CERT-FLOW: Theory — Coverage Certificates and the Certifiability Threshold

Theorem–proof companion to the CERT-FLOW preprint (engrXiv, DOI 10.31224/7306)

Krishi Attri

July 2026

1 Setting

A directed graph $G = (V, E)$ with start s and goal g. Each edge $e \in E$ has an unknown, time-varying cost $c_e(t) > 0$. Time advances in rounds of length Δ (the sensing period); at most one edge is observed per round. Observing edge e at time u returns

$$Y_e(u) \;=\; c_e(u) + \eta_{e,u},$$

where the noise variables $\eta_{e,u}$ are independent across observations. The planner stores, per edge, the last observation $\hat{c}_e$ taken at time t_e, and the age $a_e(t) = t - t_e$.

Assumption 1 (A1: bounded drift)**.** *For every e and $t' \geq t$: $|c_e(t') - c_e(t)| \leq \rho_e\,(t' - t)$, with ρ_e known or conservatively over-estimated.*

Assumption 2 (A2: TV-Lipschitz noise drift)**.** *Let F_u denote the distribution of $\eta_{e,u}$ (it may vary with time and edge within a terrain class sharing one calibration buffer). For all $u' \geq u$: $d_{\mathrm{TV}}(F_{u'}, F_u) \leq \varepsilon_{\mathrm{TV}}\,(u' - u)$.*

Assumption 3 (A3: symmetric unimodal noise)**.** *Each $\eta_{e,u}$ is symmetric about 0 with a unimodal density. (Gaussian, Laplace, and Student-t all qualify; no moment assumptions are made — Student-t with two degrees of freedom is admissible.)*

Scores and intervals. When edge e is re-observed at time u (previous estimate $\hat{c}_e$ from time t_e, age $a = u - t_e$), the planner records the *drift-adjusted nonconformity score*

$$R \;=\; \big|\, Y_e(u) - \hat{c}_e \,\big| \;-\; \rho_e\, a. \tag{1}$$

Writing $\delta = c_e(u) - c_e(t_e)$ (so $|\delta| \leq \rho_e a$ by A1) and η, η' for the fresh and previous noise, $Y_e(u) - \hat{c}_e = \delta + \eta - \eta'$, hence

$$|\eta - \eta'| - 2\rho_e a \;\leq\; R \;\leq\; |\eta - \eta'|. \tag{2}$$

Scores enter a single rolling calibration buffer $\{(R_i, u_i)\}_{i=1}^{n}$ shared across edges. With *data-independent* age weights $w_i = \rho_w^{\,t-u_i}$ (normalized $\tilde{w}_i = w_i/(1 + \sum_j w_j)$, $\tilde{w}_{n+1} = 1/(1 + \sum_j w_j)$), the planner computes the weighted conformal quantile

$$q_t(\alpha) \;=\; \mathrm{Quantile}_{1-\alpha}\Big(\textstyle\sum_i \tilde{w}_i\, \delta_{R_i} + \tilde{w}_{n+1}\, \delta_{+\infty} \Big),$$

and forms, for margin factor $\lambda \geq 1$ (implementation: `latent_margin`),

$$\ell_e(t) = \hat{c}_e - \lambda q_t - \rho_e a_e(t), \qquad u_e(t) = \hat{c}_e + \lambda q_t + \rho_e a_e(t), \tag{3}$$

clipped below at the cost floor.

The weights w_i depend only on observation *times* (data-independent, as the non-exchangeable guarantee requires); with one observation per round the u_i are deterministic round indices even though *which edge* is observed is adaptive. Adaptivity does, however, affect which score law lands in the buffer; Lemma 1's independence hypothesis prices this via thinning and leave-one-out (Honest accounting, item 1).

Definition 1 (staleness correction)**.** *Let $a_{\max}$ bound every calibration and test age. Then*

$$\Delta_{\text{stale}}(t) \;=\; \sum_{i=1}^{n} \tilde{w}_i \, \min\Big(1,\; 2\big[\varepsilon_{\text{TV}}(t-u_i) + \varepsilon_{\text{TV}}\big((t-u_i)+a_{\max}\big)\big]\Big).$$

Each score involves two *noise draws — the fresh draw at u_i and the previous draw at $u_i - a_i$ — and the second draw's gap to the test score's previous draw (at $t - a_e$) is not bounded by $t - u_i$; the two-term form prices both draws. It supersedes the single-draw form $\sum_i \tilde{w}_i \min(1, 2\varepsilon_{\text{TV}}(t-u_i))$ used by the implementation's reported claim line, which undercounts by at most a factor 2 plus a $2\varepsilon_{\text{TV}} a_{\max}$ term — below reported precision at the swept ε_{TV} and the small calibration ages the sensing loop maintains.*

Definition 2 (calibration-age slack)**.** $\langle \rho a \rangle_{\text{cal}} := \max_i \rho_{e_i} a_i$ *over the calibration buffer, and*

$$\pi_{\text{cal}} \;:=\; \sup_x \, \big[\mathbb{P}(W \leq x) - \mathbb{P}(W \leq x - 2\langle \rho a \rangle_{\text{cal}})\big],$$

the largest mass the score law places on an interval of length $2\langle \rho a \rangle_{\text{cal}}$, where W is either the pure pair score $|\eta - \eta'|$ or the single-draw magnitude $|\eta|$ (both are bounded by $4 f_{\max} \langle \rho a \rangle_{\text{cal}}$ under the bounded-density addition; Lemma 8). It vanishes as calibration re-observations happen at small ages.

2 Per-edge coverage (T1)

Lemma 1 (weighted conformal validity of the score)**.** *Fix edge e and time t, and let R_{n+1} denote the drift-adjusted score* (1) *of a (hypothetical) re-observation of e at time t. If the calibration scores' underlying noise pairs and the test pair are independent draws whose distributions satisfy A2 (the thinned buffer of Honest-accounting item 1, with the test edge's most recent score excluded — leave-one-out — so the test pair shares no draw with the buffer), then*

$$\mathbb{P}\big(R_{n+1} \leq q_t(\alpha)\big) \;\geq\; 1 - \alpha - \Delta_{\text{stale}}(t) - \pi_{\text{cal}}.$$

Proof. The drift-adjusted score R_i is *not* a fixed measurable function of its noise pair: the drift realization δ_i sits inside the absolute value in (1) and varies across calibration scores (different edges, ages, drift paths), so $R_i = f_i(\text{pair})$ with $f_i \neq f_{n+1}$, and TV data-processing cannot be applied to the R-scores directly (a magnitude bound $|\delta| \leq \rho a$ does not bound TV). We therefore pass to the *pure pair scores* $S_i := |\eta - \eta'|_i$, which are one fixed function of the pair for every i. By the sandwich (2), $R_i \geq S_i - 2\langle \rho a \rangle_{\text{cal}}$ for every calibration score and $R_{n+1} \leq S_{n+1}$; the weighted quantile is coordinatewise monotone, so $q_t(\alpha; R\text{-buffer}) \geq q_t(\alpha; S\text{-buffer}) - 2\langle \rho a \rangle_{\text{cal}}$. Hence

$$\mathbb{P}\big(R_{n+1} > q_t(\alpha)\big) \;\leq\; \mathbb{P}\big(S_{n+1} > q_t(\alpha; S) - 2\langle \rho a \rangle_{\text{cal}}\big) \;\leq\; \mathbb{P}\big(S_{n+1} > q_t(\alpha; S)\big) + \pi_{\text{cal}}$$

(Definition 2, with $W = |\eta - \eta'|$). Now apply Barber, Candès, Ramdas and Tibshirani (2023, Theorem 2) to the S-scores: $\mathbb{P}(S_{n+1} > q_t(\alpha; S)) \le \alpha + \sum_i \tilde{w}_i \, d_{\mathrm{TV}}(S(Z), S(Z^i))$, and under independence their Lemma 1 bounds each term by $2\, d_{\mathrm{TV}}(S_i, S_{n+1})$. Each S-score involves two draws: the fresh one at u_i (gap $t - u_i$ to the test's fresh draw) and the previous one at $u_i - a_i$, whose gap to the test score's previous draw (at $t - a_e$) is at most $(t - u_i) + a_{\max}$. A2 plus subadditivity of TV over product measures and data-processing under the common map $|\cdot - \cdot|$ give $d_{\mathrm{TV}}(S_i, S_{n+1}) \le \varepsilon_{\mathrm{TV}}(t - u_i) + \varepsilon_{\mathrm{TV}}\big((t - u_i) + a_{\max}\big)$; with the swap constant 2 and the cap at 1, summing yields exactly $\Delta_{\mathrm{stale}}(t)$ of Definition 1. □

Theorem 1 (T1a: observable coverage, $\lambda = 1$)**.** *Under A1–A2 and the buffer hypothesis of Lemma 1, for any edge e and time t, a re-observation $Y = Y_e(t)$ satisfies*

$$\mathbb{P}\Big(Y \in \big[\hat{c}_e - q_t - \rho_e a_e(t),\ \hat{c}_e + q_t + \rho_e a_e(t)\big]\Big) \ \ge\ 1 - \alpha - \Delta_{\mathrm{stale}}(t) - \pi_{\mathrm{cal}}.$$

The π_{cal} deduction is not optional here: it is the price of the S-score sandwich in Lemma 1's proof, so T1a is not π_{cal}-free.

