**

- direct physical-embodiment claims in the literal measurement sense;
- mechanistic physical causality claims;
- claims that proxy variables are physical observables by analogy alone.

Tier B embodiments may support serious operational claims, but they do not inherit Tier A status by analogy alone.

## 8.3 Tier C: Interpretive-only uses

In Tier C, the proxies lack a tether, the pipeline is not reproducible, or the uncertainty and falsification conditions are missing.

**Allowable claims.** Qualitative framing, exploratory monitoring, and conceptual interpretation only.

**Not allowable.** Predictive irreversibility claims, deficit-acceleration claims, strong policy comparisons, and direct physical-embodiment language.

**Critical boundary sentence.** The Mena Dominance Law is one law. Direct physical-embodiment claims in a specific deployment attach only to Tier A embodiments. Tier B and Tier C do not inherit Tier A status by analogy.

## 9 Core Operators for Analysis and Control

The framework uses a hierarchy of operators.

**Level 1: Core operators.** The indispensable objects are the dominance margin, declared boundaries, and first-crossing times.

**Level 2: Conditional but central operators.** Dose, Grace, and restoration burden are central where the embodiment supports them, but they require additional specification.

**Level 3: Derived operators and governance constructions.** The Mena Horizon, layered gating, and horizon-triggered policy structures are derived objects. They are useful and often powerful, but they do not sit at the same logical level as the margin itself.

### 9.1 Postulates and propositions

The framework also adopts a minimal formal layer. These statements are not presented as deep theorems of substrate physics. They are framework postulates and propositions governing admissible use of the dominance trace.

**Postulate 1 (Declared boundary).** A bounded system eligible for MDL analysis admits a declared viability boundary together with an operational detection rule.

**Postulate 2 (Commensurate governing channel).** An admissible embodiment defines  $P_X(t)$  and  $L_X(t)$  in commensurate units over a declared observation or estimation window.

**Postulate 3 (Trajectory primacy).** The framework acts on the time-tracked trace  $\Delta_X(t)$ , not on a static snapshot alone.

**Proposition 1 (Threshold-crossing well-posedness on sampled traces).** Given a declared sampling cadence, threshold basis, and hold rule, first observed crossings of warning, critical, or irreversibility thresholds are operationally well-defined on sampled traces without requiring smoothness of  $\Delta_X(t)$ .

**Proposition 2 (Dose monotonicity under sign-consistent operator).** If  $\phi(\cdot) \geq 0$  and is nondecreasing on  $[0, \infty)$ , then larger or longer deficit exposure cannot produce a smaller dose under

$$D_X(t) = \int_{t_0}^t \phi(\max(0, -\Delta_X(\tau))) d\tau.$$

**Proposition 3 (Indeterminacy inside the uncertainty band).** If  $|\Delta_X(t)| \leq \epsilon_X$ , then boundary-adjacent claims at that instant are indeterminate unless additional resolution or evidence is declared.

**Proposition 4 (Inadmissibility of boundary-shifted policy comparisons).** A claimed policy advantage is not MDL-grade evidence if boundary definitions, violation rules, cadence, or end-point definitions differ across compared policies, unless boundary redesign is itself the declared object of study.

## 9.2 Margin as headroom

The operational meaning of  $\Delta_X(t)$  is headroom in the governing channel. Standard trajectory summaries include

$$\text{Time in deficit} = \int \mathbf{1}[\Delta_X(t) < 0] dt, \quad (4)$$

$$\text{Deficit area} = \int \max(0, -\Delta_X(t)) dt, \quad (5)$$

and, for a warning margin  $\tau > 0$ ,

$$\text{Near-boundary exposure} = \int \max(0, \tau - \Delta_X(t)) dt. \quad (6)$$

A static margin says whether the system is acceptable now. A trajectory margin says how close the system is to crossing a boundary, how long it has been near the boundary, and whether the state is becoming harder to recover.

## 9.3 Dose

A general nonnegative dose operator is

$$D_X(t) = \int_{t_0}^t \phi(\max(0, -\Delta_X(\tau))) d\tau, \quad \phi(\cdot) \geq 0. \quad (7)$$

The universal structural requirement is sign consistency: larger or longer deficit exposure must yield larger dose. Beyond that, the operator is embodiment-defined. A linear  $\phi$  yields deficit area. A thresholded  $\phi$  can ignore noise-floor deficits. A convex  $\phi$  can penalize severe deficits more heavily than shallow ones.

## 9.4 Grace

Grace is the finite recoverability region in which admissible intervention can still restore margin meaningfully.

