The Mena Dominance Law: Viability-Governed Control Under

Experimental Battery Runs, Risk Tool Evidence, and Simulation Validation

Arlex Orlando Murcia Mena*

January 2026

Abstract

This paper evaluates the Mena Dominance Law kernel in applied settings, using the time-tracked margin

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

where $P_X(t)$ is recoverable potential and $L_X(t)$ is imposed load. We test whether *viability-governed control*—policies that actively avoid sustained dominance deficit under fixed boundary rules—changes usable work, deficit exposure, and boundary-adjacent recovery outcomes.

First, we report paired battery-discharge experiments on commercial Li-ion packs (four physical packs; one model has two units; $N_M = N_R = 57$), comparing constant-current baselines against a viability-governed four-step current schedule near a fixed cutoff. Under identical cutoff rules, the governed policy increases delivered energy and charge to cutoff (median +5.82% Wh; +4.80% Ah), extends runtime (median +63.38%), and reduces initial voltage sag (median −4.13%). Standardized post-cutoff recovery voltages shift lower under the governed policy; we interpret this as a boundary-adjacent recovery signature under fixed termination logic, not as a degradation claim.

Second, we examine StARS as a sociotechnical risk instrument for routing corrective action between agent-channel stress and structure-channel fragility. Third, using CAISO OASIS time series, we simulate single-layer and dual-layer grid governors, which in this simulation reduce deficit exposure versus baseline and show a clear tradeoff between curtailment and residual deficit time/area under reserve-headroom admissibility gating.

Taken together, the results show that under fixed boundary definitions, dominance-aware control yields measurable policy-conditional differences. The battery experiments provide the primary executed physical evidence in this paper; the StARS and CAISO sections provide secondary operational and simulation validation.

Keywords: dominance margin, viability-governed control, load shaping, battery discharge, recoverability, operational decay, standardized recovery voltage, risk-log evidence, simulation validation

*Patent Pending: Portions of the risk-governance methodology (agent–architecture viability assessment) and systems/methods for viability-based control of complex adaptive environments.

Contents

Data and artifact repository

All public-source datasets used in this paper (as-downloaded archives), derived simulation tables, tool screenshots, and reproducibility artifacts are hosted at:

`https://menadominancelaw.com/`

The repository includes original-source citations and download dates, immutable copies of source files under the exact filenames referenced in this paper, and all generated simulation outputs (CSV), including sensitivity sweeps and Monte Carlo ensembles.

1 Scope and Relationship to Paper 1

Paper 1 defines the dominance kernel, including the time-tracked margin $\Delta_X(t)$, regime interpretation, and the embodiment contract (commensurate Potential/Load definitions with measurable or estimable time traces). This paper is deliberately narrower: it evaluates whether dominance-governed control produces measurable benefits under fixed constraints.

Accordingly, Paper 2 does not claim new electrochemical physics, new institutional dynamics, or a new degradation law. It treats dominance as an instrument signal and tests the practical proposition that a controller which respects margin trajectory (and avoids prolonged deficit exposure) yields reproducible improvements in operational outcomes. The experimental and simulation designs isolate policy differences while holding boundary rules constant (e.g., cutoff thresholds, termination logic, measurement cadence, and recovery protocol). Where later sections introduce layered controllers, these are presented as embodiment-specific control constructions rather than additional core-law axioms.

The evidence in this paper is not flat in kind. The battery section provides the primary executed physical evidence under fixed boundary rules. The StARS section provides instrument-level operational evidence in a meta-operational embodiment. The CAISO section provides real-data simulation validation in an infrastructure embodiment. These sections therefore support different grades of claim and should not be read as carrying identical evidentiary weight.

2 Kernel Summary

The diagnostic used throughout is the dominance margin:

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

Each embodiment supplies commensurate definitions of $P_X(t)$ (recoverable operational capacity) and $L_X(t)$ (imposed operational demand) in matched units and over consistent estimation windows. Paper 1 treats the general trajectory logic (dominance regimes, Grace structure, and deficit-sensitive deterioration claims where applicable). This paper does not re-derive those constructs. Instead, it evaluates embodiment-specific operational proxies tied to recoverability and usable work, including standardized recovery voltage, runtime, and energy delivered to a fixed cutoff under controlled conditions.

3 Battery Experiments: Dominance-Governed Load Control Under Fixed Cutoff Rules

This section provides the primary executed physical evidence in Paper 2, using an electrochemical embodiment. We evaluate whether a dominance-governed controller changes *usable work delivered*

and *boundary-adjacent recovery outcomes* when boundary rules are held fixed. The plant is a set of commercial cordless-tool lithium-ion battery packs. The boundary is a pack-specific cutoff voltage V_{cut} held constant across policies within each physical pack. The actuator is the discharge current commanded by an electronic load.

The central design choice is to separate *boundary conditions* from *load policy*. For every run we enforce the same termination rule—stop at first detection of $V_{\text{pack}}(t) \leq V_{\text{cut}}$—and we hold the measurement cadence and recovery observation window constant when available. The only difference between conditions is how current is scheduled as the pack approaches the boundary.

Policies compared. Two discharge policies are compared:

- **R (regular baseline):** constant-current discharge without dominance feedback, terminated at the fixed cutoff V_{cut} for that pack.
- **M (Mena; viability governor):** a **four-step** quantized reduction in discharge current ("drop-down amps") as the pack approaches the fixed cutoff boundary, implemented as a **monotone step-down schedule** in ET5406A+ Battery-Test (BATT) mode.

In this embodiment, "drop-down amps" is the actuator because current is the controllable input that can be reduced near the fixed boundary without altering the battery pack or cutoff rule.

Dominance proxy in the battery embodiment. For the experiments, the real-time viability observable is the *terminal voltage margin to the fixed cutoff*:

$$\widehat{\Delta}_{\text{batt}}(t) := V(t) - V_{\text{cut}}.$$

Small positive $\widehat{\Delta}_{\text{batt}}$ indicates proximity to the boundary under load; the M policy reduces current to avoid prolonged near-boundary (and post-boundary) exposure.

Important clarification (derived from the Paper 1 battery embodiment). Although the ET5406A+ executes M using voltage thresholds, the *four-step schedule parameters* were not chosen arbitrarily. For each pack we derived the step currents and thresholds from the Paper 1 battery embodiment by (i) fixing V_{cut} per pack, (ii) estimating an effective use-resistance R_{use} from early baseline step behavior, and (iii) computing an admissible-current scale

$$I_{\text{max},0} \approx \frac{V_{\text{rest},0} - V_{\text{cut}}}{R_{\text{use}}},$$

where $V_{\text{rest},0}$ is the pre-load rest/low-current proxy at the start of discharge. The M currents are then selected as fixed fractions of this admissible-current scale for that pack (high / mid / near-boundary glide), and the ET5406A+ voltage thresholds are chosen to place those current reductions in the observed knee region of that pack's voltage curve while preserving the same fixed cutoff rule. This makes the M schedule *pack-specific* because $V_{\text{rest},0}$, V_{cut}, and R_{use} differ by battery type.

What the tests measure. Each run records $V(t)$ and $I(t)$ from load enable until cutoff, yielding: (i) delivered energy to cutoff $Wh_{\text{to cut}}$ and delivered charge $Ah_{\text{to cut}}$ as primary measures of usable work under fixed boundary rules; (ii) runtime to cutoff as an operational metric; and (iii) when available, a standardized post-discharge recovery voltage V_{rec} measured with the load removed for a nominal 60 s window.

Interpretation constraints. Because the M policy intentionally reduces current near cutoff, discharge current and power are not matched between M and R. Runtime is therefore a *policy-mediated* outcome. Accordingly, the primary cross-policy comparators are delivered energy (Wh) and delivered charge (Ah) to the same cutoff rules, and recovery behavior is treated as a boundary-adjacent recovery signature within this protocol rather than degradation evidence.

Organization. The remainder of this section specifies the test hardware, packs, environmental/rest controls, signals and endpoint definitions, and then reports pooled and per-pack comparisons between M and R.

3.1 Test hardware, charging, and instruments

Electronic load. All discharge tests were executed using a **Yertai ET5406A+** DC electronic load operated in current-controlled modes (constant-current and Battery-Test CC mode with multi-step cutoffs). An inline **15 A fuse** was placed **in series on the positive lead** between the pack and the load for protection. Additional wiring photos and device screenshots are archived in the artifact repository `https://menadominancelaw.com/`.

Charging. All packs were charged using their **original manufacturer chargers**. No third-party chargers or modified charging protocols were used.

Temperature instrumentation (auxiliary). Temperature was recorded using a **TEMPer1F** logger and a **TXFM75** temperature sensor. The sensor was placed in the *same physical location* on the pack for all runs and secured with electrical tape to reduce placement variability. Ambient temperature was not regulated; a subset of runs were re-executed on different days to replace invalid trials, which further increased day-to-day variation. Temperature is therefore used only as a secondary descriptive signal and is reported via paired complete-case summaries.

3.2 Packs and dataset size

Four packs were tested (one model has two physical units):

- **DCB205**: DeWalt 20V MAX 5Ah pack.
- **R840040_A**: Ridgid 18V 4Ah pack (Battery A).
- **R840040_B**: Ridgid 18V 4Ah pack (Battery B).
- **KB224-03**: Kobalt 24V MAX 2Ah pack.

Only the Ridgid R840040 model is represented by two physical packs (A and B); the DeWalt DCB205 and Kobalt KB224-03 entries are single-pack embodiments in this dataset.

A total of 114 usable runs were collected, balanced between conditions:

$$N_{\mathrm{M}} = 57, \qquad N_{\mathrm{R}} = 57.$$

Pack counts are balanced by design: DCB205 (15/15), KB224-03 (14/14), R840040_A (14/14), R840040_B (14/14). A small number of initially recorded runs were flagged invalid (e.g., missing segments or instrumentation artifacts) and were replaced by reruns; all counts and outcomes reported below refer to the final usable set.

3.3 Environmental conditions and rest protocol

Experiments were executed across real ambient conditions spanning approximately

$$68^\circ\text{F to } 100^\circ\text{F}.$$

Ambient temperature was *intentionally not regulated*: tests were executed across multiple ambient-temperature conditions to evaluate whether the observed policy effects persist under different temperatures. Both policies (M and R) were run across the same range of operating conditions encountered during data collection (including reruns), and temperature is treated as a descriptive/context variable rather than a controlled factor.

To reduce confounding from recent use and transient thermal/polarization effects, a strict rest protocol was applied:

- **Pre-run rest:** each pack rested ≥ 1 hour after use before a test.
- **Post-run rest:** each pack rested ≥ 1 hour after a test before charging or re-running.

3.4 Signals, cadence, termination, and recovery window

Each run logs pack terminal voltage $V(t)$ and current $I(t)$ from the electronic load. When available, temperature $T(t)$ is logged from the TEMPer1F/TXFM75 channel. Sampling cadence is determined by the logger (typical cadence $\sim$2 s).