Proof. The event $R_{n+1} \le q_t$ is exactly $|Y - \hat{c}_e| \le q_t + \rho_e a_e(t)$. Apply Lemma 1. □

Theorem 1 is the standard conformal statement: it certifies the *measurable* quantity. The oracle in our experiments, however, scores coverage of the *latent* cost $c_e(t)$. The gap between the two is one fresh noise draw, and closing it provably costs one more quantile margin.

Lemma 2 (Anderson domination)**.** *Let η be symmetric unimodal (A3) and independent of a random variable D. Then for every $x \ge 0$, $\mathbb{P}(|D + \eta| \le x) \le \mathbb{P}(|\eta| \le x)$, and consequently $|\eta - \eta'| \succeq_{\mathrm{st}} |\eta|$ for independent noise draws η, η'.*

Proof. Anderson's lemma: for a symmetric unimodal density and any constant d, $\mathbb{P}(|d + \eta| \le x) \le \mathbb{P}(|\eta| \le x)$; integrate over the law of D. The stochastic domination follows by conditioning on η'. □

Theorem 2 (T1b: latent-cost coverage, $\lambda = 2$)**.** *Under A1–A3 and the buffer hypothesis of Lemma 1, for any edge e and time t,*

$$\mathbb{P}\Big(c_e(t) \in \big[\hat{c}_e - 2q_t - \rho_e a_e(t),\ \hat{c}_e + 2q_t + \rho_e a_e(t)\big]\Big) \ \ge\ 1 - 2\alpha - 2\Delta_{\mathrm{stale}}(t) - 2\pi_{\mathrm{cal}},$$

with π_{cal} as in Definition 2. The π_{cal} term appears twice*: once through T1a (Step 1) and once through the fresh-noise quantile step (Step 3).*

Proof. Write $D := c_e(t) - \hat{c}_e$ and let η be the fresh noise of a hypothetical re-observation at t, independent of D and — by the leave-one-out convention of Lemma 1 — of the calibration buffer (without leave-one-out, the stored draw η'_e of edge e's most recent calibration score would be shared). Step 1: by Theorem 1, $|D + \eta| \le q_t + \rho_e a_e(t)$ with probability at least $1 - \alpha - \Delta_{\mathrm{stale}} - \pi_{\mathrm{cal}}$. Step 2: by the triangle inequality, $|D| \le |D + \eta| + |\eta|$, so it suffices to bound $\mathbb{P}(|\eta| > q_t)$. Step 3: by (2) each calibration score satisfies $R_i \ge |\eta - \eta'|_i - 2\langle \rho a \rangle_{\mathrm{cal}}$, and by Lemma 2 $|\eta - \eta'| \succeq_{\mathrm{st}} |\eta|$; therefore the $(1 - \alpha)$ weighted quantile of the scores under-shoots the $(1 - \alpha)$ quantile of $|\eta|$ by at most $2\langle \rho a \rangle_{\mathrm{cal}}$, up to the $\alpha + \Delta_{\mathrm{stale}}$ conformal slack on the pure pair scores (the S-score route inside Lemma 1's proof, applied to the event $|\eta| > q_t$ through the score domination). Collecting: $\mathbb{P}(|\eta| > q_t) \le \alpha + \Delta_{\mathrm{stale}} + \pi_{\mathrm{cal}}$. A union bound over Steps 1 and 3 gives $|D| \le 2q_t + \rho_e a_e(t)$ with the stated probability. □

Remark 1 (implementation vs. theorem)**.** *The deployed default is $\lambda = 1$ (Theorem 1 semantics); Tier-0 measures latent coverage* 0.998–1.000 *against claims of* ~ 0.65*, consistent with the slack chain in Theorem 2 being loose in its favor (Anderson domination contributes a $\sqrt{2}$ scale factor for Gaussian-like noise; Bonferroni and worst-case drift widening add more). $\lambda = 2$ (`latent_margin`) is the provable mode; all soundness claims in the paper are stated for it, all empirical tables report both.*

Corollary 1 (path-level certificate)**.** *Let $L_{\max} \le |V|-1$ bound the number of edges of any simple s–g path, let $\mathcal{C} \ni P_{\mathrm{lb}}$ be the candidate set, $E_{\mathcal{C}} = |\bigcup_{P' \in \mathcal{C}} \mathrm{edges}(P')|$, and set $\alpha_{\mathrm{edge}} = \alpha'/(L_{\max} + E_{\mathcal{C}})$. Let $\mathrm{LB} = \sum_{e \in P_{\mathrm{lb}}} \ell_e(t)$ over the ℓ-shortest path and $\mathrm{UB} = \min_{P' \in \mathcal{C}} \sum_{e \in P'} u_e(t)$. Under the assumptions of Theorem 2 with $\lambda = 2$,*

$$\mathbb{P}\big(\mathrm{LB} \le \mathrm{OPT}(t) \le \mathrm{UB}\big) \;\ge\; 1 - 2\alpha' - 2(L_{\max} + E_{\mathcal{C}})\,\big(\Delta_{\mathrm{stale}}(t) + \pi_{\mathrm{cal}}\big),$$

where $\mathrm{OPT}(t)$ is the cost of the true optimal path under $c(\cdot, t)$. With the single candidate $\mathcal{C} = \{P_{\mathrm{lb}}\}$ (i.e. UB $= \sum_{P_{\mathrm{lb}}} u_e$*), $L_{\max} + E_{\mathcal{C}} \le 2L_{\max}$, giving $\alpha_{\mathrm{edge}} = \alpha'/(2L_{\max})$ and level $1 - 2\alpha' - 4L_{\max}(\Delta_{\mathrm{stale}} + \pi_{\mathrm{cal}})$. Two accounting points the earlier statement missed: (i) the lower bound must hold on the* unknown *optimum's edges and the upper bound on the* UB*-attaining candidate's edges — different paths in general, so one $L_{\max}$-edge budget cannot pay both sides (T1b's proof consumes both of its events for either side, so no one-sided bound at half the rate is available); (ii) the* min *over $\mathcal{C}$ is a data-dependent selection (the minimizer is chosen* because *its u-sum is small), so simultaneous upper coverage of* all *candidate edges is required — hence $E_{\mathcal{C}}$, not the minimizer's length. The deployed planner operates at the realized-L level α'/L, exact for the tracked candidate paths and conservative for P^* only through interval slack; this operating claim is validated against ground truth (coverage* 1.000*, every condition). The `strict_lb_alpha` mode implements the $L_{\max}$-level constant (measured: valid* 76.5%*, coverage* 1.000*, claims anneal to* ≈ 0.13*); the theorem-exact two-sided budget above is stricter still — the honest price of airtightness is $\alpha_{\mathrm{edge}} = \alpha'/(2L_{\max})$, not $\alpha'/L_{\max}$.*

Proof. Take the intersection, over the $\le L_{\max} + E_{\mathcal{C}}$ distinct edges of $P^* \cup \bigcup_{\mathcal{C}}$, of the two-sided T1b events at level α_{edge}; each fails with probability at most $2\alpha_{\mathrm{edge}} + 2\Delta_{\mathrm{stale}} + 2\pi_{\mathrm{cal}}$, and the union bound gives the stated level. On this event: $\mathrm{LB} \le \sum_{e \in P^*} \ell_e \le \sum_{e \in P^*} c_e(t) = \mathrm{OPT}(t)$ for the true-optimal P^* — the first inequality because P_{lb} minimizes ℓ-cost — and $\mathrm{OPT}(t) \le \sum_{e \in P'} c_e(t) \le \sum_{e \in P'} u_e$ for *every* candidate P', hence $\mathrm{OPT}(t) \le \mathrm{UB}$. Covering all candidate edges (not only the realized minimizer's) is what makes the data-dependent min harmless. □

Remark 2. *The Bonferroni factor is the dominant practical cost (warm-up consumes* 25–48% *of rounds in Tier-0/ablation runs). Replacing it with a sum-aware nonconformity score à la Luo–Zhou (2024) under non-exchangeable weights is the open stretch problem; nothing in Sections 2–3 depends on which is used.*

3 The certifiability threshold (T2′)

Throughout this section fix a round length Δ, sensing rate one edge per round, margin factor λ, and write $\bar{q}$ for an upper bound on λq_t over the horizon considered and $q_{\min} \ge 0$ for a lower bound (under nondegenerate noise $q_{\min} > 0$ at any fixed α).

Lemma 3 (gap decomposition)**.** *At any round,* $\mathrm{UB} - \mathrm{LB} \le \sum_{e \in P_{\mathrm{lb}}} \big(u_e - \ell_e\big) = \sum_{e \in P_{\mathrm{lb}}} \big(2\lambda q_t + 2\rho_e a_e(t)\big)$ *(before cost-floor clipping, which only shrinks widths).*

Proof. P_{lb} is in the UB candidate set, so $\mathrm{UB} \le \sum_{e\in P_{\mathrm{lb}}} u_e$, while $\mathrm{LB} = \sum_{e\in P_{\mathrm{lb}}} \ell_e$. □

Theorem 3 (T2′a: achievability)**.** *Suppose the ℓ-shortest path P (L edges, drift rates $\le \bar\rho$) is stable over the horizon, and recall $\bar q$ bounds λq_t (λ is absorbed into $\bar q$). The policy that re-observes the edges of P in round-robin order sustains, after L burn-in rounds,*

$$\mathrm{UB} - \mathrm{LB} \;\le\; 2L\bar q \;+\; \bar\rho\,\Delta\,L(L-1) \qquad \textit{at every post-observation instant}$$

(evaluated at the start of a round, before that round's observation, the in-round shift $\Delta_0 \le \Delta$ adds $+2\bar\rho L\Delta_0 \le 2\bar\rho L\Delta$). With lazy pre-widening at horizon B rounds the bound gains $+2\bar\rho\,\Delta BL$. Hence the target gap ε is sustainable whenever $\varepsilon \ge 2L\bar q + \bar\rho\Delta L(L-1)\,[+\,2\bar\rho\Delta BL]$.