A response-based form useful for control is

$$G_X(t; R, \Delta t) = \Delta_X(t + \Delta t | R) - \Delta_X(t + \Delta t | R_0), \quad (8)$$

where  $R$  is an admissible intervention and  $R_0$  is a baseline policy.

Grace exists at time  $t$  if there exists some admissible  $R$  such that

$$G_X(t; R, \Delta t) \geq g_{\min, X},$$

where  $g_{\min, X}$  exceeds both measurement noise and operational irrelevance.

Grace is an operational statement about remaining recoverability under admissible intervention. It is not a metaphysical reserve. It is a measurable claim about what admissible action can still accomplish.

## 9.5 First-crossing times

Let  $\tau_{w,X}$ ,  $\tau_{c,X}$ , and  $\tau_{r,X}$  denote warning, critical, and irreversibility thresholds in the embodiment. Then the standard first-crossing times are

$$t_{w,X} = \inf\{t : \Delta_X(t) \leq \tau_{w,X}\}, \quad (9)$$

$$t_{c,X} = \inf\{t : \Delta_X(t) \leq \tau_{c,X}\}, \quad (10)$$

$$t_{r,X} = \inf\{t : \Delta_X(t) \leq \tau_{r,X}\}. \quad (11)$$

These are embodiment-level predictive outputs and must never be reported without a declared threshold basis and uncertainty context.

Under sampled data, these outputs remain operationally well-defined as first observed crossings of a declared thresholded trace. Smoothness is not required for the reporting of threshold crossings, but cadence and hold assumptions must be declared.

## 9.6 Mena Horizon

In budget-dose embodiments, the margin becomes a conditional countdown. Under locally constant load rate,

$$T_M \approx \frac{\Delta_E(t)}{\dot{E}_0}. \quad (12)$$

Under locally linear load rate,

$$T_M = \frac{-\dot{E}_0 + \sqrt{\dot{E}_0^2 + 2\alpha \Delta_E(t)}}{\alpha}, \quad \alpha \geq 0. \quad (13)$$

The Mena Horizon is a conditional time-to-boundary under declared local assumptions. Reported horizon values must always be accompanied by the local load model under which they were computed.

The framework treats  $T_M$  as a derived control object, not as a replacement for first-crossing logic. It is useful because it turns headroom into a control-ready countdown, but its validity is local and assumption-dependent.

## 9.7 Restoration burden

Where the embodiment supports it, dose may be linked to restoration cost, restoration time, capacity loss, or another observable burden. That is the operational bridge to Reciprocal Decay. Not every embodiment is required to claim this coupling, but any embodiment that does must declare the observable burden signature, control the relevant confounds, and state the falsification rule.

# 10 Design Patterns for Dominance-Governed Control

Once  $\Delta_X(t)$  is treated as a control signal, recurrent control and governance patterns appear across embodiments. These are not new axioms. They are reusable design patterns extracted from the law plus the embodiment contract.

## 10.1 Pattern 1: Boundary-respecting load shaping

**Intent:** reduce effective load as the system approaches a declared boundary while keeping the boundary fixed.

**General form:**

$$u(t) = f(\Delta_X(t)),$$

where the actuator progressively unloads, throttles, or reshapes demand as  $\Delta_X(t) \rightarrow 0^+$ .

**Use cases:** battery current reduction near cutoff, thermal throttling near a thermal floor, throughput shaping near an error-budget boundary, extraction-rate shaping near a reservoir floor.

## 10.2 Pattern 2: Dual-channel routing

**Intent:** prevent corrective pressure from being applied to the wrong channel.

In some systems the practical problem is not only low margin, but misidentified cause. The framework supports routing logic by defining margins on separate channels and using their relative dominance to determine where intervention should land.

**General proxy form:**

$$\Delta_{\text{route}}(t) = P_A(t) - L_B(t),$$

or another signed differential that decides whether pressure should be routed to channel  $A$ , channel  $B$ , or both.

**Use case:** sociotechnical systems where agent stress and structural fragility must not be conflated.

## 10.3 Pattern 3: Layered admissibility gating

**Intent:** prevent surplus in one layer from masking depletion in another decisive layer.

**General form:**

$$\Delta_{\text{eff}}(t) = \Delta_{\text{primary}}(t) \cdot \prod_{i \in \mathcal{C}} \sigma(k_i \Delta_i(t)), \quad (14)$$

where  $\mathcal{C}$  indexes critical secondary constraints and  $\sigma$  is a smooth gate. These layered operators are framework-level governance constructions rather than additional core-law axioms. They specify how multiple declared margins may be combined for conservative admissibility in multi-constraint systems.

