

Contents

Consequence-Equilibrium Attention: the shape, and the programme that builds it	4
Abstract	4
1. Introduction	4
2. The shape	7
2.0 Occupancy — cited before the shape is named	7
2.1 Definitions	8
2.2 Propositions	9
2.3 What the shape computes	13
2.4 The topological layer (restricted)	13
2.5 The cost law and the kernel path	14
2.6 Lean targets (build order; every [M] is pending <code>lake build</code> ; a missing Mathlib lemma moves it to [S]; nothing is cited as proved before it builds — L-LEAN, P-11)	14
2.7 Limits carried to §8	15
3. Why softmax cannot occupy this ground — exactly as the theorems license it	15
3.1 Obstruction 1 — per-row independence (definitional)	16
3.2 Obstruction 2 — hops need depth (skyline unconditional, lower bounds conditional)	16
3.3 The DET-class relocation — what “one operator” may claim	16
3.4 Obstruction 3 — composition	17
3.5 Obstruction 4 — non-negativity, WITHDRAWN as written	17
3.6 D-1, stated with the right predictor and the right census	17
3.7 The skyline table — the only place a “cannot” sentence may stand	17
3.8 The honest sentence	18
3.9 Limits carried to §8	18
4. What the record proved, measured and refuted	19
4.1 Proved	19
4.2 Measured	26
4.3 Refuted, by mechanism	32
4.4 Where the campaign stands	40
5. THE PROGRAMME — the development plan for Consequence-Equilibrium Attention	44
5.0 The shape, and the claim the programme can at best earn	44
5.1 How to use this plan	45
5.2 The laws in force	47
5.3 The phases and their milestone cards	50
5.4 The critical path, and the evenings count derived from it	81
5.5 The predictions ledger, with counters	82

5.6 The standing attacks	87
5.7 The cost table by phase, and the Kaggle gate	91
5.8 The Lean ledger	93
5.9 The leap-dossier discipline	96
5.10 What dies if what	97
5.11 Deliverables, the return report, and how to keep this plan current	99
6. The reference apparatus — beds, floors, skylines, controls, metrics, census, binds, certificates, manifest, prices	100
A.0 What the six refutations decided, in one table	101
A.1 Beds	102
A.2 Labels and heads	106
A.3 Information floors (L-FLOOR)	110
A.4 The native skyline and the fellow approximators (R-SKY, D-1)	112
A.5 Metrics (C11 adopted, with the refuters' repairs)	116
A.6 Control arms per bed	118
A.7 Domain census (L-DOM) — printed per draw batch, before any arm is trained; 0 % on any line blocks the bed (D-4)	119
A.8 Binds, with planted negatives and the $O(1)$ failure each must produce	120
A.9 Certificates — what δ , on what object, in vector units	135
A.10 The identity manifest — field list (against L-2, READ V20_R15_LEAP_LEDGER.md:23)	139
A.11 The pricing rule	140
A.12 Deleted (unanswered KILL) — one line each	145
A.13 Limits (collected once)	146
7. Prior art and occupancy	146
7.0 How to read the table	147
7.1 The occupancy table	147
7.2 The delta, narrowly	158
7.3 The record's own prior-art defect	159
8. Limits	160
Appendix A. The two ledgers	161
A.1 Proposition ledger — the shape, judged	161
A.2 Bind ledger — binds, controls, censuses, certificates, kills	184
1. Binds	185
2. Controls and skylines	193
3. Census lines (L-DOM)	195
4. Certificates	196
5. Kills (pre-registered, thresholds frozen here)	197
6. Predictions with counters (L-SIGN, D-7) — filed before P3	201

7. Tally	202
Appendix B. Vocabulary — the field’s name for each construction	203
B.1 The operator and its read	203
B.2 The boundary rows and what they read	204
B.3 Interventions and decisions	205
B.4 The gate, the segmentation and the topology	206
B.5 The instruments and the statistics	206
B.6 Two words this campaign used in two senses, and what each cost	207

Consequence-Equilibrium Attention: the shape, and the programme that builds it

Invented by Teerth Sharma · <https://github.com/teerthsharma/resolvent> · teerths57@gmail.com

Front matter, coordinator, 2026-09-03 (revised after the author's clarification that the body of this paper is the development plan). Evidence classes: RUN (executed on this box this session), READ (path:line at HEAD 207e7b9), CITED [V] (primary page fetched, title matched; bib key), DERIVED (steps shown). The assembler keeps them.

Abstract

This paper is a development programme, and it states first what it is a programme for. Self-attention reads a context as one convex mixture per query, each query's read independent of every other's — the right operator for the next token and, on the one task shape every bed in this repository asked (a scalar at one position), the operator a single attention layer is provably Bayes-optimal for (Marion et al. 2025; Duranthon et al. 2026). Fifteen rounds of this campaign raced softmax on that ground and could not win, for that reason (MISTAKES.md D-1). What the author wants is different ground: the *consequence* of an intervention on a context, the *safest move* under several constraints at once, and the *next transient state* toward equilibrium. The shape that carries those three is Consequence-Equilibrium Attention (CEQ): the causal resolvent read $O = (1 - \gamma)P(I - \gamma P)^{-1}V$ of Fagnou et al. (ChaCAL, EMNLP 2024), taken over the three-corner operator P this repository proved in Lean, with two things no source in eight prior-art sweeps holds together — absorbing boundary rows for $K \geq 2$ constraint sets and a goal set, so that $K + 1$ value channels through one triangular solve read the reach-avoid committors of every position; and an interventional re-solve, so that the displacement of the jointly determined state under $\text{do}(a)$ is the consequence of a candidate move and the safest move is the one that reaches the goal before any constraint. At $\gamma = 0$ the read is softmax bitwise (RUN); the record's path-product corner is the resolvent of the sub-diagonal chain (RUN, exact); the committor bed's label is the resolvent read with absorbing rows (RUN, exact); the Neumann tail of a row-stochastic P equals $\gamma^{K+1}/(1 - \gamma)$, a certificate with no sum over the context length (RUN; planted negative fires). No CEQ arm has been trained. The body of this paper is therefore the plan: a directed graph of small, priced, pre-registered milestones on the certified RTX 4060, each with the number that passes it, the number that kills it, the control that would kill it (ChaCAL at the same γ without boundaries; three softmax skylines at matched parameters), the failure mechanism it is designed against, and its deliverable — ordered cheapest-decisive-first so the author can develop it one evening at a time.

Keywords: attention mechanisms · resolvent · absorbing Markov chains · committor · reach-avoid safety · interventions · research programme · pre-registration · Lean 4 · persistent homology · negative results

1. Introduction

1.1 What this document is

It is a plan with its justification attached. Section 2 defines the shape exactly and lists what is proved, run, cited and open about it. Section 3 says why per-row softmax cannot occupy the

ground the shape is built for, citing each theorem only as far as its hypotheses reach. Section 4 compresses fifteen rounds of record into the facts the plan rests on and the mistakes it is designed against. **Section 5 is the programme** — the reason the document exists — written as milestone cards a single developer can pick up in any order the dependency graph allows. Section 6 is the apparatus the cards refer to: beds, labels, floors, metrics, controls, binds, certificates, prices. Section 7 is the occupancy table. Section 8 collects the limits once.

1.2 The north star, in the author’s words

The repository’s immutable north star (`CEQ_V20_R15_CONTRACT.md:42-45`, `READ`) asks for *attention equal to self-attention on its own ground, built from softmax and AdamW, and capable on ground they cannot occupy — predicting the next state toward equilibrium, not the next token*. On 2026-09-03 the author sharpened it:

the main attention mechanism i want is causality for ai to learn consequences for it to predict safest move under multiple constraint something self attention is bad at and any other benefit coming from next transient state phase equilibrium predictor based attention

and, the same day, fixed what this document is for:

the main goal of this research paper was to make the whole future plan so i can slowly develop it

Three capabilities are named, and every milestone in Section 5 serves one of them:

1. **Consequences.** Given a context and an intervention on it — a token rewritten, a move taken — what happens to everything else: the displacement of the whole configuration, not the next symbol.
2. **The safest move under several constraints at once.** Given candidate moves, which one keeps every constraint, and by how much.
3. **The next transient state, the phase, the equilibrium.** Where the context is going: which basin, whether it is at a decision point, which transition.

1.3 The thesis in one paragraph

Take the record’s three-corner operator $W_{\beta,qk,g}$ (`ceq/arm_smprime.py`, `lean/CEQ/V16Domain.lean`, `READ`): softmax at $\beta = 1$, linear attention at $\beta = 0$, the exact path product at $\beta = 0$ with the content term off, machine-checked and measured distinct. Solve one lower-triangular system, $z = (I - \gamma W)^{-1}V$, and read $O = (1 - \gamma)Wz$. At $\gamma = 0$ this is softmax bitwise. At $\gamma \uparrow 1$ over a row-stochastic W with some rows made absorbing — a goal set and K constraint sets, the BOS position declared inside one of them because it is absorbing by construction — it is the fundamental matrix of an absorbing chain: with the value channel $V = \mathbf{1}_{\mathcal{A}_k}$ the read at a transient position is $\mathbb{E}[\gamma^{T_k-1}]$, the hitting-time transform of set k , whose limit is the committor; the $K + 1$ reads sum to one. Over the record’s strictly causal chain at $\gamma = 1$ it is the path product the record already banked: corner 3 was a resolvent all along. A candidate move is an intervention on the context whose consequence is the difference of two solves — one extra solve, rank-one when a single row moves. The parameter γ is a horizon dial from the next token to the equilibrium. That is the shape. Its operator is ChaCAL’s; its base is the record’s; its boundary rows, its interventional channel, its decision rule, its certificates and its proofs are what the sweeps report as held by no one together.

1.4 What self-attention is bad at, as theorems rather than a slogan

Each statement names exactly what its source licenses; two of the four fail at this repository’s own geometry and are cited only for the beds that satisfy them.

- **Per-row independence** (definition). One softmax step computes $O_i = \sum_{j \leq i} P_{ij} V_j$; the read at i never constrains the read at j . A jointly determined configuration is not what a per-row mixture computes in one read.
- **Hops need depth.** k content-dependent hops need depth $\lceil \log_2 k \rceil + 2$ (Sanford–Hsu–Telgarsky 2024, Thm 4.2, [105]); the lower bound is conditional on the 1-vs-2-cycle conjecture (Cor. 4.3). Connectivity needs logarithmic depth ([103]). ChaCAL reaches the same $\lceil \log_2(n+1) \rceil$ bound independently. Depth is not necessary at linear width ([139]), so every comparison matches width too.
- **Composition in one layer.** A single attention layer cannot compose two functions on a large domain ([92], Thm 1, unconditional) — but its inequality is not satisfied at $s = 64$, $d = 16$ ($384 < 544$), so it licenses nothing at the record’s geometry.
- **Optimality on the wrong ground.** One attention layer attains Bayes risk on single-location regression ([77], asymptotic, `erf` gate; [38] for softmax proper). Every bed the campaign ran was that task (`scale/negation_scope.py:286–304`, READ).

Two facts bound what the plan may ever claim. The exact resolvent is a determinant-class computation (Cook 1985; $NL \subseteq DET \subseteq NC^2$), at or above the class where log-precision constant-depth transformers sit ([79]): CEQ does not evade the depth law, it relocates the log-depth into one linear solve, and “depth one” here means one parameter set. And a looped or deeper softmax stack computes the same resolvent by iteration ([126]) with truncation error exactly $\gamma^{K+1}/(1-\gamma)$ after K loops (Prop. 2.x), so every claim is made at matched depth, width and parameters, and “beats softmax” is not a sentence this paper writes.

1.5 How the programme is organised

Section 5’s milestone cards each carry: prerequisites and the cards they can run beside (the record’s dependency law D-1); what to build, as a specification; what to prove; what to measure, with the bed, the seed count and the statistic; the PASS number; the KILL number and what dies with it; the price in GPU-minutes on the certified RTX 4060, labelled fitted, assumed, or run; the MISTAKES.md mechanism it is designed against; the deliverable; and an estimate in evenings. Phases end in a gate the author decides alone with the numbers in hand. The five cheapest decisive cards are marked, and a first-evening list opens the section. Nothing in the plan launches on Kaggle without the author’s explicit yes.

1.6 What this document is not

It is not a result: no CEQ arm has been trained, and the record’s last deciding measurement crossed a one-hop threshold on five seeds of eight, straddled it on the interval, and that threshold was later shown not to be an information floor. It is not a novelty claim for the resolvent, the committor, the successor representation, the triangular solve, the attention sink or the discount factor — each has an owner, cited before it is named. Every number in it names what it was compared against and which file or command produced it; struck constants of the record do not appear; sources not reached at the primary text are marked [U] and used as pointers, never as theorems.

1.7 Contributions

1. The shape, exactly defined, with its containment of softmax and of the record’s path-product corner machine-checked and run (§2).
 2. Boundary rows as the capability mechanism, with the sum-to-one identity that forces the goal set and the BOS convention to exist, and the honest statement that absorption redirects mass and never subtracts (§2).
 3. The interventional channel with the order-of-operations proposition in both halves (§2).
 4. An exact certificate with no sum over the context length, resolving the record’s open impossibility (§2).
 5. The programme: a dependency graph of small, priced, pre-registered milestones with predictions and counters, three fellow approximators in every table, the cheapest killer beside every claim, and the evenings count derived from the critical path (§5).
 6. The apparatus the programme measures with, designed against sixty-six named failure mechanisms (§6), and eight corrections to the record found while writing (§4, §7).
-

2. The shape

JUPITER (MYCROFT), judge of the theory, 2026-09-03, HEAD 207e7b9. The paper’s §2 as final prose. Inputs: the brief, the notes, both corrections, the six `sections/`, the eight `sweep/`, `references.bib` (canonical keys only), the three designs and all six refutations (`refute_{theory,instrument,falsify}_{occvac,math}.md`; none is missing). Rule: every KEEP adopted, every REPAIR applied in the refuter’s wording where right, every KILL deleted unless answered with a number; where refuters disagree the ruling carries its RUN. Evidence classes: `RUN[J]` — one numpy float64 draw by this judge (seed 0, $s = 32$, $\gamma = 0.6$, causal softmax P with the diagonal, sink $\{0\}$, goal $\{5\}$, constraints $\{9, 10\}$, $\{15\}$; computed and printed only); `RUN[coord]/RUN[M]/RUN[I]/RUN[F]` — the coordinator’s, MARS’s, SATURN’s, VENUS’s runs as their files record them; `READ path:line`; `CITED [V]/[U]` by canonical key; `DERIVED` with steps. Every proposition carries its `MISTAKES.md` mechanism and a status in `{LEAN, CLASSICAL, DERIVED, RUN, OPEN}`. Nothing here is a result; the paper’s body is the plan (`THESIS_CORRECTIONS_2.md §0`) and this section is its load-bearing wall.

2.0 Occupancy — cited before the shape is named

The read operator on a causal row-stochastic softmax matrix, solved as a triangular system and reverting to standard attention at $\gamma = 0$, is **ChaCAL** ([42], EMNLP 2024, Eq. 5, Eq. 7, Thm 1; `CITED [V]`; 0 files in the tree cite it — a V-7-on-search finding for §3). One detail decides two controls below: ChaCAL removes the diagonal inside the inverse (`refute_falsify_math.md §3.15, [V-fetched]`), so its resolvent is a finite regime- N sum and its read is **sub-stochastic** — `RUN[J]` row sums 0.400 ... 0.765 against 1.000 diagonal-kept. Its blockwise evaluation is [143] `[V]`. The discounted Neumann closure of a content-adaptive attention with learnable per-head γ , read as the fundamental matrix of an absorbing chain with one leak state, is [99] `[V]` (ReLU Frobenius base, non-causal). Multi-hop PageRank on attention: [124], [43], [140] (lineage [27]); the implicit resolvent step of an attention flow [24]; the exact resolvent layer on a fixed graph [45]; $(I - \gamma P)^{-1}$ as policy evaluation, successor representation, Katz/PageRank index [14], [36], [64], [20]. The three corners: [120] ($\beta = 1$), [63] ($\beta = 0$, $g \equiv 0$), [34] Def. 3.1 (the 1-semiseparable mask); the chunked triangular solve [136] Eq. 10; no finite-state dual for row-softmax [58] `[V]`.

Boundary rows: $N = (I - Q)^{-1}$, $B = NR$ — [65] `[V-cat]` (the bib carries the same book as [66] and [67]; the assembler collapses the three) and [47] Thm 11.6 `[U]`; clamped labelled nodes as

absorbing boundary with a harmonic read *as a learning layer* — [145] [V], [144], [129], [13], [5]; the discrete Dirichlet committor — [82], [39], [40], [37]; one absorbing row inside attention (`do()` with the token fixed and forbidden from attending) — [61] [V]; the attention *sink*, a column device distinct from a row condition — [130], [50], [98]. Decision: reach-avoid probability — [114], [1], [62]; several target sets in one linear system with a verdict — [8]; Chebyshev [119], lexicographic [132], max–min [89], CMDP [2], argmin over evaluated values [14], a move from a linear first-exit solve [115] / [116], a transformer choosing the lowest predicted violation probability [60] (NEAR-MISS), safety filter [57] / [17], the discount that buys a contraction [44] / [56]. Intervention: row surgery on a linear SCM [91], [112]; consequence as equilibrium displacement [86], [16], first-order [18] §7.3; rank-one re-solve [111], [51]; fundamental-matrix perturbation [109]; barrier re-solve on a fixed chain [94] Eq. 5; the SCM as the fixed point of a causally masked transformer with `do()` by clamp-and-re-solve [106] [V]; order of intervention and equilibration [35]; oracle-scored effect with a no-change floor [118], [74]. Neumann tail: textbook ([83], [55]). Segmentation mechanism: [75], [134], [59]; converse [138] Thm 1. Topology: thresholded attention barcodes [71], [72]; intervention effects on persistence [68]; directed stability [117], [29]; Mapper [113] (records [V], landing [U]), [21]; cluster covers [28], [100], [69], [141]; the persistence-derived CSR schedule in an attention kernel — the author’s own [110] [V].

The delta, stated narrowly. No fetched source combines the following in one attention operator, and each part is owned above (eight sweeps, queries recorded; absence is bounded by those queries, V-7): (e) $K \geq 2$ absorbing constraint sets, a goal set and a declared sink as identity rows of a *content-dependent, causal, row-stochastic* read; (g) the interventional re-solve with the displacement as a jointly scored vector label at two stated prices; (h) the committor vector and the safest-move rules over candidate moves placed in the context; (i) a printed certificate with a planted negative and the $1/(1 - \gamma)$ mask amplification; (j) a learnable discount in a causal LM read pinned by a boundary-corrected LR test; (k) machine-checked containment and the zero-gate iff; (l) restricted to the F0 endpoint; (m) the cover-to-schedule path. The name is the repository’s, **Consequence–Equilibrium Attention (CEQ)**; its operator is ChaCAL’s read with the diagonal kept in the inverse, on the record’s three-corner base.

2.1 Definitions

Definition 1 (base family). For s positions with logits $qk_{ij} = q_i \cdot k_j$, magnitudes $m_k \in [0, 1]$, phases θ_k , and a switch $\beta \in \{0, 1\}$,

$$W_{\beta, qk, m, \theta}[i, j] = \mathbb{K}[j \leq i] \text{pathProd}(m, \theta)_{ij} e^{qk_{ij}} / Z_i^\beta, \quad \text{pathProd}(m, \theta)_{ij} = \prod_{k=j+1}^i m_k e^{i\theta_k},$$

READ `lean/CEQ/V16Domain.lean:92–97`, `ceq/arm_smprime.py:144`. The exponential prefix-scan form of `V16Domain.lean:366–378` (Hop) is its restriction to $m > 0$ (`prefix_logit_mask_restated` clause 5, :221) and carries no zero (`no_prefix_scan_represents_a_zero_gate`, :165). Corners: $\beta = 1, m \equiv 1$ softmax; $\beta = 0, m \equiv 1$ linear attention; $\beta = 0$, QK off, the path product (`three_corners_containment` :433, `corners_are_distinct` :445). Ruling: the base carries the multiplicative gate (`refute_theory_math` item 1, `refute_theory_occvac` §1.2 row 9) so that Proposition 6 has a subject. Mechanism: V-25, P-7.

Definition 2 (two regimes, one solve). Regime S: $\beta = 1$, $W = P$ row-stochastic with $P_{ii} > 0$ for $i \geq 1$ and $P_{00} = 1$; $\gamma \in [0, 1)$; $(I - \gamma P)^{-1} = \sum_{t \geq 0} (\gamma P)^t$ by $\|\gamma P\|_\infty = \gamma < 1$. Regime N: $W = A$ strictly causal ($A_{ij} = 0$ for $j \geq i$), any γ , $A^s = 0$ (`Nilpotent.pow_card_eq_zero` :77, `occupancy_is_exact_inverse` :96). The boundary is a theorem (`one_not_nilpotent` :105), and an absorbing row is not `StrictlyLower`: regime N admits no absorbing row, and at $\gamma = 1$ one makes $I - A'$ singular (`RUN[M] det(I - A') = 0.0`). The north star lives in regime S, corner 3 in regime N; they do not meet on a boundary row. Mechanism: V-25, C8.

Definition 3 (boundary rows). Disjoint position sets: the sink $\mathcal{A}_{\text{sink}} \ni 0$, the goal $\mathcal{A}_0$, the constraints $\mathcal{A}_1, \dots, \mathcal{A}_K$; $\mathcal{A} = \bigsqcup \mathcal{A}_\bullet$, $T = [s] \setminus \mathcal{A}$, the query position in T . $P_{a,\cdot} = e_a$ for $a \in \mathcal{A}$; canonical form $P = \begin{pmatrix} Q & R \\ 0 & I \end{pmatrix}$, $Q = P_{TT}$, $R_k = P_{T\mathcal{A}_k}$. Two construction facts (THESIS_CORRECTIONS_2.md §1, re-run RUN[J]): **F1** row 0 of any causal softmax at $\beta = 1$ is e_0 , so position 0 is absorbing whether or not declared; undeclared, $\det(I - Q) = 0.0$. **F2** walks descend, so a boundary set after the query has committor exactly 0.0; every set precedes the query. Ruling on BOS (refute_theory_math C8(b), refute_falsify_math §3.11 against design_theory §1.3): BOS is its **own sink set with value 0 on every indicator channel**, not a goal member — otherwise “reach the goal” collapses to “descend to 0 without a constraint” and the two rules of Definition 7 become one by construction (V-3, C4); the sink share (“fell off the prompt”) is printed beside every label. The value-zero column sink of V15Fork.Asink (V15Fork.lean:67-70) and the sink row are two distinct conditions that coexist on any causal softmax (sweep_resolvent.md §2.14; P-7). Mechanism: V-25, V-8, V-12, D-3.

Definition 4 (state, mixing matrix, read).

$$z(\gamma) = (I - \gamma W)^{-1} V, \quad \Pi_\gamma = (1 - \gamma) W (I - \gamma W)^{-1}, \quad O(\gamma) = \Pi_\gamma V.$$

The read carries the $(1 - \gamma)$ factor. For any row-stochastic W — boundary rows or not — the bare $W(I - \gamma W)^{-1}$ has every row sum $1/(1 - \gamma)$ (RUN[J], RUN[M]: 2.5 at $\gamma = 0.6$ with and without the sets); the factor is a property of the class, not of the boundary rows. Mechanism: V-17, V-23.

Definition 5 (horizon dial). $\gamma \in [0, 1)$ is one learnable scalar per head: the discount of [14], learnable as in [99], the $\gamma \uparrow 1$ limit of [65]. At $\gamma = 1$ with any absorbing row $I - P$ is singular (RUN[M]); the endpoint is a limit (Proposition 9).

Definition 6 (interventional channel, two objects). A candidate move a is an intervention. (i) On the *oracle’s* latent chain: a row clamp $P'_{v_a,\cdot} \leftarrow e_{u_a}$, rank one, $u^\top \mathbb{K} = 0$. (ii) On the *arm’s* operator: a token rewrite at position i , which changes q_i and k_i , hence every softmax row $j \geq i$ through qk_{ji} and Z_j — RUN[J] at $i = 12$, $s = 32$: $\text{rank}(\Delta P) = 20 = s - i$. Two objects, two prices (Proposition 7); design_theory §1.6 (“supported on the intervened rows”) is KILLED for (ii), kept for (i). Mechanism: M-8, P-8.

Definition 7 (labels and rules). Per move, the reach-avoid vector $(q^{(\text{sink})}, q^{(0)}, q^{(1)}, \dots, q^{(K)})(a)$ with $\sum = 1$ on T (Proposition 8). Two rules as columns of one bed (V-1): $a^\star = \arg \max_a q^{(0)}(\text{do } a)$ (reach the goal before any constraint) and the Chebyshev $a^\dagger = \arg \min_a \max_{k \geq 1} q^{(k)}(\text{do } a)$; lexicographic as a third column; the disagreement fraction printed. Two heads: (H-q) the committor vector by the **exact triangular solve at $\gamma = 1$ on the transient block**, $\hat{q} = (I - \hat{Q})^{-1} \hat{R}_k \mathbb{K}$; (H-z) the state/displacement channel at the trained $\hat{\gamma}$. Ruling (refute_instrument_occvac FATAL-2, refute_falsify_math item 16): an undiscounted label and a discounted read at $\hat{\gamma} < 1$ pull γ in opposite directions — q to 1% at $\tau = 8$ needs $\hat{\gamma} \geq 0.99857$, $1/(1 - \hat{\gamma}) \approx 697$ (RUN[F]) — so the committor head does not carry γ ; the discounted read $E_i[\gamma^{\tau-1} \mathbb{K}_k]$ is a separate registered label at a bed constant γ_{env} ([44]) for the horizon-dial bet only. Mechanism: V-17, D-2, D-3, M-20.

2.2 Propositions

Proposition 1 (parity at $\gamma = 0$, with a rejection region). For every causal W and V , $O(0) = WV$; in IEEE-754 forward substitution at $\gamma = 0$, $z = V$ bitwise and the read is the same matmul as attention. *Proof.* $(I - 0 \cdot W) = I$; $0 \cdot x = 0$, $v - 0 = v$, $v/1 = v$ exact for finite entries. $\square$ Carrier V16Domain.corner_softmax :394; target gamma_zero_is_softmax [M]. *Evidence.* RUN[coord] torch.equal(O(0), PV)=True at $s = 64, d = 16$; rejection region $\max |O(0.5) - O(0)| = 2.3002850040264393$; the record’s own softmax corner is 1.110223×10^{-16} off ceq/lm.py on 19/64 entries (READ V16_ARM_SMPRIME.md:266-293), so “bitwise” is against

the lane's own routine. *Battery* (`refute_theory_occvac §2.1`): the plant “non-causal W ” is deleted (it passes at $\gamma = 0$, `RUN[M] array_equal True`); kept: $\gamma \neq 0$ and $\beta = 0$ at the softmax corner (`max |gap| > 0.5`, `READ V16_ARM_SMPRIME.md:336-339`); declared-empty row: “ChaCAL at the same γ ” — a reproduction of ChaCAL's Eq. 5 at $\gamma = 0$, not a contribution. Mechanism: V-24, V-3. Status: **RUN + LEAN (corner)**.

Proposition 2 (corner 3 is a resolvent, regime N). For A strictly lower bidiagonal with $A_{i,i-1} = a_i$: $[(I - A)^{-1}]_{ij} = \prod_{k=j+1}^i a_k$ for $j \leq i$, 0 above. *Proof.* One walk per (i, j) of length $i - j$; $A^s = 0$; $\sum_{m < s} A^m = (I - A)^{-1}$ (`occupancy_is_exact_inverse`), entry (i, j) receiving the $m = i - j$ term. $\square$ Consequence: BED-M's label $y_{s-1} = ((I - A)^{-1}b)_{s-1}$ (`READ scale/negation_scope.py:286-304`) is $z(1)_{s-1}$, so any arm containing corner 3 reproduces BED-M's label as its own forward — BED-M is contained, not won. *Evidence.* `RUN[coord]` `max |G - (I - A)^{-1}| = 0.0`, last row vs `equilibrium_oracle` $6.217248937900877 \times 10^{-15}$, $A^{64} = 0$ exactly. Mechanism: D-2. Status: **LEAN** (two halves) + target `pathprod_is_chain_resolvent [S]`.

Proposition 3 (the committor is the resolvent read with absorbing rows — classical). P row-stochastic including its absorbing rows, $\rho(Q) < 1$, $V = \mathbb{K}_{\mathcal{A}_k}$; for $i \in T$,

$$(1 - \gamma)z(\gamma)_i = E_i[\gamma^{\tau_k} \mathbb{K}_k], \quad O(\gamma)_i = E_i[\gamma^{\tau_k-1} \mathbb{K}_k], \quad \lim_{\gamma \uparrow 1} O(\gamma)_i = q_i^{(k)} = [(I - Q)^{-1} R_k \mathbb{K}_k]_i,$$

and, when the declared sets exhaust the absorbing states, $\sum_{\bullet} q^{(\bullet)} = \mathbb{K}$ on T ; at $i \in \mathcal{A}_k$ the read is 1, not γ^{-1} . *Proof.* Exchange sums in $z_i = \sum_t \gamma^t P_i(\tau_k \leq t, k)$; first-step analysis; monotone convergence; $\sum_k R_k \mathbb{K} = \mathbb{K} - Q \mathbb{K}$. $\square$ Every clause is Kemeny–Snell; the *read* form on an attention P is the composition. Discounted reads sum below one: `RUN[I]` state-form `[0.136, 0.340]`, read-form `[0.227, 0.567]` at $\gamma = 0.6$, ratio exactly γ ; the delay share printed beside every safest-move reading is $1 - \sum_k E[\gamma^{\tau-1} \mathbb{K}_k]$. *Evidence.* `RUN[J]/RUN[M]`: read vs $E[\gamma^{\tau-1}]$ to 8.33×10^{-17} ; BED-1's real sets $A = [0], B = [1], |T| = 9$: resolvent read vs `bed["q"]` 0.0, residual 1.04×10^{-17} , $\rho(Q) = 0.9408612510154677$; Kirchhoff $< 10^{-10}$ (`READ MATHEMATICS.md:542-600`). *Ruling on the 3.8×10^{-7} line.* `refute_theory_math` item 8 derived a universal gap $\geq 1 - \gamma$; the derivation dropped the indicator — $q_i - E_i[\gamma^{\tau} \mathbb{K}_k] = E_i[(1 - \gamma^{\tau}) \mathbb{K}_k] \geq (1 - \gamma)q_i$, small where q_i is. `RUN[J]` at $\gamma = 1 - 10^{-6}$ on channel $\{15\}$: $\max_T = 1.42 \times 10^{-7}$. The line stands with $(1 - \gamma)q_i$ printed beside it. *Read-side plant* (ADD): declare $\mathcal{A}_k$ on the wrong set $\Rightarrow q$ moves by $O(1)$. Mechanism: D-2, V-12, V-25. Status: **CLASSICAL, RUN on the arm's P** ; targets `committor_is_resolvent_read` (a) `[M]`, (b-causal) `[M]` via Proposition 10, (b-chain) `[S]`, `hitting_time_transform [D]`.

Proposition 4 (Neumann tail: attained on the class, and what a mask does). (i) $P \geq 0$ row-stochastic (identity rows included), $\gamma \in [0, 1)$, $K \geq 0$:

$$\left\| (I - \gamma P)^{-1} - \sum_{k \leq K} (\gamma P)^k \right\|_{\infty} = \frac{\gamma^{K+1}}{1 - \gamma} \quad (\text{equality});$$

for $\|P\|_{\infty} \leq 1$ the same quantity is $\leq$ the right side; for $\|P\|_{\infty} > 1$ the bound *can* fail (a nilpotent A with $\|A\|_{\infty} = 2$ satisfies it at $K = s - 1$, `RUN[M]` 1.98×10^{-10} vs 3.68×10^{-5}). *Proof.* The tail is entrywise non-negative with row sums $\sum_{k > K} \gamma^k$; the ∞ -norm of a non-negative matrix is its largest row sum ([83], [55]). $\square$ On the vector: $\leq \gamma^{K+1} \|V\|_{\infty} / (1 - \gamma)$, never the bare δ (V-17). `err = bound` is a declared V-3 identity; the bind is the planted non-stochastic P — quote the convergent plant (rows 1.5, $\gamma = 0.6$, $K = 2$: 7.29 vs 0.54, `RUN[M]`); the coordinator's $\gamma = 0.7$ plant (119.37 vs 1.143) is a divergent series ($\gamma \cdot 1.5 = 1.05$) and is labelled so. (ii) **Mask amplification** (ADD, `refute_instrument_math` row 8): for a sparsified P_m with dropped row mass $\varepsilon = \|P - P_m\|_{\infty}$, $\|O_{\text{full}} - O_{\text{mask}}\|_{\infty} \leq \varepsilon \|V\|_{\infty} / (1 - \gamma)$, attained up to a constant: `RUN[J]` $s = 3$, $\gamma = 0.9$, $\varepsilon = 0.1$: 0.5263157894736843 against the naive 0.1. Every F1 union bound carries $1/(1 - \gamma)$; the exact solve on an F1 mask is refused (L-CERT). *Census.* `rowsum( $\hat{P}$ )` on every trained cell; at $\hat{\beta} \neq 1$ rows sum 1.31 ... 10.29 (`READ V16_ARM_SMPRIME.md:28-32`) and no certificate is

printed (V-25). The $\approx 7.6\text{e-}05$ vector bound of `sec_cost.md §4.x.3` carries [ASSUMED $\|v\|_\infty \approx 5$]. Mechanism: V-3, V-10, V-17, V-24, L-CERT. Status: **DERIVED (textbook) + RUN**; targets `neumann_truncation_bound` [S], `neumann_tail_attained` [S], `mask_amplification` [S].

Proposition 5 (one triangular solve; Π_γ is a mixture; the support does not move).

(a) For causal W , $I - \gamma W$ is lower-triangular with diagonal $1 - \gamma W_{ii} \in [1 - \gamma, 1)$ in regime S (row 0 and every absorbing row attain $1 - \gamma$; `RUN[M]` min 0.4 at $\gamma = 0.6$) and $= 1$ in regime N; z is a forward substitution costing $+s^2d/2$ MACs and depth s **over** the s^2d causal MACs of the softmax head it contains at $\gamma = 0$ (+50%; +25% if $QK^\top$ is counted dense) — an increment, never a stand-alone figure (`sec_cost.md §4.x.1`, M-8). (b) In regime S, $\Pi_\gamma = (1 - \gamma) \sum_t \gamma^t P^{t+1}$ is row-stochastic and non-negative, with or without absorbing rows (`RUN[F]` row sums 1.0000000000000000), so $O_i \in \text{hull}\{V_j : j \leq i\}$ for every γ : the record’s Q2/W3 hull bound applies verbatim, and labels outside the hull need the value rescale of `V15Fork.Asink_computes_chain` or regime N; committors live in $[0, 1]$. (c) On a dense causal softmax, $\text{supp } \Pi_\gamma = \text{supp } P$ with or without boundary rows (`RUN[J]` `True`; lower-triangle minimum on transient rows $2.1 \times 10^{-3} > 0$): the resolvent adds **no support** in regime S; absorbing rows change *weights* downstream (mass redirected to a); the “transitive closure” sentence has content only under F0 zeros. The claim “identity rows change the support” (`design_falsify §3.1`) is **KILLED** (`refute_falsify_math` item 13, `refute_instrument_math` row 5). *Evidence.* `RUN[coord]` `solve_triangular` vs dense inverse $1.7763568394002505 \times 10^{-15}$; `RUN` (NEPTUNE, certified RTX 4060, float32, $n = 2048, s = 64, d = 16$; producer owed to `scripts/k_cert.py`): `solve+Pz fwd+bwd` 2.514 ms vs `PV` 1.473 ms vs one Neumann hop 3.001 ms — truncation never wins in MACs. Kernel boundary: no finite-state dual for softmax P ([58]); the exact path is the full $s \times s$ solve or Zhao et al.’s $\tilde{O}(n^{4/3}d)$. Mechanism: M-8, M-3, V-17, P-7. Status: **DERIVED + RUN**; targets `lower_triangular_isUnit`, `diag_one_sub_smul_pos` [M], `resolvent_is_triangular_solve`, `mixing_matrix_rowStochastic` [S].

Proposition 6 (segmentation is exact, and a zero gate is a boundary condition).

(i) In the Definition-1 base, $m_c = 0$ gives $W_{ij} = 0$ for all $j < c \leq i$ (`pathProd_eq_zero_iff :129`, no hypothesis); then $I - \gamma W$ is block lower-triangular with a zero sub-block and inverts blockwise; an identity row respects every cut, so absorbing rows are consistent with the block structure (`refute_instrument_math` row 15). In regime S an exact zero is a $-\infty$ logit — F0 by mask, not by gate — and the iff has no producing theorem on the softmax corner (`no_prefix_scan_represents_a_zero_gate`); `segmentation_blockdiag` [M] is stated for `StrictlyLower` and `resolvent_fromBlocks` [M] is the half that transfers. (ii) **ADD** (`refute_theory_math` item 11, `RUN[J]`): after a cut at c the row c has window $\{c\}$ and renormalises to $P_{cc} = 1.0$ — a new undeclared absorbing state; with the sets of Definition 3 and $c \notin \mathcal{A}$, $\rho(Q) = 1.0$ and the committor solve is singular; declaring c absorbing severs every boundary set before the cut from every $i \geq c$ ($q = 0.0$, `RUN[M]`); segmentation and the reach-avoid read coexist only if every segment carries its own sink, goal and constraint sets, with the cut position a registered dial. (iii) The dividend is stated in pair-count units, $D = s(s+1)/\sum_m L_m(L_m+1)$: BED-M reads $31.04\times$ from the live-pair fraction 0.0322 (`READ V16_ARM_SMPRIME.md:518-522`; the source line’s pair counts disagree by $2\times$, flagged); (C5)’s $s^2/\sum L_m^2$ would read 58.5 (V-17); a corpus property, $1\times$ on the softmax corner. Mechanism: L-CERT, V-25, D-3, V-22, V-24 (the 10^{-300} plant reads F0 by underflow after ≥ 80 gates at 0.5; the safe plant is a -30 logit, $e^{-30} = 9.36 \times 10^{-14}$, `refute_instrument_math` row 14). Status: **LEAN (annihilation) + DERIVED + RUN**; targets `cut_makes_segment_head_absorbing` [M], `masked_softmax_blockdiag` [S].

Proposition 7 (consequence as a displacement field).

(i) For $(P, V) \rightarrow (P', V')$, $\Delta z = (I - \gamma P')^{-1}(\Delta V + \gamma \Delta P z)$ exactly — one extra solve (four lines; [18] §7.3 is its linearisation). (ii) For lower-triangular P' with $P'_{<i,\cdot} = P_{<i,\cdot}$, $V'_{<i} = V_{<i}$: $z'_{<i} = z_{<i}$ and the suffix re-solve equals the full re-solve — by *triangularity*, not nilpotency (`RUN[F]` 5.6×10^{-17} on the non-nilpotent softmax corner); consequences propagate forward only, $\Delta z_0 = 0$ on every draw. (iii) For a **row clamp** (Definition 6(i)) $P' = P + e_i u^\top$, $u^\top \mathcal{K} = 0$, $M = (I - \gamma P)^{-1}$: $\Delta z = \gamma(Me_i)(u^\top z)/(1 - \gamma u^\top Me_i)$ with denominator $(1 - \gamma p'_{ii})/(1 - \gamma p_{ii}) > 0$ for every $\gamma < 1$, absorbing rows included (`RUN[I]`

1.45×10^{-15} ; RUN[F] 0.511918); the column Me_i costs one substitution ($s^2/2$ MACs), the update $O(sd)$; m candidates cost $m \cdot s^2/2$ against the solve's $s^2d/2$, ratio $m/d = 0.5$ at (8, 16). For a **token rewrite** (Definition 6(ii)) ΔP has rank $s - i$ and the price is the suffix re-solve $(s - i)^2d/2$ — pricing the arm at the oracle's rank-one rate is M-8 literally. (iv) $V \equiv \mathbb{K} \Rightarrow \Delta z \equiv 0$ exactly; in floats $\leq 10^{-15}$ (RUN[F] 1.1×10^{-16} ; a bitwise 0.0 is a one-code-path accident); the region is a Gaussian V (0.1096 at $s = 32$, 1.127 at $s = 64$). (v) A value-only move leaves every $q^{(k)}$ constant across moves — the $m(K + 1)$ -RHS route is legal for Δz only. *Metric* (VENUS): position-matched per-coordinate NRMSE over coordinates $\geq i_{\min}$, field cosine, magnitude ratio, harmonic residual $r(\hat{z}) = \|(I - \gamma P_{\text{env}})\hat{z} - V\|_{\infty}/\|V\|_{\infty}$ with $\sigma_{\min}$ and $\|I - \gamma P_{\text{env}}\|_{\infty}$ printed (its units), the sign column kept (softmax reads 0.807843, READ MATHEMATICS.md:396-406); McNemar on the sign column only, paired t /Wilcoxon on the cosine. Mechanism: V-24, D-5, V-26, M-8, P-8, V-17. Status: **DERIVED + RUN**; targets `displacement_identity` [M], `causal_forward_only` [S], `sherman_morrison_row` [S] (docstring: row clamp only).

Remark (EMC, retired to a remark). For $\gamma < 1$ the map $x \mapsto V + \gamma P'x$ is a contraction with one fixed point, so “settle then intervene” and “intervene then solve” coincide for every row-stochastic P' , cyclic or not (RUN[J] 8.9×10^{-16} on a dense feedback P ; RUN[I] 1.33×10^{-15} ; RUN[F] 4.4×10^{-16}). [35] Thm 1 concerns a reduced model whose equilibrated form hides feedback; an explicit linear fixed point has no such gap. The “planted feedback instance” exhibits loss of forward-only propagation ($\Delta z_{<i} \neq 0$ on a non-causal P) and is filed as a **triangularity** bind under Proposition 7(ii). [85] / [101] (mechanism [U]) remain the external argument that a cached resolvent needs the re-solve under a transition change. Ruling: two refuters KEEP-with-label, two KILL; the KILL stands because the rejection region was an artefact of a wrong algorithm (V-24, V-3).

Proposition 8 (the safest move; degeneracy; the reach-avoid identity with a sink). With Definition 3, $q^{(\text{sink})} + q^{(0)} + \sum_{k \geq 1} q^{(k)} = \mathbb{K}$ on T (Proposition 3; RUN[I] [0.9999999999999993, 1.0]). *Degeneracy lemma.* If the constraint sets exhaust the absorbing states, $\sum_{k \geq 1} q^{(k)} = \mathbb{K}$ and $\max_k q^{(k)} \geq 1/K$ by pigeonhole — “avoid every constraint” is unattainable (V-12 generalised). Its planted negative is “BOS declared inside a constraint set”; the plant “drop $\mathcal{A}_0$ ” cannot run (dropping BOS makes $I - Q$ singular, `refute_instrument_math` row 3), and the coordinator's 0.605 was computed *with* the goal counted. *Floors.* The zero-information floor on the argmin is chance, $1 - \max_a \pi(a^*) = 1 - 1/m$ uniform: 0.75/0.875/0.9375 at $m = 4/8/16$; the weak Fano $1 - \ln 2 / \ln m$ (0.5/0.6667/0.75) is struck as a floor (an arm at error 0.60, $m = 4$, would print above it while worse than chance); tight Fano $(\ln m - \ln 2) / \ln(m - 1)$ (0.6309/0.7124/0.7679, RUN[J]) only where $I(X_{\leq k}; a^*) > 0$ is computed; $m = 2$ has floor 0.5. The zero-hop view carries the move tokens and membership flags, so $I(X_{\leq 0}; a^*) > 0$ and the plug-in is computed (V-25 on a floor). *Census.* Label `sd` over the admitted query region (the prefix before the first constraint is constant); every class frequency of a^* in (0.05, 0.95) (the “argmin-unique” gate was inverted); rule-disagreement fraction > 0 ; discard count; $q^{(\text{sink})}$ apart from $q^{(0)}$. *Bed.* The environment chain must lie inside the arm's operator class or the identity licenses nothing (`refute_theory_occvac` §2.4, `refute_instrument_occvac` FATAL-1): an undirected Rips/`bed_1` chain has `SymmSupport` and no causal $\hat{P}$ equals it (one line; target `lowerTriangular_ne_symmSupport`); a causal chain recoverable from edge tokens makes the label the arm's own resolvent (D-2). Ruling: BED-S's environment is a random DAG in token order with self-loops, adjacency rows as multi-hot node features so $\hat{P} = P_{\text{env}}$ is representable at $d_{\text{model}} \geq n_{\text{nodes}}$; the VOID list is pre-registered — `shape` — `softmax` and `shape` — `skyline` are reproduction-vs-non-reproduction contrasts; creditable are `shape` — `ChaCAL-diag-with-sink` (boundary mechanism), `shape` — `Neumann-K` (exactness) and $\|\hat{P} - P_{\text{env}}\|_{\infty}$ per seed (identification); leak clause C2(a) is dropped (the graph is the bed's input), C2(c) kept; `bed_1/E4'` remain oracle cross-checks (Kirchhoff at $K = 2$). Hazard carried to Limits: [84]. Mechanism: V-8, V-12, V-25, D-2, D-3, M-2, V-10, L-FLOOR. Status: **OPEN** (no cell; discharged by PLAN items S-1 to S-4 of the ledger).

Proposition 9 (the horizon dial). $\gamma = 0$: next step, softmax bitwise (P1); $0 < \gamma < 1$:

the hitting-time transform $E[\gamma^{\tau-1}]$ (P3); $\gamma \uparrow 1$: absorption probabilities, basin membership, the $q = \frac{1}{2}$ isocommittor for $K = 2$ (`ceq/beds/bed_1.py:5-8, :382, [40]`); corner 3 is $\gamma = 1$ on a nilpotent chain (P2). *Pinning*. Ruling 10's $\Lambda = 2[LL(\hat{\gamma}) - LL(\gamma = 0)]$ on held-out, null and held-out set declared for γ (Ruling 10' was stated for β); with $\gamma \in [0, 1)$ the null sits on a boundary and the 95% point is $\frac{1}{2}\chi_0^2 + \frac{1}{2}\chi_1^2 = 2.7055$, not 3.8415 (RUN[J]; the citation is owed, not in `references.bib`); the $|\hat{\gamma}| < 0.05$ rule is deleted (two verdicts for one cell); the mirror kill $\hat{\gamma} > 0.99$ prints $1/(1 - \hat{\gamma})$ beside every δ (READ V15_R1.md:56). Mechanism: V-9, M-2, V-17. Status: **DERIVED (corollary)**; instrument OPEN.

Proposition 10 (invertibility on the causal class is a diagonal read). For the arm's causal P with finite logits, $Q = P_{TT}$ is lower-triangular, $\rho(Q) = \max_{i \in T} P_{ii}$ (RUN[J] 0.6926596893360386 both), and $I - Q$ is invertible **iff** $0 \in \mathcal{A}$ ($\det(I - Q) = 0.0$ undeclared): “absorption a.s.” is the single census line “BOS declared”, and `committor_is_resolvent_read(b)` is [M] via `lower_triangular_isUnit`; the Perron certificate is for the oracle's non-causal chain only (SymmSupport, $\rho(Q) = 0.9409$). The $\|Q\|_\infty$ Neumann bound on the committor is vacuous ($\|Q\|_\infty = 0.958$, bound 19.2 at $K = 4$ vs error 2.5×10^{-3} , RUN[I]); the committor head is the exact solve and claims no truncation certificate. Mechanism: V-10, V-25, L-CERT. Status: **DERIVED + RUN**; `isUnit_one_sub_of_perron` split into (b-causal) [M] and (b-chain) [S].

Proposition 11 (regime boundary). The softmax corner keeps $P_{ii} > 0$, so γP is not nilpotent (`one_not_nilpotent`) and “all s hops by nilpotency” holds in regime N only. The discriminating evidence is $\max_{i \geq 1} (\gamma P)_{ii}^{32} = 6.27 \times 10^{-13}$ beside $\gamma^{32} = 7.96 \times 10^{-8}$ attained at $(0, 0)$ (RUN[M]; the design's number was V-4). The masked-diagonal route is not a matrix identity ($\|(I - \gamma P_m)^{-1} - \sum_{t < s} (\gamma P_m)^t\| = 1.99 \times 10^{-7}$); it is exact on a value channel iff $V_0 = 0$, which the Definition-3 sink satisfies on every indicator channel (target `masked_sink_finite_sum` [S]). Default: keep the diagonal and carry Proposition 4. Mechanism: V-25, P-3, V-4. Status: **LEAN + RUN**.

2.3 What the shape computes

Consequence: Δz (Proposition 7), one extra solve, scored as a field. Committor vector: $\hat{q}$ by the exact solve on $\hat{Q}$ (Propositions 3, 10). Safest move: the two rules of Definition 7 with their disagreement fraction; the threshold form $\max_k \hat{q}^{(k)} + \delta \|V\|_\infty \leq \delta_{\text{thr}}$ bounds the solve, never the model error $\hat{P} \neq P_{\text{env}}$ — no move is “admitted” by it (V-17). Next transient state / phase / equilibrium: the dial (Proposition 9); “phase” is metastable membership ([95], [97], [41]).

2.4 The topological layer (restricted)

- (a) The F0 endpoint is a theorem (Proposition 6) on gated corners or masked P ; on the parity corner the $\varepsilon = 0$ influence graph is one weak component on 100% of draws ($\Pi_{i0} > 0$), so the β_0 barcode row is KILLED on BED-S (three refuters agree) and survives only as the thresholded barcode with a permutation null on BED-M (corner 3, zero gates on 3 of 3 values): NOT MEASURED - needs a digraph β_0 instrument (`beta0_interleaving` consumes point clouds, READ `ceq/certs/topological.py:476-505`); Turner's hypothesis is stated as a predicate on the filtered object first. (b) Isocommittor surfaces: BED-1's $q = \frac{1}{2}$ guard inside the read; NOT MEASURED. (c) Mapper cover $\rightarrow$ tiles $\rightarrow$ schedule: the do-nothing control is both the 0D-salience schedule and Zhao et al.'s blockwise resolvent; the F1 union certificate carries $1/(1 - \gamma)$ (Proposition 4(ii)); the dense control is not runnable at $n = 2048$, $s = 4096$ (≈ 137 GB against 7.996 GiB), so n is stated per s ; `kernels#22` has no backward. Mechanism: V-3, V-25, M-15, M-2, V-9, L-CERT, P-4, P-8.

2.5 The cost law and the kernel path

Per head, MACs, increments over the softmax head contained at $\gamma = 0$ (`sec_cost.md` §4.x.9): serial substitution $+s^2d/2$, depth s , $O(s^2)$ memory (P explicit; no bf16 path for `solve_triangular` [U]); chunked $+s^2d/2 + sCd/2 + sC^2/3$, depth s/C , NOT MEASURED - needs a chunked kernel (pattern [136]); Neumann $+Ks^2d/2$, depth K , never wins in MACs; F0-segmented $+\sum_m L_m^2d/2$; CSR two-stage $+units \cdot B^2d/2$, forward-only. Measured increment $2.514 - 1.473 = 1.041$ ms/step at $n = 2048$ (2.312 at 4096, 4.542 at 8192), a per-op floor under a dispatch gap of $2.0 \times -6.6 \times$ (`READ scale/m3_flops.py:101-121`). One bed-cell pair on the softmax corner ≈ 34 s; seven arms $7 \times 8 \times 1.680 + 4.0 \approx 98$ s ≈ 1.6 GPU-min (`refute_instrument_math` row 28 corrects the design's 4); the corner-3 base (≈ 133 s) has no certificate domain ($\hat{\beta} = 0$ rows sum 1.31–10.29). Determinism: `solve_triangular` bitwise forward and backward over 8 repeats under strict mode, no torch-documented guarantee; `cumsum` raises; the shape on the softmax corner is the first *gated wing* whose training step runs under `use_deterministic_algorithms(True)`. Circuit class: DET ([31], §3.3). Must-fire checks: parity at $\gamma = 0$ with the $\gamma = 0.5$ region; solve vs dense inverse with a non-triangular plant; certificate in vector units with the convergent plant; segmentation `torch.equal` zeros with the -30 plant; CSR fill-in guard; three-outcome determinism; cell price within the law-vs-measured band. Mechanism: M-8, M-3, L-CERT, V-17, P-1, P-8, V-16, V-23.

2.6 Lean targets (build order; every [M] is pending lake build; a missing Mathlib lemma moves it to [S]; nothing is cited as proved before it builds — L-LEAN, P-11)

#	target	grade	carries	refusal shipped	mechanism
1	<code>gamma_zero_is_softmax</code>	Softmax	P1	<code>gamma_half_is_softmax</code> witness	V-24
2	<code>bos_row_is_absorb</code> at $\beta = 1$	Multiping	F1	at $\beta = 0$ row 0 is $e^{qk_{00}}$	V-25
3	<code>later_boundary</code>	Unreachable	F2	—	V-8
4	<code>softmax_corner</code> ($0 < \gamma$, lower-triangular)	Not nilpotent	P11	<code>pow_card_eq_zero</code>	V-25
5	<code>lower_triangular</code> <code>diag_one_sub_smul_pos</code> (interval $[1 - \gamma, 1)$)	[M] isUnit,	P5a, P10	<code>zero_diag_not_Must</code>	M-8
6	<code>resolvent_from_blocks</code> ; <code>segmentation_blockdiag</code> (StrictlyLower)	[M] blocks;	P6(i)	<code>tiny_gate_does_LoGER</code> at a -30 logit	L-6, P-11
7	<code>cut_makes_segment</code> <code>cut_severs_boundary_sets</code>	[M] head absorb (ing),	P6(ii),	—	V-25, D-3
8	<code>displacement_identity</code>	[M] identity	P7(i)	<code>const_value_zero</code>	V-24, D-1
9	<code>lowerTriangular</code>	[M]ne_symmSupp	P8 (D-2)	<code>zero_not_a_count</code> pattern	D-2, example
10	<code>committor_is_resolvent_read</code> (a), (b-causal)	[M] resolvent_read	P3, P10	—	D-2
11	<code>subdiag_pow_enrich</code> ; <code>pathprod_is_chain_resolvent</code>	[S]ry,	P2	<code>forward_map_fix</code> at general n	D-13

#	target	grade	carries	refusal shipped	mechanism
12	neumann_truncation_bound; neumann_tail_attained; mask_amplification ($\varepsilon/(1-\gamma)$)	[6]	P4	nonstochastic_V-break-10 (rows 3/2, $\gamma \cdot 3/2 < 1$); the $s = 3$ witness	V-24, V-10 L-CERT
13	resolvent_is_triangular_solve mixing_matrix_rowStochastic	[6]	P5	bare_read_row_Vum7 $= 1/(1-\gamma)$	V-24, V-10
14	causal_forward_Sonly; suffix_resolve_eq_full (prefix invariance); sherman_morrison_row (row clamp only)	[6]	P7	dense_P_displays_backward; token_rewrite_rank remark	V-24, V-10 L-CERT
15	isUnit_one_sub_Sof_perron (b-chain); reach_avoid_sum_one (sets exhaust the absorbing states)	[6]	P3, P8	unit_needs_rho_V125, V-12 no_goal_forces_max_ge_inv_K	V-24, V-10 L-CERT
16	masked_sink_finite_sum ($V_0 = 0$)	[6]	P11	—	V-25
17	hitting_time_transform; f1_cantelli_union	[6]	P3, L-CERT	—	—

Rows 1–10 are one evening of independent [M] items; 11–16 name their dependencies; 17 is off the critical path. Price: 0 GPU-s each.

2.7 Limits carried to §8

Every RUN[J] number is one numpy float64 draw at $s = 32$, $\gamma = 0.6$, seed 0, CPU, an identity or counterexample check, not a statistic; other planets' runs are quoted at their own geometries. The NEPTUNE timings are one box, float32, a per-op floor, producer owed to `scripts/k_cert.py`. The ChaCAL diagonal convention rests on one HTML fetch (`refute_falsify_math` §3.15) and is re-read against the PDF before typesetting. Theorem numbers inside cited sources inherit the sweeps' [U] marks; [47] is [U]; the boundary-null citation is owed. Proposition 8 has no cell, no realised sd, no measured t^* ; its DAG bed is written, not run. The barcode and Mapper rows have no instrument. Lean grades are statements about statements; nothing was compiled. No code file, no git write, no external fetch by this judge.

3. Why softmax cannot occupy this ground — exactly as the theorems license it

JUPITER (MYCROFT), 2026-09-03, HEAD 207e7b9. Evidence classes as in `judge/sec_shape.md`. Every obstruction sentence below carries a theorem number, the

model class it is stated for, its hypotheses as a predicate on the geometry the record ran ($s = 64$, $d = 16$, $h = 1$, float32 so $p = 32$), and the fraction of that geometry it admits. A “cannot” sentence outside the skyline table of §3.7 is refused. Mechanism throughout: P-10 (a source’s intro cited as its theorem), V-25 (a theorem whose hypothesis no draw satisfies), P-3 / P-11 (a conditional bound stated as settled), P-8 / V-17 (“depth” in two units).

3.1 Obstruction 1 — per-row independence (definitional)

One softmax step is $O_i = f(P_i, V)$: the read at i never depends on the read at j . The shape’s z is the unique fixed point of the γ -contraction $z = V + \gamma Pz$ in $\|\cdot\|_\infty$ (Contraction.rowStochastic_perron :118, weighted_contraction :72), so the reads are mutually constrained. No citation licenses more than this; the nearest published fixed points are per-query ([97]), and the sharpness failure of the row mixture is [121]. What makes it measurable is the harmonic residual $r(\hat{z})$ of Proposition 7 with its two unit factors printed — a $10\times$ ratio at equal marginal error is possible only if one arm’s error lies in the slow mode, which is a statement about what the control mislearns, not “joint consistency” (refute_falsify_occvac §3.12). The claim “joint determination in one read” is UNTESTED (READ MATHEMATICS.md:106-127) and stays a prediction with the counter “ $r_{\text{softmax}}/r_{\text{shape}} \leq 2$ ” (SPLIT in (2, 10)). Mechanism: V-26, V-17, D-7.

3.2 Obstruction 2 — hops need depth (skyline unconditional, lower bounds conditional)

Skyline (unconditional). [105] Thm 4.2: a causally masked transformer computes hop_k at depth $\lceil \log_2 k \rceil + 2$, $m = O(1)$, $H = 1 - 3/5/7$ at $t^* = 2/8/32$. [81] Thm 2: connectivity at $\lceil \log_2 n \rceil$ unrolls of a uniform block. [42] Thm 1: $\lceil \log_2(\text{depth}(G) + 1) \rceil$ layers for entity tracking — the same law with different constants, cited together. *Lower bounds.* [105] Cor. 4.3, $\Omega(\log k)$, **conditional** on the one-vs-two-cycle conjecture, $k = \Theta(N^\xi)$, $mH = O(k^{1-\epsilon})$; [103] Thm 3/19, same conjecture; [25] Thm 1.1, **unconditional** for decoder-only softmax at depth L against L -sequential composition, hypothesis $Hdp \leq n^{2^{-4L}}$: at $L = 1$, $n = 64$, $64^{1/16} = 1.2968$ against $Hdp = 512$ — vacuous at every geometry the record ran; [104] Thm 1, unconditional at $t^* = 2$, hypothesis $hmp = \Omega(n)$: $512 \geq 64$, not violated, silent here. *Width.* [139]: depth is not necessary at linear width; the matched control fixes width as well as depth, and at the record’s geometry $d_{\text{model}} = 16 < s = 64$ the wide skyline is not instantiated at matched parameters — said, not left blank. *Circuit class.* [79] Thm 2: log-precision constant depth $\subseteq$ uniform TC^0 ; connectivity L-complete; consequences conditional on class separations. The clause “linear systems P-complete” is dropped: rational linear systems are in $DET \subseteq NC^2$ ([31]); P-completeness is linear *inequalities* (refute_theory_math item 13). *What the four license together:* an asymptotic, lineage argument that one softmax parameter set cannot compose $t^* \geq 2$ content-dependent hops — and nothing at the record’s geometry. *The missing reduction.* hop_k is pointer chasing in a token-defined graph; the shape’s reachability is in the graph P the layer computes; “every hop_k instance is a committor instance of some P ” is plausible and NOT FOUND (sweep_expressivity.md §3.5). Until written, obstruction 2 is a lineage argument, and the skyline depth $\lceil \log_2 t^* \rceil + 2$ for BED-S is a choice **by analogy** with Thm 4.2, which does not cover BED-S’s task. Mechanism: P-10, V-25, P-3, P-4.

3.3 The DET-class relocation — what “one operator” may claim

The exact committor solves $(I - Q)q = R\mathbb{K}$; the exact z is $(I - \gamma P)^{-1}V$. Inversion and iterated product are DET-complete, $NL \subseteq DET \subseteq NC^2$ ([31]), and a triangular inverse is DET-hard (iterated product embeds in $(I - A)^{-1}$ for block sub-diagonal A). So the resolvent as a circuit sits

at or above the class the conditional bounds place reachability in: **the shape relocates the log-depth from the parameter stack into the linear solve** — $O(s)$ sequential rounds by substitution, or recursive block inversion at $O(\log^2 s)$ circuit depth / $O(\log s)$ matmul rounds with $O(s^3)$ -class work, exact in regime N and a certified Neumann truncation at $K = 2^k - 1$ in regime S (`refute_theory_math` item 14, `refute_instrument_math` row 24; NOT MEASURED - needs a parallel-prefix kernel timing). “Depth-1” in this paper means **one attention parameter set** (P, γ); the circuit-depth reading is disclaimed in this paragraph. Mechanism: P-8, V-17.

3.4 Obstruction 3 — composition

[92] Thm 1 (identifier confirmed; unconditional, communication complexity): a single H -head softmax layer at p bits errs on $f(g(x))$ with probability $\geq R/(3n \log n)$, $R = n \log n - H(d+1)p > 0$. At $n = 64$: $64 \cdot 6 = 384$ against $1 \cdot 17 \cdot 32 = 544$ — not satisfied, vacuous at $s = 64$, $d = 16$, float32; it bites at $n \gtrsim 100$ or $p = 16$. [70] Thm 3.4 removes precision (infinite precision, asymptotic $n^{\Omega(1)}$). “A consequence *is* a composition” ($z' = R(P')V'$ after $P' = S(P, a)$) is a reduction sketch, not a theorem about the committor. Strassen attention is a NEAR-MISS: two hops per layer by arity, not all hops by inversion. Mechanism: V-25, P-10.

3.5 Obstruction 4 — non-negativity, WITHDRAWN as written

For row-stochastic P the mixing matrix Π_γ is non-negative and stays non-negative under absorbing rows (Proposition 5b), so $\partial O_i / \partial V_j \geq 0$ survives the resolvent and `ceq/attention.py`’s min-entry test reads exactly 0 on the shape — by the identity rows and the mask, a theorem instance, not a reading (RUN[F]: 151 exact zeros from the identity rows, lower-triangle min 2.6×10^{-4} without them). Absorption **redirects** mass: a walk that hits $a \in \mathcal{A}_k$ stays at a , so the weight on every earlier token the walk would have reached through a goes to a instead; on a dense causal softmax the *support* of Π_γ does not change (Proposition 5c), only the weights on rows downstream of a . The paper writes “a boundary row captures the mass that would have reached earlier positions”, never “veto”, “subtract” or “negative influence”; the signed programme is not revived (`RESEARCH.md`, `PROGNOSIS.md`); DoFormer’s single fixed row is the NEAR-MISS named beside “boundary rows”. Mechanism: P-7, V-23, V-3.

3.6 D-1, stated with the right predictor and the right census

[77] Cor. 2 (erf predictor, $d \rightarrow \infty$, $L = o(d)$) and [38] Prop. 4.2 (softmax proper): one layer is Bayes-optimal on single-location regression $Y = X_{J_0}^\top v^\star + \xi$. The record ran $L/d = 4.00$ (`READ MATHEMATICS.md:1044–1057`), the wrong direction; and the theorems’ label model admits **0 of 3** record beds — BED-M’s label is a product $a_{s-1} \cdots a_{s-t^\star} b_{s-1-t^\star}$ of $t^\star + 1$ features at $t^\star + 1$ positions, not single-location regression for any $t^\star \geq 1$ (`refute_theory_occvac` §1.2). D-1 is therefore a *mechanism*: no bed in the record leaves the scalar-readout class (`out[:, s-1]`, `READ MISTAKES.md:677–708`), and the SLR theorems are the nearest owned regime, cited as such. No bed in this paper scores a scalar at one position: the labels are $z, \Delta z$, the reach-avoid tensor, the argmin. Single-effect estimation is also softmax’s ground ([142]), which is why the sign column of Proposition 7 is kept and not claimed. Mechanism: D-1, P-10, V-25.

3.7 The skyline table — the only place a “cannot” sentence may stand

Every bed carries every row; a claim is at matched depth and parameters (4,769, recounted per arm because a $[m, K+1]$ head or a vector readout changes it — Ruling 3) against the fellow approximators, and the skylines at unmatched depth/width/decode length are read beside it (`R-SKY`, `READ CEQ_V16_CONTRACT.md:209`).

arm	class	setting to fix before the arena	what it decides
depth-1 softmax, same head	matched	4,769 params, same readout	the per-row control; D-1
ChaCAL-diag (the shape with $\mathcal{A} = \emptyset$, diagonal kept)	matched	same $\hat{\gamma}$	the boundary-row mechanism; identity half of B-E1
ChaCAL-published (diagonal removed in the inverse, sub-stochastic read)	matched	same $\hat{\gamma}$	the published object; a planted negative for bitwise parity at $\gamma > 0$
ChaCAL-diag + sink token	matched	same $\hat{\gamma}$	the C2 kill: boundary rows vs a column sink
InfSA-style Neumann read, no boundaries	matched	$K = 16$; cell ≈ 5.1 s	exactness vs truncation
cached-mixture $O = \hat{P}_{\text{base}}(I - \gamma \hat{P}_{\text{base}})^{-1} V_{\text{int}}$	matched	$\hat{P}$ frozen from the un-intervened context	whether the re-solve is needed ([85])
$\lfloor \log_2 t^* \rfloor + 2$ softmax stack	skyline	depth 3/5/7; ≈ 7.6 s per cell [ASSUMED]	Thm 4.2 by analogy
wide constant-depth stack	skyline	width = n ([139]); NOT MEASURED	depth is not necessary
chain-of-thought decoder	skyline	step count fixed ([80]); NOT MEASURED	the autoregressive route
MuZero-style value head	skyline (consequence beds)	return, not state	a planner need not predict the state

[126] and [131] [V] (abstract level; theorem numbers not read, [U]) show linear-then-softmax transformers implement TD policy evaluation layer by layer, so the deeper stack computes the *same* resolvent by iteration: the looped block $z^{(t+1)} = V + \gamma P z^{(t)}$ run K times is the truncated Neumann sum with matrix-residual ∞ -norm exactly $\gamma^{K+1}/(1 - \gamma)$ and vector error at most $\gamma^{K+1}\|V\|_{\infty}/(1 - \gamma)$ (Proposition 4; [133], [45] power iteration). A nonlinear loop is a DEQ ([7]) and out of scope. “Beats softmax” and “beats native” ([136], [34] on the linear corner) are banned sentences; the shape’s separate claims are exactness with a printed $\delta\|V\|_{\infty}$, one-read consistency as a **cost** statement (one solve, depth s), and the boundary-row mechanism, each stated as a property with its counter-prediction. Mechanism: D-1, R-SKY, D-7, V-17, P-4.

3.8 The honest sentence

One *parameterised* operator (one P , one γ), not one parallel round; its hops are exact and certified (Propositions 2, 4) rather than learned layer by layer; on a bed whose environment chain is inside the arm’s class the contrast against softmax is a reproduction-vs-non-reproduction contrast and is pre-registered VOID as a capability number (Proposition 8); what is creditable is the boundary mechanism against ChaCAL-diag-with-sink, exactness against the Neumann read, and the identification $\|\hat{P} - P_{\text{env}}\|_{\infty}$. Under the record’s calibration (7 of 8 optimistic, READ V16_CALIBRATION.md:96–100) the counter is the point estimate for every row above.

3.9 Limits carried to §8

Every obstruction clause is asymptotic, and every unconditional one is vacuous at $s = 64$, $d = 16$, float32; the inequalities are printed so a reader sees it. The $\text{hop}_k \rightarrow \text{committor}$ reduction is

unwritten. The TD-policy-evaluation skyline claim rests on two abstracts. The skyline depths are by analogy, the skyline prices are assumed or unmeasured, and the two non-depth skylines have no cell. No sentence here is a result.