Proof. After burn-in, at the post-observation instant of any round the ages of the L edges are exactly a permutation of $\{0, \Delta, \dots, (L-1)\Delta\}$, and summing $2\bar\rho\sum_{j<L} j\Delta = \bar\rho\Delta L(L-1)$. At the start of a round every age carries an additional shift $\Delta_0 \le \Delta$, contributing the displayed $2\bar\rho L\Delta_0$ — a term the bound must carry if quoted at round starts. Apply Lemma 3. The pre-widening term adds $2\rho_e B\Delta$ per edge by construction of the cache (Lemma 5). □

Theorem 4 (T2′b: impossibility for this construction)**.** *Let $C \subseteq E$ be a set of edges with drift rates $\ge \rho_{\min}$ such that every s–g path contains at least m edges of C. Consider the certificate of Section 1 (interval widths $2\lambda q + 2\rho_e a_e$, measured* before *cost-floor clipping, which can shrink $u-\ell$ below that width), and suppose* every *observation — including any initial mapping — is one-per-round (no synchronized free initial map; equivalently, restrict to rounds past the age-separation horizon of any synchronized initialization). Then for every sensing policy and every round after the first,*

$$\mathrm{UB} - \mathrm{LB} \;\ge\; 2m\,q_{\min} \;+\; \rho_{\min}\,\Delta\,\frac{m(m-1)}{2}\cdot 2 \;=\; 2m\,q_{\min} + \rho_{\min}\Delta\,m(m-1).$$

In particular $\varepsilon < 2mq_{\min} + \rho_{\min}\Delta m(m-1)$ is never certifiable by this construction, and $\varepsilon < 2mq_{\min}$ is not certifiable by it even with unbounded sensing rate. The statement bounds the gap of this *certificate family; it is not an information-theoretic impossibility for every conceivable sound certificate.*

Proof. Let P_u attain $\min_P \sum_{e\in P} u_e$. Since $\mathrm{LB} = \min_P \sum_{e\in P}\ell_e \le \sum_{e\in P_u}\ell_e$,

$$\mathrm{UB} - \mathrm{LB} \;\ge\; \sum_{e\in P_u}(u_e - \ell_e) \;\ge\; \sum_{e\in P_u\cap C}\big(2q_{\min} + 2\rho_{\min}a_e\big)$$

(the second inequality uses unclipped widths), and $|P_u \cap C| \ge m$ by assumption. With all observations one-per-round, at most one edge's age resets per round, so among any m edges the j-th smallest age is at least $(j-1)\Delta$ (this is where a synchronized initial map would break the argument: all ages equal at $t=0$); hence $\sum_e a_e \ge \Delta m(m-1)/2$ over the m freshest edges of $P_u \cap C$. Substitute. □

Corollary 2 (regimes)**.** *(i) Static-but-unknown ($\rho \equiv 0$): the ρa width terms vanish, and although q_t itself need not be monotone (a new score can grow the quantile), $q_t \to q_\infty$ almost surely by weighted empirical-quantile convergence for iid scores at a continuity point of the score law; hence the gap along the round-robin path converges and the loop terminates with an ε-certificate after finitely many observations, almost surely, for any $\varepsilon > 2Lq_\infty$. (ii) $\rho \to 0, q \to 0$ recovers the deterministic prove-optimal-or-infeasible stopping rule (Traversing Mars) as a degenerate case. (iii) The deployed greedy policy with the age-triggered backstop inherits Theorem 3: the backstop forces each P_{lb} edge to be re-observed at least once per $\lceil \kappa_{\mathrm{slack}} L\rceil$ rounds, giving the same bound with Δ replaced by $\kappa_{\mathrm{slack}}\Delta$.*

4 The sum-aware upper certificate (T4)

The Bonferroni split $\alpha_{\text{edge}} = \alpha'/L$ in Corollary 1 prices simultaneous control of L edges into every per-edge quantile; its costs are the dominant practical burdens (warm-up $n_0 \approx L/\alpha'$; UB noise floor $\approx L\, q_{\alpha'/L}$). The upper bound, however, only ever concerns *one* path — the incumbent — and for a single path the right object is the distribution of the *sum* of deviations, whose $(1-\alpha')$ quantile scales as $\sqrt{L}$, not L. The lower bound must hold uniformly over all paths (including the unknown optimum) and keeps the per-edge construction; the certificate becomes asymmetric.

Assumption 4 (A4: shared noise family)**.** *Edges feeding one calibration buffer have observation noise drawn from a common family (terrain class); deviations across distinct edges are independent. (The pooled buffer of Section 1 already assumes this implicitly; we make it explicit because T4 leans on it harder.)*

Construction. Alongside the absolute scores (1), record *signed* deviations $D = Y_e(u) - \hat{c}_e$, so $|D - (\eta - \eta')| \le \rho_e a$ by A1. Partition the signed buffer (newest first) into blocks of L consecutive samples; each block contributes the sum $G_b = \sum_{i \in b} D_i$ with weight $w_b = \min_{i \in b} \rho_w^{t-u_i}$ (a data-independent function of ages). Let $M_L(\alpha)$ be the weighted $(1-\alpha)$ quantile of the G_b with test mass at $+\infty$, and define the sum-aware upper bound on a path P with L edges:

$$\mathrm{UB}_{\text{sum}}(P) \;=\; \sum_{e \in P} \hat{c}_e \;+\; \lambda\, M_L(\alpha') \;+\; \sum_{e \in P} \rho_e\, a_e(t).$$

Lemma 4 (block symmetry and domination)**.** *Under A3–A4, sums of L independent symmetric unimodal noises are symmetric unimodal (Wintner: symmetric unimodality is closed under convolution), and for independent copies, $\sum_{i \le L}(\eta_i - \eta_i') \succeq_{\text{st}} \sum_{i \le L} \eta_i'$ by Lemma 2 applied at the sum level. Moreover one-sided tails of the symmetric block sums dominate one-sided tails of $\sum_i \eta_i'$ at the same level.*

Theorem 5 (T4: fixed-path sum-aware upper coverage)**.** *Let P be a path with L edges* chosen independently of the deviations entering UB_{sum}*, and let M_L be formed with each incumbent edge's most recent score excluded from the buffer (leave-one-out: thinning fixes calibration-internal pair sharing, but the test quantity $-\sum_{e \in P} \eta_e'$ is built from the* stored *draws of P's edges — the very draws that entered those edges' most recent scores). Under A1–A4, with $\Delta^{(L)}_{\text{stale}}$ the block-level analogue of Definition 1 (per-block TV $\le$ sum of member terms, by subadditivity over the product, each member term in the corrected two-draw form), and $\lambda = 1$ for the observable / $\lambda = 2$ for the latent statement as in Theorems 1–2:*

$$\mathbb{P}\Big(\textstyle\sum_{e \in P} c_e(t) \;\le\; \mathrm{UB}_{\text{sum}}(P) \Big) \;\ge\; 1 - \lambda\,\alpha' - \lambda\,\Delta^{(L)}_{\text{stale}}(t) - \pi^{(L)}_{\text{cal}},$$

where $\pi^{(L)}_{\text{cal}} := \sup_x [\mathbb{P}(W_L \le x) - \mathbb{P}(W_L \le x - 2\max_b \sum_{i \in b} \rho_i a_i)]$ is the largest mass an interval of length $2\max_b \sum_{i \in b} \rho_i a_i$ receives under the law W_L of the pair-sum $\sum_{i \in b}(\eta - \eta')_i$ — the block-level analogue of Definition 2. The margin satisfies $M_L = \Theta(\sqrt{L})$ for light-tailed noise versus $\Theta(L\, q_{\alpha'/L})$ for the Bonferroni bound.