All runs terminate at first detection of the cutoff condition:

$$V_{\text{pack}}(t) \leq V_{\text{cut}},$$

with V_{cut} declared per pack and held constant across policies within that pack.

When available, post-discharge recovery is recorded with the load removed for a standardized window (nominally 60 s). Define the standardized recovery voltage:

$$V_{\text{rec}} \equiv V(t_{\text{cut}} + 60\text{ s}). \tag{2}$$

For recovery-rise reporting we also define $V_{\text{rec},0}$ as the first recorded voltage sample after load removal, and the 60 s recovery rise as $\Delta V_{\text{rec}} := V_{\text{rec}} - V_{\text{rec},0}$. When the logger cadence does not land exactly at 60 s, V_{rec} is taken as the nearest recorded sample to $t_{\text{cut}} + 60$ s.

3.5 Policies compared (R vs. M) and the "drop-down amps" implementation

R policy (Regular baseline). Regular runs use constant-current discharge without dominance feedback, terminated at the same fixed cutoff V_{cut} used in the corresponding M runs for that pack.

M policy (Mena; four-step viability governor in ET5406A+ BATT mode). The M policy is implemented on the ET5406A+ using **Battery Test (BATT)** $\rightarrow$ **CC** with multiple cutoffs enabled. The controller is a **monotone step-down schedule** driven by terminal-voltage thresholds (no step-up/hysteresis):

$$I(t) = \begin{cases} I_1, & V(t) > V_1, \\ I_2, & V_1 \geq V(t) > V_2, \\ I_3, & V_2 \geq V(t) > V_3, \\ I_4, & V_3 \geq V(t) > V_{\text{cut}}, \\ 0\ \text{(stop)}, & V(t) \leq V_{\text{cut}}, \end{cases} \qquad I_1 > I_2 > I_3 > I_4.$$

In this study the fourth step is a near-boundary *glide-slope* floor and is held constant across M runs for a given pack (implementation floor: $I_4 = 2.0$ A in our ET5406A+ profiles). The voltage thresholds (V_1, V_2, V_3) and currents (I_1, I_2, I_3) are pack-specific and were derived from the Paper 1 battery embodiment (admissible-current scaling from $V_{\mathrm{rest},0}$, V_{cut}, and R_{use}, then snapped to the ET5406A+ step-threshold interface). The resulting fixed schedules used in this study are listed in Table 1.

Continuous viability governor (software implementation). The four-step ET5406A+ profiles in Table 1 are a *quantized hardware approximation* of a continuous viability-governed controller. In the underlying battery embodiment, the admissible discharge current at time t is scaled from the estimated near-rest voltage and use-resistance via

$$I_{\max}(t) \approx \frac{V_{\mathrm{rest}}(t) - V_{\mathrm{cut}}}{\hat{R}_{\mathrm{use}}(t)}.$$

A software implementation can apply this ceiling at each sample (e.g., once per second), clipping the externally requested current $I_{\mathrm{demand}}(t)$ to $I_{\max}(t)$ whenever the demand would violate the dominance margin. Appendix B.1, Algorithm 1, lists pseudocode for this continuous battery viability governor. The ET5406A+ step-down schedules used in this study were chosen to approximate the same admissible-current rule under the device's multi-step Battery-Test interface.

Interpretation of runtime under M. Because M intentionally steps current down near cutoff, discharge current and power are not matched between M and R. Consequently, *runtime is policy-mediated*, while delivered energy (Wh) and delivered charge (Ah) are treated as the primary cross-policy comparators under fixed cutoff rules.

Table 1: Pack-specific M schedules used in this study (ET5406A+ BATT-CC; four-step drop-down with $I_4 = 2.0$ A floor).

Pack	**M schedule (voltage window → current)**
DCB205 (5S)	$V > 18.00$: 7.50 A; $18.00 \geq V > 17.00$: 5.00 A; $17.00 \geq V > 16.50$: 3.00 A; $16.50 \geq V > 16.00$: 2.00 A; Stop at $V \leq 16.00$ V.
KB224-03 (6S)	$V > 21.60$: 7.50 A; $21.60 \geq V > 20.40$: 5.00 A; $20.40 \geq V > 19.80$: 3.00 A; $19.80 \geq V > 19.20$: 2.00 A; Stop at $V \leq 19.20$ V.
R840040_A (5S)	$V > 18.00$: 7.50 A; $18.00 \geq V > 17.00$: 5.00 A; $17.00 \geq V > 16.50$: 3.00 A; $16.50 \geq V > 16.00$: 2.00 A; Stop at $V \leq 16.00$ V.
R840040_B (5S)	$V > 18.00$: 7.50 A; $18.00 \geq V > 17.00$: 5.00 A; $17.00 \geq V > 16.50$: 3.00 A; $16.50 \geq V > 16.00$: 2.00 A; Stop at $V \leq 16.00$ V.

3.6 Endpoints and inclusion rules

Primary endpoints (all runs). Pairs are defined within each physical pack by matched run index (e.g., the kth M run vs. the kth R run for that same pack); complete-case means both runs in the pair contain the required segment.

Primary endpoints depend only on the load-active segment and therefore use all runs ($N_M = N_R = 57$):

- runtime to cutoff t_{cut} (s);
- delivered energy to cutoff:

$$Wh_{\text{to cut}} = \int_0^{t_{\text{cut}}} V(t)\,|I(t)|\,dt; \tag{3}$$

- delivered charge to cutoff:

$$Ah_{\text{to cut}} = \int_0^{t_{\text{cut}}} |I(t)|\,dt;$$

- initial sag η_0 (rest proxy minus first under-load voltage sample), where the rest proxy is the final recorded voltage sample immediately prior to load enable and the first under-load sample is the first recorded voltage after load enable.

Secondary endpoints (paired complete-case). Secondary endpoints are computed using paired complete-case inclusion to preserve fairness: a run index contributes to a metric only if both M and R provide the required data.

- Temperature endpoints (e.g., $T_{\max}$): use only pairs with valid temperature traces in both M and R.
- Recovery endpoints (e.g., V_{rec}, ΔV_{rec}): use only pairs with valid 60 s recovery windows in both M and R.

Missing temperature or recovery segments are treated as missing measurements rather than battery anomalies and do not affect the primary endpoints.

3.7 Outcomes (pooled across all packs)

This subsection reports pooled outcomes across all four packs (balanced N across policies).

Primary pooled results (all runs; medians; $N_M = N_R = 57$).

- **Runtime (policy-mediated):** median increases from 1756 s (R) to 2869 s (M), a gain of +1113 s (+63.38%).
- **Delivered energy:** median increases from 64.39 Wh (R) to 68.14 Wh (M), a gain of +3.75 Wh (+5.82%).
- **Delivered charge:** median increases from 3.650 Ah (R) to 3.825 Ah (M), a gain of +0.175 Ah (+4.80%).
- **Initial sag:** median decreases from 0.968 V (R) to 0.928 V (M), a reduction of 0.040 V (−4.13%).

Secondary outcomes (paired complete-case; medians).

- **Temperature (n=57 pairs):** median peak temperature difference $T_{\max,M} - T_{\max,R} = -0.75°\text{C}$; approximately 62% of pairs are cooler under M. Because tests were conducted across uncontrolled ambient conditions (including rerun days), this is treated as descriptive rather than causal evidence.
- **Recovery (n=54 pairs):** median recovery-rise difference $(\Delta V_{\text{rec},M} - \Delta V_{\text{rec},R}) = -0.137$ V (M lower), and median end-of-recovery difference $(V_{\text{rec},M} - V_{\text{rec},R}) = -0.568$ V (M lower). These shifts are consistent with longer/deeper operation before termination under M and are not interpreted as degradation evidence.

3.8 MDL Yield Parameter (battery)

To make the embodiment-level tradeoff explicit under this battery protocol, we define the MDL yield parameter as the ratio of percent gain in delivered energy to percent gain in runtime:

$$Y_{\text{batt}} := \frac{\Delta Wh\%}{\Delta t\%}. \tag{4}$$

Using the pooled medians reported above,

$$\Delta Wh\% = 5.82, \qquad \Delta t\% = 63.38,$$

so

$$Y_{\text{batt}} = \frac{5.82}{63.38} \approx 0.0918. \tag{5}$$

Operationally, this means the pooled-median delivered-energy gain is about 0.092 percentage points for each 1 percentage point of runtime extension under this fixed-cutoff protocol.

This yield parameter is a protocol-specific summary of the observed tradeoff in the present battery experiment; it is not presented as a universal constant, an electrochemical efficiency, or a fundamental physical parameter. Because runtime is policy-mediated under M through deliberate current reduction near cutoff, runtime should not be interpreted by itself as an efficiency gain. The primary cross-policy comparators remain delivered energy $Wh_{\text{to cut}}$ and delivered charge $Ah_{\text{to cut}}$ under identical cutoff rules.

3.9 Positioning relative to Model Predictive Control (MPC)

A likely comparison is Model Predictive Control (MPC), since both approaches can shape current as a function of proximity to a boundary (e.g., cutoff). The distinction is architectural: in this work the policy is an explicit, deterministic governor computed from measured/estimated dominance margin (voltage headroom to cutoff) and any accumulated deficit exposure state used by the embodiment. The resulting action is produced in closed form (or via a small rule map) and is therefore constant-time and microcontroller-deployable without an online optimizer.

By contrast, MPC is defined by receding-horizon optimization: at each update it solves an optimization problem over a future horizon using a predictive model of system dynamics and constraints, applies the first control move, then re-solves at the next step. MPC can enforce hard constraints and can also represent soft constraints via slack variables and penalties; however, doing so requires an explicit predictive model and a chosen objective/penalty structure. Importantly, these approaches are not mutually exclusive: the dominance governor can wrap any baseline controller (including MPC) by vetoing or throttling actions that would drive sustained boundary proximity or deficit exposure under the embodiment's measurement model.

3.10 Outcome tables (per-pack and pooled)

Table 2: Per-pack primary endpoint medians (M vs. R) with median difference and percent change relative to R.

Pack	Endpoint	M median	R median	Δ (M–R)	% vs R
DCB205	Delivered charge $Ah_{\text{to cut}}$ (Ah)	4.299	4.139	+0.159	+3.85%
DCB205	Delivered energy $Wh_{\text{to cut}}$ (Wh)	77.65	74.49	+3.17	+4.25%
DCB205	Initial sag η_0 (V)	0.599	0.604	-0.005	-0.83%
DCB205	Runtime t_{cut} (s)	3134	2050	+1084	+52.88%
KB224-03	Delivered charge $Ah_{\text{to cut}}$ (Ah)	2.195	2.061	+0.134	+6.52%
KB224-03	Delivered energy $Wh_{\text{to cut}}$ (Wh)	50.03	46.91	+3.12	+6.65%
KB224-03	Initial sag η_0 (V)	0.979	0.971	+0.008	+0.85%
KB224-03	Runtime t_{cut} (s)	2794	1566	+1228	+78.41%
R840040_A	Delivered charge $Ah_{\text{to cut}}$ (Ah)	3.447	3.268	+0.179	+5.48%
R840040_A	Delivered energy $Wh_{\text{to cut}}$ (Wh)	60.74	57.77	+2.96	+5.13%
R840040_A	Initial sag η_0 (V)	1.312	1.349	-0.037	-2.76%
R840040_A	Runtime t_{cut} (s)	2918	1755	+1163	+66.26%
R840040_B	Delivered charge $Ah_{\text{to cut}}$ (Ah)	3.827	3.651	+0.175	+4.80%
R840040_B	Delivered energy $Wh_{\text{to cut}}$ (Wh)	68.44	64.48	+3.96	+6.14%
R840040_B	Initial sag η_0 (V)	0.868	0.938	-0.070	-7.43%
R840040_B	Runtime t_{cut} (s)	2918	1713	+1205	+70.35%

Table 3: Recovery outcomes by pack (paired complete-case). $V_{\text{rec},0}$ is the first recorded sample after load removal; V_{rec} is the nearest sample to $t_{\text{cut}} + 60$ s; $\Delta V_{\text{rec}} = V_{\text{rec}} - V_{\text{rec},0}$. Medians are reported; Δ columns are (M–R).