4. What the record proved, measured and refuted

*The plan's justification, compressed: this section keeps only what §5's cards rest on or are designed against, and every OPEN item names the card that discharges it. Evidence classes — **RUN** executed on this box this session (2026-09-03, HEAD 207e7b9, branch v17k-gate0, torch 2.5.1 float64, Lean leanprover/lean4:v4.7.0); **READ** path:line at that HEAD; **CITED** [V]/[U]; **DERIVED** with steps shown. Struck constants are named by the quantity they were, never by their numeral; the numerals live in STRUCK.md, rendered from tests/loop/test_no_struck_constant_ships.py:47, and a document reprinting one fails that test (READ STRUCK.md:6-10).*

4.1 Proved

The inventory, RUN this session. `grep -cE '^(theorem|lemma) ' lean/CEQ/*.lean` returns **169** declarations across **12** files (Contraction 5, Nilpotent 4, Occupancy 3, OracleSeparation 12, OrbitBound 5, Refcount 10, V15 23, V15Fork 11, V15Kernel 30, V15Phase 26, V15Source 7, V16Domain 33; the import root contributes 0). `lake build` exits **0** on a warm cache; the record calls a warm green *weak evidence* and backs it by force re-elaboration (READ V16_LEAN_DOMAIN.md:394-399), which this session did not repeat, so card **J-L18** rebuilds cold. `grep -n '\bsorry\b' lean/CEQ/*.lean` returns six hits, every one the prose phrase “No sorry” in a file header (V15.lean:41, V15Fork.lean:44, V15Kernel.lean:72, V15Phase.lean:72, V15Source.lean:60, V16Domain.lean:75): **zero** sorry terms. The axiom set [propext, Classical.choice, Quot.sound], with no sorryAx, is **READ** from README.md:101-104 and declaration-by-declaration for V16Domain from V16_LEAN_DOMAIN.md:417-465; Mathlib is pinned at lean/lake-manifest.json:7. Mechanism: **P-11** (MISTAKES.md:1597, a contract citing its own [M]-tagged item as settled) — the paper cites only declarations that build, and §4.3 correction (3) names three that do not.

Table 4.1 — the load-bearing declarations. Statements transcribed from the .lean text, not paraphrased from a report.

declaration	file:line	exact hypotheses	certifies	measured instance (producer)	domain census, BED-M / BED-K
Nilpotent.pow_nilpotent_zero	77	StrictlyLower A (:46: $A_{ij} = 0$ whenever $i \leq j$)	A strictly causal operator is nilpotent at the cardinality — no sign, magnitude or norm hypothesis.	matrix_power(A,64) max-abs 0.0 on make_equilibrium_diagonal(8,64,24) (RUN)	BED-M full (the chain is strictly nilpotent) BED-K empty
Nilpotent.occupancy_escape	96	StrictlyLower A, real matrices	The finite n -term occupancy sum is the two-sided inverse of $1 - A$, exactly.	$(I - A)^{-1}$ equals ceq/arm_smprime.py:144 path_product to max-abs 0.0 (RUN); tests/w2/test_w2_nonnormal.py:124 asserts $\rho(A) < 10^{-12}$ (READ)	as above
Nilpotent.one_nilpotent_bound	105	Weakens StrictlyLower	Weakening to lower-triangular destroys the theorem: $1^k \neq 0$.	— (a refusal; its instance is the softmax corner's diagonal)	binds wherever the shape keeps the diagonal (regime S)
Occupancy.occupancy_escape	52	[Ring R]; none on A	$(1 - A) \sum_{k \leq N} A^k = 1 - A^N$: the truncation residual is exactly A^N , with no convergence hypothesis.	at $K = 16$, $\gamma = 0.7$: residual 0.007754350466241788 against $\gamma^{17}/(1 - \gamma) =$ 0.007754350466241788 (RUN)	vacuous by design; the census belongs to its instance. occupancy_eq_inverse_of_ is its $A^N = 0$ corollary

declaration	file:line	exact hypotheses	certifies	measured instance (producer)	domain census, BED-M / BED-K
Contraction.weighted_rowStochastic_perron	Contraction.weighted_rowStochastic_perron:118	$\text{Domain70, Certificate}(\rho M w_i, \text{and } \gamma P \text{ carries such a certificate at weight } \mathbf{1} \text{ and rate } \gamma \text{ — the sup-norm contraction the shape's certificate rests on.})$ $\text{nonneg, } w_{\text{pos}}, \text{dominates});$ then $\text{RowStochastic } P, 0 \leq \gamma$	$ (\rho M w_i)_i \leq \rho M w_i, \text{ and } \gamma P \text{ carries such a certificate at weight } \mathbf{1} \text{ and rate } \gamma \text{ — the sup-norm contraction the shape's certificate rests on.}$ Every chain coefficient is a product of consecutive gates along the path.	causal softmax rows read 1 to 2.220446×10^{-16} (READ V16_ARM_SMPRIM) 10^{-15} (RUN)	3 of 3 (the hypothesis is on P , not on the corpus); expander_expands_12 (:151) refuses reading
V15.chain_pathV15_problem	V15_problem:74	$m \leq n; a, b$ arbitrary	Every chain coefficient is a product of consecutive gates along the path.	last row against scale/negation_inscope.py:286 equilibrium_oracle , max-abs $6.217248937900877 \times 10^{-15}$ (RUN)	BED-M full; BED-K scope.py:286
V15.prefix_logV15_mask	V15_mask:128	$\forall k, 0 < a_k; j \leq i$	An exponential prefix scan of log-gates reproduces the path product.	—	1 of 3 BED-M values (bedM_overlap_old_two = 1, V16Domain.lean:302 , by decide); BED-K empty universal
V15Fork.Asink_V15Fork	V15Fork.lean:81	none on a	The telescope head's rows sum to 1 unconditionally: the reweighting is a change of units.	—	

declaration	file:line	exact hypotheses	certifies	measured instance (producer)	domain census, BED-M / BED-K
V15Fork.Asink_V15Fork.Asink_chain	V15Fork.Asink_chain:140	$a_k \neq 1$	One causal row-stochastic softmax head with key-side bias $s_j = \log(1 - a_j)$ and a value-zero BOS sink reproduces the chain label exactly.	6.661338147750929 of 3; the 10^{-16} at $s = 8$, excluded 7.549516567451064 at 10^{-15} at $s = 64$ (RUN +1, where the results/v15_r1β is shown) ARM PHASE reads 9.15513359704447535288 10^{-16} , ARM (READ SMPRIME V16_ARM_SMPRIME.md, verdict (h)) 5.919777×10^{-16} (READ workdonenewseal.md:91-96)	2 of 3; the excluded point is the band value +1, where the reads (READ V16_ARM_SMPRIME.md, verdict (h))
V15Fork.no_row_stochastic_chain	V15Fork.no_row_stochastic_chain:145	No row-sum 1; exactness on all b	No row-stochastic head reproduces the chain on unrescaled drives: the label leaves the convex hull of its own drives.	on $b \equiv 1$ the label exceeds 1 (chain_one_gt_one:169)	the two-branch escape it refuses is V-24 in proof form
V16Domain.three_domains_containment, qk	V16Domain.three_domains_containment, qk:430	arbitrary	Softmax ($\beta = 1$), linear attention ($\beta = 0$) and the path product ($\beta = 0$, QK off) are three settings of one operator.	$\beta = 1$ bitwise against softmaxAttn; row sums 1.000000 at $\beta = 1$ against [1.312192, ..., 10.293107] at $\beta = 0$ (READ V16_ARM_SMPRIME.md:28-32)	3 of 3
V16Domain.corners_containment	V16Domain.corners_containment:445	(a decidable witness)	The containment is not decoration: the corners differ at (1, 0).	$ c_1 - c_2 = 4.472918$, $ c_1 - c_3 = 1.144938$, $ c_2 - c_3 = 5.335671$ (READ :333)	3 of 3

declaration	file:line	exact hypotheses	certifies	measured instance (producer)	domain census, BED-M / BED-K
V16Domain.path_product_zero_info	V16Domain.path_product_zero_info:169		A zero gate annihilates the path product exactly, and only a zero gate does.	96.78% of BED-M causal pairs annihilate; annihilation MCC 1.000000 corpus-alone (READ V16_ARM_SMPRIME.md:519-527)	3 of 3 (bedM_overlap_new_two = 3, :304); BED-M proper 2 of 2 (:310)
V16Domain.no_prefix_scan_carries_a_zero	V16Domain.no_prefix_scan_carries_a_zero:165	$[j+1, i], m_k = 0$	No gate exponential prefix scan — none — can carry a zero gate.	the exp_scan nan on 133,120 of 133,120 causal pairs at $n = 128$, $s = 64$ (READ workdone.newseal.md:122)	3 of 3
V15Kernel.first_order_delay_arbitrary, not affine	V15Kernel.first_order_delay_arbitrary, not affine:140		No scalar-state first-order recurrence delays by $d \geq 1$.	the fitted first-order recurrence reaches $R^2 = -0.000170$ on the delay bed, where the same fitting code recovers AR(1) at $R^2 = 1.000000$ (READ V15_N4_BEDK.md:194)	BED-K covers $d \geq 1$; $d = 0$ is a legal cell and <i>necessarily</i> outside, exhibited by V16Domain.delay_zero_is (:339)
V15Kernel.delay_forces_state_injective + linear_first_order_cannot_delay_beyond_state_dim,	V15Kernel.delay_forces_state_injective + linear_first_order_cannot_delay_beyond_state_dim:253	type, no linear A, B, C with $k < e + 1$	Delay forces an injective state dim, and linear delay d needs state dimension at least d .	bracketed in-file by delay_realizable_at_dimension_d (:235): a shift register delays at dimension $d + 1$	BED-K

declaration	file:line	exact hypotheses	certifies	measured instance (producer)	domain census, BED-M / BED-K
V15Kernel.first_order_kernel	V15Kernel.lean:504	power-law	The Grünwald–Letnikov power-law kernel is not geometric, so no affine first-order recurrence reproduces it exactly.	the registered box is $\alpha = H - \frac{1}{2} \in (0, \frac{1}{2})$; the <i>approximation</i> gap is open — the bed fits at $R^2 = 0.604$ ($H = 0.75$) and 0.755 ($H = 0.9$) (READ workdone/newseal.md:189)	BED-K; debt D-APPROX , discharged by no card and carried in Limits
V15Source.source_separation	V15Source.lean:811	first-order difference A ; $r = \text{occ}(A)h$	Inverting the resolvent costs one mat-vec and a subtraction, $O(\text{nnz } A)$, against a forward map that fills in.	propagated-field error 8.882×10^{-16} , two planted sources (READ V15_X35P_SOURCE.md:33); $\text{nnz}(1 - A) = 251$ against $\text{nnz}(W) = 1703$ at $n = 128$, $d = 5$ (READ V15Source.lean:40–42)	an X35' instrument, not a bed
OracleSeparationOracleSeparation	OracleSeparation.lean:166	is a simulation A ; Nonneg Q ; SymmSupport Q ; $0 < Q_{ij}$	An undirected-walk transient block is never a strictly causal operator, at the operator level, before any label is drawn.	BED-1's own chain, $\text{bed_1.build}(T=0.25, \text{jitter}=0.05, \text{seed}=11)$: Nonneg True, SymmSupport True, $\rho(Q) = 0.9408612510154677$, $\max Q^{11} = 0.47739216729378814 \neq 0$ (RUN)	BED-1; D-2 (MISTAKES.md:710) is a theorem

records that dropping $0 < Q_{ij}$ leaves the statement *true* — the zero matrix is non-negative, symmetric-support and nilpotent — a vacuous control caught in proof form, and the template every Lean card follows (**J-L0** to **J-L18**).

4.2 Measured

Table 4.2a — the identity binds. A bind whose rejection region is empty is **V-24** (MISTAKES.md:1658), so every row ships a planted mutilation failing at $O(1)$; the binds run on BED-M’s real support $\{-1, 0, +1\}$ rather than a toy draw, which is **V-25**.

bind	residual	bar	planted negatives (residual)	source
oracle gates give the label, path-product corner	5.919777×10^{-16} , real part exactly 0.000×10^0	10^{-6}	drop_phase 1.934830; drop_magnitude 0.466267; beta_one 1.335288; exp_scan nan on 133,120/133,120	READ V16_ARM_SMPRIME.md:25,195-
label bind, arm_pl telescope head	$6.661338147750939 \times 10^{-6}$ 10^{-16} at $s = 8$; $7.549516567451064 \times 10^{-15}$ at $s = 64$		drop_key_bias 0.9749459015140511; drop_value_rescale 0.9165274652308163; drop_bos_sink 1; half_key_bias 0.48449311856985267 — four for four	RUN #results/v15_r1.jsonl; READ V15_ARM_PL.md:183-200
label bind, arm_phase	$9.155133597044475 \times 10^{-6}$ 10^{-16} at $s = 8$; nan on BED-M’s band ($ a = 1$ gives $\log(1-1) = -\infty$)		the band failure <i>is</i> the planted negative, and is not adjusted away	READ V15_ARM_PHASE.md:36
softmax corner against its own softmaxAttn	bitwise ; against ceq/lm.py::Attention("softmax_x") 1.110223×10^{-16} on 19/64 entries — a fused kernel subtracting the row max, a named mechanism and not a tolerance	bitwise	wrong switch: $\beta = 0$ at the softmax corner reads $\max \text{gap} > 0.5$	READ V16_ARM_SMPRIME.md:27,266-

bind	residual	bar	planted negatives (residual)	source
DAG resolvent against brute-force path sums	2.7755575615628914k1 $\times 10^{-16}$ 10 ⁻¹⁷ (light gates); 3.5527136788005009 $\times$ 10 ⁻¹⁵ absolute, 2.1572667249528595 $\times$ 10 ⁻¹⁶ relative (heavy)		the two routes share no code — dense LU against enumeration of increasing vertex sequences, a V-3 guard	READ V15_ARM_PL.md:153-177
$\gamma = 0$ parity of the shape's read	torch.equal returns True against <i>PV</i>	bitwise	rejection region at $\gamma = 0.5$: max-abs 2.3002850040264393	RUN
Neumann truncation certificate	equality to 10 ⁻¹⁵ at $K \in$ {1, 2, 4, 8, 16}, $\gamma = 0.7$, $s = 64$	attained	a non-stochastic P (rows scaled to 1.5) reads 119.37486584159647 against the bound 1.1433333333333329	RUN
corner 3 is the sub-diagonal resolvent	max-abs 0.0 entrywise; label 6.217248937900877 $\times$ 10 ⁻¹⁵	exact	$A^s = 0$ exactly	RUN
committor is the resolvent read with absorbing rows	max-abs 0.0 on BED-1's real sets ($A = [0]$, $B = [1]$, $ T = 9$); harmonic residual 1.0408340855860843 $\times$ 10 ⁻¹⁷	exact	a must-fire perturbation drives max $ Lq $ to 1.000000 $\times 10^{-6}$, a ratio of 9.608 $\times 10^{10}$	RUN ; READ V15_BED1.md:123-124
triangular solve against the dense inverse	1.7763568394002505k0 $\times 10^{-12}$ 10 ⁻¹⁵ at $\gamma = 0.5$, $s = 64$, $d = 16$		— (the cost claim, not the identity, carries the risk; rule 8 below)	RUN

The two-branch form $\text{softmax}(q)V_1 + \lambda XV_2$ passes parity bitwise for the honest gate, for Gaussian noise, **and for the label itself** — an empty rejection region, filed **V-24** (READ `workdonenewseal.md:124-126`). The plan inherits the battery, not the claim: cards **S-20** to **S-27**, every plant journalled as a FOUND cell with its own kind (**S-28**, against **V-7**, :117).

Table 4.2b — the deciding measurement R1, and its re-take on the certified device. The row that matters is the split, never the mean: a marginal standing in for a joint claim is **V-26** (MISTAKES.md:2190).

quantity	CPU, results/v15_r1.jsonl	CUDA re-take, results/v17k_r4_retake.jsonl
cell	BED-M, e3_t2, $t^* = 2$, $s = 64$, $d = 24$, $d_{\text{model}} = 16$, $n_{\text{train}} = 2048$, 150 steps, 8 threads, $N = 8$	the same, device="cuda", instrument hash 5d41a63d...9a309 on all 24 cells, deterministic_algorithms: True, warn_only: True
$\text{floor}_1 = \sqrt{(t^* - 1)/t^*}$	0.7071067811865476	the same
arm_pl	mean 0.829151, sd 0.253673, CI95 [0.617075, 1.041227], $\hat{h} = 0.625017$, crosses False	mean 0.830200, sd 0.255414, CI95 [0.616669, 1.043731], $\hat{h} = 0.621535$, crosses False
softmax	mean 0.951767, sd 0.011824, CI95 [0.941881, 0.961652]	mean 0.952561, sd 0.011452, CI95 [0.942987, 0.962135]
arm_smprime	no CPU counterpart exists	mean 0.822193, sd 0.250321, $\hat{h} = 0.647997$; 1 of 8 below the floor (seed 2 at 0.20391993939877656)
the bimodal split	crossed 5 of 8 (0.634002 to 0.662021; gate- R^2 0.97 to 0.99; $\hat{a}_{\text{max}}$ 1.10 to 1.51); NO READING 3 of 8 (1.113403, 1.139404, 1.152430; gate- R^2 0.011, 0.627, 0.047; $\hat{a}_{\text{max}}$ 20.31, 49.66, 285.07); nothing between 0.663 and 1.113	identical membership; zero verdict sign flips across 16 paired cells
paired contrast (PL minus softmax)	−0.122616, sd 0.257560, resolution $\Delta = t_{.975,7} \text{sd}/\sqrt{8} = 0.215326$	−0.122360, sd 0.260159, $\Delta = 0.217498$; the device moves the contrast by $+2.56 \times 10^{-4}$, +1.0% 9.522×10^{-3} at arm_pl seed 7 (1.139404 → 1.148927), $4.06 \times$ the thread floor; two other non-learning seeds read 1.505×10^{-4} and 6.329×10^{-5} , so “did not learn” is not the mechanism
worst CPU-to-CUDA delta	—	1.505×10^{-4} and 6.329×10^{-5} , so “did not learn” is not the mechanism
identical-seed floor	—	$\delta_{\text{nrms}} = 0.0$ bitwise on 6 of 6 pairs, results/v17k_r4_floor.jsonl; the seed-to-seed spread is at least 10^4 times larger
cost	1358.3225734233856 s for 16 cells	2.79 GPU-min measured against 2.65 projected (+5.3%); peak 0.9739 GiB, 12.2% of 7.996 GiB

The mean 0.829151 is a value no seed produced. That is why the plan reports every capability number as a split with n^+ , the seed-agreement count, beside it (**V-26**; card **S-61**, the adjudicator specification written before data). The recorded verdict is NOT CROSSED with the interval straddling, and two campaign firsts sit inside the negative: $\hat{h} = 0.625$ against a nine-cell

ceiling of 0.389, and five seeds below a floor no prior arm had crossed at any cell. Deduplicated across the three journals (**RUN**, 43 rows to 40 unique on (**kind**, **seed**)), **13 of 40** banked cells sit below floor₁, 13 read $\hat{h} > 1.0$, and $s = 64$ on 40 of 40.

The three corners, measured distinct. $|c_1 - c_2| = 4.472918$, $|c_1 - c_3| = 1.144938$, $|c_2 - c_3| = 5.335671$ in float64, against **corners_are_distinct**; the switches move the operator by 0.673101 (g) and 3.522037 (QK), and β , g , QK are all **nn.Parameters**. Mechanism **V-2** (:58) — a containment whose corners coincide is decoration, so the distance is measured, not asserted.

Table 4.2c — the device certificate (RTX 4060 Laptop, sm_89). Capability 8.9, 8,585,216,000 B (7.996 GiB), torch 2.5.1+cu121, Python 3.11.9, Windows-10, eight threads, matmul TF32 off, CUBLAS_WORKSPACE_CONFIG=:4096:8 (**RUN** results/k_cert_local.json).

law or constant	value	R^2 / verdict	note
softmax throughput	s/step = $\exp(-12.18515596) n^{0.99625107}$	0.9999975	5 points, $n \in \{2048, \dots, 32768\}$
arm_smprime throughput	s/step = $\exp(-10.01874582) n^{1.00258069}$	0.99999987	3 points; at $n = 16384$ and 32768 the allocator reserves 10.578 and 13.969 GiB from a 7.996 GiB card and pages over PCIe, and admitting them bends the exponent to 1.2341 at $R^2 = 0.976933$. A throughput law fitted through swap is not a throughput law (P-8 , :387)
C_RESIDUAL fp32	module 18, re-solved 17.874	0.996373, CONFIRMED, -0.7%	—
C_OPERATOR fp32	module 3.9, re-solved 3.823	0.996373, CONFIRMED, -2.0%	—
bf16-autocast residual	module 2.2, re-solved 2.383	0.999830, WRONG , $+8.3\%$, optimistic	at the Q3 chunk shape 2.715 against 2.738 GiB ($+0.85\%$), max_batch 45 either way — no decision moves, so it is reported and not repaired
the complex arm (arm_phase)	measured over predicted 1.842; C_OPERATOR 7.50 at 8 B/element, $4.29 \times$ softmax's	—	the module breaks exactly where it is new; C_RESIDUAL NOT IDENTIFIED at $s = 64$

law or constant	value	R^2 / verdict	note
determinism, forward	<code>cumprod</code> has a deterministic CUDA kernel, so <code>arm_smprime</code> 's forward is bitwise: $\max \Delta = 0.0$ over 8 repeats	—	<code>cumsum</code> does not , and it is the whole of <code>arm_pl</code> and <code>arm_phase</code>
determinism, backward	the hole moved and did not close: autograd differentiates <code>cumprod</code> through <code>cumsum</code> , so the strict flag raises; gradients are bitwise (0.0, 8 repeats) only with the flag off	—	READ <code>COSTS.md:149-154</code>
determinism, run to run	$\delta_{\text{nmse}} = 0.0$ on 6 of 6 identical-seed pairs under <code>warn_only=True</code> ; the flag changed no value, 6 of 6 flag-OFF cells reproducing flag-ON bitwise (<code>==</code> , not <code>allclose</code>)	—	Ruling 1
cpu against cuda <code>cumprod</code>	11/64 entries move by at most 5.551115×10^{-17} (sequential against parallel scan); zeros identical on both devices ; on BED-M's own support 0/289 entries move	—	association order does not matter where it would change a verdict
bar re-certification	<code>BAR CALIBRATED</code> on four rungs and both devices; worst δ $8.580024779547557 \times 10^{-8}$ against a 10^{-6} tolerance (8.58%, headroom $5.83\times$ to the 50 % HALT line); <code>oracle</code> exactly 0 everywhere	—	30 doubles bit-identical in IEEE-754 hex against the previous revision

Mechanisms: **M-8** (:544, pricing every arm at one arm's rate) — one law per arm with its R^2 ; **V-22** (:1140, a constant certified under conditions the reading does not reproduce) — the flag regime is journalled on the header; **P-8** — the R^2 gate refuses paged points. What the certificate does **not** contain is an exponent in s : $s = 64$ on 40 of 40 cells, and the per-cell timer is un-synchronised host wall clock whose strongest correlate is run order, $\rho = +0.7029$, $p = 0.0024$,

above the gate correlation +0.5197 it would have to be separated from. That is **D-3** (:727), discharged only by card **S-66** — the *s*-sweep with **synchronize()** and randomised order, priced at 206–537 GPU-s, band only, with the run-order confound *surviving* those edits unless a fifth (blocked order) is made, which no office has priced.

The calibration column. Round 11: **9 checked, 9 adverse** on the filed rows (at least 17 checked, 10 adverse, with the confirmations the table omits). Signs: **7** optimistic, **1** pessimistic (a sizing repair already made, reported as outstanding), **1** unsigned. The one-sided sign test on the eight signed rows gives 7/8, $p = 0.0352$ — direction established at $\alpha = 0.05$ — and the Wilson 95 % interval on the optimism fraction is [0.5291, 0.9776]: it excludes 0.5 and fixes nothing else, so it **licenses an ordering only**. The round’s own headline “six wrong, all six optimistic” gives 5/6, $p = 0.1094$ on the honestly signed rows: right about the direction, wrong about its evidence. The single sub-census with a real denominator reads 11 checked, 4 wrong, 36.4%, Wilson [0.152, 0.646], so 9/9 is a count and not a rate. The discount rule §5 adopts verbatim: **D-CALIB-1** the counter is the point estimate; **-2** a bare prediction is blocked, never discounted; **-3** the licence is a *sign*, never a size, so **no numeric shrink factor is authorised**; **-4** the cheapest refutation of the optimistic half runs first; **-5** the row is appended whether or not it flatters. Mechanisms **D-7** (:2037) and **M-2** (:451); cards **V-0, V-1, V-19**.

Numbers this paper must not use. The registry holds **12** struck constants (READ STRUCK.md:18–33, rendered at aa82df7), named here by quantity.

struck quantity (registry rows)	why it was struck	what the paper writes instead
the M2 decay exponent and its R^2 (1–2)	a floor = 1e-6 artefact; the two candidate repairs disagree (R^2 0.9990 against 0.9662)	delete the claim ; no replacement is published
the “live rows only” K1 slope and both interval endpoints (3–5)	UNVERIFIED — the interval exists in no .py , .json , .jsonl or .txt , and the only live producer emits a different number and exits 1	the as-computed slope -0.4137 , CI $[-0.4579, -0.3704]$, reproducing from scale/arm_a_k1.py
the M2 slope as first reported (6)	contradicted by measurement, and propagated in a [RUN] voice	the shipped operator’s reading, carried in §4.3 as a <i>verdict</i> and not as an exponent
the M5 tail norms at $s = 128$ and $s = 512$ (7–8)	FABRICATED — an 1,800-setting sweep produced neither	measured 0.880500 (hops = 2), 0.882030 (hops = 4) at 128; 1.292741 at 512
the U1/N3 pilot rank correlation and both CI edges (9–11)	NO PRODUCER HAS EVER EXISTED — nine named functions defined in no commit on any ref	nothing; the item is void
the Karcher residual in float64 (12)	asserted [RUN] from a throwaway script whose own output was nan	7.481×10^{-9} , 8.155×10^{-9} , 8.405×10^{-9} at published settings; 7.307×10^{-13} , 7.958×10^{-13} , 8.405×10^{-13} at --tol 1e-15 --steps 400

The named stale traps, with their corrected figures. “No arm crosses floor₁” becomes 6 of 24 at it.2 and **13 of 40** at it.14 (C1, C15); “6 of 34 cells violate the floor” becomes **13 of 40** (**RUN**); “12 of 16 cross” becomes “12 of 16 sit below the floor, and the pooled verdict is **crosses: false**”, the flip turning on seed 9 alone (C13). $p = 6.730 \times 10^{-4}$ was **unpaired** — the honest paired figure on the banked record is $p = 0.012821$, a factor of 19, repaired by measurement at

it.10 to **8 of 9 against 0 of 9**, paired, over three draws (C14). Every GPU-second and cost ratio derived from `scripts/v15_r1.py`'s `secs` is void, the K-cert laws being the replacement (C17). The “5.8× wall-clock gap” and softmax at 71.32 s per 150 steps do not reproduce: the same arm, corpus, thread pin, steps and statistic re-read 12.3–13.6 s on a quiet host, and what survives is E-core placement plus host contention, 41–83% of the gap in log terms with the residual unexplained. The “13.6× CPU-to-GPU decision” becomes 11.94× at the point estimate and 29.0× at the pessimistic end. $\sum_j |W_{ij}| = 1.000000$ as the softmax-corner certificate equals $Z_i^{1-\beta}$ and is blind to g ; write $\sum_j \operatorname{Re} W_{ij}$, at 2.220446×10^{-16} from 1. “A pole precisely on the unit circle” is `torch.clamp(u, 0, 1)` at `ceq/arm_smprime.py:113` — a ceiling, not a converged pole (C9). “The arms return a state distribution” is false: 0 of 40 banked cells journal one. And “beats softmax” on any BED-M cell is banned by **R-SKY**; the paired contrast with its resolution statement is what the paper writes.

4.3 Refuted, by mechanism

Table 4.3a — the dead programmes, each with the number and interval that killed it.

programme	the number	the control / interval	source
the signed strictly-causal path sum (sgate)	sign-flip rate 0.16511, 0.02732, 0.00000, 0.00000 at $s = 8, 32, 128, 512$ — a decay of order $1/s$, killed against a bar of −0.3 by a factor of four under every floor tried	softmax control exactly 0.0000, 0/1024 flips. The fitted exponent is not carried : it rests on two nonzero points, a different geometry reads −0.958 at R^2 0.9990, and the audit’s third reading reconciles with none of its own five rates	READ PROGNOSIS.md:40–63, 357–364
the same, downstream	COGS-gen 0/512 against softmax 15/512, one-sided Fisher $p = 2.75 \times 10^{-5}$; in-distribution 0.7734 against 0.9258	3,652,096 parameters matched on both arms; zero_success_upper_bound (512) = 0.005834	READ RESEARCH.md:178–194
pivot routing	dense slope −1.088 against pivot −1.298: routing makes the decay worse . The paper carries the ordering, not the exponents	at $s = 8$ routing restricted nothing — $A[:, P] A[P, :]$ equals AA to 1.86×10^{-9}	READ PROGNOSIS.md:44–45, 79–85

programme	the number	the control / interval	source
the exclusion confound beneath it	the maximum legal pivot is $s - 2$ while the label $a_{s-1}b_{s-2}$ sits on the hidden position; a +100 perturbation moves softmax by 101.6983 and the pivot path by 3.263746	lifting it: twin_plus 0.938728 against a threshold of 0.871391, which is 0.015610 worse than the twin; the inert version selected $s - 1$ in 0/32 draws (V-9)	READ MATHEMATICS.md:329-347; READ D1.md:385-398
settling the <i>mixture</i> <i>weights</i> over pivots	settled – twin = −0.002959, exact 95 % CI [−0.042903, +0.031557] over all 3125 paired resamples with 126 atoms	settled sd 0.064106 against twin 0.016547, 3.874×; argmax – softmax = −0.118456, CI [−0.134115, −0.102786]	READ MATHEMATICS.md:355-380
the hop-2 term on the softmax operator	filed blind, 4/4 held: 0.976488 (sd 0.004039, CI [0.973932, 0.979036]) against softmax 0.975371, a contrast of +0.001117; $\hat{h} = 0.372$	five independent routes agree : a K sweep; a γ sweep in which no γ beats 0; deflation; a Wiener route at 979× attenuation; and a gated hop at 0.966692, worse than both	READ V13_PREDICTION_HOP2.md:35-46; READ workdonenew.md:266-294
VGPE, path-ordered non-abelian transport as positional encoding	occupied outright by PaTH §2.1, $A_{ij} \propto$ $\exp(k_j^\top (\prod_s H_s) q_i)$, with RoPE recovered as $H_s = R$	the Hankel rank is d regardless of alphabet	CITED [V] [137]; READ PRIOR_ART.md:742-778
the original parity clause, “ $g \equiv 0$ gives bitwise standard attention”	row i sums to $i + 1$; smallest witness $i = 1$, row (1, 1), sum 2	gate_zero_row_sum , gate_zero_not_stochastic , repaired by the telescope at 2.2×10^{-16}	READ workdonenew.md:47-124
“the normalizer obstructs path products” X_{35} hidden-cause inference	false: it dissolves uniquely at $\gamma_j = 1/(1 - a_j)$ occupied since 1993 : Basseville and Nikiforov §7.2.4 in closed form, equation for equation, with both must-fires discharged by the 1993 equations on the first attempt	the additive-logit repair computes y_i/R_i to 4.44×10^{-16} the learned-model form is [U] and is the record’s own analogy; correction (5) below	READ workdonenew.md:118-124 CITED [U] [10]; READ README.md:214-219

programme	the number	the control / interval	source
closed-magnitude phase gates	occupied: S4D's ReLU variant reads $ \bar{A} = 1.0$ exactly on 32.93% of a standard sample, in 2022	$m = 0$ is attainable too — modReLU, at 14.70%	CITED [U] [48]; READ V15_X36_PRIOR_ART.md:396
R1, the deciding measurement	the CI [0.617075, 1.041227] straddles floor ₁ ; 5 of 8 cross, 3 of 8 are NO READING	$1/(1 - \hat{a}_{\max})$ is undefined at all eight seeds; the capability clause was not earned	READ V15_R1.md:51-56
floor ₁ read as an information floor	violated by 13 of 40 banked cells (12 arm _{p1} , 1 arm _{smprime}) — a lower bound is never violated, and this one is	it is the locus of $\hat{h} = 1$: a one-hop capability threshold . The real BED-M floor is the exact oracle at 0.0, and no annex theorem predicts either wing's distance to it (best 0.203920, modally about 0.85)	READ V20_R15_THEORY_TABLE.md:69-91; C11, C15
Q6, the state metric	F4, the domain is empty : both wings return $[n]$ scalars and 0 of 40 cells journal a distribution	marginal W_1 reads 0.0 on the oracle's own values permuted while NRMSE reads 1.421901; metric and bed rank two predictors in opposite orders by $14.465\times$	READ V20_R15_THEORY_TABLE.md:211-223
the M6-against-M2 tension (IMPOSSIBLE.md I1)	$w(A) = 1.499315 > 1$, so $2w^h$ grows from 2.999 to 10.11 while $\ A^h\ $ falls from 2.838 to 0.210 over $h = 1..4$; every eigenvalue guard vacuous, the spectrum being $\{0\}$	resolved by Proposition 4 , not repaired: with a row-stochastic P the normalizer pays for a denominator-free certificate carrying no sum over s , and the sign capability bought by dropping the normalizer is not recovered — it is replaced by absorbing rows	READ IMPOSSIBLE.md:34-76; ledger row P4(i)

Three readings the table cannot carry. First, pivot routing's deeper one: *pivot_signed* was *pivot_unsigned* wearing a name — at the harness geometry the operator is entrywise non-negative with minimum entry exactly 0.0, so the paired test that credited sign compared two non-negative operators, and every M2 and S2 headline was measured on `tgate`, which ships nowhere. The negative on signed attention is therefore also a negative on the record's ability to say which operator it had measured; `tests/loop/test_measured_operator_is_shipped.py` exists because of it, and card **S-01** extends the identity manifest so a cell cannot be journalled without

its boundary condition. Second, the hop-2 mechanism transfers: *the label composes values along paths; the arms composed weights across positions* — the algebraic content of the shape’s second identity, and the reason §2 does not add a hop to softmax but makes the resolvent the operator. Third, the native skyline: DeltaNet’s WY matrix $T = (I + \text{tril}(\text{diag}(\beta)KK^\top, -1))^{-1}\text{diag}(\beta)$ is the exact inverse of a strictly lower-triangular, content-dependent, signed operator, shipping in `fla/ops/utis/solve_tril.py` (CITED [V] [136]; decay gate [135]; semiseparable framing [34]). “Beats native” is banned for the reason “beats softmax” is — a fellow approximator of the same class at the same cost — so what the shape may claim is a Δ_{sky} column with an interval plus the *mechanism* difference: DeltaNet and SSD carry no absorbing rows, no committor read, no re-solve under $\text{do}(a)$. That difference is written NOT FOUND – **sweep owed**, never “novel”.

The taxonomy. MISTAKES.md carries **66** mechanism headings (**RUN grep -cE '^### [VPMD]-'** MISTAKES.md; the record’s own 65 omits the sub-entry **V-14a** at :770) in four classes, ordered by the cost the record assigns them: **V**, vacuous controls, a control that cannot fail (26); **P**, provenance, a number with no live producer or a claim true once (11); **M**, measurement, where the instrument measured but not the thing the verdict names (21); **D**, design, where the answer was fixed before any data arrived (7). The largest single entry is **D-1**, *the largest one in the repository, and it subsumes most of the null results*.

Table 4.3b — the ten most likely to recur here, the rule each imposes, and its card.

#	mechanism	why the shape invites it	the rule it imposes	card
1	D-2 (:710) the oracle is the arm’s own resolvent	the shape’s state <i>is</i> $z = (I - \gamma P)^{-1}V$, so a bed labelled $(I - \gamma^* P^*)^{-1}V^*$ makes every contrast a reproduction check	the label is generated by a latent environment chain P^* the arm never sees; the VOID-contrast list is registered before any cell runs; $\ \hat{P} - P_{\text{env}}\ _\infty$ is a learnability reading, never a capability one	S-13, V-9

#	mechanism	why the shape invites it	the rule it imposes	card
2	D-1 (:677) racing a proven optimum	every bed in the tree scored one real at position $s - 1$	no bed scores a scalar at one position: the label is the vector z^* , the displacement field Δz , or the argmin over candidate moves; the regime in which the softmax skyline is optimal is stated with the L/d geometry printed	S-11, S-16
3	V-24 (:1658) an identity bind with an empty rejection region	$\gamma = 0$ giving PV is proved by one fact and says nothing about P , the boundary sets, or the solve	every bind ships the mutilation battery — drop the absorbing rows, drop the solve, replace P with noise, substitute the label for z — admissible only if each mutilation fails at $O(1)$ with counts printed	S-20..S-27, V-4
4	V-8 (:136) / V-12 (:189) the boundary sets make the label constant	one absorbing target collapsed a label to 1.11×10^{-14} ; a one-component graph gave label sd 0.0	the builder prints at construction $\text{sd}(q^{(k)}) > 0$ for every k , class frequencies inside $(0.05, 0.95)$, the argmin-uniqueness fraction, the disagreement fraction, the discard count	S-12

#	mechanism	why the shape invites it	the rule it imposes	card
5	V-25 (:1954) a theorem whose hypothesis no draw satisfies	the certificate needs P non-negative row-stochastic with $\gamma < 1$; the record's own $0 < a_k$ theorems admitted 1 of 3	a domain census beside every gating theorem — the measured support of γ and of $\text{rowsum}(P)$ on drawn cells, at the quantifier level the theorem uses; 0% blocks the gate	J-L18, S-42
6	M-13 (:1210) / M-9 (:580) a margin or α the design cannot reach	parity needs $N = 70$ for TOST power 0.80, and at $N = 8$ two bit-identical arms return NO VERDICT	parity is claimed by identity bind with rule 3's battery, never by TOST at $N < 70$ (paired $N \approx 36$ only once sd_d is measured — currently NOT MEASURED); every seed CI prints n^+	S-65, V-18
7	M-21 (:2101) / M-18 (:1483) a diagnostic whose value the corpus fixes	the gate was written into the context for the whole campaign, and a corpus-alone probe reads gate $R^2 = 1.000000$ with no arm	the contract names the <i>discrimination</i> , never the statistic; a statistic is admitted only after three readings — corpus-alone ceiling, zero-step floor, headroom printed; $\text{Var}(z) = 0$ on the corpus blocks registration	S-14, S-15

#	mechanism	why the shape invites it	the rule it imposes	card
8	D-3 (:727) / P-8 (:387) a dial that does not vary, a price that is an upper bound	$s = 64$ on 40 of 40 cells with the exponent unidentified; a headline of about 29.6 h was 5.9 h	the cost law is measured at $s \in \{64, 256, 1024, 4096\}$ with synchronize() and randomised order; every headline price carries its direction and the implementation it assumes	S-66, V-21
9	D-7 (:2037) / L-SIGN a prediction without its counter	nine contract statements, seven optimistic, $p = 0.0352$	every prediction ships a counter of equal specificity, filed and hashed before the first arena cell; the calibration column is appended whether or not it flatters	S-60, V-1, V-19
10	V-3 (:72) / V-10 (:168) an assertion that is an identity of its own construction	the Neumann bound is <i>attained</i> , not slack: for non-negative row-stochastic P the tail's row sums are exactly $\gamma^{K+1}/(1 - \gamma)$ (DERIVED), so “error at most bound” is a row-sum identity	before any inequality enters a verdict, state what set each side ranges over; where the right side is computed from the left's hypotheses alone, label it <i>definitional</i> in the output; the certificate is measured on the shipped mask in vector units, never on the truncated series	S-41, S-43, V-20

Three more sit one step behind: **V-17** (:855, a threshold out of its units — a committor threshold is stated in NRMSE or floor units, never raw feature distance); **M-15** (:1289, a statistic without its sharpness-matched null — any β_0 persistence number needs its random-operator null first, card **S-71**); **M-2** (:451, a threshold refitted after the data — every threshold in §5 is frozen with provenance beforehand, which is why the round declined to pick its own tail ruling).

The eight corrections to the record found while writing this paper. Each is a defect in the repository's own documents, found by reading them against their sources, and filed here

rather than by editing a verbatim-of-record file.

1. **The contract’s M10 lineage identifiers are all misattributed.** arXiv:1905.12200 is Bruel-Gabrielsson et al., 1904.09378 is PersLay ([22]), 2011.05804 is Corcoran and Deng ([32]); the intended papers are Hofer 1906.09003 ([53]), Moor 1906.00722 ([87]), Carrière 2010.08356 ([23]). **P-5** (:343) compounded by **P-10** (:1382).
2. **The Cheeger line is stated on a non-reversible chain.** MATHEMATICS.md:455 cites Levin–Peres–Wilmer Thm 13.10 ([73], read at p. 183: $\Phi_\star^2/2 \leq \gamma_{\text{gap}} \leq 2\Phi_\star$), whose hypothesis is **reversibility**. A causal P with $P_{i0} > 0$ has $P_{0i} = 0$, so the line is a **V-25** exposure as written. It is restated on the transient block’s lazy symmetrisation $Q_s := (Q + Q^\top)/2$, whose Cheeger bound controls that surrogate’s spectral gap and **not** $\rho(Q)$, and moved to the oracle cross-check chains (BED-1, E4′) where reversibility holds. Card **J-D6**.
3. **Three Lean declarations the contract tags [M] do not exist.** CEQ_V20_R15_CONTRACT.md:179, :185, :213, :260 cite Lean **#18**, **#19**, **#22** as machine-checked; **RUN**, `grep -rn 'segment' lean/CEQ/` returns one hit, a prose comment at V16Domain.lean:53, and no declaration carries that content. The [M] is the contract’s *grade*, not a build, and the F0 instance the same line prints $(8.9 \times 10^{-16}, 595\times)$ has **no located producing file**. **P-11** (:1597) realised; the targets are re-registered as work in cards **J-L6** and **J-L7**.
4. **THEORY.md:22 calls the successor operator new.** The row reads *Successor operator* $(I - \gamma P)^{-1}$ / *new* / *exists, mislabelled* — its audit column already contradicts its claim column, and the operator is Dayan 1993 ([36], whose eq. 3.1 is $(I - Q)^{-1}$ with **no** γ , per Occupancy.lean:6–17). **P-3** (:316).
5. **arXiv:2604.25655 does not cite Basseville and Nikiforov.** “The learned-model form of §7.2.4” is the record’s own analogy, not the source’s claim, and is carried [U]. **P-10**.
6. **The journal’s head phase table is stale.** It still reads *it.15 DONE*, *it.16 next* for Phase C and *not started* for Phases D and E while the body carries entries through *it.35* — **P-3 sitting inside the block built to catch P-3**. The plan reads its start state from the *it.35* entry.
7. **Rulings 11 and 12 have no text in the tree**, although CEQ_V20_R15_CONTRACT.md:59 binds “Rulings 1–12”: **RUN**, `grep -rniE "ruling 1[12]\b" *.md` returns zero hits, and the same holds for L-GRADE, cited once with no rubric anywhere. The plan states the rubric it uses and marks it a reconstruction; it cites neither as if it had text.
8. **The scoreboard is 2 of 44 and the leap call is still owed.** Both points are the *it.4* freeze; the theory table’s +6 was refused twice and carried unclaimed through *it.35*; the round’s single leap call was **not spent**; *it.36* to *it.45* are unspent.

The bibliographic traps. The audit merged eight sweep bibliographies into **408** canonical entries (508 parsed, 51 same-key duplicates dropped, 49 same-identifier entries folded through `bib_aliases.md`) and re-ran the identifier check directly against arXiv and Crossref, the audit service having returned “not subscribed”: 278 of 278 arXiv identifiers resolved, 91 of 92 DOIs resolved, and all seven title-agreement misses were subtitle truncations of the same works. The one unresolved entry is [113], a Eurographics identifier absent from Crossref, carried [U]. Four traps were caught by hand: DOI 10.1007/978-1-4684-9455-6 is *Denumerable Markov Chains* and **not Finite Markov Chains**, so it must not be attached to the Kemeny–Snell entry; arXiv:2501.00663 is Titans, not test-time regression (2501.12352); arXiv:2302.11294 is not Sander et al. (2302.01425); Contreras Arredondo et al. carry the journal title *Learning the committor without collective variables*. One further trap was found **RUN** this session by scanning the merged file for duplicate titles: **three** groups survive the identifier-level fold — [2] with [3], [9] with [10], and [65] with [67] ([66] being a third key for the Springer printing under its appendix subtitle). All are books carrying neither an arXiv identifier nor a Crossref DOI — precisely the 38-entry class the audit says it could not check — so the fold is *known* incomplete there and the assembler picks one key per work before typesetting. None of this establishes that a source supports the sentence it is cited for: that is the [V-eq] mark, carried per citation, and the record’s rule stands that theorem

and equation numbers read through a rendered page are re-checked against the compiled PDF before publication.

4.4 Where the campaign stands

The scoreboard. After thirty-five iterations of a forty-five-iteration round the carried score is **2 of 44** — both points the it.4 wing freeze, nothing else banked. The denominator is 44 and not the contract’s 48 because the *winner’s cost at most one-tenth of the incumbent’s bonus (+4)* was shown unreachable by either surviving wing: on both timing bases W3 costs **more** than softmax, $1.078\times$ and $1.059\times$. The theory table’s +6 was frozen at it.14 and ruled NOT FIT for the leap, refused at it.10 and again at it.11, carried unclaimed; the +12 is gated on a single unmade author ruling; the annex’s +4 is forfeit. The single leap call was **not spent**, for the reason on the record: *calling it now would consume the round’s single shot on a table whose header miscounts its own body, whose published grading rule contradicts its applied one, and three of whose cells the gate cannot read*. Phase E has not begun.

The frozen wings, and the struck one. FROZEN-N = 2; FREEZE-SHA256 = fbf17e07e6cab495bd1fdbb4392ab88e92f5b0610e40d2b8a283d8bcb9c542ff over eight normalised wing|clause|path:line|anchor rows — the digest covers which wings, lines and anchor strings, and deliberately **not** the content of the cited files, which is why the anchor check is a separate node re-run at every HEAD. W1 is arm_smprime, the path product $G_{ij} = \prod_k m_k e^{i\theta_k}$ with β , QK and g switches (ceq/arm_smprime.py:144); W3 is arm_pl, a real prefix scan, key-side only, over a value-zero BOS sink (ceq/arm_pl.py:88). Wing identifiers are held at their it.1 values so the strike of **W2** is visible as a gap rather than hidden by renumbering. The discriminator that struck it, stated so it can be run: *a wing is FOUND if and only if results/ holds at least one journalled record whose kind field is the wing’s arm name; otherwise it is NAMED — it exists in source and in prose, and the ledger has never seen it*. The census by kind over every results/**/* .jsonl: arm_pl 191, softmax 182, arm_smprime 182, arm_phase 0. The counter is calibrated on both sides — it must read 0 for arm_phase and above 0 for softmax — so the zero is a measurement and not a broken search (V-7, :117). **W2 is STRUCK and RETIRED to UNPRODUCED, not refuted**: the route back is one make_arm branch, an optimizer path and a field emit, at 5.114 s per cell on the arm_pl basis, roughly 41 to 130 GPU-s for eight cells. *The W2 gate is under half a minute of wall clock on this box. It was never a cost decision*. Every arm the plan proposes enters under this discriminator (card S-28).

The eight gate cells, both offices, no coordination. The theory table grades two wings against Q1–Q6: twelve cells, one F0, five F1, one F2, two F3, three F4. The gate’s population is **eight**, not the nine its own heading claims — the ninth row is Q1/W1 at F0, *not a failure* — and one office rules Q4’s two F3s a single harness fact, so the gate graded **seven distinct failures on eight cells**.

cell	grade	JUPITER	MARS	field, and the killer where one is named
Q1/W3	F1	LEAPABLE	LEAPABLE	realization theory (Hankel/Kronecker); formal verification of recurrence invariants by induction on a telescoping product
Q2/W3	F1 + const	TERMINAL	TERMINAL	the bound $\text{err}_i \geq \text{dist}(t_i, \text{hull})$ — row-stochastic implies a convex combination for every (g, s, q, k) , a nonexistence at any size; killer, one (g, s, q, k) with $Z_i \neq 1$
Q3/W1	F2	LEAPABLE	LEAPABLE, conditional on an instrument swap	transfer-operator / Koopman spectral theory, attached to <code>lambda_hat_live</code> at :384 and not <code>lambda_hat</code> at :383 — one bit
Q3/W3	F1 + const	LEAPABLE	TERMINAL	bifurcation theory / gradient-flow convergence against “causal direction unidentified from $n = 8$ ”; killer, a declaration predicting the sign of $\hat{\lambda}$ from initialisation alone
Q4/W1	F3	TERMINAL	TERMINAL	the exponent in S is unidentified at $n = 1$, and no theorem supplies a slope from one point

cell	grade	JUPITER	MARS	field, and the killer where one is named
Q4/W3	F3	TERMINAL	TERMINAL	the same harness fact; brute_force_path_sums at ceq/arm_pl.py:304 is $O(2^S)$ and unguarded — an implementation guard, not a theorem
Q5/W1	F1 + const	LEAPABLE, field contested	LEAPABLE, field corrected	approximation theory / Kolmogorov n -width — not the table’s rate-distortion, on both readings
Q5/W3	F1 + const	LEAPABLE	TERMINAL	finite-mixture inference against “a mean over a bimodal population is not a statement about either mode”; killer, a simultaneous component-wise coverage theorem from pooled bootstrap draws

Tallies: 5 LEAPABLE and 3 TERMINAL against 4 (one conditional) and 4; agreement on 6 of 8 tokens, both offices against the table on Q5/W1’s field. **The two divergences are a missing amendment, not a judgement call.** Q3/W3 and Q5/W3 are exactly the cells whose verdict column carried a hybrid — *LEAPABLE, but about 6 GPU-s buys it outright* and *LEAPABLE, and it is a SCORING RULE, not a theorem* — placed in the gate-class and field columns, where the ruling that a hybrid verdict is not a verdict does not reach. Beneath them: the ledger publishes one grading rule (LEAPABLE when the gap is a missing statement **or a missing measurement**) while a ruling resolved a cell to TERMINAL against that rule — *the round has published one grading rule and applied another, twenty-one iterations, it.14 to it.35*. Three further cells the gate cannot see at all (Q2/W1, Q6/W1, Q6/W3), one filing a **refutation** inside a token meaning *unattempted*; that ruling is in its seventeenth iteration, unruled. Mechanisms **V-14a** (:770, the scope test that condemns every refusal guard — the role must be declared) and **M-20** (:1723, a pre-registration predicting both outcomes in two sections that never met).

The laws in force, one line each. **D-1:** work is a DAG, parallel dispatch only on nodes with no shared repository state. **D-2:** planet names label responsibilities inside documents, never concurrent processes. **D-3:** no autonomous loop mounts until the previous loop’s cause of death is one sentence, and the iteration count comes from the DAG’s critical path. **D-4:** contracts

scheduled behind an unreached round are staged, not started. **L-DOM**: every theorem that gates a run ships a domain census, and no overlap means decoration (pays for V-25). **L-SIGN**: a counter-prediction of equal specificity beside every prediction, with a calibration column across rounds (pays for D-7). **L-DIAG**: a contract prescribes what a diagnostic must distinguish, never which statistic does it (M-18). **L-FLOOR**: every capability number ships beside its information floor, so “how good” reads as distance-to-floor — and floor₁ is not one. **L-CERT**: every mask ships its certificate, F0 exact or F1 with δ printed; a mask without one is refused. **L-EQ**: [V] is inadmissible for a load-bearing statement, and [V-eq] requires the theorem with hypotheses plus one numeric instance run. **L-LEAN**: the arm is not trained before its identity theorems are green. **L-G2**: journals never move and are never deleted, superseded cells staying with a supersede marker. **FOUND-not-NAMED**: a wing named rather than found is struck. **R-SKY**: the native skyline is read beside every bed with a Δ_{sky} column, and “beats softmax” is not licensed where softmax is a fellow approximator. **Ruling 1**: CUDA determinism with `warn_only=True`, bitwise for replay and every deciding forward cell, training compared against a measured noise floor. **Ruling 3**: matched parameters, a 0.032% residual counting as matched, exact counts in every table header, no re-architecting to close it. **Ruling 10'**: “pinned” is decided by $\Lambda = 2[\text{LL}(\beta_{\text{final}}) - \text{LL}(\beta \equiv 1)]$ on held-out data, with the minimum detectable departure printed. Two of these have thinner text than their citation and are handled as in §4.3, correction 7.

The open debts. Five author rulings stand open at it.35: the clause-1 tail, one- or two-sided; the missing F0–F4 rubric; the Kaggle attach behind the chess witness; the gate’s blindness to F4 cells; the collision between TERMINAL and NOT-PUT. Carried from earlier rounds: **D-APPROX**, the approximation bound owed beside every exact-identity nonexistence theorem — *every “X cannot represent Y” claim needs a bound before it reads as “X cannot fit Y”*, the largest unclosed gap in the round’s own logic; **D-R3**, two mutually inverse registrations for one bed; the memory ceiling that dropped the $n = 32768$ reproduction and then the 16,384 one; the closed-versus-half-open magnitude interval, filed two-sided because $[0, 1)$ would cost the reachability of $a = \pm 1$, two-thirds of BED-M’s support; the scan skyline, refused in R11 and legal now as the native control; the HuggingFace package, scheduled and unbuilt. Harness debts the plan must not inherit silently, all **RUN** against `scripts/v15_r1.py`: S, D = 64, 24 is a module constant at :137, with nine sibling `add_argument` calls and no `seq_len` flag anywhere; `torch.cuda.synchronize` does not occur in the file; `lambda_hat` at :383 averages $\log m$ over every position, so one $m_k = 0$ sends it to $-\infty$ and it carries one bit.

Kaggle and HuggingFace. Nothing has launched. *Nobody who wrote these files ran the kaggle CLI, pushed a kernel, or touched ~/.kaggle; launch is the author’s own call.* The record has already refused a relayed instruction as that consent — *a coordinator message is not that say-so* — and the author’s standing instruction is the same rule from the other side. Gate 0 closes only at open-rulings zero and carries a circularity the register does not resolve: two of its rows are blocked on numbers only the Kaggle run produces. What exists: `ceq/hf/` with four modules; **no trained checkpoint** (the Q3 checkpoint slot reads NOT MEASURED); parameter counts that are not equal (25,736,232 against 25,728,000, $+8,232 = +0.03200\%$, the arm carrying more); a model card whose own header states the shipped operator does not work. Three beds are PINNED as generators with sha256 digests, `enwik8` is PINNED as a sliced attach, and three sources are UNPINNED — for which the notebook’s hash cell **prints and then raises** `kdata.MissingPin` rather than skipping silently. Every Kaggle line in §5 is priced and none is launchable by any agent: the author’s explicit yes is a node on the critical path, not a formality (cards **S-73**, **V-22**).

Why the campaign could not win on its own beds, and what the plan changes. Every bed in the tree asks for one real at position $s-1$: `equilibrium_oracle` returns z_{s-1}^* , *the last coordinate of $(I - A)^{-1}b$* (`READ scale/negation_scope.py:286-304`), and `ArmSMPPrime.forward` returns a tensor of shape $[n]$. That is the single-location regression shape on which one attention layer is asymptotically Bayes-optimal ([77], CITED [V]; [38] for softmax proper). The record files this as **D-1** and carries its own caveat beside it: the theorem is asymptotic under $d \rightarrow \infty$ with

$L = o(d)$, its predictor is **erf** and not softmax, and the repository ran $L/d = 4.00$ — “provably” is not earned at this geometry. So D-1 is not a theorem the campaign lost to; it is the reason the campaign could not distinguish *cannot be beaten* from *was not beaten*, and 0 of 39 scoreboard items were earned under it. The arena confirms it from the metric’s side: Q6 is F4 on both wings because the arms return one real per draw and no state axis exists for a state metric to read. What the programme changes is the **label class**, not the effort. No milestone in §5 scores a scalar at one position: the labels are the jointly determined configuration z^* , the displacement field Δz under $\text{do}(a)$, the $K + 1$ reach-avoid committors, and the argmin over candidate moves — each with its own information floor (L-FLOOR), its own zero-hop guard (**S-14**), its own leak guard firing both ways (**S-15**), and its own VOID-contrast list registered before a cell runs (**S-13**). Three of those cost 0 GPU-s. Whether the shape is worth anything is not settled by this section and is not settled by this paper; it is settled by cards **S-12**, **S-52** and **S-62**, in that order, and each of them can kill it.

5. THE PROGRAMME — the development plan for Consequence-Equilibrium Attention

MYCROFT (chair), 2026-09-03, HEAD 207e7b9, branch v17k-gate0. This section is the body of the paper (THESIS_CORRECTIONS_2.md §0: the deliverable is the plan) and is written to be read standalone, repeatedly, over months, by one developer working alone in evenings on the certified RTX 4060 Laptop (7.996 GiB, torch 2.5.1+cu121, READ COSTS.md:53-63). It merges the six roadmaps — plan/plan_jupiter.md (Lean, J-), plan_saturn.md (instrument, S-), plan_mercury.md (price, M-), plan_mars.md (attack, R-), plan_venus.md (prediction, V-), plan_neptune.md (systems, N-) — deduplicated by content, with every planet’s card kept and its aliases printed on the card’s header line. The judged sections (judge/sec_shape.md, judge/sec_obstructions.md, judge/proposition_ledger.md, judge/sec_apparatus.md, judge/bind_ledger.md) outrank the roadmaps and the designs wherever they differ; where a roadmap’s number was overruled, the card carries the ruling.

Evidence classes on every load-bearing claim. RUN — executed on this box this session or in a named planet’s session, carried with its geometry; READ path:line at HEAD 207e7b9; CITED [V] (primary page fetched, title matched) or [U] (reached through an index) by canonical references.bib key; DERIVED with the steps shown; [FITTED] with its R-squared, [ASSUMED] with its reason (READ COSTS.md:16-20). No number appears without one of these. Struck constants of STRUCK.md appear nowhere. Nothing here is a result: no CEQ arm has been trained, BED-S has no cell, and every BED-S figure below is a floor formula or a design constant.

Two id spaces collide and are disambiguated once: CONTRACT.md D-1 is the dependency law (work is a DAG), MISTAKES.md D-1 is the design mechanism (racing a proven optimum). Mechanisms cited as bare D-n / M-n / P-n / V-n are MISTAKES.md’s 66 headings (RUN grep -cE '^### (V|P|M|D)-[0-9]+' reads 66, including the lettered V-14a at MISTAKES.md:770); laws are cited by their L- name and rulings by number.

5.0 The shape, and the claim the programme can at best earn

The shape, in one sentence. Consequence-Equilibrium Attention is one causal, row-stochastic attention operator read through a single triangular solve,

$$z = (I - \gamma W)^{-1} V, \quad \Pi_\gamma = (1 - \gamma) W (I - \gamma W)^{-1}, \quad O = \Pi_\gamma V,$$

over the record's Lean-checked three-corner base $W_{\beta,qk,m,\theta}$ (READ `lean/CEQ/V16Domain.lean:92-97`, READ `ceq/arm_smprime.py:144`), with a declared value-zero sink set $\mathcal{A}_{\text{sink}} = \{0\}$, a goal set $\mathcal{A}_0$ and $K \geq 2$ constraint sets $\mathcal{A}_1, \dots, \mathcal{A}_K$ written as identity rows, so that one solve returns the jointly determined state z , its displacement Δz under an intervention, and the reach-avoid committor $q^{(\bullet)}$ of every transient position; at $\gamma = 0$ the read is softmax bitwise (RUN[coord] `torch.equal True` at $s = 64$, $d = 16$; rejection region $\max |O(0.5) - O(0)| = 2.3002850040264393$).

The claim the programme can at best earn, in one sentence. On BED-S — a registered bed whose environment chain is a random DAG in token order, hence inside the arm's own operator class — CEQ reads the reach-avoid vector below the exact restricted-view floor and the displacement field above depth-1 softmax's field cosine at matched depth and parameters, with an exact certificate ($\delta = 0$ on the solve route), a trained $\hat{\gamma}$ that is MOVED under the boundary null $\frac{1}{2}\chi_0^2 + \frac{1}{2}\chi_1^2$ (95 per cent point 2.7055, RUN[J]), and a separation from ChaCAL-with-a-sink-token that a column device does not reproduce — and it can never earn “beats softmax” on any bed where a deeper softmax stack is a fellow approximator that computes the same resolvent by iteration ([126], [131] [V]; R-SKY, READ `CEQ_V16_CONTRACT.md:209`).

Everything the programme can lose is written down beside what it can win: §5.10's outcome trees name the sentence each ending licenses and the replacement route it owes.

5.1 How to use this plan

The phases and their gates. Seven phases. Each ends in a decision the author takes alone, from printed numbers, without an agent.

phase	goal	the gate at its end	the numbers that decide it
0 free work: theorems, manifests, specifications, censuses	no arm trains before its identity theorems are green (L-LEAN, READ <code>CEQ_V15_CONTRACT.md:51-58</code>)	lake build exit 0, zero sorry, axioms [proptext, Classical.choice, True sound] only; 5/5 manifest drift flips move the hash; every obstruction sentence carries its vacuity inequality	
1 BED-S construction and admission	D-4 admission: a bed not admitted produces no reading	every label sd > 0.05 over the admitted region; every argmin class frequency in (0.05, 0.95); rule-disagreement fraction > 0; $0 \in \mathcal{A}_{\text{sink}}$ and every set before the query on 100 per cent of draws	

phase	goal	the gate at its end	the numbers that decide it
2 binds through the front door, the smoke test, the device certificate	V-14 / V-24: every plant fails at $O(1)$ as a FOUND cell with its own kind	the 0.9749/0.9165/1.000/0.4845 pattern (READ workdonenewseal.md:114-122); torch.equal True at $\gamma = 0$; cell price inside the band [−1.8%, +14.4%] of 1.680 s	
3 the arena at the design point	the realised paired sd fixes the MDE row; every arm is FOUND, not NAMED	realised paired sd against the pilot's 0.050146; MDE ₈ from the noncentral- t bisection (0.039827 at sd 0.034451, RUN[VENUS], RUN[MARS] 0.0398266)	
4 verdicts, the remaining certificates, the dossier	the five bets scored with their counters; K-E1 fires or escalates	per-bet PASS / KILL / SPLIT with the deciding interval; Λ per seed against 2.7055 and $\ln 4096 = 8.318$	
5 the cost law, memory, kernels	no price is quoted at any $s > 64$ before the sweep runs (K-9)	the fitted exponent's CI in s ; the run-order Spearman ρ against the record's +0.7029	
6 the Kaggle gate and the LM cell	the author's explicit yes is a node, not a formality	every local certificate line filled; open rulings listed by name beside the request	

DAG conventions. Every card names its **prerequisites** and what it is **independent-of** (CONTRACT.md D-1: work is a DAG; parallel dispatch only on nodes with no shared repository state). A card marked **parallel-safe** touches no file another live card touches and may be picked up in any order the prerequisites allow; **results/*.jsonl** are append-only per lane (L-G2), so two cards writing different lanes are parallel-safe, two cards writing one lane are not. A card's **price** is tagged [FITTED] with its R-squared, [RUN], [ASSUMED] with its reason, DERIVED with its arithmetic, or NOT MEASURED with the instrument it needs. **Evenings** are the author's time, never GPU time. * marks a cheapest-decisive card. Aliases on a header line name the planets' own card ids for the same content, so a reader of one roadmap can find its card here.

The five cheapest decisive cards. Each settles something alone, at 0 GPU-s or seconds, and none blocks another (the coordinator's order, judge/bind_ledger.md §7):

1. **J-L18** — the ten [M] Lean targets build with their refusals and #print axioms (0 GPU-s). Decides whether the paper may write “machine-checked” for component (k).
2. **S-01** — the extended identity manifest with its one-unit drift plant (0 GPU-s). Decides whether any cell in the lane can be attributed at all (L-2: 0 of 24 attributable, READ V20_R15_LEAP_LEDGER.md:23).

3. **S-12** — the BED-S domain census on 512 draws (0 GPU-s). Decides admission, which is the record's dominant failure shape made free (D-4, M-3, V-8).
4. **S-20** — the ChaCAL-diag smoke test at $\gamma = 0.9$ on BED-M with the B-J plant battery (≈ 1.7 s). Decides whether the solve path runs, is deterministic, and prices as predicted.
5. **S-61** — the γ -pinning likelihood-ratio instrument specified with the boundary null, the held-out set and the seed rule (0 GPU-s). Without it Bet C cannot be scored.

FIRST EVENING — **five cards that each settle something at 0 GPU-s or seconds.** J-L18 (with J-L0, J-L1, J-L3 written the same sitting), S-01, S-12 (with S-11 written first), S-20, S-61. A sixth, independent of all five and long owed, is **S-52**: the ≈ 6 GPU-s capped run at seeds 2, 3, 7 on BED-M that the record priced at it.7, it.8, it.9 and it.35 and never took (READ V20_R15_IT35_JUPITER.md:157-159). The evening's total GPU cost is under 20 seconds.

How to read a KILL. A KILL line is the point estimate, not the tail. The record's only scored prediction census reads 9 checked, 9 adverse, 7 of 8 signed rows optimistic, one-sided sign test $p = 0.0352$ (READ V16_CALIBRATION.md:96-100), so D-CALIB-1 makes the counter the planning estimate for every unscored card. That rule is why the plan is ordered cheapest-refutation-first and why every phase can end honestly.

5.2 The laws in force

One line each, with its source line (sections/sec_state.md §S.4), and the mechanism each pays for. These bind every card below.

law	one line	source	pays for
D-1 dependency	work is a DAG; parallel dispatch only on nodes with no shared repository state; git is the shared state	READ CONTRACT.md:27-33	the independent-of field on every card
D-2 skills are modes	planet names label responsibilities inside documents, never concurrent processes	READ CONTRACT.md:35-38	the J-/S-/M-/R-/V-/N- prefixes here
D-3 loop gate	the iteration count comes from the DAG's critical path, never chosen by an agent	READ CONTRACT.md:40-44	§5.4's evenings count
D-4 order	contracts that schedule work behind an unreached round are staged, not started	READ CONTRACT.md:55-57	the phase-1 admission gate
L-DOM	every theorem that gates a run ships a domain census; no overlap means decoration	READ CEQ_V16_CONTRACT.md:54-56	V-25 (READ MISSES.md:1954-1956)

law	one line	source	pays for
L-SIGN	a counter-prediction of equal specificity beside every prediction; a calibration column across rounds	READ CEQ_V16_CONTRACT.md:50-51	D-7 (READ M-51 TAKES.md:2037-2039)
L-DIAG	a contract prescribes what a diagnostic must distinguish, never which statistic does it	READ CEQ_V16_CONTRACT.md:61-66	M-18 (READ M-51 TAKES.md:2103)
L-FLOOR	every capability number ships beside its information floor, so “how good” reads as distance-to-floor	READ CEQ_V20_R15_CONTRACT.md:60-65	C15: floor_1 is not a floor (READ V20_R15_JOURNAL.md:51)
L-CERT	every sparsity mask ships its certificate, F0 exact or F1 with δ printed; a mask without one is refused	READ CEQ_V20_R15_CONTRACT.md:64-67	V-10, V-17
L-EQ	[V] is inadmissible for a load-bearing statement; [V=eq] needs the theorem, its hypotheses and one numeric instance	READ CEQ_V15_CONTRACT.md:51-55	P-10
L-LEAN	the arm may not be trained before its identity theorems are green	READ CEQ_V15_CONTRACT.md:57-58; READ CEQ_V16_CONTRACT.md:50-51	P-11
L-FIRST	predictions and counters are filed before each run — listed as standing; its definitional text was not located this session	READ CEQ_V16_CONTRACT.md:49, :191	M-2
L-G2	journals never move and are never deleted; a superseded cell keeps a supersede marker	READ V17K_RULINGS.md:62-64	P-3
FOUND-not- NAMED	a wing named rather than found is struck; FOUND iff results/ holds a record with its kind	READ CEQ_V20_R15_CONTRACT.md:49, :266	P-4

law	one line	source	pays for
R-SKY	the native skyline is read beside every bed; “beats softmax” is not licensed where softmax is a fellow approximator	READ CEQ_V16_CONTRACT.md:209	MISTAKES.md D-1
Ruling 1	determinism: CUDA with <code>warn_only=True</code> ; bitwise for replay and every deciding forward cell; CPU-strict refused	READ V17K_RULINGS.md:39-45	M-10, M-16
Ruling 2 / 2a	β learnable and logged per instance; no sentence transfers across corners without a bind at the corner it describes	READ V17K_RULINGS.md:47-54, :275-325	V-22
Ruling 3	matched parameters: a 0.032 per cent residual is matched; exact counts in every table header; never re-architect to close it	READ V17K_RULINGS.md:56-59	M-8
Ruling 6	the flight envelope 6a-6f: stale is two polls, tier precedence, OOM keyed per shape, deciding-cell list frozen at launch	READ V17K_RULINGS.md:71-81	V-15
Ruling 7	a bed is carried as generator plus seed plus hash and regenerated in-notebook; the hash assertion is the provenance	READ V17K_RULINGS.md:83-86	P-1
Ruling 8	Q2 dropped: 10.578 GiB against the 4060’s 7.996 GiB	READ V17K_RULINGS.md:328-339	P-8
Ruling 9	R1-prime cells fall under the measured-floor clause; no equality statement without its floor printed beside it	READ V17K_RULINGS.md:354-379	L-FLOOR

law	one line	source	pays for
Ruling 10-prime	“pinned” by likelihood ratio on held-out, with the minimum detectable departure printed; supersedes the ulp criterion	READ V17K_RULINGS.md:389-436	M-20
the R15 kills	named-not-found is struck; a mask without a certificate is refused; a capability number without its floor is not a number; a leap output acted on before its instance runs is struck	READ CEQ_V20_R15_CONTRACT.md:265-273	§5.9

Two laws thinner than their citation, recorded so the plan does not lean on them. L-GRADE (the F0-F4 rubric) is cited once and no rubric text exists in the tree (READ V20_R15_THEORY_TABLE.md:22-24); Rulings 11 and 12 are cited as binding and have no text (RUN grep, zero hits outside the two citing lines). This plan states the grading it uses and marks it a reconstruction (P-11).