Proof. $\sum_{e \in P}(c_e(t) - \hat{c}_e) = \sum_e \delta_e - \sum_e \eta_e'$ with $|\sum \delta_e| \le \sum_e \rho_e a_e$ (A1). The blocks G_b bracket independent sums $\sum_{i \in b}(\eta - \eta')_i$ within the block's $\sum \rho a$ slack (2) (the sandwich costs the $\pi^{(L)}_{\text{cal}}$ deduction, exactly as in Lemma 1); by A4 and leave-one-out the test quantity $-\sum_e \eta_e'$ is an independent draw of the dominated-side sum, with distribution drifting from each block's by at

most the block TV (Lemma 1 applied at the block level — blocks are the exchangeable units; disjointness of blocks gives independence across calibration units, and thinning handles the within-edge pair sharing as before). Lemma 4 converts block-sum quantiles into one-sided bounds on $\sum \eta'$; the $\lambda = 2$ latent step repeats Theorem 2's triangle argument at the sum level. We note an unclaimed improvement: since the latent sum contains no fresh noise at all, a single-event route (conformal validity plus domination applied directly to $-\sum \eta'$) is valid at level $1-\alpha'-\Delta^{(L)}_{\text{stale}}-\pi^{(L)}_{\text{cal}}$, i.e. $\lambda = 1$ suffices for the latent *sum*; the displayed $\lambda = 2$ level replays T1b's two-event triangle and is conservative. We keep $\lambda = 2$ for uniformity with Theorem 2. □

Remark 3 (selection bias is real and measurable)**.** *T4 requires P to be chosen independently of the deviations. The planner's incumbent is* not*: it minimizes estimated costs, so its $\hat{c}_e$ are biased low (winner's curse) and the fixed-path guarantee degrades. Measured: in a noise-dominated static regime, naive application to the selected incumbent drops empirical coverage from* 1.000 *(Bonferroni) to* 0.823 *— still above the claimed* 0.65*, but the slack is consumed by an uncontrolled mechanism. The deployed protocol therefore applies T4 only through a* freshness gate*: the sum-aware bound is used only when every edge of the standing incumbent has been re-observed since the path last changed. Observations taken after the selection event are independent of it, so Theorem 5 applies conditionally on the gate. (Two-line formalization: the gate event is itself measurable with respect to pre-refresh data — it asks only whether every incumbent edge has been re-observed since the path change — and on $\{$gate open at $t\}$ every $\hat{c}_e$ entering UB_{sum} postdates the selection, so conditioning on the gate and the selection leaves the post-selection noise laws intact and the fixed-path argument applies verbatim.) Empirically the gate (plus κ-hysteresis, which stabilizes the incumbent and hence opens the gate) recovers coverage to* 0.916–0.966 *while keeping the tightening (gap* −43% *and certified fraction* 95% → 100% *in the low-noise static regime; no effect under strong drift, where age widths dominate and the gate rarely opens — consistent with the $\Theta(\sqrt{L})$ analysis applying to the noise floor only).*

5 Supporting lemmas for the implementation

Lemma 5 (pre-widening soundness)**.** *Fix a refresh time t_0 and horizon $B\Delta$, and suppose the cache invariant $q_{\text{used}} \geq \lambda q_s$ holds for all rounds $s \in [t_0, t_0 + B\Delta]$ (enforced by rebuild-on-growth with headroom). Then the cached metrics $\hat{\ell}_e = \hat{c}_e - q_{\text{used}} - \rho_e(a_e(t_0) + B\Delta)$ and $\hat{u}_e = \hat{c}_e + q_{\text{used}} + \rho_e(a_e(t_0) + B\Delta)$ satisfy $\hat{\ell}_e \leq \ell_e(s)$ and $\hat{u}_e \geq u_e(s)$ for every s in the window. Consequently every certificate computed from cached metrics is valid whenever the per-round certificate is, with gap inflated by at most $2\rho_e B\Delta + 2(q_{\text{used}} - \lambda q_s)$ per edge (the second term is the quantile headroom the cache invariant maintains; it vanishes only when $q_{\text{used}} = \lambda q_s$).*

Proof. For $s \in [t_0, t_0 + B\Delta]$: $a_e(s) = a_e(t_0) + (s - t_0) \leq a_e(t_0) + B\Delta$ (ages only reset on observation, which expires the cache entry). Monotonicity of (3) in q and a gives both inequalities; the gap inflation is the difference of the two substitutions — $2\rho_e B\Delta$ from the age and $2(q_{\text{used}} - \lambda q_s)$ from the quantile headroom. □

Lemma 6 (κ-hysteresis safety)**.** *If the executed incumbent is chosen among candidates whose u-cost is within σ of* UB*, then its true cost satisfies $c($*incumbent$) \leq \text{UB} + \sigma$ on the certificate event, and the reported* (LB, UB) *are unchanged. Hysteresis therefore costs at most σ of certified execution quality and cannot affect coverage.*

Proof. Immediate: the candidate's u-cost upper-bounds its true cost on the coverage event, and the certificate reports min over candidates regardless of which is executed. □

Remark 4 (T3)**.** *Repair cost scaling with the locally affected region is inherited from LPA*/D* Lite (Koenig & Likhachev 2002, 2004); our contribution is only Lemma 5, which restores locality for age-driven metric drift. Empirics:* `docs/results/tier1-latency.md`*.*

6 Limitation closures

Lemma 7 (A1-violation robustness)**.** *Suppose A1 fails on some edge-rounds: for edge e at query time t, let $\nu_e(t) = \mathbb{P}\big(|c_e(t) - c_e(t_e)| > \rho_e a_e(t)\big)$ be the violation probability since the last observation. Then every coverage statement of Sections 2–3 holds with an additional deduction: per-edge coverage loses at most $\nu_e(t)$, and the path certificate loses at most $\sum_e \nu_e(t)$ over the certifying path's edges.*

Proof. Couple the true process with a modified process that satisfies A1 everywhere and coincides with the true one outside the violation event V (replace each violating increment by its $\rho_e a$-clipped version). The coverage statements hold for the modified process by the original arguments, and the two processes — hence their scores, quantiles, and intervals — differ only on V, an event of probability at most the stated violation mass; a union bound adds it. (Conditioning on V^c instead of coupling would distort the score law the conformal step uses, which is why the coupling is the correct formalization.) □

Remark 5 (violations visible to calibration are absorbed)**.** *Lemma 7 is worst-case: it treats violations as invisible. Violations that occur during* re-observed *intervals enter the drift-adjusted scores and inflate the conformal quantile, so calibration absorbs them at width cost. Measured on replayed METR-LA traffic: empirical per-edge A1-violation rates of* 5%*,* 25%*, and* 49% *(drift bound at the $p95/p75/p50$ rate quantiles) all yield path coverage* 1.000 *— far above the lemma's pessimistic deduction — because the violating increments are in-distribution for the score buffer.*

Lemma 8 (explicit π_{cal} bound)**.** *If, in addition to A3, the noise density is bounded by $f_{\max}$, then $\pi_{\text{cal}} \le 4 f_{\max} \langle \rho a \rangle_{\text{cal}}$ for both score laws of Definition 2.*

Proof. π_{cal} is the W-mass of an interval of length $2\langle \rho a \rangle_{\text{cal}}$. For $W = |\eta|$ the density is at most $2f_{\max}$. For $W = |\eta - \eta'|$: the difference $\eta - \eta'$ is a convolution, so its density is also bounded by $f_{\max}$, and the absolute value doubles it — again at most $2f_{\max}$. □

Remark 6 (gated T4 has no uncontrolled residual)**.** *With the freshness gate, rounds where the gate is closed use the Bonferroni upper bound, whose guarantee is unconditional; gate-open rounds are conditionally valid by the post-selection data-splitting argument of Remark 3. No selection-bias mass is left uncontrolled — the earlier "residual" applied only to the ungated construction.*

Remark 7 (online drift estimation)**.** *Estimating ρ online (a pooled quantile of observed rates $|Y - \hat{c}|/a$, which observation noise inflates conservatively) replaces the worst-case linear envelope with the realized drift scale; the conformal layer covers the tail exactly as in the misspecification analysis above. Measured: coverage is unchanged (*1.000*) while gaps tighten* 1.7× *(synthetic) and* 2.4× *(METR-LA) versus the supplied worst-case bound — the estimator automates the drift-aggressiveness dial.*

Remark 8 (the churn floor, diagnosed)**.** *Under drift the realized gap runs above the T2′ floor by a residual factor (≈* 1.6× *after focused sensing, target stabilization, online ρ, and rate feedback). The mechanism is structural: unsensed edges' lower bounds* fall *with age, so optimism attracts P_{lb} to the stalest region of the graph; the sensing rate must outpace this attraction over the whole graph, not the drift on any one path. A sound uniform lower bound cannot ignore stale-cheap regions —*

the optimum could hide there — so the residual is the price of soundness, not an implementation artifact.

Remark 9 (toward a uniform sum-aware lower bound)**.** *The $\sqrt{L}$ tightening of T4 cannot transfer to the lower bound for free: with m edge-disjoint s–g paths the LB must hold simultaneously for m independent sums, forcing a per-path level of order α'/m in the worst case. The best improvement available in principle is therefore block margins at level α'/m over a disjoint-path cover (margin $\sim \sqrt{L}\, q_{\alpha'/m}$) versus Bonferroni's $L\, q_{\alpha'/(mL)}$ — a $\sqrt{L}$-type gain with the union factor intact. Formalizing the cover construction for general graphs is open; the disjoint-paths argument bounds what any construction can achieve.*

7 The lower bound cannot be sum-aware (T5)

T4 tightens the upper certificate to a $\Theta(\sqrt{L})$ margin; the open question was whether the *lower* bound — which must hold uniformly over all paths, including the unknown optimum — admits a comparable construction. The answer is no, by more than a construction failing: no valid uniform lower bound can beat the per-edge union bound by more than logarithmic factors.