Pack	n	$V_{\text{rec},0}$ **M**	$V_{\text{rec},0}$ **R**	Δ	V_{rec} **M**	V_{rec} **R**	Δ	ΔV_{rec} **M**	ΔV_{rec} **R**	Δ
DCB205	13	16.170	16.580	-0.410	16.278	16.849	-0.571	0.115	0.279	-0.164
KB224-03	14	19.870	20.390	-0.520	20.343	20.794	-0.451	0.463	0.392	+0.072
R840040_A	13	16.240	16.710	-0.470	16.472	17.096	-0.624	0.220	0.359	-0.139
R840040_B	14	16.130	16.620	-0.490	16.392	16.977	-0.585	0.258	0.362	-0.104
Pooled	54	16.190	16.640	-0.450	16.409	16.977	-0.568	0.225	0.362	-0.137

Table 4: Primary and secondary outcomes (pooled; medians). Primary endpoints use all runs. Secondary endpoints use paired complete-case inclusion.

Endpoint	**M median**	**R median**	**Δ (M–R)**	**% vs R**
Runtime t_{cut} (s)	2869	1756	+1113	+63.38%
Delivered energy $Wh_{\text{to cut}}$ (Wh)	68.14	64.39	+3.75	+5.82%
Delivered charge $Ah_{\text{to cut}}$ (Ah)	3.825	3.650	+0.175	+4.80%
Initial sag η_0 (V)	0.928	0.968	−0.040	−4.13%
Temp pairs: $n = 57$ paired complete-case. Median $(T_{\max,M} - T_{\max,R}) = -0.75$°C; ∼62% cooler under M (descriptive only).				
Recovery pairs: $n = 54$ paired complete-case. Median recovery-rise diff ≈ −0.137 V (M lower). Median end-recovery diff ≈ −0.568 V (M lower).				

3.11 Longitudinal degradation validation of Reciprocal Decay (proposed multi-cycle test)

The single-run results above establish that a dominance-governed discharge policy can increase usable work delivered under fixed cutoff rules. To directly test the *Reciprocal Decay* claim (deficit-adjacent exposure steepens loss of recoverable capacity over repeated operation), a longitudinal multi-cycle protocol is required.

Current status. Longitudinal battery-cycle testing relevant to the Reciprocal Decay claim is in progress. Because those results are not yet complete, they are not included as evidence in this paper.

Near-boundary dose metric (operational). Because the termination rule is fixed at $V(t) \le V_{\text{cut}}$, sustained operation occurs primarily *near* the boundary rather than deep below it. Define a warning margin $\tau_w > 0$ (in volts) and the near-boundary dose:

$$\mathcal{D}_w := \int_0^{t_{\text{cut}}} \max\bigl(0, \tau_w - \widehat{\Delta}_{\text{batt}}(t)\bigr)\, dt,$$
$$\widehat{\Delta}_{\text{batt}}(t) = V(t) - V_{\text{cut}}.$$

This produces a commensurate exposure measure (V·s) for how long and how tightly the run operated against the boundary.

Proposed protocol. This longitudinal test requires additional identical packs per model to support an R-vs-M split. For each model, acquire ≥ 2 matched physical packs per arm (R baseline vs. M governed), then run ≥ 100 charge–discharge cycles per arm under the same chargers, rest protocol, cutoff rules, and comparable ambient conditions. With the present dataset (one DeWalt pack, one Kobalt pack, and two Ridgid packs), only the Ridgid model can be split across arms (A vs. B), yielding $n = 1$ per arm and therefore serving as a pilot rather than a definitive degradation-rate comparison. At fixed cycle intervals (e.g., every 10 cycles), measure (i) delivered energy/charge to cutoff under a standardized reference discharge, (ii) an internal-resistance proxy (e.g., initial sag η_0 under a fixed current step or a step-estimated $\hat{R}_{\text{use}}$), and (iii) the per-cycle near-boundary dose $\mathcal{D}_w$.

Expected signature if Reciprocal Decay holds. If Reciprocal Decay holds in this embodiment, higher cumulative dose $\sum \mathcal{D}_w$ should correspond to faster growth of resistance proxy and faster loss of delivered capacity. Under this hypothesis, the R arm (higher dose) would be expected to degrade faster than the M arm (lower dose), converting the claim from "controller improves single-run usable work" to "dose-coupled decay accumulates over cycles under fixed boundary rules."

3.12 Replication checklist and dataset summary

Battery file-level reproducibility. Battery results were computed from the paired run logs shipped in `MDL`$_b$*attery.zip*, *withbaselinerunslabeled*`Test_R_*`*andgovernedrunslabeled*`Test_M_*`, *pairedwithineach*

Table 5: Replication checklist: equipment, protocol, and environment.

Item	Specification
Electronic load	Yertai ET5406A+ (CC and Battery-Test CC multi-step mode)
Protection	Inline 15 A fuse on positive lead (pack → load)
Charging	Original manufacturer chargers (no third-party charging)
Pre-run rest	≥ 1 hour after use
Post-run rest	≥ 1 hour after test before charge / next run
Ambient conditions	~ 68°F to 100°F across collection days (incl. reruns)
Temperature logging	TEMPer1F + TXFM75 (fixed placement; auxiliary)
Signals logged	$V(t)$, $I(t)$; optional $T(t)$
Termination rule	Stop at first $V_{\text{pack}}(t) \leq V_{\text{cut}}$ (per pack)
Recovery endpoint	$V_{\text{rec}} = V(t_{\text{cut}} + 60\text{ s})$ when available
Policies	R: baseline constant-current; M: four-step drop-down amps (BATT-CC)
Notes	Runtime is policy-mediated under M; Wh/Ah are primary cross-policy comparators

Table 6: Replication checklist: packs and run counts (balanced by design).

Pack	Pack identity	Runs (M/R)
DCB205	DeWalt 20V MAX 5Ah	15/15
KB224-03	Kobalt 24V MAX 2Ah	14/14
R840040_A	Ridgid 18V 4Ah (Battery A)	14/14
R840040_B	Ridgid 18V 4Ah (Battery B)	14/14
ALL	Pooled	57/57

4 StARS: Dominance, Grace, and Decay in a Sociotechnical Control Instrument

This section provides instrument-level operational evidence for dominance-guided routing in a meta-operational embodiment under a declared proxy framework. It is not presented as direct physical-system proof. Its purpose is to show that a declared dominance construction can stably route corrective action between agent-channel stress and structure-cha fragility, while preserving the Paper 1 logic of grace, deficit, and burden within a tethered operational frame.

4.1 Purpose and embodiment contract

StARS (*Structural Alignment Risk Scoring*) is an operational instrument that instantiates the Mena Dominance Law inside sociotechnical systems. Its purpose is to prevent the dominant institutional error: applying corrective pressure to the wrong channel. StARS decomposes misalignment into (i) agent-channel stress and (ii) structure-channel fragility, then routes correction according to dominance, and escalates when *grace closes* and *decay becomes irreversible*.

A StARS evaluation is defined over an entity and a measurement window:

$$(\texttt{Entity ID},\ \texttt{Window start},\ \texttt{Window end}).$$

In hospital use, the entity may be a hospital, region, or unit, and the window may be monthly/quarterly/annual.

4.2 Dominance margin and routing rule

StARS operationalizes the dominance kernel:

$$\Delta_{\text{org}}(t) = ASL(t) - CV(t), \tag{6}$$

where ASL is **Agent Stress Load** and CV is **Code Vulnerability**. In StARS this is displayed as *Classic Dominance* (or dominance margin). ASL and CV are computed from indicator vectors as defined in Section 4.4.

Grace (state-based). In this embodiment, **grace** is the available viability headroom carried by the dominance margin:

$$\texttt{Grace state} \equiv \Delta_{\text{org}}(t). \tag{7}$$

Thus $\Delta_{\text{org}}(t) > 0$ indicates grace-open (viable headroom), $\Delta_{\text{org}}(t) \approx 0$ indicates boundary, and $\Delta_{\text{org}}(t) < 0$ indicates dominance deficit (grace closed/closing without intervention).

Correction routing (mandate). Define a dominance threshold $t_{\text{dom}} > 0$. StARS routes correction by:

$$\texttt{Mandate}(t) = \begin{cases} \textbf{Agent Stabilization} & \Delta_{\text{org}}(t) > t_{\text{dom}}, \\ \textbf{Structural Correction} & \Delta_{\text{org}}(t) < -t_{\text{dom}}, \\ \textbf{Dual-Path Intervention} & |\Delta_{\text{org}}(t)| \le t_{\text{dom}}. \end{cases} \tag{8}$$

The mandate is a prohibition rule: if $\Delta_{\text{org}} < 0$ (structure-dominant), increasing pressure on agents is prohibited as the primary fix; if $\Delta_{\text{org}} > 0$ (agent-dominant), adding structural complexity is prohibited as the primary fix.

4.3 Calibration to commensurate operational rates

The 0-100 indicators in StARS are an interface scale. To satisfy the embodiment contract in strictly commensurate units, StARS admits a calibration layer that maps index values to operational rates expressed in the same unit.

Calibration functions. Let $\kappa_A(\cdot)$ and $\kappa_C(\cdot)$ map the index scores (ASL, CV) into a common operational unit, e.g. *expected nursing hours lost per month*:

$$P_{\texttt{org}} := \kappa_A(ASL), \qquad L_{\texttt{org}} := \kappa_C(CV),$$

with $P_{\texttt{org}}$ and $L_{\texttt{org}}$ in identical units (hours/month). The commensurate dominance margin is then:

$$\Delta_{\texttt{org}}^{(\text{unit})} = P_{\texttt{org}} - L_{\texttt{org}}.$$

Example linear tether (placeholder). A minimal tether uses declared maxima:

$$\kappa_A(ASL) = H_{\max}\frac{ASL}{100}, \qquad \kappa_C(CV) = H_{\max}\frac{CV}{100},$$

where $H_{\max}$ is a declared worst-case monthly burden scale for the entity (e.g. total staff hours at risk per month). More realistic calibrations can use nonlinear maps (e.g. piecewise or logistic) fitted to observed outcomes such as overtime hours, sick time, incident remediation time, and near-miss rates.