**Interpretation:** a system does not get to spend one layer's surplus to offset another layer's hidden deficit when that hidden layer is boundary-decisive.

**Use cases:** grid adequacy gated by reserve headroom, throughput gated by thermal headroom, model quality gated by error budget, output gated by containment capacity.

## 10.4 Pattern 4: Horizon-triggered mode switching

**Intent:** use time-to-boundary as a policy trigger rather than waiting for boundary violation.

**General structure:**

- enter warning mode when  $T_M < T_{\text{warn}}$ ;
- enter protective mode when  $T_M < T_{\text{protect}}$ ;
- escalate when  $T_M$  approaches zero or Grace is exhausted.

**Use cases:** return-to-home triggers in power-limited missions, reserve protection, emergency throttling, clinical escalation logic, or autonomous safety-mode switching.

## 10.5 Pattern 5: Dose-aware recovery governance

**Intent:** distinguish states with identical instantaneous margin but different stress histories.

Two systems can have the same  $\Delta_X(t)$  now while differing sharply in accumulated deficit exposure. A controller that sees only current margin treats them as equivalent. A controller that also sees  $D_X(t)$  can distinguish recently stressed states from persistently stressed states and adjust recovery, rest, or protection accordingly.

**Use cases:** recovery scheduling, maintenance priority, cooldown enforcement, staff rotation, cache/error-budget governance, and post-event restoration planning.

## 11 Standard Reporting Metrics

These metrics are intended for comparisons under fixed boundary rules and declared actuator constraints.

Policy comparisons under fixed boundary rules should report the following.

1. **Usable work to boundary:** output delivered before violation.
2. **Time in deficit:** total time with  $\Delta_X(t) < 0$ , or below a declared warning threshold if that is the operationally relevant exposure.
3. **Deficit area or dose:**  $\int \max(0, -\Delta_X(t))dt$ , or another declared dose operator.
4. **Minimum margin:**  $\min_t \Delta_X(t)$ , with uncertainty context.
5. **Recoverability signature:** a standardized post-boundary or post-stress recovery metric.
6. **Control cost:** foregone throughput, curtailment, energy use, labor burden, latency cost, or another explicit intervention cost term.

Table 2: Standard reporting bundle for fixed-boundary comparisons.

Metric	Purpose
Usable work to boundary	Measures what the system actually delivered before violation
Time in deficit	Measures how long the system spent in a harmful region
Deficit area / dose	Measures persistence and severity of deficit exposure
Minimum margin	Shows worst-case closeness to the boundary
Recoverability signature	Captures standardized post-stress or post-boundary response
Control cost	Prevents one-sided reporting of gains without intervention cost

**Core comparison rule.** A policy-advantage claim is strongest when, under unchanged boundary rules, it improves usable work and/or recoverability while reducing deficit dose, with control cost reported explicitly.

**Why control cost matters.** Dominance-aware control often trades peak output for survivability or recoverability. Reporting only the gains while hiding the costs turns a control result into advocacy.

**Minimum reporting discipline.** A claimed improvement in a single metric is not enough. MDL-grade reporting requires a bundle view in which gains, harms, burden shifts, and intervention costs can all be seen at once.

## 12 Misuse Modes and Guardrails

### 12.1 Proxy arbitrariness

A proxy can be made to look scientific while remaining empty. The guardrail is simple: no tether, no higher-tier claim.

### 12.2 Goodhart and gaming

Once  $\Delta_X(t)$  becomes a target, users may manipulate inputs to move the score without improving actual viability. Guardrails include:

- publishing the mapping pipeline;
- retaining auditability of observables;
- reporting multiple margins when layered constraints exist;
- treating persistent near-boundary operation as a harm signal even when immediate output appears favorable.

### 12.3 Punitive ranking misuse

In institutional settings, dominance outputs should not be used as punitive rankers absent causal validation and outcome tethering. Their proper use is routing, early warning, prioritization, and support allocation under limited observability.

### 12.4 Causal overclaim

A dominance trace can support policy-conditional comparisons and boundary statements. It does not, by itself, prove universal mechanism. Strong causal language must be tied to an embodiment, a measurement pipeline, and a falsification hook.

### 12.5 Hidden-layer depletion

A positive primary margin does not guarantee real viability if another ungated layer is being silently exhausted. The guardrail is layered admissibility where appropriate.

### 12.6 Boundary drift

A favorable result can be manufactured by quietly changing the violation rule, threshold placement, sampling cadence, or endpoint definition across policies. The guardrail is boundary-faithful comparison and predeclared protocol.