The author’s supersession of 2026-09-03. On that date the author licensed a new primitive — “rethink the shape” — and instructed “do not write any code only docs” (READ BRIEF.md:11-30, :40-42). The round-15 contract’s clause *no new primitive that is not a repair* is therefore **superseded by the author for this document**, and this plan says so here rather than by editing the contract, which is verbatim-of-record and is never edited (READ CEQ_V20_R15_CONTRACT.md:3-6). What the supersession does **not** lift, because the author did not lift it: the leap’s remaining clauses (named theorems with runnable instances, predictions **and** counters, every untested claim with its cheapest killer, instances RUN before anything is acted on) and every standing law in the table above. What it changes in kind: the iteration count of this programme is derived from its own DAG’s critical path (§5.4), never inherited from the 45 iterations the R15 schedule named; and the Phase-E template — candidate built with identity binds, manifests, certificates, parameter-matched, full arena, verdict by the registered clauses, purge, prognosis (READ CEQ_V20_R15_CONTRACT.md:153-169) — is inherited with the beds and labels replaced. Mechanism: P-3 (a stale claim never retracted) — the supersession is dated, and §5.11’s corrections index is where it is amended if the author changes it.

5.3 The phases and their milestone cards

Card format, every field on every card: **id** · **title** — phase · prerequisites · independent-of · parallel-safe · aliases; then *build* (a specification, never code), *prove*, *measure*, *PASS*, *KILL* with what dies, *price*, *mechanism*, *deliverable*, *evenings*.

PHASE 0 — free work: theorems, manifests, specifications, censuses

Goal. Put every identity the paper leans on into Lean or into a written derivation, give the lane an identity manifest that refuses an undeclared field, and freeze the specifications and thresholds later phases may not move (M-2). Price of the whole phase: 0 GPU-s.

The gate. L-LEAN. lake build from lean/ exits 0, zero sorry, zero sorryAx, axioms [propext, Classical.choice, Quot.sound] only, every [M] card green; a target that does not build is [S] and is never cited as proved (K-K). The manifest’s 5/5 single-field flips move manifest_hash and an omitted declared field raises. Every obstruction sentence carries its vacuity inequality: $hmp = 512 \geq 64$, $n \log_2 n = 384 < 544$, $64^{1/16} = 1.2968$, all vacuous at $s = 64$, $d = 16$, $p = 32$. The author decides from three printed things: the build log, the flip count, the census table. **No arm trains until this gate is green.**

J-L0 * . Lay the Shape module and run the Mathlib name census — phase 0 · prereq none · independent-of every card · parallel-safe · aliases none *build* one import root lean/CEQ/Shape.lean over CEQ.Contraction, Occupancy, Nilpotent, V16Domain, OracleSeparation, with files Shape/{Parity,Triangular,Boundary,Segment,Displacement,Neumann,Committer}.lean, none above the V15Source.lean scale (201 lines, 7 declarations); shared definitions stated once — LowerTri, Absorbing P A, IsSink 0; journal fields per card: declaration, file:line, grade, #print axioms, lake build exit, Mathlib rev, census line. *prove* nothing; the census is a *grep* over the vendored tree at rev a45ae637 (2024-04-04, toolchain leanprover/lean4:v4.7.0, READ lean/lean-toolchain) for the eight Mathlib names the [M] cards rest on, each recorded [V-name] with file:line or [U]. *measure* lake build exit 0 on the empty module (the warm cache reads [1530/1531], READ workdonenewseal.md:41-45). *PASS* all eight names resolve at the pinned rev (they do, plan_jupiter.md §0). *KILL* a missing name demotes its dependent card to [S], which must then supply the lemma itself; what dies is the [M] grade, never the identity. *price* 0 GPU-s · *mechanism* P-11, P-4 · *deliverable* docs/LEAN_SHAPE_MANIFEST.md · *evenings* 1

J-L1 * . Parity at $\gamma = 0$, with its refusal — phase 0 · prereq J-L0 · independent-of J-L2 to J-L17 · parallel-safe *build* gamma_zero_is_softmax ($P(1 - 0 \cdot P)^{-1}V = PV$) and the witness gamma_half_is_not_softmax on Fin 2. *prove* [M]; Matrix.inv_one, zero_smul, sub_zero; the witness by decide or norm_num. *measure* the float instance stays RUN: torch.equal(0(0), PV) = True at $s = 64$, $d = 16$ (RUN[coord]); rejection region 2.3002850040264393 at $\gamma = 0.5$; the record’s own softmax corner is 1.110223×10^{-16} off ceq/lm.py on 19/64 entries (READ V16_ARM_SMPRIME.md:266-293), so “bitwise” is against the lane’s softmaxAttn, never a fused kernel. *PASS* both declarations build with the three-axiom set. *KILL* the witness does not close: the parity proposition ships RUN-only and Proposition 1 loses its Lean rejection region. *price* 0 GPU-s · *mechanism* V-3 (the pass is by construction, so the refusal is mandatory), V-24 · *deliverable* lean/CEQ/Shape/Parity.lean · *evenings* 1

J-L2 . Row 0 of a causal softmax is absorbing at $\beta = 1$ — phase 0 · prereq J-L0 · independent-of all but J-L10, J-L18 · parallel-safe *build* on the tree’s Hop (READ V16Domain.lean:366-378): bos_row_is_absorbing, refusal bos_row_not_absorbing_at_beta_zero (at $\beta = 0$ row 0 reads $e^{q_{k00}} = 2.0138$, RUN[M]). *prove* [M]; Finset.sum_range_one, Real.exp_ne_zero, Real.rpow_one, div_self. This ships fact F1 of THESIS_CORRECTIONS_2.md §1 as a theorem. *measure* BOS undeclared gives $\rho(Q) = 1.000000$, $\det(I - Q) = 0.0$; declared as $\mathcal{A}_{\text{sink}} = \{0\}$ gives $\rho(Q) = 0.692660$ (RUN[P], RUN[J]). *PASS* builds, and the domain census reads 100 per cent of BED-S draws with $0 \in \mathcal{A}_{\text{sink}}$. *KILL* the theorem needs $0 < Z_{\text{norm}}$ as an extra hypothesis the tree does not supply: grade [S] and census line X-2 stays a per-batch RUN. *price* 0 GPU-s · *mechanism* V-25, V-12 · *deliverable* lean/CEQ/Shape/Boundary.lean · *evenings* 1

J-L3 * . $I - \gamma P$ is a unit on the causal class, with the diagonal interval — phase

0 · prereq J-L0 · independent-of J-L1, J-L2, J-L5, J-L8, J-L9, J-L11, J-L12 · parallel-safe *build* `lower_triangular_isUnit`, `diag_one_sub_smul_pos` (diagonal of $1 - \gamma P$ in $[1 - \gamma, 1]$; the open upper bound needs $0 < P_{ii}$, which J-L5 supplies), refusal `zero_diag_not_unit`. *prove* [M] — `Matrix.det_of_lowerTriangular` (Block.lean:265) and `Matrix.isUnit_iff_isUnit_det` (NonsingularInverse.lean:151), both confirmed at the pinned rev. *measure* minimum diagonal 0.4 at $\gamma = 0.6$, attained by row 0 and every absorbing row (RUN[M]); `solve_triangular` against the dense inverse $1.7763568394002505 \times 10^{-15}$ (RUN[coord]). *PASS* builds. *KILL* the `LinearOrder` (Fin n) instance clashes: *prove* the determinant by induction, grade [S], plus one evening; the triangular branch stalls one evening and the substitution stands as RUN. *price* 0 GPU-s · *mechanism* M-8 (the shape is priced as a substitution; this theorem is what makes the substitution well-defined) · *deliverable* `lean/CEQ/Shape/Triangular.lean` · *evenings* 1

J-L4 · Consequences propagate forward only — phase 0 · prereq J-L3 · independent-of J-L5 to J-L9, J-L11, J-L12 · parallel-safe *build* `inv_lowerTri`, `later_boundary_unreachable` (F2: a set after the query has committor exactly zero), `causal_forward_only`, refusal `dense_P_displaces_backward`. *prove* [M] via `Matrix.blockTriangular_inv_of_blockTriangular` (Block.lean:346); `causal_forward_only` is one line routed through `inv_lowerTri` rather than Sherman-Morrison. *measure* $\max |\Delta z_{<12}| = 0.0$ on the row clamp at $i = 12$ (RUN[P]); 5.6×10^{-17} on the non-nilpotent softmax corner (RUN[F]); the triangularity plant (non-causal P) reads 0.0761 and 0.0868 (RUN[M]). *PASS* builds. *KILL* the `IsUnit-to-Invertible` cast does not elaborate: grade [S] and *prove* `inv_lowerTri` by J-L6’s block formula; the forward-only claim stays a RUN with its plant. *price* 0 GPU-s · *mechanism* V-8, V-24 · *deliverable* `lean/CEQ/Shape/Triangular.lean` · *evenings* 1

J-L5 · The softmax corner is not nilpotent (the regime boundary) — phase 0 · prereq J-L0 · independent-of everything but J-L18 · parallel-safe *build* `lowerTri_pow_diag`, `softmax_corner_not_nilpotent`; the refusal is already in the tree (`Nilpotent.pow_card_eq_zero` :77, `one_not_nilpotent` :105). *prove* [M]; `Matrix.mul_apply`, `Finset.sum_eq_single`, `pow_pos`. *measure* $\max_{i \geq 1} (\gamma P)_{ii}^{32} = 6.27 \times 10^{-13}$ beside $\gamma^{32} = 7.96 \times 10^{-8}$ attained at (0, 0) (RUN[M]) — the discriminating pair the judge fixed after a V-4. *PASS* builds. *KILL* none foreseeable; if the Fin order must be unfolded, plus one evening. Nothing measurable dies: the RUN pair stands and the default “keep the diagonal, carry the Neumann certificate” is unaffected. *price* 0 GPU-s · *mechanism* V-25, P-3 · *deliverable* `lean/CEQ/Shape/Triangular.lean` · *evenings* 1

J-L6 * · A zero gate splits the resolvent into blocks — phase 0 · prereq J-L0 · independent-of J-L1 to J-L5, J-L8 to J-L17 · parallel-safe *build* `segmentation_blockdiag` on `StrictlyLower`, `resolvent_fromBlocks`, the chain corollary `chain_zero_gate_cuts` [S], refusal `tiny_gate_does_not_cut` at a -30 logit. *prove* [M] for the first two (`inv_fromBlocks_zero_21_of_isUnit_iff`, `SchurComplement.lean`:202, confirmed); the induction is `pow_entry_zero`’s. *measure* `segmentation zeros array_equal` on the gated corner (RUN[M]); the -30 plant ($e^{-30} = 9.36 \times 10^{-14}$) leaves the block non-zero; the 10^{-300} plant is struck as float64 underflow. Domain census: BED-M admits a zero gate on 3 of 3 values (READ `V16Domain.lean`:304); BED-S at the softmax corner reads 0 of N and the theorem is silent there. *PASS* builds. *KILL* `fromBlocks`’ `Sum m n` indexing does not match the Fin n cut: ship the entrywise statement only and demote the block inverse to a remark; the F0 row keeps its `torch.equal` RUN and no [M] tag. *price* 0 GPU-s · *mechanism* L-CERT, P-11, V-25, V-2 · *deliverable* `lean/CEQ/Shape/Segment.lean` · *evenings* 1

J-L7 · A cut makes the segment head absorbing and severs the sets behind it — phase 0 · prereq J-L6, J-L2 · independent-of J-L8 to J-L17 · parallel-safe *build* `cut_makes_segment_head_absorbing`, `cut_severs_boundary_sets`, on a gated hop `HopGated` defined once with the product taken directly (the tree’s Hop carries no zero, `no_prefix_scan_represents_a_zero_gate` :165). *prove* [M] both; objects as J-L2 and J-L6.

measure after a cut at c : $P_{cc} = 1.0$, $\rho(Q) = 1.0$ with $c \notin \mathcal{A}$ (RUN[J]); once c is declared, $q^{(0)} = q^{(1)} = 0.0$ past the cut (RUN[M]). *PASS* builds. *KILL* defining *HopGated* forces re-proving *three_corners_containment* for it: plus one evening, still [M]; if abandoned, segmentation and the committor channel are kept apart by rule rather than by theorem. *price* 0 GPU-s · *mechanism* V-25, D-3, V-8 · *deliverable* `lean/CEQ/Shape/Segment.lean` · *evenings* 1

J-L8 * · The displacement identity with its planted zero — phase 0 · prereq J-L0 · independent-of everything but J-L14, J-L18 · parallel-safe *build* `displacement_identity` and refusal `const_value_zero_displacement`. *prove* [M]; `Matrix.sub_mulVec`, `smul_mulVec_assoc`, `mulVec_sub`; four lines. *measure* residual 2.897×10^{-16} (RUN[P]), 1.2×10^{-15} (RUN[I]), 1.36×10^{-15} (RUN[F]), 1.03×10^{-15} (RUN[J]); $V \equiv \mathbf{1}$ gives $\max |\Delta z| \leq 10^{-15}$ (RUN[F] 1.1×10^{-16} ; a bitwise zero is a one-code-path accident); Gaussian V gives 0.1096, 0.363, 1.127 on three draws — the rejection region. *PASS* builds. *KILL* none foreseeable; nothing dies, the three independent RUNs stand. *price* 0 GPU-s · *mechanism* V-24, D-5 · *deliverable* `lean/CEQ/Shape/Displacement.lean` · *evenings* 1

J-L9 · A causal operator is never the oracle's undirected chain (D-2 in regime S) — phase 0 · prereq J-L0 · independent-of everything but J-L18 · parallel-safe *build* `lowerTriangular_ne_symmSupport`, extending `OracleSeparation.oracle_ne_resolvent` (:166, `StrictlyLower` only) to the diagonal-kept class; refusal pattern `zero_not_a_counterexample` (`OracleSeparation.lean:193`). *prove* [M], one line: `SymmSupport` gives $0 < Q_{ji}$, `LowerTri` gives $P_{ji} = 0$. *measure* BED-1's chain instantiates the hypotheses: `Nonneg True`, `SymmSupport True`, $\rho(Q) = 0.9408612510154677$, $Q^{11} \neq 0$ (RUN[coord]). *PASS* builds. *KILL* none. If it fails, the ruling that BED-S's environment must be a DAG in token order stands on the RUN alone rather than on a theorem. *price* 0 GPU-s · *mechanism* D-2 (READ MISTAKES.md:710), V-25 · *deliverable* `lean/CEQ/Shape/Committer.lean` · *evenings* 1

J-L10 · The committor read and its causal invertibility — phase 0 · prereq J-L3, J-L2 · independent-of J-L5 to J-L9, J-L11, J-L12 · parallel-safe *build* `committer_is_resolvent_read` (a), `isUnit_one_sub_transient_causal` (b-causal), `transient_diag_lt_one_of_bos_declared`, refusal `bos_undeclared_singular`. *prove* [M] all three; the last is the positivity argument $P_{ii} = 1 - \sum_{j < i} P_{ij} < 1$ once $P_{i0} > 0$. *measure* $\rho(Q) = \max_{i \in T} P_{ii} = 0.692660$ (RUN[P]); BED-1's real sets $A = [0]$, $B = [1]$, $|T| = 9$: resolvent read against `bed["q"]` 0.0, harmonic residual $1.0408340855860843 \times 10^{-17}$ (RUN[coord]); the must-fire perturbation drives the residual to 10^{-6} , ratio 9.6×10^{10} . *PASS* builds. *KILL* the strict-positivity hypothesis fails on gated corners (a zero gate zeroes P_{i0}): the theorem is stated for the softmax corner only, census line X-3 stays a RUN on gated corners, and the head ships on a per-batch determinant check. *price* 0 GPU-s · *mechanism* D-2, V-10, V-16 · *deliverable* `lean/CEQ/Shape/Committer.lean` · *evenings* 1

J-L11 · Corner 3 is the sub-diagonal resolvent, entrywise — phase 0 · prereq J-L0 · independent-of everything but J-L6's chain corollary and J-L18 · parallel-safe *build* `subdiag`, `subdiag_strictlyLower`, `subdiag_pow_entry`, `pathprod_is_chain_resolvent`. *prove* [S]; the induction is `pow_entry_zero`'s with one surviving term; objects `Finset.sum_eq_single`, `Finset.prod_Ico_succ_top`. The $n = 3$ instance is in the tree (`V15Source.forward_map_fills_in :168`). *measure* $\max |G - (I - A)^{-1}| = 0.0$; last row against `equilibrium_oracle` $6.217248937900877 \times 10^{-15}$; $A^{64} = 0$ exactly (RUN[coord]). *PASS* builds. *KILL* the Ico product bookkeeping does not close: split the two cases, plus one evening; the containment stays a 0.0 RUN and `shape - corner-3` on BED-M remains VOID by registration rather than VOID by theorem. *price* 0 GPU-s · *mechanism* D-2, MISTAKES.md D-1 · *deliverable* `lean/CEQ/Shape/Chain.lean` · *evenings* 2

J-L12 · The Neumann certificate, its attainment, and the mask amplification — phase 0 · prereq J-L0 · independent-of J-L1 to J-L11 · parallel-safe *build* pointwise in the style of `Contraction.lean`, no norm instance: `resolvent_sup_bound`, `isUnit_one_sub_smul`,

neumann_truncation_bound, neumann_tail_attained, mask_amplification [S], refusal nonstochastic_breaks_bound at a **convergent** $\gamma_0 = 3/5$ with rows $3/2$. *prove* [S] all; the K -fold iterate of **weighted_contraction** is the only induction. *measure* equality to the printed digits at $(\gamma, K) = (0.6, 2)$: 0.540000; $(0.7, 4)$: 0.5602333; $(0.7, 16)$: 7.754350×10^{-3} (RUN[P]); the convergent plant reads 7.2900 against bound 0.5400; the coordinator's $\gamma = 0.7$ plant (119.37 against 1.143) is a divergent partial sum ($\rho(\gamma P) = 1.05$) and is labelled so; the mask witness at $s = 3$, $\gamma = 0.9$, $\varepsilon = 0.1$ reads 0.5263157894736843 against the naive 0.1 (RUN[J]). *PASS* builds, and the paper prints $\delta \|V\|_\infty$, never a bare δ , with $1/(1 - \hat{\gamma})$ beside it. *KILL* RowStochastic (P^k) is missing as a lemma: plus one evening; if the amplification constant does not close, ship the $s = 3$ witness as [M] and the general bound as [D]. If the card dies entirely, every δ is printed from the textbook ([83], [55]) with the RUN plant and no arena cell is lost. *price* 0 GPU-s · *mechanism* V-3, V-10, V-24, V-17, L-CERT · *deliverable* lean/CEQ/Shape/Neumann.lean · *evenings* 2

J-L13 · The resolvent is one forward substitution, and Π_γ is a mixture — phase 0 · prereq J-L3, J-L12 · independent-of J-L5 to J-L9, J-L11 · parallel-safe *build* fwdSub, resolvent_is_triangular_solve, mixing_matrix_rowStochastic, refusal-shaped bare_read_row_sum = $1/(1 - \gamma)$. *prove* [S] all three; non-negativity of the inverse follows from $(1 - \gamma P)^{-1} = \text{occupancy}(\gamma P, N) + (1 - \gamma P)^{-1}(\gamma P)^N$ with both terms non-negative for every N . *measure* $1.7763568394002505 \times 10^{-15}$ solve against dense inverse (RUN[coord]); Π_γ row sums 1.0000000000000000 with and without absorbing rows (RUN[F]); bare row sums $2.5 = 1/(1 - 0.6)$ (RUN[J], RUN[M]); support of Π_γ equals support of P on a dense causal softmax (RUN[J]). *PASS* builds. *KILL* the well-founded recursion does not elaborate: state the characterisation $M_{ii}z_i = v_i - \sum_{j < i} M_{ij}z_j$ instead ([M]) and drop the executable definition; the cost law then stands on the RUN. *price* 0 GPU-s · *mechanism* M-8, V-17, V-23 · *deliverable* lean/CEQ/Shape/Triangular.lean, Shape/Neumann.lean · *evenings* 2

J-L14 · The suffix re-solve and the rank-one row clamp — phase 0 · prereq J-L4, J-L8 · independent-of J-L5 to J-L7, J-L9, J-L11 to J-L13, J-L15 to J-L17 · parallel-safe *build* suffix_resolve_eq_full, sherman_morrison_row (docstring: a row clamp only; a token rewrite is rank $s - i$), sm_denominator_pos, refusal dense_P_breaks_suffix_resolve. *prove* [S]; the algebra is [111] and [51] [V] re-proved on Fin n; the denominator equals $(1 - \gamma p'_{ii})/(1 - \gamma P_{ii}) > 0$ for every $\gamma < 1$. *measure* suffix against full 5.6×10^{-17} on the non-nilpotent corner (RUN[F]); closed form against re-solve $1.2339847026143769 \times 10^{-15}$ with denominator 1.9 (RUN[I]); 0.511918 (RUN[F]); m columns cost $m/d = 0.5$ of one solve at $(8, 16)$, SPLIT band $(0.5 \times, 8 \times)$ filed. *PASS* builds. *KILL* the Matrix.of row update does not simp: restate with Matrix.vecMulVec, same content; if abandoned, the move price is a DERIVED band with the RUN identities. *price* 0 GPU-s · *mechanism* M-8, P-8, V-24, D-5 · *deliverable* lean/CEQ/Shape/Displacement.lean · *evenings* 1

J-L15 · The reach-avoid conservation row and the degeneracy lemma — phase 0 · prereq J-L10 · independent-of J-L5 to J-L9, J-L11 to J-L14 · parallel-safe *build* reach_avoid_sum_one (hypothesis: the sets exhaust the absorbing states), no_goal_no_sink_forces_max_ge_inv_K, isUnit_one_sub_of_perron (b-chain, the oracle's non-causal Q only), refusal unit_needs_rho_lt_one. *prove* [S]; [66] and [47] Thm 11.6 [U] own the identity $B = NR$ — the card machine-checks the packaging. *measure* the four reads sum on T to $[1.0000000000000000, 1.0000000000000001]$ (RUN[P]); with BOS inside a constraint set and no goal, the two committors sum to 1.0000000000000000 and $\min_T \max_k q^{(k)} = 0.548718 \geq 1/2$ — the lemma's instance. The coordinator's 0.605 was computed with the goal counted and is not this lemma's instance. *PASS* builds. *KILL* Finset submatrix indexing does not close in one file: state on the full Fin n with Absorbing rows, plus one evening; if abandoned, the row is printed with its residual on every batch (V-23) and the lemma is DERIVED. *price* 0 GPU-s · *mechanism* V-12, V-8, V-23, V-25 · *deliverable* lean/CEQ/Shape/Committor.lean · *evenings* 2

J-L16 · The masked-sink finite sum, and BED-1's solve as the committor

read — phase 0 · prereq J-L12, J-L10 · independent-of J-L5 to J-L9, J-L11, J-L14 · parallel-safe *build* `masked_sink_finite_sum` ($V_0 = 0$), refusal `goal_channel_breaks_finite_sum`, `bed1_committor_eq` unpacking `READ ceq/beds/bed_1.py:173-174, :188-198`. *prove* [S] both. *measure* the masked route reads 4.4×10^{-16} on a value-zero channel against 1.99×10^{-7} on $\mathbf{1}_{\mathcal{A}_0}$ with $0 \in \mathcal{A}_0$ (RUN[M]); BED-1’s harmonic residual $1.0408340855860843 \times 10^{-17}$ at $T = 0.25$, jitter 0.05, seed 11. *PASS* builds. *KILL* the unpacking needs the bed’s kernel as a Lean object: state it for an abstract diagonally dominant L ([S]); if abandoned, the paper keeps the diagonal and Proposition 4 everywhere, which is already the default. *price* 0 GPU-s · *mechanism* V-25, V-4, D-2 · *deliverable* `lean/CEQ/Shape/Committor.lean` · *evenings* 1

J-L17 · Register the two deferred targets, statements only — phase 0 in place, phase 5 in grade · prereq J-L0 · independent-of all · parallel-safe *build* `Shape/Deferred.lean`, **not** imported by `CEQ.lean` so the root build stays sorry-free: `hitting_time_transform` and `f1_cantelli_union`, both [D] — this Mathlib has no finite-chain hitting-time object (RUN[P] `grep`, zero files; the stopping-time machinery at `Mathlib/Probability/Martingale/Basic.lean:508-510` is for filtrations of a measure space). *prove* nothing this round; the content is J-D3. *measure* read against $E[\gamma^{\tau-1}]$ to 8.33×10^{-17} , Monte-Carlo-free, by the finite-sum identity (RUN[J], RUN[M]). *PASS* the file’s statements type-check. *KILL* a statement needs an object the tree cannot name: write it in J-D3’s finite-sum form; what dies is the Lean statement, never the derivation. *price* 0 GPU-s · *mechanism* P-11, P-4 · *deliverable* `lean/CEQ/Shape/Deferred.lean` · *evenings* 1

J-L18 * · The build gate, the axiom print and the domain census — phase 0 gate · prereq every [M] card it certifies (J-L1 to J-L10 at minimum) · independent-of the [S] cards it does not list · **not** parallel-safe · aliases M-0.1, M-0.4, R-01, V-2 *build* (i) lake build from `lean/` with the new imports in `CEQ.lean`; (ii) a `#print` axioms block per declaration in the pattern of `V16Domain.lean:586-594`; (iii) a domain census file in the `bedM_overlap_*` style (`READ V16Domain.lean:302-310`) giving, per theorem, the corpus and the fraction of it satisfying the hypotheses: `gamma_zero_is_softmax` 3/3; `bos_row_is_absorbing` 3/3 at $\beta = 1$, 0/3 at $\beta = 0$; `segmentation_blockdiag` 3/3 on BED-M, 0/N on BED-S’s softmax corner (silent); `softmax_corner_not_nilpotent` 0/3 on BED-M, N/N on BED-S; `lowerTriangular_ne_symmSupport` 1/1 on BED-1; `neumann_*` N/N at $\hat{\beta} = 1$, 0/N at $\hat{\beta} \neq 1$ (rows sum 1.31 to 10.29, `READ V16_ARM_SMPRIME.md:28-32`). *prove* the census rows by `decide` where the value set is finite, by a printed RUN where it is a drawn batch. *measure* lake build exit code; `grep -c '^theorem\|^lemma'` on the new files beside the record’s 169 over 12 files; the axiom list per declaration. *PASS* exit 0, zero sorry, zero `sorryAx`, the three-axiom set, every [M] card green; a census line of 0 per cent demotes that theorem to decoration on that bed (L-DOM). *KILL* (K-K) any [M] target not building is [S] and is never cited as proved; what dies is the paper’s “machine-checked” sentence for component (k), narrowed to the rows that build — and **no arm trains until this reads green**. *price* 0 GPU-s (CPU minutes) · *mechanism* P-11, L-LEAN, L-DOM, V-25, V-16 · *deliverable* `docs/LEAN_SHAPE_CENSUS.md`, `lean/CEQ/Shape/Census.lean` · *evenings* 1

J-D1 · The EMC remark, with both halves and the ruling — phase 0 · prereq J-L0’s notation · independent-of every card · parallel-safe *build* a remark, never a proposition (`judge/bind_ledger.md` B-EMC is KILL-deleted). The half that holds: for $\gamma < 1$ and any row-stochastic P' , cyclic or not, $x \mapsto V + \gamma P'x$ is a γ -contraction in $\|\cdot\|_\infty$, so the post-intervention fixed point is unique and “settle then intervene” equals “intervene then solve”. The half where Dash applies: [35] Thm 1 concerns a reduced model whose equilibrated form hides feedback; an explicit linear fixed point has no such gap; the case where it returns is P' recomputed from z , a DEQ ([7]), out of scope. What survives as a bind is J-L4’s triangularity plant. The external argument for the re-solve is [85] and [101] (mechanism [U]). *prove* uniqueness in four lines from the contraction (DERIVED). *measure* 8.9×10^{-16} on a dense feedback P (RUN[J]), 1.33×10^{-15} (RUN[I]), 4.4×10^{-16} (RUN[F]). *PASS* the remark carries the three numbers and no proposition

number. *KILL* a planted feedback instance reading above 0.1 on the shape's class — impossible for $\gamma < 1$; if it ever reads so, the algorithm is wrong, as the design's “violation” was (a suffix-only solve on a non-triangular P). *price* 0 GPU-s · *mechanism* V-24, D-7, V-3 · *deliverable* docs/CEQ_SHAPE.md §4.2 Remark · *evenings* 1

J-D2 · Derive reach-avoid sum-to-one with a sink, and the degeneracy lemma — phase 0 · prereq J-D1's notation; states J-L15 · independent-of all Lean cards · parallel-safe *build* Proposition 8's clauses in full: (i) with $\mathcal{A} = \mathcal{A}_{\text{sink}} \sqcup \mathcal{A}_0 \sqcup \bigsqcup_k \mathcal{A}_k$ exhausting the absorbing states and $\rho(Q) < 1$, $\sum_{\bullet} q^{(\bullet)} = \mathbf{1}$ on T , because $\sum_{\bullet} R_{\bullet}\mathbf{1} = \mathbf{1} - Q\mathbf{1}$ and $N(I - Q)\mathbf{1} = \mathbf{1}$; (ii) if the constraint sets alone exhaust absorption, $\max_k q^{(k)} \geq 1/K$ by pigeonhole; (iii) the discounted reads sum below one, so the delay share $1 - \sum_k E[\gamma^{\tau-1}\mathbf{1}_k]$ is printed beside every safest-move reading. Owners cited before the object: [66], [47] Thm 11.6 [U], [114], [119], [44]; [84] carried to Limits. *prove* DERIVED, steps as above. *measure* the read-form sums [0.227, 0.567] at $\gamma = 0.6$ with the state-form/read-form ratio exactly γ (RUN[I]); the conservation row [0.9999999999999993, 1.0]. *PASS* census lines X-5, X-6, X-7 read the printed row on every batch. *KILL* a batch failing by more than 10^{-12} means a set was left undeclared, and the solve must raise rather than return (V-16). *price* 0 GPU-s · *mechanism* V-12, V-8, V-23, V-25 · *deliverable* docs/CEQ_SHAPE.md §4.2 Prop. 8 · *evenings* 1

J-D3 · Derive the $E[\gamma^{\tau}]$ identity without a probability space — phase 0 · prereq J-L10's statement · independent-of all · parallel-safe *build* Proposition 3 in a finite-sum form the Lean tree can eventually carry: for $i \in T$ and $V = \mathbf{1}_{\mathcal{A}_k}$, $z_i = \sum_{t \geq 0} \gamma^t (P^t \mathbf{1}_{\mathcal{A}_k})_i$; with $\mathcal{A}_k$ absorbing, $(P^t \mathbf{1}_{\mathcal{A}_k})_i = \sum_{u \leq t} f_i(u)$ with $f_i(u) \geq 0$ the first-passage mass at step u ; exchanging sums gives $(1 - \gamma)z_i = \sum_u \gamma^u f_i(u)$, the read is $O_i = \sum_u \gamma^{u-1} f_i(u)$ by first-step analysis, and $\lim_{\gamma \uparrow 1} O_i = q_i^{(k)}$ by monotone convergence of a non-negative series. The gap line: $q_i - E_i[\gamma^{\tau}\mathbf{1}_k] = E_i[(1 - \gamma^{\tau})\mathbf{1}_k] \geq (1 - \gamma)q_i$, small where q_i is. At $i \in \mathcal{A}_k$ the read is 1, not γ^{-1} . *prove* DERIVED; this overrules *refute_theory_math* item 8's universal $\geq 1 - \gamma$ gap. *measure* $\max_T = 1.42 \times 10^{-7}$ at $\gamma = 1 - 10^{-6}$ on channel {15} (RUN[J]); read against $E[\gamma^{\tau-1}]$ to 8.33×10^{-17} ; BED-1's real sets 0.0. *PASS* the derivation carries the $(1 - \gamma)q_i$ line and both RUNs. *KILL* none; the identity is classical and its owners are cited before the object is named. *price* 0 GPU-s · *mechanism* D-2, V-25, P-10 · *deliverable* docs/CEQ_SHAPE.md §4.2 Prop. 3 · *evenings* 1

J-D4 · The DET-class statement, and what “depth-1” means — phase 0 · prereq none · independent-of all · parallel-safe *build* the displayed statement: the exact committor solves $(I - Q)q = R\mathbf{1}$, the exact z is $(I - \gamma P)^{-1}V$; inversion and iterated product are DET-complete with $\text{NL} \subseteq \text{DET} \subseteq \text{NC}^2$ ([31] [V]); a triangular inverse is DET-hard (J-L11's *subdiag_pow_entry* is that embedding with scalar blocks); rational linear **equalities** are in DET and the P-complete problem is linear **inequalities**. So the shape relocates the log-depth from the parameter stack into the solve: $O(s)$ sequential rounds by substitution, or recursive block inversion at $O(\log^2 s)$ circuit depth and $O(\log s)$ matmul rounds with $O(s^3)$ -class work — NOT MEASURED, needs a parallel-prefix kernel timing. “Depth-1” means one attention parameter set, (P, γ) ; the circuit-depth reading is disclaimed in the same paragraph. Beside it the skyline facts and their vacuity inequalities. *prove* DERIVED. *measure* $64^{1/16} = 1.2968$ against $Hdp = 512$; $512 \geq 64$; $384 < 544$. *PASS* every obstruction sentence carries its theorem number, model class and inequality (census X-15). *KILL* (K-10) any “cannot” sentence outside the skyline table is struck at assembly. *price* 0 GPU-s · *mechanism* P-8, V-17, P-10, V-25, P-3 · *deliverable* docs/CEQ_SHAPE.md §5.3 · *evenings* 1

J-D6 · Restate the Cheeger line on a symmetrised surrogate — phase 0 · prereq none · independent-of all · parallel-safe *build* MATHEMATICS.md:455 cites [73] Thm 13.10, whose bound holds for a **reversible** P ; the shape's causal P is not reversible (a lower-triangular P with $P_{i0} > 0$ has $P_{0i} = 0$), so the line is V-25 as written. The restatement: the honest surrogate is the transient block's lazy symmetrisation $Q_s = (Q + Q^{\top})/2$ on T , whose Cheeger bound controls the surrogate's gap and **not** $\rho(Q)$; the quantity the hop ladder truncates is $\rho(Q_{\text{env}})$, read off the ladder

and never predicted from a spectrum (0.9964 against 0.9985 on LargestJoin_S2Rips_1024, READ PRIOR_ART.md:703-707); on the DAG substrate Q is lower-triangular with $\rho(Q) = \max_T P_{ii}$, so there is no bottleneck spectrum to bound and the Cheeger line applies to the oracle cross-check chains (BED-1, E4-prime) only. *prove* DERIVED; the reversibility hypothesis is quoted once, under fifteen words. *measure* the E4-prime bottleneck reading 1372.50 is that bed's own constant, marked V-22, never a BED-S bar (READ scale/e4_harmonic.py:247-263). *PASS* the restated line names the surrogate, the object bounded and the bed it applies to. *KILL* any $g \leq 2\phi$ on a causal P without a stated symmetrisation is struck (K-10 type). *price* 0 GPU-s · *mechanism* V-25, P-10, V-17, V-22 · *deliverable* docs/CEQ_SHAPE.md §4.4 and a MATHEMATICS.md §7 correction pointer · *evenings* 1

J-D7 · The mask-amplification and certificate-units derivation — phase 0 · prereq J-L12's statement · independent-of all · parallel-safe *build* Proposition 4(ii) in full with the $s = 3$ closed-form instance ($\gamma = 0.9$, $\varepsilon = 0.1$, $V = e_0$, P_{10} dropped and row 1 renormalised): $\|O_{\text{full}} - O_{\text{mask}}\|_{\infty} = 0.5263157894736843$ against $\varepsilon/(1 - \gamma) = 1.000$, with the naive $\varepsilon = 0.1$ violated by $5.26\times$; the rule that every F1 union bound carries $1/(1 - \gamma)$; the units rule ($\delta \|V\|_{\infty}$ with $\|V\|_{\infty}$ printed — the $\approx 7.6 \times 10^{-5}$ of sec_cost.md §4.x.3 carried $\|V\|_{\infty} \approx 5$ [ASSUMED]); the census guard $\|\hat{P}\|_{\infty} > 1$ implies no certificate; and the sentence that the exact solve on an F1 mask is refused for fill-in. *prove* DERIVED; mask_amplification [S] is J-L12's. *measure* the $s = 3$ instance; a random $s = 32$ draw did **not** violate the naive bound (RUN[J]), which is why the plant is the closed form and never a random draw (V-10). *PASS* written with the instance. *KILL* none this round; the row stays dormant until a mask ships (D-4). *price* 0 GPU-s · *mechanism* V-10, V-17, L-CERT, V-25 · *deliverable* docs/CEQ_SHAPE.md §4.2 Prop. 4(ii) · *evenings* 1

S-01 * · Extend the identity manifest so a cell cannot be journalled without its boundary condition — phase 0 · prereq none · independent-of every card · parallel-safe · aliases M-0.2, R-02, N-04 *build* the field list of judge/sec_apparatus.md §A.10 for the shape lane, extending CONFIG_FIELDS (READ scale/identity_manifest.py:67-71, which carries beta but not qk, g, gamma): beta, qk, g; gamma (init, final, Λ , verdict); diag_convention in {kept, removed}; committor_route in {dirichlet_gamma1, discounted_gamma}; route in {solve_triangular, neumann_K, segmented, csr} with K , chunk C , tile B ; boundary_sets, goal_set, sink_set as sorted lists with a sha256 each; K , m , t_{star} ; S , D , d_{model} , n_{train} , n_{eval} , steps; seed, rng_plan; device, threads, dtype, torch_version, cublas_workspace, deterministic_regime; instrument_hash, manifest_hash; producer_cmd; delta_vec, delta_bare, one_over_lmg, V_{inf} ; the census block; floor_exact, floor_zeroinfo, floor_fano_k1, ceiling_hop_k; sky_depth, sky_width, sky_cot, chacal_gamma, chacal_diag, params_per_arm; void_contrasts; supersedes. Invariant: a missing declared field is a **refusal**, not an absent entry — the record's reason for reporting (stored manifests would be invalidated, READ :146-148) does not apply to a lane with no stored manifests. *prove* nothing; an instrument. *measure* the drift plant: flip each of beta / qk / g / gamma / boundary_sets by one unit and read manifest_hash; the pattern is test_moving_one_citation_by_one_line_fires_both_binds (READ V20_R15_WING_MANIFEST.md:148-152). *PASS* 5/5 flips move the hash and an omitted sink_set raises. *KILL* any flip leaving the hash unchanged, or an omission returning absent: the lane has no identity, no cell may be journalled, and every certificate printed under that field is unattributable (L-2's recurrence: 0 of 24 attributable, READ V20_R15_LEAP_LEDGER.md:23). *price* 0 GPU-s · *mechanism* L-2, P-1, V-16, M-10, M-16 · *deliverable* docs/apparatus/MANIFEST_SHAPE_LANE.md · *evenings* 1

S-02 · The identity-script rule: raise on a missing key, never default — phase 0 · prereq none · independent-of every card · parallel-safe *build* a written rule for every identity check in the lane: bed keys are read by name and a missing key raises; no silent default. The filed instance: a scratch identity script asked bed_1.build for basin_A/basin_B, which it does not return (its keys are A, B, READ ceq/beds/bed_1.py:167), defaulted silently, and printed 0.858

for an identity that reads 0.0 on the real sets. *prove* nothing. *measure* must-fire: a misspelt key raises. Must-not-fire: the real keys return $\max |q - (I - Q)^{-1}R1| = 0.0$ (RUN[coord]). *PASS* an exception on the misspelt key and 0.0 on the real one. *KILL* a number returned on a misspelt key: every identity number in the lane is suspect until the rule is enforced. *price* 0 GPU-s · *mechanism* V-16, V-3 · *deliverable* docs/apparatus/IDENTITY_SCRIPT_RULES.md · *evenings* 0.5

S-03 · Print the obstruction vacuity inequalities beside every obstruction citation — phase 0 · prereq none · independent-of every card · parallel-safe · aliases M-0.3 (obstruction half), R-03 *build* a one-page census of the three unconditional bounds as predicates on the record’s geometry ($s = 64, d = 16, h = 1, p = 32$): $hmp = 512 \geq 64$ ([104] Thm 1), $n \log_2 n = 384 < 544 = H(d+1)p$ ([92] Thm 1), $64^{1/16} = 1.2968$ against $Hdp = 512$ ([25] Thm 1.1) — all vacuous here; the conditional bound ([105] Cor. 4.3, the one-versus-two-cycle conjecture) named conditional in the same sentence. *prove* DERIVED arithmetic. *measure* the fraction of drawn BED-S cells whose geometry satisfies each hypothesis. *PASS* every obstruction sentence carries its line and its [V]/[U] mark; 0 per cent admitted means the sentence is stated as asymptotic lineage only. *KILL* (K-10) an obstruction stated without its line is struck at assembly. *price* 0 GPU-s · *mechanism* V-25, P-10, P-3 · *deliverable* docs/apparatus/OBSTRUCTION_CENSUS.md · *evenings* 0.5

M-0.3 · Freeze the price ledger — phase 0 · prereq none · independent-of every card · parallel-safe *build* docs/PRICE_LEDGER.md carrying §5.7’s basis with provenance: the two fitted throughput laws with their R-squared and their fitting range, the solve increment as a per-op floor with the $2.0\times$ to $6.6\times$ dispatch gap beside it, the unit prices, the residency model with its [ASSUMED] operator-constant band, and the rule that no price crosses the device boundary (V-22). *prove* nothing. *measure* nothing; the ledger is read, not computed. *PASS* every price carries a tag ([MEASURED], [FITTED] with R-squared, [RUN], [ASSUMED] with its reason, DERIVED with its arithmetic) and a producer. *KILL* any price without a tag or a producer is struck; a price quoted at $s > 64$ before S-66 runs is struck (K-9). *price* 0 GPU-s · *mechanism* P-1, P-8, M-8, V-22 · *deliverable* docs/PRICE_LEDGER.md · *evenings* 1

N-03 · The kernel-path must-fire battery, as a test specification — phase 0 · prereq none · independent-of every card · parallel-safe *build* one test-module specification, CPU float64, each check with its honest half, its planted negative and the $O(1)$ failure the plant must produce: (1) parity `torch.equal(0(0), PV)` against the lane’s own `softmaxAttn`, rejection 2.3002850040264393 at $\gamma = 0.5$; (2) `solve_triangular` against the dense inverse at $\leq 1.8 \times 10^{-15}$, plant a dense M passed with `upper=False`; (3) the certificate in vector units with $\|V\|_\infty$ printed, plant rows scaled to 1.5 at $\gamma = 0.6$, $K = 2$ reading 7.29 against 0.54 (convergent); (4) segmentation zeros by `torch.equal` on a masked column, plant a -30 logit that leaves the block non-zero; (5) CSR fill-in: a schedule dropping tile (k, l) while keeping $(k, m), (m, l)$ makes the exact solve’s (k, l) block non-zero and the exact route is refused; (6) the three-outcome determinism table is read from N-01, never re-run here; (7) a missing key raises (S-02). *prove* nothing; the theorems behind (1), (2), (4) are cited only once they build (L-LEAN). *measure* the failure set, never the pass count (`sec_measured.md` M.9). *PASS* every honest half passes and every plant fails at $O(1)$ with counts printed. *KILL* a plant that passes strikes that check from the battery and the paper’s corresponding must-fire sentence. *price* 0 GPU-s · *mechanism* V-24, V-3, V-2, V-16, V-17 · *deliverable* docs/CEQ_KERNEL_PATH_BATTERY.md · *evenings* 1

R-15 · Recount parameters per arm; fix the widths of the unmatched skylines — phase 0 · prereq none · independent-of every card · parallel-safe *build* a per-arm parameter count printed in every table header (Ruling 3); the 4,769 of `READ CEQ_V20_R15_CONTRACT.md:119` recounted, never assumed, once a $[m, K+2]$ head or a vector readout is attached; the wide constant-depth skyline’s width fixed at n_{nodes} ([139] [V]) with the sentence that at $d_{\text{model}} = 16 < s = 64$ the matched-parameter instance does not exist; the chain-of-thought decoder’s step count fixed at t^* ([80] [V]). *prove* nothing. *measure* the count per arm and the residual mismatch. *PASS* every matched arm within 0.032 per cent of the reference count (`READ MODEL_CARD.md:76-79`). *KILL* a “matched” contrast above that residual is moved to the skyline column and never credited;

an arm re-architected to close the gap is refused (Ruling 3). *price* 0 GPU-s · *mechanism* M-8, MISTAKES.md D-1, R-SKY · *deliverable* docs/plan/PARAM_TABLE.md · *evenings* 1

R-17 · Freeze every guard partition before any itinerary is read — phase 0 · prereq none · independent-of every card · parallel-safe *build* the $q = \frac{1}{2}$ isocommittor guards (READ ceq/beds/bed_1.py:5-8, :382) and any BED-S basin partition fixed from the oracle’s committor **before** any itinerary is scored, with the partition’s sha256 in the manifest; the Pesin-deficit reading kept as the diagnostic it is (the record’s contract guards had deficit 0.009654, 11.9 per cent of h ; the state partition 0.000233, READ V15_BED1.md:248-272) and never used to select guards by argmin. *prove* nothing. *measure* the deficit of the frozen partition; the guard-crossing balance (the record’s 21 net on 65,481 events). *PASS* the partition hash precedes the first itinerary record in the journal’s order. *KILL* a partition whose hash postdates an itinerary it scores: every itinerary statistic on it is struck as a threshold refitted to the data it judges (M-2). *price* 0 GPU-s · *mechanism* M-2, M-18, M-19 · *deliverable* docs/plan/GUARD_FREEZE.md · *evenings* 1

S-61 * · The γ -pinning likelihood-ratio instrument, with the boundary null — phase 0 · prereq none · independent-of every card · parallel-safe · aliases J-D9, M-4.3 (spec half), R-11, V-8, ledger J-1 *build* the specification: $\Lambda = 2[\text{LL}(\hat{\gamma}) - \text{LL}(\gamma \equiv 0)]$ on a declared held-out set of $n_{\text{eval}} = 4096$ (Ruling 10-prime restated for γ , which was stated for β); with γ trained on $[0, 1)$ the null sits on the parameter boundary, so Λ is asymptotically $\frac{1}{2}\chi_0^2 + \frac{1}{2}\chi_1^2$ whose 95 per cent point is the χ_1^2 0.90 quantile 2.705543, **not** 3.841459; PINNED at $\Lambda \leq 2.7055$, MOVED at $\Lambda > \ln n_{\text{eval}} = 8.318$, the interval verdict between; the seed rule “MOVED on at least 6 of 8” and the SPLIT band 3/8 to 5/8 declared here; the ablation $(I - \hat{\gamma}\hat{P})^{-1} \rightarrow I$ at trained weights must move NRMSE by more than one seed sd; the mirror kill $\hat{\gamma} > 0.99$ on at least 6 of 8 prints $1/(1 - \hat{\gamma})$ beside every δ ; the $|\hat{\gamma}| < 0.05$ rule is **deleted** (two verdicts for one cell, M-20). The citation for the boundary null is **owed** — not in references.bib, not fetched — and the paper may not cite it by name until it is. *prove* nothing; the quantile is RUN[J], RUN[MARS], RUN[MERCURY] to the printed digits. *measure* the instrument’s own plants on synthetic data: a generator at $\gamma_{\text{env}} = 0$ must read PINNED on 8/8 seeds and one at $\gamma_{\text{env}} = 0.6$ must read MOVED on 8/8. *PASS* both plant counts 8/8. *KILL* either fails: Bet C cannot be scored and is VOID until the instrument is repaired; a threshold chosen after Λ is seen is M-2 and strikes the verdict. *price* 0 GPU-s for the spec, seconds for the plants · *mechanism* Ruling 10-prime, V-9, M-2, M-20, V-17, V-24 · *deliverable* docs/apparatus/INSTRUMENT_GAMMA_LR.md · *evenings* 1

V-0 · Open the calibration ledger and carry the R11 column into it — phase 0 · prereq none · independent-of every card · parallel-safe *build* a ledger with one row per bet: bet, source in {author, leap}, prediction, counter, deciding number, frozen PASS, frozen KILL, realised, verdict in {HOLDS, COUNTER, SPLIT, VOID}, sign of miss, date, producer command; the R11 nine rows carried in verbatim as the first block (READ V16_CALIBRATION.md:40-48), each with its sign; the five D-CALIB rules printed at its head. *prove* nothing; an instrument. *measure* the running one-sided sign test after every scored row; the column starts at 7/8 optimistic, $p = 0.0352$, Wilson 95 per cent [0.5291, 0.9776]. *PASS* the ledger exists with the nine rows and the five rules. *KILL* a row filed without its counter is refused at the file, never discounted (D-CALIB-2); what dies is the row, not the bet. *price* 0 GPU-s · *mechanism* D-7, M-2, P-1 · *deliverable* docs/CALIBRATION.md · *evenings* 1

PHASE 1 — BED-S: specification, census, admission

Goal. Build the one bed on which the shape’s three capabilities have a rejection region, and decide from a printed census whether it is admissible — before a GPU-second is spent. The record’s dominant failure is a bed that dies at construction (D-4, M-3, V-8); this phase makes that death free.

The gate (D-4 admission). The author reads one census block and admits or reroutes. Admit requires, on 512 draws at the design point $t^* = 8$, $m = 8$, $K = 2$, $s = 64$: row sums within 10^{-12} of one including absorbing rows; $0 \in \mathcal{A}_{\text{sink}}$ on 100 per cent; $\rho(\hat{Q}) = \max_{i \in T} \hat{P}_{ii} < 1$ on 100 per cent; every boundary set before the query and the query in T on 100 per cent; every per-coordinate label sd over the **admitted query region** above 0.05; every argmin class frequency in $(0.05, 0.95)$; rule-disagreement fraction above zero; the sink share printed apart from the goal share; the discard count printed. Any sd = 0, any class outside the band, disagreement zero, or t^* unplaceable inside the causal window kills admission (K-E2, K-P): no BED-S reading exists, the reroute is the jittered **bed_1** landscape at $K = 2$ with B as goal used as an **oracle cross-check only**, and N is repriced from whatever realised sd the first admitted batch shows. Price of the whole phase: 0 GPU-s, except S-14's single zero-step forward.

S-10 · Register the placement convention: BOS is a value-zero sink set, boundary sets precede the query — phase 1 · prereq none · independent-of S-01 to S-03, J-, *M-0.3 · parallel-safe* build* the two construction facts as a written convention. **F1**: row 0 of any causal softmax at $\beta = 1$ is e_0 , so position 0 is absorbing whether or not it is declared; undeclared, the transient block is singular. **F2**: walks descend, so a boundary set after the query is unreachable. BED-S therefore declares $\mathcal{A}_{\text{sink}} = \{0\}$ with value 0 on every indicator channel — **not** a goal member, since with $0 \in \mathcal{A}_0$ “reach the goal” collapses to “descend to 0 without hitting a constraint” — and places the sink, $\mathcal{A}_0$ and $\mathcal{A}_1 \dots \mathcal{A}_K$ all before the query, the query in T . *prove* both facts are theorem instances once J-L2 and J-L4 build; until then they are RUN. *measure* with BOS undeclared, $\rho(Q) = 1.000000$ and $\det(I - Q) = 0.0$ — the solve must raise; with $\mathcal{A}_{\text{sink}} = \{0\}$ declared, $\rho(Q) = 0.692660$, the conservation row holds to 4.4×10^{-16} and the sink share reads $[0.362, 1.000]$ on that draw; a set at position 20 read from query 12 gives $q = 0.0$ exactly (discard), from query 25 gives $q = 0.1033$ (RUN[SATURN]). *PASS* the four readings reproduce on the bed's own generator at $s = 64$. *KILL* a declared-sink draw with $\rho(Q) = 1$: the generator has a second undeclared absorbing row (a zero gate at c , J-L7) and census line 14 must be re-read before admission. *price* 0 GPU-s · *mechanism* V-25, V-12, V-8, D-3 · *deliverable* docs/beds/BED_S_PLACEMENT.md · *evenings* 0.5

S-11 · Write the BED-S generator specification — phase 1 · prereq S-10 · independent-of S-01 to S-03, S-50 to S-53, every J- card · parallel-safe · aliases M-1.1, R-04 (spec half), V-3 (spec half), ledger S-1 *build* the environment is a random directed graph on the s token positions whose edges point to **earlier** positions — a DAG in token order, so a walk from the query descends as the arm's causal $\hat{P}$ does — with self-loops only on declared absorbing positions and P_{env} row-stochastic including them. Node tokens carry the out-adjacency as a multi-hot feature, the membership flags of the absorbing sets, and **nothing derived from any solve**; the transition matrix is a deterministic function of the edge tokens, so leak clause (a) is dropped for BED-S and clause (c) is kept as the kill. Candidate moves are m **query-side row clamps** $P_{\text{env}}[v_a, :] \leftarrow e_{u_a}$, each rank one with $u^\top \mathbf{1} = 0$ (the graph-surgery precedent, READ scale/negation_scope.py:718-732); a token rewrite is rank $s - i$, a different object priced as a suffix re-solve. Label per move: the reach-avoid tensor $(q^{(\text{sink})}, q^{(0)}, q^{(1..K)})$ at $\gamma = 1$ on the transient block, the full vector on T retained for the residual; $a^* = \arg \max_a q^{(0)}$, the Chebyshev $\arg \min_a \max_{k \geq 1} q^{(k)}$ and the lexicographic form as **columns**; z^* and $\Delta z(a)$ by the displacement identity on P_{env} at a bed constant γ_{env} . Dials with registered supports: $t^* \in \{2, 8, 32\}$ (a **tolerance** dial read off the hop ladder, never predicted from a spectrum), $K \in \{2, 3, 4\}$ ($K = 1$ is illegal: the committor is constant to 1.11×10^{-14}), $m \in \{4, 8, 16\}$ default 8, $|\mathcal{A}_\bullet| \in \{1, 2\}$, $s \in \{64, 256, 1024, 4096\}$. Oracle: the Dirichlet solve of **bed_1.committor**'s form (READ ceq/beds/bed_1.py:188-198) on P_{env} , never on $\hat{P}$; the Kirchhoff second route covers $K = 2$ single-node sets only, the multi-node extension is NOT MEASURED, needs the grounded Laplacian extended. *prove* nothing here; the identities are Phase 2's. *measure* nothing yet; the spec names every field S-01's manifest consumes and every census line S-12 prints. *PASS* the spec states the oracle runs on P_{env} and names the D-4 registration triple on the same commit. *KILL* any sentence handing the arm P_{env} as an input

channel other than the edge tokens (D-2), or a spec in which $K = 1$ is legal (V-12). *price* 0 GPU-s · *mechanism* D-2, V-25, V-12, V-8, D-3, M-8 · *deliverable* docs/beds/BED_S_SPEC.md · *evenings* 2

S-12 * · Run the BED-S domain census on 512 draws and decide admission — phase 1 · prereq S-11 · independent-of S-01 to S-03, S-50 to S-53, every J- card · parallel-safe · aliases M-1.2, R-04, V-3, ledger S-2, P1.2 *build* the fifteen census lines of `judge/sec_apparatus.md` §A.7 printed per draw batch as one JSON block, before any arm is trained: (1) row stochasticity including absorbing rows; (2) $0 \in \mathcal{A}_{\text{sink}}$ on 100 per cent; (3) $\rho(\hat{Q}) < 1$ as a diagonal read — a V-10 gate, printed, never counted as evidence; (4) every set before the query, query in T ; (5) per-coordinate label sd over the **admitted query region** (the prefix before the first constraint reads $q^{(0)} \approx 1$ and would inflate it), sink share printed apart from $q^{(0)}$; (6) every argmin class frequency in $(0.05, 0.95)$, ties within 10^{-9} discarded and counted; (7) rule-disagreement fraction above zero; (8) the corpus-alone probe of S-15; (9) $I(s_0; a^*) = 0$ by plug-in; (10) $\text{Var}(\Delta z) > 0$ over coordinates $\geq i_{\min}$ with the zero-coordinate fraction printed; (11) the move census (which moves change P against V only; $u^\top \mathbf{1} = 0$ and a non-negative edited row on 100 per cent; the Sherman-Morrison denominator printed as a check); (12) and (13) filled on trained cells; (14) the exact-zero-gate count (expected 0 of N on the softmax corner, so the F0 theorems are silent); (15) S-03’s lines. *nrmse* returns **nan** on a constant label and **nan** $>= 1.0$ is **False** (READ MISTAKES.md:149–153), so the sd print is not optional. *prove* nothing; F1 and F2 appear as the two structural lines. *measure* 512 draws at $t^* = 8$, $m = 8$, $K = 2$, $s = 64$. *PASS* every admitted-region sd above 0.05; every class frequency in $(0.05, 0.95)$; disagreement above zero; lines 1 to 4 at 100 per cent; the discard fraction printed. *KILL* (K-E2, K-P) any sd = 0; any class outside the band; disagreement zero (then only one rule may be named); $\rho(Q) = 1$ on any draw (BOS undeclared: the solve must raise, V-16); t^* unplaceable. The bed is not admitted, **no arm runs**, every card from S-14 to S-65 is VOID, and §5.10 tree C is the paper. *price* 0 GPU-s; the oracle is m factorisations with $K + 2$ right-hand sides, $|T|^3/3 = 5.76 \times 10^8$ MACs at $|T| = 1200$, under a second per draw in float64 (DERIVED) · *mechanism* V-8, V-12, V-25, V-10, D-3, D-4, M-3 · *deliverable* results/bed_s_census.json1 and docs/beds/BED_S_CENSUS.md (the admission decision, signed by the author) · *evenings* 1

S-13 · Register the VOID-contrast list and the identification metric before any cell — phase 1 · prereq S-11 · independent-of S-12 · parallel-safe · aliases M-1.3, R-16, ledger S-3 *build* because the arm’s operator class contains the oracle’s chain, **shape - softmax** and **shape - skyline** on the committor head are reproduction-versus-non-reproduction contrasts and are **VOID as capability numbers**. Creditable: **shape - ChaCAL-diag**, **shape - ChaCAL-published**, **shape - ChaCAL-with-sink-token**, **shape - InfSA-Neumann-16**, **shape - 0-hop MLP** and **shape - 1-hop softmax** on the argmin, and $\|\hat{P} - P_{\text{env}}\|_\infty$ per seed as a **learnability** reading. The model is E1’s registration comment (READ scale/negation_scope.py:1033–1041). *prove* J-L9 is the D-2 separation that fails on an undirected substrate and holds on the DAG. *measure* (at the arena) $\|\hat{P} - P_{\text{env}}\|_\infty$ per seed, sorted. *PASS* the list is in the manifest field **void_contrasts** on every cell, in a commit dated before the first arena cell. *KILL* (K-D2) $\|\hat{P} - P_{\text{env}}\|_\infty < 10^{-3}$ on at least 6 of 8 seeds: the arm copied the environment, every contrast against a non-reproducing arm is a copy-versus-no-copy statement, and the reading is filed as identification, never as capability. *price* 0 GPU-s · *mechanism* D-2, M-7, V-10 · *deliverable* docs/beds/BED_S_VOID_LIST.md · *evenings* 0.5

S-14 · The zero-hop guard: the query carries s_0 only, and the untrained arm reads chance — phase 1 · prereq S-11 · independent-of S-12, S-13 · parallel-safe · aliases M-1.4, ledger P1.3, X-9 *build* s_0 drawn uniformly from T independently of a^* ; the plug-in $I(s_0; a^*)$ printed with its sampling error; the zero-step RED gate applied unchanged — the untrained arm within $1/m$ plus or minus its Clopper-Pearson half-width on the argmin, committor NRMSE $\geq 1 - \text{GATE_TOL}$ with $\text{GATE_TOL} = 10^{-3}$ (READ COSTS.md:137). Precedent: BED-M’s first builder leaked $1/(t^* + 1)$ at zero hops and aborted three of five rungs; $b[s - 1] = 0$ restored it (READ scale/negation_scope.py:399–415). *prove* nothing. *measure* must-fire: a query token carrying $q^{(0)}(a^*)$ pushes the untrained read below $1 - \text{GATE_TOL}$. Must-not-fire: the honest

query reads at chance within the interval. *PASS* both halves. *KILL* the honest query reads below the gate: the label is legible at zero hops, the bed is a copy task, and it is struck before any number. *price* ≤ 1.524 s, one zero-step forward [FITTED] · *mechanism* V-10, D-5, V-24 · *deliverable* results/bed_s_zerohop.jsonl · *evenings* 0.5

S-15 · The leak guard, firing both ways — phase 1 · prereq S-11 · independent-of S-12 to S-14 · parallel-safe · aliases M-1.5, R-05, ledger P1.4, X-8 *build* a corpus-alone linear probe from the token features to q at order 0 on the admitted draws; the planted positive is a move token carrying its own q ; the planted negative is the strictly-local feature set E4-prime uses (degree, ball sizes to radius 5, absorbing-endpoint-in-ball flags, READ scale/e4_harmonic.py:220-244), whose bars are E4-prime’s and are **re-derived for BED-S** before use (V-22). *prove* nothing. *measure* the coefficient of determination on the honest probe and on the planted leak. *PASS* honest below 0.5 and planted at or above 0.99 — the detector fires when the leak is planted, calibrated both ways. *KILL* honest at or above 0.5: the one-read claim is void on this bed (leak kill (c)); planted below 0.99: the detector is blind and no leak verdict is admissible. *price* 0 GPU-s (CPU probes) · *mechanism* M-21, V-24, V-7, V-22, D-2 · *deliverable* results/bed_s_leak.jsonl · *evenings* 1

S-16 · Fix the metrics and the floors before any arm exists — phase 1 · prereq S-11 · independent-of S-12 to S-15 · parallel-safe *build* the committor head scored in the Fisher-Rao coordinate $\varphi(p) = 2 \arcsin \sqrt{p}$ per entry, position-matched NRMSE on φ (mean and max), plus the harmonic residual $r(\hat{q}) = \|(I - Q_{\text{env}})\hat{q} - R_{\text{env}}\mathbf{1}\|_{\infty}$ — a **score, never a loss** (D-2) — printed beside $\sigma_{\min}(I - \gamma P_{\text{env}})$ and $\|I - \gamma P_{\text{env}}\|_{\infty} \leq 1 + \gamma$, its units. Argmin accuracy with Clopper-Pearson ([30]) and McNemar paired on identical draws ([78]). The z and Δz channels by the position-matched per-coordinate NRMSE vector, the field cosine and the magnitude ratio over coordinates $\geq i_{\min}$ with the masked fraction printed (an arm emitting $\Delta \hat{z} \equiv 0$ makes the cosine undefined and is **refused**, not scored, V-16), the sign column kept because softmax owns it (0.807843, READ MATHEMATICS.md:396-406), McNemar on the sign column only and a paired t or Wilcoxon on the cosine with its **own** realised sd. Floors (L-FLOOR): the exact oracle at 0.0 with the hop-ladder ceilings printed per batch; the argmin zero-information floor $1 - \max_a \hat{\pi}(a^*)$, which at uniform $m = 8$ is 0.875; the tight Fano $(\ln m - I - \ln 2) / \ln(m - 1)$, which at $m = 8$ with $I = 0$ reads 0.7124, used **only** where $I(X_{\leq k}; a^*) > 0$ is computed; the weak Fano 0.6667 retired as a floor (it sits 0.208 below chance); **floor_1** nowhere. Refused metrics: pooled first-order Wasserstein (the permuted oracle scores 0.0 against NRMSE 1.421901), Procrustes, position-free optimal transport. *prove* nothing. *measure* nothing; the file is hashed into the prediction file of S-60 before the first arena cell. *PASS* the metric file exists and is hashed before S-62. *KILL* any metric added or removed after the first arena cell (M-2); that cell’s numbers are then not quotable. *price* 0 GPU-s · *mechanism* M-2, V-26, V-17, L-FLOOR, V-16 · *deliverable* docs/apparatus/METRICS_AND_FLOORS.md · *evenings* 1

S-17 · The Δz channel’s own census and the cached-mixture plant classes — phase 1 · prereq S-11 · independent-of S-12 to S-16 · parallel-safe *build* for the z and Δz lane — its **own** journal, because a vector label voids the record’s published softmax baseline (READ MISTAKES.md:701-708) — the intervened position drawn uniformly; $\text{Var}(\Delta z) > 0$ per coordinate $\geq i_{\min}$; per-draw flags **move_changes_P** against **move_changes_V_only** so the cached-mixture control can be scored on the two plant classes separately; the no-change predictor’s error printed beside every Δz number ([118] [V] owns that floor). *prove* nothing. *measure* the two plant classes’ counts. *PASS* both classes non-empty, at least 64 draws each. *KILL* a corpus with no P -changing move: the interventional channel has no rejection region against the cached mixture and card S-33 cannot be scored. *price* 0 GPU-s · *mechanism* V-8, D-5, V-24, M-1 · *deliverable* results/bed_s_dz_census.jsonl · *evenings* 0.5

R-06 · Register the zero-hop, one-hop and hand-rule controls at construction — phase 1 · prereq S-11 · independent-of S-13 to S-17 · parallel-safe *build* the argmin controls (0-hop per-position MLP on the candidate’s context row, 1-hop softmax, majority, random at

1/m, predict-the-mean) plus R-24’s closed-form rule, all registered before any arm trains; the exact zero-information floor from the realised class distribution; $I(X_{\leq 0}; a^*)$ by plug-in over the zero-hop view **including** the move tokens and membership flags. *prove* nothing. *measure* the untrained arm at 1/m within its interval; committor $\text{NRMSE} \geq 1 - \text{GATE_TOL}$ at step 0. *PASS* the untrained arm reads chance and $I(X_{\leq 0}; a^*)$ is printed, with the tight Fano used only where it exceeds zero. *KILL* (K-H2, scored at S-26) the 0-hop MLP within the Clopper-Pearson half-width of the shape at $N = 8$, or McNemar $p > 0.05$ against the 1-hop softmax: BED-S is a third static task and is struck before any number is quoted (C8; the record’s static beds gave an iterating arm nothing to compute toward, READ D1.md:286–291). *price* 0 GPU-s at registration; ≤ 1.524 s per control cell when trained [FITTED] · *mechanism* V-10, D-5, C8 · *deliverable* docs/plan/CONTROLS_REGISTERED.md · *evenings* 1

R-07 * · Census the sink-escape degeneracy of the Chebyshev rule and the exact-zero tie it creates — phase 1 · prereq S-11 · independent-of S-13 to S-17, R-06 · parallel-safe *build* the attack, filed by no office before: with $\mathcal{A}_{\text{sink}} = \{0\}$ declared and moves realised as clamps, a clamp into **any** absorbing position gives $\max_{k \geq 1} q^{(k)} = 0$ exactly, so the Chebyshev column is minimised equally by every move that sends the walk to the sink or to the goal — the safest move under Chebyshev is “fall off the prompt” — and the 10^{-9} tie rule then discards the draw, which empties the bed. Two census lines are added to S-12: (i) the fraction of admitted draws on which $a^\dagger$ has $q^{(\text{sink})}(a^\dagger) > q^{(0)}(a^\dagger)$; (ii) the discard fraction attributable to exact-zero Chebyshev ties. The construction fix is registered **now**, before the number: clamp targets restricted to transient positions ($u_a \in T$) so no move is a direct absorption, and a third column, the conditioned Chebyshev $\arg \min_a \max_{k \geq 1} q^{(k)}(a)/(1 - q^{(\text{sink})}(a))$, printed beside the other two. *prove* DERIVED: for $u \in \mathcal{A}_j$ the clamp gives $q^{(j)}(v) = 1$ and $q^{(k)}(v) = 0$ for $k \neq j$ exactly, so the Chebyshev value of every absorbing-target move is zero unless $j \geq 1$. *measure* lines (i) and (ii) on 512 draws, with and without the restriction. On the judge’s draw at query 20, over the 20 clamp targets below it, the Chebyshev minimum is 0.000, attained by 9 of 20 targets, of which 7 carry sink share above 0.5 and 2 goal share above 0.5 (RUN[M]). *PASS* with the restriction, line (i) below 0.05 and line (ii) below 0.05. *KILL* line (i) at or above 0.5 on the unrestricted bed: the Chebyshev column as written is struck and only the goal rule and the conditioned form are named; line (ii) at or above 0.5: the discard rule slices the bed to nothing and admission is refused until the restriction is in the spec. *price* 0 GPU-s · *mechanism* V-12 in a new guise, V-8, V-10, V-5 · *deliverable* docs/plan/SINK_ESCAPE.md and two columns in results/bed_s_census.jsonl · *evenings* 1

R-08 · Count no-op moves and correct the chance floor to the effective m — phase 1 · prereq S-11 · independent-of R-05 to R-07 · parallel-safe *build* the attack: on a causal P_{env} a clamp at a row above the query cannot change any committor at the query, and a clamp at a row unreachable from s_0 is a no-op for the same reason — F2 applied to **moves**, which no office filed (F2 covers boundary sets). A bed with $m = 8$ of which k are no-ops has effective $m - k$ candidates, so the printed chance floor $1 - 1/m = 0.875$ and the tight Fano 0.7124 are too **low**, and a shape reading above the printed floor may be at chance on the effective set. The census line: per move, $\max_k |q^{(k)}(\text{do } a) - q^{(k)}|$ at the query, a move below 10^{-12} counted a no-op; m_{eff} per draw; both floors recomputed at m_{eff} ; the spec restricted to clamp rows at or below s_0 and reachable from it, with the fraction of draws touched by the restriction printed. *prove* DERIVED from triangularity, and by J-L4’s **later_boundary_unreachable** extended to a clamp row. *measure* m_{eff} on 512 draws; clamps at rows 21, 25, 31 move the query’s four committors by 5.551×10^{-17} , zero to rounding (RUN[M]). *PASS* $m_{\text{eff}} = m$ on at least 95 per cent of admitted draws after the restriction. *KILL* $m_{\text{eff}} < m$ on more than 5 per cent of draws in any cell quoting a floor: every argmin accuracy in that cell is re-floored at m_{eff} and this census line becomes an admission line. *price* 0 GPU-s · *mechanism* L-FLOOR, V-10, V-17, D-3 · *deliverable* docs/plan/EFFECTIVE_M.md and one column in results/bed_s_census.jsonl · *evenings* 1

R-24 · Register the clamp-target membership rule as a closed-form zero-hop control — phase 1 · prereq S-11 · independent-of R-05 to R-08; consumed by R-06 and S-62 ·

parallel-safe *build* the attack: the goal rule is decided at zero hops whenever some clamp target lies in the goal set, because a clamp into $\mathcal{A}_0$ gives $q^{(0)} = 1$ exactly and the goal flag of u_a is a token feature. The registered 0-hop control is a **trained** MLP; a closed-form rule R0 — “pick the move whose clamp target carries the goal flag; else the move whose target carries no constraint flag and sits earliest” — costs nothing, needs no seed, and reads its accuracy at construction. R0 and its two ablations (goal-flag only; earliest-transient only) are computed as oracle-free columns on every draw batch, with the fraction of draws on which R0 is defined. *prove* DERIVED: a clamp into $\mathcal{A}_0$ yields $q^{(0)} = 1$ and is the unique argmax unless another move also clamps into $\mathcal{A}_0$; on the judge’s draw $q^{(0)}(v) = 1.0000$ exactly for $u = 5 \in \mathcal{A}_0$ (RUN[M]). *measure* R0 accuracy against the oracle argmin with Clopper-Pearson on 512 draws, before and after R-07’s restriction. *PASS* R0 accuracy within the interval of $1/m$ after the restriction. *KILL* (frozen) R0 accuracy at or above $\text{acc}_{\text{shape}} - \text{MDE}_8$ on at least 6 of 8 arena seeds, or R0 above the exact zero-information floor by more than its interval on the unrestricted bed: the argmin label is a flag lookup, BED-S’s argmin head is struck, and only the committor-vector head remains. *price* 0 GPU-s · *mechanism* V-10, D-5, M-21, M-18 · *deliverable* docs/plan/FLAG_RULE.md and three columns in results/bed_s_census.jsonl · *evenings* 1

R-20 · File the one-step rule as a safety filter, not a policy, with the multichain hazard tested on the oracle — phase 1 · prereq S-11 · independent-of R-05 to R-08, R-24 · parallel-safe *build* on 512 admitted draws, the oracle-only comparison of the one-step rule against the two-step optimum (apply a , re-solve, take the best second clamp) on the same P_{env} ; the vocabulary fixed as a one-step, most-restrictive **safety filter** whose value is the goal committor ([57], [17] [V]), the aggregations as columns ([119], [132] [V]), and [84] [V] (Bellman’s principle can fail for safety-constrained multichain MDPs) carried in Limits beside the rule. *prove* DERIVED: on a DAG with declared absorbing sets every transient state is absorbed with probability one ($\rho(Q) = \max_T P_{ii} < 1$), so the chain is absorbing; the hazard lives in the policy-level object, which the paper does not claim. *measure* the fraction of draws on which the one-step and two-step optima disagree. *PASS* the paper never writes “optimal move” or “policy”, and the disagreement fraction is printed. *KILL* disagreement above 0.5 on 512 draws: the phrase “safest move” is replaced by “one-step safety filter” in every sentence and the sequential problem is filed as future work with the citation. *price* 0 GPU-s (oracle solves) · *mechanism* P-7, V-17, P-10 · *deliverable* docs/plan/SAFETY_FILTER.md, results/bed_s_onestep_vs_twostep.jsonl · *evenings* 1

PHASE 2 — the binds through the front door, the smoke test, the device certificate

Goal. Prove on the bed’s **real draw** that every identity the paper leans on holds and every planted negative fails at $O(1)$, that the solve path runs deterministically at the predicted price, and that the device certificate has a live producer. Nothing here reads a capability number.