Interpretation. This calibration layer preserves the interface convenience of 0-100 indicators while making the dominance computation physically commensurate when required by the application. If calibration data are unavailable, StARS remains an ordinal routing instrument (as used in this paper's robustness sweeps), and $\Delta_{\texttt{org}}$ should be interpreted as a standardized index rather than a literal unit-valued margin.

4.4 Deviation Diagnostic Engine (DDE): ASL/CV vector construction

StARS computes ASL and CV as weighted sums of indicator sliders on a 0-100 scale.

ASL vector. Let

$$\mathbf{a} = [a_{\texttt{burn}},\, a_{\texttt{moral}},\, a_{\texttt{rule}}] \in [0,100]^3,$$

corresponding to Burnout/Exhaustion, Moral Injury/Distress, and Rule Bending (Volition), with weights $\mathbf{w}_A = [0.35, 0.35, 0.30]$:

$$ASL = \mathbf{w}_A^\top \mathbf{a} = 0.35\, a_{\texttt{burn}} + 0.35\, a_{\texttt{moral}} + 0.30\, a_{\texttt{rule}}. \tag{9}$$

CV vector. Let

$$\mathbf{c} = [c_{\texttt{proto}},\, c_{\texttt{kpi}},\, c_{\texttt{inc}},\, c_{\texttt{gaps}},\, c_{\texttt{ctrl}}] \in [0,100]^5,$$

corresponding to Protocol Complexity, Conflicting KPIs, Unit Incident Rate, Policy Gaps, and Control Failures, with weights $\mathbf{w}_C = [0.25, 0.20, 0.20, 0.20, 0.15]$:

$$CV = \mathbf{w}_C^\top \mathbf{c} = 0.25\, c_{\texttt{proto}} + 0.20\, c_{\texttt{kpi}} + 0.20\, c_{\texttt{inc}} + 0.20\, c_{\texttt{gaps}} + 0.15\, c_{\texttt{ctrl}}. \tag{10}$$

Worked example (UI-consistent). Using the screenshot inputs $\mathbf{a} = [72, 55, 46]$ and $\mathbf{c} = [32, 16, 23, 92, 28]$,

$$ASL = 0.35(72) + 0.35(55) + 0.30(46) = 58.25 \approx 58.3, \tag{11}$$

$$CV = 0.25(32) + 0.20(16) + 0.20(23) + 0.20(92) + 0.15(28) = 38.4. \tag{12}$$

$$RI_{\texttt{StARS}} = \tfrac{1}{2}(58.3 + 38.4) = 48.35 \approx 48.3.$$

Severity (magnitude) index. Define the StARS severity index as

$$RI_{\texttt{StARS}} = \frac{ASL + CV}{2}. \tag{13}$$

Unlike Δ_{org}, which routes correction, $RI_{\texttt{StARS}}$ quantifies magnitude.

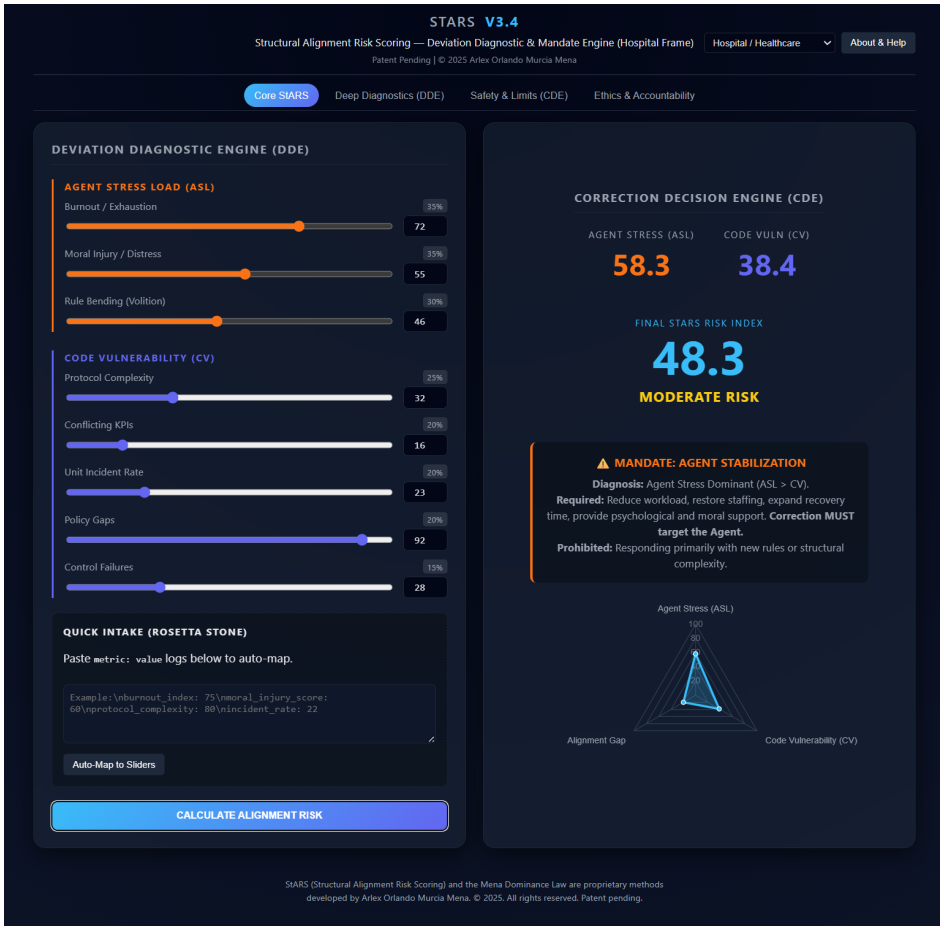

Figure 1: StARS v3.4 (Hospital/Healthcare) dashboard showing ASL and CV inputs, calculated risk index (48.3, Moderate), and the routed mandate (Agent Stabilization).

4.5 Attribution and intervention projection (stress release vs. structural correction)

StARS is designed to (i) localize which inputs are driving risk and (ii) project how specific interventions shift dominance routing and severity. Because ASL and CV

are linear in their indicator vectors, both attribution and counterfactual projections are closed-form.

Component attribution (where the risk comes from). Define component contributions as the weighted terms:

$$ASL_{\texttt{burn}} := 0.35\, a_{\texttt{burn}}, \quad ASL_{\texttt{moral}} := 0.35\, a_{\texttt{moral}}, \quad ASL_{\texttt{rule}} := 0.30\, a_{\texttt{rule}},$$

$$CV_{\texttt{proto}} := 0.25\, c_{\texttt{proto}}, \quad CV_{\texttt{kpi}} := 0.20\, c_{\texttt{kpi}}, \quad CV_{\texttt{inc}} := 0.20\, c_{\texttt{inc}},$$

$$CV_{\texttt{gaps}} := 0.20\, c_{\texttt{gaps}}, \quad CV_{\texttt{ctrl}} := 0.15\, c_{\texttt{ctrl}}.$$

Then $ASL = \sum ASL_i$ and $CV = \sum CV_j$. A practical attribution report is obtained by ranking the set

$$\{ASL_{\texttt{burn}}, ASL_{\texttt{moral}}, ASL_{\texttt{rule}}, CV_{\texttt{proto}}, CV_{\texttt{kpi}}, CV_{\texttt{inc}}, CV_{\texttt{gaps}}, CV_{\texttt{ctrl}}\}$$

from largest to smallest. The largest terms identify the dominant drivers of severity; the sign of $\Delta_{\texttt{org}} = ASL - CV$ identifies the dominant *channel* (agent vs. structure) for routing.

Intervention projection (release stress vs. fix structure). Let an intervention plan propose changes to the raw indicators:

$$\Delta\mathbf{a} = [\Delta a_{\texttt{burn}}, \Delta a_{\texttt{moral}}, \Delta a_{\texttt{rule}}], \qquad \Delta\mathbf{c} = [\Delta c_{\texttt{proto}}, \Delta c_{\texttt{kpi}}, \Delta c_{\texttt{inc}}, \Delta c_{\texttt{gaps}}, \Delta c_{\texttt{ctrl}}],$$

where negative values represent reductions (improvements). Because ASL and CV are linear,

$$ASL' = ASL + 0.35\,\Delta a_{\texttt{burn}} + 0.35\,\Delta a_{\texttt{moral}} + 0.30\,\Delta a_{\texttt{rule}},$$

$$CV' = CV + 0.25\,\Delta c_{\texttt{proto}} + 0.20\,\Delta c_{\texttt{kpi}} + 0.20\,\Delta c_{\texttt{inc}} + 0.20\,\Delta c_{\texttt{gaps}} + 0.15\,\Delta c_{\texttt{ctrl}}.$$

The projected dominance routing and severity are then:

$$\Delta'_{\texttt{org}} = ASL' - CV', \qquad RI'_{\texttt{StARS}} = \frac{ASL' + CV'}{2}.$$

This provides an immediate stress-release test: if the plan reduces agent-channel stress (negative Δa_i) but does not reduce structural fragility (small Δc_j), RI may improve while $\Delta_{\texttt{org}}$ remains structure-dominant, implying the correct routing remains structural even after stress relief.

Minimal change to exit a regime (optional planning rule). Given a dominance threshold t_{dom}, the condition for a *Dual-Path* regime is $|\Delta'_{\texttt{org}}| \le t_{\text{dom}}$. Therefore, any plan must satisfy:

$$-t_{\text{dom}} \;\le\; \Delta_{\texttt{org}} + (\mathbf{w}_A^\top \Delta\mathbf{a}) - (\mathbf{w}_C^\top \Delta\mathbf{c}) \;\le\; t_{\text{dom}},$$

where $\mathbf{w}_A = [0.35, 0.35, 0.30]$ and $\mathbf{w}_C = [0.25, 0.20, 0.20, 0.20, 0.15]$. This makes "release stress" vs. "structural correction" quantitatively comparable in the same routing variable.

4.6 Grace (response-based) and Decay (reciprocal burden)

StARS distinguishes *grace as headroom* (the instantaneous dominance margin) from *grace as response* (the observed restoration of margin after corrective action). It also tracks *decay* as the compounding restoration burden that grows disproportionately under sustained deficit exposure.

Response-based grace. If an intervention R is applied at time t and the entity is rescored after a window Δt, the response-based grace is:

$$G_{\mathtt{org}}(t) = \Delta_{\mathtt{org}}(t + \Delta t \mid R) - \Delta_{\mathtt{org}}(t). \tag{14}$$

Grace closes when interventions no longer restore dominance sufficiently, i.e., when $G_{\mathtt{org}}(t) < g_{\min}$ for a declared minimum restoration threshold $g_{\min}$.

