TRACKING THE BEST IMPUTER UNDER DRIFT: CONTINUAL COMPETENCE ROUTING FOR NON-STATIONARY TIME SERIES

ABSTRACT

Time-series imputation is typically deployed statically: an imputer is selected offline and retained as the stream evolves. Yet even when all imputers remain frozen, their *relative competence* can change with missingness structure and data regime. We formulate this problem as *continual competence routing* under sparse reconstruction feedback and introduce *Continual Competence Routing* (CCR), which tracks time-varying competence using multi-scale memories of audited reconstruction evidence while leaving the underlying imputers fixed. Our analysis identifies a fundamental distinction between coverage and freshness. Under a linear relative-competence model, arbitrarily strong all-history geometric coverage can remain stale after a change in expert ordering, whereas sufficiently informative post-change coverage controls fresh competence-estimation error and supports recovery of the current ordering. Across clinical and general time-series streams, frozen expert libraries exhibit substantial competence-oracle headroom and recurrent reversals in the best imputer. Fresh geometry is informative on ETTh1 and Weather but weaker on clinical streams. Under sparse audits, multi-scale routing improves over static deployment in several settings but does not uniformly beat the best static or fixed-memory baseline. Successful continual imputation therefore depends on evidence that is not merely abundant, but relevant to the current competence regime.

1 INTRODUCTION

Time-series imputation is usually deployed statically: an imputer is selected offline and retained throughout the subsequent stream. This protocol can fail even when every imputer is frozen, because heterogeneous imputers rely on different structural assumptions and their *relative competence* can change with missingness geometry, cross-variable dependence, or data regime. The deployment problem is therefore not only which imputer performs best on average, but which frozen imputer is currently competent.

We study this problem as *continual competence routing*. At round t, a router selects from a fixed expert library before receiving current feedback. Naturally missing entries provide no reconstruction target, so competence is observed only through sparse post-decision audits: controlled holdouts reveal reconstruction losses after the routing decision and can affect only future rounds. The experts remain fixed; only the routing state adapts.

The central difficulty is that competence evidence has an *age*. An all-history estimator can have excellent geometric coverage of the current context while relying on labels generated under an obsolete competence regime. Conversely, fewer recent observations can better identify the current ordering. We formalize this distinction under a linear relative-competence model: strong historical coverage can remain arbitrarily stale after a change in ordering, whereas informative post-change coverage controls current competence-estimation error. In short, *coverage is not freshness*.

Motivated by this separation, CCR maintains competence estimates over multiple memory horizons, updates them only from revealed audits, and combines the resulting estimates with adaptive scale weights. A conservative certification rule separates evidence favoring a candidate from evidence sufficient to replace the deployed expert. We evaluate this framework on clinical and general multivariate time-series streams under interventions specified independently of their eventual effect on expert ordering.

Our contributions are:

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1. We formulate continual imputation as tracking the time-varying relative competence of heterogeneous frozen imputers under sparse post-decision reconstruction audits, and introduce CCR as a multi-scale routing framework for this setting.
2. We prove an information-theoretic stale–fresh separation. Even in environments where $G_t^{\mathrm{all}} \succeq \varepsilon^{-1} I$ holds with arbitrarily high probability, every causal router has post-reset expected audit excess at least $\Delta(1-p)^j$ in one of two environments that remain observationally indistinguishable until fresh feedback arrives. In contrast, under the stated post-reset relative-competence model, a known-reset fresh ridge selector satisfies a high-probability excess bound $r_t^{\mathrm{audit}}(I_t^+) \le b_{\tau,t}^{\mathrm{rel}}(\delta)$; once this radius is smaller than the current competence margin, the current winner is recovered exactly.
3. We construct a controlled competence-drift benchmark that separates interventions from their effects on expert ordering and show that frozen libraries exhibit substantial oracle headroom and recurrent reversals, while the value of fresh evidence and adaptive routing is strongly regime- and representation-dependent.