## 13 Falsifiable Claim Classes

This paper supports several classes of claims, each falsifiable within properly instrumented embodiments.

### 13.1 Claim A: Policy-conditional improvement

Under fixed boundary rules and with at least one non-trivial actuator,

Policies that reduce sustained deficit exposure will, on average, improve usable work and/or recoverability relative to policies that ignore dominance, after accounting for control cost.

This claim is falsified by well-instrumented counterexamples in which dominance-aware control fails to improve the declared outcomes despite comparable conditions and non-trivial actuation.

### 13.2 Claim B: Dose and restoration burden

In embodiments with measurable restoration cost, restoration time, or capacity-loss signatures,

Greater cumulative deficit dose is associated with greater restoration burden, after controlling for confounds.

This claim is falsified when the declared burden does not track cumulative deficit exposure once measurement quality and confounds are handled properly.

### 13.3 Claim C: Layered-gating advantage

In layered systems,

Policies that respect both primary and secondary margins via admissibility gating will exhibit fewer or less severe boundary violations than policies that optimize the primary margin alone.

This claim is falsified by controlled counterexamples showing that secondary gating adds no meaningful protective value under the declared embodiment.

### 13.4 Claim D: Grace exhaustion as a measurable transition

In embodiments with response-based intervention testing,

There exists a finite region in which admissible intervention materially restores margin, and beyond which comparable intervention fails to do so.

This claim is falsified if no measurable transition between recoverable and non-recoverable regimes can be identified despite adequate instrumentation.

### 13.5 Claim E: Embodiment failure condition

For any embodiment claiming MDL-grade standing,

A declared operational boundary must not be violated while the measured dominance trace remains reliably positive beyond uncertainty, unless the failure is explicitly shown to arise from an omitted decisive layer or a broken mapping.

This claim is falsified when the system crosses its declared viability boundary while  $\Delta_X(t)$  remains stably positive outside the indeterminate band, with no admissible repair to the mapping.

## 13.6 Failure modes of the framework

The framework is most exposed when its own admissibility conditions fail. The principal failure modes are:

1. boundary violation while the measured dominance trace remains reliably positive beyond uncertainty;
2. no measurable recoverability transition where Grace exhaustion is claimed;
3. no observable dose-burden association where deficit-sensitive deterioration or restoration burden is claimed;
4. no protective value from layered gating where layered protection is claimed;
5. disappearance of claimed policy advantage under fixed-boundary replication.

These failure modes do not show that all uses of the Mena Dominance Law are invalid. They show that a specific embodiment, claim class, or framework construction has failed under its declared conditions.

## 14 Research Agenda

### 14.1 Physical embodiments

The most important next step is not broader analogy; it is more Tier A work. Candidate directions include battery life-cycle degradation under dose-aware control, thermal-limited computing, robotics cells, reservoir operations, manufacturing drift, bridge and structural reserve tracking, and other systems with measurable actuator-mediated boundary behavior.

### 14.2 Meta-operational embodiments

Operational embodiments remain valuable when outcome tethering is real. Relevant domains include incident-response workload governance, cyber containment, cloud service degradation control, network error budgets, emergency services dispatch, and institutional routing systems.

### 14.3 Methodological agenda

Several methodological standards should become normal.

- Standard Embodiment Cards for every reported study.
- Uncertainty-aware dominance reporting, including indeterminate bands.
- Fixed-boundary policy comparisons with explicit control-cost reporting.
- Longitudinal dose-versus-capacity-loss protocols where Reciprocal Decay is claimed.
- Explicit declaration of actuator limits and decision cadence.
- Public replication artifacts whenever feasible.

## 14.4 The real job of future work

Future work should prioritize stronger embodiments, better instrumentation, sharper proposition-level formalization, and cleaner falsification protocols rather than broad analogical extension.

## 15 Conclusion

Paper 1 established the Mena Dominance Law. Paper 2 supplied executed embodiment evidence under fixed boundary rules. This paper states the discipline required to use the law without reducing it to slogan, score, or analogy.

The central object remains

$$\Delta_X(t) = P_X(t) - L_X(t),$$

but the law is not the algebraic identity alone. Scientific use requires a declared boundary, commensurate mapping, a reproducible pipeline, uncertainty declaration, dose definition where relevant, admissible intervention where control is claimed, and fixed-boundary comparison.

A law stated in domain-invariant form is easy to overextend unless embodiment discipline is enforced. This paper supplies that discipline.