Reciprocal decay (restoration burden). Let deficit exposure be:

$$D(t) = \max(0, -\Delta_{\mathtt{org}}(t)). \tag{15}$$

StARS models decay as a cumulative burden state $B(t)$ with reciprocal growth near the boundary:

$$B(t + \Delta t) = \lambda B(t) + \frac{D(t)}{\epsilon + \max(\Delta_{\mathtt{org}}(t), 0)}, \qquad 0 \le \lambda < 1,\ \ \epsilon > 0. \tag{16}$$

This form enforces the operational signature: equal deficit exposures produce larger restoration burden when the system has little positive headroom, and burden compounds when deficit persists.

4.7 Hospital simulation validation: real-data scenario sweeps and Monte Carlo

Hospital simulations are computational runs executed on real hospital operational datasets and are reported here as operational robustness checks for the StARS instrument, not as direct physical-embodiment validation. Each run fixes: (i) a mapping from raw metrics to 0-100 indicators; (ii) channel weights; (iii) a dominance threshold t_{dom}. The tool then computes $(ASL, CV, \Delta_{\mathtt{org}}, RI_{\mathtt{StARS}}, \mathtt{Mandate})$ for each entity-window record.

Robustness is evaluated by:

- **Threshold sweeps:** $t_{\mathrm{dom}} \in \{0.05, 0.10, 0.15\}$.
- **Weight variants:** agent-heavy vs structure-heavy reweighting (weights remain convex and sum to one).
- **Monte Carlo sensitivity:** random weight vectors sampled from Dirichlet distributions and thresholds sampled from $U[0.05, 0.20]$, producing hundreds to thousands of distinct runs on the same real dataset.

Across runs, we report mandate stability (fraction of records in each mandate class) and severity distribution stability (mean and upper quantiles), demonstrating that dominance-governed routing is robust to reasonable reparameterization.

Simulations executed and artifacts

We executed hospital-frame simulations on three publicly available operational datasets, each represented as an entity-window table: (i) Australia (state-year) using ED utilization and resource proxies; (ii) England (region-quarter) using overnight bed availability and occupancy (KH03); (iii) New Zealand (DHB-year) using publicly funded discharges, injury discharges, and procedure mix proxies. For each dataset we generated (a) a baseline StARS scoring file and (b) sensitivity outputs.

Deterministic scenario sweeps. For each dataset we ran a five-scenario sweep:

$$\texttt{baseline},\ \texttt{low\_threshold},\ \texttt{high\_threshold},\ \texttt{agent\_heavy},\ \texttt{structure\_heavy},$$

where the threshold variants use $t_{\mathrm{dom}} \in \{0.05, 0.10, 0.15\}$ and the weight variants shift mass toward ASL or CV while preserving convexity. These runs produce scenario-indexed mandate classifications and risk scores for every entity-window record.

Monte Carlo sensitivity runs. In addition, we executed 1000 Monte Carlo runs per dataset by sampling channel weights from Dirichlet distributions and sampling t_{dom} uniformly from $[0.05, 0.20]$. Each Monte Carlo run recomputes $(ASL, CV, \Delta_{\mathtt{org}}, RI_{\mathtt{StARS}}, \mathtt{Mandate})$ across all records. We summarize robustness by the distribution of mandate shares across runs and the distribution of risk statistics (mean and upper quintiles).

4.8 Data and artifact repository

For transparency and reproducibility, all MDL evidence artifacts for this paper are hosted at https://menadominancelaw.com/, including: (i) as-downloaded public-source data archives, (ii) processed/derived tables and simulation outputs (CSV), (iii) tool screenshots and figures, and (iv) reproducibility metadata (source citations, download dates, and checksums).

The StARS tool UI is available at https://starsframework.com/; representative StARS screenshots used in this paper are archived in the MDL repository above for permanence.

4.9 Public-data simulation drivers (non-patient operational proxies)

To demonstrate that StARS can be executed using *public, non-patient operational data*, we ran three proof-of-mechanism simulations: (i) Australia (AIHW; state-year ED utilization with a staffing-capacity anchor) [1, 2], (ii) England (NHS KH03; region-quarter beds open overnight and occupancy) [4], and (iii) New Zealand (publicly funded hospital and specialist services event data; DHB-year, FY 2018-19) [5]. These runs do *not* claim direct measurement of burnout, moral injury, protocol conflict, policy gaps, or control failures. They instantiate StARS using workload/strain and structural-leakage/capacity *proxies* that are available in public reporting.

The simulation output tables used below are exported as CSV in the artifact repository [6]: stars_internal_simulation_australia_state_year.csv, stars_internal_simulation_uk_beds_reg stars_internal_simulation_nz_dhb_2018_2019.csv, stars_hospital_sensitivity_australia_state_yea stars_hospital_sensitivity_uk_beds_region_quarter.csv, stars_hospital_sensitivity_nz_dhb_2018_ and stars_mc_summary_1000runs.csv.

Ethical guardrail (non-punitive use). Because public proxies can be misused, outputs in this section are presented as *instrument behavior under limited observability*, not as a ranking, performance scorecard, or justification for punitive action.

4.9.1 Normalization (robust min–max with clipping)

For each metric x computed over all entity-window records in a dataset, define $q_{\text{lo}} = Q_{0.05}(x)$ and $q_{\text{hi}} = Q_{0.95}(x)$ and apply:

$$N(x) = \text{clip}\left(\frac{x - q_{\text{lo}}}{q_{\text{hi}} - q_{\text{lo}}},\, 0,\, 1\right).$$

Capacity-deficit proxies are defined as inverted capacity terms (e.g., $n_{\text{Def}} := 1 - N(\texttt{Capacity})$).

Complete-case rule (Australia staffing anchor). The Australia state-year export contains 7 rows with missing staffing-anchor fields (and therefore missing CV, Δ, and RI). All Australia summaries and tables below use *complete-case* rows only ($n = 56$) by excluding records with missing RI.

Scale note (public-data proxies vs. StARS UI). In the public-data simulations we compute normalized proxies $ASL, CV \in [0,1]$ and $\Delta = ASL - CV$. *These public-proxy composites are not the DDE slider vectors*; they are dataset-specific proxy drivers used to exercise the routing kernel under limited observability. To align with the StARS UI convention (0-100), define $ASL^{(100)} = 100\,ASL$, $CV^{(100)} = 100\,CV$, and $\Delta^{(100)} = 100\,\Delta$. Accordingly, $t_{\text{dom}} = 0.1$ on $[0,1]$ corresponds to $t_{\text{dom}} = 10$ on $[0,100]$. We report ASL, CV, Δ on $[0,1]$ and report the magnitude index

$$RI_{\text{StARS}} = 100 \cdot \frac{ASL + CV}{2}$$

on a 0-100 scale.

4.9.2 Proxy definitions (dataset-specific driver mappings)

Table 7 summarizes the proxy mappings used to compute ASL (strain/workload channel) and CV (structural fragility channel). All coefficients below are declared *a priori* for the simulation; they are not fit to outcomes.

Table 7: Public-data StARS simulation drivers (proxy mapping; non-patient operational data).

Dataset	**Entity–window**	ASL **proxy (0–1)**	CV **proxy (0–1)**
Australia (AIHW)	state–year	$0.6\,N(\text{EDrate}) + 0.4\,N(\text{AHLUrate})$	$0.4\,N(\text{LUrate}) + 0.3\,N(\text{LUshare}) + 0.3\,\big(1 - N(\text{StaffRate})\big)$
England (KH03)	region–quarter	$0.5\,N(P) + 0.5\,N(O)$	$0.6\,\big(1 - N(A)\big) + 0.4\,N(P)$
New Zealand (FY 2018–19)	DHB–year	$0.5\,N(T) + 0.3\,N(P) + 0.2\,N(S^{inj})$	$0.35\,N(C^{ICDsub}) + 0.35\,N(C^{block}) + 0.15\,N(H^{proc}) + 0.15\,N(S^{inj})$

4.9.3 Dominance differential, mandate rule, and illustrative post-intervention projection

For each entity-window record:

$$\Delta = ASL - CV, \qquad \text{Mandate} = \begin{cases} \texttt{Agent Stabilization} & \Delta > t_{\text{dom}}, \\ \texttt{Structural Correction} & \Delta < -t_{\text{dom}}, \\ \texttt{Dual-Path Intervention} & |\Delta| \le t_{\text{dom}}. \end{cases}$$

Unless otherwise stated, results below use $t_{\text{dom}} = 0.1$ on the normalized scale.

To illustrate operational meaning (*not* a validated causal effect), we also compute an illustrative "post-intervention" projection by applying a simple targeted reduction to the *dominant channel* in proxy space: Agent Stabilization: $ASL' = 0.85\,ASL$, $CV' = CV$; Structural Correction: $CV' = 0.85\,CV$, $ASL' = ASL$; Dual-Path: $ASL' = 0.90\,ASL$, $CV' = 0.90\,CV$, then recompute Δ' and RI'_{StARS} and reclassify the mandate.

4.9.4 Executed outputs (baseline)

Table 8 reports baseline mandate shares and mean risk index in each dataset (complete-case rows only). Australia uses $n = 56$ complete-case state-years; England includes $n = 28$ region-quarters (7 regions $\times$ 4 quarters, FY 2024-25); and New Zealand includes $n = 21$ DHBs (FY 2018-19).

Table 8: Baseline public-data StARS outputs (complete-case). Mandate shares are percentages of entity–window records.

Dataset	n	Agent %	Structural %	Dual %	$\overline{RI}$	$\overline{RI}$ (post)
Australia (AIHW ED+resources)	56	12.5	85.7	1.8	35.3	31.7
England (KH03 beds)	28	46.4	35.7	17.9	52.0	46.8
New Zealand (DHB discharges)	21	4.8	66.7	28.6	47.5	42.8

Top-risk entity–windows (baseline). Tables 9-11 list the five highest RI_{StARS} records in each dataset under the baseline configuration. These are included for reproducibility and inspection; they are not presented as "worst" institutions.

Table 9: Australia (AIHW) baseline: top-5 by RI_{StARS} (state–year; complete-case).

State	Year	ASL	CV	Δ	RI	Mandate
NT	2022	1.000	0.674	0.326	83.7	Agent Stabilization
NT	2023	1.000	0.669	0.331	83.5	Agent Stabilization
NT	2021	1.000	0.665	0.335	83.2	Agent Stabilization
NT	2020	0.991	0.628	0.363	81.0	Agent Stabilization
NT	2018	0.966	0.638	0.328	80.2	Agent Stabilization

Table 10: England (KH03) baseline: top-5 by RI_{StARS} (region–quarter).

Region	**Q**	**Period**	*ASL*	*CV*	Δ	*RI*	**Mandate**
South West	Q4	January–March 2025	0.511	1.000	-0.489	75.5	Structural Correction
South West	Q3	October–December 2024	0.501	1.000	-0.499	75.1	Structural Correction
South East	Q4	January–March 2025	0.721	0.735	-0.014	72.8	Dual-Path Intervention
East Of England	Q4	January–March 2025	0.463	0.867	-0.404	66.5	Structural Correction
East Of England	Q3	October–December 2024	0.439	0.855	-0.417	64.7	Structural Correction

Table 11: New Zealand baseline: top-5 by RI_{StARS} (DHB; FY 2018–19).