2 Related Work

Time-series imputation. Time-series imputation has evolved from recurrent models that explicitly encode missingness, such as GRU-D and BRITS (Che et al., 2018; Cao et al., 2018), toward architectures exploiting attention, relational structure, and generative modeling. Representative methods include SAITS (Du et al., 2023), GRIN and SPIN (Cini et al., 2022; Marisca et al., 2022), conditional-diffusion approaches such as CSDI and PriSTI (Tashiro et al., 2021; Liu et al., 2023), and ImputeFormer (Nie et al., 2024). Empirical studies further show that relative performance varies across datasets, missingness mechanisms, rates, and gap structures (Kazijevs and Samad, 2023; Wang et al., 2025; Poette et al., 2026). These works primarily improve the imputer itself; we instead keep a heterogeneous library frozen and study how to track changes in its *relative competence* after deployment.

Adaptive and mixture-of-experts imputation. Mixture-of-experts approaches adapt the contribution or selection of specialized components according to the input or regime (Atman and Nannuru, 2025; Qin et al., 2026). CCR likewise exploits heterogeneous expertise, but uses a different deployment interface: its experts are complete, independently trained imputers whose parameters remain fixed online. The router neither jointly trains the experts nor learns a differentiable mixture of their predictions; it updates their estimated competence from sparse post-decision reconstruction audits. Our contribution therefore does not rest on multiple experts or adaptive routing alone.

Online expert selection and partial feedback. Non-stationary expert and bandit learning studies how to track changing losses or rewards, including dynamic tracking of the best expert (Wei et al., 2016). Label-efficient prediction instead restricts how often outcomes may be queried (Cesa-Bianchi et al., 2005). More recently, contextual active model selection chooses among pretrained models while actively querying labels (Liu et al., 2024), and constrained-feedback non-stationary bandits impose an explicit feedback budget (Li and Li, 2025). These are close online-learning abstractions, but the feedback interface differs in continual imputation: naturally missing entries do not reveal reconstruction targets; an audit is a controlled holdout revealed only after routing; and the tracked quantity is the context-dependent relative competence of frozen full imputers. CCR is therefore a sparse-feedback contextual routing problem with a reconstruction-specific observation mechanism, rather than a claim of a new generic expert-advice or feedback-budget primitive.

Freshness under non-stationarity. Windowing, discounting, and adaptive forgetting are standard responses to non-stationarity in sequential estimation and decision making (Russac et al., 2019; Cheung et al., 2022; Jacobsen and Cutkosky, 2024); these mechanisms are not themselves our novelty claim. Our focus is the distinction between *coverage* and *freshness*: historical competence observations may geometrically cover the current context while carrying labels from an obsolete expert ordering. We formalize this distinction by showing that arbitrarily strong all-history coverage can coexist with a post-change information-theoretic routing lower bound, whereas realized post-change directional coverage yields a high-probability relative-competence error bound. Thus, CCR

studies the intersection of frozen heterogeneous imputers, context-dependent competence, sparse post-decision reconstruction audits, and non-stationary evidence relevance.

3 Problem and Method

3.1 Continual competence routing

Let $\mathcal{M} = \{1, \dots, K\}$ be a library of imputers frozen before the online stream. At round t, the router observes context Z_t, forms $x_t = \Phi(Z_t) \in \mathbb{R}^d$, and selects $I_t \in \mathcal{M}$ using only information available before the current feedback. For expert m, define the evaluation competence

$$\mu_{m,t}^{\mathrm{eval}}(x_t) = \mathbb{E}[L_{m,t}^{\mathrm{eval}} \mid \mathcal{G}_t], \qquad m_t^{\mu} \in \arg\min_m \mu_{m,t}^{\mathrm{eval}}(x_t),$$

where $\mathcal{G}_t$ is the information available at decision time. Unlike static model selection, the competence oracle m_t^{μ} may change with the stream. Because naturally missing values provide no reconstruction target, feedback comes from exogenous post-decision audits $A_t \sim \mathrm{Bernoulli}(p)$. When $A_t = 1$, the full audit-loss vector $L_t^{\mathrm{audit}} = (L_{1,t}^{\mathrm{audit}}, \dots, L_{K,t}^{\mathrm{audit}})$ is revealed after I_t has been chosen and may affect only future decisions. Audit and evaluation targets are disjoint. We assume their relative competences are calibrated up to an expert-independent shift:

$$\sup_t \inf_{c \in \mathbb{R}} \max_m \left| \mu_{m,t}^{\mathrm{audit}} - \mu_{m,t}^{\mathrm{eval}} - c \right| \le \delta_{\mathrm{audit}}.$$