Theorem 6 (T5: uniform LB impossibility)**.** *Consider the layered graph with L layers of width w (all w^L layer-paths present), a prior $c_e \sim_{iid} \mathcal{N}(\mu, \sigma^2)$, and one observation $Y_e = c_e + \eta_e$, $\eta_e \sim \mathcal{N}(0, \sigma^2)$ per edge. Fix $\alpha < 1/4$ and any $c < \sqrt{2} - 1$. There is a $w_0(c)$ such that for all $w \geq w_0$: any estimator $\mathrm{LB}(Y)$ with $\mathbb{P}(\mathrm{LB} \leq \mathrm{OPT}) \geq 1 - \alpha$ satisfies*

$$\mathbb{P}\big(\,\widehat{\mathrm{OPT}} - \mathrm{LB} \;\geq\; c\, L\, \sigma \sqrt{\ln w}\,\big) \;\geq\; 1 - \alpha - o(1) \qquad (L \to \infty),$$

and consequently $\mathbb{E}\big[(\widehat{\mathrm{OPT}} - \mathrm{LB})_+\big] \geq (1 - \alpha - o(1))\, c\, L\sigma\sqrt{\ln w}$, where $\widehat{\mathrm{OPT}}$ is the shortest path under the posterior-mean costs. (The positive-part / probability form is essential: an unconstrained expectation bound is gameable — an estimator may spend its α miscoverage budget spiking LB upward on a small set and drive $\mathbb{E}[\widehat{\mathrm{OPT}} - \mathrm{LB}]$ to $-\infty$.) Per-edge Bonferroni achieves slack $O\big(L\sigma\sqrt{\ln(wL/\alpha)}\big)$: the asymmetric certificate (sum-aware UB, per-edge LB) is order-optimal up to a $\sqrt{1 + \ln L / \ln w}$ factor.

Proof. The proof compares the polymer (first-passage) constants of the posterior-mean field and the true-cost field; no greedy-vs-posterior coupling is needed. Unconditionally, the posterior means $m_e := \mathbb{E}[c_e \mid Y] = (\mu + Y_e)/2$ are iid $\mathcal{N}(\mu, \sigma^2/2)$ (since $Y_e \sim \mathcal{N}(\mu, 2\sigma^2)$), and the true costs c_e are iid $\mathcal{N}(\mu, \sigma^2)$.

Polymer constant for an iid $\mathcal{N}(\mu, s^2)$ edge field. Upper bound by the greedy path: layer by layer, the w outgoing edges carry fresh iid values, and the layer minimum has mean $\mu - s\sqrt{2\ln w}\,(1 - o_w(1))$; the sum of L independent layer minima concentrates within $O(s\sqrt{L})$ of its mean, so $\min_P \sum_{e \in P} \xi_e \leq L\mu - Ls\sqrt{2\ln w}\,(1 - o(1))$ with probability $\to 1$. Lower bound by the first moment: there are w^L paths, each with sum $\sim \mathcal{N}(L\mu, Ls^2)$, so for any $\varepsilon > 0$, $\mathbb{P}\big(\min_P \sum_P \xi_e \leq L\mu - Ls\sqrt{2\ln w}(1 + \varepsilon)\big) \leq w^L \exp\big(-L \ln w\,(1+\varepsilon)^2\big) \to 0$. Hence with probability $\to 1$ the minimum is $L\mu - Ls\sqrt{2\ln w}\,(1 + o(1))$.

Comparison. Applying this with $s = \sigma/\sqrt{2}$ to the posterior-mean field and with $s = \sigma$ to the true costs:

$$\widehat{\mathrm{OPT}} = L\mu - L\sigma\sqrt{\ln w}\,(1 + o(1)), \qquad \mathrm{OPT} = L\mu - L\sigma\sqrt{2\ln w}\,(1 + o(1)),$$

each with probability $\to 1$, so $\widehat{\mathrm{OPT}} - \mathrm{OPT} = (\sqrt{2} - 1)\, L\sigma\sqrt{\ln w}\,(1 + o(1))$ with probability $\to 1$; the $w_0(c)$ threshold absorbs the $o_w(1)$ terms in the layer-minimum constant (a universal constant valid down to $w = 2$ is not claimed).

Conclusion. Any estimator with $\mathbb{P}(\mathrm{LB} \le \mathrm{OPT}) \ge 1-\alpha$ satisfies, on the intersection of the coverage event and the two polymer events, $\widehat{\mathrm{OPT}} - \mathrm{LB} \ge \widehat{\mathrm{OPT}} - \mathrm{OPT} \ge (\sqrt{2}-1-o(1))\, L\sigma\sqrt{\ln w}$; the intersection has probability $\ge 1-\alpha-o(1)$. The positive-part expectation bound follows from Markov's inequality in reverse (restrict the expectation to the event). The matching upper bound is Corollary 1 with $\alpha_{\mathrm{edge}} = \alpha/|E|$, $|E| \le 2Lw^2$. □

Remark 10 (what the simulation measured)**.** *The previously reported simulation (median deficit* 0.71–$0.80 \times L\sigma\sqrt{\ln w}$ *for* $w \in [10, 50]$*) measures the* greedy D-deficit $-\sum_{P_g} D_e$ *used by an earlier proof sketch — a quantity that conflates the two polymer constants. The quantity the repaired proof controls is* $\widehat{\mathrm{OPT}} - \mathrm{OPT}$*, whose constant is* $\sqrt{2}-1 \approx 0.41$*; the old number is retained only as a check on the per-layer greedy constant, not as an estimate of* c_0*.*

Remark 11. *The mechanism is selection: with exponentially many candidate paths, some path's plausible downside is* linearly *deep, and a sound LB must respect it. The upper side escapes because it prices one chosen path; the lower side cannot, because the optimum chooses adversarially. This finally explains the asymmetry of the certificate as a theorem rather than a limitation.*

8 Decision-uniform certificates (T6)

The coverage statements of Sections 2–3 are *per-round marginal*: each round's interval contains the optimum with the stated probability, but a policy that acts whenever the certificate reads "certified" effectively selects rounds, and across a T-round trajectory the probability that *some acted-on* certificate failed can approach $T\alpha'$.

Remark 12 (per-round time-uniformity is impractical, quantifiably)**.** *The standard repair — replace the conformal quantile with a time-uniform quantile confidence sequence — costs an anytime inflation. A stitched construction (DKW at level* $\delta_k = \delta/k(k{+}1)$ *on doubling epochs* $n \in [2^k, 2^{k+1})$*; union over* k*) gives the time-uniform band*

$$\sup_x |\hat{F}_n(x) - F(x)| \;\le\; \varepsilon_n \;=\; \sqrt{\tfrac{\ln(2/\delta)+2\ln(\log_2 2n+1)}{2n}} \qquad \textit{simultaneously for all } n,$$

so the time-uniform quantile sits at level $\alpha_e - \varepsilon_n$*, which is positive only when* $n \gtrsim \varepsilon^{-2}$*. (Strictly, DKW is a fixed-n bound; uniformity* within *each doubling epoch needs a maximal-inequality version, which inflates the constants slightly and leaves the order-of-magnitude conclusion unchanged.) At Bonferroni levels (*$\alpha_e \approx 0.011$ *for* $L \approx 18$*,* $\alpha' = 0.2$*) this requires* $n \gtrsim 63{,}000$ *calibration scores — two orders of magnitude beyond any realistic rolling buffer. Per-round uniformity is not where the budget should go.*

Theorem 7 (T6: decision-uniform validity for predictable schedules)**.** *Call a round a* decision instant *when the certificate is acted on (sensing stops; the robot departs; an* ε*-claim is reported to an operator). Suppose the decision indicators* $A_r \in \{0,1\}$ *are* predictable from data independent of the conformal calibration buffer *— fixed before deployment, or driven by an auxiliary stream (mission clock, operator requests) independent of the scores — with* $\sum_r A_r \le N_{\mathrm{dec}}$ *almost surely, and every certificate is constructed at claim level* α'/N_{dec} *(implementation:* `decision_uniform`*). Then, under the assumptions of the corresponding per-round statement,*

$$\mathbb{P}\big(\textit{every acted-on certificate in the mission is valid}\big) \;\ge\; 1-\alpha'-\textstyle\sum \Delta_{\mathrm{stale}}\textit{-terms over the acted-on rounds},$$

i.e. trajectory-level validity where the trajectory consumes it, for schedules that do not themselves read the certificate.

Proof. Let I_r be the invalidity event of round r's certificate. Then $\mathbb{P}(\exists r : A_r I_r) \le \mathbb{E}[\sum_r A_r \mathbf{1}\{I_r\}] = \sum_r \mathbb{E}[A_r]\,\mathbb{P}(I_r \mid A_r = 1)$. By the independence hypothesis $\mathbb{P}(I_r \mid A_r = 1) = \mathbb{P}(I_r) \le \alpha'/N_{\mathrm{dec}} + (\Delta_{\mathrm{stale}}\text{-terms})_r$, so the sum is at most $(\alpha'/N_{\mathrm{dec}})\,\mathbb{E}[\sum_r A_r] + \sum_{\mathrm{acted}}(\Delta\text{-terms}) \le \alpha' + \sum_{\mathrm{acted}}(\Delta\text{-terms})$. The independence step is where the hypothesis earns its keep: without it, $\mathbb{P}(I_r \mid A_r)$ is not controlled by the marginal theorem. □