DHB	*ASL*	*CV*	Δ	*RI*	**Mandate**
Canterbury	0.750	1.000	-0.250	87.5	Structural Correction
Waitemata	0.650	0.962	-0.312	80.6	Structural Correction
Waikato	0.633	0.951	-0.318	79.2	Structural Correction
Auckland	0.617	0.953	-0.336	78.5	Structural Correction
Counties Manukau	0.613	0.945	-0.332	77.9	Structural Correction

4.9.5 Sensitivity sweeps and Monte Carlo robustness

Threshold sweep. We evaluate threshold sensitivity using $t_{\mathrm{dom}} \in \{0.05, 0.10, 0.15\}$ on the normalized scale, holding proxy weights fixed (baseline mapping). Table 12 reports mandate shares under each threshold (complete-case where applicable).

Table 12: Threshold sweep for the mandate rule ($t_{\text{dom}} \in \{0.05, 0.10, 0.15\}$ on $[0, 1]$). Entries are mandate shares (%) of entity–window records.

Dataset	t_{dom}	**Agent %**	**Structural %**	**Dual %**
Australia (AIHW ED+resources)	0.05	12.5	87.5	0.0
Australia (AIHW ED+resources)	0.10	12.5	85.7	1.8
Australia (AIHW ED+resources)	0.15	12.5	78.6	8.9
England (KH03 beds)	0.05	57.1	35.7	7.1
England (KH03 beds)	0.10	46.4	35.7	17.9
England (KH03 beds)	0.15	42.9	28.6	28.6
New Zealand (DHB discharges)	0.05	9.5	71.4	19.0
New Zealand (DHB discharges)	0.10	4.8	66.7	28.6
New Zealand (DHB discharges)	0.15	0.0	61.9	38.1

Monte Carlo robustness. `We executed 1000 Monte Carlo runs per dataset (exported summary in stars_mc_summary_1000runs.csv), sampling proxy-weight variants and dominance thresholds. Table 13 reports mean mandate shares and mean risk statistics across runs (risk statistics on 0-100 scale).`

Table 13: Monte Carlo robustness summary (1000 runs per dataset). Entries are means across runs. Risk statistics are reported on a 0–100 scale.

Dataset	**Runs**	**Agent %**	**Structural %**	**Dual %**	$\overline{RI}$	$RI_{0.90}$
Australia (AIHW ED+resources)	1000	9.6	72.5	17.9	38.1	69.1
England (KH03 beds)	1000	40.3	32.9	26.8	52.4	72.1
New Zealand (DHB discharges)	1000	13.1	53.8	33.1	44.9	74.8

4.9.6 Limitations (public-data proxy scope)

`These simulations are constrained by public operational reporting. Public datasets do not directly measure burnout, moral injury, protocol conflict, policy gaps, or control failures; therefore` ASL `and` CV `are driven by workload/strain and structural leakage/capacity proxies. Outputs are interpreted as a proof-of-mechanism demonstration that dominance routing remains well-defined and reproducible under limited observability, not as a causal model of institutional dynamics or a high-stakes ranking instrument.`

5 Grid Simulation Validation: Dual-Layer Viability Governor on CAISO OASIS

This section provides real-data simulation validation for an embodiment-specific layered controller applied in an infrastructure embodiment. The physical layer governs instantaneous adequacy using the dominance margin $\Delta_{\text{grid}}(t)$ computed from CAISO system totals. The structural layer governs admissibility using operating reserve headroom derived from CAISO ancillary-service reserve reporting. The dual-layer governor implements the principle: *even when power balance is favorable, low reserve headroom can constrain admissible operation.*

Scope limitation. This validation uses system-level adequacy and reserve-headroom proxies (ENE_SLRS totals and AS_OP_RSRV) and does not model line thermal limits, N-1 contingency violations, or topology constraints. The aim is to isolate dominance-governed policy effects on sustained deficit exposure under fixed boundary definitions using public time series [3]. This section therefore supports a simulation-validation claim, not a full physical-grid causal claim.

5.1 Data sources and window

We use CAISO public OASIS exports for a common 24-hour window:

- **ENE_SLRS (DAM):** CAISO totals for generation, imports, exports, and total system load (hourly).
- **AS_OP_RSRV:** operating reserves maintained (*percent*), resampled to hourly mean and aligned to ENE_SLRS timestamps.

Window: 2025-01-13 07:00 UTC $\rightarrow$ 2025-01-14 07:00 UTC (24 hourly points after alignment). These files are archived in the artifact repository under their original CAISO filenames, along with the derived simulation outputs (CSV).

5.2 Embodiment contract and dominance traces

System-adequacy layer: operational dominance. Define physical Potential and Load using CAISO totals:

$$\begin{aligned} P_{\text{grid}}(t) &= \texttt{ISO\_TOT\_GEN\_MW}(t) + \texttt{ISO\_TOT\_IMP\_MW}(t) \\ &\quad - \texttt{ISO\_TOT\_EXP\_MW}(t), \end{aligned} \tag{17}$$

$$L_{\text{grid}}(t) = \texttt{ISO\_TOT\_LOAD\_MW}(t). \tag{18}$$

The actuator is curtailable load $u(t) \in [0, u_{\max}]$ (interpretable as demand response / non-essential shedding):

$$L_{\text{eff}}(t) = L_{\text{grid}}(t) - u(t). \tag{19}$$

The physical dominance margin is:

$$\Delta_{\text{grid}}(t) = P_{\text{grid}}(t) - L_{\text{eff}}(t). \tag{20}$$

Structural layer: admissibility from reserve headroom. Let $R(t)$ be the hourly mean of operating reserves (percent) from AS_OP_RSRV. Define a reserve floor $r_{\min}$ and the structural margin:

$$\Delta_{\texttt{arch}}(t) = R(t) - r_{\min}, \qquad r_{\min} = 7.5\%. \tag{21}$$

We convert structural margin to a smooth admissibility gate:

$$\begin{aligned} g(t) &= \sigma(k\,\Delta_{\texttt{arch}}(t)) \\ &= \frac{1}{1+\exp(-k\,\Delta_{\texttt{arch}}(t))}, \qquad k = 4. \end{aligned} \tag{22}$$

The reserve floor $r_{\min}$ and gate slope k are declared policy parameters (not fitted); they represent a conservative reserve-headroom boundary and a smooth approximation to a hard admissibility gate.

Dual-layer governed viability. Define the governed (true) viability margin as:

$$\Delta_{\texttt{true}}(t) = \Delta_{\texttt{grid}}(t) \cdot g(t). \tag{23}$$

When reserves are ample ($\Delta_{\texttt{arch}}$ high), $g(t) \approx 1$ and $\Delta_{\texttt{true}}$ tracks physical adequacy. When reserve headroom compresses, $g(t)$ reduces admissible operation even if raw power balance is favorable.

5.3 Policies and controller implementation

We compare three policies under identical boundary definitions (same data, same time window, same thresholds):

- **Baseline:** no curtailment $u(t) = 0$.
- **Single-layer MDL:** MDL acts on $\Delta_{\texttt{grid}}(t)$.
- **Dual-layer MDL:** MDL acts on $\Delta_{\texttt{true}}(t)$.

MDL control law (dominance + burden + grace). Let $\Delta_\star(t)$ denote the controlled margin: $\Delta_{\texttt{grid}}(t)$ (single-layer) or $\Delta_{\texttt{true}}(t)$ (dual-layer). We implement a burden state $B(t)$ that accumulates under sustained deficit exposure:

$$B_{t+1} = \lambda_B B_t + k_B\,\max(0, -\Delta_{\texttt{grid}}(t)), \qquad \lambda_B = 0.5,\ \ k_B = 0.002. \tag{24}$$

A grace/hysteresis rule defines a protective regime: enter protect mode after $N_{\texttt{enter}} = 1$ deficit hour; exit only after $N_{\texttt{exit}} = 2$ consecutive hours above $\Delta_{\texttt{exit}} = 500$ MW. The actuation is:

$$\begin{aligned} u(t) = \mathrm{clip}\Big(& a(\texttt{mode})\max(0, -\Delta_\star(t)) \\ & + b_B\,B(t),\ \ 0,\ \ u_{\max}\Big), \end{aligned} \tag{25}$$

with $u_{\max} = 1500$ MW, $a_{\texttt{normal}} = 0.8$, $a_{\texttt{protect}} = 1.1$, and $b_B = 250$.

Interpretive guardrail. The reserve-headroom gate and the dual-layer combination are embodiment-specific control constructions for this grid simulation. They are not additional core-law axioms; they are the declared way this embodiment combines physical adequacy and reserve admissibility for conservative decision logic.

5.4 Evaluation metrics

We summarize the dominance-trace outcomes using:

- **Time in deficit (hours):** $\sum \mathbb{1}[\Delta_{\text{grid}}(t) < 0]\Delta t$,
- **Deficit area (MWh):** $\sum \max(0, -\Delta_{\text{grid}}(t))\Delta t$,
- **Shedding energy (MWh):** $\sum u(t)\Delta t$,
- **Minimum margin (MW):** $\min_t \Delta_{\text{grid}}(t)$,
- **Gate statistics:** $\min_t g(t)$.

5.5 Results (real-data simulation)

Across the 24-hour window, baseline operation exhibits sustained physical deficit exposure. Both MDL policies reduce deficit exposure substantially versus baseline via curtailment. Single-layer MDL yields the lowest deficit hours and deficit area in this window. Dual-layer MDL introduces reserve-headroom admissibility gating, which reduces curtailment energy relative to single-layer MDL while leaving higher residual deficit time/area; this evidences an explicit tradeoff between curtailment and deficit exposure under fixed boundary definitions.

Table 14: CAISO dual-layer grid simulation results (24h window; ENE_SLRS + AS_OP_RSRV).

Policy	Deficit hrs	Deficit area (MWh)	Shedding (MWh)	min Δ_{grid} (MW)
Baseline	18	18318.12	0.00	-1687.84
MDL (single)	5	603.62	18852.76	-289.36
MDL (dual)	9	1584.53	17786.78	-471.75

Structural layer summary. Using $r_{\min} = 7.5\%$ and $k = 4$, the admissibility gate satisfies $\min_t g(t) = 0.4005$ over the window (i.e., reserves approach the floor but do not collapse to zero admissibility). This gate modulates the governed margin $\Delta_{\text{true}}(t)$ and changes when the controller prioritizes curtailment.

Reproducibility artifacts. Reproducibility outputs include the aligned hourly time series for $P_{\text{grid}}(t)$, $L_{\text{grid}}(t)$, $R(t)$, $\Delta_{\text{arch}}(t)$, $g(t)$, $\Delta_{\text{grid}}(t)$, $\Delta_{\text{true}}(t)$, and $u(t)$ exported as CSV under the filenames referenced in the artifact repository.