This condition is deliberately relative rather than absolute. Define $r_t^{\mathrm{audit}}(i) = \mu_{i,t}^{\mathrm{audit}} - \min_m \mu_{m,t}^{\mathrm{audit}}$ and $r_t^{\mathrm{eval}}(i) = \mu_{i,t}^{\mathrm{eval}} - \min_m \mu_{m,t}^{\mathrm{eval}}$. Writing $a_{m,t} = \mu_{m,t}^{\mathrm{eval}} - \mu_{m,t}^{\mathrm{audit}}$ and $\omega_t = \max_m a_{m,t} - \min_m a_{m,t}$, Proposition 11 shows that

$$r_t^{\mathrm{eval}}(I_t) \le r_t^{\mathrm{audit}}(I_t) + \omega_t, \qquad \omega_t = 2 \inf_{c \in \mathbb{R}} \max_m |a_{m,t} - c|.$$

Hence an expert-independent shift between audit and evaluation risks does not change the routing problem, while the stated condition contributes at most $2\delta_{\mathrm{audit}}$ to evaluation excess. Appendix A gives the formal filtration and audit–evaluation interface.

3.2 Multi-scale competence estimation

A fixed memory faces a bias–adaptation tradeoff: long histories provide more audits but can retain stale competence evidence, whereas short histories adapt faster but may be poorly supported. CCR therefore maintains horizons $\mathcal{H} = \{h_1, \dots, h_J\}$. For the trailing window $W_t(h)$,

$$G_t(h) = \lambda I + \sum_{s \in W_t(h)} A_s x_s x_s^{\top}, \qquad \widehat{\theta}_{m,t}^{(h)} = G_t(h)^{-1} \sum_{s \in W_t(h)} A_s x_s L_{m,s}^{\mathrm{audit}},$$

with

$$\widehat{R}_{m,t}^{(h)} = x_t^{\top} \widehat{\theta}_{m,t}^{(h)}, \qquad u_t(h) = x_t^{\top} G_t(h)^{-1} x_t.$$

Here $u_t(h)$ measures geometric support for the current context using audited evidence retained at horizon h. Its all-history analogue u_t^{all} is used diagnostically; Section 4 shows why small historical leverage need not imply fresh competence information.

3.3 Adaptive memory scales

Let $I_t^{(h)} \in \arg\min_m \widehat{R}_{m,t}^{(h)}$ and let $\ell_{h,t} \in [0,1]$ be its normalized audit loss, using a normalization fixed before test evaluation. Before round t,

$$q_t(h) = \frac{w_{h,t}}{\sum_{g \in \mathcal{H}} w_{g,t}}.$$

After audit revelation,

$$\widetilde{\ell}_{h,t} = \frac{A_t}{p} \ell_{h,t}, \qquad w_{h,t+1} = w_{h,t} \exp(-\eta \widetilde{\ell}_{h,t}).$$

The practical router aggregates the horizon-specific predictions as

$$\bar{R}_{m,t} = \sum_{h \in \mathcal{H}} q_t(h) \widehat{R}^{(h)}_{m,t}.$$

Thus unaudited rounds generate no artificial competence feedback, while the scale weights adapt to several candidate freshness ranges without assuming a known change point.