Remark 13 (what is *not* guaranteed: certificate-triggered acting)**.** *The deployed semantics — act whenever the certificate reads "certified" (*$\mathrm{UB}-\mathrm{LB} \le \varepsilon$*) — violates the predictability hypothesis: acting then selects rounds* by *the certificate, and the selection is positively correlated with miscoverage (a certified read requires a small conformal quantile and small ages, and conditionally on the buffer the miscoverage probability is larger exactly when the quantile is smaller). Weighted split conformal provides no conditional-on-buffer validity, so the union bound of Theorem 7 does not compose for such schedules, and the worst-case failure over a T-round mission scales like* $\min(1, T\alpha'/N_{\mathrm{dec}})$ *rather than α': a mission acting only on atypically-small-score buffers concentrates its decisions on the high-miscoverage rounds. Honest routes for certificate-triggered acting are (a) horizon spending (α'/T, which surrenders the width advantage), (b) the anytime-valid betting/e-process machinery of Theorem 12, which is designed for data-dependent stopping and is the natural closure (not yet carried out here), or (c) an explicit assumption $\mathbb{P}(\textit{invalid} \mid \textit{acted, buffer}) \le \alpha'/N_{\mathrm{dec}}$, which makes the gap visible rather than closing it. Theorem 7 guarantees nothing for certificate-triggered schedules.*

Remark 14. *The width price is a quantile at $\alpha'/(LN_{\mathrm{dec}})$ instead of α'/L — supportable when the effective sample size exceeds $LN_{\mathrm{dec}}/\alpha'$ (e.g.* 250 *at L=10, N_{dec}=5, α'=0.2), versus the* $63k$ *of the per-round-uniform route. Rounds between decisions remain monitored at the marginal level; only consumption pays the spending factor.*

9 The churn-measured floor (T7)

Under drift the realized gap of the deployed planner exceeded the T2′ floor by a residual factor; the diagnosis (optimism attracts P_{lb} to the stalest region) is now quantified by the *churn set* $\mathcal{K}(t) = \bigcup_{\tau \in [t-W,t]} \mathrm{edges}(P_{\mathrm{lb}}(\tau))$, $K = |\mathcal{K}|$, tracked online over a sliding window.

Theorem 8 (T7: churn-aware sustainability)**.** *Assume the stationary regime, made precise as: the sliding window includes the current round (so $P_{\mathrm{lb}}(t) \subseteq \mathcal{K}(t)$) and $\mathcal{K}$ is stable over one full rotation. If sensing rotates over a set containing $\mathcal{K}$ at rate k per round, then — at post-observation instants, as in Theorem 3 — every edge of every $P_{\mathrm{lb}}(t)$ has age at most $\lceil K/k \rceil\,\Delta$, and*

$$\mathrm{UB} - \mathrm{LB} \;\le\; 2L\lambda\bar{q} \;+\; 2\bar{\rho}\,\Delta\,L\,\lceil K/k \rceil,$$

i.e. T2′ with K in place of the instantaneous path length (the $L\times$ max-age form is coarser than T2′'s exact age sum, but valid). The floor reported by the planner (and the adaptive rate k) now uses the realized $\widehat{K}$, making attainability declarations honest under churn.

Remark 15 (measured resolution; rotation refuted)**.** *Empirically the better policy is not to chase the churn set but to* suppress *it: focused sensing on a hysteresis-stabilized path keeps $\widehat{K} \approx L$ (measured: K:* $59 \to 11$ *at $\rho = 0.05$), recovering the original floor, while rotating over the full churn set spreads observations thin (same certification,* +20% *sensing). The deployed policy therefore senses the focused path, tracks $\widehat{K}$, and feeds it to the floor and the rate only. Certification in the test regime improved* 5.6% $\to$ 36.7% *across the sequence of churn-directed changes, at coverage* 1.000 *throughout.*

10 Non-exchangeable round two: sum-level pricing and observability (T8)

Corollary 1 prices a path by an α'/L union over its edges. Two constructions replace the union with a single level-α' quantile of a scalar, and both survive the age weights of Lemma 1.

LP-shift staleness (optional). The TV-Lipschitz correction Δ_{stale} (Definition 1) can be replaced by a Lévy–Prokhorov worst-case quantile (Aolaritei et al., 2025): under an LP ambiguity set of local radius ε and global mass ρ around the calibration law, the worst-case $(1-\alpha)$ quantile is $\text{Quant}_{1-\alpha+\rho}(\cdot)+\varepsilon$ and the worst-case coverage of a threshold q is $F(q-\varepsilon)-\rho$. This raises the *quantile level* rather than deducting from the *claim level*, so it stays above the effective-sample-size annealing floor while additionally pricing edges never observed in the buffer (the ρ mass). It is offered as an alternative staleness model; A2/Δ_{stale} remains the default.

Theorem 9 (T8a: group-sum upper certificate)**.** *Fix a path P with L edges chosen independently of the calibration deviations. Let $\{G_j\}$ be group-sum scores formed by drawing, for each calibration group j, a* pairwise-disjoint *L-subset of signed per-edge deviations and summing them, with data-independent age weights. Let $M(\alpha')$ be the weighted $(1-\alpha')$ quantile of $\{G_j\}\cup\{+\infty\}$. Then, under A1–A4 with the λ-margin convention of Theorems 1–2,*

$$\mathbb{P}\Big(\textstyle\sum_{e\in P} c_e(t) \le \sum_{e\in P}\hat{c}_e + \lambda M(\alpha') + \sum_{e\in P}\rho_e a_e(t)\Big) \ge 1-\lambda\alpha'-\lambda\Delta^{(L)}_{\text{stale}}(t)-\pi^{(L)}_{\text{cal}},$$

with $M(\alpha')=\Theta(\sqrt{L})$ for light-tailed noise.

Proof. Identical in structure to Theorem 5: the disjoint group sum is the calibration unit (disjointness gives independence across units, exactly as T4's blocks), weighted-conformal validity (Lemma 1) applies at the group level, and the $\lambda=2$ latent step repeats Theorem 2's triangle argument at the sum. This is the drift-retrofitted, age-weighted form of the symmetric-calibration group sum of Luo–Zhou (2024). For *overlapping* groups the implementation exposes a slack δ (the maximum pairwise edge-sharing across groups, a count made explicit by symmetric calibration); δ is not a probability and no coverage deduction has been derived for it, so the overlapping variant is reported as a heuristic and is not covered by this theorem — the measured width numbers use it only where disclosed. □

Theorem 10 (T8b: joint per-edge (max-score) pricing)**.** *Fix a path P with L edges chosen independently of the calibration deviations (as in T4/T8a — the incumbent is optimizer-selected, so the freshness-gate discussion of Remark 3 applies verbatim here). Let Q be the weighted $(1-\alpha')$ quantile of the per-block* maximum *signed score over disjoint length-L blocks of the buffer, with test point $+\infty$, and let the test scores s_e be the drift-adjusted magnitudes evaluated at the query time t (the same post-observation convention as T1). Then $\bigcap_{e\in P}\{s_e\le Q\}=\{\max_e s_e\le Q\}$ holds with probability $\ge 1-\alpha'-\Delta^{(L)}_{\text{stale}}(t)$ under block exchangeability of the thinned buffer (the age-weighted variant carries the same staleness deduction as every other weighted-conformal statement in this paper; a π-type term enters if the drift adjustment is sandwiched as in Lemma 1), so $Q+\rho_e a_e \ge |Y_e-\hat{c}_e|$ simultaneously for all $e\in P$ at that level — the same per-edge magnitude Bonferroni certifies, calibrated jointly rather than by the α'/L union.*

Proof. The set identity is immediate; validity is the module's own split-conformal quantile with the $\cup\{+\infty\}$ test point, so soundness rests on no max-score-specific constant, and the age-weighted variant inherits Lemma 1 — including its staleness and calibration-age deductions, which the

displayed level now carries (Jung, 2026). Block exchangeability (not per-edge exchangeability) is the standing assumption, and the max-of-L needs $\sim L\times$ the samples per block. □

Remark 16 (which functional wins is empirical, and two-signed)**.** *T8a and T8b both remove the union factor, but their* width *differs by which functional of the block is calibrated. On long real-traffic paths the buffer holds few length-L blocks, so the block-*max *quantile of T8b is coarse and* loses *to Bonferroni (+24.7%/+25.1% measured), while the sum functionals win — −23.6% for the T4 block quantile and −26.6% for the group-sum of T8a. The theorems are soundness statements; the width verdict is in the main text (width experiment). Reported, not hidden: under* independent *edges Bonferroni is already tight and neither helps; the joint price pays only under positive edge correlation (measured* 16.5% *narrower at correlation* 0.9*, L*=20*).*

Observability (the pinned coverage, made testable). The certificate covers at 1.000 in every measured condition, which makes the claim sound but its slack unobservable. A weighted conformal test martingale (Prinster et al., 2025) bets on the weighted conformal p-values of the score stream; its wealth M_t is a nonnegative supermartingale under the weighted-exchangeability null, so Ville gives $\mathbb{P}(\sup_t M_t \geq 1/\delta) \leq \delta$ — a false-alarm-controlled *validity monitor*. Because a plain martingale can random-walk toward zero over a long null and miss a late change, we pair it with a Shiryaev–Roberts statistic $R_t = (1 + R_{t-1})e_t$ ($\mathbb{E}[R_t] = t$ under the null), which restarts implicitly and detects post-null shifts. Conformal e-values (Gauthier et al., 2025) provide the same evidence in mergeable form ($\mathbb{E}[E] \leq 1$ under the null; the average merge is valid under arbitrary dependence, the safe cross-edge combiner). None of this changes the certificate — the $(\mathrm{LB}, \mathrm{UB}, \mathrm{confidence})$ stream is byte-identical with monitoring on or off — so the observability is free.