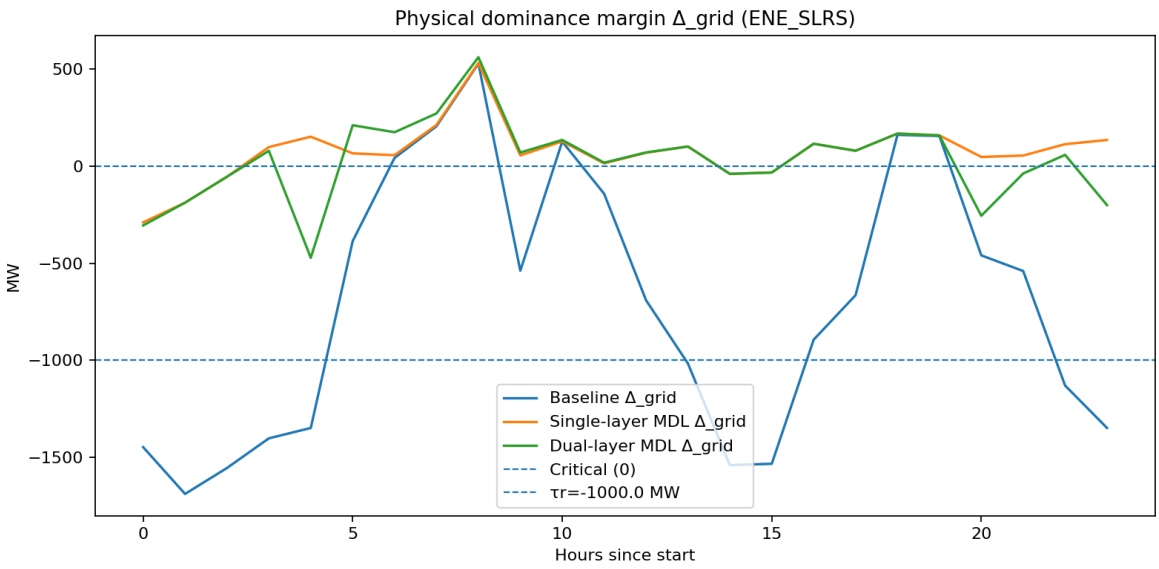

Figure 2: Physical dominance margin $\Delta_{\text{grid}}(t)$ (CAISO ENE_SLRS) under baseline, single-layer MDL, and dual-layer MDL. The MDL policies reduce sustained deficit exposure under identical boundary definitions; the dual-layer governor modulates the response using reserve-headroom admissibility.

5.6 MDL Yield Parameter (grid)

For the CAISO simulation, we define the MDL yield parameter as avoided deficit area per unit of shedding energy:

$$Y_{\text{grid}} = \frac{D_0 - D_1}{S}, \tag{26}$$

where D_0 is the baseline deficit area, D_1 is the policy deficit area, and S is total shedding energy over the same fixed evaluation window.

Using Table 14, the single-layer policy gives

$$D_0 = 18318.12, \qquad D_1 = 603.62, \qquad S = 18852.76,$$

hence

$$Y_{\text{grid}}^{\text{single}} = \frac{18318.12 - 603.62}{18852.76} \approx 0.9396. \tag{27}$$

For the dual-layer policy,

$$D_0 = 18318.12, \qquad D_1 = 1584.53, \qquad S = 17786.78,$$

hence

$$Y_{\text{grid}}^{\text{dual}} = \frac{18318.12 - 1584.53}{17786.78} \approx 0.9408. \tag{28}$$

Thus, in this 24-hour window, both governed policies convert shedding into avoided deficit area at roughly 0.94 MWh of deficit-area reduction per 1 MWh of shedding. The dual-layer controller preserves slightly more avoided-deficit-area per unit shed, while accepting higher residual deficit exposure than the single-layer controller.

This is a domain- and regime-specific yield defined on deficit-area and shedding-energy metrics; it is not numerically comparable to the battery yield unless benefit and cost are first mapped into a common canonical form.

6 Threats to validity and scope limits

Battery experiments. (1) Power is not matched across policies: M intentionally reduces current near cutoff, so runtime is policy-mediated. (2) Ambient temperature is uncontrolled; although both policies were executed across the observed temperature range, residual confounding may remain. (3) Sample size in terms of physical packs is small (four packs total, two of which are the same model). The paired-run count is large, but pack-level generalization is limited. (4) Recovery-voltage shifts are protocol-defined signatures and are not treated as degradation evidence. A multi-cycle longitudinal design is required for degradation-rate claims.

StARS instrument evidence. StARS results in this paper are validation of instrument behavior (routing logic, stability under sensitivity), not causal proof about hospital outcomes. The tool can be calibrated into commensurate operational units when calibration data exist; absent that, it is an ordinal routing instrument.

Grid simulation. The CAISO validation uses system-level adequacy and operating-reserve proxies and does not model network topology, N-1 contingencies, unit commitment constraints, or locational constraints. The goal is to isolate dominance-governed policy effects on deficit exposure under fixed boundary definitions using public time series.

7 Conclusion

This paper evaluates the operational claim of the Mena Dominance Law across three evidence classes under fixed boundary rules: (i) battery discharge experiments, which provide the primary executed physical evidence; (ii) StARS instrument-level operational evidence, showing how dominance, grace, and burden logic can route corrective action in a sociotechnical control setting; and (iii) real-data grid simulation validation using CAISO OASIS time series.

Across these evidence classes, the consistent pattern is policy-conditional: when the boundary is held fixed and the actuator is the load policy or correction routing, dominance-aware control is associated with reduced sustained deficit exposure. In the battery embodiment, that pattern is associated with increased usable work to a fixed cutoff and with measurable shifts in boundary-adjacent recovery metrics within the declared protocol. In the StARS instrument, it is associated with robust routing behavior. In the grid simulation, it is associated with reduced deficit hours and deficit area. The time-integration (dose) formulation and the Mena Horizon T_M further convert the dominance signal into a predictive countdown to boundary violation under declared load models, enabling proactive rather than reactive control.

Finally, Appendix A clarifies scope: the MDL kernel is not limited to the executed systems in this paper. Any bounded system with measurable headroom, measurable demand, and a real actuator can instantiate the same contract. For complex deployments, layered governance helps prevent the dominant operational error: spending a primary surplus while silently exhausting a hidden constraint that ultimately determines viability.

A Illustrative Extensions Beyond the Executed Evidence

This appendix clarifies scope. The executed evidence in this paper is limited to the battery experiments, the StARS instrument analyses, and the CAISO grid simulation. The materials below are *illustrative control-ready extensions* of the same dominance-governed logic. They are included to show portability of the framework under the embodiment contract; they are not additional executed evidence claims for Paper 2 and should not be read as carrying the same evidentiary status as the battery experiments or the reported simulations.

Minimal embodiment checklist. A control-ready extension is admissible only when it declares:

1. a hard boundary condition;
2. commensurate definitions of Potential and Load;
3. a measurement or estimation cadence;
4. an actuator or intervention class that can modify effective load, effective support, or admissible operation.

A.1 Catalog of illustrative control-ready extensions

Table 15: Illustrative extensions beyond the executed evidence in this paper. These are embodiment templates, not simulations executed here.

Domain	**Boundary (hard constraint)**	**Potential / Budget**	**Load / Dose**	**Actuator** $u(t)$
Autonomous drones (mission viability)	Return-to-home feasible; stall/wind envelope	Energy remaining E_{rem} (Wh)	Energy-to-home $\int P_{\text{req}}(t)\,dt$ (Wh)	Speed, altitude, route, payload drop, loiter cancel
Spacecraft / satellites (ops)	Propellant reserve; thermal limits; link margin floor	Fuel budget + thermal headroom	Consumption + dissipation dose	Safe-mode, duty cycling, attitude, comm scheduling
Cyber containment (SOC)	Max blast radius; time-to-contain SLA	Containment throughput	Spread / alert burden and backlog dose	Isolation, segmentation, throttling, automation
Cloud SRE / error budget governance	Latency SLO; error budget floor	Error budget + capacity headroom	Request rate, fault rate, error accumulation	Rate limiting, feature flags, brownouts, autoscale
Traffic networks	Queue spillback; gridlock onset; max delay	Throughput capacity	Arrival demand and congestion dose	Signal timing, ramp metering, reversible lanes
Reservoir / water operations	Minimum storage; flood release cap	Storage headroom	Demand / inflow imbalance dose	Release scheduling, ration tiers, pumping rules
Manufacturing drift / yield	Out-of-spec rate; CpK floor; scrap ceiling	Process headroom to spec	Drift / noise accumulation and rework burden	Setpoints, line speed, inspection cadence, holds
Airspace flow management	Separation minima; sector capacity; fuel reserves	Sector capacity + per-flight margin	Demand rate + weather complexity dose	Metering, reroutes, ground stops, miles-in-trail
Robotics cells	Collision risk; torque / thermal limits	Compute / thermal headroom + safety margin	Task load + uncertainty dose	Speed limits, replanning, torque limiting, pauses

A.2 Three illustrative instantiations

A.2.1 Time-limited missions: drone return-to-home horizon

A drone mission is bounded by return feasibility: the energy required to return to home must not exceed remaining energy under current wind and routing conditions. A budget-dose dominance form may be written as

$$\Delta_E(t) = E_{\text{rem}}(t) - \int_t^{t+T_{\text{home}}} P_{\text{req}}(\tau)\,d\tau, \tag{29}$$

where P_{req} is the power required to maintain flight and overcome wind along a chosen route and altitude profile. The actuator bundle $u(t)$ (speed setpoint, altitude, route replanning, loiter cancellation, payload drop) modifies $P_{\text{req}}(\cdot)$ and therefore the dose. Under a locally constant load approximation,

$$T_M \approx \frac{\Delta_E(t)}{P_{\text{req}}(t)}. \tag{30}$$

A.2.2 Containment governance: cyber incident response

Containment is a bounded system where failure is not only a breach but uncontrolled spread. A rate-form dominance may be written in matched units such as hosts per hour:

$$\Delta_R(t) = P_R(t) - L_R(t), \tag{31}$$

where $P_R(t)$ is containment throughput and $L_R(t)$ is effective spread pressure. When backlog and restoration burden matter, the cumulative uncontained spread becomes a deficit dose that predicts when the SOC saturates. The actuator bundle includes segmentation, isolation rules, credential rotation, service throttling, and automation policy.

A.2.3 Process viability: manufacturing drift and yield collapse

Manufacturing viability is bounded by out-of-spec rate and process capability. A budget-dose form may be declared as

$$\Delta_E(t) = E_{\text{spec}} - \int_{t_0}^{t} \dot{E}_{\text{drift}}(\tau)\, d\tau, \tag{32}$$

where E_{spec} is a declared spec budget and $\dot{E}_{\text{drift}}$ is a drift / variance rate inferred from sensor data and yield outcomes. The actuator bundle includes setpoint adjustment, line-speed changes, tool recalibration, inspection cadence, and hold / rework policies.