3.4 Certification and deployment

Let $\widehat{m}_{1,t}$ and $\widehat{m}_{2,t}$ be the best and second-best experts under $\bar{R}_{m,t}$, with $\widehat{\Delta}_t = \bar{R}_{\widehat{m}_{2,t},t} - \bar{R}_{\widehat{m}_{1,t},t}$. For practical uncertainty scores $C_{m,t}$, the leading candidate is certified when

$$\widehat{\Delta}_t > \alpha\big(C_{\widehat{m}_{1,t},t} + C_{\widehat{m}_{2,t},t}\big).$$

These scores are motivated by the geometry in Section 4, but are not claimed to be formal anytime confidence radii. Certification does not automatically authorize a switch. If J_{t-1} is the incumbent and the certified candidate differs from it, deployment changes only when

$$\bar{R}_{J_{t-1},t} - \bar{R}_{\widehat{m}_{1,t},t} > \alpha\big(C_{J_{t-1},t} + C_{\widehat{m}_{1,t},t}\big) + \gamma_{\mathrm{sw}}.$$

A candidate blocked by this condition cannot become the fallback certified state without subsequently passing the deployment rule. Full state recursion and implementation details are given in Appendix F.

3.5 Evaluation

The primary metric is normalized competence pseudo-regret relative to the time-varying competence oracle:

$$\mathrm{nCPR} = \frac{\sum_t \left(\widetilde{\mu}^{\mathrm{eval}}_{J_t,t} - \widetilde{\mu}^{\mathrm{eval}}_{m^{\mu}_t,t}\right)}{\sum_t \widetilde{\mu}^{\mathrm{eval}}_{m^{\mu}_t,t} + \epsilon_{\mathrm{norm}}},$$

where $\widetilde{\mu}^{\mathrm{eval}}$ is estimated offline and never exposed to the router. We additionally report oracle-match accuracy, harmful-switch rate, switch count, and adaptation or recertification delay when applicable.

Throughout the primary experiments, the experts and representation Φ remain frozen. Only competence memories, scale weights, and routing/deployment state adapt online; implementation and hyperparameters are reported in Appendix F.

4 Why Competence Evidence Must Be Fresh

Sparse feedback is not the only difficulty in continual competence routing: past feedback may no longer describe the current expert ordering. We isolate this distinction under an idealized linear relative-competence model. The analysis is mechanism-level; complete assumptions and proofs are deferred to Appendices A, B, and C.

4.1 Historical coverage can be stale

Let

$$G^{\mathrm{all}}_t = \lambda I + \sum_{s<t} A_s x_s x_s^{\top}, \qquad u^{\mathrm{all}}_t = x_t^{\top} (G^{\mathrm{all}}_t)^{-1} x_t.$$

In a stationary regime, small u^{all}_t indicates that the current context is well supported by previous audited observations. Under competence drift, however, the design records *where* competence was observed, not whether those competence labels remain valid.

Theorem 1 (Historical coverage can be arbitrarily stale)**.** *For any $p \in (0,1]$, $\lambda > 0$, $\varepsilon > 0$, and $\Delta \in (0, 1/2]$, there exist two bounded-loss piecewise-stationary linear environments with opposite post-reset winners and common pre-reset observable laws such that, with arbitrarily high probability,*

$$G^{\mathrm{all}}_t \succeq \varepsilon^{-1} I$$

after the reset. Nevertheless, at post-reset offset j, every causal router has expected audit excess at least

$$\Delta(1-p)^j$$

in one of the two environments.

Thus arbitrarily strong historical coverage does not identify the current competence state. The factor $(1-p)^j$ captures the probability that no distinguishing post-reset feedback has yet arrived: additional old evidence cannot replace information from the current regime.

4.2 Fresh coverage controls routing error

For analysis, suppose a reset occurs at a known time τ. Assume exact stationary post-reset relative competence,

$$\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}} = x_t^\top v_{ij}^+, \qquad \|v_{ij}^+\|_2 \le S_{\mathrm{rel}}, \qquad t \ge \tau,$$

with zero relative approximation residual and conditionally σ_{rel}-sub-Gaussian contrast noise. Define

$$G_{t,\tau}^+ = \lambda I + \sum_{s=\tau}^{t-1} A_s x_s x_s^\top, \qquad u_{t,\tau}^+ = x_t^\top (G_{t,\tau}^+)^{-1} x_t.$$

Unlike G_t^{all}, this design contains only competence evidence from the current regime.