11 The two-tier certificate: a-priori shrink is impossible (T9)

The observability layer suggests a tempting move: shrink the radius while the alarm is quiet. We show the a-priori version of this is impossible under drift, and give the a-posteriori version that is sound.

Theorem 11 (T9: no a-priori shrink from windowed evidence under drift)**.** *Let the radius be shrunk to $k(q_t + \rho a)$ with $k < 1$ chosen as any measurable function of the observed score window $\{R_i\}_{i \leq t}$. Then for every $\alpha < 1/4$ there is an environment in the drift class A1–A2 (indeed a pair of environments indistinguishable on the score stream) on which the shrunk per-edge interval's next-round observable miscoverage exceeds α on every round where the certificate radius is finite: no such rule retains a distribution-free next-round coverage guarantee $\geq 1 - \alpha$ over A1–A2. (A3 is not assumed here, which is what the construction exploits; see the closing remark.)*

Proof. We exhibit the adversary explicitly. Single edge e, re-observed every round (age $a = \Delta$ at every query). Noise: iid two-point, $\eta = \pm B$ with probability $1/2$ each, $B > \Delta\rho$ — symmetric, time-invariant (so A2 holds with any $\varepsilon_{\mathrm{TV}}$), but not unimodal (A3 is not assumed by this theorem). Drift at the A1 boundary with hidden direction: $c_e(t) = c_0 + s\rho t$ with $s \in \{+1, -1\}$ fixed but unknown; A1 holds with equality. Write $g := \eta - \eta' \in \{0, \pm 2B\}$ with probabilities $\{1/2, 1/4, 1/4\}$.

The window is uninformative. Each score is $R = |s\rho\Delta + g| - \rho\Delta$, taking values 0 (at $g = 0$), $2B$ (at $g = s \cdot 2B$), and $2B - 2\rho\Delta$ (at $g = -s \cdot 2B$), with probabilities $1/2, 1/4, 1/4$ — *the same law for $s = +1$ and $s = -1$* (the drift adjustment removes exactly $\rho\Delta$ and the noise is symmetric). The score stream is therefore iid with a distribution independent of s: any measurable rule $k(\{R_i\}_{i \leq t})$ has the same law in both environments, and no rule can even detect the drift direction.

Any shrink misses an atom. Whenever the certificate radius is finite — i.e. the weighted quantile does not land in the $+\infty$ test atom — $q_t \le \max_i R_i \le 2B$ deterministically. The next-round observable deviation is $|Y_{t+1} - \hat{c}_e| = |s\rho\Delta + g|$, which equals $2B + \rho\Delta$ exactly when $g = s \cdot 2B$: an atom of probability $1/4$ sitting exactly at the un-shrunk radius's maximal value. Since $k < 1$ and $q_t \le 2B$,

$$k\,(q_t + \rho\Delta) \;<\; q_t + \rho\Delta \;\le\; 2B + \rho\Delta,$$

so the shrunk interval excludes the atom and mis-covers with probability at least $1/4 > \alpha$, on every finite-radius round, in *both* environments. (If the radius is infinite the interval is vacuous and no shrink exists to license.) The un-shrunk interval, by contrast, covers the atom exactly when $q_t = 2B$ — which the weighted quantile delivers once the buffer's empirical mass at the top atom plus the test mass $\tilde{w}_{n+1}$ exceeds α, as happens eventually, almost surely, at any $\alpha < 1/4$. □

Remark 17 (scope of the construction, and detection delay)**.** *The two-point adversary places noise mass* exactly *at the un-shrunk radius; under a bounded-density addition (A3 plus $f_{\max}$), the same construction with mass concentrated just inside the radius degrades any shrink to $k \le 1 - \gamma$ by a γ-dependent constant instead — the impossibility is quantitative there, absolute here. The detection-based patch fares no better, by the standard change-point trade-off (an alarm with bounded false-alarm rate has unbounded worst-case detection delay — folklore we use qualitatively, not as a proved ingredient): the pre-alarm rounds are exactly the rounds a coverage consumer needed guaranteed. The CIA collapse (*$0.95 \to 0.20$ *under staleness, measured) is the empirical face of the same refusal.*

Theorem 12 (a-posteriori licensed radius)**.** *Let $x_i(k) = \mathbf{1}\{|Y_i - \hat{c}_i| > k(q + \rho a_i)\}$ be the shrunk-interval violation indicator on the fresh score stream, and let $\mathcal{K}$ be a finite grid of candidate shrink factors fixed in advance. For each $k \in \mathcal{K}$, a betting confidence sequence (Waudby-Smith–Ramdas, 2024) at level $\delta/|\mathcal{K}|$ yields a time-uniform upper bound* $\mathrm{UCB}_t(k)$ *on the running average of the* conditional *violation rates $\bar{\mu}_t(k) := \frac{1}{t}\sum_{i \le t} \mathbb{E}[x_i(k) \mid \mathcal{F}_{i-1}]$ (the WSR sequence controls conditional means of a bounded adapted stream, not the raw empirical mean): $\mathbb{P}(\exists t : \mathrm{UCB}_t(k) < \bar{\mu}_t(k)) \le \delta/|\mathcal{K}|$, and by a union bound over the grid the guarantee holds simultaneously for all $k \in \mathcal{K}$ at level δ. Hence at any t the licensed radius $k^\star(q + \rho a)$ with $k^\star = \max\{k \in \mathcal{K} : \mathrm{UCB}_t(k) \le \alpha'\}$ carries the honest, anytime-valid, self-revoking statement: "over this deployment so far, the shrunk interval's conditional mis-coverage averaged $\le \alpha'$ with $1 - \delta$ validity."*

Remark 18 (the two tiers are different objects, and the labels say so)**.** *Tier-1 is the a-priori distribution-free* $[\mathrm{LB}, \mathrm{UB}]$ *(the product; never shrunk, bit-identical with the license enabled). Tier-2 is $k^\star(q + \rho a)$ with the a-posteriori claim of Theorem 12. Measured on real METR-LA:* -62.4% *width at* 0.51% *shadow miscoverage against true* OPT *(target $\alpha' = 0.20$; the license floor $k = 0.5$ binds on* 82% *of rounds). Deployment reading: safety gates read Tier-1; resource allocation may read Tier-2.*

12 Team and multi-agent certificates (T10)

Theorem 13 (T10a: additive team certificate)**.** *Let N agents share one drifting-cost graph and* one *conformal edge-price store, so every edge's age a_e is global. Let each agent's per-round certificate $(\mathrm{LB}_i, \mathrm{UB}_i, \mathrm{conf}_i)$ be produced from that shared store. Then*

$$\mathrm{LB}_{\mathrm{team}} := \textstyle\sum_i \mathrm{LB}_i \;\le\; \mathrm{OPT}_{\mathrm{team}} \;\le\; \sum_i \mathrm{UB}_i =: \mathrm{UB}_{\mathrm{team}}$$

with probability $\ge \mathrm{conf}_{\mathrm{team}} := \max(0, 1 - \sum_i (1 - \mathrm{conf}_i))$, where $\mathrm{OPT}_{\mathrm{team}} = \sum_i \mathrm{OPT}_i$ is the separable team optimum. Exactness caveat: what is exact is the decomposition $\mathrm{OPT}_{\mathrm{team}} = \sum_i \mathrm{OPT}_i$ — *a*

modeling identity holding by definition of the uncoupled objective, not a probabilistic claim; each LB_i *itself remains strictly conservative.*

Proof. Each summand is sound over the shared store by the single-agent guarantee (Corollary 1, whose repaired constant conf_i inherits); the age a_e is global, so no agent sees a staler edge than the store records. The team objective separates, $\mathrm{OPT}_{\mathrm{team}} = \sum_i \mathrm{OPT}_i$, so summing the per-agent inequalities gives the bracket; the confidence follows from a union bound over the N miscoverage events (valid under arbitrary dependence, including the shared store), floored at 0. □

Remark 19 (why only the additive bound)**.** *A congestion-coupled joint certificate (*$c_e(m) = \phi_e(1+\beta_e m)$*) with a decision-focused shared-sensing allocator is tighter on synthetic bottlenecks (gap ratio additive/joint up to* 1.78 *at* N=8*) but* looser *on the real METR-LA network (ratio* 0.90–0.91 *— additive* $\approx 10\%$ *tighter), because route diversity lets agents avoid the few congested edges. The joint model is falsified as a deployable object and not shipped; the additive bound is the survivor. Separately, age-binning the per-edge widths tightens the additive gap* 2.1–2.4× *at coverage* 1.000*, while a joint block-conformal team quantile over-shoots coverage (*0.86 *at* $N \geq 2$*, minimum* 0.68*) by the same selection bias that gates T4 — so the joint quantile is a diagnostic, never the operating point.*

Certified MAPF (the corridor lift). Lift the certificate to N agents whose realized edge *durations* drift. Each agent's low-level state is (vertex, certified arrival window $[lo, hi]$); the certified duration window of edge e used no later than hi is $[\hat{\ell}_e, \hat{u}_e]$, priced at the *latest* certified use (staleness widening is monotone in age, so the window at hi contains the window at any earlier use). The per-side conventions matter for soundness and are pinned here: realized durations are integers; $\hat{u}_e = \lceil \kappa(\hat{c}_e + q + \rho_e(hi - t_{\mathrm{obs}})) \rceil$ with $\kappa \geq 1$ (rounding *up*; $\hat{u}$ is never clipped downward — a horizon cap triggers abstention, not truncation); $\hat{\ell}_e = \lfloor \hat{c}_e - q - \rho_e(hi - t_{\mathrm{obs}}) \rfloor$ with the *unscaled* radius (applying κ to the lower side would raise it — anti-conservative), rounded *down* and clipped only below at the cost floor (sound for positive latent durations). Conflict-based search branches on overlapping certified windows.