A.3 Layered governance in multi-constraint systems

Many real deployments are governed by multiple margins rather than a single global one. Vehicles, aircraft, satellites, data centers, and large infrastructures often have several coupled subsystems, each with its own boundary and actuator. A layered controller may compute a declared margin vector

$$\boldsymbol{\Delta}(t) = \big[\Delta_1(t), \Delta_2(t), \ldots, \Delta_n(t)\big],$$

and combine those margins conservatively. One example is a gated effective margin:

$$\Delta_{\text{eff}}(t) = \Delta_{\text{primary}}(t) \cdot \prod_{i \in \mathcal{C}} \sigma\big(k_i\, \Delta_i(t)\big), \tag{33}$$

where $\mathcal{C}$ indexes critical layers and $\sigma(\cdot)$ is a smooth gate. A stricter example is

$$\Delta_{\text{eff}}(t) = \min_i \Delta_i(t). \tag{34}$$

These layered operators are interpretive governance constructions rather than additional core-law axioms. They specify how multiple declared margins may be combined for conservative decision logic in multi-constraint systems.

Operationally, the point is simple: one layer's surplus does not erase another layer's hard boundary.

B Algorithm Listing (Replication Reference)

This appendix provides compact pseudocode for the dominance-governed controllers used in the executed experimental and simulation sections. Parameters are declared in-text in the relevant sections; the logic below is sufficient to reproduce the control traces given the exported time series.

Scope of the replication listings. The pseudocode in this appendix is tied to the executed evidence of Paper 2:

- the continuous battery viability governor corresponding to the battery embodiment;
- the single-layer and dual-layer grid controllers corresponding to the CAISO simulation.

The illustrative extensions in Appendix A are not additional executed algorithms for this paper.

Battery-controller purpose. The battery controller implements a *continuous* viability ceiling on discharge current at each sample (e.g., once per second), using the admissible-curre rule

$$I_{\max}(t) \approx \frac{V_{\mathtt{rest}}(t) - V_{\mathtt{cut}}}{\hat{R}_{\mathtt{use}}(t)}.$$

Given a fixed cutoff $V_{\mathtt{cut}}$, it estimates a near-rest voltage $V_{\mathtt{rest}}(t)$ and an effective use-resistance $\hat{R}_{\mathtt{use}}(t)$ online, computes the current ceiling $I_{\max}(t)$, and clips the demanded current to this ceiling. The four-step ET5406A+ drop-down schedules used in the physical experiments are quantized approximations of this per-sample viability governor.

Battery-controller inputs.

- Measured terminal voltage $V_{\mathtt{term}}[k]$ and current $I_{\mathtt{meas}}[k]$ at discrete sample index k.
- Fixed cutoff boundary $V_{\mathtt{cut}}$ for the pack.
- External current demand $I_{\mathtt{demand}}[k]$.

Battery-controller parameters.

- $R_{\min}$: minimum admissible use-resistance.
- $I_{\mathtt{rest,max}}$: current magnitude below which a sample is treated as near-rest for updating $V_{\mathtt{rest}}$.
- $I_{\mathtt{R,min}}$: minimum current magnitude used when updating $\hat{R}_{\mathtt{use}}$.
- $\alpha_V \in (0,1)$: smoothing factor for the near-rest voltage proxy.
- $\alpha_R \in (0,1)$: smoothing factor for the resistance estimate.
- $I_{\mathtt{hw,max}}$: hardware current ceiling.
- $\tau_{r,\mathtt{batt}} \leq 0$: optional irreversibility / violation threshold on the dominance margin.

Battery-controller outputs.

- Commanded discharge current $I_{\mathrm{cmd}}[k]$.
- Dominance margin trace $\Delta_{\mathrm{batt}}[k] = V_{\mathrm{term}}[k] - V_{\mathrm{cut}}$.
- Optional violation flag $S_{\mathrm{batt}}[k] = \mathbf{1}[\Delta_{\mathrm{batt}}[k] \le \tau_{r,\mathrm{batt}}]$.
- Optional logs of $V_{\mathrm{rest}}[k]$, $\hat{R}_{\mathrm{use}}[k]$, and $I_{\max}[k]$.

B.1 Battery viability governor (continuous per-sample control)

Algorithm 1: Pseudocode for a continuous battery viability governor. The controller estimates near-rest voltage and effective use-resistance online, computes an admissible-current ceiling under a fixed cutoff boundary, clips demanded current to that ceiling, and logs the commanded trace and dominance margin.

```
Algorithm: MDL_Battery_Viability_Governor (continuous per-sample control)

Inputs (arrays indexed by sample k = 1..K):
  V_term[k] # measured terminal voltage (V)
  I_meas[k] # measured current (A), sign convention arbitrary
  I_demand[k] # requested discharge current magnitude (A), assume >= 0
  V_cut # fixed cutoff boundary voltage (V)

Parameters / constants:
  V_open_init # initial near-rest/open-circuit voltage estimate (V)
  R_init # initial use-resistance estimate (ohms)
  R_min # minimum admissible resistance (ohms)
  I_rest_max # if |I_meas| <= I_rest_max, treat as near-rest sample (A)
  I_R_min # minimum current magnitude for resistance update (A)
  alpha_V # smoothing factor for V_rest_hat, 0 < alpha_V <= 1
  alpha_R # smoothing factor for R_use_hat, 0 < alpha_R <= 1
  I_hw_max # hardware current ceiling (A)
  tau_r_batt # optional violation threshold on delta (V), typically <= 0

Initialization:
  V_rest_hat = V_open_init
  R_use_hat = R_init

  I_cmd_trace = []
  I_max_trace = []
  V_rest_log = []
  R_use_log = []
  Delta_batt = []
  S_batt = []

For each sample k = 1..K:

  V = V_term[k]
  I = I_meas[k]
  I_req = max(0, I_demand[k])

  # 1) Update near-rest voltage estimate when current is small
  if abs(I) <= I_rest_max:
      V_rest_hat = (1 - alpha_V) * V_rest_hat + alpha_V * V
```

```
  # 2) Update use-resistance estimate from observable sag under load
  I_mag = max(abs(I), I_R_min)
  sag = V_rest_hat - V
  if (abs(I) >= I_R_min) and (sag > 0):
      R_sample = sag / I_mag
      R_use_hat = (1 - alpha_R) * R_use_hat + alpha_R * R_sample

  # 3) Compute admissible-current ceiling from dominance headroom
  R_eff = max(R_use_hat, R_min)
  headroom = V_rest_hat - V_cut
  I_max = headroom / R_eff
  I_max = clip(I_max, 0.0, I_hw_max)

  # 4) Command is demand clipped by admissible ceiling
  I_cmd = min(I_req, I_max)

  # 5) Dominance proxy and optional violation flag
  delta = V - V_cut
  violation = 1 if (delta <= tau_r_batt) else 0

  # 6) Log traces
  I_cmd_trace.append(I_cmd)
  I_max_trace.append(I_max)
  V_rest_log.append(V_rest_hat)
  R_use_log.append(R_use_hat)
  Delta_batt.append(delta)
  S_batt.append(violation)

Return:
  I_cmd_trace, Delta_batt, S_batt,
  plus optional logs: I_max_trace, V_rest_log, R_use_log
```

B.2 Grid MDL controller (single-layer and dual-layer)

Grid-controller inputs. Hourly series: $P_{\mathtt{grid}}(t)$, $L_{\mathtt{grid}}(t)$, and, for the dual-layer case, reserve series $R(t)$.

Grid-controller parameters. Declared parameters: $u_{\max}$, $r_{\min}$, k, (λ_B, k_B), $(a_{\mathtt{normal}}, a_{\mathtt{protect}}, b_B)$, and grace hysteresis $(N_{\mathtt{enter}}, N_{\mathtt{exit}}, \Delta_{\mathtt{exit}})$.

Grid-controller outputs. Curtailment $u(t)$ and the resulting margins $\Delta_{\mathtt{grid}}(t)$ and $\Delta_{\mathtt{true}}(t)$.

Interpretive note. In the dual-layer case, the reserve-headroom gate is an embodiment-specific control construction for this paper's grid simulation. It is not a new core-law axiom; it is the declared way this particular embodiment combines physical adequacy and reserve admissibility for conservative decision logic.

Algorithm 2: Pseudocode for the grid MDL controller used in the single-layer and dual-layer simulations.

```
Algorithm: MDL_Grid_Controller (single-layer or dual-layer)

Initialize:
  B = 0
  mode = NORMAL
```

```
  deficit_streak = 0
  exit_streak = 0

For each timestep t = 1..T:
  if DUAL_LAYER:
      delta_arch = R[t] - r_min
      g = 1 / (1 + exp(-k * delta_arch))
  else:
      g = 1

  delta_grid_nom = P_grid[t] - L_grid[t]
  delta_star = delta_grid_nom * g

  D = max(0, -delta_grid_nom)
  B = lambda_B * B + k_B * D

  if delta_grid_nom < 0:
      deficit_streak += 1
      exit_streak = 0
  else:
      deficit_streak = 0
      if delta_grid_nom > Delta_exit:
          exit_streak += 1
      else:
          exit_streak = 0

  if mode == NORMAL and deficit_streak >= N_enter:
      mode = PROTECT
  if mode == PROTECT and exit_streak >= N_exit:
      mode = NORMAL

  a = a_protect if mode == PROTECT else a_normal
  u_cmd = a * max(0, -delta_star) + b_B * B
  u[t] = clip(u_cmd, 0, u_max)

  L_eff = L_grid[t] - u[t]
  delta_grid[t] = P_grid[t] - L_eff
  delta_true[t] = delta_grid[t] * g

Return u, delta_grid, delta_true
```

References

[1] Australian Institute of Health and Welfare (AIHW). *Use of emergency departments for lower urgency care, 2020-21 and 2021-22*. 2024. Available at: https://www.aihw.gov.au/reports/primary-health-care/use-eds-lower-urgency-care-2020-21-and-2021-22/contents/about. Accessed 2026-03-03.

[2] Australian Institute of Health and Welfare (AIHW). *Hospital resources 2023-24: tables (Excel workbook)*. 2024. Available at: https://www.aihw.gov.au/getmedia/98f5bcd7-34dd-48af-aab4-dbeab291b3f4/hospital-resources-tables-2023-24.xlsx.aspx. Accessed 2026-03-03.

[3] California Independent System Operator (CAISO). *OASIS: Open Access Same-Time Information System*. 2026. Available at: https://oasis.caiso.com/. Accessed 2026-03-03.

[4] NHS England. *Bed availability and occupancy data: overnight (KH03)*. 2025. Available at: https://www.england.nhs.uk/statistics/statistical-work-areas/bed-availability-and-occupancy/bed-data-overnight/. Accessed 2026-03-03.

[5] New Zealand Ministry of Health. *Publicly funded hospital and specialist services event data and statistics*. 2019. Available at: https://www.health.govt.nz/health-system-data-resources/health-statistics/health-statistics-and-data-sets/publicly-funded-hospital-and-specialist-services-event-data-and-statistics. Accessed 2026-03-03.

[6] Arlex Orlando Murcia Mena. *Mena Dominance Law: data and artifact repository (Paper 2 evidence bundle)*. 2026. Available at: https://menadominancelaw.com/. Accessed 2026-03-03.