Theorem 2 (Fresh-geometry recovery)**.** *Under the assumptions above, with probability at least $1-\delta$, simultaneously for all $t \ge \tau$, the fresh ridge selector I_t^+ satisfies*

$$r_t^{\mathrm{audit}}(I_t^+) \le \sigma_{\mathrm{rel}} \sqrt{\log \frac{\det G_{t,\tau}^+}{\lambda^d} + 2\log \frac{K(K-1)}{2\delta}} \sqrt{u_{t,\tau}^+} + \lambda S_{\mathrm{rel}} \left\| (G_{t,\tau}^+)^{-1} x_t \right\|_2 .$$

Under the audit–evaluation calibration condition, evaluation excess incurs at most an additional $2\delta_{\mathrm{audit}}$.

Together, Theorems 1 and 2 formalize *coverage is not freshness*. All-history uncertainty can be arbitrarily small while its competence labels are obsolete, whereas post-change directional coverage controls current routing error and, under a margin, the current winner.

The reset τ is an analytical device and is not given to CCR. The practical router instead maintains multiple finite memories and adapts their weights from observed audits. The theory motivates this use of fresh competence evidence; it does not claim that the practical multi-scale router is equivalent to the known-reset estimator or provide an end-to-end dynamic-regret guarantee.

5 Benchmark

We evaluate continual competence routing on streams in which the imputers remain frozen while missingness or data regimes change. The benchmark separates three questions: whether relative competence changes, whether fresh evidence is more informative than accumulated evidence, and whether the resulting competence changes can be tracked under sparse feedback.

5.1 Datasets and controlled non-stationarity

The clinical panel contains MIMIC-IV and PhysioNet/CinC 2019 hospital B; ETTh1 and Weather form the primary general time-series panel, with Electricity as an exploratory extension. The frozen expert library is $\{\mathrm{Linear}, \mathrm{SAITS}, \mathrm{T1}\}$ (Appendix F.1). For every dataset, candidate imputers and the router representation are fitted using training/validation data and then frozen throughout the online evaluation stream.

We generate controlled interventions affecting missingness intensity (D1), gap geometry (D2), affected channels (D3), cross-variable dependence (D4), data state or distribution (D5), and recurrent

regimes (D6). Intervention schedules and parameters are fixed independently of their eventual effect on expert ordering.

An observable change need not induce competence drift. We therefore assign, retrospectively, an evaluation-only taxonomy (D7) indicating whether an intervention preserves or reverses the relative competence ordering. D7 never determines intervention generation or routing. Ordering-preserving episodes are retained as negative controls, separating observable distributional change from decision-relevant competence change.

5.2 Sparse causal audits

At round t, the router observes the context and frozen expert predictions, selects an expert, and only then receives the audit indicator A_t. When $A_t = 1$, the full audit-loss vector becomes available for future updates. Evaluation losses, competence-oracle identities, reversal labels, and future regime information are never available to the online router.

Audit masks are pre-generated and shared across competing methods, so paired routers receive identical opportunities for competence feedback. For clinical streams, naturally missing values are never treated as ground truth: controlled audit and evaluation masks are applied only to source locations with observed measurements.

5.3 Research questions

The benchmark addresses three research questions. **RQ1** asks whether relative competence actually drifts: we measure competence-oracle headroom over static deployment and characterize recurrent changes in the identity and margin of the competent expert. **RQ2** asks whether fresh evidence is more informative than historical coverage: we compare fresh and all-history geometry as diagnostics of current oracle mismatch, providing an empirical counterpart to the stale–fresh separation in Section 4. **RQ3** asks whether competence can be tracked under sparse feedback: we compare static, fixed-memory, multi-scale, and certified routing using nCPR as the primary metric, complemented by oracle-match and switching diagnostics, and examine how routing behavior changes as audit feedback becomes sparser.

5.4 Paired evaluation

All primary comparisons are paired: competing routers receive the same stream, intervention realization, frozen expert predictions, controlled targets, audit activations, and evaluation targets. Evaluation-only quantities are computed only after online decisions are frozen. Statistical resampling is performed at the stream or intervention-episode level rather than treating individual time points as independent observations.

Dataset preprocessing, clinical cohort construction, intervention generators, audit construction, and the complete statistical protocol are provided in Appendices E–G.