Under that discipline, margin becomes headroom, deficit becomes dose, Grace becomes measurable recoverability, and time-to-boundary becomes a control-ready signal. Those objects are the operational content of the Mena Dominance Law once an embodiment is properly built.

This paper closes the trilogy by stating the conditions under which future embodiments may make strong, bounded, and falsifiable claims under the framework.

<https://menadominancelaw.com/>

## A Notation and Glossary

Table 3: Key symbols and terms used in the Mena Dominance Law framework.

Symbol / Term	Meaning
$P_X(t)$	Recoverable potential in system $X$
$L_X(t)$	Imposed load in system $X$
$\Delta_X(t)$	Dominance margin $P_X(t) - L_X(t)$
$\Delta_R(t)$	Rate-form dominance margin
$\Delta_E(t)$	Budget-dose dominance
Boundary	Declared viability constraint and detection rule
$\epsilon_X$	Uncertainty or resolution band for $\Delta_X(t)$
Dose $D_X(t)$	Time-integrated nonnegative deficit exposure
Grace	Finite recoverability region under admissible control
$G_X(t; R, \Delta t)$	Response-based Grace gain
$T_M$	Mena Horizon: conditional time-to-boundary estimate
Actuator $u(t)$	Control input that modifies load, potential, or architecture
Layered gating	Effective-margin rule gated by secondary margins
Boundary-faithful comparison	Policy comparison performed under unchanged boundary rules

## B Embodiment Card (Minimum Reporting Template)

Table 4: Embodiment Card for replication and falsification.

Item	Required content
Boundary	Violation definition; detection rule; thresholds if used
Potential $P_X(t)$	Definition; units; estimator/pipeline; cadence
Load $L_X(t)$	Definition; units; estimator/pipeline; cadence
Margin $\Delta_X(t)$	Construction; logging; uncertainty band $\epsilon_X$
Observables	Sensors, signals, or data sources
Dose	Dose operator and rationale if used
Actuator(s) / intervention frame	Control variables, intervention frame, and admissibility constraints if control or recoverability is claimed
Outcomes	Predeclared endpoints under fixed boundary rules
Control cost	Resource cost of intervention or protection
Falsifiability	Explicit failure mode for the embodiment mapping
Tier claim sought	Tier A, Tier B, or Tier C claim class

## C Minimum Implementation Checklist

1. Declare the boundary and the rule for detecting violation.

2. Choose the governing channel or channels.
3. Define  $P_X(t)$  and  $L_X(t)$  in commensurate units.
4. Specify the pipeline from observables to  $P_X(t)$ ,  $L_X(t)$ , and  $\Delta_X(t)$ .
5. Declare uncertainty and treat  $|\Delta_X(t)| \leq \epsilon_X$  as indeterminate.
6. Define dose if persistence of deficit matters.
7. Identify admissible actuators or an intervention-relevant evaluation frame if control or recoverability is claimed.
8. Hold boundary rules fixed during policy comparison.
9. Report usable work, deficit time/area/dose, recoverability, and control cost.
10. State what observation would falsify the embodiment mapping.

## References

- [1] A. O. M. Mena. *The Mena Dominance Law of Operational Decay: A Fundamental Law of Bounded Viability for Constraint-Bearing Systems*. Preprint, 2026.
- [2] A. O. M. Mena. *The Mena Dominance Law: Viability-Governed Control Under Experimental Battery Runs, Risk Tool Evidence, and Simulation Validation*. Preprint, 2026.
- [3] J.-P. Aubin. *Viability Theory*. Birkhäuser, Boston, 1991.
- [4] C. S. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973.
- [5] H. K. Khalil. *Nonlinear Systems*. 3rd ed., Prentice Hall, Upper Saddle River, NJ, 2002.
- [6] N. Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, Cambridge, MA, 2011.
- [7] A. M. Lyapunov. *The General Problem of the Stability of Motion*. Translated by A. T. Fuller. CRC Press, London, 1992. Original memoir published in 1892.
- [8] R. E. Melchers. *Structural Reliability Analysis and Prediction*. 2nd ed., John Wiley & Sons, Chichester, 1999.
- [9] M. A. Miner. Cumulative damage in fatigue. *Journal of Applied Mechanics*, 12(3):A159–A164, 1945.
- [10] E. Hollnagel, D. D. Woods, and N. Leveson, eds. *Resilience Engineering: Concepts and Precepts*. Ashgate, Aldershot, 2006.
- [11] International Organization for Standardization. *ISO 2394:2015: General Principles on Reliability for Structures*. ISO, Geneva, 2015.