Theorem 14 (T10b: certified-MAPF soundness)**.** *Assume the per-side conventions above and* single-visit *plans (each agent traverses any edge at most once, so its use time* τ *is determined by its predecessor edges and is independent of that edge's noise; revisits would break this). Let the priced universe be all agent–edge pairs* $\mathcal{U} = \{1, \ldots, N\} \times E$ *— fixed by the instance,* not *by the returned plan — and set* $\alpha_{\mathrm{edge}} = \alpha_{\mathrm{team}}/(N|E|)$*. Let* $\mathcal{E}$ *be the event that for every* $(i, e) \in \mathcal{U}$*, the realized duration of* e *at any plan-measurable use time* $\tau \leq hi$ *deviates from* $\hat{c}_e$ *by at most* $q + \rho_e(\tau - t_{\mathrm{obs}})$ *(which places it inside the certified window priced at* hi*, both sides). Then* $\mathbb{P}(\mathcal{E}) \geq 1 - \alpha_{\mathrm{team}} - \sum_{\mathcal{U}}(\Delta_{\mathrm{stale}} + \pi_{\mathrm{cal}})$*-terms and, on* $\mathcal{E}$*: (C1) the realized execution of the returned plan* Π *is collision-free; (C2)* $\sum_i \mathrm{LB}_i \leq \mathrm{OPT}_{\mathrm{team}} \leq \mathrm{cost}(\Pi) \leq \sum_i \mathrm{UB}_i$*; and (C3) if no jointly certified conflict-free plan is found within budget the planner returns* ABSTAIN*.*

Proof. $\mathbb{P}(\mathcal{E}^c) \leq \sum_{(i,e)\in\mathcal{U}} \big(\alpha_{\mathrm{edge}} + (\Delta_{\mathrm{stale}} + \pi_{\mathrm{cal}})\text{-terms}\big)$ by Lemma 1 per pair and a union bound; per-pair validity at the random use time τ uses the single-visit hypothesis ($\tau \perp$ edge-e noise). The union bound runs over a universe fixed *before* plan selection — this repairs the earlier budget over the returned plan's $\leq N\bar{L}$ priced durations, which is a data-dependent set (CBS selects Π to minimize certified cost, favouring low-$\hat{u}$ corridors: exactly the winner's curse this paper measures and gates for T4, under which summing per-edge marginal miscoverage over the selected set does not bound $\mathbb{P}(\mathcal{E}^c)$). Two selection-free alternatives with tighter budgets: re-observe every plan edge after selection and before commitment (a freshness-gate analogue), or restrict to $\mathcal{U}$ = the search-explored

priced set when that set is provably plan-independent. On $\mathcal{E}$, induction over each agent's path (base window $[t_0, t_0]$; a traverse adds a covered $[\hat{\ell}, \hat{u}]$ — the deviation bound at τ, monotone widening to hi, and the ceil/floor conventions give containment; a sync-wait departs at a deterministic $T \geq$ the certified latest arrival) shows every realized occupancy lies inside its certified window; CBS returns only nodes whose certified windows are pairwise disjoint, so realized occupancies cannot intersect (C1). Realized cost $\leq \sum_i \mathrm{UB}_i$ edge-wise on $\mathcal{E}$. For (C2)'s left inequality, any conflict-free plan costs at least $\sum_i$ (agent i's *unconstrained* ℓ-shortest cost) $= \mathrm{LB}_{\mathrm{team}}$, which requires $\hat{\ell}_e \leq$ realized duration on the edges of each agent's unconstrained realized-optimal path — edges generally *not* in Π; because $\mathcal{E}$ covers every pair in $\mathcal{U}$, these lower events are on-event (the earlier plan-only $\mathcal{E}$ did not contain them). Budget exhaustion triggers (C3). □

Remark 20 (honest scope: soundness yes, completeness no, α inert, and the sub-floor label)**.** *Window-CBS* completeness *and* optimality *are not claimed: interval branching can exclude jointly-feasible schedules (the continuous-CBS subtlety), which costs success rate, never soundness. Measured at P0 scale (600 runs), the certificate realizes* 0 *collisions everywhere against* 0–100% *for point-estimate MAPF, but the per-edge budget sits below the* finite-sample support floor*: when* $\alpha_{\mathrm{edge}} < \tilde{w}_{n+1} \approx 1/(n_{\mathrm{eff}} + 1)$ *the* $(1 - \alpha_{\mathrm{edge}})$ *weighted quantile lands in the* $+\infty$ *test atom — the licensed window is infinite — and the implementation's max-score fallback licenses only per-edge miscoverage* $\approx \tilde{w}_{n+1}$*, not* α_{edge}*. At the P0 numbers (*$\alpha_{\mathrm{team}} = 0.1$*,* $N\bar{L} = 192$ *priced plan durations,* $\sim$672 *age-decayed scores):* $\alpha_{\mathrm{edge}} \approx 5.2 \times 10^{-4}$ *against a supportable* $\approx 1.5 \times 10^{-3}$*, so the supportable team level is* $\approx 1 - 192 \times 1.5 \times 10^{-3} \approx 0.71$*, not* 0.90 *— and that arithmetic still counts only the returned plan's durations; the selection-free* $N|E|$ *budget of the repaired theorem is weaker still. The reported team confidence must therefore be annealed to the supportable level in sub-floor cells (or the planner must return infinite windows and abstain); the P0 tables carry the honest label "nominal* $\alpha = 0.1$*; supportable team level* ≈ 0.71*". Empirical collisions are zero either way — the label, not the behavior, was wrong. q floors to the max score in* 100% *of certified cells and certified*(α) *is bit-identical to worst-case inflation — the probabilistic knob is inert until sum-level per-agent pricing (T8a at level* α/N*) lifts the level above the floor; note that T8a per agent reintroduces the fixed-path hypothesis at team scale, so the freshness-gate discussion of Remark 3 must be replayed there. This is the same width disease of the main-text width experiment, amplified by the team union factor.*

Honest accounting of assumptions and gaps

1. **Score independence.** Lemma 1 treats calibration scores as independent across observations. Scores sharing an edge share the "previous" noise draw of a pair; strictly, consecutive scores on the same edge are one-dependent. Thinning the buffer to disjoint observation pairs (the 2nd, 4th, ... observation of each edge) removes this at a per-edge factor-2 sample cost; implemented as `thinned_scores` and used, together with $\lambda = 2$, as the provable mode. Even thinned, the *test* score of edge e shares the stored draw η'_e with e's most recent calibration score; the provable statements therefore carry the leave-one-out convention (exclude that score when forming the quantile — Lemma 1, Theorem 5), an $O(\tilde{w}_i)$ correction that is now priced rather than ignored. Empirically the cost is small under route-critical sensing, which spreads observations across edges so most contribute only their (already disjoint) first pair. Relatedly, in unknown-terrain mode the *first* observation of an edge is never scored: a score requires a real previous observation, not a prior (`EdgeBelief.observed`); without this, prior error would contaminate the calibration distribution.

2. **The π_{cal} constant.** Theorem 2's π_{cal} is left as a distribution-dependent constant; under

A3 with noise scale σ_η it is $O(\langle \rho a \rangle_{\text{cal}}/\sigma_\eta)$. The sensing loop keeps calibration ages $\leq (L-1)\Delta$ (Theorem 3), making it second-order in all our regimes. The implementation tracks the realized $\langle \rho a \rangle_{\text{cal}}$ (`CertPlanner.cal_rho_a_max`) so the constant is reportable per run rather than asserted.

3. **Union bound over the optimum (GAP-A, re-repaired).** The corollary now carries the rigorous two-sided constant $\alpha_{\text{edge}} = \alpha'/(L_{\max} + E_{\mathcal{C}})$ — $\alpha'/(2L_{\max})$ for the single-candidate form — pricing both the unknown optimum's edges and every candidate edge (the min over candidates is a data-dependent selection). The earlier "resolved" constant $\alpha'/L_{\max}$ double-counted: one $L_{\max}$-edge budget cannot pay both sides. The deployed realized-L level remains documented as the operating approximation (empirically covered at 1.000 everywhere), and `strict_lb_alpha` implements the $L_{\max}$-level mode — one step short of theorem-exact.

4. **The drift model enters twice** (A1 in widths, A2 in Δ_{stale}); they are distinct assumptions and misspecification of each is swept separately in Tier-0.

5. **Clipping and observable semantics.** The search metrics clip ℓ_e at a positive cost floor; this is sound for the *latent* certificate (true costs are positive, so raising a lower bound below them loses nothing) but *invalid for Theorem 1's observable claim*: the observable $Y = c + \eta$ can be negative under heavy-tailed noise, and testing observables against clipped intervals manufactures spurious miscoverage concentrated exactly in heavy-left-tail regimes (measured: a 3.7× apparent edge-level break under Student-t noise that disappears entirely against unclipped intervals, while Gaussian noise and right-skewed noise show none — the left-tail fingerprint). Coverage events (ACI feedback, audits) must therefore be evaluated against the unclipped interval; the clip lives only inside the search metrics.