6 Experiments

We organize the experiments around three questions introduced in Section 5: whether frozen experts exhibit decision-relevant competence drift, whether fresh geometry captures this drift better than all-history geometry, and whether multi-scale routing can exploit it under sparse feedback. These questions separate the existence of competence drift from its observability and from the ability of a practical router to act on it.

6.1 RQ1: Frozen imputers exhibit substantial competence drift

We first ask whether continual routing is necessary at all. If a single offline-selected expert remains close to the competence oracle throughout the stream, adaptive deployment has little room to improve.

Across the primary benchmark panel, competence-oracle headroom ranges from 7.68% to 11.67% on MIMIC-IV, 13.48% to 17.44% on PhysioNet-2019B, 9.08% to 17.99% on ETTh1, and 17.21% to

Table 1: **Competence drift in the frozen expert library.** Range of competence-oracle headroom over static selection across the reported intervention families. The primary analysis regards 5% relative headroom together with recurring strong reversals as evidence of decision-relevant competence drift. Electricity is an exploratory extension.

Dataset	Oracle headroom	Evidence of drift
MIMIC-IV	7.68–11.67%	strong
PhysioNet-2019B	13.48–17.44%	strong
ETTh1	9.08–17.99%	strong
Weather	17.21–26.71%	strong
Electricity	0.003–10.40%	heterogeneous

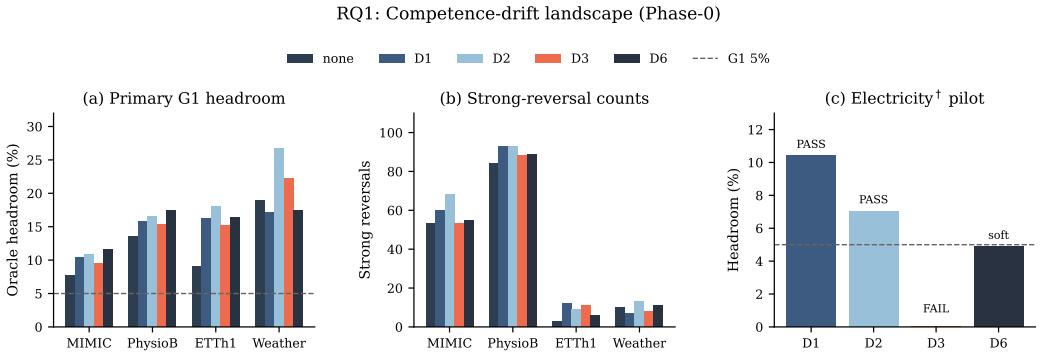


Figure 1: **A frozen library can have a moving winner.** Competence-oracle headroom and strong-reversal counts under controlled interventions. The dashed line denotes the predeclared 5% headroom criterion. The primary clinical and general time-series datasets exhibit both meaningful headroom and repeated changes in expert ordering; Electricity is more heterogeneous across intervention families.

26.71% on Weather (Table 1). These gaps coincide with recurring strong changes in expert ordering, showing that although every imputer is frozen, the identity of the competent model is not.

Electricity provides a useful counterpoint: competence headroom is heterogeneous across intervention families, with some changes in data regime producing little decision-relevant competence change. This supports measuring competence effects after intervention rather than assuming them from the intervention design.

Overall, RQ1 establishes the empirical premise of continual routing: static selection leaves substantial attainable competence unexploited across all four primary datasets, with recurring changes in which frozen expert is preferable. Per-family headroom, strong-reversal counts, and the Electricity pilot appear in Tables 3–4 (Appendix G.1).

6.2 RQ2: Fresh geometry is informative, but not universally so

We next test the empirical implication of Theorems 1–2: recent competence observations may remain informative about the current ordering when all-history geometry is dominated by earlier regimes. We compare fresh and all-history leverage as diagnostics of incorrect routing decisions using

$$\Delta\mathrm{AUC} = \mathrm{AUC}_{\mathrm{fresh}} - \mathrm{AUC}_{\mathrm{all}}.$$

The strongest separation occurs on ETTh1 and Weather. Under full auditing ($p = 1$), fresh geometry improves failure-discrimination AUROC by $+0.0967$ and $+0.1052$, respectively.