The gate (V-14, V-24). Every plant of S-20 to S-27 exists as a FOUND cell under results/ with its own kind and **manifest_hash** — a plant that lives only in a script’s print is NAMED, and the bind is inadmissible. Every honest half passes; every plant fails at $O(1)$ with counts printed, in the pattern 0.9749/0.9165/1.000/0.4845 (READ workdonenewseal.md:114–122). The smoke cell reads inside the law-versus-measured band $[-1.8, +14.4]$ per cent of 1.680 s, or is re-labelled the per-op floor it was declared to be. The author decides one thing: whether the arena may start.

M-2.0 · Build the shape arm as a registered kind — phase 2 · prereq none (building is not training; L-LEAN gates M-2.2, not this) · independent-of every Phase-1 card · not parallel-safe with S-20 *build* one module exposing the corners **beta/qk/g**, a learnable $\gamma \in [0, 1)$ initialised at 0, **boundary_sets/goal_set/sink_set** as identity rows, **diag_convention**, **route**, **committor_route**, the read $O = (1 - \gamma)P(I - \gamma P)^{-1}V$ with the factor carried ([42] Eq. 5

[V]), and two heads: (H-q) the exact triangular solve at $\gamma = 1$ on the transient block and (H-z) the state and displacement channel at $\hat{\gamma}$; plus the **make_arm** branch, the optimiser path and the journal emit — the W2 lesson was that a three-line branch was the whole gap (READ V20_R15_WING_MANIFEST.md:115-123). *prove* nothing. *measure* the arm produces a journalled record with its **kind** under **results/** (FOUND, not NAMED). *PASS* one FOUND record. *KILL* a **kind** with zero records at the end of Phase 2 is STRUCK, which is the R15 kill applied to this lane. *price* 0 GPU-s to build; the first record is S-20's · *mechanism* FOUND-not-NAMED, P-4, L-2 · *deliverable* the author's module plus the FOUND line in docs/PLAN_STATUS.md · *evenings* 3

S-20 * · B-J parity at $\gamma = 0$ and the ChaCAL-diag smoke test at $\gamma = 0.9$ on BED-M — phase 2 · prereq S-01, S-02, M-2.0, J-L18 rows 1 and 5 (L-LEAN) · independent-of every Phase-1 card · not parallel-safe with M-2.0 · aliases M-2.2, R-12, V-7, N-05, ledger P2.1, B-J, B-P5 *build* **torch.equal**(0(0), PV) against the lane's own **softmaxAttn**; then one 150-step ChaCAL-diag cell at fixed $\gamma = 0.9$ (ChaCAL's own setting, [42] App. C) on BED-M e3_t2, $n = 2048$, $s = 64$, $d_{\text{model}} = 16$, under **use_deterministic_algorithms**(True) with the cuBLAS workspace exported **before** process start (set in-process it does not take, READ V17_R4_RETAKES_PRICE.md:178-181), timed with **torch.cuda.synchronize()** bracketing. Plants: $\gamma = 0.5$ (RUN[coord] 2.3002850040264393 at $s = 64$; RUN[SATURN] 1.6143 at $s = 32$); $\beta = 0$ at the softmax corner (gap above 0.5); ChaCAL-published's diagonal-removed inverse at any $\gamma > 0$; and B-P5's dense non-triangular M passed with **upper=False**. The struck plant — a non-causal W — passes bitwise at $\gamma = 0$ because $I - 0 \cdot W = I$ regardless of support (RUN[M]) and is recorded as struck. *prove* nothing here; Proposition 1 is J-L1's. *measure* **torch.equal** at $\gamma = 0$; **solve_triangular** against the dense inverse $1.7763568394002505 \times 10^{-15}$; the four plant residuals; the cell wall-clock; bitwise repeats forward and backward over 8 repeats under strict mode. *PASS* identity True; 4/4 plants at $O(1)$; the cell inside $[-1.8, +14.4]$ per cent of 1.680 s; repeats bitwise both ways. *KILL* **torch.equal** False at $\gamma = 0$ (the containment sentence dies); any plant passing (that bind is struck, V-24); the backward raising under strict mode (the “first gated wing whose training step runs strict” sentence is deleted and the lane runs **warn_only**); the cell above $2.2\times$ the law (every [FITTED + RUN] price in §5.7 is re-tagged the P-8 floor it was declared to be and re-measured before any GPU-minute figure is quoted). No capability number is read: BED-M is contained (J-L11, D-2). *price* ≈ 17.4 s for 8 seeds ($8 \times 1.681 + 4.0$), ≈ 1.7 s for a single cell [FITTED + RUN] · *mechanism* V-24, V-3, M-8, P-8, V-16, D-2 · *deliverable* results/binds/b_j_parity.jsonl with kinds shape_g0, shape_g05_plant, shape_b0_plant, chacal_pub_plant, dense_M_plant · *evenings* 1

S-21 · B-E1 boundary rows against ChaCAL-diag — phase 2 · prereq S-12 admitted, S-20 · independent-of S-22 to S-28 · parallel-safe · alias ledger P2.2 *build* identity half: with $\mathcal{A} = \emptyset$ the shape is **torch.equal** to **ChaCAL-diag**, the lane's own re-implementation with **diag_convention = kept** declared in the manifest (a declared V-3). Rejection half: with $\mathcal{A} \neq \emptyset$, $\|\Pi_{\text{shape}} - \Pi_{\text{ChaCAL-diag}}\|_{\infty}$ on rows downstream of the boundary positions exceeds a printed $O(1)$ gap — a **weight** statement, never a support statement, since on a dense causal softmax the support of Π_{γ} equals the support of P with or without boundary rows (RUN[J], rows changed = [3], RUN[M]). The struck plant — “a non-identity absorbing row breaks the conservation row” — is empty (RUN[M] 0.0; the sum uses transient rows only) and is replaced by the full- P read at $\gamma < 1$, where it reads 0.1491. The InfSA base must break the identity half: NOT MEASURED, **needs the base wired**. *prove* nothing. *measure* the identity half and the downstream gap on every admitted draw. *PASS* **torch.equal** True at $\mathcal{A} = \emptyset$ and a downstream gap of at least 0.1 at $\mathcal{A} \neq \emptyset$ on 100 per cent of admitted draws. *KILL* the identity half fails (the lane's ChaCAL-diag is not the shape at $\mathcal{A} = \emptyset$ — a P-7 object, and the “same operator” sentence dies); the gap reads below 10^{-3} on any draw (boundary rows change nothing there; census line 4 is re-read). *price* seconds · *mechanism* V-24, V-3, V-14, P-7 · *deliverable* results/binds/b_e1_boundary.jsonl · *evenings* 1

S-22 · B-E2 conservation with the sink set — phase 2 · prereq S-12 · independent-

of S-21, S-23 to S-28 · parallel-safe · alias ledger P2.3 *build* $q^{(\text{sink})} + q^{(0)} + \sum_k q^{(k)} = \mathbf{1}$ on T , printed with its residual on every batch **including when it fails** (V-23). Plants: drop BOS from every set and the solve must **raise** ($\rho(Q) = 1.000000$, $\det(I - Q) = 0.0$), never return a number; a set after the query gives $q = 0.0$ exactly and the draw is discarded by census line 4. The struck plant — “drop the goal set, so $\max_k q^{(k)} \geq 1/K$ ” — is a V-3 of the conservation row once the sink is present; the degeneracy lemma is carried as J-D2’s **proposition**, its plant being “BOS declared inside a constraint set”. *prove* J-L15 when it builds; until then DERIVED. *measure* the residual per batch: 4.4×10^{-16} (RUN[SATURN]), $[0.9999999999999993, 1.0]$ (RUN[I]), $[1.0000000000000000, 1.0000000000000001]$ (RUN[P]). *PASS* residual at or below 10^{-12} on 100 per cent; the BOS-drop raises; the after-query set discards. *KILL* a number returned on the BOS-drop (the solver silently regularised, V-16, and every committor number in the lane is suspect); residual above 10^{-8} on any declared draw (a second undeclared absorbing row). *price* 0 GPU-s · *mechanism* V-12, V-23, V-16, V-3 · *deliverable* results/binds/b_e2_conservation.jsonl · *evenings* 0.5

S-23 · B-G1 the interventional re-solve, with the triangularity plant — phase 2 · prereq S-17 · independent-of S-21, S-22, S-24 to S-28 · parallel-safe · aliases ledger P2.4, B-C6 *build* (a) $V \equiv \mathbf{1}$ implies $\max |\Delta z| \leq 10^{-15}$ (the bar is not “bitwise”: the exact zero is a code-path accident); (b) Gaussian V implies $\max |\Delta z| = O(1)$; (c) Sherman-Morrison closed form against a re-solve at or below 10^{-12} ; (d-prime) the **triangularity plant** — a non-causal P makes $\Delta z_{<i} \neq 0$, with the solve route named (`solve_triangular` gives $\Delta z_{<i} = 0.0$ and `torch.equal True`; an LU route gives 8.9×10^{-16} and `False`), and a two-row edit breaks the rank-one formula at $O(1)$ and is repaired by Woodbury; (e) the displacement identity’s residual at or below 10^{-12} . The EMC feedback plant is **deleted**: for $\gamma < 1$ the fixed point is unique on every P . *prove* J-L8 and J-L14 when they build. *measure* (a) 0.0 (RUN[SATURN]), 1.1×10^{-16} (RUN[M]); (b) 0.1096, 0.363, 1.127; (c) 1.2×10^{-15} , 8.9×10^{-16} , 4.4×10^{-16} ; (d-prime) 0.0761, 0.0868; (e) 1.03×10^{-15} , 1.2×10^{-15} , 1.36×10^{-15} . *PASS* (a), (c), (e) within bars and (b), (d-prime) at $O(1)$. *KILL* (b) reading below 10^{-6} on Gaussian V (an empty rejection region: the bind is struck and Proposition 7 drops to a remark); (d-prime) reading zero on a non-causal P (the solver is masking, not solving, and every Δz number in the lane is suspect). *price* 0 GPU-s · *mechanism* V-24, D-5, D-7, V-3 · *deliverable* results/binds/b_g1_resolve.jsonl · *evenings* 0.5

S-24 · B-G2 the vector-metric plants — phase 2 · prereq S-16 · independent-of S-21 to S-23, S-25 to S-28 · parallel-safe · alias ledger P2.5 *build* the journal carries the vector, a histogram and quantiles (the record’s Q6 census read 0 of 40); three plants on the metric itself: (i) the permuted oracle must **not** score 0; (ii) the oracle plus 0.1σ must be preferred to the permutation; (iii) $-\Delta z$ must be distinguished from Δz . *prove* nothing. *measure* the three plant readings. *PASS* 3/3. *KILL* any plant passing strikes the metric: pooled first-order Wasserstein fails (i) and (ii) by $14.465 \times$ (READ V20_R15_THEORY_TABLE.md:221), which is why it is refused up front. *price* 0 GPU-s · *mechanism* V-26, L-14, V-16 · *deliverable* results/binds/b_g2_metric.jsonl · *evenings* 0.5

S-25 · B-H1 the committor identity on the environment chain and on the arm’s $\hat{Q}$ — phase 2 · prereq S-02, S-12 · independent-of S-21 to S-24, S-26 to S-28 · parallel-safe · alias ledger P2.6 *build* on the environment chain, with BED-1 as the instrument: $q = (I - Q)^{-1} R \mathbf{1}$ against `bed_1.committor` on the **real** sets; the harmonic residual; the Kirchhoff second route at $K = 2$. On the arm’s causal $\hat{Q}$: invertible if and only if $0 \in \mathcal{A}$ (J-L10), a diagonal read $\rho(\hat{Q}) = \max_T \hat{P}_{ii}$. Plants: the must-fire perturbation; declaring $\mathcal{A}_k$ on the wrong set. *prove* J-L10 and J-L16(c) when they build. *measure* 0.0 on the real sets ($A = [0]$, $B = [1]$, $|T| = 9$, RUN[coord]); residual $1.0408340855860843 \times 10^{-17}$; the perturbation drives it to 10^{-6} , ratio 9.6×10^{10} ; Kirchhoff below 10^{-10} (READ MATHEMATICS.md:542-600); $\rho(\hat{Q}) = 0.692660$ both ways. *PASS* 0.0 on the real sets and both plants at $O(1)$. *KILL* the residual not moving under the perturbation (it is reading a cached q , V-3, and every committor identity in the lane is suspect); the arm-side solve returning a number with BOS undeclared (V-16). *price* 0 GPU-s · *mechanism* D-2, V-3, V-16 · *deliverable* results/binds/b_h1_committor.jsonl · *evenings* 0.5

S-26 · B-H2 the argmin identity and the static-task kill — phase 2 · prereq S-14, S-15, S-16, R-06 · independent-of S-21 to S-25, S-27, S-28 · parallel-safe · aliases V-5, ledger P2.7, K-H2 *build* the argmin from the oracle tensor equals the label’s argmin on 100 per cent of admitted draws (an identity of the builder, declared V-3); the rejection region is the control pair — a 0-hop per-position MLP on the candidate’s context row and a 1-hop softmax, matched at the recounted parameter count (R-15), McNemar on identical draws. *prove* nothing. *measure* the identity fraction; the 0-hop and 1-hop accuracies with Clopper-Pearson; the McNemar *p*. *PASS* identity at 100 per cent; the 0-hop MLP outside the interval half-width of the shape at $N = 8$ and McNemar $p \leq 0.05$ against the 1-hop softmax. *KILL* (K-H2) the 0-hop MLP inside the half-width, or $p > 0.05$: BED-S is a third static task and is struck before any number is quoted; what dies is every capability sentence on BED-S, and §5.10 tree C is the paper. *price* ≈ 0.5 GPU-min (two control arms, 8 seeds each, at or below 1.524 s) [FITTED] · *mechanism* V-24, V-3, C8, V-10 · *deliverable results/binds/b_h2_argmin.jsonl* · *evenings* 1

S-27 · B-K containment and segmentation, the float instances — phase 2 · prereq none (the Lean grades are Phase 0’s) · independent-of every card · parallel-safe · alias ledger P0.1 float half *build* the float instances the [M] targets must match: corner 3 against $(I - A)^{-1}$ entrywise and its last row against `equilibrium_oracle`; segmentation zeros by `torch.equal`, never `allclose`, on a masked *P*. Plants: keeping the diagonal breaks $A^s = 0$ (`Nilpotent.one_not_nilpotent`); a -30 logit leaves the block non-zero. The 10^{-300} plant is struck as float64 underflow — a path of at least 80 gates at 0.5 annihilates exactly, so “nonzero” can read F0. *prove* J-L6, J-L11 when they build. *measure* 0.0 entrywise, $6.217248937900877 \times 10^{-15}$ on the last row (`RUN[coord]`); `torch.equal` zeros; the two plants at $O(1)$; the census 3 of 3 on BED-M and 0 of N on BED-S’s softmax corner. *PASS* the three identities and both plants. *KILL* (K-K) any [M] target not building by the Lean milestone is [S] and the paper cites only declarations that build; what dies is the tag, never the float instance. *price* 0 GPU-s · *mechanism* P-11, V-25, V-2, L-LEAN · *deliverable results/binds/b_k_containment.jsonl* · *evenings* 0.5

S-28 · The front-door plant register — phase 2 · prereq S-01, S-20 to S-27 · independent-of nothing in Phase 2 · not parallel-safe *build* one table listing, per bind, the plant’s `kind`, its `manifest_hash`, the journal path, the $O(1)$ number it produced and the verdict line that consumed it; plus a census of kinds over `results/**/*jsonl` calibrated **both ways**, as the wing manifest was (one arm must read zero, the control must read above zero, `READ V20_R15_WING_MANIFEST.md:73-83`). *prove* nothing. *measure* the FOUND count per plant and the two calibration reads. *PASS* every plant of S-20 to S-27 FOUND, and the two calibration reads correct. *KILL* any plant that exists only in a script’s print (NAMED): that bind is inadmissible and its proposition drops to a remark (V-14: the control that validates the matcher and never the reach). *price* 0 GPU-s · *mechanism* V-14, V-7, P-1 · *deliverable docs/apparatus/PLANT_REGISTER.md* · *evenings* 0.5

S-52 * · The capped run at seeds 2, 3, 7 on BED-M that the record priced four times and never took — phase 2 · prereq none · independent-of every card in this plan · parallel-safe · aliases M-2.1, R-13, V-6 *build* `arm_pl` on `e3_t2` at the three NO-READING seeds (1.113403, 1.139404, 1.152430, with $\hat{a}_{\max}$ 20.31, 49.66, 285.07, `READ V15_R1.md:177-188`) with the magnitude capped at 1.0, the same thread lane and flag regime as the retake journal, `synchronize()` bracketing and randomised order. The prediction and its counter are filed first (S-60). *prove* nothing. *measure* `eval_nrmse` and $\hat{a}_{\max}$ per seed against the uncapped values, against the one-hop capability threshold 0.7071067811865476 — which is a threshold and never a floor (C15). *PASS* all three capped seeds cross the threshold: the divergence was the cause of NO READING, and the paper’s negatives section says so with the number. *KILL* any capped seed still above 1.0: the split is not the cap’s, the account of R1 keeps “3 of 8 NO READING, mechanism unresolved”, and the Q3/W3 ledger cell is TERMINAL as the record’s last verdict graded it. Either way the shape’s own sentences are untouched; what dies is a stale question, cheaply. *price* ≈ 9.65 s ($3 \times 1.884 + 4.0$) [FITTED] · *mechanism* M-6 (a partial run read as a verdict), D-6 (the repair

written and never started), P-3, C17 · *deliverable results/bedm/capped_seeds_2_3_7.jsonl · evenings* 0.5

N-01 · Add solve_triangular as the fourth determinism quantity to the device certificate — phase 2 · prereq none · independent-of every card · parallel-safe · alias ledger C-1 (first half) *build* a specification for one more entry beside `hop / forward / gradient` in `scripts/k_cert.py::determinism_at_64` (READ `scripts/k_cert.py:579-640`): a batched lower-triangular M with unit diagonal and a float32 right-hand side, the forward solve and its backward through a mean-squared error, 8 repeats, both flag regimes, single stream, workspace pinned before process start. Journal fields `solve_fwd.{bitwise,max_abs,executable}`, `solve_bwd.{...}`, `documented_guarantee`, `trsm_dispatch`. The certificate line is fixed here: forward bitwise, backward bitwise, both flag regimes, torch-documented guarantee NONE, cuBLAS dispatch [U], one box, one process, one stream. *prove* nothing; a repeatability reading, never a guarantee (V-16 forbids reading silence as a pass). *measure* $\max |\Delta|$ over 8 repeats per quantity, three outcomes distinguished (bitwise, drifting, not executable). *PASS* 0.0 forward and backward under both regimes, reproducing the cost session’s reading. *KILL* any $\max |\Delta| > 0$ under the flag: the shape’s training cells are journalled `warn_only` like the scan arms and Ruling 1’s bitwise bar is not claimed for training; a raised error on the backward: the same fallback and the “first gated wing under strict mode” sentence is deleted. *price* ≈ 1 GPU-s [RUN class] · *mechanism* P-1, V-16, V-23 · *deliverable results/k_cert_local.json* new keys plus one line in *COSTS.md* §1.6 · *evenings* 1

N-02 · Make the solve microbenchmark a producer — phase 2 · prereq N-01 · independent-of S-2x, N-19 · parallel-safe · alias M-2.3, ledger C-1 (second half) *build* a `solve_increment` block timing PV , solve plus Pz , and Neumann $K \in \{1, 2, 4, 8, 16\}$ forward and backward at $n \in \{2048, 4096, 8192\}$, $s = 64$, $d = 16$, $\gamma = 0.5$, float32, synchronize bracketed, 2 warm-ups, median of 14; journal fields `pv_ms`, `solve_ms`, `neumann_ms[K]`, `increment_ms`, `n`, `s`, `d`, `gamma`, `cublas_workspace`, `producer_cmd`, plus peak bytes above a reset baseline per route. *prove* nothing. *measure* the increment per n : the reference values are 1.041, 2.312, 4.542 ms (2.514 – 1.473, 4.659 – 2.347, 9.916 – 5.374, RUN[NEPTUNE]). *PASS* each within the clock spread of plus or minus 12 per cent, and one Neumann hop slower than the solve at every n (3.001 against 2.514 ms at $n = 2048$). *KILL* the increment at $n = 2048$ above $2.2\times$ the reference: every shape price in §5.7 is re-derived before any card past the arena runs; the solve slower than one hop: the Neumann route re-enters as a cost path and N-10 is promoted. *price* ≈ 1 GPU-min DERIVED; the whole `k_cert` rerun is 523.9 s [MEASURED] · *mechanism* P-1, P-8, M-3, M-8 · *deliverable results/k_cert_local.json* `solve_increment` block · *evenings* 1

N-07 · The training step under strict mode: the identical-seed floor for the shape — phase 2 · prereq N-01, S-20 · independent-of S-2x · parallel-safe *build* the floor run in the record’s pattern (6 identical-seed pairs, READ `V17_R4_RETAKES.md:259-286`) for the shape and the softmax control under `use_deterministic_algorithms(True)` with `warn_only=False`, journaling `deterministic_regime = strict`. *prove* nothing. *measure* the NRMSE delta on 6 of 6 pairs; whether the backward executes. *PASS* 0.0 bitwise on 6/6 for both arms. *KILL* the shape’s backward raises: the “first gated wing under strict mode” sentence is deleted and the lane runs `warn_only` (Ruling 1); softmax raises and the shape does not: the sentence stands with the control’s row printed beside it (V-23). *price* ≈ 45 s (the record’s six pairs cost 44.7 s) · *mechanism* Ruling 1, M-10, V-23 · *deliverable results/shape_floor_strict.jsonl* · *evenings* 1

S-40 · Z-EX the exact route: $\delta = 0$, the only route the committor head ships on — phase 2 · prereq S-20 · independent-of S-41 to S-44 · parallel-safe · alias ledger Z-EX, Z-COM *build* `solve_triangular` for z and O at any $\gamma < 1$ and for the committor at $\gamma = 1$ on the arm’s triangular $\hat{Q}$; determinism as **three outcomes** (bitwise, drifting, raises), with `cumsum` required to raise and a run reporting “pass” on an operator that raised refused (V-16). The committor head carries **no** truncation certificate: the Neumann δ is infinite at $\gamma = 1$ and the sup-norm substitute reads 19.2 at $K = 4$ against a true error 2.5×10^{-3} (RUN[M]); the Perron-weighted certificate needs

a weight that is NOT MEASURED (S-72). The threshold form $\max_k \hat{q}^{(k)} + \delta \|V\|_\infty \leq \delta_{\text{thr}}$ certifies the **solve, never the model**: $\hat{P} \neq P_{\text{env}}$ has no certificate and no move is “admitted” by it. *prove* J-L10, J-L13 when they build. *measure* 1.8×10^{-15} against the dense inverse; the three-outcome table from N-01. *PASS* the table filled and `route = solve_triangular` on every committor cell. *KILL* a committor cell journalled with `route = neumann_K`: no certificate exists for it and the cell is refused. *price* 0 GPU-s (the microbenchmark is N-01’s) · *mechanism* M-8, V-16, L-CERT, V-17 · *deliverable* docs/apparatus/CERT_EXACT_ROUTE.md · *evenings* 0.5

S-41 · Z-NEU the Neumann certificate in vector units, with the convergent plant — phase 2 · prereq S-01 · independent-of S-40, S-42 to S-44 · parallel-safe · alias ledger Z-NEU, B-I, P2.8 *build* for $P \geq 0$ row-stochastic including absorbing rows and $\gamma < 1$: the matrix residual **equals** $\gamma^{K+1}/(1 - \gamma)$ bare and γ^{K+1} under Π_γ — attained, hence **definitional** on the class (declared V-3; [83] owns the tail). The bind is carried by the plant and the units: rows scaled to 1.5 at $\gamma = 0.6$, $K = 2$ read 7.290 against 0.540 with $\rho(\gamma P) = 0.90$, a **convergent** series; the coordinator’s $\gamma = 0.7$ plant (119.37 against 1.143) is a divergent partial sum and is labelled so. The vector statement $\max |z_{\text{solve}} - z_K| \leq \delta \|V\|_\infty$ is printed with $\|V\|_\infty$ and with $1/(1 - \hat{\gamma})$ beside every δ (the record’s own factor was undefined at all eight R1 seeds, `READ V15_R1.md:56`). *prove* J-L12 when it builds. *measure* equality to 10^{-15} at $K \in \{1, 2, 4, 8, 16\}$ with absorbing rows (RUN[M], error minus bound in $[4.4 \times 10^{-16}, 6.7 \times 10^{-16}]$); the plant’s ratio. *PASS* equality at all five K ; the plant fails by at least $10\times$; `V_inf` and `one_over_1mg` present on every Neumann cell. *KILL* (K-I) the vector bound exceeded on any of 1,024 drawn cells, or $\delta \|V\|_\infty \geq \text{sd}(\text{label})$ (uninformative), or the row-sum identity presented as evidence about P (V-3): the certificate sentence dies and the exact route is the only one. *price* ≈ 0.04 GPU-min for the 1,024 passes DERIVED · *mechanism* V-3, V-10, V-17, V-24, L-CERT · *deliverable* results/certs/z_neu.jsonl · *evenings* 0.5

S-44 · Z-F0 the segmentation dividend in pair-count units, a corpus property — phase 2 · prereq S-27 · independent-of S-40 to S-43 · parallel-safe · aliases N-14, ledger Z-F0 *build* the dividend $D = s(s+1)/\sum_m L_m(L_m+1)$ printed with the corpus named and the segment-length distribution beside it; three non-transfer clauses printed with it: (i) on the softmax corner an exponential has no exact zero, so the dividend is $1\times$ unless a mask is declared (F0 by mask, never by gate); (ii) a cut makes the segment head a new undeclared absorbing state and severs the boundary sets behind it, so segmentation and the reach-avoid read coexist only if every segment carries its own sink, goal and constraint sets with the cut a registered dial; (iii) the dividend applies to the exact, Neumann and CSR routes alike and carries to **no other corpus** (V-22). *prove* J-L6, J-L7 when they build. *measure* BED-M’s live-pair fraction 0.0322 gives $31.06\times$ (`READ V16_ARM_SMPRIME.md:518–522`; the source line’s pair counts disagree by a factor of two and the discrepancy is flagged, the fraction is what is used); the alternative $s^2/\sum L_m^2$ would read 58.5 and is V-17; BED-S’s softmax corner prints $1\times$ and “F0 silent”. *PASS* both readings with their corpora named. *KILL* a dividend quoted for a corpus without its census, or carried across corpora (V-22); a committor cell run across an undeclared cut. *price* 0 GPU-s · *mechanism* V-22, V-17, V-25, D-3, L-CERT · *deliverable* docs/apparatus/CERT_F0_SEGMENTATION.md · *evenings* 1

S-60 · File the predictions and counters, hash the file, record the hash before the first arena cell — phase 2 · prereq S-12, S-13, S-16, S-35’s settings · independent-of S-2x binds · not parallel-safe with S-62 · alias V-1 *build* one file with the five bets of $\$5.5$, each with its prediction, a counter of equal specificity, the SPLIT band where one exists and the deciding statistic; plus S-52’s bet and the language-model γ row left blank until an LM cell runs behind the author’s yes. Under the record’s calibration (7 of 8 optimistic, $p = 0.0352$) the counter is the point estimate (D-CALIB-1) and no shrink factor is fitted (D-CALIB-3). The file’s sha256 is written into `results/arena/PREDICTIONS.sha256` **before** any arena cell exists, and every arena cell’s manifest carries `predictions_hash`. *prove* nothing. *measure* the count of rows with both halves signed against the count of rows: it must read 24/24. *PASS* the hash is recorded, every bet has a counter, and the MDE cells reference the realised-sd column S-62 will fill. *KILL* a bet without a counter (D-7); a hash recorded after the first cell’s timestamp (M-2): that cell’s numbers

are not quotable. *price* 0 GPU-s · *mechanism* D-7, L-SIGN, M-2, M-7, L-FIRST · *deliverable* docs/apparatus/PREDICTIONS_R16.md, results/arena/PREDICTIONS.sha256 · *evenings* 1

PHASE 3 — the arena on BED-S at the design point

Goal. Run every arm on byte-identical draws at $t^* = 8$, $m = 8$, $K = 2$, $s = 64$, $N = 8$ seeds deduplicated, one thread lane, and let the realised paired sd fix the minimum detectable effect that every kill formula in Phase 4 refers to. Two sequencing rules bind: the **controls run before the shape** (M-2, order stated), and **no arm trains before Phase 0’s gate is green** (L-LEAN).

The gate. Eight FOUND records per arm with the header’s flag regime present; the realised paired sd printed and the MDE row filled by the same guarded bisection that reproduces 0.039827 at sd 0.034451. If the realised sd is at or above $2.18\times$ the pilot’s 0.050146 — the M-3 precedent, where 0.109199 gives $\text{MDE}_8 = 0.126238$ — nothing in §5.5 is falsifiable at $N = 8$, N is repriced from the realised sd **before any bet is scored** (K-P), and the paper files the bed and no capability number.

S-30 · The depth-1 softmax control with the same head — phase 3 · prereq S-12 · independent-of S-31 to S-35 · parallel-safe · alias M-3.1 (first half) *build* the matched control at the recounted parameter count (Ruling 3), with the same $[m, K + 2]$ head or vector readout as the shape; it is the per-row control (Obstruction 1, definitional) and the owner of the single-location ground ([38] Prop. 4.2, whose label model admits 0 of 3 record beds). *prove* nothing. *measure* per seed: φ -NRMSE, harmonic residual with its two unit factors, argmin accuracy with its interval, $\hat{\gamma}$ and Λ where applicable; the paired sd of every contrast. *PASS* one journalled cell per bed with **params_per_arm** in the manifest. *KILL* a count differing from the shape’s by more than 0.1 per cent after the head is attached: the pair is not matched and “at matched parameters” may not be written of it. *price* 1.524 s per cell [FITTED], ≈ 16.2 s for eight seeds plus fixed cost · *mechanism* MISTAKES.md D-1, Ruling 3, R-SKY · *deliverable* results/arena/softmax_d1.jsonl · *evenings* 0.5

S-31 · The two ChaCAL arms and the sink-token arm — phase 3 · prereq S-20, S-21 · independent-of S-30, S-32 to S-35 · parallel-safe · aliases M-3.3 (part), V-15, R-19 (part) *build* **ChaCAL-diag**, the shape with $\mathcal{A} = \emptyset$ and the diagonal kept — the identity half of B-E1 is against this arm only; **ChaCAL-published**, with the diagonal removed inside the inverse ([42] Eq. 5, [V-fetched] from one HTML render and **re-read against the PDF before it is typeset**), a regime-N inverse inside a regime-S read, sub-stochastic (RUN[J] row sums 0.400 to 0.765), and the planted negative of the bitwise half; **ChaCAL-with-sink-token**, a column device standing in for the row condition, the C2 kill’s control. The “same γ ” is the shape’s **trained** $\hat{\gamma}$, so the pair is sequential and the order is stated in the manifest (M-2 in potential form). The struck control — “ChaCAL must emit chance” — is a control that must fail by construction (V-2) and is deleted. *prove* nothing. *measure* the creditable contrasts of S-13, per seed. *PASS* three FOUND arms with **diag_convention** declared on each; the shape’s committor φ -NRMSE below ChaCAL-with-sink’s by more than MDE_8 on at least 6 of 8 seeds and its residual smaller by more than $2\times$. *KILL* (K-E1, the expensive death) ChaCAL-with-sink within MDE_8 on at least 6 of 8 **and** residual within $2\times$: escalate once to $N = 16$ (S-64), never report “within TOST” at $N = 8$ (M-13). What dies if it holds at $N = 16$: the boundary-row mechanism as a **capability**; the paper becomes a property paper about a published operator, and §5.10 tree B is its sentence. *price* 1.680 s per cell [FITTED + RUN], ≈ 52 s for three arms at eight seeds · *mechanism* V-3, P-10, V-24, M-13, R-SKY · *deliverable* results/arena/chacal_{diag,pub,sink}.jsonl · *evenings* 1.5

S-32 · The InfSA-style Neumann read at $K = 16$, no boundaries — phase 3 · prereq S-12, S-41 · independent-of S-30, S-31, S-33 to S-35 · parallel-safe *build* the truncated read $\sum_{k \leq 16} (\gamma P)^k V$ with its printed $\delta_\Pi = \gamma^{17}$ and $\delta \|V\|_\infty$ ([99] [V] owns the reading); the exactness

contrast **shape** - **Neumann-16** is creditable. *prove* nothing. *measure* the gap between the exact solve and the $K = 16$ read against $\delta \|V\|_\infty$ on every draw of the cell. *PASS* one FOUND cell with **route** = **neumann_K**, $K = 16$ and **delta_vec** filled. *KILL* the exact solve not within $\delta \|V\|_\infty$ of the $K = 16$ read on any draw: either the arm is not row-stochastic on that cell (S-42) or the certificate is in the wrong units (V-17), and the exactness contrast is withdrawn. *price* ≈ 5.12 s per cell, ≈ 44.9 s at eight seeds, DERIVED from the 25.409 ms hop microbenchmark · *mechanism* MISTAKES.md D-1, L-CERT, V-17 · *deliverable* results/arena/infsa_k16.jsonl · *evenings* 0.5

S-33 · The cached-mixture arm on the two plant classes — phase 3 · prereq S-17, S-23 · independent-of S-30 to S-32, S-34, S-35 · parallel-safe · aliases M-3.5, R-25, V-16, ledger P3.3 *build* $O_{\text{cached}} = \hat{P}_{\text{base}}(I - \gamma \hat{P}_{\text{base}})^{-1} V_{\text{int}}$, with $\hat{P}$ frozen from the un-intervened context and values from the intervened one — the operational form of a cached successor representation ([85], [101], mechanism [U]). It must fail on the P -changing plants and pass on the V -only plants, or the interventional channel is not needed. *prove* nothing. *measure* field cosine and residual per plant class, paired by seed. *PASS* on V -only plants within one seed sd of the shape; on P -changing plants worse by more than one seed sd on at least 6 of 8. *KILL* (K-G1) within one seed sd on P -changing plants on at least 6 of 8: the re-solve is a per-row control wearing a name (the precedent is the round where a “signed” pivot arm was the unsigned arm renamed, READ workdonenew.md:379), component (g)’s re-solve sentence dies, and the consequence channel is repriced as a cost statement. *price* 1.680 s per cell, ≈ 35 s for the two batches · *mechanism* D-2, P-7, V-9 · *deliverable* results/arena/cached_mixture.jsonl · *evenings* 1

S-34 · The argmin controls: 0-hop MLP, 1-hop softmax, majority, random, predict-the-mean — phase 3 · prereq S-26, R-06 · independent-of S-30 to S-33, S-35 · parallel-safe *build* five arms with their own kinds, McNemar paired on identical draws; **random** reads $1/m$ and **majority** reads $\max_a \hat{\pi}(a^*)$ by construction (declared V-3, printed as the floor row). *prove* nothing. *measure* accuracy with Clopper-Pearson per arm; the two by-construction arms against their formulas. *PASS* five FOUND cells and the two formulas reproduced within their intervals. *KILL* as S-26: the 0-hop arm inside the shape’s half-width strikes the bed. *price* ≤ 1.524 s per cell · *mechanism* C8, V-10, V-3 · *deliverable* results/arena/argmin_controls.jsonl · *evenings* 0.5

S-35 · The three skylines, settings fixed before the arena; the depth-5 cell — phase 3 · prereq S-12, R-15 · independent-of S-30 to S-34 · parallel-safe · aliases M-3.4, R-23, V-14, ledger K-1 *build* the $\lfloor \log_2 t^* \rfloor + 2$ softmax stack — depth 3/5/7 at $t^* = 2/8/32$, chosen by **analogy** with [105] Thm 4.2, whose task is not the committor (the reduction is NOT FOUND, sweep_expressivity.md §3.5); the wide constant-depth stack at width n_{nodes} ([139]), which at $d_{\text{model}} = 16 < s = 64$ has no matched-parameter instance and the table says so; the chain-of-thought decoder at t^* steps ([80]). A deeper stack computes the same resolvent by iteration ([126], [131]), so “beats softmax” is banned either way (R-SKY). *prove* nothing. *measure* $\text{acc}_{\text{shape}} - \text{acc}_{\text{sky}}$ against MDE_8 . *PASS* (Bet D’s prediction) within MDE_8 : the honest control holds and the shape’s separate claims are exactness, one-read cost and the boundary mechanism. *KILL* outside in either direction, sign logged: short by more than MDE_8 says the bed needs something the hop construction does not supply; ahead by more than MDE_8 says “one operator” carries no accuracy sentence at all. A depth chosen after the arena strikes the row (M-2). *price* ≈ 7.6 s per cell [ASSUMED linear in depth], ≈ 64.8 s at eight seeds; the other two skylines NOT MEASURED with their settings journalled · *mechanism* MISTAKES.md D-1, R-SKY, P-10, V-25, D-7 · *deliverable* results/arena/skyline_depth5.jsonl, docs/apparatus/SKYLINE_SETTINGS.md · *evenings* 1

S-62 * · The eight-seed arena, one invocation, one thread lane — phase 3 · prereq S-12 admitted, S-21 to S-28, S-30 to S-35, S-60, S-61 · independent-of nothing in Phase 3 · **not** parallel-safe · aliases M-3.2, R-18, N-06, V-9 to V-13 (read off it), ledger P3.1, S-4 *build* one invocation: **device** = **cuda**, **threads** = **8**, **steps** = **150**, $n_{\text{train}} = 2048$, $n_{\text{eval}} = 4096$, $s = 64$, $d_{\text{model}} = 16$, the flag regime journalled on the header; seeds 0 to 7 deduplicated; arms **shape**, **softmax-d1**, **ChaCAL-diag**, **ChaCAL-published**, **ChaCAL-sink**, **InfSA-16**, **cached-mixture**, plus the **depth-5 skyline**; every cell carrying the S-01 manifest, **predictions_hash**, **void_contrasts**

and the census block; the run order randomised and every timed region bracketed by two `synchronize()` calls (the record's timer had run order as its strongest correlate, $\rho = +0.7029$, `READ V20_R15_JOURNAL.md:53`). *prove* nothing. *measure* per arm and seed: φ -NRMSE, the harmonic residual with $\sigma_{\min}$ and the operator norm printed, argmin accuracy with its interval, the field cosine and magnitude ratio, the sign column, $\hat{\gamma}$ and Λ , $\|\hat{P} - P_{\text{env}}\|_{\infty}$, $1/(1 - \hat{\gamma})$; then the **realised paired sd** of every scored statistic and the MDE row. *PASS* eight FOUND seeds per arm, the header's flag regime present, the MDE row filled, and the identical-seed floor re-read. *KILL* (K-P) realised sd above $2.18\times$ the pilot: reprice N before any bet is scored; what dies is the falsifiability of every $N = 8$ kill in §5.5, and the paper files the bed with no capability number. *KILL* (K-D2) $\|\hat{P} - P_{\text{env}}\|_{\infty} < 10^{-3}$ on at least 6 of 8: the arm copied the environment and every contrast on this bed is a copy-versus-no-copy statement. *price* $7 \times 8 \times 1.680 + 4.0 \approx 98$ s ≈ 1.6 GPU-min DERIVED, plus ≈ 64.8 s for the skyline [ASSUMED] and ≈ 44.9 s for InfSA-16 DERIVED, so ≈ 3.3 GPU-min in total · *mechanism* M-10, M-16, V-26, M-3, Ruling 1 · *deliverable results/arena/r16_t8_m8_k2.jsonl* · *evenings* 1

S-51 · The first BED-K cell of the arm's shape ever run — phase 3 · prereq S-01, M-2.0 · independent-of every BED-S card · parallel-safe · alias M-3.6, ledger P3.4 *build bed_k.build_delay* on its pinned generator, against the native control *hard_delay_attention*, which reproduces the delay to $(n - 1)e^{-45}$ (`READ ceq/beds/bed_k.py:311-336`); zero cells of BED-K's shape have ever run (`READ V20_R15_LEAP_LEDGER.md:323`). The Hankel ceiling is $\text{err}_1 = 0.9746794345 = \sqrt{1 - 1/20}$ at $d = 20$ (DERIVED). *prove* the shape with $\gamma > 0$ has **no** delay advantage by theorem (`V15Kernel.first_order_cannot_delay`); the reading decides only that it is not worse than one head at a fixed offset (R-SKY). *measure* NRMSE against the ceiling, $N = 8$. *PASS* the shape and the native control within MDE_8 of each other. *KILL* the shape short by more than MDE_8 on at least 6 of 8: a fixed-offset head the resolvent cannot represent at matched parameters, filed as a negative; the shape claiming a delay advantage is refused by theorem. *price* ≈ 34 s per pair · *mechanism* V-3, R-SKY, V-25 · *deliverable results/bedk/first_cell.jsonl* · *evenings* 1

S-50 · BED-M as containment: the parity bind and the truncation-law must-fire — phase 3 · prereq S-20 · independent-of S-51 to S-53 · parallel-safe *build* BED-M scores a scalar at one position and its label is the last row of corner 3's own resolvent, so **shape - corner-3** is VOID by registration (and by theorem once J-L11 builds); BED-M stays as the host of S-20's parity bind and of the truncation-law must-fire $\sqrt{(t^* - k)/t^*}$ at $k = 0, 1, 2$. *prove* J-L11. *measure* 1.000058, 0.712039, 0.0 at $t^* = 2$ (`RUN[VENUS]`); the one-hop threshold 0.7071067811865476 printed as a threshold, never a floor — 13 of 40 banked cells violate it (`RUN[WATSON]`). *PASS* both readings journalled with *void_contrasts* set. *KILL* any BED-M cell carrying a capability number for the shape is struck at assembly. *price* ≈ 34 s if the ladder is re-run · *mechanism* MISTAKES.md D-1, D-2, C15 · *deliverable results/bedm/containment.jsonl* · *evenings* 0.5

M-3.7 · Vary the dials: the one-factor-at-a-time sweep — phase 3 · prereq S-62 admitted and not killed · independent-of S-31 to S-35, S-51 · parallel-safe *build* seven design points — the centre $(t^*, K, m) = (8, 2, 8)$ plus $t^* \in \{2, 32\}$, $K \in \{3, 4\}$, $m \in \{4, 16\}$ — each a shape-plus-softmax pair with the census re-run per point; t^* is placed by the DAG's depth and read off the hop ladder, never by a kill rate (which collapses the label: the record watched a label's sd fall from 0.499989 to 0.038445, `READ scale/e4_harmonic.py:425-434`). *prove* nothing. *measure* the S-62 statistics per point; the exact zero-information floor per m (0.75, 0.875, 0.9375 at $m = 4, 8, 16$). *PASS* every point admitted by its own census and the argmin error moves with t^* (D-3: a dial that does not vary describes nothing). *KILL* a point that cannot be placed ($t^* = 32$ unreachable inside the causal window) is **recorded, never forced**; error flat across t^* means the dial is decoration and only the centre is reported. *price* ≈ 3.5 to 4.0 min DERIVED ($7\times$ one pair) · *mechanism* D-3, P-8, V-25, V-10 · *deliverable results/arena/dials.jsonl* · *evenings* 2

M-3.8 · Open the vector-label lane on BED-M — phase 3 · prereq S-20 · independent-of every BED-S card · parallel-safe *build* a separate lane and journal, because a vec-

tor label voids the record's published softmax baseline (READ MISTAKES.md:701-708); the `vector_readout` plumbing exists (READ `scale/m3_quintuple.py`:368,483); the label is the full z^* from `equilibrium_oracle` with the last-coordinate discard removed; shape against softmax with the vector readout, $N = 8$. *prove* nothing. *measure* the position-matched NRMSE vector; the harmonic residual with the chain as P_{env} ; S-24's three metric plants. *PASS* a journalled vector with its histogram and quantiles (the record's Q6 census read 0 of 40). *KILL* the permuted oracle scoring zero on the chosen metric strikes the metric (V-26); `shape - corner-3` on this lane is VOID by registration. *price* ≈ 34 s · *mechanism* MISTAKES.md D-1, M-1, V-26, D-2 · *deliverable* `results/bedm/vector_lane.jsonl` · *evenings* 1

V-17 · Bet G: a candidate move costs one column, not one solve — phase 3 · prereq S-20 · independent-of every other Phase-3 card · parallel-safe *build* the microbenchmark: after one solve at $n = 2048$, $s = 64$, $d = 16$, price $m = 8$ query-side row clamps two ways — (i) the Sherman-Morrison column route (one forward-substitution column at $s^2/2$ MACs plus an $s d$ inner product) and (ii) a fresh re-solve per candidate; synchronize bracketed, median of 14 after 2 warm-ups. *prove* `sherman_morrison_row [S]` (J-L14), a row clamp only. *measure* the measured ratio of (i) to one solve. *PASS* (prediction) route (i) at or below $0.5\times$ one solve for $m = 8$, since $m/d = 0.5$ at $(8, 16)$. *KILL* (counter) at or above $8\times$: the consequence channel is repriced as m solves in every table; SPLIT in $(0.5\times, 8\times)$ prints the band and no point. *price* seconds · *mechanism* M-8, P-8, M-3 · *deliverable* `results/arena/move_price.json` · *evenings* 1

PHASE 4 — verdicts, the remaining certificates, the dossier

Goal. Score every bet against the number frozen before the run, append the verdict without touching the prediction, and write the dossier whose brackets are all filled from journals.

The gate. `VERDICTS.jsonl` carries one line per bet with `predictions_hash`, a token in {`HOLDS`, `COUNTER`, `SPLIT`, `VOID`} and the deciding interval; the prediction file is byte-identical to its hash; the calibration column has a row per scored bet whether or not it flatters. The author then reads one of the three sentences of §5.10 and knows which paper he has.

S-63 · Record the verdict append-only; append the calibration row; do not edit the prediction — phase 4 · prereq S-62 · independent-of nothing · not parallel-safe · alias V-19 *build* the adjudicator's output appended to `results/arena/VERDICTS.jsonl`; a superseded cell gets a **supersede** marker, never a deletion (L-G2); the calibration row (checked, wrong, sign) appended to the cross-round column whether or not it flatters (D-CALIB-5); the sign of every miss logged. The adjudicator itself reads **only** journals: it refuses below 8 distinct seeds, deduplicates by seed (bit-identical duplicate rows exist in the record), reads the thread lane from the journal and never from the machine (cross-thread drift 2.345×10^{-3} ; a margin below about 4.7×10^{-3} is indefensible), computes the realised paired sd from the first eight seeds, refuses any contrast smaller than the $n = 8$ cell of that sd (M-3), checks `predictions_hash` on every cell against the recorded sha256 and **refuses to run on a mismatch**, and never writes to the prediction file. *prove* nothing. *measure* the adjudicator's own plants: a journal with 7 seeds returns INSUFFICIENT; a duplicated seed is counted once; a tampered prediction file halts the run. *PASS* the verdict line is present and the prediction file is byte-identical to its hash. *KILL* any edit to the prediction file after the hash (M-2), any verdict filed without its bet id, or any branch computing a verdict on a standard nobody registered: the verdict is struck. *price* 0 GPU-s · *mechanism* M-2, D-7, P-3, L-G2, M-9, V-5, M-10, M-16 · *deliverable* `results/arena/VERDICTS.jsonl`, `docs/apparatus/ADJUDICATOR_SPEC.md`, the calibration column · *evenings* 1.5

S-64 · The E1 escalation to $N = 16$, only under the rule — phase 4 · prereq S-63 · independent-of S-65 · parallel-safe · aliases M-4.1, R-19 (kill half), ledger P4.1 *build* if ChaCAL-with-sink is within MDE_8 of the shape on at least 6 of 8 **and** the residual ratio is within $2\times$,

escalate the pair to $N = 16$ (seeds 8 to 15, same lane, same hash); MDE_{16} from the $n = 16$ column (0.025820 at sd 0.034451). Never “within TOST” at $N = 8$: at that count two bit-identical arms return NO VERDICT (M-13). *prove* nothing. *measure* the contrast at $N = 16$ against MDE_{16} . *PASS* separated at $N = 16$ by more than MDE_{16} : the boundary-row mechanism is creditable. *KILL* not separated at $N = 16$: component (e) is a parameterisation of a column sink, the paper’s delta collapses to “ChaCAL plus a certificate plus Lean containment plus a registered bed”, and every later card runs under that sentence (§5.10 tree B). *price* ≈ 44 s ($3 \times 8 \times 1.681 + 4$) DERIVED, up to ≈ 1.1 GPU-min if two arms are re-run in full · *mechanism* M-13, M-7, D-7, M-9 · *deliverable* appended to `results/arena/r16_t8_m8_k2.jsonl` · *evenings* 0.5

S-65 · The TOST parity half on $\mathcal{A} = \emptyset$ only — phase 4 · prereq S-63, R-21 MOVED (otherwise the half is softmax against softmax) · independent-of S-64 · parallel-safe · aliases M-4.4, R-22, V-18, ledger P4.4 *build* the only contrast two one-sided tests decide: shape at $\hat{\gamma}$ against ChaCAL-diag at the same $\hat{\gamma}$ with $\mathcal{A} = \emptyset$ ([108] [V]); the margin 0.5σ is fixed **here** (M-2); the realised paired sd chooses the tier — paired at $N = 36$ if the paired sd equals σ (power 0.8014, Monte Carlo with 200,000 draws), two-sample at $N = 70$ if pairing buys nothing (0.7975 at 69, 0.8049 at 70, half-width 0.2799σ). The sequential anytime-valid form is NOT FOUND at [V], so fixed- N Schuirmann stands. *prove* nothing. *measure* the realised paired sd over σ ; the 90 per cent interval against plus or minus 0.5σ . *PASS* the interval inside the margin: the identity half of B-E1 is confirmed by equivalence, not only by bitwise identity at construction. *KILL* the interval outside: the “same operator” sentence is withdrawn at trained $\hat{\gamma}$ and narrowed to the $\gamma = 0$ bind; and any “parity” sentence resting on $N = 8$ is struck regardless (M-13). *price* 119 s paired at $N = 36$ or 228 s two-sample at $N = 70$ DERIVED, both under 4 GPU-min · *mechanism* M-13, M-9, M-2, V-5 · *deliverable* `results/arena/tost_parity.jsonl` · *evenings* 1

R-21 · Score the γ likelihood-ratio verdict, the ablation and the mirror line — phase 4 · prereq S-61, S-62 · independent-of S-64, S-65 · parallel-safe · aliases M-4.3 (verdict half), V-12, ledger P4.3 *build* nothing new; S-61’s instrument applied to the eight seeds. *prove* nothing. *measure* Λ per seed; the MOVED count; the ablation $(I - \hat{\gamma}\hat{P})^{-1} \rightarrow I$ at trained weights in seed-sd units; the support of $\hat{\gamma}$. *PASS* (Bet C) MOVED on at least 6 of 8 and the ablation moving NRMSE by more than one seed sd. *KILL* PINNED ($\Lambda \leq 2.7055$) on at least 6 of 8: the arm is softmax wearing a name on that bed, the horizon-dial sentence retires to a definition, ChaCAL-with-sink at $\hat{\gamma} = 0$ is unrunnable so S-64 is void, and the LM γ claim is retired before it is made. Mirror kill: $\hat{\gamma} > 0.99$ on at least 6 of 8 prints $1/(1 - \hat{\gamma}) > 100$ and the certificate is vacuous on the z and Δz channel. The committor head is untouched either way: it reads at $\gamma = 1$ on $\hat{Q}$. *price* 0 GPU-s (a journal read); the ablation ≤ 17.4 s · *mechanism* Ruling 10-prime, V-9, M-20, V-17, P-8 · *deliverable* `docs/plan/GAMMA_VERDICT.md` · *evenings* 1

M-4.2 · Score Bets A, B and E from the journal — phase 4 · prereq S-62 · independent-of S-64, S-65, R-21 · parallel-safe · aliases V-10, V-11, V-13 (scoring halves) *build* nothing new; the three deciding numbers with their frozen bands: the paired cosine difference against MDE_8 computed on the **cosine’s own** realised sd (V-17, never the NRMSE’s); the argmin interval against the exact restricted-view floor, with the straddle as SPLIT; the harmonic residual ratio with ≥ 10 confirming, ≤ 2 refuting and (2, 10) SPLIT, both unit factors printed. *prove* nothing. *measure* as above, paired by seed and draw. *PASS/KILL* exactly as filed in §5.5 before the first cell. A ratio of at least 10 at **unequal** marginal error is a statement about which mode the control mislearns, never about joint consistency, and is reported as such (`judge/sec_obstructions.md` §5.1). *price* 0 GPU-s · *mechanism* D-7, L-SIGN, V-26, V-17, M-2 · *deliverable* the calibration column of `docs/CALIBRATION.md` · *evenings* 1

S-42 · Z-BETA the row-sum guard on trained cells — phase 4 · prereq S-62 · independent-of S-63 to S-65 · parallel-safe · aliases ledger X-12, Z-BETA *build* census lines 12 and 13 on every trained cell: the measured support of $\hat{\beta}$ and of $\text{rowsum}(\hat{P})$; at $\hat{\beta} \neq 1$ rows sum 1.31 to 10.29 (`READ V16_ARM_SMPRIME.md:28–32`), where no row-stochastic theorem applies and the certificate field is written **refused**, never a number. *prove* nothing. *measure* the two supports per cell. *PASS* every

cell with a maximum row sum above $1 + 10^{-9}$ has `delta_vec = refused`. *KILL* a δ printed on such a cell (V-25): that certificate is struck and every sentence resting on it is withdrawn. *price* 0 GPU-s · *mechanism* V-25, V-3 · *deliverable* docs/apparatus/CERT_ROWSUM_GUARD.md · *evenings* 0.5

M-4.5 · File the leap dossier and the two honest sentences — phase 4 · prereq S-63, S-64, S-65, R-21, M-4.2 · independent-of nothing · not parallel-safe *build* the Phase-E template with the beds replaced (READ CEQ_V20_R15_CONTRACT.md:153-169): candidate built with identity binds, manifests, certificates, parameter-matched; verdict by the registered clauses; purge; prognosis. Both sentences of §5.10 with every bracket filled from a journal row or left NOT MEASURED. *prove* nothing. *measure* the count of brackets filled from `results/` rows against the count of brackets. *PASS* every number in the dossier points at a `results/` row. *KILL* a bracket filled from prose (P-1, P-2): the dossier is not filed and the paper says which bracket is empty. *price* 0 GPU-s · *mechanism* P-1, P-3, D-7, L-G2 · *deliverable* docs/CEQ_SHAPE_DOSSIER.md · *evenings* 1

PHASE 5 — the cost law, memory, and the kernels that do not exist

Goal. Make every price at $s > 64$ a measured law rather than an assertion, decide whether the serial solve is the shipped route, and write the specifications for the four instruments the paper marks NOT MEASURED so that a later evening can build them. Nothing in this phase is on the critical path to a verdict.

The gate (K-9). No price is quoted at any $s > 64$ until S-66 runs with two `synchronize()` calls and a randomised order, and the fitted exponent’s confidence interval is printed. If the interval excludes 2 toward 3, the serial solve is not the shipped route past $s = 256$, N-11 moves onto the critical path ahead of every long-context bed, and every long-context price becomes a band with direction.

S-66 · The synchronised, order-randomised s -sweep — phase 5 · prereq N-02, M-2.0 · independent-of every Phase-3 and Phase-4 card · parallel-safe · aliases M-5.1, N-08, R-14, V-21, ledger P5.1 *build* the harness edits the record priced and never made: a sequence-length flag (the constant lives at `scripts/v15_r1.py:137` and no `seq_len` occurs, `RUN[MERCURY] grep`), one substitution, two `torch.cuda.synchronize()` calls, **and** a randomised or blocked cell order — the fifth edit no office ever priced. Points $(s, n) \in \{(64, 2048), (256, 2048), (1024, 128), (4096, 8)\}$, with n declared per s from §5.7’s residency model; the dense operator at $n = 2048$, $s = 4096$ would be 137 GB against 7.996 GiB and is refused. *prove* nothing. *measure* seconds per step and peak reserved bytes per (arm, s , n), $N = 8$; the fitted exponent in s with its interval; the Spearman correlation of seconds against run index. *PASS* the exponent’s interval contains 2 and excludes 3, and the run-order correlation is insignificant (the record’s was $+0.7029$, $p = 0.0024$). *KILL* (K-9) the interval excludes 2 toward 3; the solve does not beat the quadratic attention product at the s where the paper claims it; or **any price was quoted before this ran** — those prices are struck. *price* the record’s band 206 to 537 GPU-s, band only, no point · *mechanism* D-3, P-8, C17, M-3, K-9 · *deliverable* results/cost/s_sweep.jsonl, docs/COST_LAW.md · *evenings* 3

N-09 · Where depth s becomes visible: the solve-to-product ratio across s — phase 5 · prereq S-66 · independent-of N-10 · parallel-safe *build* from S-66’s journal, the ratio of solve time to PV time per s ; the MAC prediction is plus 50 per cent on causal MACs ($3s^2d/2$ against s^2d , DERIVED), depth s against depth 1. *prove* nothing. *measure* the ratio at the four s . *PASS* ratio at or below 1.5 at $s \leq 256$ (arithmetic-bound). *KILL* ratio above 3 at any $s \leq 1024$: the triangular-solve latency chain dominates, the chunked solve (N-11) is required before any long-context row, and the “plus 50 per cent” sentence is restricted to the s where it was read. *price* inside S-66 · *mechanism* M-3, P-8 · *deliverable* docs/COST_LAW.md depth row · *evenings* 1

N-10 · The Neumann crossover, if any, across s — phase 5 · prereq S-66 · independent-of N-09 · parallel-safe *build* hop timings for $K \in \{1, 2, 4, 8, 16\}$ per s ; the hops needed for $\delta = 10^{-6}$ are $K \geq \log(\delta(1 - \gamma))/\log \gamma - 1$, which reads 19.9 at $\gamma = 0.5$ and 152 at $\gamma = 0.9$ — the cost section printed 129 at $\gamma = 0.9$ and the discrepancy is **recorded**, the formula being the deliverable. *prove* nothing; the certificate is J-L12’s, attained on the class and declared definitional. *measure* the smallest K at which K hops beat the solve, per s . *PASS* no such K at any s : “truncation never wins” holds in wall-clock as in MACs. *KILL* a crossover at some s : the Neumann route is a cost path there, its $\delta \|V\|_\infty$ is printed with $1/(1 - \hat{\gamma})$ beside it, and the F1 mask certificate wakes on that s . *price* inside S-66 · *mechanism* L-CERT, V-17, V-3, P-8 · *deliverable* docs/COST_LAW.md crossover row · *evenings* 1

M-5.2 · Refit the cost law with s as a variable and re-price every card — phase 5 · prereq S-66 · independent-of M-5.3 · parallel-safe *build* $s/\text{step} = an^bs^c$ per arm with its R-squared and the interval on c ; the increment law; the memory operator constant re-solved against the allocator, retiring the [ASSUMED] band. *prove* nothing. *measure* R-squared per law; which [ASSUMED] tags the fit closes. *PASS* every shape price carries a fitted law with R-squared at least 0.99 and no [ASSUMED] remains on a Phase-3 or Phase-4 card. *KILL* R-squared below 0.99 at any s : that s is quoted as a measured point, never as a law — the record’s own gate refused points fitted through swap (READ COSTS.md:76–79). *price* 0 GPU-s · *mechanism* M-8, P-8, V-22, P-1 · *deliverable* docs/COST_LAW.md, docs/PRICE_LEDGER.md revised with supersede markers · *evenings* 1

N-18 · Re-solve the memory operator constant for the shape — phase 5 · prereq S-20 · independent-of S-66 · parallel-safe *build* the sizing sweep at the module’s own settings over sequence lengths 128 to 2048, forward and backward, peak allocator bytes, run on the shape at the softmax corner; a two-parameter least squares separating the residual and operator constants — the record’s 17.874 and 3.823 at R-squared 0.996373 were fitted on the **signed** arm, so carrying 3.823 to the shape is M-8. *prove* nothing. *measure* the two constants for the shape with their R-squared, and the shape-to- softmax byte ratio per sequence length. *PASS* R-squared at least 0.99 and the ratio growing with s (the class is quadratic in s , as $1.44\times$ to $8.06\times$ was for the signed arm). *KILL* measured over predicted outside $[0.9, 1.1]$: the module’s formula is declared wrong for this arm, as it was for the complex arm, and the shape’s memory is quoted only from measurement. *price* ≈ 9 GPU-min DERIVED · *mechanism* M-8, V-22, P-8 · *deliverable* results/k_cert_local.json memory law, COSTS.md §1.2 · *evenings* 2

N-19 * · The bf16 gap: does solve_triangular have a bf16 CUDA path — phase 5 · prereq none · independent-of every card · parallel-safe *build* a one-line attempt on bf16 CUDA operands under autocast, recorded as executable or raising — the claim “no bf16 path in torch 2.5.1” is [U] and is settled here; then the bytes-per-element table for the shape under autocast beside the record’s (operator 3.341 B per element at R-squared 0.999830; residual 2.383 against the module’s 2.2, wrong by 8.3 per cent in the optimistic direction, READ COSTS.md:91–99). *prove* nothing. *measure* executability; the operator term for the shape under autocast. *PASS* (the expected reading) the solve raises or upcasts and the shape’s operand stays at 4.0 bytes per element, printed as its own row. *KILL* (of the [U] sentence) a bf16 path exists: the row is refitted and the 4.0 sentence is deleted. *price* seconds · *mechanism* P-3, V-22, M-8 · *deliverable* COSTS.md §1.2 autocast row · *evenings* 1

N-20 · The n that does not fit: residency of the shape at large n — phase 5 · prereq N-18 · independent-of S-66 · parallel-safe *build* the residency table for the shape: the workhorse arm reserves 10.578 GiB at $n = 16384$ against 7.996 GiB and pages over PCIe, which is why Ruling 8 dropped that reading and why the throughput law was fitted on three points rather than five; the shape on the softmax corner inherits the softmax rows (resident to $n = 32768$ at 4.908 GiB reserved) plus 8 MiB per 2048 for z plus the solve’s transients. *prove* nothing. *measure* allocated and reserved GiB per n under a training loop with synchronize. *PASS* resident at $n = 16384$. *KILL* reserved above 7.996 GiB at $n = 16384$: that reading stays dropped for the

shape as for the arm, and any law point fitted through swap is refused (P-8). *price* ≈ 3 GPU-min DERIVED · *mechanism* P-8, Ruling 8, V-22 · *deliverable* COSTS.md §1.3 · *evenings* 1

N-21 · The sizing model’s under-prediction for complex arms, and what the corner-3 base inherits — phase 5 · prereq N-18 · independent-of N-20 · parallel-safe *build* the complex-arm audit as a card: measured over predicted 1.842, operator constant 7.50 at 8 bytes per element ($4.29 \times$ softmax’s, not the $2 \times$ a “complex is two floats” argument gives), the residual constant NOT IDENTIFIED at $s = 64$. The shape on the corner-3 base inherits this and carries **no certificate** ($\hat{\beta} \neq 1$ rows sum 1.31 to 10.29). The card either re-solves the complex constants with a sweep or declares that base uncertified for sizing. *prove* nothing. *measure* measured over predicted across sequence lengths; the operator constant. *PASS* constants re-solved with R-squared at least 0.99 and the ratio stated per length. *KILL* R-squared below 0.99 or the ratio outside $[0.9, 1.1]$: the corner-3 base is journalled “uncertified for sizing” and no long-context card runs on it. *price* ≈ 5 GPU-min DERIVED (the corner-3 cell is about $9.6 \times$ softmax’s) · *mechanism* M-8, V-22, P-8, V-25 · *deliverable* COSTS.md §1.2 complex row, MODEL_CARD.md limits line · *evenings* 1

M-5.3 · Spend the 9,600-step ladder on a survivor only — phase 5 · prereq S-63 to S-65 with the shape alive on at least one creditable contrast · independent-of S-66, M-5.2 · not parallel-safe *build* the ladder (64 times the steps) for the surviving pair only. *prove* nothing. *measure* the same statistics at the ladder’s end; the crossing and residual curves. *PASS* the $N = 8$ verdict survives the ladder within its own MDE₈. *KILL* a verdict that reverses at the ladder is filed as **the ladder’s**, and the 150-step cell is marked a pilot (M-3); what dies is the 150-step sentence, not the bed. *price* 32 to 36 min per pair [FITTED] at 64 times the law · *mechanism* M-3, M-6, P-8 · *deliverable* results/arena/ladder.jsonl · *evenings* 1

N-11 · Specify and time the chunked block-triangular solve — phase 5 · prereq N-09 · independent-of N-13, N-18 · parallel-safe · alias M-5.4 *build* the recursion $M_{kk}z_k = v_k - \gamma \sum_{l < k} P_{kl}z_l$ with s/C diagonal triangular solves and one product per chunk: about $s^2d/2 + sCd/2 + sC^2/3$ MACs, depth s/C , retained blocks sC (DERIVED). Invariant: parity against the serial solve at 10^{-6} in float32; journal route = chunked with C . The intra-chunk primitive is DeltaNet’s transform pattern and is **owned**; the delta is stated narrowly: no finite-state dual exists for a softmax P ([58] [V]), so the inter-chunk term is a genuine product and the cost stays quadratic **because** the shape keeps softmax bitwise at $\gamma = 0$. *prove* nothing new; **resolvent_fromBlocks** [M] (J-L6) is the block identity. *measure* parity; wall-clock at $s \in \{1024, 4096\}$ against the serial solve for $C \in \{32, 64, 128\}$. *PASS* parity at or below 10^{-6} at every C , and at $s = 4096$ wall-clock at or below the serial solve. *KILL* slower at every C : the serial route is the shipped one and this card closes as NOT NEEDED on this device; parity failure: the kernel is refused (L-CERT — an uncertified route ships no δ). *price* NOT MEASURED, **needs a chunked kernel**; first timing is minutes once built [ASSUMED] · *mechanism* M-8, M-3, P-4, P-7 · *deliverable* docs/CEQ_CHUNKED_SOLVE.md · *evenings* 3

N-12 · The recompute backward: leaving the quadratic memory class — phase 5 · prereq N-11, N-18 · independent-of N-13 · parallel-safe *build* the backward of N-11 rebuilding off-diagonal tiles from the query and key projections instead of retaining P ; the adjoint of the solve is one upper-triangular solve plus an outer product (DERIVED); retained memory linear in s plus z . *prove* nothing. *measure* peak bytes against sequence length; gradient parity against autograd through the retained route at 10^{-5} in float32. *PASS* peak bytes grow linearly within the sizing law’s R-squared and parity holds. *KILL* bytes still quadratic: the recompute is not reached, the shape stays in the operator class, and the resident n per s from N-20 is the hard limit of every long-context card. *price* NOT MEASURED, **needs the kernel** · *mechanism* P-4, V-22, P-8 · *deliverable* docs/CEQ_CHUNKED_SOLVE.md backward section · *evenings* 2

N-13 · The published blockwise subquadratic evaluation as the occupied control — phase 5 · prereq N-09 · independent-of N-11 · parallel-safe *build* the control arm evaluating

the same operator as [143] [V] does — exact triangular solves on diagonal tiles with cross-block interaction through a reduced system — as a **fellow** at $s \in \{1024, 4096\}$ with n per s ; its own δ is unprinted in the source and is therefore **measured here** where a dense control fits, and **NOT MEASURED** where it does not. *prove* nothing; the shape’s exact block structure is J-L6’s F0 and holds only on gated corners or masked P . *measure* visited-tile fraction, wall-clock, φ -NRMSE, harmonic residual, the measured δ . *PASS* the blockwise read within MDE_8 of the dense read at lower latency, reproducing the source’s band in kind. *KILL* the shape’s F0-segmented solve not better than the blockwise control by MDE_8 at matched visited tiles: the subquadratic path is theirs, cited, and the shape claims only exactness beside it — never “beats”. *price* **NOT MEASURED**, **needs the blockwise control implemented** · *mechanism* R-SKY, MISTAKES.md D-1, V-9, P-8, L-CERT · *deliverable* results/cost/blockwise_control.jsonl · *evenings* 2

N-15 · The CSR two-stage path with the union certificate; the F1 mask bind wakes — phase 5 · prereq N-03, N-10, S-41 · independent-of N-11, N-13 · parallel-safe · aliases S-43, S-70, V-20, R-10, ledger P5.2, Z-F1 *build* stage 1, the schedule as data: a causal list of tiles per query block, with the do-nothing zero-dimensional salience schedule **always entered** (READ THEORY.md:162–165); stage 2, Neumann on the sparse P with the union certificate in vector units $(\varepsilon/(1-\gamma) + \delta_{\Pi}) \|V\|_{\infty}$ — the resolvent amplifies dropped row mass by $1/(1-\gamma)$ — and the **exact** solve on an F1 mask refused because the inverse fills in along reachability. The two structural guards stay: an empty schedule row must not return a division by zero, and a negative block index must be clamped (READ ceq/mz_kernel.py:13–21). *prove* mask_amplification [S], f1_cantelli_union [D] (J-L12, J-L17). *measure* on the shipped mask, the gap over 1,024 drawn cells against the printed bound with $\|V\|_{\infty}$ printed; the fill-in guard on a planted schedule. *PASS* no exceedance on any of 1,024 draws, $\delta \|V\|_{\infty}$ below the label sd, and the guard fires. *KILL* (K-I) one exceedance: the mask is refused (L-CERT) and the exact solve — cheaper than one hop at $s = 64$, 2.514 against 3.001 ms — is the path; $\delta \|V\|_{\infty}$ at or above the label sd: the certificate is uninformative and the row returns to dormant. *price* ≈ 2.6 s at $s = 64$ DERIVED; at $s = 4096$ **NOT MEASURED** · *mechanism* L-CERT, V-10, V-17, V-24, V-9 · *deliverable* docs/CEQ_CSR_PATH.md, results/certs/mask.jsonl · *evenings* 2

N-16 · The Mapper cover to tile quantiser — phase 5 · prereq N-15, S-44 · independent-of N-17 · parallel-safe · aliases S-70 (part), M-6.3 T-2, V-23 *build* a cover of key positions from a lens (the committor or an influence score) with parameters fixed by the Reeb-estimator rule on a **held-out** relation ([21] [V]), never on the bed (M-2); nerve edges quantised to contiguous tiles **before** they are a schedule — a gather-realised candidate set pays 2.58 ms of index-select against 0.82 ms of attention at 65,536 positions on this card (READ ceq/multizoom.py:12–16); the far field carries the only printed bound in the tree, the mean-pool coarsening bound (READ ceq/multizoom.py:38–49). Cluster covers as candidate builders are occupied and cited; the delta is the nerve of a lens cover feeding a certified resolvent, stated as such. A Mapper cover with a comparable δ is **NOT FOUND**. *prove* nothing. *measure* visited-tile fraction and φ -NRMSE against **two** do-nothing controls at matched visited tiles. *PASS* better than both by MDE_8 with the union $\delta \|V\|_{\infty}$ below the label sd. *KILL* (K-M) not better than do-nothing, or than the blockwise control: the candidate builder is the merged salience schedule, Mapper adds nothing measurable, and the corpus’s segment-length distribution is the only dividend left. *price* **NOT MEASURED**, **needs the quantiser and N-15 at long \$\$\$** · *mechanism* V-9, M-2, L-CERT, V-22, P-4 · *deliverable* docs/CEQ_MAPPER_SCHEDULE.md · *evenings* 2

N-17 · The backward through the scheduled kernel: the forward-only limit — phase 5 · prereq N-15 · independent-of N-16 · parallel-safe *build* the adjoint of the scheduled attention — the same schedule read transposed, the online-softmax statistics saved per query block, and for the resolvent stage the upper-triangular adjoint of N-12 — with both structural guards preserved. Until it exists the scheduled path is inference-side (READ THEORY.md:223–224), and the published speedup on that path is the PR page’s own number, not re-measured here. *prove* nothing. *measure* gradient parity at 10^{-5} in float32 on a dense schedule; forward and backward wall-clock against

dense attention at $s = 4096$. *PASS* parity holds and the guards fire on planted empty rows and negative indices. *KILL* parity fails: the scheduled path stays inference-only and every training row above $s = 64$ runs the chunked solve or the dense route with n per s . *price* NOT MEASURED, **needs the kernel** · *mechanism* P-4, V-16, L-CERT · *deliverable* docs/CEQ_CSR_PATH.md backward section · *evenings* 3

J-D5 · Specify the directed-stability δ for the influence barcode — phase 5 · prereq J-D4’s notation, J-L6 · independent-of every Lean card · parallel-safe · aliases S-71, M-6.3 T-1, ledger T-1, P5.4 *build* the barcode row is **killed** on BED-S — the influence graph at threshold zero is one weak component on 100 per cent of softmax draws, and the tree’s interleaving instrument consumes point clouds (READ ceq/certs/topological.py:476–505) — and survives only on gated corners with exact zeros. What is owed before any barcode is read: (i) the filtered object, the influence matrix $J_{ij} = \partial O_i / \partial V_j = (\Pi_\gamma)_{ij} \geq 0$, a **directed** weighted graph and not a metric; (ii) the connectivity notion, weak components of the thresholded digraph ([29] counts them at degree 0); (iii) the dissimilarity transform $d_{ij} = -\log J_{ij}$ so a filtration of an asymmetric function is defined; (iv) the stability constant — [117] gives bottleneck stability for asymmetric functions, and the symmetric-hypothesis stability theorems do **not** apply — the printed δ being that constant times the sup-norm perturbation, NOT MEASURED, **needs the constant evaluated on the shipped filtration**; (v) the null, row-permuted influence at the same logit scale, at least 200 permutations; (vi) the threshold-zero endpoint equal to the F0 segmentation count by torch.equal. *prove* nothing. *measure* the six-part predicate written before any barcode is computed. *PASS* the predicate is written first. *KILL* the measured barcode inside the null’s 95 per cent band on at least 6 of 8 seeds: decoration; a non-integer or aliased component count: refusal (V-16). *price* 0 GPU-s for the spec; the instrument NOT MEASURED · *mechanism* M-15, M-2, V-16, V-3, V-25, P-4 · *deliverable* docs/apparatus/DIGRAPH_BETA0_SPEC.md · *evenings* 2

S-72 · The Perron weight for a committor certificate — phase 5 · prereq S-40 · independent-of N-1x, J-D5 · parallel-safe · aliases ledger P5.7, Z-COM *build* the contraction certificate needs a weight w with $Qw \leq \rho w$ and $\rho < 1$ for the **oracle’s** non-causal chain; on the arm’s causal $\hat{Q}$ the diagonal read of J-L10 suffices and no weight is needed. Until w is measured, every committor cell ships on the exact route only. *prove* isUnit_one_sub_of_perron [S] (J-L15). *measure* the weight, if it is ever computed; until then the sup-norm substitute is recorded as vacuous (19.2 at $K = 4$ against a true error 2.5×10^{-3}). *PASS* a weight with its two inequalities printed. *KILL* a truncated committor read journalled without w is refused; what dies is any “certified committor at a truncation” sentence. *price* NOT MEASURED, **needs the Perron weight** · *mechanism* V-17, V-10, L-CERT · *deliverable* docs/apparatus/PERRON_WEIGHT.md · *evenings* 1

J-D8 · Attempt the hop-to-committor reduction, or file it NOT FOUND — phase 5 · prereq J-D4 · independent-of every card · parallel-safe *build* one page: given a hop_k instance, construct logits whose causal softmax puts mass $1 - \eta$ on the pointer target, declare the k -th target set absorbing, and show the committor into it is at least $(1 - \eta)^k$ while the committor into any other set is at most $1 - (1 - \eta)^k$, so the argmax over sets reads hop_k . What blocks a theorem: the softmax row also puts mass on itself (J-L5) and on the sink (J-L2), so the walk can stall or fall; the construction must bound the stall mass, and that bound depends on the logit scale, which the arms **learn**. *prove* either the reduction with its stall-mass bound, or NOT FOUND with the blocking step named. *measure* nothing. *PASS* either outcome written. *KILL* the reduction stated **without** the stall-mass bound (V-3, P-10); then Obstruction 2 remains a lineage argument with its vacuity lines, which is what the paper says today. *price* 0 GPU-s · *mechanism* V-3, D-2, P-10 · *deliverable* docs/CEQ_SHAPE.md §5.2 · *evenings* 2

PHASE 6 — the envelope, the package, the Kaggle gate

Goal. Make the local half of the gate complete so that, if and when the author says yes, nothing is blocked on work an agent could have done locally.

The gate. Every local certificate line filled (N-01, N-02, N-18), the envelope’s two plants firing, the package’s smoke passing, the pinned digests re-asserting on regeneration; the open rulings listed by name beside the request. **The request is a notification to the author, never a launch.** The standing rule is the author’s: no launch, upload or training start without an explicit yes.

N-22 · The flight envelope as the run’s guard, written for the shape lane locally — phase 6 · prereq S-01 · independent-of every card · parallel-safe *build* Ruling 6’s clauses as fields every shape cell carries, plus the three run-guards the record already owns: the zero-step RED gate ($\text{GATE_TOL} = 10^{-3}$; 30 of 30 shapes passed with worst margin 1.957×10^{-3}), the bar re-certification HALT line (delta over tolerance below 50 per cent; the record’s worst was 8.58 per cent with $5.83\times$ headroom), and the thread lane read from the journal and never from the machine. *prove* nothing. *measure* two planted faults — an out-of-memory-sized shape and a bar drifted past the halt line — must abort and halt respectively, each with a journal row. *PASS* both plants fire and an honest run passes both. *KILL* a plant that does not fire: the envelope is a condemning rule with no planted negative (V-15) and no launch request is made until it does. *price* seconds · *mechanism* Ruling 6, V-15, V-16, M-16, D-4 · *deliverable* docs/CEQ_ENVELOPE.md · *evenings* 1

N-23 · The package: replace the dead operator with the shape, match the count, keep the card honest — phase 6 · prereq S-20, S-01 · independent-of S-66 to N-21 · parallel-safe *build* the attention module becomes ChaCAL-diag with declared boundary rows and the manifest’s route and diag_convention fields; the smoke script asserts parity at $\gamma = 0$ against the package’s own softmax attention; the parameter count is recounted per arm with the γ scalar and any head included (Ruling 3: exact counts in every table header, never re-architect to close 0.032 per cent); the model card is rewritten limits-first, with §5.10 tree B’s sentence as its capability paragraph until a cell exists. No checkpoint ships. *prove* nothing. *measure* parity, the counts, one forward at $s = 64$ on CPU. *PASS* parity True, the counts printed and within 0.032 per cent, no checkpoint claimed. *KILL* parity fails: the package ships no shape; a card sentence with a bracket filled by anything but a FOUND cell is struck (P-1). *price* 0 GPU-s · *mechanism* P-3, Ruling 3, V-3, P-1 · *deliverable* the package edits and MODEL_CARD.md · *evenings* 2

N-24 · The Kaggle gate, local half only — phase 6 · prereq N-01, N-02, N-18, N-22, N-23 · independent-of nothing (it is the DAG’s sink) · not parallel-safe · alias M-6.1 *build* the local half of Gate 0 only: the three bed pins carried as generator, seed and hash and regenerated in-notebook with the hash asserted (Ruling 7); the corpus pin with its split; the three unpinned sources whose hash cell prints and then **raises**; the code pin, which is the snapshot the run executes and may differ from the branch tip. The Kaggle certificate slot is empty and stays empty until the author’s run; a threshold carried across the device boundary is V-22. The gate’s own circularity — its first rows are blocked on numbers only the Kaggle run produces — is the author’s ruling and is not taken here. The launch order is unchanged, and its only consumer in this programme is the LM γ cell. *prove* nothing. *measure* locally: that N-01, N-02 and N-18 have filled their certificate lines, that N-22’s plants fire, that N-23’s smoke passes, and that the pinned digests re-assert on regeneration. *PASS* every local line filled and the open rulings listed by name beside the request. *KILL* any Kaggle number pooled with a local one (V-22); any launch, upload or training start without the author’s explicit yes; a missing-pin raise bypassed. *price* 0 GPU-s locally; on Kaggle the certificate run is ≈ 8.7 GPU-min [ASSUMED] equal to the local 523.9 s · *mechanism* V-22, P-1, Ruling 7, the Kaggle-yes rule · *deliverable* docs/KAGGLE_GATE.md · *evenings* 1

S-73 · The LM γ cell on the pinned slice, behind the author’s explicit yes — phase

6 · prereq S-61, S-63, R-21 MOVED, N-24, **the author's yes** · independent-of nothing · not parallel-safe · aliases M-6.2, V-22, ledger P5.8 *build* one LM pair — shape against ChaCAL at fixed $\gamma = 0.9$ — on the pinned slice, with $\hat{\gamma}$ per seed under the boundary null; the language-model row of the predictions file is filled only then. *prove* nothing. *measure* Λ per seed; the perplexity delta against ChaCAL against MDE₈. *PASS* MOVED on at least 6 of 8. *KILL* PINNED on at least 6 of 8 — consistent with ChaCAL's own language-model result being worse than its baseline (digits [U]): the LM γ claim is retired and γ stays a bed-side dial. *price* NOT MEASURED — no LM cell of the shape has ever run and the Kaggle certificate slot is empty · *mechanism* V-22, Ruling 10-prime, V-9, the Kaggle-yes rule · *deliverable* `results/lm/gamma.jsonl` after the yes · *evenings* 2 after the yes

5.4 The critical path, and the evenings count derived from it

The path, as a chain of card ids with each card's evenings. The binding constraint is L-LEAN: the arm may not be **trained** before its identity theorems are green, and the theorems S-20 rests on are J-L1 (parity) and J-L3 (the unit), reached through J-L0 and certified by J-L18. Everything else either runs beside that chain or is shorter than it.

$$\underbrace{\text{J-L0}}_1 \rightarrow \underbrace{\text{J-L3}}_1 \rightarrow \underbrace{\text{J-L10}}_1 \rightarrow \underbrace{\text{J-L18}}_1 \rightarrow \underbrace{\text{S-20}}_1 \rightarrow \underbrace{\text{S-21}}_1 \rightarrow \underbrace{\text{S-26}}_1 \rightarrow \underbrace{\text{S-60}}_1 \rightarrow \underbrace{\text{S-62}}_1 \rightarrow \underbrace{\text{S-63}}_{1.5} \rightarrow \underbrace{\text{S-64}}_{0.5} \rightarrow \underbrace{\text{S-65}}_1 \rightarrow \underbrace{\text{M-4.5}}_1$$

$$1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1.5 + 0.5 + 1 + 1 = \mathbf{13} \text{ evenings.}$$

The one branch that lengthens it. The bed branch S-10(0.5) → S-11(2) → S-12(1) → S-16(1) = 4.5 must finish before S-21, and it runs beside the four-evening Lean gate, so it adds 0.5: **13.5, called 14 evenings to a verdict and a filed dossier.** Add M-5.3 (the ladder on a survivor, 1) for **15**; add the cost-law branch S-66(3) → M-5.2(1) if a cost law is demanded before the ladder, for **19**.

What runs beside it, and therefore costs no additional evening. M-2.0 (build the arm, 3) — it must precede S-20 but is shorter than the Lean gate. The [S] Lean set J-L11 to J-L17 (9 evenings) — it changes the paper's **grades**, never the training gate. Every J-D derivation (8 evenings, all parallel-safe). Every Phase-1 card except the four on the branch above. S-52, S-51, M-3.7, M-3.8, V-17, N-01, N-02, N-07, N-19, S-40 to S-44. All of Phase 5's kernels, which have no prerequisite in the verdict chain at all.

Reconciling the six planets' own critical-path claims. Each planet counted the serial depth of **its own slice**, and the differences are exactly the slices, not a disagreement about the DAG.

planet	its claim	what it counted	where it overlaps this plan's 14
JUPITER (Lean)	11 evenings: J-L0 → J-L3 → J-L10 → J-L13 → J-L12 → J-L15 → J-L16 → J-L18	the chain to the full [S] set, not the training gate	its first four evenings (J-L0, J-L3, J-L10, J-L18) are the training gate and are on the merged path; the other seven are beside it

planet	its claim	what it counted	where it overlaps this plan's 14
SATURN (instrument)	11 evenings: S-01 → S-11 → S-12 → S-14 → S-20 → S-21 → S-60 → S-61 → S-62 → S-63 → S-64	the bed-and-verdict chain with no Lean gate in front of it	eight of its eleven are on the merged path; it omits the four Lean evenings and folds S-26 into S-14
MERCURY (price)	12 evenings to the ladder, 15 with the cost law	M-1.1 → M-1.2 → M-1.4 → M-2.4 → M-3.1 → M-3.2 → M-3.3 → M-4.1 → M-4.4 → M-5.3	the same bed-and-arena spine under different ids; its M-2.0 (3 evenings) is beside, as here
MARS (attack)	7 evenings: R-04 → R-07 → R-09 → R-18 → R-19 → R-22	the attack chain only, with the Lean gate explicitly beside it	its six are a subset: the bed spec, one census attack, the binds, the first pair, the arena, the equivalence test
VENUS (prediction)	9 evenings: V-0 → V-1 → V-3 → V-5 → V-10 → V-15 → V-18 → V-19	the prediction-filing and scoring chain	V-1 is S-60, V-10's pair is S-62, V-15 is S-31 plus S-64, V-18 is S-65, V-19 is S-63
NEPTUNE (systems)	18 evenings: N-03 → N-05 → N-06 → N-08 → N-09 → N-11 → N-12 → N-15 → N-17 → N-16	a serial depth through kernels that do not exist ; ten of its eighteen are Phase-5 specifications	only N-05 (which is S-20) and N-06 (which is S-62's price) touch the merged path; the rest is off it by construction

The six agree on the spine — bed specification, census, binds, first pair, arena, verdict — and differ only in what they put in front of it (JUPITER: Lean; NEPTUNE: kernels) and in whether they count the arm's construction as an evening. The merged number is **14 to a verdict**, and the iteration count of any loop that runs this plan is that depth, never an agent's choice (CONTRACT.md D-3).

5.5 The predictions ledger, with counters

Filed before any arena cell exists and hashed into `results/arena/PREDICTIONS.sha256` (S-60). Every row carries a counter of equal specificity (L-SIGN), the SPLIT band where one exists, and the single number that decides between them. Under the record's calibration — 9 checked, 9 adverse, 7 of 8 signed rows optimistic, one-sided sign test $p = 0.0352$ — **the counter is the point estimate** (D-CALIB-1), the sign is the only output (D-CALIB-3), and a bare prediction is refused at the file rather than discounted (D-CALIB-2). MDE₈ everywhere means the $n = 8$ cell of the **realised** paired sd, which BED-S has not produced; the figure 0.039827 at sd 0.034451 is the record's row-2 instance and is a placeholder until S-62 fills it (M-3).

bet	prediction (the half that flatters)	counter (the point estimate)	SPLIT band	deciding number	card	price	mechanism
A conse- quence field	$\cos_{\text{shape}} - \cos_{\text{softmax}}$ $\text{MDE}_g(\cos)$ on the paired test	$\cos_{\text{within}} > \cos_{\text{softmax}}$ $\text{MDE}_g(\cos)$: a per-row mixture learns the displace- ment coordi- nate by coordi- nate as well as the joint read	none (two- sided on the sign of the miss)	the paired cosine difference against the cosine's own realised sd	S-62, M-4.2	inside the ≈ 34 s pair	MISTAKES.md D-1, V-26, V-17
B argmin floor	$\text{CP}_{\text{upper}}(\text{err}_{\text{shape}})$ $\text{floor}_{\text{exact}}(k)$ 1) at $N = 8$	$\text{CP}_{\text{lower}} \geq \text{floor}_{\text{exact}}(k)$ 1): the joint read is no better than a one-hop window licenses	the straddle $\text{CP}_{\text{lower}} < \text{CP}_{\text{upper}}$ (the R1 precedent straddled 0.7071 with [0.617075, 1.041227])	the interval against the floor	S-62, M-4.2	inside the pair	L- FLOOR, V-10, M-2
C the horizon dial	$\hat{\gamma}$ MOVED ($\Lambda > 8.318$) on at least 6 of 8 seeds, and the ablation moves NRMSE by more than one seed sd	PINNED ($\Lambda \leq 2.7055$) on at least 6 of 8: softmax wearing a name on this bed	3/8 to 5/8, or interval verdicts	Λ per seed under the boundary null	S-61, R-21	0 GPU-s plus a ≤ 17.4 s ablation	Ruling 10-prime, V-9, M-20

bet	prediction (the half that flatters)	counter (the point estimate)	SPLIT band	deciding number	card	price	mechanism
D the depth skyline	the depth-5 stack within MDE_8 of the shape's argmin accuracy — the honest control holds	short by more than MDE_8 , sign logged: the bed needs some- thing the hop con- struction does not supply	none	$\text{acc}_{\text{shape}} - \text{acc}_{\text{sky}}$	S-35	≈ 64.8 s [ASSUMED]	MISTAKES.md D-1, R-SKY, P-10
E joint consis- tency	$r_{\text{softmax}}/r_{\text{shape}}$ 10 at marginal NRMSE within MDE_8	≤ 2 : joint consis- tency is learnable by a per-row mixture and the untested sentence at MATHEMATICS.md:106-127 is with- drawn	(2, 10)	the paired ratio with $\sigma_{\min}$ and the operator norm printed	S-62, M-4.2	inside the pair	V-26, V-17, D-2
F the boundary mecha- nism (K-E1)	the shape's commit- tor φ - NRMSE below ChaCAL- with- sink's by more than MDE_8 on at least 6 of 8, residual smaller by more than $2\times$	within MDE_8 and residual within $2\times$: com- ponent (e) is a parame- terisation of a column sink	escalate once to $N = 16$, never “within TOST” at $N = 8$	the paired φ - NRMSE difference and the residual ratio	S-31, S-64	≈ 52 s, escalation ≈ 44 s	V-24, M-13, M-7

bet	prediction (the half that flatters)	counter (the point estimate)	SPLIT band	deciding number	card	price	mechanism
G the move price	the Sherman- Morrison column route at or below $0.5 \times$ one solve for $m = 8$ ($m/d =$ 0.5 at ($8, 16$))	at or above $8 \times$: a re-solve per candidate is what ships	$(0.5 \times, 8 \times)$	the measured ratio	V-17	seconds	M-8, P-8, M-3
H the cached mixture (K-G1)	the cached arm fails on P - changing plants by more than one seed sd on at least 6 of 8, and matches on V -only plants	within one seed sd on P - changing plants: the re-solve is a per-row control wearing a name	none	the two per-seed cosine gaps in seed-sd units	S-33	≈ 35 s	D-2, P-7, V-9
I the capped seeds (the record's own open item)	all three capped seeds cross the one-hop threshold: diver- gence was the cause of NO READING	at most one crosses: the split is not the cap's and the mech- anism stays un- resolved	2 of 3	the crossing count out of three	S-52	≈ 9.65 s	M-6, D-6, P-3

bet	prediction (the half that flatters)	counter (the point estimate)	SPLIT band	deciding number	card	price	mechanism
J the certificate (dor- mant)	zero ex- ceedances on 1,024 draws and $\delta \ V\ _\infty$ below the label sd	one ex- ceedance, or $\delta \ V\ _\infty$ at or above the label sd: the mask is refused	none	the ex- ceedance count and the ratio	N-15	≈ 2.6 s when a mask exists	L-CERT, V-10, V-17
K the Lean set	10/10 [M] targets elaborate with the three- axiom set	at most 9/10: a row resting on an uncon- firmed Mathlib name is demoted to [S] and never cited as proved	none	the green count out of ten	J-L18	0 GPU-s	P-11, L-LEAN
L the cost exponent	the fitted expo- nent's interval in s contains 2 and excludes 3, and the run-order correla- tion is insignifi- cant	the interval excludes 2 toward 3: depth s is visible past $s = 64$ and every long- context price is a band	none	the exponent interval	S-66	206 to 537 GPU-s, band only	D-3, P-8, C17

bet	prediction (the half that flatters)	counter (the point estimate)	SPLIT band	deciding number	card	price	mechanism
M the LM dial	$\hat{\gamma}$ MOVED on at least 6 of 8 enwik8 seeds	PINNED on at least 6 of 8, consistent with the published opera- tor's own language- model result being worse than its baseline	none	the MOVED count	S-73	NOT MEASURED, behind the author's yes	V-22, Ruling 10-prime

The scoring rule. One ledger, two sources: an **author** row is a card above; a **leap** row is any prediction inside a leap output, filed [LEAP-UNTESTED] with its counter and its cheapest killer, its instance RUN **before** the row is scored. Verdict tokens are HOLDS, COUNTER, SPLIT, VOID; a VOID row (the bed not admitted, or a prerequisite instrument failing its plants) is unscored, never a miss; a SPLIT row is scored wrong with the sign of the optimistic half. The one-sided sign test runs on the pooled column after at least eight scored rows and on the two sub-columns separately. Rows are appended, never edited; a corrected reading is a new row whose **supersedes** points at the old one (L-G2). A column that stays at or above 0.75 optimistic after eight rows makes the counter the paper's **reported** estimate for every unscored card, not merely its planning estimate.

5.6 The standing attacks

Each is the cheapest run, with its threshold frozen here (M-2), that ends a claim sentence the judged sections permit. Three of them (R-07, R-08, R-24) were filed by no office before this plan and each has an $O(1)$ instance on the judge's draw, not a rate.

attack	target claim it ends	firing number (frozen)	price	mechanism
R-01 the [M] targets do not build	"machine- checked" for component (k)	any target not building by the Lean milestone is [S]; a sorryAx anywhere deletes the row	0 GPU-s	P-11, L-LEAN

attack	target claim it ends	firing number (frozen)	price	mechanism
R-02 the manifest does not discriminate	every attribution in the lane	any one-unit flip of <code>beta/qk/g/gamma/boundary_sets</code> leaving <code>manifest_hash</code> fixed	0 GPU-s	L-2, V-16
R-03 an obstruction without its vacuity line	every “softmax cannot” sentence	any obstruction stated without its inequality; 0 per cent of the bed admitted by the hypothesis	0 GPU-s	V-25, P-10
R-04 the bed dies at construction	every BED-S sentence	any label <code>sd = 0</code> ; any argmin class frequency outside (0.05, 0.95); disagreement = 0; lines 1 to 4 below 100 per cent	0 GPU-s	D-4, V-8, V-12
R-05 the label leaks at order 0	the one-read claim on BED-S	honest probe at or above 0.5; or the planted leak below 0.99 (a blind detector)	0 GPU-s	M-21, V-24, V-7
R-06 / K-H2 the bed is a static task	every capability sentence on BED-S	the 0-hop MLP inside the shape’s interval half-width at $N = 8$, or McNemar $p > 0.05$ against 1-hop softmax	≈ 0.5 GPU-min	C8, V-10
R-07 * the Chebyshev rule is minimised by falling off the prompt	the Chebyshev column, and the bed’s admission	the safest-move-is-the-sink fraction at or above 0.5 on the unrestricted bed; or the exact-zero-tie discard fraction at or above 0.5	0 GPU-s	V-12, V-8, V-10, V-5
R-08 no-op moves inflate the effective m	every printed argmin floor	$m_{\text{eff}} < m$ on more than 5 per cent of draws in a cell quoting a floor	0 GPU-s	L-FLOOR, V-17, V-10

attack	target claim it ends	firing number (frozen)	price	mechanism
R-24 a flag lookup decides the argmin	BED-S’s argmin head	the closed-form rule at or above $\text{acc}_{\text{shape}} - \text{MDE}_8$ on at least 6 of 8, or above the zero-information floor by more than its interval at construction	0 GPU-s	V-10, D-5, M-21
R-09 a bind has an empty rejection region	the proposition that bind carries	any plant passing the honest bar; Gaussian V reading below 10^{-6} ; a non-causal P reading zero displacement before the intervened row	seconds	V-24, V-14, V-3
R-12 the solve path is not deterministic or not priced as claimed	the “first gated wing under strict mode” sentence; every [FITTED + RUN] price	strict mode raises on the backward; or the cell reads above $2.2\times$ the law	≈ 1.7 s	V-16, P-8, M-3
R-13 the capped seeds do not cross	the record’s account of R1	at most one of three crossing under the cap	≈ 9.65 s	M-6, D-6
R-14 / K-9 the timer’s strongest correlate is run order	every GPU-second in the lane	a Spearman correlation of seconds against run index comparable to the record’s $+0.7029$; or any price quoted before the sweep	band 206 to 537 GPU-s	C17, P-8, M-10
R-15 a “matched” contrast is not matched	every “at matched parameters” sentence	a per-arm count differing by more than 0.032 per cent after the head is attached	0 GPU-s	Ruling 3, M-8
R-16 / K-D2 the arm copied the environment	every capability number on BED-S	$\ \hat{P} - P_{\text{env}}\ _{\infty} < 10^{-3}$ on at least 6 of 8 seeds	inside the pair	D-2, M-7
R-17 a guard partition postdates the itinerary it scores	every itinerary statistic	the partition’s hash later than the first itinerary record	0 GPU-s	M-2, M-18

attack	target claim it ends	firing number (frozen)	price	mechanism
R-18 / K-P the realised sd is too large for $N = 8$	the falsifiability of every $N = 8$ kill	realised paired sd at or above $2.18\times$ the pilot (0.109199, giving $\text{MDE}_8 = 0.126238$)	≈ 34 s	M-3, M-9
R-19 / K-E1 a column sink reproduces the row condition	the boundary-row mechanism as a capability	ChaCAL-with-sink within MDE_8 on at least 6 of 8 and residual within $2\times$, holding at $N = 16$	≈ 98 s, escalation ≈ 44 s	V-24, M-13, M-7
R-20 the one-step rule is not a policy	the phrase “safest move”	one-step and two-step optima disagreeing on more than half of 512 oracle draws	0 GPU-s	P-7, V-17
R-21 / K-J the dial is pinned	the horizon-dial sentence, Bet C, and S-64’s runnability	PINNED on at least 6 of 8 under the boundary null; or the ablation moving NRMSE by less than one seed sd	0 GPU-s	Ruling 10-prime, M-20, V-9
R-22 parity by underpowering	any “parity” or “equivalent” sentence	such a sentence resting on $N = 8$, where the interval half-width is 0.8807σ against a 0.5σ margin	119 to 228 s	M-13, M-9, V-5
R-23 the skyline is unmatched or chosen late	Bet D’s row	a depth chosen after the arena; or a “cannot” sentence outside the skyline table	≈ 7.6 s per cell	P-10, R-SKY, K-10
R-25 / K-G1 the re-solve is not needed	component (g)’s re-solve sentence	the cached mixture within one seed sd of the shape on P -changing plants on at least 6 of 8	1.680 s per cell	D-2, P-7
R-10 / K-I the certificate is uninformative or wrong	every “certified” sentence	one exceedance on 1,024 draws; or $\delta \ V\ _\infty$ at or above the label sd	≈ 2.6 s	L-CERT, V-10, V-17

Withdrawn candidate, recorded so it is not re-found (V-7 discipline). The conditioning of the committor solve on a self-loop-saturated transient block: with the diagonal logits raised by 0, 4, 8 and 12, $\max_T P_{ii}$ reads 0.6927, 0.9919, 0.99985, 0.999997 and the sup-norm condition number 7.9, 214, 771, 810, while the float32-against-float64 solve gap stays at 2.8×10^{-8} to 3.3×10^{-8} (RUN[M]). It does not bite at $s = 32$; it survives as one printed line beside the committor head and is not a card.

5.7 The cost table by phase, and the Kaggle gate

The basis. Two fitted throughput laws on the certified card (READ COSTS.md:73-74): softmax at R-squared 0.9999975 over five points, reading 0.010165 s per step at $n = 2048$; the corner-3 base at R-squared 0.9999987 over three points (the two largest refused because they paged over PCIe), reading 0.093084 s per step. The shape on the softmax corner is the softmax law **plus the solve increment** 1.041 ms per step (RUN[NEPTUNE]: solve plus product 2.514 ms against product alone 1.473 ms, forward and backward, float32, $n = 2048$), giving 0.011206 s per step. Rules the prices obey: each arm at its own law (M-8); the solve is an **increment**, never a cross-arm ratio; the increment was measured at the shipped $s = 64$ and never scaled from a smaller s (M-3); it is a per-op floor under a $2.0\times$ to $6.6\times$ dispatch gap, so every shape price carries direction “at least” until the cell-in-band check runs (P-8); no price crosses the device boundary (V-22); the laptop clock is not stationary at plus or minus 12 per cent, so prices are quoted in GPU-minutes to two decimals.

unit	price	class
one 150-step softmax control cell, $n = 2048$, $s = 64$	1.524 s	[FITTED]
one 150-step shape cell on the softmax corner	1.681 s	[FITTED + RUN]
fixed cost per invocation	4.0 s	READ
one bed-cell pair, 8 seeds each	29.6 s in one invocation; 34 s carried (two invocations, the conservative reading the record used)	DERIVED
the pair on the corner-3 base	129 to 133 s, no certificate on this base	DERIVED
InfSA-style Neumann $K = 16$ cell	5.12 s; 44.9 s at 8 seeds	DERIVED floor
depth-5 skyline cell	7.6 s; 64.8 s at 8 seeds	[ASSUMED linear in depth]
the capped cell	1.884 s	READ
1,024 forward passes at $n = 2048$	≤ 2.6 s	DERIVED
paired equivalence test at $N = 36$ / two-sample at $N = 70$	119 s / 228 s	DERIVED
the 9,600-step ladder, one pair	32 to 36 min	[FITTED] at 64 times
the s -sweep	206 to 537 GPU-s, band only, no point	READ
Lean targets, manifests, censuses, verdicts from journals	0 GPU-s	—

phase	cards	GPU cost	class
0	J-L0 to J-L18, J-D1 to J-D8, S-01 to S-03, M-0.3, N-03, R-15, R-17, S-61, V-0	0	—
1	S-10 to S-17, R-06, R-07, R-08, R-24, R-20	≤ 16 s (0.27 min)	DERIVED
2	M-2.0, S-20 to S-28, S-52, N-01, N-02, N-07, S-40, S-41, S-44, S-60	≈ 107 s (1.8 min); +524 s if the whole device certificate is rerun	DERIVED, N-02 [ASSUMED]
3	S-30 to S-35, S-62, S-51, S-50, M-3.7, M-3.8, V-17	≈ 8.7 to 9.2 min	DERIVED, S-35 [ASSUMED]
4	S-63, S-64, S-65, R-21, M-4.2, S-42, M-4.5	≈ 3.0 to 4.8 min	DERIVED
local decision total, phases 0 to 4		≈ 14 to 16 GPU-min (≈ 25 min with the full certificate rerun)	—
5	S-66, N-09, N-10, M-5.2, N-18 to N-21, M-5.3, N-11 to N-17, J-D5, S-72, J-D8	≈ 35 to 45 min for the priced rows; the kernels are NOT MEASURED	READ band plus [FITTED] at 64 times
local total, phases 0 to 5		≈ 0.8 to 1.0 GPU-h	—
6	N-22, N-23, N-24, S-73	0 locally; on Kaggle ≈ 8.7 GPU-min [ASSUMED] for the certificate, the LM cell NOT MEASURED	behind the yes

The whole falsification programme to a verdict costs under a quarter of a GPU-hour on the certified box. The record’s dominant failure — a bed dying at construction — is a 0 GPU-s decision at S-12, and the expensive death (K-E1) is a ≈ 34 s reading escalated once for ≈ 44 s.

What fits, and what does not. Bytes are modelled as

$$\text{bytes} \approx 1.256 \left[C_{\text{OP}} n s^2 \cdot 4 + 2 n s d \cdot 4 + 17.874 n s d \cdot 4 \right],$$

with 1.256 the worst reserved-over-allocated ratio, the residual constant 17.874 at R-squared 0.996373, and $C_{\text{OP}} \in [3.823, 5.823]$ — the signed arm’s fitted value plus one to two retained tensors, [ASSUMED] until N-18 measures it — against a budget of $0.90 \times 7.996 = 7.20$ GiB. At the carried upper constant the largest power-of-two n per s is 32768 at $s = 64$, 2048 at $s = 256$, 128 at $s = 1024$ and 8 at $s = 4096$; the dense operator at $n = 2048$, $s = 4096$ is 137 GB and runs nowhere. So n is declared per s on every long-context card, and the corner-3 base, at 8 bytes per element, halves every n and is not swept.

The Kaggle gate. Nothing in this plan launches. The gate’s local half (N-24) is: the device certificate lines filled by N-01, N-02 and N-18; the envelope’s two plants firing (N-22); the package’s smoke passing (N-23); the pinned bed digests re-asserting on regeneration (Ruling 7); the corpus pin with its split; the three unpinned sources whose hash cell prints and then **raises**; and the

open rulings listed by name beside the request. The gate’s own circularity — its first rows are blocked on numbers only the Kaggle run produces — is the author’s ruling and is not taken here. Its only consumer in this programme is S-73, the LM γ cell. **The rule, from the author’s own side and this plan’s standing memory: notify first, and no launch, upload or training start without an explicit yes.** A threshold carried across the device boundary is V-22 and is refused; a Kaggle number pooled with a local one is struck.

5.8 The Lean ledger

Grades: [M] a Mathlib route is confirmed at the pinned revision and the statement is expected to elaborate; [S] a supporting lemma must be supplied in-file; [D] deferred, no object exists in this Mathlib. **Every [M] is pending lake build; nothing is cited as proved before it builds** (L-LEAN, P-11). The eight Mathlib names the [M] rows rest on were found by `grep` in the vendored tree at rev `a45ae637` this session; a card introducing a further name is [S] until that name is checked.

#	target	grade	depends on	what it licenses	what dies if it fails
1	<code>gamma_zero_is_softmax</code> with <code>gamma_half_is_not_softmax</code>	[M]	J-L0	“CEQ contains softmax at $\gamma = 0$ by theorem” (the equation is ChaCAL’s; the machine check is the record’s)	the Lean rejection region of parity; Proposition 1 reads RUN-only (2.3003)
2	<code>bos_row_is_absorbing</code> at $\beta = 1$, with the $\beta = 0$ refusal	[M]	J-L0	census line X-2 as a theorem instance; J-L10 may take “BOS declared” as its only hypothesis	BED-S admission needs the determinant check on every batch instead
3	<code>later_boundary_unreachable</code>	[M]	J-L3	F2 as a theorem: a set after the query has committor exactly zero	census line X-4 stays a RUN with its plant

#	target	grade	depends on	what it licenses	what dies if it fails
4	<code>softmax_corner[M]</code>	<code>mbt_nilpotent</code>	J-L0	“all s hops by nilpotency” is a regime-N sentence only; the default keeps the diagonal	nothing measurable; the RUN pair (6.27×10^{-13} against 7.96×10^{-8}) stands
5	<code>lower_triangular[M]</code> <code>diag_one_sub_smul_pos</code>	<code>isUnit</code>	J-L0	Propositions 5(a) and 10; every card below that inverts a triangular matrix	the triangular branch stalls one evening; the substitution stands as RUN (1.8×10^{-15})
6	<code>resolvent_fromBlocks</code> <code>segmentation_blockdiag</code>		J-L0	the contract’s F0 line finally has a declaration; certificate Z-F0 at $\delta = 0$ is a theorem instance	the F0 row keeps its <code>torch.equal</code> RUN and no [M] tag
7	<code>cut_makes_segment</code> <code>cut_severs_boundary_sets</code>	<code>mt_head_absorb</code>	<code>diag</code> , J-L2	the dial rule that every segment carries its own sink, goal and constraint sets	segmentation and the committor channel are kept apart by rule, not by theorem
8	<code>displacement_identity</code> with <code>const_value_zero_displacement</code>		J-L0	Proposition 7(i) as LEAN, cited beside the linearisation and the rank-one form	nothing; three independent RUNs stand at 10^{-15}
9	<code>lowerTriangular[M]</code>	<code>the_symmSupport</code>	J-L0	the DAG substrate ruling as a theorem consequence: no causal operator equals an undirected chain	the substrate ruling stands on the RUN alone

#	target	grade	depends on	what it licenses	what dies if it fails
10	committor_is_resolvent_read (a), (b-causal)	[S]	J-L3, J-L2	Proposition 10: “absorption almost surely” is the single census line “BOS declared”	the head ships on a per-batch determinant check
11	subdiag_pow_ently, pathprod_is_chain_resolvent	[S]	J-L0	Proposition 2 in both halves; BED-M is contained by theorem and shape - corner-3 is VOID by theorem	the containment stays a 0.0 RUN, VOID by registration
12	neumann_truncation_bound, neumann_tail_attained, mask_amplification	[S]	J-L0	Proposition 4(i) and (ii); certificate rows Z-NEU and Z-F1	every δ is printed from the textbook with the RUN plant; no arena cell is lost
13	resolvent_is_trigonangular_solve mixing_matrix_rowStochastic	[S]	J-L3, J-L12	Proposition 5(a) to (c); the hull bound applies verbatim; the cost line prices a substitution that exists	the cost law stands on the RUN
14	causal_forwardonly, suffix_resolve_eq_full, sherman_morrison_row	[S]	J-L4, J-L8	Proposition 7(ii) and (iii); the safest-move price is a theorem about the oracle’s clamp, never carried to a token rewrite	the price is a DERIVED band with the RUN identities

#	target	grade	depends on	what it licenses	what dies if it fails
15	isUnit_one_sub[8]f_perron, reach_avoid_sum_one		J-L10	Proposition 8's conservation row and degeneracy lemma; the goal set and declared sink are forced by theorem	the row is printed with its residual per batch and the lemma is DERIVED
16	masked_sink_finite_sum		J-L12, J-L10	"BOS is a value-zero sink set" as a theorem, not a convention	the paper keeps the diagonal and Proposition 4 everywhere, already the default
17	hitting_time_transform, f1_cantelli_union		J-L0	nothing yet; the content is carried as J-D3 and J-D7	nothing; the statements live outside the root build

Rows 1 to 10 are one evening each of independent [M] work; rows 11 to 16 name their dependencies; row 17 is off the critical path. Price: 0 GPU-s each. The gate that consumes them all is J-L18, and its census column is what turns a green build into a licence: a theorem whose hypotheses no draw of a bed satisfies is **decoration on that bed** (L-DOM), which is why `segmentation_blockdiag` reads 3/3 on BED-M and 0/N on BED-S's softmax corner and the paper says so in both places.

5.9 The leap-dossier discipline

The author's supersession of 2026-09-03 licensed a new primitive; it did not lift the leap's other clauses. This is what a single leap call would receive, and the rule that governs what may be done with what comes back.

What the call receives. (i) The shape's definition and the eleven propositions with their status tokens and their `MISTAKES.md` mechanisms, as `judge/sec_shape.md` states them. (ii) The Lean ledger of §5.8 with each row's grade and each row's refusal. (iii) The apparatus: beds with the F1 and F2 conventions, labels and heads, floors, skylines and fellow approximators, control arms, metrics, census lines, binds with their planted negatives, certificates in vector units, the identity manifest field list, and the pricing rule. (iv) The predictions ledger of §5.5 with counters and `SPLIT` bands, and the calibration column (7 of 8 optimistic, $p = 0.0352$). (v) The standing attacks of §5.6 with frozen thresholds. (vi) The **NOT FOUND** list and the eight prior-art sweeps' queries, so absence is bounded by what was searched (V-7). (vii) The named fields the record's own leap ledger already owes — realisation theory, Koopman spectral theory, bifurcation and gradient-flow convergence, structured sparsity scheduling, finite-mixture inference,

anti-concentration, error-statistical severity, implicit bias — with the field-naming rule: naming a field and then saying what is wanted from it is compliant, naming the want and calling it a field is not.

The rule. A leap output is filed [LEAP-UNTESTED] with, per claim, its counter of equal specificity and its cheapest killer. **Its instances run before its text is believed**, and a leap output acted on before its instance runs is **struck** — that is the R15 kill, unchanged by the supersession (READ CEQ_V20_R15_CONTRACT.md:265–273). A second leap call in the same round is a breach. An absorbed component without a licensing theorem is dropped. The leap is scored in the same calibration column as the author, in its own sub-column, on instances-run and predictions-falsified — never on how interesting the text was.

What a dossier row looks like when it is complete (M-4.5): the candidate built with identity binds that each failed their plants at $O(1)$; the manifest with every declared field present; the certificates in vector units with $\|V\|_\infty$ and $1/(1 - \hat{\gamma})$ printed; the parameter count per arm; the verdict by the clauses registered before the run; the purge, with the manifest committed **before** any deletion; and the prognosis naming one arm, its dossier, its cost, its certificate, its floors and its one sentence. A bracket filled from prose rather than from a **results/** row voids the row (P-1, P-2).

5.10 What dies if what

Three outcome trees. Each names the sentence it licenses and the replacement route it owes. Under D-CALIB-1 the middle tree is the planning estimate.

Tree A — every prediction holds (the half the calibration column says is less likely). J-L18 reads 10/10; S-12 admits; S-26 strikes nothing; the arena reads the arm identified in $[10^{-3}, 10^{-2})$, $\hat{\gamma}$ MOVED on at least 6 of 8, the residual ratio at least 10, the cosine above softmax’s by more than MDE_8 , the argmin interval below the exact floor; S-31 and S-64 separate the shape from ChaCAL-with-sink; S-35 lands inside MDE_8 ; V-17 reads at or below $0.5\times$; S-65 declares equivalence at $N = 36$. *Sentence licensed:* “Consequence-Equilibrium Attention is ChaCAL’s causal resolvent read on the record’s Lean-checked three-corner base with a declared sink, a goal set and K constraint sets as boundary rows; on BED-S at $t^* = 8$, $m = 8$, $K = 2$, $N = 8$, parameters recounted per arm, it reads the reach-avoid vector with argmin error $[e]$ (interval $[l, u]$) below the exact restricted-view floor $[f]$, harmonic residual ratio $[r] \geq 10$ at equal marginal error, and displacement cosine $[c]$ above depth-1 softmax’s by $[\Delta] > MDE_8$; ChaCAL at the same $\hat{\gamma} = [g]$ ($\Lambda = [\Lambda]$, MOVED) with a sink token reads above the floor; the depth-5 skyline reaches the same argmin accuracy within MDE_8 , so the separate advantages are exactness ($\delta = 0$ on the exact route), one-read cost (one solve, depth s , plus 50 per cent causal MACs) and the boundary-row mechanism.” Every bracket is a number BED-S has not produced. *Route owed:* none; the ladder (M-5.3) is spent and the price is re-quoted upward.

Tree B — the median outcome (every counter is the point estimate). J-L18 reads 9/10, one row demoted to [S] on an unconfirmed Mathlib name. S-12 admits after one repair of the sink share — the fix’s own risk, since the declared sink read $[0.362, 1.000]$ on the one draw run. S-26 does not strike. The arena reads the arm identified, $\hat{\gamma}$ PINNED on at least 6 of 8, the residual ratio SPLIT in $(2, 10)$, the cosine within MDE_8 , the argmin interval straddling the floor. S-31 and S-64 are then void, because ChaCAL at $\hat{\gamma} = 0$ is softmax. S-35 lands inside MDE_8 ; V-17 is SPLIT; S-65 needs $N = 70$. *Sentence licensed:* “Consequence-Equilibrium Attention is a re-parameterisation of ChaCAL’s causal resolvent read with absorbing boundary rows on the record’s three-corner base. Nine of its identities are machine-checked and the tenth is [S] with its numeric instance; the committor head is the exact solve at $\gamma = 1$, matching the bed’s own Dirichlet solve to 0.0 on its real sets; BED-S is registered with an exact oracle at 0.0, the exact

zero-information floor 0.875 at $m = 8$, and a printed paired sd [sd]. At $N = 8$ the trained $\hat{\gamma}$ is not distinguishable from zero under the boundary null, depth-1 softmax matches the marginal error and the displacement cosine within MDE_8 , and the residual ratio is SPLIT. No capability sentence is licensed; the contributions are the identities, the certificate discipline, the bed with its floors, the controls and the negatives.” *Routes owed*: **retire** the horizon-dial sentence to a definition (R-21) and the LM γ claim (S-73) before it is made; **reprice** “one read” as a cost statement (Bet E) and the argmin to $N = 16$ (about 68 s) before any argmin sentence; **retire** “consequence as a field” to a label class (Bet A); **keep** regime N at $\gamma = 1$ (corner 3, exact without a dial) and the constant-value plant. This is the sentence the author reads first.

Tree C — the shape dies at its weakest component, before a GPU-second. BED-S does not admit at the design point: the sink share dominates, or the admitted-region sd of the goal committor reads at or below 0.05, or an argmin class frequency leaves (0.05, 0.95), or R-07’s Chebyshev ties empty the bed, or R-24’s flag rule reads above the floor at construction. Everything from S-14 to S-65 is VOID and unscored. *Sentence licensed*: “BED-S at the design point does not admit; the paper files the generator specification, the census with its printed failure line, the floors, the controls and the pricing rule, and no capability number.” *Routes owed, in order*: **reroute the dial** — move t^* to 2 or raise the graph depth, since the DAG substrate bounds t^* by the causal window, and re-run the 0 GPU-s census; **reroute the substrate** — the jittered `bed_1` landscape at $K = 2$ with B as goal, as an **oracle cross-check only** (it is undirected with $\rho(Q) = 0.9409$, and no causal operator equals it, so it is never the shape lane’s substrate); **reprice** N from whatever realised sd the first admitted batch shows — at the $2.18\times$ precedent, $\text{MDE}_8 = 0.126238$ and nothing in Phase 2 is falsifiable at $N = 8$, so $N = 16$ is the first honest count.

The smaller deaths, one line each, with what survives.

kill fires at	what dies	what survives, and at what price
J-L18 red (K-K)	the “machine-checked” sentence for the rows that failed; training does not start	every float instance as RUN; the arena waits one evening
S-12 census (0 GPU-s)	BED-S at $t^* = 8$	the reroute above, repriced from its own sd
S-26 / R-06 (K-H2, ≈ 0.5 GPU-min)	BED-S as a capability bed	the label class registered with printed floors; no number
R-07 (0 GPU-s)	the Chebyshev column as written	the goal rule and the conditioned column, with the restriction in the spec
R-24 (0 GPU-s)	BED-S’s argmin head	the committor-vector head alone
S-20 (about 1.7 s)	the determinism sentence, or every [FITTED + RUN] price	the corner-3 base at 133 s per pair, with no certificate, and the paper says why
S-62 K-D2 (≈ 34 s)	every capability number on BED-S	the identification reading; the property paper
R-21 PINNED (0 GPU-s)	the equilibrium sentence; S-64 and S-65 become void	γ as a bed-side dial; regime N at $\gamma = 1$, exact without a dial
S-64 K-E1 (≈ 44 s)	component (e) as a mechanism	ChaCAL plus a certificate plus Lean containment plus a registered bed — tree B’s sentence

kill fires at	what dies	what survives, and at what price
S-33 K-G1 (≈ 35 s)	component (g)’s re-solve sentence	the displacement identity as algebra; the channel as a cost statement
Bet E at or below 2	“joint determination in one read”	the residual as a diagnostic; one solve, depth s , as a cost statement
S-41 / N-15 K-I (≈ 2.6 s)	the mask	the exact solve, cheaper than one hop at $s = 64$
S-66 K-9 (band)	every s -scaled price and the “plus 50 per cent” wall-clock sentence	the MAC law as a MAC law; measured points quoted as points
N-24 (the yes withheld)	nothing local	the whole local programme; the Kaggle rows stay NOT MEASURED

5.11 Deliverables, the return report, and how to keep this plan current

Deliverables of the paper. docs/CEQ_SHAPE.md, docs/CEQ_SHAPE.tex, docs/CEQ_SHAPE.pdf, docs/references.bib, and docs/PLAN.md — this section, standalone.

Deliverables of the programme, by phase. Phase 0: lean/CEQ/Shape/*.lean with docs/LEAN_SHAPE_MANIFEST.md and docs/LEAN_SHAPE_CENSUS.md; docs/apparatus/MANIFEST_SHAPE_LANE.md, IDENTITY_SCRIPT_RULES.md, OBSTRUCTION_CENSUS.md, INSTRUMENT_GAMMA_LR.md; docs/PRICE_LEDGER.md; docs/CEQ_KERNEL_PATH_BATTERY.md; docs/CALIBRATION.md. Phase 1: docs/beds/BED_S_PLACEMENT.md, BED_S_SPEC.md, BED_S_CENSUS.md, BED_S_VOID_LIST.md; docs/apparatus/METRICS_AND_FLOORS.md; results/bed_s_census.jsonl, bed_s_zerohop.jsonl, bed_s_leak.jsonl, bed_s_dz_census.jsonl. Phase 2: results/binds/*.jsonl with one FOUND cell per plant; docs/apparatus/PLANT_REGISTER.md, CERT_EXACT_ROUTE.md, CERT_F0_SEGMENTATION.md; results/certs/z_neu.jsonl; results/bedm/capped_seeds_2_3_7.jsonl; results/k_cert_local.json with its two new blocks; docs/apparatus/PREDICTIONS_R16.md and results/arena/PREDICTIONS.sha256. Phase 3: results/arena/*.jsonl per arm, results/bedk/first_cell.jsonl, results/bedm/{containment,vector_lane}.jsonl. Phase 4: results/arena/VERDICTS.jsonl, docs/apparatus/ADJUDICATOR_SPEC.md, docs/plan/GAMMA_VERDICT.md, docs/CEQ_SHAPE_DOSSIER.md. Phase 5: docs/COST_LAW.md, results/cost/*.jsonl, docs/CEQ_{CHUNKED_SOLVE,CSR_PATH,MAPPER_SCHEDULE}.md, docs/apparatus/DIGRAPH_BETA0_SPEC.md, PERRON_WEIGHT.md. Phase 6: docs/CEQ_ENVELOPE.md, docs/KAGGLE_GATE.md, MODEL_CARD.md.

The report the author finds on return. One file, docs/PLAN_STATUS.md, regenerated at the end of every working session and never edited in place, with exactly these blocks and no prose beyond them:

1. **Where the DAG stands** — every card id with a token in {not started, in progress, PASS, KILL, VOID, NOT MEASURED} and, for a PASS or KILL, the deciding number with its evidence class.
2. **What was decided since the last report** — one line per gate crossed, naming the gate, the numbers that decided it and the phase it opened.

3. **What died, and what it took with it** — the row of §5.10’s table that fired, quoted, with the sentence now withdrawn and the replacement route owed.
4. **The GPU-minutes spent** against the phase’s budget in §5.7, with the class of each price.
5. **The calibration column** — rows scored, rows wrong, signs, and the running one-sided sign test on the pooled column and on the author and leap sub-columns separately.
6. **The next five cards**, cheapest-decisive first, each with its price and its evenings.
7. **Open questions for the author only** — the Kaggle yes, and any ruling the plan cannot take alone. Nothing else in the report asks a question; every other line is a decided number.

How to keep this plan current. An **append-only corrections index**, docs/PLAN_CORRECTIONS.md, one row per correction: id, date, what was wrong, the number or line that says so, the card or section corrected, supersedes. The plan’s body is never silently edited: a corrected card gets a new row and a **supersedes** pointer, exactly as journals do (L-G2), and the round-15 corrections index C1 to C40 is the pattern. The mechanism is **P-3** (a stale claim never retracted): the record’s own instance is a prompt calling a direction “the one live direction that has not yet been measured here” while a complete 305-line instrument for it sat in the tree with two of six required cells already journalled, and a document saying a run “was still running when this was written” that finished and recorded a worse number and was never updated (READ MISTAKES.md:316-328). The rule the correction index enforces is that one: **a claim of the form “not yet measured” or “still running” carries the date it was written and is checked against results/ before it is read.** Three standing consequences for this plan: every NOT MEASURED tag carries the instrument it needs and is re-checked at each report; every [ASSUMED] price carries the card that would replace it; and every [U] citation carries what would make it [V]. The debts open at the time of writing, so that they are not rediscovered as findings: the boundary-null critical value 2.7055 has no references.bib citation and is **owed**; ChaCAL’s diagonal convention rests on one HTML fetch and must be re-read against the PDF before S-31 is typeset; the adjacency-feature representability of the environment chain inside the arm’s class is unmeasured; the Kirchhoff second route covers $K = 2$ single-node sets only; the hop-to-committor reduction is NOT FOUND; and BED-S has no cell, no realised sd and no measured t^* , so every BED-S number in this plan is a floor formula or a design constant.

6. The reference apparatus — beds, floors, skylines, controls, metrics, census, binds, certificates, manifest, prices

Judged by SATURN (WATSON), 2026-09-03, against HEAD 207e7b9. Inputs: BRIEF.md, THESIS_NOTES.md, THESIS_CORRECTIONS.md, THESIS_CORRECTIONS_2.md, the six sections/sec_.md, the eight sweep/sweep_*.md, references.bib, bib_aliases.md, the three designs and all six refutations (every refuter ran; none is missing). Evidence classes: RUN (executed this session, numpy float64, seed 0, s = 32, CPU; the two one-liners are reproduced in §A.13), RUN[x] (a named planet’s run, carried with its own class), READ path:line, CITED [V]/[U] by canonical references.bib key, DERIVED (steps shown). The judge’s rule is the coordinator’s: KEEP is adopted as written, REPAIR is applied in the text below, a KILL whose reversal the refuter named and the judge can apply in one sentence is applied and marked, and a KILL whose reversal needs an instrument or a bed that does not exist is deleted from the apparatus and listed once in §A.12. This file is the plan’s reference apparatus — every table is a specification for a developer working alone on the certified RTX 4060, never a result. Every row names the MISTAKES.md mechanism it is built against (READ MISTAKES.md, 66 ### headings including V-14a, RUN[MERCURY] grep).*

A.0 What the six refutations decided, in one table

object	design	refuter verdicts	judge's disposition
BED-S substrate (undirected Rips / bed_1 chain under a causal arm)	instrument E1/H1/H2, falsify Bet B, theory P8	KILL $\times 3$ (refute_instrument_occvac FATAL-1; refute_theory_occvac §2.4 both horns; refute_falsify_math §3.13)	applied: the environment chain is a random DAG in token order (§A.1); the VOID-contrast list is written (§A.1); bed_1 / E4' demoted to oracle cross-checks
BOS undeclared; constraints after the query	falsify (predates F1/F2)	KILL (refute_falsify_occvac §5.1; refute_falsify_math §0.3)	applied: F1/F2 adopted; BOS in its own value-zero sink set $\mathcal{A}_{\text{sink}} = \{0\}$ (§A.1), which also removes the goal-collapse of refute_falsify_math §3.11
committor label at γ = 1 read by an arm at $\gamma < 1$	instrument E2/H1, falsify Bets B+C	KILL $\times 2$ (refute_instrument_occvac FATAL-2; refute_falsify_math §3.12)	applied: the committor head is the Dirichlet solve at $\gamma =$ 1 on the transient block; γ lives on the $\mathbf{z}/\Delta\mathbf{z}$ channel only; discounted reads are printed as columns (§A.2)
Fano 1 — ln 2/ln m as the zero-hop floor	beds §6.C.5, all three designs	KILL/REPAIR $\times 3$ (refute_theory_math row 16; refute_instrument_math row 18; refute_falsify_math row 14)	applied: exact zero-information floor $1 - \max_a \pi(a^*)$; thought Fano only where $I > 0$ (§A.3; RUN)
F1 mask certificate “dropped mass + Neumann δ ”	instrument I/M, falsify §3.4/§3.8	KILL (refute_instrument_math row 8, $s = 3$ counterexample)	applied: the solvent amplifies dropped mass by $1/(1$ $- \gamma)$; RUN reproduces 0.5263 vs naive 0.1 (§A.9)
ChaCAL as “identical arithmetic” control	falsify §1, instrument E1	REPAIR/KILL (refute_falsify_math §0.4 [V-fetched]: diagonal removed inside the inverse)	applied: two ChaCAL arms, ChaCAL-published and ChaCAL-diag (§A.4); ChaCAL-published's $\mathbf{A}_s$ is a planted negative for the bitwise half

object	design	refuter verdicts	judge's disposition
EMC as a proposition with a feedback plant	theory C6(iii), falsify §3.10, instrument G1	KILL $\times 3$ (<code>refute_instrument_math</code> row 11; <code>refute_falsify_math</code> row 7; <code>refute_falsify_occva</code> §3.13)	deleted as a proposition; kept as the <i>triangularity</i> bind ($\Delta z[:i] = 0$) with a non-causal P plant, and one remark citing [35]
Sherman–Morrison 0(sd) per candidate	theory P8/P9, falsify §3.2	KILL for token <code>do()</code> (<code>refute_theory_math</code> items 5/19; <code>refute_falsify_math</code> row 10)	applied : moves are query-side row clamps on the environment chain (as <code>sec_beds.md</code> §6.C.1 specifies); key-side moves are priced as m suffix re-solves (§A.11)
Row L, influence-Jacobian β_0 barcode on BED-S	instrument L, falsify §3.7	KILL $\times 3$ ($\beta_0(\varepsilon = 0) = 1$ on 100 % of softmax draws; instrument consumes point clouds)	deleted (§A.12); reversal needs a digraph β_0 instrument that does not exist
regime-N survival route for component (e)	falsify §5 rank 2	KILL (<code>refute_falsify_math</code> row 5: an absorbing row breaks <code>StrictlyLower</code>)	deleted
“m(K+1) right-hand sides when moves change only values”	theory P8	KILL (<code>refute_theory_math</code> item 17: $q^{\{k\}}$ has no V)	applied : restricted to the Δz channel
support sentence “identity rows change the support of Π_γ ”	falsify §3.1, theory P5	KILL (<code>refute_falsify_math</code> row 13; <code>refute_instrument_math</code> row 5, 23)	applied : boundary rows change <i>weights</i> on downstream rows; support is the causal triangle either way
everything else in the three designs	—	KEEP or REPAIR	applied row by row in §A.1–A.11; the ledger (<code>bind_ledger.md</code>) carries each row's status

A.1 Beds

Column conventions. *Label shape* is what the arm emits. *Oracle* is what computes the label and **where it runs** — always on the latent environment chain the bed builds, never on the arm's own P (READ MISTAKES.md:710–725, D-2). *Placement* is the F1/F2 convention of THESIS_CORRECTIONS_2.md §1. *Role* says what a reading on the bed may say about the shape.

bed	generator	label shape	oracle and where it runs	BOS and constraint placement	role in the plan; mechanism
BED-M chain (E1/E3)	<code>make_equilibrium</code> (<code>READ</code> <code>scale/negation</code> via <code>sec_beds.md</code> §6.A)	<code>sum_batch</code> <code>z*_{s-1}</code> , <code>Nscope.py:354-443</code>	<code>equilibrium_oracle</code> (<code>:286-304</code>), 443 forward scan on the latent chain; equals the last row of $(I - A)^{-1} b$ to <code>6.2e-15</code> (<code>RUN[coord]</code>)	boundary sets; <code>b[s-1]</code> = 0 (zero-hop guard)	contained, never won: the parity bind (I1) and the truncation- law must-fire only; shape – corner-3 is VOID by registration; D-1, D-2 the scalar containment case of the Δz channel; its two guards (mirror contrast, raw-label kill) are inherited by BED-S; V-24
BED-M consequence (E2)	<code>make_consequence</code> (<code>:803-850</code>)	<code>scale_batch</code> contrast at player 5	200 damped best-response sweeps with player 0 clamped (<code>:718-732</code>)	—	the planted negative for the boundary mechanism (the negation gate the signed operator lost); never a capability bed; V-24
BED-M executed programs	<code>ceq/corpus.py</code> (<code>:88-130</code>)	<code>scale_batch</code> contrast at player 5	CPython <code>exec</code> (<code>:60-65</code>), cross-checked by <code>execute_subprocess</code> (<code>:68-72</code>)	—	the planted negative for the boundary mechanism (the negation gate the signed operator lost); never a capability bed; V-24

bed	generator	label shape	oracle and where it runs	BOS and constraint placement	role in the plan; mechanism
E4' harmonic on Rips	<code>scale/e4_harmonic.py</code> (:103–176)	<code>undcappr</code> node	<code>np.linalg.solve</code> – <code>Q, r</code>) and the Kirchhoff forest ratio, agreement 9.6e-14	<code>valEdge</code> endpoints absorbing	oracle instrument and the bottleneck warning (<code>t_rel</code> ≥ <code>vol(S)/2</code> = 1372.50 on <i>its</i> 1,200-node case, not a BED-S constant — V-22); its <code>PASS_BAR/FAIL_BAR</code> are E4's and are re-derived before BED-S uses them
BED-S (new; the repaired specification)	see below	[<code>m</code> , <code>K+2</code>] reach-avoid tensor (goal, <code>K</code> constraints, sink), the <code>argmin</code> , <code>z* ∈ ℝ^{s×d}</code> , <code>Δz(a)</code> per move	the Dirichlet solve of <code>bed_1.commit</code> form on the latent DAG <code>chain P_env</code> , at $\gamma = 1$ on the transient block; <code>Δz</code> by the C6 identity on <code>P_env</code> ; Kirchhoff second route for <code>K = 2</code> single-node boundaries (<code>READ</code> <code>scale/kirchhoff.py:1–90</code>), multi-node extension NOT MEASURED – needs the grounded Laplacian extended (P-4)	<code>A_sink =</code> <code>{0}</code> (BOS, value 0), <code>A_0</code> (goal), <code>A_1..A_K</code> all before the query position ; the query in <code>T</code>	the capability bed for (e), (g), (h); admitted only when every census line of §A.7 passes (D-4)

BED-S specification (the one sentence the refuters asked for, `refute_instrument_occvac` §2.1 and `refute_theory_occvac` §2.4). The environment

is a random directed graph on the $\mathbf{s}$ token positions whose edges point to *earlier* positions (a DAG in token order, so that a walk from the query descends, exactly as the arm’s causal $\mathbf{P}$ does — RUN: with every declared set absorbing the transient block $\mathbf{Q}$ is lower-triangular, $\rho(\mathbf{Q}) = \max_{\{i \in T\}} P_{ii} = 0.692660$), with self-loops only on the declared absorbing positions. Node tokens carry the node’s out-adjacency as a multi-hot feature (so that a depth-1 bilinear logit can realise $\mathbf{P} = \mathbf{P}_{\text{env}}$ at $d_{\text{model}} \geq n_{\text{nodes}}$ — a representability claim that is NOT MEASURED), the membership flags of the absorbing sets, and nothing derived from any solve. Candidate moves are **m query-side row clamps** $\mathbf{P}_{\text{env}}[v_a, :] \leftarrow \mathbf{e}_{\{u_a\}}$ (the graph-surgery precedent of READ `scale/negation_scope.py:718-732`), each a rank-one change of $\mathbf{P}_{\text{env}}$ (`refute_theory_math §1.6`: a *token* rewrite would change every row $\geq i$ and is not rank-one; the bed therefore defines `do(a)` as the clamp). The transition matrix $\mathbf{P}_{\text{env}}$ is a deterministic function of the edge tokens by construction; `sec_refuted.md` C2’s leak clause (a) is therefore **dropped** for BED-S (the graph is the input), and clause (c) — $R^2 \geq 0.5$ on q at order 0 from a corpus-alone probe — is kept as the leak kill. Design against V-25 (the mechanism’s domain is non-empty), D-2 (the oracle’s chain and the arm’s class now overlap, so the D-2 kill $\|\mathbf{P} - \mathbf{P}_{\text{env}}\|_{\infty} < 1e-3$ on $\geq 6/8$ seeds has a rejection region and is registered as a *learnability* reading, not a capability), V-12 ($K \geq 2$ and a goal set), V-8 (sink share printed).

The VOID list for BED-S, registered before any cell (D-2, M-7). Because the arm’s class contains the oracle’s operator, `shape - softmax` and `shape - skyline` on the committor head are reproduction-versus-non-reproduction contrasts and are VOID as capability numbers. Creditable contrasts: `shape - ChaCAL-diag` and `shape - ChaCAL-published` (the boundary-row mechanism), `shape - ChaCAL-with-sink-token` (the C2 kill), `shape - InfSA-style Neumann-K` (exactness), `shape - 0-hop MLP` and `shape - 1-hop softmax` on the argmin (the C8 static-task kill), and the identification reading $\|\mathbf{P} - \mathbf{P}_{\text{env}}\|_{\infty}$ per seed.

Why BOS is a sink set and not a goal member. With $0 \in \mathcal{A}_0$ every un-absorbed walk ends at BOS and “reach the goal” collapses to “descend to 0 without hitting a constraint” (`refute_falsify_math §3.11`); the masked-diagonal route is exact only when the BOS value is 0 (`refute_theory_math` item 22). A separate value-zero sink set $\mathcal{A}_{\text{sink}} = \{0\}$ (the leak state of [99] [V], the anchor of [98] [V]) keeps $\rho(\mathbf{Q}) < 1$ (F1), keeps the goal a genuine target, and closes the conservation row: RUN with $\mathcal{A}_{\text{sink}} = \{0\}$, $\mathcal{A}_0 = \{5, 6\}$, $\mathcal{A}_1 = \{9, 10\}$, $\mathcal{A}_2 = \{15\}$ gives $q^{\{0\}} + q^{\{1\}} + q^{\{2\}} + q^{\{\text{sink}\}} = 1$ on T to $1e-15$ with sink share in $[0.362, 1.000]$ on this draw; BOS undeclared gives $\rho(\mathbf{Q}) = 1.000000$, $\det(\mathbf{I} - \mathbf{Q}) = 0$ (F1 reproduced). The sink share is a census line (§A.7) and a printed column, never hidden. The record’s value-zero BOS column sink (V15Fork.Asink, READ `lean/CEQ/V15Fork.lean:67-70`) and this row condition are two distinct facts that coexist on every causal softmax (`sweep_resolvent.md §2.14`; P-7).

A.2 Labels and heads

head	object	read at	oracle	notes; mechanism
committor / reach-avoid	$q(a) \in [0,1]^{\{K+2\}}$ per move	$\gamma = 1$ on the transient block $((I - Q)^{-1}$ $R \not\models$, a triangular solve on the arm's causal Q , invertible by <code>lower_triangular_isUnit</code> once BOS is declared, <code>refute_theory_math</code> §1.3)	Dirichlet solve on <code>P_env</code>	the arm's γ does not enter this head <code>(refute_instrument_occvac</code> FATAL-2 repair (a)); discounted reads $E[\gamma^{\tau-1}$ <code>1_k]</code> at $\gamma \in$ <code>{0.6, 0.9}</code> are printed as columns with the delay share $1 -$ Σ_k in read form (RUN: z- form/read-form ratio is exactly γ $= 0.600000$ on every reachable position; the z-form of <code>THESIS_CORRECTIONS_2.md</code> §1 is the state, not the read); V-17, D-2
argmin	$a^* = \operatorname{argmax}_a$ $q^{\{0\}}(\text{do } a)$	—	from the oracle tensor (an identity of the builder, V-3 declared)	the Chebyshev <code>argmin_a</code> <code>max_{\{k \geq 1\}}</code> $q^{\{k\}}$ ([119] [V]) and the lexicographic form ([132] [V]) are columns , with the disagreement fraction printed; the safety-filter reading is [57] [V], the threshold form [17] [V]; the nearest transformer NEAR-MISS is [60] [V]; V-1

head	object	read at	oracle	notes; mechanism
state	$\mathbf{z}^* \in \mathbb{R}^{s \times d}$	γ	$(\mathbf{I} - \gamma \mathbf{P}_{\text{env}})^{-1} \mathbf{V}$ on the environment chain at the arm's γ	<code>vector_readout</code> plumbing exists (<code>READ</code> <code>scale/m3_quintuple.py:368,</code> a separate lane and journal because a vector label voids <code>PUBLISHED_SOFTMAX_8192</code> (<code>READ</code> <code>MISTAKES.md:701-708</code>); M-1, D-1
consequence	$\Delta \mathbf{z}(\mathbf{a}) = \mathbf{z}(\text{do } \mathbf{a}) - \mathbf{z}$	γ	the C6 identity $\Delta \mathbf{z} = (\mathbf{I} - \gamma \mathbf{P}'_{\text{env}})^{-1} (\Delta \mathbf{V} + \gamma \Delta \mathbf{P} \mathbf{z})$ (<code>RUN[JUPITER]</code> 1.03e-15, <code>RUN[SATURN]</code> 1.2e-15, <code>RUN[VENUS]</code> 1.36e-15)	zero before the intervened row by triangularity (not nilpotency: <code>refute_falsify_math</code> §3.5); $\mathbf{V} \equiv \mathbb{K} \Rightarrow$ $\Delta \mathbf{z} \equiv 0$ is a must-pass identity ($\leq$ 1e-15, not “bitwise” — <code>refute_falsify_math</code> §3.7); the Gaussian- $\mathbf{V}$ region (0.1096 / 0.363 / 1.127 on three draws) is the bind; owners cited before the word “displacement”: [86], [16], [18] §7.3, [94] Eq. 5, [106] (all [V]); V-24, D-5
distributional	none ships	—	—	CRPS registered, never computed (P-1)

The read the paper prints is $\mathbf{0} = (1 - \gamma) \mathbf{P} (\mathbf{I} - \gamma \mathbf{P})^{-1} \mathbf{V}$ with the factor (ChaCAL’s Eq. 5, [42] [V]); the bare read has row sums $1/(1 - \gamma)$ **with or without** boundary rows (`RUN[MARS]` 2.5 both; V-23), so any table printing the bare form declares its row sums. Under the factor the Neumann certificate is $\delta_{\Pi} = \gamma^{K+1}$ (`refute_falsify_math` §3.1).

A.2.1 The dials (D-3: a dial that does not vary describes nothing)

dial	registered support	how it is set	what it must not be	mechanism
t^*	{2, 8, 32}	the smallest k with $\text{NRMSE}(z_k, z^*) < 1e-3$ on the hop ladder (E2's convention, E2_DIAL_TOL, READ scale/negation_scale — a tolerance dial read off the ladder, never predicted from a spectrum	the rate is $\rho(Q_{\text{env}})$, not λ_2 of an ergodic chain (measured 0.9964 vs 0.9985 on LargestJoin_S2Rips_1024, READ.py:264) PRIOR_ART.md:703-707); on the DAG substrate t^* is bounded by the causal window ($\leq i$ hops from position i , refute_theory_occvac §2.2) and is placed by the graph's depth, not by a kill rate (which collapses the label, sd 0.499989 $\rightarrow$ 0.038445, READ scale/e4_harmonic.py:425-434)	P-8, D-3
K	{2, 3, 4}	number of constraint sets; moves the number of right-hand sides and the conservation row	K = 1 (V-12: the committor is constant to $1.11e-14$)	V-12
m	{4, 8, 16}, default 8	number of candidate moves; moves the exact zero-information floor $1 - 1/m$	— ($m = 2$ admitted with floor 0.5, not default)	V-10
$ \mathcal{A}_0 , \mathcal{A}_k $	{1, 2} per set	set sizes; the Kirchhoff second route covers single-node sets at K = 2 only	—	P-4

dial	registered support	how it is set	what it must not be	mechanism
cut position c (segmentation rows only)	varies per draw	on a gated corner a zero gate at c makes position c a new undeclared absorbing state ($P_{cc} = 1, \rho(Q) = 1, \text{RUN}[\text{MARS}]$) and severs every boundary set before c from every $i \geq c$ ($q = 0.0$); BED-S therefore either never uses a zero gate on the committor channel or declares per-segment sets with the segment head in $\mathcal{A}_{\text{sink}}$ (<code>refute_theory_math</code> item 11)	a cut that is not censused per segment	D-3, V-25, V-8
γ (arm-side)	$[0, 1)$, learnable, init 0	the horizon dial on the $z/\Delta z$ channel; PINNED/MOVED by the LR test (§A.11)	a signed γ (forfeits the Neumann equality)	V-17, M-20
s	$\{64, 256, 1024, 4096\}$	the sequence length of the cost law; $s = 64$ on 40 of 40 banked cells left the exponent unidentified (<code>READ</code> <code>V20_R15_THEORY_TABLE.md:95-99</code>)	quoting any price before the sweep runs (C9)	D-3, P-8
β, qk, g	corners $\{0, 1\}$	Ruling 2a: β learnable per instance and logged	a sentence transferred across corners without a bind at the corner it describes	Ruling 2

A.3 Information floors (L-FLOOR)

Every capability number ships beside its information floor and never beside floor_1 (`READ` `V20_R15_JOURNAL.md:51, C15; 13 of 40` banked cells violate $\text{floor}_1, \text{RUN}[\text{WATSON}]).$

head	information floor	budget ceiling (printed at construction, per batch)	provenance; mechanism
committor vector, $\mathbf{z}^*$, $\Delta \mathbf{z}$	the exact oracle at 0.0	the hop ladder $\text{NRMSE}(\mathbf{z}_k, \mathbf{z}^*), \mathbf{z}_k$ $= \mathbb{Q} \mathbf{z}_{\{k-1\}} + \mathbf{R}$ (READ <code>scale/e4_harmonic.py:194-199</code>), numerical because the chain does not terminate (<code>OracleSeparation.truncation_never_exact</code>)	<code>sec_beds.md §6.C.5</code> ; M-2
$\Delta \mathbf{z}$	the exact oracle at 0.0; the no-change predictor's error printed beside it	—	[118] [V] owns the no-change floor; the paper claims only the pairing against a per-row control on byte-identical draws
argmin	exact zero-information floor $1 - \max_a$ $\pi(a^*)$ from the realised class distribution; at uniform π this is $1 -$ $1/m = 0.7500 /$ $0.8750 / 0.9375$ at $m = 4 / 8 / 16$ (RUN)	restricted-view Fano at $k \geq 1$: $P_e \geq$ $(\ln m - I(X_{\{\leq k\}};$ $a^*) - \ln 2) / \ln(m$ $- 1)$ (the tight form, [33] §2.10, theorem number [U]), used only where $I(X_{\{\leq k\}}; a^*) > 0$; I by plug-in over the full k -hop view including the move tokens and membership flags (<code>refute_instrument_ocfluar</code> §2.4)	the weak form $1 -$ $\ln 2 / \ln m = 0.5 /$ $0.6667 / 0.75$ (RUN) sits $0.25 / 0.208 /$ 0.1875 below chance and is retired as a floor; it is recorded only as the record's M11 arithmetic (0.1858 at $m = 8$, $I = 1$, 0.2229 at $m = 32$, I $= 2$ reproduce). $m =$ 2 is admitted (exact floor 0.5) but the default stays $m = 8$; V-10, V-17
BED-K	Hankel $1/d$: $\text{err}_1 =$ $0.9746794345 =$ $\sqrt{(1 - 1/20)}$ at $d =$ 20 (READ <code>V20_R15_LEAP_LEDGER.md:323</code> , DERIVED)	—	V-10 (the block is I_d with a zero row, arithmetic not evidence)
BED-M	the exact oracle at 0.0; $\text{floor}_1 =$ $\sqrt{((t^*-1)/t^*)}$ is a one-hop capability threshold , printed as such	$\sqrt{((t^*-k)/t^*)}$ at k $= 0, 1, 2$ (RUN [VENUS] $1.000058 /$ $0.712039 / 0.0$ at $t^* = 2$)	C15
certificates	none — a certificate is not a capability	—	L-CERT

A.4 The native skyline and the fellow approximators (R-SKY, D-1)

“Beats softmax” and “beats native” are banned sentences on every bed where a fellow approximator can solve it (READ CEQ_V16_CONTRACT.md:209); the shape’s separate claims are exactness (a printed δ), one-read joint consistency **as a cost statement** (one solve, depth s ; refute_falsify_occvac §1.4), and the boundary-row mechanism, stated as properties with intervals. [126] and [131] ([V]) show a softmax stack computes the same resolvent by iteration.

arm	class	depth / width / steps (fixed before the arena)	params	what it decides	price per 150-step cell ([FITTED + RUN] , sec_cost.md §4.x.8)
depth-1 softmax, same head	matched control	1 / d_model = 16 / —	4,769 (Ruling 3; the count is re-printed per arm once a [m, K+2] head is attached)	the single- location ground ([38] Prop 4.2 [V]; its SLR label model admits 0 of 3 record beds — refute_theory_occvac §1.2)	1.524 s
ChaCAL- diag	fellow approximator	1 / matched / —	matched	the shape with $\mathcal{A} = \emptyset$ (diagonal kept); the identity half of B-E1 is bitwise against this re- implementation only (V-3 declared)	1.680 s

arm	class	depth / width / steps (fixed before the arena)	params	what it decides	price per 150-step cell ([FITTED + RUN], sec_cost.md §4.x.8)
ChaCAL- published	fellow approximator	1 / matched / —	matched	$(1 - \gamma)A(I - \gamma A_s)^{-1}V$ with the diagonal removed inside the inverse ([42] [V-fetched], refute_falsify_math §0.4); regime N inside a regime-S outer read; sub- stochastic; its A_s is the planted negative for the bitwise half (must differ from the shape by $O(1)$ at every $\gamma > 0$)	1.680 s (DERIVED, same MACs)
ChaCAL- with-sink- token	the C2 kill's control	1 / matched / —	matched	whether a column sink reproduces the row condition; if it matches the committor read within MDE_8 on $\geq$ 6/8 seeds, component (e) is a parame- terisation	1.680 s

arm	class	depth / width / steps (fixed before the arena)	params	what it decides	price per 150-step cell ([FITTED + RUN], sec_cost.md §4.x.8)
InfSA-style Neumann, K = 16, no boundaries	fellow approximator	1 / matched / K = 16 hops	matched	exactness: the shape's $\delta = 0$ route against a truncated read ([99] [V])	≈ 5.1 s (DERIVED from the 25.409 ms hop microbench- mark, sec_cost.md §4.x.3)
$\lfloor \log_2 t^* \rfloor +$ 2 softmax stack	skyline (unmatched depth)	3 / 5 / 7 at $t^* = 2 / 8$ / 32 — chosen by analogy with [105] Thm 4.2 ([V]), whose hop_k task is not the committor (hop_k $\rightarrow$ committor reduction NOT FOUND, sweep_expressivity.md §3.5)	unmatched	Bet D (§A.8): reaches the shape's argmin accuracy within MDE_8	≈ 7.6 s at depth 5 (ASSUMED linear in depth)
wide constant- depth stack	skyline (unmatched width)	1 / width = n_nodes per [139] [V] / —	unmatched; at the record's geometry d_model = 16 < s = 64 so the matched- parameter instance does not exist and the table says so	depth is not necessary at linear width	NOT MEASURED
chain-of- thought decoder	skyline (unmatched decode)	1 / matched / t^* decode steps ([80] [V])	matched weights	the autore- gressive route to the hops	NOT MEASURED

arm	class	depth / width / steps (fixed before the arena)	params	what it decides	price per 150-step cell ([FITTED + RUN], sec_cost.md §4.x.8)
chunked-WY (DeltaNet / SSD)	native skyline for a resolvent- shaped read on the linear corner	—	—	[136], [34] [V]; closed to the softmax corner by [58] [V] (rank T), so the shape pays the intra-chunk primitive on the whole sequence; Δ_{sky} column, never “beats”	NOT MEASURED - needs a chunk-size sweep
MuZero-style value head	skyline for the consequence channel	—	—	a planner that never predicts the state can win on return ([107] [V]); the shape claims fidelity with a certificate, not return the C8	NOT MEASURED
0-hop per-position MLP; 1-hop softmax; majority move; random (1/m); predict- the-mean	argmin controls	—	matched	static-task kill (McNemar on identical draws)	≤ 1.524 s each

arm	class	depth / width / steps (fixed before the arena)	params	what it decides	price per 150-step cell ([FITTED + RUN], sec_cost.md \$4.x.8)
cached- mixture arm O_cached = P_base (I — $\gamma P_{\text{base}}^{-1}$ V_int	the interventional- channel control	1 / matched / —	matched	P frozen from the un-intervened context, values from the intervened one — the operational form of a cached successor rep- resentation ([85], [101], mechanism [U]); must fail on ΔP plants and pass on ΔV plants [19], [90] [V]:	1.680 s
return- conditioned DT-style arm	argmin control on deterministic beds	—	matched	a win on a deterministic bed is not about causal semantics	NOT MEASURED

A.5 Metrics (C11 adopted, with the refuters' repairs)

head	metric	declared units and nulls	mechanism
$\mathbf{z}^*$, $\Delta\mathbf{z}$	position-matched per-coordinate NRMSE vector (mean, max); the harmonic residual $\mathbf{r}(\hat{\mathbf{z}}) = \ (\mathbf{I} - \gamma\mathbf{P}_{\text{env}})\hat{\mathbf{z}} - \mathbf{V}\ _{\infty} / \ \mathbf{V}\ _{\infty}$ as the joint-consistency score (never a loss, D-2)	$\mathbf{r}$ is printed beside $\sigma_{\min}(\mathbf{I} - \gamma\mathbf{P}_{\text{env}})$ and $\ \mathbf{I} - \gamma\mathbf{P}_{\text{env}}\ _{\infty} \leq 1 + \gamma$, so that a $\geq 10 / \leq 2$ ratio between arms is in units (refute_falsify_occvac §3.12: a $10\times$ ratio at equal marginal error is a statement about which mode the control mislearns)	V-26, V-17
$\Delta\mathbf{z}$	additionally field cosine $\langle \Delta\hat{\mathbf{z}}, \Delta\mathbf{z} \rangle / (\ \Delta\hat{\mathbf{z}}\ \ \Delta\mathbf{z}\)$ and magnitude ratio; coordinates with label variance 0 (every coordinate $< \mathbf{i}_{\min}$, and coordinate 0 on every draw since row 0 is $\mathbf{e}_0$) are masked from the cosine and their fraction printed; the intervened position is drawn uniformly	an arm emitting $\Delta\hat{\mathbf{z}} \equiv 0$ makes the cosine 0/0 and is refused , not scored (V-16)	V-8, V-16
$\Delta\mathbf{z}$ sign column	sign fidelity with Clopper–Pearson ([30] [V]); McNemar ([78] [V]) on the sign column only , paired on byte-identical draws	the cosine is continuous: paired $\mathbf{t} /$ Wilcoxon with its own realised sd, never McNemar and never the NRMSE MDE (refute_falsify_math §3.14)	V-17

head	metric	declared units and nulls	mechanism
committor vector	Fisher–Rao coordinate $\varphi(\mathbf{p}) = 2 \arcsin$ $\sqrt{p}$ per entry, declared here; position-matched NRMSE on φ ; the harmonic residual with the environment’s <code>Q_env</code> , <code>R_env</code> — which requires the label to be the full vector on T , not only the <code>[m,</code> <code>K+2]</code> summary (<code>refute_instrument_occvac</code> row 14)	the record measured Euclidean/Fisher– Rao ratios 1.418962 and 1.209472 (<code>READ</code> <code>MATHEMATICS.md:69–104</code>), so choosing the coordinate after the curve would be M-2	M-2, V-26
argmin	accuracy with Clopper–Pearson; McNemar paired against each control on identical draws	the exact zero-information floor of §A.3 printed beside it	L-FLOOR
any distributional head	CRPS, registered	never computed until a head exists	P-1
refused	pooled W1 (permutation-blind: 0.0 on the permuted oracle against NRMSE 1.421901, <code>READ</code> <code>V20_R15_THEORY_TABLE.md:221</code>), Procrustes, OT without a position cost	the three metric plants of §A.8 (B-G2)	V-26, L-14

A.6 Control arms per bed

bed	mandatory controls	skylines	mechanism
BED-S committor / argmin	depth-1 softmax with the same head; ChaCAL-diag; ChaCAL-published; ChaCAL-with-sink-token; InfSA-style Neumann $K = 16$; 0-hop MLP; 1-hop softmax; majority; random; predict-the-mean	the three skylines with width and steps fixed (§A.4)	D-1, R-SKY, C8
BED-S Δz	depth-1 softmax with the vector readout; cached-mixture arm; no-change predictor	the three skylines; MuZero-style value head	D-2, D-5
BED-S z^*	depth-1 softmax with the vector readout; ChaCAL-diag	the three skylines	D-1
BED-M	softmax; <code>arm_pl</code> ; <code>arm_smprime</code> (contained)	the scan skyline (legal as the native control since R11, <code>READ workdonenewseal.md:514</code>)	D-2
BED-K	<code>hard_delay_attention</code> — (READ <code>ceq/beds/bed_k.py:311-336</code>)	—	V-3
BED-S long-context ($s \in \{256, 1024, 4096\}$)	do-nothing 0D-salience schedule ([110] [V]); [143] blockwise resolvent ([V]); k-means routing ([100]); SBM-sampled overlapping masks ([28]); dense resolvent at the same s with n declared per s (at $n = 2048$, $s = 4096$ the explicit <code>[n, s, s]</code> operator is ≈ 137 GB in float32, DERIVED $2048 \cdot 4096^2 \cdot 4$ B, against 7.996 GiB)	—	V-9, P-8

A.7 Domain census (L-DOM) — printed per draw batch, before any arm is trained; 0 % on any line blocks the bed (D-4)

1. Every drawn `P_env` row is stochastic including the absorbing rows, $|\text{rowsum} - 1| \leq 1e-12$ (the hypothesis of every row-stochastic theorem; V-25).

2. $0 \in \mathcal{A}_{\text{sink}}$ on 100 % of draws (F1; RUN: BOS undeclared gives $\rho(Q) = 1.000000$, $\det(I - Q) = 0$).
3. $\rho(Q) = \max_{i \in T} P_{ii} < 1$ on 100 % — a diagonal read on the arm’s triangular Q (RUN 0.692660 both ways), implied by line 2 for finite logits and therefore a V-10 gate; printed, not counted as evidence.
4. Every constraint set, the goal set and the sink lie before the query position; the query is in T (F2; a set after the query reads $q = 0.0$ exactly, RUN[SATURN] R10, RUN[VENUS] RUN-2; a query inside a boundary set reads its own indicator — V-8).
5. Per-coordinate label sd of q **over the admitted query region** (the prefix before the first constraint reads $q^{\{(0)\}} \approx 1$ and would inflate the sd , `refute_instrument_occvac` §2.7), each > 0 ; the sink share $q^{\{\text{sink}\}}$ printed separately from $q^{\{(0)\}}$ (RUN: sink share [0.362, 1.000] on one draw — the balance guard is at risk from the fix itself and must be watched).
6. Argmin class balance: **every class frequency of a^* over m lies in (0.05, 0.95)** (the instrument’s “argmin-unique fraction outside (0.05, 0.95)” was an inverted gate — `refute_instrument_occvac` §2.4); ties within $1e-9$ discarded and the discard count printed.
7. Reduction-disagreement fraction between $\argmax q^{\{(0)\}}$, Chebyshev and lexicographic > 0 , else only one rule is named (C4).
8. Corpus-alone linear probe from x to q at order 0 reads $R^2 < 0.5$ (leak kill (c)); the planted leak (a move token carrying its own q) must read $R^2 \geq 0.99$ — the detector fires both ways (V-24, V-7).
9. $I(s_0; a^*) = 0$ up to plug-in error; the untrained arm reads argmin accuracy within $1/m \pm CP$ and committor $NRMSE \geq 1 - GATE_TOL$ ($GATE_TOL = 1e-3$, `READ COSTS.md:137`).
10. $Var(\Delta z) > 0$ per coordinate **over coordinates $\geq i_{\min}$** ; the zero-coordinate fraction printed (V-8).
11. Fraction of moves that change P_{env} (needs a re-solve) versus only V ; fraction with $u^T \mathcal{K} = 0$ and a non-negative edited row (100 % required); the Sherman–Morrison denominator $1 - \gamma u^T M e_i = (1 - \gamma p'_{ii}) / (1 - \gamma p_{ii}) > 0$ is a theorem for $\gamma < 1$ (`refute_instrument_math` §2.7) and is printed as a check, not censused.
12. Measured support of β and of $\text{rowsum}(P)$ on every **trained** cell: Ruling 2 makes β learnable, and at $\beta \neq 1$ rows sum to 1.31 ... 10.29 (`READ V16_ARM_SMPRIME.md:28-32`), where no row-stochastic theorem and no Neumann certificate applies — $\|P\|_{\infty} > 1 \Rightarrow \text{no certificate is printed}$ (`refute_theory_occvac` §1.2).
13. Measured support of γ across seeds with the LR verdict (PINNED / MOVED / interval).
14. For the segmentation rows only: at least one draw with an exact zero gate (BED-M: 3 of 3 values, `READ lean/CEQ/V16Domain.lean:304`); on BED-S at the softmax corner the count is 0 of N and the F0 theorems are **silent** (`refute_instrument_occvac` §2.6).
15. For the obstruction theorems: the vacuity inequalities printed beside every citation — $h \cdot m \cdot p = 512 \geq 64$ ([104] Thm 1), $n \log_2 n = 384 < 544$ ([92] Thm 1), $64^{\{1/16\}} = 1.2968$ ([25] Thm 1.1) — all vacuous at $s = 64$, $d = 16$, $p = 32$; conditional bounds named conditional in the same sentence (V-25, P-10).

A.8 Binds, with planted negatives and the $O(1)$ failure each must produce

The full row-by-row ledger is `bind_ledger.md`; this table is the specification. A bind is admissible only if its honest construction passes **and** every plant fails at $O(1)$ with the counts printed (`READ workdonenewseal.md:114-122`, the 0.9749 / 0.9165 / 1.000 / 0.4845 pattern). Every plant must enter at the **front door** — `bed_s.build` $\rightarrow$ `arm.forward` $\rightarrow$ `journal row` $\rightarrow$ `verdict()` — as a planted-negative cell FOUND under `results/` with its own `kind` and `manifest_hash` (V-14; `READ V20_R15_WING_MANIFEST.md:59-63`), not as an identity check on random logits.

bind	identity half (must pass)	rejection half (each plant must fail at $O(1)$)	plants that were found empty and are struck	mechanism
B-J parity at $\gamma = 0$	<code>torch.equal(0(0), P V)</code> against the lane's own <code>softmaxAttn (RUN[coord] True)</code> ; the record's own corner is <code>1.110223e-16</code> off <code>ceq/lm.py</code> on 19/64 entries, named not rounded	$\gamma = 0.5 \Rightarrow 2.3003; \beta = 0$ at the softmax corner $\Rightarrow \max \text{gap} > 0.5$; ChaCAL-published's <code>A_s</code> $\Rightarrow O(1)$ at every $\gamma > 0$	"a non-causal W " — passes bitwise at $\gamma = 0$ because $(I - 0 \cdot W) = I$ regardless of support <code>(RUN[MARS] True)</code> ; moved to B-P5	V-24, V-3
B-E1 boundary rows	$\mathcal{A} = \emptyset \Rightarrow$ <code>torch.equal</code> with <code>ChaCAL-diag</code> (the lane's own routine; declared V-3 with the diagonal convention in the manifest)	$\mathcal{A} \neq \emptyset \Rightarrow$ <code> \Pi_shape - \Pi_ChaCAL-diag _\infty</code> on downstream rows exceeds a printed $O(1)$ gap (a <code>weight</code> statement — the support is the causal triangle either way, <code>RUN[MARS] rows changed = [3]</code>); InfSA base $\hat{A} = \text{ReLU}(QK^T)/\ \cdot\ _F$ breaks the identity half (NOT MEASURED — needs the base wired)	"a non-identity absorbing row breaks $\Sigma_k q = 1$ " — the sum uses transient rows only, <code>RUN[MARS] max q - q_identity = 0.0</code> (<code>refute_instrument_math</code> row 12); replaced by the full-P read at $\gamma < 1$ where the plant reads 0.1491	V-24, V-3

bind	identity half (must pass)	rejection half (each plant must fail at O(1))	plants that were found empty and are struck	mechanism
B-E2 conservation	$q^{\{(0)\}} + \Sigma_k$ $q^{\{(k)\}} +$ $q^{\{(sink)\}} =$ $\not\models$ on T (RUN 1e-15; [47] Thm 11.6 [U], [66] [V])	drop BOS from every set $\Rightarrow$ the solve must raise ($\rho(Q) = 1$, V-16), never return a number; a constraint after the query $\Rightarrow q =$ 0.0 exactly and the draw is discarded by census line 4	“drop $\mathcal{A}_0 \Rightarrow$ $\max_k q \geq 1/K$ ” — with the sink set present the sum over constraints alone is < 1 and the pigeonhole is a V-3 of $\Sigma = 1$ (refute_instrument_math row 3); the degeneracy lemma is carried as a proposition (no goal and no sink $\Rightarrow \max_k \geq$ $1/K$), not as a plant	V-12, V-23
B-G1 interventional re-solve	(a) $V \equiv \not\models \Rightarrow$ $\max \Delta z \leq$ 1e-15 (RUN [SATURN] 0.0, RUN [MARS] 1.1e-16 — the exact zero is a code-path accident); (c) Sherman– Morrison closed form vs re-solve $\leq 1e-12$ (1.2e-15, 8.9e-16, 4.4e-16); (e) the C6 identity residual $\leq$ 1e-12	(b) $V \sim N(0, 1)$ $\Rightarrow \max \Delta z =$ O(1) (0.1096 / 0.363 / 1.127); (d') triangularity plant: a non-causal P makes $\Delta z[:i]$ $\neq 0$ (0.0761 / 0.0868); a two-row edit breaks the rank-one formula by O(1) and is repaired by Woodbury ([51] [V]); the route is named (solve_triangular gives $\Delta z[:i] =$ 0.0 and torch.equal True; LU gives 8.9e-16, False — D-5)	the EMC “feedback plant” — for $\gamma < 1$ the fixed point is unique, so settle-then-do equals do-then-solve on every P (RUN [MARS] 1.33e-15 and 4.4e-16 on dense cyclic P); what the plant showed is loss of triangularity, filed as (d')	V-24, D-5, D-7

bind	identity half (must pass)	rejection half (each plant must fail at $O(1)$)	plants that were found empty and are struck	mechanism
B-G2 vector metric	the journal carries the vector, a histogram and quantiles (the Q6 census read 0 of 40)	(i) the permuted oracle must not score 0; (ii) oracle + 0.1σ preferred to the permutation; (iii) $-\Delta z$ distinguished from Δz	—	V-26, L-14
B-H1 committor identity	$q = (I - Q)^{-1} R \not\llcorner$ against bed_1.committor on the bed's real sets: 0.0 at $A =$ $[0]$, $B = [1]$, $ T = 9$ (RUN[coord] corrected); the script raises on a missing key (the basin_A default that read 0.858 is the filed V-16, sec_refuted.md C12)	the must-fire perturbation drives harmonic_residual from $1.04e-17$ to $1e-6$ (ratio $9.6e10$); declaring $\mathcal{A}_k$ on the wrong set moves q by $O(1)$ (refute_theory_occvac §2.1); Kirchhoff second route for $K = 2$ single-node boundaries < $1e-10$	—	D-2, V-3, V-16
B-H2 argmin	the argmin from the oracle tensor equals the label's argmin on 100 % of admitted draws (V-3 declared)	the planted leak (a move token carrying its own q) must make the corpus-alone probe read $R^2 \geq$ 0.99 ; the 0-hop MLP within the CP half-width at $N = 8$ of the shape, or McNemar $p >$ 0.05 against 1-hop softmax, strikes the bed (C8)	—	V-24, C8

bind	identity half (must pass)	rejection half (each plant must fail at O(1))	plants that were found empty and are struck	mechanism
B-P5 triangular solve	<code>solve_triangular</code> vs dense inverse 1.8e-15	a dense (non-triangular) M passed with <code>upper=False</code> must disagree with the dense inverse (the plant moved here from B-J)	—	M-8
B-M schedule	the do-nothing schedule always entered (<code>READ</code> <code>THEORY.md:162-165</code>); on a schedule that drops tile (<code>k,1</code>) but keeps (<code>k,m</code>), (<code>m,1</code>) the exact solve's (<code>k,1</code>) block is reported non-zero and the exact route refused unless a δ is printed; the empty-row NaN and negative-index loads of <code>READ</code> <code>ceq/mz_kernel.py:13-21</code> stay structural	—	— (unpriced; backward through <code>kernels#22</code> does not exist, <code>READ</code> <code>THEORY.md:223-224</code>)	V-9, L-CERT

Predictions beside kills (D-7, L-SIGN). A counter without its prediction leaves no residual to sign; the arena carries one line per row, filed before step P3 of the plan: *prediction* `<statistic>` `<direction>` `<threshold>`, *counter* the kill as written, and the band between them named SPLIT. The five bets: A (Δz field cosine above depth-1 softmax by more than the cosine's own realised MDE₈; counter: within), B (`CP_upper(err_shape) < floor_exact(k = 1)`; counter `CP_lower` $\geq$; SPLIT the straddle), C (γ MOVED on $\geq 6/8$ seeds under the boundary null; counter PINNED on $\geq 6/8$; SPLIT $3/8-5/8$), D (the depth-5 skyline within MDE₈ of the shape's argmin accuracy; counter: short by more than MDE₈, sign logged — the P2 of `sec_beds.md §6.E` that `design_falsify.md` asserted in its headline sentence and filed nowhere), E (harmonic-residual ratio ≥ 10 at marginal NRMSE within MDE₈; counter ≤ 2 ; (2, 10) SPLIT). The (j) LM γ row carries no calibration row until an LM cell runs (V-1, M-20).

A.8.1 The propositions the apparatus rests on, as the refuters left them

Every proposition is a load-bearing wall of the plan, not a result; the status token is after the six refutations. “CLASSICAL” means the identity is owned and the paper cites the owner before it names the object.

#	statement, as repaired	status	owner / carrier	mechanism
P1	at $\gamma = 0$ the read is $W \vee$, bitwise in IEEE-754 forward substitution for finite entries; rejection region $\gamma = 0.5 \Rightarrow 2.3003$	KEEP (RUN; Lean <code>gamma_zero_is_softmax</code> [M])	[42] Eq. 5 at $\gamma = 0$; <code>main.three_corners_containment</code>	V-24, V-3
P2	corner 3 is $(I - A)^{-1}$ of the sub-diagonal chain, entrywise; BED-M's label is its last row	KEEP (0.0 exact; <code>6.2e-15</code>)	<code>Nilpotent.occupancy_2</code> (BED-M, <code>V15.chain_path_produced</code> , not [34] Def. 3.1 for the mask)	D-2 (BED-M, not won)
P3	$(1 - \gamma)z_i = E_i[\gamma^{\{\tau_k\}}]$, read $0_i = E_i[\gamma^{\{\tau_k - 1\}}]$ for $i \in T$; $\lim_{\gamma \uparrow 1} = q^{\{k\}} = [(I - Q)^{-1}R_k]_i$; $\Sigma = \mathbb{K}$ on T when every absorbing state is declared	CLASSICAL ([66], [47] Thm 11.6 [U], [82], [37]); RUN on the arm's P	the <code>3.8e-07</code> evidence line of the theory design is impossible (the gap is $\geq 1 - \gamma = 1e-6$ on every transient position, <code>refute_theory_math</code> §P3) and is struck; the read-form/z-form ratio is exactly γ (RUN)	D-2, V-12, V-25, P-1

#	statement, as repaired	status	owner / carrier	mechanism
P6	in the <code>pathProd</code> base a zero magnitude <code>m_c = 0</code> makes <code>W_ij = 0</code> for $j < c \leq i$ (<code>V16Domain.pathProd_eqs_zero_iff</code>, :129), the resolvent is block-diagonal across the cut and the solve splits; the base of §1.1 of the theory design (<code>Hop, exp(scan g)</code>) cannot carry a zero (<code>path_product_corner_fails_at_a_zero_gate</code>, :424; <code>no_prefix_scan_represents_a_zero_gate</code>, :165), so on the softmax corner F_0 is by $-\infty$ mask only; a cut makes the segment head absorbing and severs the boundary sets behind it	REPAIR (base restated; the segmentation–committor incompatibility proposition with <code>RUN[MARS] P_cc = 1.0</code>, $\rho(Q) = 1.0$, $q^{\{0\}} = 1.0$, $q^{\{1\}} = 0.0$ past the cut)	[75] (mechanism), [138] Thm 1 (converse), the record’s Lean (the iff)	V-25, P-7, D-3, L-CERT

#	statement, as repaired	status	owner / carrier	mechanism
P7	what depth-1 softmax cannot do, exactly as licensed: (a) per-row form — definitional; (b) hops need depth — [105] Thm 4.2 upper (unconditional), Cor 4.3 lower (conditional); [25] Thm 1.1 and [104] Thm 1 unconditional but vacuous at $\mathbf{s} = 64, \mathbf{d} = 16$; (c) composition — [92] Thm 1 vacuous here ($384 < 544$); (d) the ground softmax owns — [38] Prop 4.2 on single-location regression, whose label model admits 0 of 3 record beds; “depth-1” means one parameter set; the exact solve is DET-class ([31]; linear <i>equalities</i> are in $\text{DET} \subseteq \text{NC}^2$, the P-complete problem is linear <i>inequalities</i> — <code>refute_theory_math</code> §3.2), $O(\log^2 \mathbf{s})$ circuit depth by recursive block inversion, $O(\log \mathbf{s})$ matmul rounds, NOT MEASURED	REPAIR	the <code>hop_k</code> → <code>committor</code> reduction is NOT FOUND; obstruction 2 is a lineage argument	P-10, V-25, V-17, P-8

#	statement, as repaired	status	owner / carrier	mechanism
P8	the safest move: label ($q^{\{0\}}$), $q^{\{1..K\}}$, $q^{\{\text{sink}\}}$ per move at $\gamma = 1$; $a^* = \text{argmax}_a$ $q^{\{0\}}$; Chebyshev and lexicographic as columns; degeneracy lemma: with no goal and no sink, $\max_k q^{\{k\}}$ $\geq 1/K$; the one-step rule is not a policy ([84] [V] in Limits)	KILL→applied (substrate, label read, floor, cost clause all repaired above)	[14], [2], [114], [1], [8], [116], [119], [57], [60]	D-2, V-12, V-8, V-17
P9	$\Delta z = (I -$ $\gamma P')^{-1}(\Delta V$ $+ \gamma \Delta P \ z)$ exactly; for a row clamp $\Delta P = e_i$ u^T , Sherman– Morrison with denominator $(1$ $- \gamma p'_{ii})/(1$ $- \gamma P_{ii}) > 0$ for every $\gamma < 1$ (RUN[MARS] 0.492623, 1.9); forward-only by triangularity; a token rewrite is rank $s - i$ and is priced as a suffix re-solve	REPAIR (scope)	[111], [51], [109], [94] Eq. 5, [18] §7.3	M-8, P-8, V-24
P10	the horizon dial: $\gamma = 0$ softmax, 0 $< \gamma < 1$ the hitting-time transform, $\gamma \uparrow 1$ absorption (a limit; $I - P$ is singular at every absorbing row at $\gamma = 1$)	KEEP (corollary); the pinning instrument as in §A.11	[14] (discount), [36] (the limit object), [99] (learnable per-head γ in attention), [44] (the safety discount)	V-9, M-2

#	statement, as repaired	status	owner / carrier	mechanism
C6	(i) the displacement identity; (ii) suffix re-solve = full re-solve when $P'_{\{i\}} = P_{\{i\}}$ and $V'_{\{i\}} = V_{\{i\}}$ (true for a softmax P because k_i , v_i enter rows $\geq i$ only); (iii) withdrawn as an EMC statement	REPAIR	[35] as a remark only	V-3, D-7

#	statement, as repaired	status	owner / carrier	mechanism
C8	regime N (StrictlyLower, nilpotent, any γ, no absorbing rows possible) and regime S (diagonal kept, γ < 1, Neumann) are a theorem boundary (Nilpotent.one_not_nilpotent); the masked-diagonal route is exact only on value channels with $V_0 = 0$ (RUN[MARS] 4.4e-16 vs 1.99e-07 on $\mathbb{K}_{\{\mathcal{A}_0\}}$ with 0 $\in \mathcal{A}_0$) — which is why BOS is a value-zero sink set; the C8 evidence 7.96e-08 = $\gamma^{\{32\}}$ was attained at (0,0) and reads the same with the diagonal masked (V-4), so the discriminating number is $\max_{\{i \geq 1\}}(\gamma P)^{\{32\}}_{\{ii\}}$ = 6.27e-13	REPAIR	V15Fork.Asink_row_V-4m, V-4, P-3 for the sink row	

A.8.2 Lean targets, restated (P0.1; grades [M]/[S]/[D]; every [M] resting on a [U] Mathlib name is [S] until lake build is green — P-11)

target	statement, as corrected	grade	refusal shipped	depends on
gamma_zero_is_softmax	$P \max(1 - (0:\mathbb{R}) \bullet P)^{-1} * V$ $= P * V$	[M]	gamma_half_is_not_softmax	

target	statement, as corrected	grade	refusal shipped	depends on
bos_row_is_absorbing	$\text{bin} = 1$, causal softmax row 0 = e_0 (at $\beta = 0$ row 0 reads $\exp(qk_00) =$ 2.0138, RUN[MARS])	[M]	—	—
later_boundary_unreachable	lower-triangular $M, j > i \Rightarrow$ $(M^{-1} *_{\nu} e_j) i$ $= 0$	[M]	—	lower_triangular_isUnit
softmax_corner_not_nilpotent	lower-triangular $P, 0 < P i i \Rightarrow$ $(\gamma \bullet P)^k i i =$ $(\gamma P i i)^k >$ 0	[M]	strict_lower_is_nilpotent (in tree)	
lower_triangular_disUnit	Unit, of $1 -$ diag_one_sub_smul_P in $[1 - \gamma,$ 1); unit	[M] $\rightarrow$ [S] until Matrix.det_of_lowerTriangular [U] is confirmed	zero_diag_not_unit	
resolvent_fromBlocks	StrictlyLower diagc $\Rightarrow$	[M] $\rightarrow$ [S]	neg30_gate_does_not_cut (inv_fromBlocks_replacesUnit_iff	
segmentation_block_diagc	occupancy A n zero across the cut; block inverse	[U])	1e-300 witness)	
cut_makes_segment_head_absorbing (new)	Unit, causal $c = 1$ and every boundary set before c unreachable from $i \geq c$	[M]	—	—
displacement_identity	$(\gamma \bullet P') *_{\nu}$ $(z' - z) =$ $(V' - V) + \gamma$ $\bullet ((P' - P) *_{\nu}$ $z)$	[M]	const_value_zero_displacement	
pathprod_is_chain_resolvent	no resolvent subdiag a equals the path product entrywise	[S]	forward_map_fills_in at general n	

target	statement, as corrected	grade	refusal shipped	depends on
neumann_truncation_bound neumann_tail_attained	on bound in $\{0, 1\}/(1-\gamma) \cdot M$; equality of row sums for $P \geq 0$ stochastic; the refusal's docstring says "can fail", not "no such bound holds"	[S]	nonstochastic_breaks_bound at a convergent γ ($\gamma \cdot 3/2 < 1$)	
resolvent_is_triangular_solve	Angular solve fwdSub M v	[S]	—	lower_triangular_isUnit
mixing_matrix_row_stochastic	$P \cdot (1 - \gamma \cdot P)$ row sums 1, entries ≥ 0	[S]	bare_read_row_sum = $1/(1-\gamma)$	Neumann
causal_forward_only	$M^{-1} * v$ e_i) j = 0 for j < i	[S]	dense_P_displaces_lower_triangular	lower_triangular_isUnit
suffix_resolve_eq_with_P'_{<i> = (renamed from emc_causal)	full P'_{<i> and V'_{<i> = V_{<i>, suffix re-solve = full re-solve	[S]	dense_P_breaks_suffisplacessolve_identity, causal_forward_only	
committor_is_resolvent_readr, (a)	Q, readr, IsUnit (1 - Q) $\Rightarrow q = (1-Q)^{-1}$ $*_v r$	[M]	—	—
(b-causal) isUnit_one_sub_transient_causal (new)	for Q with Q i i < 1 on T ($\Leftarrow 0 \in$ $\mathcal{A}$), IsUnit (1 - Q)	[M]	bos_undeclared_singular	lower_triangular_isUnit
(b-chain) isUnit_one_sub_on_topperon (the oracle's non-causal Q only)	Perron certificate supperon (the oracle's non-causal Q only)	[S]	unit_needs_rho_lt_one	
(c) bed1_committor_eq	BED-1's solve = $(1-Q)^{-1} *_v (R$ $*_v \not\llcorner_B)$	[S]	—	(a)

target	statement, as corrected	grade	refusal shipped	depends on
lowerTriangular_ne_symmSupport (new)	lower-triangular P is never a SymmSupport Q with an off-diagonal positive entry — the D-2 separation for regime S (oracle_ne_resolvent covers StrictlyLower only)	[M]	zero_not_a_counterexample pattern	
reach_avoid_sum_one	ones of P sum to 1, IsUnit (1-Q), the sets exhaust the absorbing states $\Rightarrow \Sigma_k (1-Q)^{-1} *_{\nu} (R_k *_{\nu} \not\models) = \not\models$	[S]	no_goal_no_sink_forces_max_ge_inv_K	
sherman_morrison	low-rank-one displacement for a row clamp; docstring: a token intervention is rank s - i	[S]	—	lower_triangular_isUnit
masked_sink_finite (new)	the sum 0 $\Rightarrow$ occupancy $(\gamma \bullet P_m) n *_{\nu} V = (1 - \gamma \bullet P_m)^{-1} *_{\nu} V$ for the sink-row masked operator	[S]	goal_channel_breaks_finite_sum (V_0 = 1)	
hitting_time_transform	Reform, probability	[D]	—	—
f1_cantelli_union	clause; the F1 union bound with the $1/(1-\gamma)$ amplification			

A.9 Certificates — what δ , on what object, in vector units

route	object certified	δ (operator)	vector-unit statement printed	planted negative	status; mechanism
exact solve (<code>solve_triangular</code>)	<code>z</code> , 0 at any γ the committor at $\gamma = 1$ on the arm's triangular <code>Q</code>	0 (F0)	<code>torch.equal</code> on zero blocks; <code>max z_solve</code> – <code>z_dense </code> = <code>1.8e-15</code>	a dense <code>M</code> with <code>upper=False</code>	live on every bed; <code>lower_triangular_isUnit</code> [<code>M</code>]; M-8
Neumann at K hops, $P \geq$ 0 row- stochastic incl. absorb- ing rows, $\gamma <$ 1	<code>z</code> (bare) / <code>Π_γ V</code> (read)	matrix residual = $\gamma^{\{K+1\}}/(1$ – $\gamma)$ bare, = $\gamma^{\{K+1\}}$ for <code>Π_γ</code> — definitional (V-3 declared)	<code>max z_solve</code> – <code>z_K </code> $\leq \delta$ · $\ V\ _\infty$ with $\ V\ _\infty$ printed; $1/(1$ – $\gamma)$ printed beside every δ (the record's $1/(1 -$ <code>â_max)</code> was undefined at all eight R1 seeds, <code>READ</code> <code>V15_R1.md:56)</code>	rows 1.5 at γ = 0.6: 7.29 vs 0.54 (convergent)	live only on Neumann routes; never wins in MACs ($K < 1$); V-17, L-CERT

route	object certified	δ (operator)	vector-unit statement printed	planted negative	status; mechanism
F1 mask (dropped row mass ε) on a Neumann route	0_mask	$\varepsilon/(1 - \gamma)$ plus the truncation δ_{Π} — the resolvent amplifies dropped mass (RUN: $s = 3$, $\gamma = 0.9$, $\varepsilon =$ 0.1 gives $\ 0_{full} -$ $0_{mask}\ _{\infty}$ $= 0.5263$, $\varepsilon/(1 - \gamma) =$ 1.000 , naive $\varepsilon = 0.1$ violated $5.26\times$; a random $s =$ 32 draw did not violate the naive bound, which is why a bind on random draws would pass the wrong certificate — V-10)	$(\varepsilon/(1 - \gamma)$ $+ \delta_{\Pi}) \cdot$ $\ V\ _{\infty}$ over $1,024$ drawn cells	the $s = 3$ closed-form instance	dormant until row M ships a mask; the exact solve on an F1 mask is refused (fill-in); L-CERT

route	object certified	δ (operator)	vector-unit statement printed	planted negative	status; mechanism
F0 segmentation (exact zero blocks)	the solve per segment	0	<code>torch.equal;</code> dividend D = $s^2/\Sigma L_m^2$ (C5) — BED-M's $31.06\times$ is the pair-count ratio $s(s+1)/\Sigma$ $L_m(L_m+1)$ (<code>refute_theory_math</code> item 26) and a corpus property (V-22); on the softmax corner the dividend is $1\times$	−30 logit	live on gated corners only; on BED-S at the softmax corner the theorems are silent
committor at $\gamma = 1$ on a truncated route	q	none — the Neumann δ is ∞ at $\gamma = 1$; the sup-norm substitute with $\ Q\ _\infty$ = 0.958 reads 19.2 at $K = 4$ against a true error 2.5e-03 (<code>RUN[MARS]</code>); the Perron- weighted certificate (<code>Contraction.PerronCertificate,</code> <code>isUnit_one_sub_of_perron</code> [S]) needs a weight w that is NOT MEASURED	—	—	the committor head therefore ships on the exact route only ($\delta = 0$); the H2 threshold form is $\max_k$ $q^{\{(k)\}} +$ $\delta \cdot \ V\ _\infty \leq$ δ_{thr} with $\ V\ _\infty = 1$ for indicator values and δ = 0 on the exact route — and it certifies the solve, never the model ($P \neq P_{env}$ has no certificate; no move is “admitted” by it)

route	object certified	δ (operator)	vector-unit statement printed	planted negative	status; mechanism
coarsening (far field)	a coarse-far schedule	the mean-pool bound of READ ceq/multizoom.py:38-49	in $\ A - \tilde{A}\ _\infty$	—	the only printed far-field bound in the tree; a Mapper cover with a comparable δ is NOT FOUND V-25
β -census guard	every certificate above	—	$\ P\ _\infty > 1$ $\Rightarrow$ no certificate printed	—	

A.10 The identity manifest — field list (against L-2, READ V20_R15_LEAP_LEDGER.md:23)

The record's `CONFIG_FIELDS` (READ `scale/identity_manifest.py:67-71`) carry `beta` but not `qk`, `g`, `gamma`, and the manifest reports a missing declared field as `absent` outside the hash (:141-151). In the shape lane a missing declared field is a **refusal**. The manifest ships its own planted negative: flipping any one of `beta` / `qk` / `g` / `gamma` / `boundary_sets` by one unit must move `manifest_hash` (the one-line-drift test of READ `V20_R15_WING_MANIFEST.md:148-152` is the pattern). A cell is FOUND iff `results/` holds a record with the arm's kind.

field	support	why it is identity	mechanism
<code>beta</code> , <code>qk</code> , <code>g</code>	$\beta \in \{0,1\}$ at the corners (learnable per Ruling 2a); <code>qk</code> switch; gate-on flag + parametrisation id	the corner (<code>beta_one_row_is_one</code> , READ <code>lean/CEQ/V16Domain.lean:509</code>); effect sizes 3.522037 (QK), 0.673101 (g)	L-2
<code>gamma</code> : init, final γ , Λ , verdict	$\gamma \in [0, 1)$	row J; the parameterisation decides the LR null (§A.11)	C5, V-17
<code>diag_convention</code>	{kept, removed}	ChaCAL-published vs ChaCAL-diag; regime S vs N inside the inverse	C8, V-3
<code>committor_route</code>	{ <code>dirichlet_gamma1</code> , <code>discounted_gamma</code> }	which object the committor head reads	FATAL-2
<code>route</code>	{ <code>solve_triangular</code> , <code>neumann_K</code> , <code>segmented</code> , <code>csr</code> } + <code>K</code> + chunk <code>C</code>	δ is meaningless without it	L-CERT, P-8
<code>boundary_sets</code> , <code>goal_set</code> , <code>sink_set</code>	sorted position lists + sha256 of each	the mechanism under test; F1	E1, E2, V-12

field	support	why it is identity	mechanism
K, m, t_star (tolerance dial, E2_DIAL_TOL)	ints	the dials that must vary	D-3
S, D, d_model, n_train, n_eval, steps	ints	S = 64 on 40/40 was the D-3 instance	D-3, C9
seed, rng_plan	int + RNG_PLAN dict (READ scale/identity_manifest.py:78)	reproduction	P-1
device, threads, dtype, torch_version, cublas_workspace, deterministic_regime	strings/ints	a cell's identity includes its thread count and flag regime (READ workdonenew.md:363-366; READ V17_R4_RETAK.md:14-16)	M-10, M-16, Ruling 1
instrument_hash, manifest_hash	sha256	which scorer read the cell; which arm produced it (READ scripts/v15_r1.py:174,819-830)	Ruling 5, V-3
producer_cmd	the exact one-liner or script invocation	a number with no live producer is P-1	P-1
delta_vec, delta_bare, one_over_lmg, V_inf	floats	the certificate in vector units with its inputs printed	L-CERT, V-17
census	the §A.7 block for the row	a theorem without its census is decoration	V-25, V-8
floor_exact, floor_zeroinfo, floor_fano_k1, ceiling_hop_k	floats	every capability number beside its floor	L-FLOOR
sky_depth, sky_width, sky_cot, chacal_gamma, chacal_diag, params_per_arm	settings and counts	Δ_{sky} mandatory; Ruling 3 per-arm count	R-SKY, Ruling 3
void_contrasts	the §A.1 list	pre-registered before any cell	D-2, M-7
supersedes	pointer	journals never move	L-G2, P-3

A.11 The pricing rule

1. **N = 8 minimum, deduplicated by seed** (REQUIRED_SEEDS = 8, READ scale/it11_verdict.py:59); at N = 5 the finest two-sided p is 0.0625 (M-9, V-5).
2. **Minimum detectable effect**, paired, $\alpha = 0.05$ two-sided, power 0.80, noncentral t (RUN[VENUS], re-verified RUN[MARS] to 0.0398266):

paired sd	provenance	n = 8	n = 16
0.010101	BED-M $t^* = 2$, n = 2048 seed sd (READ MISTAKES.md:1096)	0.011677	0.007570
0.034451	record row 2	0.039827	0.025820
0.050146	negation_scope pilot sd (M-3)	0.057971	0.037583
0.109199	realised sd, $2.18\times$ the pilot	0.126238	0.081841

BED-S's paired sd is **NOT MEASURED**; the first eight seeds fix it and the row is filled before any prediction is scored; any contrast smaller than the $n = 8$ cell of the **realised** sd may not be claimed (M-3). The MDE is in NRMSE units; the cosine and the φ -NRMSE carry their own realised sd (V-17).

3. **Parity.** At $N = 8$ parity is claimed by identity bind only, never by TOST (M-13). The equivalence half (shape at γ vs ChaCAL-diag at the same γ , $\mathcal{A} = \emptyset$) is the only contrast TOST decides (the E1 kill fires or a difference is detected; no branch consumes TOST there — M-7). The $N = 70$ figure is the **two-sample** requirement (0.7975 at 69, 0.8073 at 70, half-width 0.2799σ , RUN[VENUS], RUN[MARS]); the arena is **paired**, where the margin is in units of **sd_d** and power 0.80 is reached at $N = 36$ if **sd_d** = σ (Monte Carlo 200,000 draws, RUN[MARS]) and at $N = 70$ only if pairing buys nothing. Both are priced; the realised **sd_d** chooses ([108] [V]).
4. **Thread-count floor** $2.345e-3$ (READ scale/it11_verdict.py:133-137); a margin below $\approx 4.7e-3$ is indefensible; the lane is read from the journal, never from the machine (M-10, M-16). CUDA **cumprod** backward has no deterministic kernel (READ COSTS.md:143-156); **solve_triangular** ran bitwise forward and backward over 8 repeats under strict mode on this card with no torch-documented guarantee (RUN[NEPTUNE]), so the softmax-corner shape is the first **gated** wing whose training step runs under **use_deterministic_algorithms(True)** (the control softmax was never tested — V-23).
5. **Per-cell prices on the certified RTX 4060 Laptop** (7.996 GiB, torch 2.5.1+cu121; laws [FITTED] READ COSTS.md:73-74, solve increment [RUN] sec_cost.md §4.x.3; the increment is a per-op floor with the $2.0\times-6.6\times$ dispatch gap beside it, P-8; the clock is not stationary, $\pm 12\%$, so prices are quoted in GPU-minutes to two decimals, M-8):

item	price	class
one 150-step cell, shape on the softmax corner, $n = 2048$, $s = 64$	1.680 s	FITTED + RUN
the softmax control cell	1.524 s	FITTED
one bed-cell pair, 8 seeds each, + 4.0 s fixed	≈ 34 s	FITTED + RUN
the pair on the corner-3 base (arm_smprime)	≈ 133 s; no certificate on this base (rows sum 1.31 ... 10.29 at $\beta = 0$, V-25)	FITTED + RUN
the eight-seed arena at $t^* = 8$, $m = 8$, $K = 2$, seven arms on the softmax corner	$7 \times 8 \times 1.680 + 4.0 \approx 98$ s ≈ 1.6 GPU-min (the instrument's " ≈ 4 GPU-min" was a $2.4\times$ over-book)	DERIVED
InfSA-style $K = 16$ cell	≈ 5.1 s	DERIVED floor
depth-5 skyline cell	≈ 7.6 s	ASSUMED

item	price	class
paired TOST at N = 36 / two-sample at N = 70, one pair	$36 \times 3.204 + 4 \approx 119 \text{ s}$ / $70 \times 3.204 + 4 \approx 228 \text{ s}$	DERIVED
BED-S oracle per draw	m factorisations with K + 2 right-hand sides each, $ T ^3/3$ = 5.76e8 MACs (1.15e9 FLOPs under the house convention) at $ T = 1200$, under a second in float64, paid once per draw	DERIVED
a candidate move (query-side clamp) after the first solve	one forward-substitution column ($s^2/2$ MACs) + an s·d inner product; m = 8 columns cost m/d = 0.5× one solve at d = 16 (RUN[MARS] DERIVED-1); prediction 0.5×, counter ≥ 8× (a re-solve per candidate), SPLIT between	DERIVED
key-side (token) moves	m suffix re-solves, $\approx (s - i)^2 d/2$ MACs each	DERIVED
identity checks, Lean targets, census, manifests	0 GPU-s	—
the s-sweep of C9 ($S \in \{64, 256, 1024, 4096\}$, synchronize(), randomised order)	NOT MEASURED - needs the --seq-len flag and two synchronize() calls; the record's band 206-537 GPU-s (READ V20_R15_THEORY_TABLE.md:104-109)	band only
the full 9,600-step ladder, one pair	$\approx 36 \text{ min}$, spent only on a survivor	FITTED ×64

The whole falsification programme on BED-S at the design point is $17 + 98 + 119 + 2.6 \text{ s} \approx 4.0 \text{ min}$ of GPU time plus free CPU identities and Lean files (DERIVED from the rows above).

6. **The LR test for γ** (Ruling 10', READ V17K_RULINGS.md:389-436, restated for γ against $\gamma = 0$ with the null and the held-out set declared): if γ is trained on $[0, 1)$, $\gamma = 0$ is a boundary and the null of $\Lambda = 2[LL(\gamma) - LL(\gamma \equiv 0)]$ is $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ with 95 % point 2.706, not 3.841 (RUN[MARS]; the citation for the boundary null is **owed** — not in references.bib, not fetched); MOVED at $\Lambda > \ln n_{\text{eval}}$ ($\ln 4096 = 8.318$); the interval verdict between. The paper trains $\gamma \in [0, 1)$ (a signed γ would forfeit the Neumann equality). The $|\gamma| < 0.05$ rule is **deleted** (two verdicts for one cell, M-20). The mirror kill $\gamma > 0.99$ on $\geq 6/8$ stays on the $z/\Delta z$ channel with $1/(1 - \gamma)$ printed.
7. **L-FLOOR.** A number without its floor is refused at assembly.

A.11.1 Order and dependencies (D-1 dependency law; iteration count from the DAG's critical path, D-3)

step	items	price	independent of	gate at the end of the step
P0	Lean targets of §A.8.2 (P0.1); the manifest with its drift plant (P0.2); the vacuity lines of §A.7 item 15 (P0.3)	0 GPU-s	each other	L-LEAN: no arm is trained before its identity theorems are green; a target that does not build is [S], never a claim
P1	bed_s.build specification with the DAG substrate, sink/goal/constraint placement and query-side clamps (P1.1); the census of §A.7 on 512 draws (P1.2); the zero-hop and untrained-arm reads (P1.3); the corpus-alone probes with their planted leak (P1.4)	0 GPU-s	P0 (it may run first)	D-4: a bed not admitted produces no reading; the author decides admission from the printed census alone
P2	the binds of §A.8 on BED-S's real draw through the front door (B-J, B-E1, B-E2, B-G1, B-G2, B-H1, B-H2, B-K, B-P5), each with its planted-negative cell FOUND under results/	seconds	— (needs P1 admitted)	V-14, V-24: every plant fails at O(1) with counts printed

step	items	price	independent of	gate at the end of the step
P3	the eight-seed arena at $t^* = 8$, $m = 8$, $K = 2$, seven arms on the softmax corner (P3.1); the depth-5 skyline (P3.2); the cached-mixture arm on $\Delta P/\Delta V$ plants (P3.3); BED-K's native control (P3.4)	≈ 1.6 GPU-min $+ \approx 7.6$ s	P3.2–P3.4 of each other	the realised paired sd fixes the MDE row (M-3); Bets A–E scored with their counters and SPLIT bands
P4	the E1 kill at MDE ₈ and, if not fired and not separated, the N = 16 escalation (P4.1); the residual ratio and cosine verdicts (P4.2); the γ -pinning LR verdict with the boundary null (P4.3); the paired/two-sample TOST on the $\mathcal{A} = \emptyset$ parity half only (P4.4)	≈ 119 –228 s for TOST	P4.1–P4.3 of each other	the two honest sentences of <code>design_falsify.md</code> §6, with “six” corrected to “five” machine-checked identities in the median sentence

step	items	price	independent of	gate at the end of the step
P5	deferred instruments: the s -sweep with synchronize() and randomised order (P5.1, 206-537 GPU- s band); a mask on the CSR path and the dormant certificate rows (P5.2); the long-context schedule with n per s (P5.3); a digraph β_0 instrument (P5.4); the width and CoT skylines (P5.5); the chunk-size sweep (P5.6); the Perron weight for the committor certificate (P5.7); the LM γ cell behind the author's explicit Kaggle yes (P5.8)	NOT MEASURED each	each other	none is on the critical path; the ladder (≈ 36 min) is spent only on a survivor of P4

The critical path is $P1 \rightarrow P2 \rightarrow P3 \rightarrow P4$; P0 runs beside it; the whole falsification programme at the design point is under 4.0 GPU-min (DERIVED, §A.11 item 5). The record's dominant failure shape — a bed that dies at construction (D-4, M-3, V-8) — is now a 0 GPU-**s** decision the author takes from the P1 census; the expensive death (ChaCAL-with-sink matching the committor read) is a ≈ 34 s reading escalated to $N = 16$, and what survives it is a property paper about a published operator with five machine-checked identities, a certificate discipline and a registered bed, which the paper is written to remain honest in.

A.12 Deleted (unanswered KILL) — one line each

- The EMC proposition and its “feedback plant” (theory C6(iii), falsify §3.10, instrument G1(d)): for $\gamma < 1$ the post-intervention fixed point is unique, so both orders coincide on every P (RUN[MARS] 1.33e-15, 4.4e-16 on dense cyclic P); retained as one remark citing [35], [86], [123] and, for why a cached resolvent needs the re-solve, [85] / [101].
- Row L, the influence-Jacobian β_0 barcode on BED-S: $\beta_0(\varepsilon = 0) \in \{1, \mathbf{s}\}$ on 100 % of softmax draws; **beta0_interleaving** consumes point clouds (READ ceq/certs/topological.py:476-505); the Turner hypothesis is unchecked on a similarity. Reversal needs a digraph β_0 instrument and a gated corner: NOT MEASURED. The descriptive owners are cited in the prior-art table regardless ([71], [72], [68] [V]).

- The regime-N survival route for component (e) (falsify §5 rank 2): an absorbing row breaks `StrictlyLower` (`rank(I - A') = 7 of 8`, `RUN[MARS]`).
- The $m(K+1)$ -right-hand-sides sentence for the committor channel (theory P8): $q^{\{k\}}$ has no V .
- The $|\gamma| < 0.05$ pinning rule (instrument J).
- The `1e-300` gate plant (instrument K): float64 underflow.
- The “ChaCAL must emit chance” control (instrument H1): a control that must fail by construction is V-2.
- The $31.06\times$ segmentation dividend as a fallback bet (falsify §3.8): $1\times$ on the softmax corner; carried only as BED-M’s corpus property.
- Marginal W1, Procrustes, position-free OT as state metrics.
- `floor1` as an information floor anywhere.

A.13 Limits (collected once)

Every `RUN` in this file is one numpy float64 draw at `s = 32`, seed 0, Gaussian logits, on this CPU — an identity or counterexample check, never a statistic; the two one-liners were: (i) `L = rng.standard_normal((32,32)); L[triu(1)] = -inf; P = softmax(L)`, rows {0} sink, {5,6} goal, {9,10}, {15} constraints set to `e_i`; `Q = P[T,T]`; `N = (I - Q)^{-1}`; `q = N P[T, S]` $\not\models$ per set; `z = (I - 0.6P)^{-1} \not\models_S`; `z`-form $(1 - \gamma)z$, read-form $(1 - \gamma)Pz$; (ii) the `s = 3` mask instance `P = [[1,0,0], [ε,1-ε,0], [0,1,0]]`, $\gamma = 0.9$, $\varepsilon = 0.1$, `V = e_0`, with `P_10` dropped and row 1 renormalised; plus the Fano arithmetic. BED-S has no cell, no realised sd, no measured `t*`, no `I(X_{≤k})`; `a*`; every BED-S number is a floor formula or a design constant. The DAG-in-token-order substrate is a specification; its adjacency-feature representability at `d_model ≥ n_nodes` is unmeasured. The Kirchhoff second oracle covers `K = 2` single-node boundaries only. The committor head carries no truncation certificate until a Perron weight exists; the F1 mask certificate is dormant until a mask ships. The paired-TOST `N = 36` is a Monte-Carlo figure at `sd_d = σ`; BED-S’s `sd_d` chooses. The boundary-null critical value 2.706 is stated without a `references.bib` citation and is marked owed. Theorem numbers inside cited sources are the sweeps’ HTML readings and inherit their [U] marks; ChaCAL’s diagonal-removed inverse is [V-fetched] by one refuter from the HTML render and should be re-read against the PDF before it is typeset. Three Kemeny–Snell keys coexist in `references.bib` ([65], [66], [67]) and the assembler collapses them to one. No code file, no git write, no Kaggle contact, no external fetch was made by this judge.

7. Prior art and occupancy

JUPITER (MYCROFT), 2026-09-03, HEAD 207e7b9. Inputs read in full: the brief; both corrections; the outline; `sec_intro_draft.md`; the eight sweeps `sweep_{resolver,linrec,expressivity,causality,safety,topology,occupied,methods}.md`; `judge/sec_shape.md` §4.0, whose occupancy paragraph this section must agree with and go wider than; `bib_audit.md`; `bib_aliases.md`; `references.bib` (canonical keys only — an alias key is never printed). Evidence classes: `RUN` (executed on this box this session), `READ path:line`, `CITED [V]` (abs page, DOI/registry record, catalogue record or proceedings page fetched by the named sweep this session and the title matched), `CITED [U]` (reached through a search index, a secondary paper, or memory), `DERIVED` (steps written out). No number without provenance. Struck constants (`STRUCK.md`, 12 entries) do not appear.

The rule this section exists to serve is the brief’s rule 7: *occupied components are cited before they are named, and absence is recorded as NOT FOUND with the query that bounds it, never as “novel”*. It is the same rule the record breaks in §7.3.

7.0 How to read the table

Verdicts. **OCCUPIED** — one named source states, in print, the equation or mechanism with the same object on both sides. **NEAR-MISS** — a named source states a neighbouring object and the table says which clause is missing. **NOT FOUND** — no fetched source states it, and the row carries the queries that failed; the totals in §7.2 bound how much that absence is worth (**MISTAKES.md** V-7, a search structurally incapable of finding anything read as absence, **MISTAKES.md:117**).

Marks. **[V]** and **[U]** are per source, as the sweep that found it recorded them. **[V]** here is a title match against a fetched identifier page; it is **not** an equation-level check. Where a sweep read the body and transcribed the equation the row says so; those are the only rows on which an equation number is load-bearing.

Boxes. Six of the shape’s structural properties were fixed *before* the first fetch (**sweep_resolvent.md** §0, against M-2, a threshold refitted to the data it judges): (a) **P** is content-dependent, (b) causally masked, (c) row-stochastic softmax, (d) the read uses the exact resolvent, (e) absorbing rows on named sets, (f) softmax parity at $\gamma = 0$; and four more — (g) the interventional re-solve, (h) the safest-move decision, (i) a printed truncation certificate, (j) learnable γ . The delta letters in §7.2 are those boxes plus (k) machine-checked containment and the zero-gate iff, (l) the influence barcode restricted to the F0 endpoint, and (m) the cover-to-schedule path.

7.1 The occupancy table

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
1	the read operator $O = (1 - \gamma)P(I - \gamma P)^{-1}V$ on causal row-stochastic softmax P	OCCUPIED	[42] [V] , Eq. 5, EMNLP 2024 (body read)	Discounted path sum over the causal softmax matrix, written as one closed-form read with $\gamma \in [0, 1)$ fixed at 0.9.	No boundary rows, no committor, no intervention, no truncation bound; diagonal removed inside the inverse, so the read is sub-stochastic

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
2	the triangular solve	OCCUPIED	[42] [V] Eq. 7; [136] [V] Eq. 10; [143] [V]	Solve $(I - \gamma A)Y =$ $(1 - \gamma)AV$ by forward substitution rather than inverting; DeltaNet solves the same unit-lower- triangular system per chunk.	Nothing on the solve itself. CEQ's I5 is a reproduction, not a contribution
3	$\gamma = 0$ parity with softmax	OCCUPIED	[42] [V] (body: $\gamma = 0$ reverts to standard attention)	At zero discount the layer is the ordinary attention matmul.	The <i>machine- checked</i> containment of softmax, linear attention and the path product in one family; that is the record's Lean and is NOT FOUND elsewhere (letter k)
4	the discount / horizon dial γ	OCCUPIED	[14] [V]; [36] [V]; [45] [V]; learnable per head [99] [V]	The discount of policy evaluation; APPNP's teleport α ; InfSA learns one γ per head through a sigmoid.	A learnable discount inside a causal language- model read, pinned by a boundary- corrected likelihood- ratio test (letter j)

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
5	the row- stochastic base and its three corners	OCCUPIED	[120] [V]; [63] [V]; [34] [V] Def. 3.1	Row-softmax; kernel-feature linear attention; the 1- semiseparable cumulative- product mask that is the record's path product.	That the three are one family with distinct corners, proved rather than asserted (letter k, <code>V16Domain.three_corners_</code>
5b	why softmax leaves that lineage	OCCUPIED	[58] [V]	Row-softmax over $QK^\top$ has rank T even at rank-1 logits, so it admits no finite-state dual.	Nothing. This is a gift: it <i>proves</i> the exit point the record had asserted
6	the chunked / blockwise cost path	OCCUPIED	[143] [V] ($\tilde{O}(n^{4/3}d)$); [136] [V]; [12] [V]	Exact triangular solves on diagonal tiles with cross-block traffic routed through a reduced system; chunk tiling for linear-RNN kernels.	An exact block structure from a zero gate with a printed $\delta = 0$, against a learned block- plus-residual approxima- tion (letter l's sibling; §7.1 row 15)

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
7	absorbing rows as boundary conditions	OCCUPIED on fixed graphs; NOT FOUND in a causal content-dependent attention operator	[145] [V]; [144] [V]; [129] [V]; [13] [V]; [5] [V]; one row inside attention [61] [V]	Clamp labelled nodes, take the harmonic extension elsewhere; partially absorbing walks with per-node absorption; DoFormer fixes an intervened token and forbids it attending.	$K \geq 2$ constraint sets <i>plus</i> a goal set <i>plus</i> a declared sink, as identity rows of a causal row-stochastic content-dependent read (letter e). Queries: resolvent 17/22, linrec Q11/Q17, occupied S4/S5, methods S11
8	the fundamental matrix and absorption probabilities	OCCUPIED	[65] [V-cat]; [47] [U], Thm 11.6 (§11.2 body read by the safety sweep)	$N = (I - Q)^{-1}$, $B = NR$, expected absorption times $N\mathbf{1}$, in canonical form.	Nothing on the identity — the record concedes zero delta at V13_TIER6_PRIOR_ART.md:3 (READ). The <i>use</i> as a layer's read is elsewhere in this table
9	the committor / reach-avoid reading	OCCUPIED, twice, for one object	[82] [V]; [39] [V]; [40] [V]; [37] [V]; [114] [V]; [1] [V]	The discrete Dirichlet problem $(Lq)_i = 0$ with $q _A = 0$, $q _B = 1$; the harmonic/electrical reading; reach-avoid probability by dynamic programming.	The naming identity itself (reach-avoid $\equiv$ committor) is stated in neither literature as far as the safety sweep reached; the paper states it as a remark and claims nothing. The read inside a layer stays open (letter h)

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
10	the safest-move decision rule	OCCUPIED as names; NOT FOUND as a read	[119] [V]; [89] [V]; [132] [V]; [2] [V-cat]; [14] [V]; [8] [V-cat]; [115] [V] / [116] [V]; NEAR-MISS [60] [V]; framing [57] [V], [17] [V]	Chebyshev (worst- objective) action selection; max-min and lexicographic orders; K cost functionals on one chain; several target sets in one linear system with a verdict; a move from a linear first-exit solve; a transformer picking the lowest predicted violation probability.	The scores being committors read out of the layer's own operator, over candidate moves placed in the context (letter h). Queries: safety 2/5/10/12/13/17/18/19, occupied S8, expressivity S15

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
11	the interventional re-solve and its rank-one update	OCCUPIED as algebra; NOT FOUND as a trained read	[91] [V]; [112] [V]; [111] [V]; [51] [V]; [109] [V]; [94] [V] Eq. 5; [86] [V]; [16] [V]; [18] [V] §7.3; [106] [V]	<i>do</i> ($\cdot$) as row surgery on $x =$ $(I - B)^{-1}e$; rank-one inverse update; perturbation of the fundamental matrix; the barrier update of a default repre- sentation; consequence as displacement of an equilibrium; a causally masked transformer whose <i>do</i> () is clamp-and-re- solve.	The displacement Δz as a jointly scored vector label at two stated prices — row clamp $O(sd)$ against token rewrite $(s - i)^2 d/2$ (letter g). Queries: resolvent 19, linrec Q15, causality Q6/Q9, methods S10
12	the order-of- intervention question	OCCUPIED as a theorem; NOT FOUND inside an attention model	[35] [V] Def. 4, Thms 1–2 (PDF read by the causality sweep); [123] [V]; [15] [V]	When “equilibrate then manipulate” differs from “manipulate then equilibrate”, and when a reduced equilibrium model hides the feedback that makes it differ.	Nothing: for $\gamma < 1$ the map $x \mapsto V + \gamma P'x$ is a contraction with one fixed point, so the orders coincide (RUN[J] 8.9×10^{-16}) and the paper retires EMC to a remark rather than claiming a bind

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
13	equilibrium as the layer's output	OCCUPIED	[7] [V]; [49] [V]; [76] [V]; [46] [V]; [97] [V]; [54] [V]	A layer whose output is the root of $z = f(z, x)$, differentiated implicitly; algorithmic targets as equilibria; softmax attention as one Hopfield update.	All of these fixed points are nonlinear or per query . CEQ's is linear in z , which is what makes the certificate and the one-pass solve available; Hopfield's is the per-row independence obstruction, not its repair
14	the attention barcode	OCCUPIED	[71] [V]; [72] [V] (preprint, abs marks it not accepted); [26] [V]; [93] [V]; [102] [V]; [11] [V]	Threshold a head's attention matrix over a sweep, read H_0/H_1 persistence as classification features; zigzag over the layer axis; a two-subgraph topological divergence.	The filtering function : a filtration by the influence Jacobian of a resolvent read, with a stability δ from a <i>directed</i> theorem, whose $\varepsilon = 0$ endpoint is the exact segmentation (letter l). Queries: topology 4/8/16

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
15	segmentation as an exact block certificate	NEAR- MISS on mechanism; NOT FOUND as a certificate	mechanism [75] [V] (body read), [134] [V], [59] [V]; converse [138] [V] Thm 1	A data- dependent cumulative decay on softmax logits, $D_{ij} = c_i - c_j$ with $c =$ $\text{cumsum } \log f$; a fixed block mask confines reach at every depth.	The iff : an exponential prefix scan cannot carry an exact zero (<code>no_prefix_scan_represen</code> so the annihilation certificate is the record's and is NOT FOUND outside its Lean (letter k). Queries: topology 7/14
16	the truncation certificate $\gamma^{K+1}/(1 - \gamma)$	OCCUPIED as mathe- matics; NOT FOUND as an instrument	[83] [V]; [55] [V]; global-bound NEAR-MISS [99] [V] ($\ S_L\ \leq$ $\gamma/(1 - \gamma)$)	The Neumann tail of a matrix with $\ P\ _\infty \leq 1$; the spectral- radius convergence condition for a content- adaptive attention series.	A per- K bound printed beside a truncated attention read , with a planted non- stochastic negative and the $1/(1 - \gamma)$ mask amplification (letter i). No attention paper fetched prints any δ . Queries: resolvent 3, linrec Q6/Q14, methods S4

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
17	intervention effects on persistence	OCCUPIED	[68] [V]	Treatment effects defined on persistence- diagram summaries of the outcome, with a doubly robust estimator and a shape test.	Nothing on the estimand. If CEQ ever reads Δz through a barcode it is an instance of this and says so
18	the Mapper cover	OCCUPIED	[113] [U on the landing page; dblp conf/spbg/Singh and one further record agree]; [21] [V]; cluster covers [28] [V], [100] [V], [69] [V], [141] [V]	Lens, overlapping cover, partial MCO7, nerve graph; Mapper as a consistent Reeb-graph estimator with parameters fixable in advance; sampled/hashed/ means covers deciding who attends to whom.	A nerve-of- a-lens-cover as the candidate builder for a <i>resolvent</i> read, cover parameters fixed by the Reeb- estimator rule, candidate set carrying the Neumann δ (letter m). Queries: topology 6/15
19	the persistence- derived CSR schedule	OCCUPIED BY THE AUTHOR	[110] [V] (PR page and files tab fetched; merged 2026-07-28)	A forward-only Triton kernel consuming a causal CSR block schedule built from sink blocks, local windows and a 0D- persistence salience over key-block centroids.	The consumer: the schedule feeds a softmax kernel, forward-only, with no backward and no truncation certificate. The cover and the certified consumer are letter m

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
20	the depth law	OCCUPIED, and conditional where it is a lower bound	[105] [V] Thm 4.2 / Cor. 4.3; [103] [V]; [81] [V] Thm 2; [42] [V] Thm 1; [139] [V]; ceiling [79] [V]; class [31] [V]	k hops at depth $\lfloor \log_2 k \rfloor + 2$ (upper, uncon- ditional); $\Omega(\log k)$ conditional on one-vs-two- cycle; $\lfloor \log_2(n+1) \rfloor$ for entity tracking; log depth suffices for connectivity; depth unnecessary at linear width; log-precision constant depth $\subseteq \text{TC}^0$; NL $\subseteq \text{DET} \subseteq$ NC ² .	Nothing. CEQ does not evade the law; it relocates the log-depth into one linear solve, which is a DET-class computation. Every comparison therefore fixes depth <i>and</i> width, and three fellow approxima- tors sit in every table
21	single- location optimality (why the record could not win)	OCCUPIED	[77] [V] Cor. 2 (erf gate, asymptotic $d \rightarrow \infty$, $L = o(d)$); [38] [V] Prop. 4.2 / Cor. 4.3 (softmax proper); [142] [V]	An attention layer attains Bayes risk on scalar single- location regression; softmax beats linear attention exponentially in the spike; balancing weights are a per-query mixture.	Nothing. This is D-1's owner and the reason the plan's label class is chosen <i>away</i> from it. The record's own geometry has $L/d = 4.00$, the wrong direc- tion for Cor. 2 (MATHEMATICS.md §17.5)

#	component of the shape	verdict	owner(s), canonical key, mark	what it owns (under 25 words)	what it leaves open relative to CEQ
22	the statistical instruments	OCCUPIED	[108] [V]; [30] [V]; [78] [V]; [88] [V]; [4] [V]; [128] [V]; [122] [V]; [125] [V]; [96] [V]; [52] [V]; [6] [V]; floors [33] [V], [127] [V]	Two one-sided tests; exact binomial and score intervals; the paired- proportions test; ordered alternatives; pool-adjacent- violators; the martingale maximal inequality and e-processes; bounded- increment concentration; Fano and the Welch bound.	Nothing. The plan uses them and claims none. The <i>pairing</i> of a sign test against a per-row control on byte-identical draws is the record’s and is NOT FOUND outside it (causality Q10)

Agreement with §4.0. Every owner in `judge/sec_shape.md` §4.0 appears above under the same canonical key and the same verdict; this table adds rows 5b, 6, 12, 14, 15, 16, 17, 19, 20, 21, 22 and splits row 7’s fixed-graph owners from its attention near-miss. No row contradicts §4.0.

7.1.1 The cite-first sentences

Rule 7 requires that an occupied component be citable in one sentence *before* it is named. These are those sentences; the assembler places each at the first occurrence of its component.

1. **Read operator.** “The read is that of ChaCAL ([42], EMNLP 2024, Eq. 5), with the diagonal kept inside the inverse so that the mixing matrix stays row-stochastic.”
2. **Triangular solve.** “ChaCAL solves it as a lower-triangular system (Eq. 7); the per-chunk form of the same primitive is DeltaNet’s ([136], Eq. 10).”
3. **Discount.** “The discount is Bellman’s ([14]), the operator Dayan’s successor representation ([36]); learning one per head is InfSA’s ([99]).”
4. **Base family.** “The three corners are Vaswani et al.’s softmax, Katharopoulos et al.’s linear attention, and the 1-semiseparable mask of [34] Def. 3.1.”
5. **Absorbing rows.** “Clamping labelled positions and reading the harmonic extension is [145]; one absorbing row inside attention is [61].”
6. **Fundamental matrix.** “ $N = (I - Q)^{-1}$ and $B = NR$ are Kemeny–Snell ([65]); [47] Thm 11.6 states the same, and is [U].”
7. **Committer.** “The committer is the discrete Dirichlet problem of [82]; its control-theoretic name is the reach-avoid probability ([114]).”
8. **Decision rule.** “Selecting by the worst of several scores is Chebyshev scalarisation ([119]); as a safety filter it is [57].”
9. **Re-solve.** “The rank-one re-solve is Sherman–Morrison ([111]); on a fixed chain with a barrier it is [94] Eq. 5.”

10. **Equilibrium layer.** “A layer defined as a fixed point is [7]; on a causally masked transformer with *do()* by clamp-and-re-solve it is [106].”
11. **Barcode.** “Persistence of a thresholded attention graph is [71]; an intervention read through persistence summaries is [68].”
12. **Cover and schedule.** “Mapper is [113], its estimator theory [21]; a persistence-derived CSR schedule inside an attention kernel is the author’s own merged [110].”
13. **Depth law.** “ k hops need depth $\lfloor \log_2 k \rfloor + 2$ ([105] Thm 4.2; the matching lower bound is conditional, Cor. 4.3); the same law is reached independently for entity tracking by [42] Thm 1.”
14. **Optimality.** “One attention layer attains Bayes risk on single-location regression ([77]; for softmax proper, [38]).”
15. **Certificate.** “The Neumann tail bound is textbook ([83]); no attention paper fetched prints one.”

7.2 The delta, narrowly

No fetched source combines the following in one attention operator, and every part of every clause is owned above. **(e)** $K \geq 2$ absorbing constraint sets, a goal set, and a declared sink as identity rows of a content-dependent, causal, row-stochastic read — bounded by `sweep_resolvent` searches 17 and 22, `sweep_linrec` Q4/Q11/Q17, `sweep_occupied` S4/S5 (S5 returned **0** hits containing “committor” in a transformer context), `sweep_methods` S11 and `sweep_topology` 5/13; the nearest things found are InfSA’s single uniform absorption leak, DoFormer’s one clamped row, and the attention *sink*, which is a column device and not a row condition. **(g)** the interventional re-solve with the displacement as a jointly scored vector label at two stated prices — bounded by `sweep_resolvent` 19, `sweep_linrec` Q15, `sweep_causality` Q1/Q3/Q6/Q9/Q16 and `sweep_methods` S10; Q9 returned nothing at all on interventions in deep-equilibrium layers. **(h)** the committor vector and the two safest-move rules over candidate moves placed in the context — bounded by `sweep_safety` 2/5/10/12/13/17/18/19 (query 14 established that the word “committor” does not occur in the control literature), `sweep_occupied` S8, `sweep_expressivity` S15 and `sweep_causality` Q11/Q15. **(i)** a printed certificate with a planted negative and the $1/(1 - \gamma)$ mask amplification — bounded by `sweep_resolvent` 3, `sweep_linrec` Q6/Q14 and `sweep_methods` S4; the inequality is textbook and is claimed only as an instrument, never as mathematics. **(j)** a learnable discount in a *causal language-model* read pinned by a boundary-corrected likelihood-ratio test — bounded by `sweep_resolvent` 2 and 3 and `sweep_linrec` Q1/Q16 (Q16 and `sweep_occupied` S12 each returned **0** hits for the successor representation used as an attention operator); ChaCAL fixes $\gamma = 0.9$, InfSA learns it and is non-causal vision. **(k)** machine-checked containment of the three corners and the zero-gate iff — bounded by `sweep_topology` 7 and 14 and `sweep_expressivity` S16; the nearest statements are FoX’s cumulative-log gate, which the record’s own theorem excludes at an exact zero, and the block-mask converse of [138] Thm 1. **(l)** the influence-Jacobian barcode restricted to the F0 endpoint, with a directed stability δ — bounded by `sweep_topology` 4/8/16. **(m)** the cover-to-schedule path, Mapper nerve to CSR blocks to a certified resolvent read — bounded by `sweep_topology` 6/15/2. **The totals.** Across the eight sweeps, **143** numbered queries were logged (resolvent 22, linrec 18, expressivity 18, causality 18, safety 25, topology 17, occupied 12, methods 13), plus five bibliographic-repair searches in the methods sweep; `references.bib` carries **408 canonical entries** (RUN, `grep -c "^@" references.bib` = 408, this session), of which the audit resolved **278/278** arXiv identifiers and **91/92** DOIs against the registries, with 38 entries carrying neither identifier; the five sweeps that print an entry-level [U] count report **13** in total (methods 9, occupied 2, topology 1, causality 1, expressivity 0), and the other three mark [U] per field rather than per entry, so no merged [U] count exists. Absence here is exactly as strong as those 143

queries and no stronger, and two of the sweeps paid for their absences with planted positives — `sweep_topology` §1.1 shows the same instrument returning the seed papers it was required to find, and `sweep_occupied` §4 plants the record’s own dead signed programme as a lineage the table must return red.

7.3 The record’s own prior-art defect

The paper that owns CEQ’s read operator has been in print since October 2024 and the campaign never cited it. The mechanism is a filter, and the filter is written down. At round 1 the record’s prior-art trawl concluded: “*Every multi-hop propagation found (APPNP 1810.05997, GDC 1911.05485, MAGNA 2009.14332) requires a non-negative matrix; every signed attention found ... is single-hop*” (READ DONE_ARCHIVE_ROUND1.md:676–680). The clause “requires a non-negative matrix” was a *disqualifier*: the programme then alive was the signed strictly-causal path sum, so a multi-hop operator built on a non-negative base was recorded as inapplicable and dropped. Every candidate in the class ChaCAL belongs to was therefore discarded by construction, and the same page later concedes that the survey sentence built on that filter was false (READ DONE_ARCHIVE_ROUND1.md:1046–1050). When the signed programme died (RESEARCH.md, PROGNOSIS.md, per BRIEF.md §3), the filter was never re-run without its sign criterion. The shape is now itself a non-negative multi-hop operator, which is to say the record’s search was, for the whole life of the campaign, structurally incapable of finding the paper that occupies its read.

The measurement of the gap is a grep. `grep -r1Ei "ChaCAL|Fagnou|2410\05565" --include='*.md'` over the repository returns **0 files** (RUN, 2026-09-03, this box; confirmed twice); a full-tree `grep -r1Ei "2410\05565"` returns exactly **one** file, `docs/references.bib` — the bibliography this paper assembles. 2603.00175 (InfSA, which put “fundamental matrix of an absorbing Markov chain” into an attention abstract in February 2026) appears in **2** .md files, both inside the dead signed programme’s archive, where it was read as a signed near-miss and its absorbing-chain paragraph was not read (READ DONE_ARCHIVE_ROUND1.md:1317–1318).

This is MISTAKES.md V-7 — *a search structurally incapable of finding anything, read as absence* (MISTAKES.md:117) — applied to a **literature** search rather than to a code search, which is the form the taxonomy had never recorded. V-7’s existing instances condemn grep patterns and test selectors that could not match; this instance condemns an inclusion criterion in a prior-art census. It compounds **P-3** (a stale claim never retracted, MISTAKES.md:316): THEORY.md:22 still calls the successor operator new, and THEORY.md:225–235 still marks Dayan 1993, Ramsauer 2020 and Bai et al. 2019 unverified, all three of which the sweeps have now fetched. **The rule the plan installs against it:** every prior-art census keeps *one mechanism-keyed query per component beside every name-keyed query*, and no census criterion may name the programme’s current hypothesis. ChaCAL was found by `sweep_expressivity` S6 and `sweep_resolvent` search 2 — queries built from *triangular solve* and *resolvent*, the shape’s mechanism — and not by any query naming an operator family; the name-keyed query `sweep_resolvent` 18 (“chain and causal attention ... citing papers”) returned only ChaCAL’s own pages after it had already been found. The census’s planted positive is its own refuted lineage: a table that cannot return a red on the signed programme is not calibrated.

One interval in the sweeps is not supported by anything in the tree. `sweep_occupied`.md §0.1 says ChaCAL was in print “eleven months before the brief”; the identifier is arXiv 2410.05565, submitted October 2024, twenty-three months before 2026-09-03, and the tree’s own history runs only 2026-08-25 to 2026-09-01 (RUN, `git log --format=%ad --date=short`), so no in-tree date supports either figure. The paper prints the identifier’s date and no interval.

8. Limits

Collected once, here, and nowhere else.

There is no result. No CEQ arm has been trained and no capability number exists; the paper reports a plan and the identities that justify it. Proposition 8 — the safest move — has no cell, no realised label standard deviation and no measured hop depth, and the barcode and Mapper rows have no instrument at all (`beta0_interleaving` consumes point clouds, `READ ceq/certs/topological.py:476-505`). The record’s last deciding measurement crossed a one-hop threshold on five seeds of eight, straddled it on the interval, and that threshold was afterwards shown not to be an information floor (13 of 40 cells violate it, `V20_R15_THEORY_TABLE.md §0.2`). Every number here is an identity check, a price, or a citation.

Theorems cited beyond their hypotheses license nothing, so the hypotheses are printed. The one-layer composition bound of [92] Thm 1 requires $n \log n > H(d+1)p$; at $s = 64$, $d = 16$, one head, float32 this reads $384 < 544$ (DERIVED, $\log_2$), so the inequality is **vacuous** at the record’s geometry and the theorem is cited for its asymptotic shape with those numbers beside it. The same discipline applies to [105] Cor. 4.3 and [103] Thm 3/19, both conditional on the one-vs-two-cycle conjecture, and to [25] Thm 1.1, whose parameter budget $n^{2^{-4L}}$ reads 1.30 at $L = 1$, $n = 64$. The Cheeger restatement is owed: the record writes $g \leq 2\varphi$ at `MATHEMATICS.md:455`, while Levin–Peres Thm 13.10 ([73], read at printed p. 183 by the methods sweep) states $\Phi_\star^2/2 \leq \gamma \leq 2\Phi_\star$ **for a reversible chain**; the shape’s P is causal and not reversible, so the sentence must be restated on a named symmetrised or lazy surrogate, or replaced by a non-reversible bottleneck bound, and until it is, no conductance sentence is written about P . A second hazard is carried, not resolved: [84] [V] warns that Bellman optimality can fail in multichain constrained MDPs with several unsafe sets — the class the safest-move rule inhabits — so the rule is stated as a one-step filter over an evaluated committor vector and no card claims policy-level optimality.

Determinism and cost carry their own caveats. `solve_triangular` was observed bitwise identical forward and backward over eight repeats under `use_deterministic_algorithms(True)`, but torch documents no such guarantee for it, so that is an observation on one box and one version, not a contract; `cumsum` raises under the same setting, which is why the segmentation route is a mask and not a scan. The timings inherit an unsynchronised timer and a run-order confound — the arms were not interleaved — so every price is a **per-op floor** under a measured dispatch gap of $2.0\times$ to $6.6\times$ (`READ scale/m3_flops.py:101-121`). The measured increment of the solve over the softmax head it contains is $2.514 - 1.473 = 1.041$ ms/step at $n = 2048$ (2.312 at 4096, 4.542 at 8192), float32, one box, producer owed to `scripts/k_cert.py`; the chunked and CSR paths are NOT MEASURED - needs a chunked kernel, and the dense control is not runnable at $n = 2048$, $s = 4096$ (≈ 137 GB against 7.996 GiB), so n is stated per s everywhere.

The bibliography is verified at the identifier, not at the equation. The audit resolved 278 of 278 arXiv identifiers and 91 of 92 DOIs with no identifier paired to a wrong title; the unresolved one is [113], 10.2312/SPBG/SPBG07/091-100, absent from Crossref (Eurographics digital library), carried [U] with two independent records agreeing on title, authors and year. Thirty-eight entries carry neither identifier — books, nineteenth- and twentieth-century papers, unregistered proceedings — and rest on catalogue, ISBN, zbMATH or DBLP marks; nine of those were never fetched at all (Cantelli 1928, Bonferroni 1936, Frobenius 1912, Neumann 1877, Kantorovich–Rubinstein 1958, Grünwald 1867, Letnikov 1868, Kramers 1927, Littlewood–Offord 1943) and carry no load-bearing sentence, a textbook statement standing in where one is needed. **The sweeps’ [V] marks are abs-page title matches, not equation-level [V=eq] marks.** Only a handful of sources were read at the body — ChaCAL, InfSA, DeltaNet, SSD, Hu et al., FoX, Dash 2005, Bottou §7.3, Grinstead–Snell §11.2, Levin–Peres Thm 13.10 — and every other equation number, theorem number and hypothesis here was read through a rendered page or a

summarising fetch and **must be re-checked against the compiled PDF before typesetting**, ChaCAL’s diagonal convention included, since it rests on one HTML fetch and decides two of the paper’s controls. That is not a formality; it is the record’s own hazard. Extracting the LRU paper’s §3.3 with `pypdf` returned text whose Type-1 encoding mapped ∞ to the glyph 1, turning an open interval into a closed one — *“the single sentence that decides that node is the one the extractor corrupts, and it corrupts it into the claim under test”* (`READ_workdonenewseal.md:468–473`; `V15_X36_PRIOR_ART.md:45–56`). The remedy there was to read every ML source from arXiv LaTeX source, and it is the remedy owed here. One internal inconsistency, recorded: `sweep_expressivity`’s header says sixteen searches where its own table and count section say eighteen; eighteen is used.

The provenance of this document. It was produced by two multi-agent passes, both cut by a session cap, and the reader is owed what was lost. In the **first** pass two of three design documents reached the first synthesis and no refutations reached it at all: `design_instrument.md` (525 lines) and `design_falsify.md` (fourteen bets with counters and killers) survive as inputs, `design_theory.md` was never written (`THESIS_CORRECTIONS_2.md` §2, `READ`; `PAPER_OUTLINE.md` §8), and the theory design was produced fresh in the second pass. In the **second** pass the plan synthesis was the casualty and was re-run separately. So this section was judged against six refutations, all present, but the plan it points at was assembled on a second attempt rather than in one pass, and no single agent saw every input at once. The judge’s `RUN[J]` numbers are each one numpy float64 draw at $s = 32$, $\gamma = 0.6$, seed 0, CPU — identity and counterexample checks, not statistics — and other planets’ runs are quoted at their own differing geometries. No planet wrote a code file or made a git write, so nothing here is reproduced by running this document; the reproduction path is the plan’s first-evening list, and it starts at zero GPU-seconds.

Appendix A. The two ledgers

The proposition ledger and the bind ledger are the paper’s audit trail. Every proposition of §2 appears in the first with its status, its evidence class and number, the failure mechanism it is designed against, the verdict each of the six refuters returned on it, and the ruling that settled it; anything left `OPEN` names the §5 card that discharges it. Every bind, control, census, certificate and kill of §6 appears in the second with the same fields and the card that builds it. Neither is a summary of the sections: they are the record of what survived adversarial review and what was deleted, and they are the fastest way to check that a sentence in the body is carried by something.

A.1 Proposition ledger — the shape, judged

JUPITER (MYCROFT), 2026-09-03, HEAD 207e7b9. One row per proposition, definition or remark of `judge/sec_shape.md` and `judge/sec_obstructions.md`. Refuter columns: T-occ / T-math (`refute_theory_*`), I-occ / I-math (`refute_instrument_*`), F-occ / F-math (`refute_falsify_*`); `-` means the refuter did not address the item. Status $\in \{\text{LEAN, CLASSICAL, DERIVED, RUN, OPEN}\}$. `RUN[J]` = this judge’s one-liner (seed 0, $s = 32$, $\gamma = 0.6$). `PLAN` items name what discharges an `OPEN`; their full specification is the plan’s, not this ledger’s.

evidence											PLAN
id	statement (short)	status	num- ber	mechanism	T-	T-	I-	F-	F-	ruling	item (OPEN only)
					iso	math	I-occ	math	occ		
D1	base fam- ily with the mul- ti- plica- tive gate; Hop is its $m > 0$ re- stric- tion	LEAN	V16Domain. READ; 25, cor- ners :433, :445	Min.lean P-7 (cor- ners 1–2 own- ers; base)	REPAIR, REPAIR (base has no zero)	REPAIR, REPAIR :366–378	—	—	—	both re- pairs ap- plied: base rede- fined with pathProd; vaswani, katharopoulos, dao cited	—
D2	two regimes; no ab- sorb- ing row in regime N; $\gamma = 1$ sin- gular with one	LEAN + RUN	one_not_nilpotent :105; 25, RUN[MC8 det($I - A'$) = 0.0	KEEP (add the regime- N sen- tence)	REPAIR- (rows 13, 24, 26 con- fla- tions)	REPAIR- (rows 13, 24, 26 con- fla- tions)	—	KILL of “regime N sur- vival for (e)”	adopted: regime N can- not carry bound- ary rows; design_falsify §5 rank- 2 sur- vival route deleted for (e)	—	

id	statement (short)	status	evidence + num- ber	T- mechan- ism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
D3	boundary rows; F1 BOS ab- sorb- ing; F2 sets pre- cede the query; BOS in its own sink set	RUN	RUN[J] V- det($I-25$, Q) = V-8, 0.0 V-12, un- D-3 de- clared; F2 $q =$ 0.0; RUN[I] R6/R7/R10	REPAIR (sink vs row "same ob- ject")	KEEP + 2 addi- tions	KEEP	KEEP	REPAIR (adopt F1/F2)	REPAIR (§3.11: goal col- lapses to "de- scend to 0")	Ruled: — sepa- rate sink set with value 0 (T- math C8(b), F- math §3.11) over design_theory's BOS- in- goal; query in T printed	—
D4	read car- ries (1 — γ); bare row sum $1/(1-\gamma)$ is a class prop- erty	RUN	RUN[J],V- RUN[M] 7, 2.5 V-23 with/without sets	REPAIR (word- ing)	KEEP (word- ing)	—	REPAIR (R12 min- entry)	—	REPAIR (choose once)	wording— adopted; one ob- ject cho- sen ($\Pi_\gamma V$)	—
D5	horizon dial; own- ers named	CITED	bellmarP-7 roffo, kemeny	KEEP	—	—	—	REPAIR (own- ers)	—	owners — ap- pended	—

id	statement (short)	status	evidence + num- ber	T- mechan- ism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
D6	two inter- ven- tions: row clamp (ora- cle, rank 1) vs to- ken rewrite (arm, rank $s-i$)	RUN	RUN[J] rank 20 at $i =$ 12; RUN[M] 20 of 32; RUN[F] 6 of 8	M-8, P-8 —	KILL (§1.6 sup- port)	—	—	—	REPAIR (scope)	KILL an- swered by split- ting the defi- ni- tion; both prices stated	—
D7	two rules as columns; two heads; com- mit- tor head at $\gamma =$ 1 ex- act solve; dis- counted label sepa- rate at γ_{env}	DERIVE + RUN	RUN[F] $\hat{\gamma} \geq 17$, 0.99857 for 1 % at $\tau =$ 8	V- 17, D-2, D-3, M- 20 —	—	KILL (FATAL 2)	KILL (row 16)	—	KILL (item 16)	KILL an- swered: heads split (I- occ op- tion (a) for H-q; F- math op- tion (a) for the dial bet)	—

id	statement (short)	status	evidence +							ruling	PLAN item (OPEN only)
			num- ber	T- mechanism	T- math	I-occ	I- math	F- occ	F- math		
P1	parity at $\gamma =$ 0 bit- wise; bat- tery re- paired	RUN + LEAN	RUN[covered] True; 24, 2.3002850340264393 at $\gamma =$ 0.5	REPAIR (delete causal plant)	KEEP	—	REPAIR (bit- wise con- ven- tion)	—	—	plant deleted; “Cha- CAL same γ ” filed as declared- empty; a re- pro- duc- tion of Cha- CAL Eq. 5	—
P2	corner 3 is the sub- diagonal resol- vent; BED- M con- tained	LEAN	RUN[covered] 0.0; 6.217248937900877 $\times$ 10^{-15}	KEEP	KEEP	—	—	—	—	kept ver- ba- tim	—

id	statement (short)	status	evidence +			T- mechanism	T- math	I- I-occ	F- math	F- occ	F- math	ruling	PLAN item (OPEN only)
			num- ber										
P3	commit = resol- vent read with ab- sorb- ing rows; $E[\gamma^{\tau-1}]\ll_k$; sum- to- one when sets ex- haust	CLASSICAL RUN	1 real sets 0.0; resid- ual 1.04×10^{-17} ; $\rho(Q) = 0.9408612510154677$	D-2, V-12, V-25	REPAIR (sta- tus to- ken; oracle out- side hy- poth- esis read plant)	KEEP iden- ti- ties; RE- P3 3.8×10^{-7}	—	REPAIR (read is $E[\gamma^{\tau-1}]$, not $E[\gamma^{\tau}]$)	—	KEEP iden- tity / RE- PAIR state- ment ($i\in\mathcal{A}_k$ reads 1)	status — CLAS- SI- CAL; read form fixed; T- math's $\geq$ $1-\gamma$ bound over- ruled — RUN[J] max gap 1.42×10^{-7} ; cor- rect bound ($1-\gamma$) q_i ; new tar- get lowerTriangular_ne_symm		

id	statement (short)	status	evidence +			I-occ	I-math	F-occ	F-math	ruling	PLAN item (OPEN only)
			num- ber	T- mechan-	T- math						
P4(i)	Neumann tail equal- ity on the class; “> 1 can fail”	DERIVED + RUN	EDN[cover-3] equal- ity to 10^{-15} ; plant $\gamma =$ 0.6: 7.29 vs 0.54	V-3, V-10, V-24 Meyer; AS- SUME tag; $\hat{\beta}$ cen- sus)	REPAIR KEEP equal- ity; RE- PAIR clause	—	KEEP (row 6)	KILL as BED- S bind (\$3.4)	KEEP equal- ity / RE- PAIR ob- ject	textbook- cited; sta- tus DE- RIVED; the $\gamma =$ 0.7 plant la- belled di- ver- gent; row- sum cen- sus added; the bind is dor- mant until a mask ships	—
P4(ii)	mask am- plifi- ca- tion $\varepsilon\ V\ _\infty/(1-\gamma)$	DERIVED + RUN	EDN[J] L- 0.5263157894736843 vs naive 0.1	V-10	—	—	KILL of the naive union bound (row 8)	—	—	KILL adopted as a new propo- si- tion; ev- ery F1 line car- ries $1/(1-\gamma)$	—

id	statement (short)	status	evidence +							ruling	PLAN item (OPEN only)
			num- ber	T- mechanism	T- math	I-occ	I- math	F- occ	F- math		
P6(i)	zero gate $\Rightarrow$ ex- act block- diagonal resol- vent; F0 by mask in regime S; ab- sorb- ing rows re- spect cuts	LEAN + DE- RIVED	pathPrda_eq_z :129; RUN[MY-25 array_equal ze- ros; 10 ⁻³⁰⁰ plant	— CERT, (base; 31.06 $\times$ flag)	— REPAIR (base)	—	KEEP (row 15); RE- PAIR (10 ⁻³⁰⁰ un- der- flow, row 14)	REPAIR (\$3.6 [M] on [U] names) S)	REPAIR (item 6, iff not on regime S)	all ap- plied: iff marked [S] on regime S; -30 plant; [M] pend- ing build	—
P6(ii)	a zero gate is a bound- ary con- di- tion: $P_{cc} =$ 1, $\rho(Q) =$ 1 un- de- clared; declar- ing sev- ers sets be- fore the cut	RUN	RUN[J] V- $P_{cc} =$ 25, 1.0, V-8, $\rho(Q) =$ D-3 1.0; RUN[M] $q =$ 0.0	—	KILL as com- posed (item 11)	—	—	—	—	KILL	— an- swered with the num- ber: new propo- si- tion; seg- ments carry their own sets; cut posi- tion a dial

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I- I-occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
P6(iii)	dividend in pair- count units, $31.04 \times$ on BED- M; $1 \times$ on the soft- max cor- ner	DERIVED	ED6_ARM_SMP frac- tion 0.0322	17, V-22	REPAIR (flag) (item 26 units)	REPAIR —	—	REPAIR ($\$3.11$ fall- back reads $1 \times$)	—	units — cor- rected; cor- pus prop- erty; no bet on it	—
P7(i)	$\Delta z =$ ($I -$ $\gamma P'$) ⁻¹ ($\Delta P z$)	DERIVED + RUN	EDUN $1.2 \times$ 10^{-15} ; RUN[F] $2.0 \times$ 10^{-16}	V- 24, D-5	REPAIR (cite Bot- tou, Vakalis)	KEEP (own- ers)	KEEP (own- ers)	KEEP (own- ers)	KEEP (own- ers)	owners cited; kept	—
P7(ii)	suffix re- solve = full re- solve by tri- an- gu- lar- ity; forward- only; $\Delta z_0 =$ 0	DERIVED + RUN	EDUN $5.6 \times$ 10^{-17} on the non- nilpotent cor- ner; RUN[I] equal True on solve_triangular	V-3 (de- clared), V-24	KEEP (la- bel V-3)	KEEP (hy- poth- esis re- stated)	REPAIR (Var(Δz_0)) 0 cen- sus)	REPAIR (note (con- stant row co- ords masked)	REPAIR (“by nilpo- tency” wrong) masked)	REPAIR hypothesis re- stated; route named; cen- sus over co- ords $\geq$ $i_{\min}$; “by tri- an- gu- lar- ity”	—

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I-occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
P7(iii)	Sherman Morrison at a row clamp, de- nom- ina- tor > 0; col- umn costs $s^2/2$; to- ken do() is a suf- fix re- solve	DERIVED RUN	RUN[I] M-8, 1.45 × P-8 10 ⁻¹⁵ , de- nom- ina- tor 1.9; RUN[F] 0.511918; ratio $m/d =$ 0.5	—	KILL for to- ken do() (item 19)	—	KEEP (row 9)	REPAIR (price band 0.5 × vs 8 ×)	KEEP alge- bra / RE- PAIR scope	KILL an- swered by D6's split; both prices; band (0.5 ×, 8 ×) filed SPLIT	—
P7(iv)	$V \equiv$ $\not\models \Rightarrow$ $\Delta z \leq$ 10 ⁻¹⁵ ; Gaus- sian V is the re- gion	RUN	RUN[F]V- 1.1 × 24, 10 ⁻¹⁶ ; D-5 0.1096 / 1.127	—	—	—	KEEP with toler- ance (row 30)	REPAIR (one- liners recorded)	REPAIR ("ex- actly 0.0")	exactly — 0.0" re- placed by $\leq$ 10 ⁻¹⁵	—
P7(v)	value- only moves leave q con- stant; $m(K +$ 1) RHS route is for Δz only	DERIVED	RUN con- tains no V	V-8	—	KILL (item 17)	—	—	—	KILL adopted; clause re- stricted	—

id	statement (short)	status	evidence +	T- mechanism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
			num- ber								
P7 met- ric	position matched NRMSE, co- sine, resid- ual with unit fac- tors; Mc- Ne- mar on sign only	DERIVED	permutated- ora- cle $W_1 =$ 0.0 at NRMSE 1.421901 (V20_R15_THEORY_TABLE.md:221)	—	—	KEEP (G2 plants)	—	REPAIR (\$3.12 units)	REPAIR (item 18 Mc- Ne- mar)	Both ap- plied	—
R- EMC	EMC re- tired to a re- mark; feed- back plant re- named a tri- an- gu- lar- ity bind	RUN	RUN[J] V- 8.9×10^{-16} on cyclic P ; RUN[I] 1.33×10^{-15} ; RUN[F] 4.4×10^{-16}	KEEP (la- bel V-3)	KEEP	—	KILL (row 11)	KILL (\$3.13)	KILL (item 7)	KILL stands (3 vs 2; the “re- gion” was a wrong algo- rithm); re- mark with dash, momennejad, russek	—

											evidence	PLAN
											+	item
id	statement	status	num-ber	T-mechanism	T-math	I-occ	I-math	F-occ	F-math	ruling	(OPEN only)	
P8 lemma	degeneracy $\max_k q^{(k)} \geq 1/K$ when con- straints ex- haust ab- sorp- tion; plant = BOS in a con- straint set	DERIVED	RUN[I] V-0.605 12, was V-3 with the goal counted; constraints- only min max = 0.0	—	KEEP	KEEP	—	KILL of the “drop $\mathcal{A}_0$ ” plant (row 3)	—	lemma — kept with ex- act hy- poth- esis; plant re- placed	—	
P8 floor	zero- information floor is chance 1 — 1/ m ; weak Fano struck; tight Fano where $I >$ 0; $I(X_{\leq 0}; a^*) >$ 0 com- puted	RUN	RUN[J] V-0.75/0.875/0.9375 vs V-17, 0.5/0.667/0.75; tight FLOOR 0.7124 at $m =$ 8	—	REPAIR	REPAIR	REPAIR	REPAIR	KILL	adopted: chance is the floor; $m =$ 2 not re- fused	—	

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
P8 cen- sus	label sd over the ad- mit- ted re- gion; class fre- quen- cies in (0.05, 0.95); dis- agree- ment > 0; $q^{(\text{sink})}$ apart	DERIVED	R9 sd in- flated by the con- stant pre- fix	V-8, V-12	REPAIR- (pre- dicted fail- ure mode)	REPAIR- (in- verted gate; R9 re- gion)		REPAIR- (F1/F2 in cen- sus)	REPAIR- (§3.11 cen- sus)	all ap- plied	—

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
P8 bed	environmental = random DAG in to- ken or- der with self- loops; adja- cency multi- hot; VOID list; C2(a) dropped, C2(c) kept	OPEN	F-B DE- RIVED (Symm vs causal sup- port); RUN[coord] BED- 1 $\rho(Q) =$ 0.9409	D-2, V-25, V-10, Support horns)	—	KILL (FATAL 1)	REPAIR (row 26: which chain)	REPAIR (Bet B)	REPAIR (item 17)	KILL an- swered by I- occ's DAG + T- occ's VOID list; adopted ver- ba- tim	S-1 write bed_s spec (DAG in to- ken or- der, sink/goal/constraint be- fore the query, adja- cency fea- tures)); S-2 cen- sus run on 512 draws (sd, class fre- quen- cies, dis- agree- ment, dis- cards, $q^{(\text{sink})}$), PASS iff sd > 0.05 and ev- ery class in (0.05, 0.95), 0 GPU- s; S-3 VOID list and

											PLAN
											item
											(OPEN
											only)
id	statement	status	evidence	+	T-	T-	I-	F-	F-		
	(short)		num-	mechanism	isoc	math	I-occ	math	occ	ruling	
P8	E1	DERIVED	MDE	M-	REPAIR-	REPAIR-	REPAIR-	REPAIR-	—	MDE-	
kill	kill	+	0.0398213,	(§5	(TOST (row	based				—	
	by	RUN	at sd	M-9,	kill	tier	20			kill;	
	$n =$		0.03445M-7,	un-	has	paired				TOST	
	8		re- V-17	reach-	no	N)				branch	
	MDE		pro-	able)	branch)					stated;	
	of		duced;							paired	
	the		power							N	
	re-		0.7975 $\rightarrow$							NOT	
	alised		0.8073							MEA-	
	sd;		at							SURED	
	TOST		69 $\rightarrow$							until	
	($N =$		70							sd_d	
	70										
	two-										
	sample;										
	paired										
	$N \approx$										
	36 if										
	$sd_d =$										
	σ ,										
	NOT										
	MEA-										
	SURED)										
	de-										
	cides										
	the										
	par-										
	ity										
	half										
	only										

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I- I-occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
P9	horizon dial corol- lary; LR null at the bound- ary 2.7055; $ \hat{\gamma} <$ 0.05 rule deleted; $1/(1-\hat{\gamma})$ printed	DERIVED + RUN	BDN[J] V-9, $\chi_1^2(0.90)M=2$, 2.7055; V-17 V17K_RULINGS.md:389-436	KEEP	KEEP	REPAIR (J: delete 0.05)	REPAIR (row 21 bound- ary null)	REPAIR (null de- clared; seed count)	—	applied; boundary null cita- tion owed (not in bib)	J-1 R in- stru- ment spec with the bound- ary null, held- out set, seed count “MOVED on $\geq$ 6/8”; SPLIT band named
P10	$\rho(Q) =$ $\max_T P_{\#};$ $I-Q$ in- vert- ible iff $0 \in$ $\mathcal{A};$ com- mit- tor head by ex- act solve, no Neu- mann cer- tifi- cate on it	DERIVED + RUN	BDN[J] V- 0.6926596893360386 both; V-25, RUN[I] L- $\ Q\ _{\infty}$ CERT bound 19.2 vs error $2.5 \times$ 10^{-3}	—	KEEP (ad- di- tion 1)	—	REPAIR (row 2), KILL (row 16)	—	REPAIR (item 3)	adopted; H1/H2 cer- tifi- cate cell reads “ex- act solve, $\delta =$ 0 to round- ing”, not a Neu- mann δ	—

id	statement (short)	status	evidence +		T- mechanism	T- math	I- occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
			num- ber	mechanism								
P11	regime bound- ary; V-4 evi- dence fixed; masked route ex- act iff $V_0 = 0$	LEAN + RUN	RUN[M] 6.27×10^{-13} vs 7.96×10^{-8} ; 1.99×10^{-7} masked	V- 25, P-3, V-4	REPAIR (C8 evi- dence)	KILL (masked sen- tence as ma- trix iden- tity)	—	—	—	—	evidence— re- paired; masked route re- stated as a value- channel iden- tity; de- fault keeps the diag- onal	—
O1	per- row inde- pen- dence — def- ini- tional; resid- ual as the mea- sur- able	DERIVED	—	V- 26, D-7	KEEP	—	—	—	REPAIR (units)	—	kept with the unit fac- tors	—

id	statement (short)	status	evidence +	T- mechanism	T- math	I- I-occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
			num- ber								
O2	depth law — sky- line un- con- di- tional, lower bounds con- di- tional; vacu- ity printed; re- duc- tion NOT FOUND; P- complete clause dropped; $O(\log^2 s)$	CITED DE- RIVED 64 ^{1/16} 1.2968; 512 $\geq$ 384 < 544	$\Rightarrow$ P-10, V-25, P-8	KEEP (§3.1– 3.4) / RE- PAIR (sky- line ta- ble)	REPAIR (items 13, 14)	REPAIR (Thm 4.2 by anal- ogy; width/CoT fixed)	REPAIR (row 24)	REPAIR (Thm 4.2)	REPAIR (item 12)	all ap- plied	—
O3	composi- — Peng Thm 1 vacu- ous at $s =$ 64; Kozachin- skiy asymptotic	CITED 384 < 544	V- 25, P-10	KEEP	KEEP	—	—	—	—	kept	—

id	statement (short)	status	evidence	T- mechanism	T- math	I-occ	I- math	F- occ	F- math	ruling	PLAN
			num- ber								item (OPEN only)
O4	non-negativity with-drawn; redi-rec-tion; sup-port un-changed	RUN	RUN[F]P-7, 151 zeros from iden-tity rows; RUN[J] sup-port equal	KEEP V-23, V-3	—	—	REPAIR (row 5)	REPAIR (\$3.5)	KILL (item 13)	wording— adopted: “ear-lier posi-tions”, weights not sup-port	—
O5	D-1 with soft-max proper; label model ad-mits 0 of 3 beds	CITED + READ	MATHEMATICS. $L/d =$ 4.00 P-10, V-25	REPAIR (P7(d))	REPAIR 1057	—	—	REPAIR (rank 4 past SLR)	—	applied —	—

id	statement (short)	status	evidence + num- ber	T- mechanism	T- math	I- I-occ	I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
O6	skyline table with three fel- lows, ChaCAL- diag, ChaCAL- published, sink, In- fSA, cached- mixture, MuZero	CITED	—	D-1, R- SKY, D-7	REPAIR- (one ta- ble)	REPAIR- (sky- lines on G1/G2/G12, 4,769 re- count)	REPAIR- (row 13 diag- H2, con- ven- tion)	REPAIR- (miss- ing Bet D)	REPAIR- (item 19 two Cha- CAL arms)	table writ- ten; the sky- line bet dic- tion of sec_beds.md §6.E) is filed with its counter- side in ei- ther di- rec- tion”, ≈ 7.6 s per cell [AS- SUMED]	K-1 sky- line bet filed: pre- dic- tion “depth- within the MDE”, counter “out- side

id	statement (short)	status	evidence +					I- math	F- occ	F- math	ruling	PLAN item (OPEN only)
			num- ber	mechan- ism	T- math	T- math	I- occ					
C- cost	cost law; +50% causal MACs; ≈ 34 s pair; 7 arms ≈ 98 s; pro- ducer owed	FITTE + RUN	Dec_cost §4.x.8; F- math §3.16 arith- metic re- derived	M-8, P-8, P-1	REPAIR (pro- ducer owed)	—	KEEP (\$6)	REPAIR (rows 28, 29)	KEEP	KEEP	applied (98 s; MACs vs FLOPs)	C-1 add determinism_at_6 fourth quan- tity and the solve mi- crobench- mark to scripts/k_cert.p (the one code edit owed, not done here)

id	statement (short)	status	evidence	T- mechanism	T- math	I-occ	I- math	F- occ	F- math	ruling	PLAN
			+ num- ber								item (OPEN only)
C- topo	barcode row killed on BED- S; F0 end- point theo- rem only; Map- per with 1/(1- γ) union bound and n per s	OPEN	RUN[F]V-3, $\beta_0(\varepsilon =$ 0) = 1 on ev- ery soft- max draw; 137 GB at $n =$ 2048, $s =$ 4096	— V-25, M- 15, P-4, L- CERT	—	KILL (L), RE- PAIR (M)	REPAIR KILL (row 22), RE- PAIR (row 8)	—	KILLs adopted the layer is re- stricted to what has a theo- rem	T-1 di- graph β_0 in- stru- ment spec (con- nec- tiv- ity no- tion, dis- simi- lar- ity trans- form, Turner pred- i- cate) be- fore any bar- code is read; T-2 Mapper- to- tile quan- tiser spec; both NOT MEA- SURED	

id	statement (short)	status	evidence	T- mechanism	T- math	I- I-occ	I- math	F- occ	F- math	ruling	PLAN
			+ num- ber								(OPEN item only)
L- Lean	20 tar- gets; [M] pend- ing build; four state- ments re- paired; three added	OPEN	sec_proved,md §2.8; re- futer §3/§5 ta- bles	Added,md LEAN, gets V-25 2, 4, 14, 18; add 9)	REPAIR REPAIR (4, 16, 6; add 7, 19) add 9)	—	—	REPAIR (§3.9 [M] on [U])	—	all ap- plied	L-1 first evening: tar- gets 1–10, lake build, #print axioms; a sorry de- motes the row

Tally. Propositions and definitions judged: 40 rows. Kept unchanged: 6 (P2, P5(b), P7(i) algebra, O1, O3, C-cost arithmetic). Repaired in the refuters’ wording: 27. Killed as written and answered with a number (a new proposition or a split definition replaces them): 5 — D6 support, P4(ii) naive union bound, P5(c) support claim, P6(ii) composed segmentation, P7(v) value-only moves. Killed and deleted: 2 — the EMC proposition (retired to a remark) and the β_0 barcode row on BED-S. Refuter overruled with a RUN: 1 — `refute_theory_math` item 8’s universal $\geq 1 - \gamma$ gap (RUN[J] 1.42×10^{-7} ; the bound is $(1 - \gamma)q_i$).

Open items and the plan nodes that discharge them. S-1..S-4 (BED-S specification, census, VOID list, first pair), J-1 (LR instrument with the boundary null), K-1 (skyline bet), C-1 (the one code edit owed), T-1/T-2 (topological instruments), L-1 (the Lean evening). Independent nodes: L-1, S-1, J-1, T-1, T-2 may be picked up in any order; S-2 needs S-1; S-3 needs S-1; S-4 needs S-2 and S-3 and L-1 (L-LEAN); K-1 needs S-4’s realised sd; C-1 needs nothing.

Citations owed (not in references.bib, therefore not cited). The boundary likelihood-ratio null; the reach-avoid LP papers the safety sweep did not fetch; the 2026 anytime-valid equivalence-testing preprints. The ChaCAL diagonal convention is [V-fetched] by one HTML read and is re-read against the PDF before it is typeset.

A.2 Bind ledger — binds, controls, censuses, certificates, kills

*SATURN (WATSON), 2026-09-03, HEAD 207e7b9. Companion to sec_apparatus.md. Status tokens: **KEEP** (adopted as designed), **REPAIR** (adopted with the refuter’s correction applied; the correction is in the “what changed” column), **KILL→applied** (the refuter’s reversal was one sentence and is applied), **KILL→deleted** (reversal needs an instrument or bed that does not exist; the row is out of the apparatus and appears only here). Every row names the **MISTAKES.md** mechanism by symbol and the **PLAN** item that builds it. **PLAN** item ids are proposed for docs/PLAN.md: P0.* Lean and manifests (0 GPU-s), P1.* bed specification and census (0 GPU-s), P2.* identity binds on BED-S’s real draw (seconds), P3.* the eight-seed arena (≈ 1.6 GPU-min),*

P4. MDE / TOST / LR verdicts, P5.* deferred instruments (NOT MEASURED). Items marked first evening are the five the author can settle alone in one sitting (THESIS_CORRECTIONS_2.md §0).*

1. Binds

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-J	parity at γ = 0, <code>torch.equal(0(0),</code> PV) against the lane's <code>softmaxAttn</code>	REPAIR	the “non-causal W” plant passes bitwise (RUN[MARS] True, <code>refute_theory_occvac</code> §2.1) — moved to B-P5; ChaCAL- published's <code>A_s</code> added as a plant (<code>refute_falsify_math</code> §0.4)	$\gamma = 0.5 \Rightarrow 2.3003;$ $\beta = 0 \Rightarrow$ $> 0.5;$ <code>A_s</code> $\Rightarrow O(1)$ at any $\gamma > 0$	V-24, V-3	P2.1 (first evening: the ChaCAL smoke test at $\gamma = 0.9$ on BED-M runs this bind's plant)

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-E1	boundary rows vs ChaCAL- diag, identity half at $\mathcal{A} =$ $\emptyset$, rejection half on down- stream rows	REPAIR	identity half is against the lane's own re- implementation with diag_convention declared (V-3 declared); rejection half is a weight gap, not a support gap (RUN[MARS] rows changed = [3]); plant (iii) “non- identity absorbing row breaks $\Sigma_k q =$ 1” struck — empty (RUN[MARS] 0.0, refute_instrument_math row 12) and replaced by the full-P read at $\gamma <$ 1 (0.1491)	$\mathcal{A} \neq \emptyset \Rightarrow$ $\ \Pi_{\text{shape}}$ — $\Pi_{\text{ChaCAL-diag}}\ _{\infty}$ $= 0(1)$ down- stream; in a base breaks the identity half (NOT MEASURED — needs the base wired)	V-24, V-3, V-14	P2.2

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-E2	conservation $q^{\{0\}} + \Sigma_k$ $q^{\{k\}} + q^{\{sink\}}$ $= \not\models$ on T	REPAIR	sink set $\mathcal{A}_{sink} = \{0\}$ added (RUN sum $1e-15$, sink share $[0.362, 1.000]$); “drop $\mathcal{A}_0$ $\Rightarrow \max_k \geq 1/K$ ” plant struck (V-3 of $\Sigma = 1$, refute_instrument_math row 3) and carried as the degeneracy proposi- tion (no goal, no sink $\Rightarrow$ $\max_k \geq 1/K$)	drop BOS from every set $\Rightarrow$ the solve raises $(\rho(Q) = 1.000000,$ det = 0, V-16); a set after the query $\Rightarrow q$ = 0.0 exactly $\Rightarrow$ discard	V-12, V-23, V-16	P1.2, P2.3

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-G1	interventional REPAIR re-solve: V $\equiv \not\models \Rightarrow$ $\Delta z \equiv 0$; Sherman– Morrison vs re-solve; C6 identity; forward- only		“bitwise $0.0 \rightarrow \leq$ $1e-15$ (refute_false_of_fmath §3.7); the solve route is named (LU gives $8.9e-16$, equal False; solve_triangular $0.0, True$ — refute_instrumental row 10); the EMC feedback plant struck and replaced by the trian- gularity plant	Gaussian V $\Rightarrow$ $\max \Delta z =$ (0.1096 / $0.363 /$ 1.127); non-causal $P \Rightarrow$ $\Delta z[:i] \neq$ 0 (0.0761 0.0868); — two-row rank-one by O(1), repaired by Woodbury	V-24, D-5, D-7	P2.4
B-G2	vector- metric plants (permuted oracle, $+0.1\sigma$, $-\Delta z$)	KEEP	—	permuted oracle scores $0 \Rightarrow$ metric struck; $+0.1\sigma$ not preferred $\Rightarrow$ struck; $-\Delta z$ indistin- guishable $\Rightarrow$ struck	V-26, L-14	P2.5

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-H1	committor identity q $= (I - Q)^{-1} R$ vs <code>bed_1.committor</code> on real sets (0.0)	REPAIR	the identity is checked on the environ- ment chain (BED-1 as instrument) and separately on the arm's triangular Q (invertible by <code>lower_triangular_isUnit</code>); the script raises on a missing key (<code>sec_refuted.md</code> C12); a wrong-set plant added (<code>refute_theory_occvac</code> §2.1)	must-fire perturba- tion $1.04e-17$ $\rightarrow 1e-6$ (ratio $9.6e10$); wrong $\mathcal{A}_k$ $\Rightarrow q$ moves O(1); Kirchhoff < $1e-10$ at K $= 2$	D-2, V-3, V-16	P2.6
B-H2	argmin from the oracle tensor equals the label's argmin (V-3 declared); leak plants	KEEP	—	planted leak token $\Rightarrow$ probe R^2 ≥ 0.99 ; E4' local- feature decoder above FAIL_BAR (re-derived for BED-S, V-22)	V-24, V-3	P1.4, P2.7

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-I	Neumann tail equality (defini- tional) + planted non- stochastic P + shipped- mask measure- ment in vector units	REPAIR, dormant	the $\gamma =$ 0.7 plant (119.37 / 1.143) is a divergent partial sum, labelled so; the $\gamma =$ 0.6 plant (7.29 / 0.54, convergent) is quoted; $\ V\ _\infty$ must be printed (the $\approx$ 7.6e-05 carried $\approx$ 5 [ASSUMED]); dormant until row M ships a mask (refute_instrument_occvac §2.6, D-4)	rows 1.5 fail the bound; on the shipped mask exceedance on any of 1,024 draws; $\delta \cdot \ V\ _\infty$ $\geq$ sd(label) $\Rightarrow$ uninfor- mative	V-3, V-10, V-17, L-CERT	P5.2

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-K	[M] Lean targets build; corner 3 vs (I − A) ^{-1} 0.0; seg- mentation zeros by torch.equal	REPAIR	[M] grades resting on [U] Mathlib names are [S] until lake build is green (P-11); the 1e-300 plant struck (underflow, refute_instrument_math row 14) and replaced by −30 logit; on BED-S at the softmax corner the F0 theorems are silent (0 of N draws)	keeping the diagonal breaks A ^s = 0; −30 logit leaves the block non-zero	P-11, V-25, V-2	P0.1 (first evening)
B-P5	solve_triangle vs dense inverse 1.8e-15	KEEP (plant received from B-J)	—	a dense M with upper=False must disagree	M-8	P2.1
B-M	CSR fill-in guard; do-nothing schedule always entered; empty-row NaN and negative- index loads structural	KEEP, unpriced	—	a dropped (k,1) with (k,m),(m,1) kept ⇒ exact solve refused unless δ printed	V-9, L-CERT	P5.3

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-C6	the dis- placement identity Δz $= (I - \gamma P')^{-1}(\Delta V + \gamma \Delta P z)$	KEEP	three inde- pendent RUNs $1.03e-15$ $/ 1.2e-15$ $/ 1.36e-15$	carried by B-G1's plants	V-24	P2.4
B-EMC	“settle then do = do then settle” with a feedback plant	KILL →deleted	edique fixed point for $\gamma < 1$ on every P (RUN[MARS] $1.33e-15$, $4.4e-16$ on dense cyclic P); the “violation” was a suffix-only solve on a non- triangular P; one remark survives	—	V-24, D-7, V-3	—
B-L	influence- Jacobian β_0 barcode, $\varepsilon = 0$ endpoint = segmenta- tion count	KILL →deleted	ed ($\varepsilon = 0$) $\in \{1, s\}$ on 100 % of softmax draws; beta0_interleaving consumes point clouds; Turner hypothesis unchecked	—	V-3, V-25, P-4	P5.4 (NOT MEASURED - needs a digraph β_0 instrument)

id	object	status	what changed (refuter, evidence)	planted negatives and required O(1) failure	mechanism	PLAN item
B-manifest	flipping any of beta/qk/g/gamma/boundary-coordinator_route, pattern by one unit moves manifest_hash; a missing declared field is a refusal	REPAIR	fields diag_convention, sink_set, (READ producer_cmdV20_R15_WING_MANIFEST.md:148-152) V_inf, void_contrasts added	the one- line-drift use, pattern	L-2, P-1, V-16	P0.2 (first evening)

2. Controls and skylines

id	arm	status	what changed	mechanism	PLAN item
C-SM1	depth-1 softmax, same head, matched params (count re-printed per arm)	KEEP	Ruling 3 count per arm (refute_instrument_occvac §2.5)	D-1, Ruling 3	P3.1
C-CHD	ChaCAL-diag (the shape with $\mathcal{A} = \emptyset$, diagonal kept)	REPAIR	split from “ChaCAL”	V-3, R-SKY	P3.1
C-CHP	ChaCAL- published (A_s diagonal removed inside the inverse, [V-fetched])	REPAIR	new arm; sub- stochastic; planted negative for B-J’s bitwise half	P-10, V-24	P3.1
C-CHS	ChaCAL- with-sink- token	KEEP	the “same γ ” is the shape’s trained γ , so the pair is sequential and the order is stated (M-2 in potential form)	V-24	P3.1

id	arm	status	what changed	mechanism	PLAN item
C-INF	InfSA-style Neumann $K = 16$, no boundaries	KEEP	—	D-1	P3.1
C-SKY-D	$\lfloor \log_2 t^* \rfloor + 2$ softmax stack (3/5/7)	REPAIR	depth chosen by analogy with Thm 4.2; the <code>hop_k</code> $\rightarrow$ <code>committor</code> reduction is NOT FOUND; carried by Bet D	P-10, V-25, D-7	P3.2
C-SKY-W	wide constant- depth stack, width = <code>n_nodes</code>	REPAIR	width fixed; at <code>d_model = 16</code> < <code>s = 64</code> the matched instance does not exist and the table says so	P-4	P5.5 (NOT MEASURED)
C-SKY-C	chain-of- thought decoder, <code>t*</code> steps	REPAIR	step count fixed	P-4	P5.5 (NOT MEASURED)
C-NAT	chunked-WY (DeltaNet / SSD), linear corner	KEEP	closed to the softmax corner by [58]	R-SKY, M-8	P5.6 (NOT MEASURED – needs a chunk-size sweep)
C-MZ	MuZero-style value head for the consequence channel	KEEP	—	D-1	P5.5
C-0H, C-1H, C-MAJ, C-RND, C-MEAN	argmin controls	KEEP	McNemar on identical draws; the C8 kill	C8, V-10	P3.1
C-CACHE	cached- mixture <code>0_cached =</code> <code>P_base(I –</code> <code>γP_{base}^{-1})V_int</code>	REPAIR	given an operational definition (<code>refute_instrument_occvac</code> §2.4); Momen- nejad/Russek mechanism [U]	P-7, D-2	P3.3

id	arm	status	what changed	mechanism	PLAN item
C-DT	return-conditioned DT-style arm on deterministic beds	KEEP	—	D-7	P5.5
C-CH-EMIT	“ChaCAL must emit chance”	KILL → deleted	a control that must fail by construction (V-2)	V-2, V-10	—
C-DENSE	dense resolvent at $s \in \{256, 1024, 4096\}$	REPAIR	n declared per s (137 GB at n = 2048, s = 4096)	P-8	P5.3
C-BEDK	hard_delay_attenuation on BED-K	KILL	—	V-3	P3.4

3. Census lines (L-DOM)

id	line	status	what changed	mechanism	PLAN item
X-1	row-stochastic P_{env} incl. absorbing rows	KEEP	—	V-25	P1.2 (first evening: the BED-S census on 512 draws)
X-2	$0 \in \mathcal{A}_{sink}$	REPAIR	sink set, not goal member	V-25, V-12	P1.2
X-3	$\rho(Q) = \max_{diag} < 1$	REPAIR	a diagonal read, V-10 gate, printed not counted (RUN 0.692660)	V-10	P1.2
X-4	boundary sets before the query; query in T	REPAIR	the query-in-T clause added (refute_theory_math §1.3)	V-8	P1.2
X-5	label sd over the admitted query region; sink share printed	REPAIR	region restriction (refute_instrument_occvac §2.7); sink share separate from goal	V-8	P1.2
X-6	argmin class frequencies in (0.05, 0.95); discards	REPAIR	the inverted “argmin-unique fraction” gate replaced	V-10, D-3	P1.2

id	line	status	what changed	mechanism	PLAN item
X-7	reduction-disagreement fraction > 0	KEEP	—	C4, V-1	P1.2
X-8	corpus-alone probe to q at order $0 <$ 0.5 ; planted leak ≥ 0.99	REPAIR	clause (a) (rows of P_{env}) dropped for BED-S — the graph is the input	V-24, V-7	P1.4
X-9	$I(s_0; a^*) =$ 0 ; untrained arm at chance	KEEP	—	V-10, D-5	P1.3
X-10	$Var(\Delta z) > 0$ over coordinates $\geq$ i_{min}	REPAIR	coordinate 0 and the prefix are zero on every draw	V-8	P1.2
X-11	move census (ΔP vs ΔV ; $u^T \not\models 0$; non-negative row)	KEEP	denominator positivity is a theorem, printed as a check	V-25	P1.2
X-12	β and $rowsum(P)$ on trained cells	REPAIR (new)	$\ P\ _{\infty} > 1$ $\Rightarrow$ no certificate printed	V-25	P3.1
X-13	γ support with LR verdict	KEEP	boundary null 2.706	V-17	P4.3
X-14	exact zero gate present (segmentation rows)	KEEP	0 of N on BED-S softmax corner $\Rightarrow F0$ theorems silent	V-25	P1.2
X-15	obstruction vacuity inequalities printed	KEEP	$512 \geq 64$, $384 < 544$, $64^{\{1/16\}} =$ 1.2968	V-25, P-10	P0.3

4. Certificates

id	object	status	what changed	mechanism	PLAN item
Z-EX	exact solve, δ $= 0$	KEEP	—	M-8	P2.1

id	object	status	what changed	mechanism	PLAN item
Z-NEU	Neumann K hops, $\delta = \gamma^{\{K+1\}}/(1-\gamma)$ bare / $\gamma^{\{K+1\}}$ for Π_γ , vector units, $\ \mathbf{V}\ _\infty$ printed, $1/(1-\gamma)$ beside	REPAIR	factor under Π_γ ; $\ \mathbf{V}\ _\infty$ printed; definitional label	V-17, V-3, L-CERT	P2.8
Z-F1	F1 mask: $\varepsilon/(1-\gamma) +$ δ_Π , vector units	KILL → applied	the resolvent amplifies dropped mass (RUN 0.5263 vs naive 0.1, $5.26\times$)	V-10, L-CERT	P5.2 (dormant until a mask ships)
Z-F0	segmentation, $\delta = 0$, dividend as a pair-count ratio, corpus property	REPAIR	$31.06\times$ is $s(s+1)/\Sigma$ $L_m(L_m+1)$, not s^2/Σ L_m^2 ; $1\times$ on the softmax corner	V-17, V-22	P0.1
Z-COM	committor at $\gamma = 1$: no Neumann certificate; exact route only; Perron weight NOT MEASURED	KILL → applied	H1/H2 inherit no δ ; the H2 threshold form certifies the solve, never the model	V-17, V-10, L-CERT	P5.7 (NOT MEASURED – needs the Perron weight w)
Z-COARSE	multizoom mean-pool bound for the far field	KEEP	—	L-CERT	P5.3
Z-BETA	$\ \mathbf{P}\ _\infty > 1$ $\Rightarrow$ no certificate	REPAIR (new)	—	V-25	P3.1

5. Kills (pre-registered, thresholds frozen here)

id	kill	status	threshold	what changed	mechanism	PLAN item
K-E1	ChaCAL- with-sink matches the committor read	REPAIR	within MDE_8 of the realised paired sd on $\geq 6/8$ seeds and residual within $2\times$ $\Rightarrow$ (e) is a parameteri- sation; escalate to $N = 16$, never “within TOST” at $N = 8$	TOST at N $= 8$ has no passing branch (M-13); TOST decides only the $\mathcal{A}$ $= \emptyset$ parity half	M-13, M-7, D-7	P4.1
K-E2	label degeneracy	REPAIR	any $\text{sd}(\mathbf{q}^{\{k\}})$ $= 0$ over the admitted region; any class frequency outside (0.05, 0.95); dis- agreement 0	inverted gate fixed	V-8, V-12	P1.2
K-G1	cached- mixture within one seed sd on $\Delta\mathbf{P}$ plants on $\geq 6/8$	KEEP	—	—	C5-type, D-2	P3.3
K-G2	field cosine within the cosine’s own MDE_8 of matched- depth softmax on $\geq 6/8$	REPAIR	cosine sd, not NRMSE sd (V-17)	V-17, D-1	P4.2	

id	kill	status	threshold	what changed	mechanism	PLAN item
K-H1	residual ratio ≤ 2 $\Rightarrow$ “joint determina- tion in one read” withdrawn; (2, 10) SPLIT	REPAIR	$\sigma_{\min}, \ I$ — $\gamma P_{\text{env}}\ _{\infty}$ printed beside r	V-17, V-26	P4.2	
K-H2	0-hop MLP within the CP half-width at $N = 8$, or McNemar p > 0.05 vs 1-hop $\Rightarrow$ BED-S is a static task, struck before any number	REPAIR	“the Fano floor’s resolution” replaced by the CP half-width	C8, V-10	P3.1	
K-D2	$\ P -$ $P_{\text{env}}\ _{\infty}$ $< 1e-3$ on $\geq 6/8 \Rightarrow$ a copy; registered as a learnability reading	KILL $\rightarrow$ applied	and an empty rejection region on the undirected substrate; non-empty on the DAG substrate	D-2, V-10	P3.1	

id	kill	status	threshold	what changed	mechanism	PLAN item
K-J	γ PINNED ($\Lambda \leq 2.706$) on $\geq 6/8 \Rightarrow$ softmax wearing a name; ablation (I — $\gamma P)^{-1}$ $\rightarrow I$ moving NRMSE by < 1 seed sd; mirror $\gamma >$ 0.99 on $\geq$ 6/8 $\Rightarrow$ certificate vacuous on the $z/\Delta z$ channel	REPAIR	$ \gamma <$ 0.05 deleted; boundary null; committor head unaffected by γ	M-20, V-17, V-9	P4.3	
K-I	exceedance on any of 1,024 draws; $\delta \cdot \ V\ _\infty$ $\geq$ $\text{sd}(\text{label})$	REPAIR, dormant	—	V-10, V-17	P5.2	
K-K	any [M] target not building by the Lean milestone $\Rightarrow$ [S]	KEEP	—	P-11, L-LEAN	P0.1	
K-M	Mapper schedule not better than do-nothing (and Zhao 2026 blockwise) by MDE_8 at matched visited tiles, in φ -NRMSE	REPAIR	metric named; second do-nothing control added	V-9, V-17	P5.3	

id	kill	status	threshold	what changed	mechanism	PLAN item
K-P	BED-S admission: realised sd $2.18\times$ the pilot $\Rightarrow$ reprice N; t^* unplaceable $\Rightarrow$ reroute	REPAIR	F1/F2 added to the cause list at 0 GPU-s; BED-M's 0.034451 and E4's 1372.50 marked V-22, not BED-S bars	D-4, M-3, V-22	P1.2, P3.1	
K-9	s-sweep exponent CI excludes 2 toward 3; any price quoted before the run	KEEP	—	D-3, P-8	P5.1	
K-10	any obstruction stated without its vacuity line; 0 % admitted on the bed	KEEP	—	V-25, P-10	P0.3	
K-11	any sentence calling the resolvent new	KEEP	—	P-7	assembly	
K-EMC	—	KILL $\rightarrow$ deleted		—	—	
K-L	—	KILL $\rightarrow$ deleted		—	—	

6. Predictions with counters (L-SIGN, D-7) — filed before P3

bet	prediction	counter	SPLIT band	deciding number	PLAN item
A	$\text{cos_shape} - \text{cos_softmax} > \text{MDE}_8(\text{cos})$	$\leq \text{MDE}_8(\text{cos})$	none (two-sided on the sign of the miss)	paired t on the cosine, N = 8	P3.1 / P4.2

bet	prediction	counter	SPLIT band	deciding number	PLAN item
B	$\text{CP_upper}(\text{err_shape}) \geq \text{CP_lower} < \text{floor_exact}(k)$ $\text{floor_exact}(k = 1)$ $\text{CP_upper} = 1)$	$\text{CP_lower} \geq \text{CP_lower} < \text{floor_exact}(k)$ $\text{CP_upper} = 1)$	$\text{CP_lower} < \text{CP_upper}$	the interval vs the floor	P3.1 / P4.2
C	γ MOVED on $\geq 6/8$ seeds	PINNED on $\geq 6/8$	3/8-5/8	Λ per seed	P4.3
D	depth-5 skyline within MDE_8 of the shape's argmin accuracy	short by $> \text{MDE}_8$ (signed)	none	$\text{acc_shape} - \text{acc_sky}$	P3.2
E	$\text{r_softmax} / \text{r_shape} \geq 10$ at marginal NRMSE within MDE_8	≤ 2	(2, 10)	the ratio, paired by seed and draw	P4.2
(j) LM γ	no row until an LM cell runs	—	—	—	P5.8 (Kaggle: the author's explicit yes is a node)

7. Tally

status	count
KEEP	27
REPAIR	40
KILL→applied	4 (F1 mask certificate, committor certificate, D-2 kill's substrate, the support sentence folded into B-E1)
KILL→deleted	6 (EMC bind, EMC kill, row L bind, row L kill, "ChaCAL must emit chance", the regime-N survival route — the last is recorded in <code>sec_apparatus.md</code> §A.12 without a ledger row)

The five first-evening items, in the coordinator's order: **P0.1** the [M] Lean targets and their refusals (`gamma_zero_is_softmax`, `lower_triangular_isUnit`, `resolvent_fromBlocks`, `segmentation_blockdiag` on `StrictlyLower`, `displacement_identity`, `bos_row_is_absorbing` with $\beta = 1$ stated, `softmax_corner_not_nilpotent` with $0 < \gamma$ and lower-triangularity stated, `later_boundary_unreachable`), 0 GPU-s; **P0.2** the extended manifest with its drift plant, 0 GPU-s; **P1.2** the BED-S generator spec with the DAG substrate, sink/goal/constraint placement and the census on 512 draws printing label sd, class frequencies, sink share, discard count — PASS if every sd > 0.05 , every class frequency in $(0.05, 0.95)$, $0 \in \mathcal{A}_{\text{sink}}$ and every set before the query on 100 %; **P2.1** the ChaCAL-diag smoke test at $\gamma = 0.9$ on BED-M as the solve-path check (≈ 1.7 s); **P4.3** the γ -pinning LR-test instrument spec with

the boundary null 2.706 and the held-out set declared. The ~6 GPU-s capped run at seeds 2, 3, 7 that the record priced three times and never took (READ V20_R15_IT35_JUPITER.md:157-159) is the sixth, on BED-M, and is independent of every row above.

Appendix B. Vocabulary — the field’s name for each construction

Every object in this paper was built or reached inside the repository before it was named. This appendix gives the standard term for each, the field the term comes from, and what the term commits its user to, so that a sentence about this work can be read by someone who has never opened the repository. It is a naming table, not a claim table: a term appearing here says the construction belongs to a known family, never that the family’s results transfer to it. Where two fields use one word for different objects, both are given, because that ambiguity has cost this campaign real numbers.

B.1 The operator and its read

the construction	the field’s term	field of origin	what the term commits you to
$z = (I - \gamma P)^{-1}V$	resolvent of P at $1/\gamma$; equivalently the Neumann series $\sum_t (\gamma P)^t$	operator theory, numerical linear algebra	convergence needs $\ \gamma P\ < 1$ or nilpotency; the object is linear in V
the same, for row-stochastic P and a reward V	policy evaluation ; $(I - \gamma P)^{-1}$ is the successor representation	reinforcement learning ([14], [36])	γ is a discount; the entries are expected discounted visit counts, so they are non-negative and sum to $1/(1 - \gamma)$
the same on a graph adjacency	personalised PageRank / Katz index	network science ([64], [20], [45])	teleport probability $1 - \gamma$; the read is a centrality unless V carries content
$\Pi_\gamma = (1 - \gamma)P(I - \gamma P)^{-1}$	the discounted occupancy kernel ; a Markov kernel in its own right	Markov chains	it is row-stochastic, so the read is a convex mixture and cannot leave the value hull
solving $(I - \gamma P)z = V$ by substitution rather than inverting	forward substitution on a triangular system	numerical linear algebra	cost $\Theta(s^2 d)$, sequential depth s ; the inverse is never formed
computing it in blocks	the chunked or UT-transform / WY form	linear-attention kernels ([136], [34])	the intra-chunk term is a small dense triangular inverse; inter-chunk is a matmul

the construction	the field's term	field of origin	what the term commits you to
γ swept from 0 to 1	the discount , or the horizon ; $\gamma \uparrow 1$ is the undiscounted limit	RL, Markov chains	at $\gamma = 1$ with an absorbing row the resolvent of the full matrix does not exist; the limit is taken on the transient block

B.2 The boundary rows and what they read

the construction	the field's term	field of origin	what the term commits you to
a row set to e_a	an absorbing state ; the set is an absorbing class ; imposing it is a Dirichlet boundary condition	Markov chains; potential theory	once entered, never left; the chain is then a killed or absorbed walk
the block form $\begin{pmatrix} Q & R \\ 0 & I \end{pmatrix}$	canonical form of an absorbing chain; Q is the transient block	Markov chains ([65])	$\rho(Q) < 1$ is <i>absorption almost surely</i> , and it is a hypothesis, not a given
$N = (I - Q)^{-1}$	the fundamental matrix	Markov chains	N_{ij} is the expected number of visits to j from i before absorption
$B = NR$, its column k	the absorption probability ; for two sets, the committor or splitting probability ; in control, the reach-avoid probability	Markov chains; transition-path theory ([82]); stochastic reachability ([114])	the columns sum to one across <i>all</i> absorbing sets, which is why a goal set must exist
the same quantity solved as $Lq = 0$ with q clamped on the boundary	the discrete Dirichlet problem ; q is a harmonic function ; the machine-learning instance is label propagation	potential theory ([37]); semi-supervised learning ([145])	harmonicity is on the <i>interior</i> only; the boundary values are data
the set $\{q = \frac{1}{2}\}$	the isocommittor surface ; in chemistry, the transition state	transition-path theory ([40])	it is a level set of a function, not a partition of the state space

the construction	the field’s term	field of origin	what the term commits you to
a position that absorbs mass without being asked to	an attention sink (a <i>column</i> device) — distinct from an absorbing <i>row</i>	transformer analysis ([130])	a sink receives weight; an absorbing row emits none. The two coexist and are not each other

B.3 Interventions and decisions

the construction	the field’s term	field of origin	what the term commits you to
rewriting a row of P and re-solving	an intervention , $\text{do}(a)$; on a linear system, graph surgery on a structural causal model	causal inference ([91], [112])	the model must be the data-generating one for the word “causal” to be earned; on a learned P it is an intervention on the <i>model</i>
Δz after that rewrite	the total effect ; for an equilibrium, the equilibrium displacement	causal inference; comparative statics ([86], [16])	it is defined relative to a stated pre-intervention state
updating the solve after a rank-one row change	the Sherman–Morrison formula; for several rows, Woodbury ; on a Markov chain, fundamental-matrix perturbation	numerical linear algebra ([111]); Markov chains ([109])	the update is exact, not an approximation; it costs one solve’s worth of back-substitution
whether “intervene then settle” equals “settle then intervene”	equilibration–manipulation commutability	causal modelling of equilibria ([35])	it can fail; it holds here by triangularity, which is a property of the operator and not of the world
$\arg \max_a q^{(0)}(\text{do } a)$	a safety filter or least-restrictive filter ; the value is a reach-avoid value function	safe control ([57], [44])	it is a decision rule over <i>evaluated</i> candidates, not a learned policy
$\arg \min_a \max_k q^{(k)}$	Chebyshev scalarisation of a multi-objective problem	multi-objective optimisation ([119])	the max-form and the weighted-sum form give different optima; the choice is a modelling decision and must be declared

the construction	the field’s term	field of origin	what the term commits you to
ranking constraints instead of maxing them	lexicographic ordering; bounding each is a constrained MDP	multi-objective RL ([132]); CMDPs ([2])	Bellman’s principle can fail on multichain safety-constrained problems ([84])

B.4 The gate, the segmentation and the topology

the construction	the field’s term	field of origin	what the term commits you to
a multiplicative gate m_k hitting exactly zero	a reset gate ; the segment it opens is a chunk boundary	gated linear attention ([134], [75])	implementations that carry the gate as $\log m$ cannot represent the exact zero; that is the repository’s own theorem
the resulting zero cross-block	block-diagonal structure; the solve is decoupled	linear algebra	it is exact, not sparse-approximate; approximate masks fill in under the inverse
filtering an influence graph by a threshold and counting components	the β_0 barcode of a filtration ; the construction is persistent homology	topological data analysis	stability constants for <i>directed</i> networks differ from the classical ones ([117])
covering a space by overlapping cells and taking the nerve	the Mapper graph; the object it estimates is a Reeb graph	topological data analysis ([113], [21])	the cover’s parameters must be fixed before the data, or the summary is chosen rather than measured
the record’s own S^2 point clouds	a Vietoris–Rips complex; the parameter sweep crosses a percolation transition	computational topology	connectivity is β_0 ; nothing about β_1 follows from a 1-skeleton

B.5 The instruments and the statistics

the construction	the field’s term	field of origin
the running evidence value that licenses stopping at any time declaring two arms equivalent rather than failing to distinguish them	an e-process ; the guarantee is Ville’s inequality ; the family is anytime-valid inference equivalence testing , specifically TOST ([108])	

the construction	the field’s term	field of origin
the smallest effect a design can see	the minimum detectable effect ; its inputs are power and the paired standard deviation	
an exact interval on a proportion	the Clopper–Pearson interval; the paired test on two arms is McNemar’s	
the floor below which no predictor can go	Fano’s inequality (for a discrete answer); rate–distortion (for a continuous one)	
a lower bound on error from the label’s own entropy at zero information	the zero-information or majority-class floor	
a bound with no sum over the sequence length	a uniform or dimension-free bound	

B.6 Two words this campaign used in two senses, and what each cost

“**Causal.**” In a transformer, *causal* names the triangular mask: position i reads only $j \leq i$. In causal inference, *causal* names invariance under intervention. The repository’s operator is causal in the first sense throughout, and the second sense is earned only on a bed where the label is generated by an intervention on a stated model. Every sentence in this paper that uses the second sense names the bed. This is the term-slide the campaign paid for most often.

“**State.**” In linear algebra a state is a vector; in Markov-chain language a state is a *position with a transition law*; in state-space models the state is the recurrence’s carried summary, whose *dimension* is a capacity bound. The repository’s “next state toward equilibrium” is the Markov sense; its `d_model` is the linear-algebra sense; the delay theorems of `V15Kernel` are about the third. A sentence that slides between them proves nothing, and the record has one such sentence per round.

“**Depth.**” Depth in a transformer is the number of parameter blocks; depth in circuit complexity is the length of the longest computation path. The shape has one parameter block and a solve whose sequential depth is the sequence length. Both are said, in the same paragraph, every time either is claimed.

References

- [1] Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008. doi: 10.1016/j.automatica.2008.03.027. URL <https://doi.org/10.1016/j.automatica.2008.03.027>. [V] Crossref record fetched (Elsevier landing page unreadable). Owns: probabilistic safety of controlled DTSHS as a multiplicative-cost dynamic programme; maximal safe sets at a given probability level.
- [2] Eitan Altman. *Constrained Markov Decision Processes*. Stochastic Modeling. Chapman & Hall/CRC, 1999. ISBN 9780849303821. URL https://openlibrary.org/books/OL97592M/Constrained_Markov_decision_processes. [V-cat] Open Library record fetched (Routledge landing page 403). Owns: the CMDP – maximise expected reward subject to K expected-cost constraints; occupation-measure LP; Lagrangian duality.

- [3] Eitan Altman. *Constrained Markov Decision Processes*. Stochastic Modeling. Chapman & Hall/CRC, Boca Raton, 1999. ISBN 0-8493-0382-6. doi: 10.1201/9781315140223. URL <https://doi.org/10.1201/9781315140223>. [V] Crossref DOI record fetched 2026-09-03 for the Routledge 2021 digital reissue (e-ISBN 9781315140223, title and series matched). The 1999 print ISBN is from memory – [U] on that field. Found by widening search S3.
- [4] Miriam Ayer, H. D. Brunk, G. M. Ewing, W. T. Reid, and Edward Silverman. An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics*, 26(4):641–647, 1955. doi: 10.1214/aoms/1177728423. URL <https://doi.org/10.1214/aoms/1177728423>. [V] Crossref DOI record fetched 2026-09-03. The original pool-adjacent-violators construction; Barlow et al. 1972 is the textbook, not the origin.
- [5] Amitoz Azad. Learning label initialization for time-dependent harmonic extension. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI 2022)*, 2022. doi: 10.48550/arXiv.2205.01358. URL <https://arxiv.org/abs/2205.01358>.
- [6] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, 19(3):357–367, 1967. doi: 10.2748/tmj/1178243286. URL <https://doi.org/10.2748/tmj/1178243286>. [V] Crossref DOI record fetched 2026-09-03 (title, volume, issue, year matched; the record carried no page field – pages 357–367 are from memory, [U] on that field).
- [7] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, 2019. doi: 10.48550/arXiv.1909.01377. URL <https://arxiv.org/abs/1909.01377>.
- [8] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. The MIT Press, 2008. ISBN 9780262026499. URL <https://openlibrary.org/isbn/9780262026499>. [V] Open Library ISBN record fetched 2026-09-03 (title, MIT Press, 2008-05-31); author keys on the record were not expanded to names, names from memory [U]. Reachability probabilities as a linear equation system (Ch.10) from memory [U].
- [9] Michèle Basseville and Igor V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, Englewood Cliffs, NJ, 1993. ISBN 0-13-126780-9. URL <https://people.irisa.fr/Michele.Basseville/kniga/>.
- [10] Michèle Basseville and Igor V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, Englewood Cliffs, NJ, 1993. ISBN 0-13-126780-9. URL <https://people.irisa.fr/Michele.Basseville/kniga/>. [V] author-hosted page fetched 2026-09-03 and the PDF table of contents read this session: Chapter 7 'Additive Changes in Linear Models' p.209; Section 7.2 'Statistical Approach' p.214; Subsection 7.2.4 'State-Space Models' p.234.
- [11] Alexandra Bazarova, Andrei Volodichev, Aleksandr Yugay, Andrey Shulga, Alina Ermilova, Konstantin Polev, Julia Belikova, Rauf Parchiev, Dmitry Simakov, Maxim Savchenko, Andrey Savchenko, Serguei Barannikov, and Alexey Zaytsev. Hallucination detection in LLMs with topological divergence on attention graphs. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)*, 2025. URL <https://arxiv.org/abs/2504.10063>. [V] abs page fetched 2026-09-03; title matched; venue as printed on the abs page.
- [12] Maximilian Beck, Korbinian Pöppel, Phillip Lippe, and Sepp Hochreiter. Tiled flash linear attention: More efficient linear RNN and xLSTM kernels. *arXiv preprint arXiv:2503.14376*, March 2025. doi: 10.48550/arXiv.2503.14376. URL <https://arxiv.org/abs/2503.14376>.

- [13] Ahmed Begga, Francisco Escolano, Miguel Angel Lozano, and Edwin R. Hancock. Diffusion-jump GNNs: Homophiliation via learnable metric filters, 2023. URL <https://arxiv.org/abs/2306.16976>.
- [14] Richard Bellman. A Markovian decision process. *Indiana University Mathematics Journal (Journal of Mathematics and Mechanics)*, 6(4):679–684, 1957. doi: 10.1512/iumj.1957.6.56038. URL <https://doi.org/10.1512/iumj.1957.6.56038>. [V] Crossref DOI record fetched 2026-09-03. Policy evaluation $v = (I - \gamma P)^{-1} r$ and the argmin over actions.
- [15] Tineke Blom, Stephan Bongers, and Joris M. Mooij. Beyond structural causal models: Causal constraints models. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI 2019)*, 2019. URL <https://arxiv.org/abs/1805.06539>. [V] arXiv abs fetched 2026-09-03.
- [16] Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, 2021. URL <https://arxiv.org/abs/1611.06221>. [V] arXiv abs fetched 2026-09-03; journal reference as listed on the abs page.
- [17] Javier Borquez, Kaustav Chakraborty, Hao Wang, and Somil Bansal. On safety and liveness filtering using hamilton-jacobi reachability analysis, 2023. URL <https://arxiv.org/abs/2312.15347>. [V] abs fetched. Owns: the "least-restrictive filter" – the maximal set of controls that keep the safety value non-negative at a state, with the nominal action projected onto it.
- [18] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X. Charles, D. Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14:3207–3260, 2013. URL <https://www.jmlr.org/papers/v14/bottou13a.html>. [V] JMLR landing page and arXiv abs fetched 2026-09-03; PDF text read this session: Section 7 'Equilibrium Analysis', Section 7.3 'Estimating the equilibrium response'.
- [19] David Brandfonbrener, Alberto Bietti, Jacob Buckman, Romain Laroche, and Joan Bruna. When does return-conditioned supervised learning work for offline reinforcement learning?, 2022. URL <https://arxiv.org/abs/2206.01079>. [V] abs fetched. Owns: RCSL recovers the optimal policy only under assumptions strictly stronger than dynamic programming needs, with MDP counterexamples.
- [20] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117, 1998. doi: 10.1016/S0169-7552(98)00110-X. URL [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X). [V] Crossref DOI record fetched 2026-09-03. PageRank; the personalised form $\alpha (I - (1-\alpha)P)^{-1} e$ is the resolvent of a row-stochastic P .
- [21] Mathieu Carrière, Bertrand Michel, and Steve Oudot. Statistical analysis and parameter selection for mapper. *Journal of Machine Learning Research*, 19, 2018. URL <https://arxiv.org/abs/1706.00204>. [V] abs page fetched 2026-09-03; title matched (arXiv submission 2017). JMLR volume seen in a search-index URL (jmlr.org/papers/volume19/17-291) – [U].
- [22] Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. PersLay: A neural network layer for persistence diagrams and new graph topological signatures, 2019. URL <https://arxiv.org/abs/1904.09378>. [V] abs page fetched 2026-09-03; title matched. CEQ_V20_R15_CONTRACT.md:223 attributes this id to Moor et al. (topological autoencoders); the id resolves to PersLay. AISTATS-2020 venue not fetched – [U].

- [23] Mathieu Carrière, Frédéric Chazal, Marc Glisse, Yuichi Ike, and Hariprasad Kannan. Optimizing persistent homology based functions. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, 2021. URL <https://arxiv.org/abs/2010.08356>. [V] abs page fetched 2026-09-03; title matched (arXiv submission 2020). ICML 2021 venue from memory – [U]; the abs page states no venue.
- [24] Benjamin Paul Chamberlain, James Rowbottom, Maria Gorinova, Stefan Webb, Emanuele Rossi, and Michael M. Bronstein. Grand: Graph neural diffusion. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, 2021. URL <https://arxiv.org/abs/2106.10934>. [V] arXiv abs and ar5iv HTML fetched 2026-09-03; implicit step $(I - \tau \bar{A}(x^{(k)}))x^{(k+1)} = x^{(k)}$ read from the body. ICML 2021 venue from memory [U].
- [25] Lijie Chen, Binghui Peng, and Hongxun Wu. Theoretical limitations of multi-layer transformer. *arXiv preprint*, 2024. doi: 10.48550/arXiv.2412.02975. URL <https://arxiv.org/abs/2412.02975>. [V] abs fetched 2026-09-03 (v1 4 Dec 2024). Theorem 1.1: for constant $L \leq O(\log \log n)$, an L -layer decoder-only softmax transformer cannot solve L -sequential function composition when $H^*d^*p \leq n^{2^{(-4L)}}$; unconditional; $p \geq \log n$. Corollary 1.2: an $(L+1)$ -layer model solves it with polylog parameters. Corollary 1.3: $O(\log L)$ -layer encoder vs L -layer decoder separation, unconditional.
- [26] Daniil Cherniavskii, Eduard Tulchinskii, Vladislav Mikhailov, Irina Proskurina, Laida Kushnareva, Ekaterina Artemova, Serguei Barannikov, Irina Piontkovskaya, Dmitri Piontkovski, and Evgeny Burnaev. Acceptability judgements via examining the topology of attention maps. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 2022. URL <https://arxiv.org/abs/2205.09630>. [V] abs page fetched 2026-09-03; title matched.
- [27] Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. Adaptive universal generalized pagerank graph neural network. In *International Conference on Learning Representations (ICLR 2021)*, 2021. URL <https://arxiv.org/abs/2006.07988>. [V] arXiv abs fetched 2026-09-03 (ICLR 2021 stated on the page).
- [28] Sungjun Cho, Seonwoo Min, Jinwoo Kim, Moontae Lee, Honglak Lee, and Seunghoon Hong. Transformers meet stochastic block models: Attention with data-adaptive sparsity and cost, 2022. URL <https://arxiv.org/abs/2210.15541>. [V] abs page fetched 2026-09-03; title matched. NeurIPS 2022 venue seen only in a search-index snippet – [U].
- [29] Samir Chowdhury and Facundo Mémoli. Persistent path homology of directed networks, 2017. URL <https://arxiv.org/abs/1701.00565>. [V] abs page fetched 2026-09-03; title matched. SODA 2018 venue from memory – [U].
- [30] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. doi: 10.1093/biomet/26.4.404. URL <https://doi.org/10.1093/biomet/26.4.404>. [V] OUP landing page fetched 2026-09-03; title matched.
- [31] Stephen A. Cook. A taxonomy of problems with fast parallel algorithms. *Information and Control*, 64(1–3):2–22, 1985. doi: 10.1016/S0019-9958(85)80041-3. URL [https://doi.org/10.1016/S0019-9958\(85\)80041-3](https://doi.org/10.1016/S0019-9958(85)80041-3). [V] Crossref record for the DOI fetched 2026-09-03, title/author/volume/pages matched. Cited for the class DET (integer determinant / matrix inversion) and its position NL subset-of DET subset-of NC2 – the class the shape’s exact resolvent solve belongs to (DERIVED in the sweep, not a theorem of this paper about attention).

- [32] Padraig Corcoran and Bailin Deng. Regularization of persistent homology gradient computation, 2020. URL <https://arxiv.org/abs/2011.05804>. [V] abs page fetched 2026-09-03; title matched. CEQ_V20_R15_CONTRACT.md:224 attributes this id to Carrière et al. 2021; the id resolves to Corcoran and Deng.
- [33] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Hoboken, NJ, 2 edition, 2006. ISBN 978-0-471-24195-9. doi: 10.1002/047174882X. URL <https://doi.org/10.1002/047174882X>. [V] Crossref DOI record fetched 2026-09-03 (print ISBN 9780471241959, e-ISBN 9780471748823); the publisher’s 30-page sample PDF was fetched and read: its copyright page reads ISBN-13 978-0-471-24195-9, Copyright 2006; its table of contents lists “Fano’s Inequality” in Chapter 2 and “Gaussian Source” in the rate-distortion chapter.
- [34] Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, 2024. URL <https://arxiv.org/abs/2405.21060>. [V] arXiv abs fetched 2026-09-03 (ICML 2024 stated).
- [35] Denver Dash. Restructuring dynamic causal systems in equilibrium. In *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS 2005)*, volume R5 of *Proceedings of Machine Learning Research*, pages 81–88, 2005. URL <https://proceedings.mlr.press/r5/dash05a.html>. [V] PMLR landing page fetched 2026-09-03; PDF text read this session: Definition 4 (EMC property), Theorem 1 (EMC violation), Theorem 2 (EMC obedience).
- [36] Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993. doi: 10.1162/neco.1993.5.4.613. URL <https://doi.org/10.1162/neco.1993.5.4.613>.
- [37] Peter G. Doyle and J. Laurie Snell. *Random Walks and Electric Networks*. Number 22 in Carus Mathematical Monographs. Mathematical Association of America, 1984. URL <https://arxiv.org/abs/math/0001057>. [V] arXiv abs fetched. Owns: hitting probability as the harmonic function with Dirichlet data (escape probability, effective resistance). The repository’s Kirchhoff cross-check (MATHEMATICS.md section 11) is this book’s identity.
- [38] O. Duranthon, P. Marion, C. Boyer, B. Loureiro, and L. Zdeborová. Statistical advantage of softmax attention: Insights from single-location regression. In *International Conference on Learning Representations (ICLR 2026)*, 2026. doi: 10.48550/arXiv.2509.21936. URL <https://arxiv.org/abs/2509.21936>.
- [39] Weinan E and Eric Vanden-Eijnden. Towards a theory of transition paths. *Journal of Statistical Physics*, 123(3):503–523, 2006. doi: 10.1007/s10955-005-9003-9. URL <https://doi.org/10.1007/s10955-005-9003-9>.
- [40] Weinan E and Eric Vanden-Eijnden. Transition-path theory and path-finding algorithms for the study of rare events. *Annual Review of Physical Chemistry*, 61:391–420, 2010. doi: 10.1146/annurev.physchem.040808.090412. URL <https://doi.org/10.1146/annurev.physchem.040808.090412>. [V] Crossref record fetched (Annual Reviews landing page 403). Owns: the review statement of TPT; committor as the reaction coordinate; isocommittor surfaces.
- [41] Yotam Erel, Olaf Dünkel, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Amit H. Bermano. Attention (as discrete-time markov) chains. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025. URL <https://arxiv.org/abs/2507.17657>. [V] arXiv abs (NeurIPS 2025) and HTML fetched 2026-09-03; post-hoc DTMC reading, power iteration and a damped stationary vector (TokenRank); vision only.

- [42] Erwan Fagnou, Paul Caillon, Blaise Delattre, and Alexandre Allauzen. Chain and causal attention for efficient entity tracking. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024, Main)*, 2024. doi: 10.18653/v1/2024.emnlp-main.731. URL <https://arxiv.org/abs/2410.05565>. Eq. (5): $Y = (1 - \gamma) A(I - \gamma A)^{-1} V$, $\gamma \in [0, 1)$, A causal softmax; Eq. (7): solved as the triangular system $(I - \gamma A)Y = (1 - \gamma)AV$.
- [43] Aosong Feng, Irene Li, Yuang Jiang, and Rex Ying. Diffuser: Efficient transformers with multi-hop attention diffusion for long sequences, 2022. URL <https://arxiv.org/abs/2210.11794>. [V] arXiv abs and HTML fetched 2026-09-03; PPR-style $\sum_k \alpha(1 - \alpha)^k A^k$ over a sparse softmax attention on sequences, power-iterated K steps, bidirectional.
- [44] Jaime F. Fisac, Neil F. Lugovoy, Vicenç Rubies-Royo, Shromona Ghosh, and Claire J. Tomlin. Bridging hamilton-jacobi safety analysis and reinforcement learning. In *International Conference on Robotics and Automation (ICRA)*, pages 8550–8556, 2019. doi: 10.1109/ICRA.2019.8794107. URL <https://doi.org/10.1109/ICRA.2019.8794107>. [V] Crossref record fetched. Owns: time-discounting of the minimum-payoff-over-time safety problem so that the Bellman operator contracts – the same role $\gamma < 1$ plays for the shape’s Neumann certificate.
- [45] Johannes Gasteiger, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. In *International Conference on Learning Representations (ICLR 2019)*, 2019. doi: 10.48550/arXiv.1810.05997. URL <https://arxiv.org/abs/1810.05997>. Author formerly listed as Klicpera. Propagation $Z = \alpha(I - (1 - \alpha)\hat{A})^{-1}H$.
- [46] Dobrik Georgiev, JJ Wilson, Davide Buffelli, and Pietro Liò. Deep equilibrium algorithmic reasoning. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024. doi: 10.48550/arXiv.2410.15059. URL <https://arxiv.org/abs/2410.15059>.
- [47] Charles M. Grinstead and J. Laurie Snell. *Introduction to Probability*. American Mathematical Society, Providence, RI, second revised edition, 1997. ISBN 978-0-8218-0749-1. URL <https://bookstore.ams.org/iprob/>. Ch. 11 §11.2, Definition 11.3 $N = (I - Q)^{-1}$; Theorem 11.6 $B = NR$.
- [48] Albert Gu, Ankit Gupta, Karan Goel, and Christopher Ré. On the parameterization and initialization of diagonal state space models. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022. doi: 10.48550/arXiv.2206.11893. URL <https://arxiv.org/abs/2206.11893>.
- [49] Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. Implicit graph neural networks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 2020. URL <https://arxiv.org/abs/2009.06211>. [V] arXiv abs fetched 2026-09-03. NeurIPS 2020 venue from memory [U].
- [50] Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. When attention sink emerges in language models: An empirical view. In *International Conference on Learning Representations (ICLR 2025)*, 2025. URL <https://arxiv.org/abs/2410.10781>. [V] arXiv abs fetched 2026-09-03 (ICLR 2025 spotlight stated).
- [51] William W. Hager. Updating the inverse of a matrix. *SIAM Review*, 31(2):221–239, 1989. doi: 10.1137/1031049. URL <https://doi.org/10.1137/1031049>. [V] Crossref DOI record fetched 2026-09-03. The survey of Sherman–Morrison–Woodbury updates (Woodbury’s 1950 Princeton memorandum itself has no DOI and was not fetched).

- [52] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830. URL <https://doi.org/10.1080/01621459.1963.10500830>. [V] Crossref DOI record fetched 2026-09-03.
- [53] Christoph Hofer, Roland Kwitt, Mandar Dixit, and Marc Niethammer. Connectivity-optimized representation learning via persistent homology. In *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)*, 2019. URL <https://arxiv.org/abs/1906.09003>. [V] abs page fetched 2026-09-03; title matched. ICML 2019 venue from memory – [U]; the abs page states no venue.
- [54] Benjamin Hoover, Yuchen Liang, Bao Pham, Rameswar Panda, Hendrik Strobelt, Duen Horng Chau, Mohammed J. Zaki, and Dmitry Krotov. Energy transformer. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, 2023. doi: 10.48550/arXiv.2302.07253. URL <https://arxiv.org/abs/2302.07253>.
- [55] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition, 2013. ISBN 978-0-521-54823-6. doi: 10.1017/CBO9781139020411. URL <https://doi.org/10.1017/CBO9781139020411>. [V] Crossref DOI record fetched 2026-09-03 (ISBNs 9780521839402, 9780521548236, e-ISBN 9781139020411; Crossref ‘issued’ reads 2012, the print copyright is 2013).
- [56] Kai-Chieh Hsu, Vicenç Rubies-Royo, Claire J. Tomlin, and Jaime F. Fisac. Safety and liveness guarantees through reach-avoid reinforcement learning. In *Robotics: Science and Systems (RSS)*, 2021. URL <https://arxiv.org/abs/2112.12288>. [V] abs fetched. Owns: the discounted reach-avoid Bellman backup (max of the constraint value and min of next-step values) with a contraction property; reach-avoid Q-learning.
- [57] Kai-Chieh Hsu, Haimin Hu, and Jaime Fernández Fisac. The safety filter: A unified view of safety-critical control in autonomous systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 2023. URL <https://arxiv.org/abs/2309.05837>. [V] abs fetched. Owns: the umbrella name "safety filter" – a modular monitor + intervention that overrides a nominal action when a safety value says so; covers HJ, CBF, predictive and learned filters.
- [58] Jerry Yao-Chieh Hu, Xiwen Zhang, Ali ElSheikh, Weimin Wu, and Han Liu. On structured state-space duality. *arXiv preprint arXiv:2510.04944*, October 2025. doi: 10.48550/arXiv.2510.04944. URL <https://arxiv.org/abs/2510.04944>. revised 23 Dec 2025.
- [59] Sukjun Hwang, Brandon Wang, and Albert Gu. Dynamic chunking for end-to-end hierarchical sequence modeling, 2025. URL <https://arxiv.org/abs/2507.07955>. [V] abs page fetched 2026-09-03; title matched.
- [60] Ashkan B. Jeddi, Nariman L. Dehghani, and Abdollah Shafieezadeh. Lyapunov-based uncertainty-aware safe reinforcement learning, 2021. URL <https://arxiv.org/abs/2107.13944>. [V] abs fetched. Owns: a Transformer memory plus a risk-averse action rule that picks the action with the lowest predicted constraint-violation probability; the probability comes from an uncertainty model, not from a resolvent in the layer.
- [61] Alireza Karbalayghareh, Evan Paull, and Andrea Califano. DoFormer: causal transformer for gene perturbation. bioRxiv preprint, 2026. URL <https://www.biorxiv.org/content/10.64898/2026.05.02.722054v1>. [V] bioRxiv landing page fetched 2026-09-03 (posted 2026-05-04). Found by widening search S10. Implements do(.) inside attention by fixing the intervened token’s value and forbidding it from attending – an absorbing row, for gene perturbation, without a resolvent.

- [62] Nikolaos Kariotoglou, Maryam Kamgarpour, Tyler Summers, and John Lygeros. The linear programming approach to reach-avoid problems for Markov decision processes, 2014. URL <https://arxiv.org/abs/1411.5925>.
- [63] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are RNNs: Fast autoregressive transformers with linear attention. *arXiv preprint arXiv:2006.16236*, June 2020. doi: 10.48550/arXiv.2006.16236. URL <https://arxiv.org/abs/2006.16236>.
- [64] Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1): 39–43, 1953. doi: 10.1007/BF02289026. URL <https://doi.org/10.1007/BF02289026>. [V] Crossref DOI record fetched 2026-09-03.
- [65] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. D. Van Nostrand, Princeton, NJ, 1960. ISBN 9780442043285. URL <https://archive.org/details/finitemarkovchai0000unse>. [V-cat] Internet Archive catalogue record fetched (Van Nostrand 1960, 210 pp.). Springer UTM reprint 1976, ISBN 9780387901923, reached only via search index [U]. Owns: canonical form, fundamental matrix $N = (I - Q)^{-1}$, $B = NR$.
- [66] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains: With a New Appendix “Generalization of a Fundamental Matrix”*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1976. ISBN 0-387-90192-2. [V] Google Books catalogue record (vid=ISBN9780387901923) fetched 2026-09-03: title with the appendix subtitle, both authors, Springer New York, Undergraduate Texts in Mathematics, ISBN 0387901922 / 9780387901923 matched; that record is dated 1983 while the booksellers date the reprint 1976 – [U] on the year field only. First edition: Van Nostrand, Princeton, 1960 – two 1960/1961 Crossref review records titled “Finite Markov Chains” (Bush, doi:10.2307/2309264; White, doi:10.1086/222959) were returned by a Crossref bibliographic query. CAUTION: the Springer DOI 10.1007/978-1-4684-9455-6 resolves to “Denumerable Markov Chains” (Kemeny, Snell, Knapp, GTM 40, 1976), a different book; do not attach it to this entry.
- [67] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. Springer-Verlag, 1976. ISBN 9780387901923. URL <https://openlibrary.org/isbn/9780387901923>. [V] Open Library ISBN record fetched 2026-09-03 (title, authors, Springer-Verlag; that record lists a 1983 printing). 1976 edition year from memory [U]. Fundamental matrix $N = (I - Q)^{-1}$, absorption $B = NR$ (Ch. III) from memory [U].
- [68] Kwangho Kim and Hajin Lee. Topological causal effects, 2026. URL <https://arxiv.org/abs/2603.02289>. Journal-ref on abs page: Proceedings of ICLR 2026.
- [69] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *International Conference on Learning Representations (ICLR 2020)*, 2020. URL <https://arxiv.org/abs/2001.04451>. [V] abs page fetched 2026-09-03; title matched; venue as printed on the abs page.
- [70] Alexander Kozachinskiy, Felipe Urrutia, Hector Jimenez, Tomasz Steifer, Germán Pizarro, Matías Fuentes, Francisco Meza, Cristian B. Calderon, and Cristóbal Rojas. Strassen attention, split vc dimension and compositionality in transformers. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2501.19215. URL <https://arxiv.org/abs/2501.19215>. [V] abs fetched 2026-09-03 (v3 25 Sep 2025); comments: accepted at NeurIPS 2025. Earlier versions carried the title "Strassen Attention: Unlocking Compositional Abilities in Transformers Based on a New Lower Bound Method" (search index). Theorem 3.2: $\text{size}(T) = \text{split-VC}(f)^{\Omega(1)}$ for one-layer softmax transformers at INFINITE precision. Theorem 3.4: $\text{size } n^{\omega(1)}$ cannot do function composition. Theorem 3.6 (Match3), 3.8 (binary relation composition). Strassen attention: $a_i = \sum_j, k a_{ijk}(v_j \circ v_k)$, $O(n^{\omega(d)})$.

- [71] Laida Kushnareva, Daniil Cherniavskii, Vladislav Mikhailov, Ekaterina Artemova, Serguei Barannikov, Alexander Bernstein, Irina Piontkovskaya, Dmitri Piontkovski, and Evgeny Burnaev. Artificial text detection via examining the topology of attention maps. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021. URL <https://arxiv.org/abs/2109.04825>. [V] abs page fetched 2026-09-03; title matched.
- [72] Laida Kushnareva, Dmitri Piontkovski, and Irina Piontkovskaya. Betti numbers of attention graphs is all you really need, 2022. URL <https://arxiv.org/abs/2207.01903>. [V] abs page fetched 2026-09-03; title matched. The abs page records the manuscript as submitted to the TDA and Beyond workshop at NeurIPS 2020 and not accepted; it is cited as a preprint.
- [73] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, second edition, 2017. ISBN 9781470429621. doi: 10.1090/mbk/107. URL <https://doi.org/10.1090/mbk/107>. [V] Crossref DOI record fetched 2026-09-03 (authors as registered: Levin, Peres; Wilmer credited as contributor to the second edition); Theorem 13.10 read on p. 183 of the authors’ PDF (pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf): for REVERSIBLE P with spectral gap γ and bottleneck ratio Φ^* , $\Phi^* \gamma^2/2 \leq \gamma \leq 2 \Phi^*$, attributed there to Sinclair–Jerrum (1989) and Lawler–Sokal (1988).
- [74] Yufeng Lin and Jialu Zhang. ScratchWorld: Evaluating if world models compute executable consequences, 2026. URL <https://arxiv.org/abs/2606.31689>. [V] arXiv abs fetched 2026-09-03.
- [75] Zhixuan Lin, Evgenii Nikishin, Xu Owen He, and Aaron Courville. Forgetting transformer: Softmax attention with a forget gate. In *International Conference on Learning Representations (ICLR 2025)*, 2025. URL <https://arxiv.org/abs/2503.02130>. [V] abs page and HTML full text fetched 2026-09-03; title matched; venue as printed on the abs page.
- [76] Juncheng Liu, Kenji Kawaguchi, Bryan Hooi, Yiwei Wang, and Xiaokui Xiao. Eignn: Efficient infinite-depth graph neural networks. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, 2021. URL <https://arxiv.org/abs/2202.10720>. [V] arXiv abs fetched 2026-09-03 (NeurIPS 2021 stated). The closed-form equation itself was not readable this session (PDF binary, ar5iv conversion failed) [U].
- [77] Pierre Marion, Raphaël Berthier, Gérard Biau, and Claire Boyer. Attention layers provably solve single-location regression. In *International Conference on Learning Representations (ICLR 2025)*, 2025. doi: 10.48550/arXiv.2410.01537. URL <https://arxiv.org/abs/2410.01537>.
- [78] Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947. doi: 10.1007/BF02295996. URL <https://doi.org/10.1007/BF02295996>. [V] Cambridge Core landing page fetched 2026-09-03; title matched.
- [79] William Merrill and Ashish Sabharwal. The parallelism tradeoff: Limitations of log-precision transformers. *Transactions of the Association for Computational Linguistics*, 2023. doi: 10.48550/arXiv.2207.00729. URL <https://arxiv.org/abs/2207.00729>. [V] abs fetched 2026-09-03 (v4 26 Apr 2023); comments: accepted at TACL. Theorem 2: log-precision constant-depth transformers are simulable by logspace-uniform TC0. Corollaries under TC0 $\neq P$ (stated as $L \neq P$ family): cannot solve linear equalities $Ax=b$ (P -complete), graph connectivity (L -complete), universal CFG recognition; asymptotic in n .

- [80] William Merrill and Ashish Sabharwal. The expressive power of transformers with chain of thought. *arXiv preprint*, 2024. doi: 10.48550/arXiv.2310.07923. URL <https://arxiv.org/abs/2310.07923>. [V] abs fetched 2026-09-03 (v5 11 Apr 2024, ICLR camera-ready per comments). Log CoT steps: little gain; linear steps: all regular languages; polynomial steps: P.
- [81] William Merrill and Ashish Sabharwal. A little depth goes a long way: The expressive power of log-depth transformers. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2503.03961. URL <https://arxiv.org/abs/2503.03961>. [V] abs fetched 2026-09-03 (v3 5 Nov 2025); comments: NeurIPS 2025. Theorem 2: a universal transformer unrolled $\text{ceil}(\log_2 n)$ times solves connectivity (directed or undirected, STCON) on n vertices; $18 + 2 \text{ceil}(\log_2 n)$ layers; $p \geq c \log n$; averaging-hard attention. Fixed-depth impossibility conditional on $\text{TC0} \neq \text{NC1}$.
- [82] Philipp Metzner, Christof Schütte, and Eric Vanden-Eijnden. Transition path theory for Markov jump processes. *Multiscale Modeling & Simulation*, 7(3):1192–1219, 2009. doi: 10.1137/070699500. URL <https://doi.org/10.1137/070699500>.
- [83] Carl D. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, Philadelphia, 2000. ISBN 978-0-89871-454-8. doi: 10.1137/1.9780898719512. URL <https://doi.org/10.1137/1.9780898719512>. [V] Crossref DOI record fetched 2026-09-03 (ISBN 9780898714548).
- [84] Rahul Misra, Rafał Wisniewski, and Carsten Skovmose Kallesøe. On Bellman’s principle of optimality and reinforcement learning for safety-constrained Markov decision process, 2023. URL <https://arxiv.org/abs/2302.13152>.
- [85] Ida Momennejad, Evan M. Russek, Jin H. Cheong, Matthew M. Botvinick, Nathaniel D. Daw, and Samuel J. Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, 2017. doi: 10.1038/s41562-017-0180-8. URL <https://doi.org/10.1038/s41562-017-0180-8>. [V] CrossRef record fetched 2026-09-03; the transition-revaluation finding is attributed from a search snippet only [U].
- [86] Joris M. Mooij, Dominik Janzing, and Bernhard Schölkopf. From ordinary differential equations to structural causal models: the deterministic case. In *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI 2013)*, 2013. URL <https://arxiv.org/abs/1304.7920>. [V] arXiv abs fetched 2026-09-03.
- [87] Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, 2020. URL <https://arxiv.org/abs/1906.00722>. [V] abs page fetched 2026-09-03; title matched; venue as printed on the abs page (arXiv submission 2019).
- [88] Ellis Batten Page. Ordered hypotheses for multiple treatments: a significance test for linear ranks. *Journal of the American Statistical Association*, 58(301):216–230, 1963. doi: 10.1080/01621459.1963.10500843. URL <https://doi.org/10.1080/01621459.1963.10500843>. [V] Crossref DOI record fetched 2026-09-03 (T&F landing page returned 403). Owner of Page’s L (MATHEMATICS.md section 12).
- [89] Giseung Park, Hyunyoung Nam, Woohyeon Byeon, Amir Leshem, and Youngchul Sung. Constrained multi-objective reinforcement learning with max-min criterion. In *International Conference on Machine Learning (ICML)*, 2026. URL <https://arxiv.org/abs/2605.31388>. [V] abs fetched. Owns: max-min (worst-objective) criterion with explicit constraints in MORL – the policy-level cousin of `argmin_a max_k`.
- [90] Keiran Paster, Sheila McIlraith, and Jimmy Ba. You can’t count on luck: Why decision transformers and rvs fail in stochastic environments, 2022. URL <https://arxiv.org/abs/>

- 2205.15967. [V] abs fetched. Owns: the negative – return-conditioning imitates lucky trajectories in stochastic environments; ESPER conditions on cluster-average return instead.
- [91] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, second edition, 2009. doi: 10.1017/CBO9780511803161. URL <https://doi.org/10.1017/CBO9780511803161>.
 - [92] Binghui Peng, Srini Narayanan, and Christos Papadimitriou. On limitations of the transformer architecture, 2024. URL <https://arxiv.org/abs/2402.08164>.
 - [93] Ilan Perez and Raphael Reinauer. The topological BERT: Transforming attention into topology for natural language processing, 2022. URL <https://arxiv.org/abs/2206.15195>. [V] abs page fetched 2026-09-03; title matched.
 - [94] Payam Piray and Nathaniel D. Daw. Linear reinforcement learning in planning, grid fields, and cognitive control. *Nature Communications*, 12:4942, 2021. doi: 10.1038/s41467-021-25123-3. URL <https://doi.org/10.1038/s41467-021-25123-3>. [V] Crossref DOI record and Europe PMC full text (PMC8368103) fetched 2026-09-03; Eq. 4 $\exp(v^*) = MP \exp(r)$ with M over nonterminal states and Eq. 5 $M = M_{\text{old}} + M_B$ read from the body.
 - [95] Jan-Hendrik Prinz, Hao Wu, Marco Sarich, Bettina Keller, Martin Senne, Martin Held, John D. Chodera, Christof Schütte, and Frank Noé. Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics*, 134(17):174105, 2011. doi: 10.1063/1.3565032. URL <https://doi.org/10.1063/1.3565032>.
 - [96] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4), 2023. doi: 10.1214/23-STS894. URL <https://doi.org/10.1214/23-STS894>. [V] arXiv abs (2210.01948, math.ST) fetched 2026-09-03 AND Crossref DOI record 10.1214/23-STS894 fetched (title, authors, Statistical Science 38(4), 2023 matched; the registry record carried no page field, so pages are omitted rather than guessed).
 - [97] Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, Victor Greiff, David Kreil, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. Hopfield networks is all you need. In *International Conference on Learning Representations (ICLR 2021)*, 2021. doi: 10.48550/arXiv.2008.02217. URL <https://arxiv.org/abs/2008.02217>.
 - [98] Yuval Ran-Milo. Attention sinks are provably necessary in softmax transformers: Evidence from trigger-conditional tasks. *arXiv preprint arXiv:2603.11487*, March 2026. doi: 10.48550/arXiv.2603.11487. URL <https://arxiv.org/abs/2603.11487>.
 - [99] Giorgio Roffo, Hazem Abdelkawy, Nilli Lavie, and Luke Palmer. Self-attention and beyond the infinite: Towards linear transformers with infinite self-attention, 2026. URL <https://arxiv.org/abs/2603.00175>. Eq. (7), §3.2: $\check{C} = (I - \gamma A)^{-1} - I$, $\gamma \in (0, 1/\rho(A))$; §3.3 absorbing-chain interpretation with a single absorbing state; A is $\text{ReLU}(QK^\top)$ Frobenius-normalised, not row-stochastic, not causal.
 - [100] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 2020. URL <https://arxiv.org/abs/2003.05997>. [V] abs page fetched 2026-09-03; title matched; venue as printed on the abs page.

- [101] Evan M. Russek, Ida Momennejad, Matthew M. Botvinick, Samuel J. Gershman, and Nathaniel D. Daw. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLOS Computational Biology*, 13(9), 2017. doi: 10.1371/journal.pcbi.1005768. URL <https://doi.org/10.1371/journal.pcbi.1005768>. [V] CrossRef record fetched 2026-09-03.
- [102] Shreyas N. Samaga, Gilberto Gonzalez Arroyo, and Tamal K. Dey. HalluZig: Hallucination detection using zigzag persistence, 2026. URL <https://arxiv.org/abs/2601.01552>. [V] abs page fetched 2026-09-03; title matched.
- [103] Clayton Sanford, Bahare Fatemi, Ethan Hall, Anton Tsitsulin, Mehran Kazemi, Jonathan Halcrow, Bryan Perozzi, and Vahab Mirrokni. Understanding transformer reasoning capabilities via graph algorithms. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024. doi: 10.48550/arXiv.2405.18512. URL <https://arxiv.org/abs/2405.18512>. OpenReview AfzbDw6DSp.
- [104] Clayton Sanford, Daniel Hsu, and Matus Telgarsky. One-layer transformers fail to solve the induction heads task. *arXiv preprint*, 2024. doi: 10.48550/arXiv.2408.14332. URL <https://arxiv.org/abs/2408.14332>. [V] abs fetched 2026-09-03. Theorem 1: a one-layer softmax transformer with h heads, embedding m , p bits solving induction heads over a 3-symbol alphabet at length n has $h \cdot m \cdot p = \Omega(n)$; unconditional, via INDEX communication complexity. Two-layer upper bound $h = O(1)$, $m = O(1)$, $p = O(\log n)$.
- [105] Clayton Sanford, Daniel Hsu, and Matus Telgarsky. Transformers, parallel computation, and logarithmic depth. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, 2024. doi: 10.48550/arXiv.2402.09268. URL <https://arxiv.org/abs/2402.09268>. Theorem 4.2 depth $\lceil \log_2 k \rceil + 2$; Corollary 4.3 $\Omega(\log k)$ conditional on the 1-vs-2-cycle conjecture.
- [106] Meyer Scetbon, Joel Jennings, Agrin Hilmkil, Cheng Zhang, and Chao Ma. A fixed-point approach for causal generative modeling, 2024. URL <https://arxiv.org/abs/2404.06969>. [V] arXiv abs and HTML full text fetched 2026-09-03: Definition 2.3 (fixed-point SCM), Section 2.4 (interventions by clamping), Definition 4.1 (causal attention), Algorithm 2 (d-fold composition).
- [107] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 2020. doi: 10.1038/s41586-020-03051-4. URL <https://arxiv.org/abs/1911.08265>. [V] arXiv abs fetched 2026-09-03; DOI as listed on the abs page.
- [108] Donald J. Schuirmann. A comparison of the Two One-Sided Tests Procedure and the Power Approach for assessing the equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics*, 15(6):657–680, 1987. doi: 10.1007/BF01068419. URL <https://doi.org/10.1007/BF01068419>. [V] Crossref DOI record fetched 2026-09-03 (Springer landing page returned a login redirect).
- [109] Paul J. Schweitzer. Perturbation theory and finite markov chains. *Journal of Applied Probability*, 5(2):401–413, 1968. doi: 10.2307/3212261. URL <https://doi.org/10.2307/3212261>. [V] Crossref record fetched. Owns: perturbation of the fundamental matrix of a finite chain under a change of transition matrix – the field name for "displacement of the resolvent under an intervention on P ".
- [110] Teerth Sharma. Add topology-derived sparse attention kernel. triton-lang/kernels pull request #22, merged 2026-07-28, 2026. URL <https://github.com/triton-lang/>

- kernels/pull/22. [V] PR page and files tab fetched 2026-09-03: state merged; files kernels/topology_sparse_attention.py, test/test_topology_sparse_attention.py, benchmarking/topology_sparse_attention.py (+146), benchmarking/benchmark_utils.py (+8/-1), kernels/___init___py (+10); forward-only Triton kernel consuming a causal CSR block schedule built from sink blocks, local-window blocks and a persistence-style salience over key-block centroids; the PR page states up to 3.48x versus dense CSR on its measured sequences.
- [111] Jack Sherman and Winifred J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1): 124–127, 1950. doi: 10.1214/aoms/1177729893. URL <https://doi.org/10.1214/aoms/1177729893>.
 - [112] Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(72): 2003–2030, 2006. URL <https://www.jmlr.org/papers/v7/shimizu06a.html>. [V] JMLR landing page fetched 2026-09-03; the model is written $x = Bx + e$, solved as $x = (I - B)^{-1}e$.
 - [113] Gurjeet Singh, Facundo Mémoli, and Gunnar Carlsson. Topological methods for the analysis of high dimensional data sets and 3d object recognition. In *Eurographics Symposium on Point-Based Graphics (SPBG 2007)*, pages 91–100, 2007. doi: 10.2312/SPBG/SPBG07/091-100. URL <https://doi.org/10.2312/SPBG/SPBG07/091-100>. [U] the DOI landing page (diglib.eg.org) returned HTTP 403 and Crossref returned 404 for this DOI form on 2026-09-03; title, authors, year and DOI were confirmed against two independent bibliographic records (dblp conf/spbg/SinghMC07 and Semantic Scholar CorpusId 5703368), both fetched. Marked [U] because the primary page itself was not read.
 - [114] Sean Summers and John Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010. doi: 10.1016/j.automatica.2010.08.006. URL <https://doi.org/10.1016/j.automatica.2010.08.006>.
 - [115] Emanuel Todorov. Linearly-solvable markov decision problems. In *Advances in Neural Information Processing Systems 19 (NIPS 2006)*, 2006. URL <https://papers.neurips.cc/paper/3002-linearly-solvable-markov-decision-problems>. [V] NeurIPS proceedings page fetched 2026-09-03.
 - [116] Emanuel Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009. doi: 10.1073/pnas.0710743106. URL <https://doi.org/10.1073/pnas.0710743106>. [V] Crossref DOI record and Europe PMC abstract fetched 2026-09-03 (PMID 19574462). The first-exit linear system with terminal boundary values is standard LMDP material not read from this body this session [U]; see hierarchical-LMDP abs (arXiv:1603.03267) for the first-exit definition.
 - [117] Katharine Turner. Rips filtrations for quasi-metric spaces and asymmetric functions with stability results. *Algebraic & Geometric Topology*, 19:1135–1170, 2019. URL <https://arxiv.org/abs/1608.00365>. [V] abs page fetched 2026-09-03; title matched; journal, volume and pages as printed on the abs page.
 - [118] Donna Vakalis. The intervention gap in latent world models, 2026. URL <https://arxiv.org/abs/2608.29998>. [V] arXiv abs fetched 2026-09-03; submitted 2026-08-30, four days before this sweep.
 - [119] Kristof Van Moffaert, Madalina M. Drugan, and Ann Nowé. Scalarized multi-objective reinforcement learning: Novel design techniques. In *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 191–199, 2013. doi:

- 10.1109/ADPRL.2013.6615007. URL <https://doi.org/10.1109/ADPRL.2013.6615007>. [V] Crossref record fetched. Owns: the Chebyshev (weighted max-norm) scalarisation as an action-selection rule in MORL – the field name for "max over K constraint scores".
- [120] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, June 2017. doi: 10.48550/arXiv.1706.03762. URL <https://arxiv.org/abs/1706.03762>. NeurIPS 2017 [U]; venue not shown on the fetched abs page.
- [121] Petar Veličković, Christos Perivolaropoulos, Federico Barbero, and Razvan Pascanu. Softmax is not enough (for sharp size generalisation). *arXiv preprint*, 2025. doi: 10.48550/arXiv.2410.01104. URL <https://arxiv.org/abs/2410.01104>. [V] abs fetched 2026-09-03 (v3 30 May 2025); comments: to appear at ICML 2025. Softmax circuits must disperse as the number of items grows; adaptive temperature as mitigation.
- [122] Jean Ville. *Étude critique de la notion de collectif*. Thèse, Faculté des Sciences de Paris, Paris, 1939. URL http://www.numdam.org/item/THESE_1939__218__1_0/. [V] numdam record fetched 2026-09-03; title, author, year and page count matched.
- [123] Mark Voortman, Denver Dash, and Marek J. Druzdzel. Learning causal models that make correct manipulation predictions with time series data. In *Proceedings of the Workshop on Causality: Objectives and Assessment at NIPS 2008*, volume 6 of *Proceedings of Machine Learning Research*, pages 257–266, 2010. URL <https://proceedings.mlr.press/v6/voortman10a.html>. [V] PMLR landing page fetched 2026-09-03.
- [124] Guangtao Wang, Rex Ying, Jing Huang, and Jure Leskovec. Multi-hop attention graph neural network. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI 2021)*, 2021. URL <https://arxiv.org/abs/2009.14332>. [V] arXiv abs (IJCAI 2021) and arXiv HTML fetched 2026-09-03; personalised PageRank with $\theta_i = \alpha(1-\alpha)^i$ on the row-softmaxed GAT attention, truncated at $K \in [3, 10]$ hops.
- [125] Hongjian Wang and Aaditya Ramdas. The extended Ville’s inequality for nonintegrable nonnegative supermartingales, 2023. URL <https://arxiv.org/abs/2304.01163>. [V] arXiv abs fetched 2026-09-03; v3 posted 2024-10-08; no journal reference on the page. The record cited this id as [U] (METHODS.md:344); it is now title-matched.
- [126] Jiuqi Wang, Ethan Blaser, Hadi Daneshmand, and Shangdong Zhang. Transformers can learn temporal difference methods for in-context reinforcement learning, 2024. URL <https://arxiv.org/abs/2405.13861>. [V] arXiv abs fetched 2026-09-03 (v4 2025-02-24; no journal reference on the page). Found by widening search S12.
- [127] L. Welch. Lower bounds on the maximum cross correlation of signals. *IEEE Transactions on Information Theory*, 20(3):397–399, 1974. doi: 10.1109/TIT.1974.1055219. URL <https://doi.org/10.1109/TIT.1974.1055219>. [V] Crossref DOI record fetched 2026-09-03 (IEEE landing page returned 403). Owner of the Welch bound (MATHEMATICS.md:812).
- [128] Edwin B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212, 1927. doi: 10.1080/01621459.1927.10502953. URL <https://doi.org/10.1080/01621459.1927.10502953>. [V] Crossref DOI record fetched 2026-09-03 (T&F landing page returned 403).
- [129] Xiao-Ming Wu, Zhenguo Li, Anthony M. So, John Wright, and Shih-Fu Chang. Learning with partially absorbing random walks. In *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, 2012. URL <https://proceedings.neurips.cc/paper/2012/hash/512c5cad6c37edb98ae91c8a76c3a291-Abstract.html>.

- [130] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. In *International Conference on Learning Representations (ICLR 2024)*, 2024. URL <https://arxiv.org/abs/2309.17453>. [V] arXiv abs fetched 2026-09-03 (ICLR 2024 stated).
- [131] Zixuan Xie, Xinyu Liu, Claire Chen, Shuze Daniel Liu, Rohan Chandra, and Shangdong Zhang. Beyond linear attention: softmax transformers implement in-context reinforcement learning, 2026. URL <https://arxiv.org/abs/2605.07333>. [V] arXiv abs fetched 2026-09-03. Found by widening search S12. Claims the layerwise forward pass of a softmax transformer equals iterative weighted-softmax TD updates, i.e. policy evaluation of the resolvent by depth.
- [132] Hsin-Jung Yang, Zhanhong Jiang, Prajwal Koirala, Qisai Liu, Cody Fleming, and Soumik Sarkar. Lexisafe: Offline safe reinforcement learning with lexicographic safety-reward hierarchy. In *ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 2026. URL <https://arxiv.org/abs/2602.17312>. [V] abs fetched. Owns: lexicographic ordering of several safety costs above reward, offline.
- [133] Liu Yang, Kangwook Lee, Robert Nowak, and Dimitris Papailiopoulos. Looped transformers are better at learning learning algorithms. In *International Conference on Learning Representations (ICLR 2024)*, 2024. doi: 10.48550/arXiv.2311.12424. URL <https://arxiv.org/abs/2311.12424>.
- [134] Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. *arXiv preprint arXiv:2312.06635*, December 2023. doi: 10.48550/arXiv.2312.06635. URL <https://arxiv.org/abs/2312.06635>.
- [135] Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated delta networks: Improving Mamba2 with delta rule. *arXiv preprint arXiv:2412.06464*, December 2024. doi: 10.48550/arXiv.2412.06464. URL <https://arxiv.org/abs/2412.06464>. ICLR 2025 per search index [U].
- [136] Songlin Yang, Bailin Wang, Yu Zhang, Yikang Shen, and Yoon Kim. Parallelizing linear transformers with the delta rule over sequence length. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024. doi: 10.48550/arXiv.2406.06484. URL <https://arxiv.org/abs/2406.06484>.
- [137] Songlin Yang, Yikang Shen, Kaiyue Wen, Shawn Tan, Mayank Mishra, Liliang Ren, Rameswar Panda, and Yoon Kim. PaTH attention: Position encoding via accumulating householder transformations. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025. doi: 10.48550/arXiv.2505.16381. URL <https://arxiv.org/abs/2505.16381>.
- [138] Zhibo Yang. Locality does not imply reachability: Boundary repair in block-sparse causal attention, 2026. URL <https://arxiv.org/abs/2606.02680>. [V] abs page fetched 2026-09-03; title matched.
- [139] Gilad Yehudai, Clayton Sanford, Maya Bechler-Speicher, Orr Fischer, Ran Gilad-Bachrach, and Amir Globerson. Depth-width tradeoffs in algorithmic reasoning of graph tasks with transformers. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2503.01805. URL <https://arxiv.org/abs/2503.01805>. [V] abs fetched 2026-09-03 (v3 17 Jun 2026). NeurIPS 2025 spotlight attested by search index only [U]. With width linear in n, constant depth suffices for a host of graph tasks; some tasks need quadratic width. Removes "log depth" as a necessity when width is allowed to grow linearly.

- [140] Chaohao Yuan, Zhenjie Song, Ercan Engin Kuruoglu, Kangfei Zhao, Yang Liu, Deli Zhao, Hong Cheng, and Yu Rong. Paraformer: A generalized pagerank graph transformer for graph representation learning. In *Proceedings of the 19th ACM International Conference on Web Search and Data Mining (WSDM 2026)*, 2026. doi: 10.48550/arXiv.2512.14619. URL <https://arxiv.org/abs/2512.14619>.
- [141] Jingyang Yuan, Huazuo Gao, Damai Dai, Junyu Luo, Liang Zhao, Zhengyan Zhang, Zhenda Xie, Y. X. Wei, Lean Wang, Zhiping Xiao, Yuqing Wang, Chong Ruan, Ming Zhang, Wenfeng Liang, and Wangding Zeng. Native sparse attention: Hardware-aligned and natively trainable sparse attention, 2025. URL <https://arxiv.org/abs/2502.11089>. [V] abs page fetched 2026-09-03; title matched.
- [142] Jiaqi Zhang, Joel Jennings, Agrin Hilmkil, Nick Pawlowski, Cheng Zhang, and Chao Ma. Towards causal foundation model: on duality between causal inference and attention, 2023. URL <https://arxiv.org/abs/2310.00809>. [V] arXiv abs fetched 2026-09-03.
- [143] Hangyue Zhao, Paul Caillon, Erwan Fagnou, and Alexandre Allauzen. Structured-sparse attention for entity tracking with subquadratic sequence complexity, 2026. URL <https://arxiv.org/abs/2605.22476>. [V] arXiv abs and HTML fetched 2026-09-03; defines $\mathcal{S}_\gamma(A) = (1 - \gamma)A(I - \gamma A)^{-1}$ and evaluates it blockwise in $O(n^{4/3}d)$.
- [144] Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *Advances in Neural Information Processing Systems 16 (NIPS 2003)*, 2003. URL <https://papers.nips.cc/paper/2506-learning-with-local-and-global-consistency>. [V] NeurIPS proceedings page fetched 2026-09-03. The closed form $(I - \alpha S)^{-1}Y$ is from memory of the paper’s body, not read this session [U].
- [145] Xiaojin Zhu, Zoubin Ghahramani, and John D. Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the Twentieth International Conference on Machine Learning (ICML 2003)*, pages 912–919, 2003. URL <https://dblp.org/rec/conf/icml/ZhuGL03.html>.