The aggregate evidence is more nuanced. Across the pooled time-series panel, the mean difference is $+0.0325$ at $p = 1$, with a 95% confidence interval of $[-0.0432, 0.1140]$; at $p = 0.2$, it is $+0.0100$, with interval $[-0.0802, 0.0981]$. The clinical effects are smaller: MIMIC-IV ranges from -0.0179 under the base full-audit context to $+0.0173$ under sparse auditing, while PhysioNet-2019B gives $+0.0054$ at $p = 1$ and $+0.0147$ at $p = 0.2$.

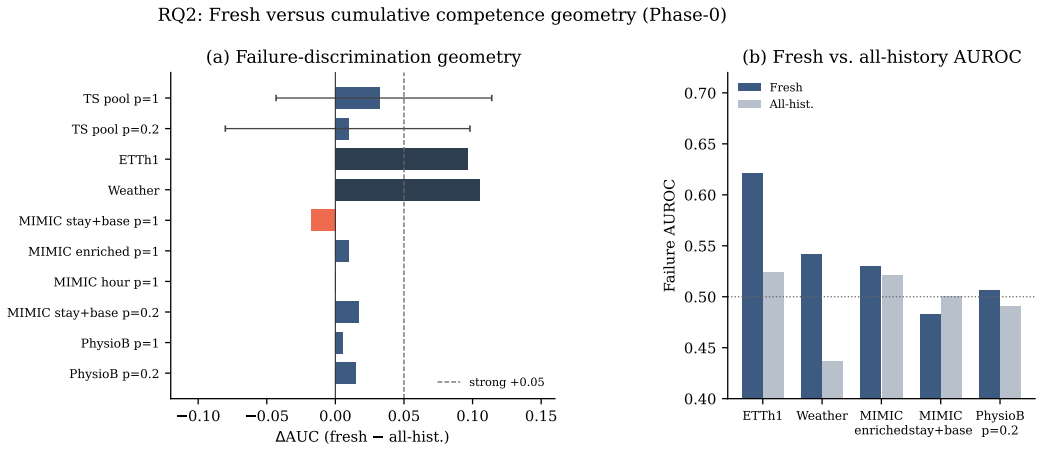


Figure 2: **Fresh and cumulative coverage need not carry the same decision information.** Difference in failure-discrimination AUROC, $\Delta\mathrm{AUC} = \mathrm{AUC}_{\mathrm{fresh}} - \mathrm{AUC}_{\mathrm{all}}$, together with representative AUROCs. Fresh geometry is substantially more informative on ETTh1 and Weather, whereas the clinical effect is smaller and sensitive to the context and audit protocol.

Table 2: **Routing performance measured by nCPR (lower is better).** Methods are compared only within the same protocol; full-audit and sparse-audit blocks are not pooled. Bold indicates the lowest reported nCPR within each row.

Audit	Dataset	Static	Fixed	CCR-MS	CCR-Cert
$p = 1$	ETTh1	**0.031**	0.179	0.132	**0.031**
$p = 1$	Weather	**0.142**	0.676	0.509	**0.142**
$p = 1$	MIMIC-IV	0.0777	0.0784	**0.0776**	0.0777
$p = 1$	PhysioNet-2019B	**0.1494**	0.1574	0.1517	0.1497
$p = 0.2$	ETTh1	0.0498	**0.0277**	0.0280	–
$p = 0.2$	Weather	0.1563	0.0180	**0.0151**	–
$p = 0.2$	Electricity	0.0552	0.0045	**0.0039**	–

Thus, the experiments support the distinction predicted by the theory without implying universal dominance of fresh leverage. Freshness is strongly informative on ETTh1 and Weather but weak in the current clinical representations, identifying representation-dependent freshness as an important boundary of the approach. Cell-level $\Delta\mathrm{AUC}$ values, including clinical protocol variants, are in Table 6 (Appendix G.2).

6.3 RQ3: Multi-scale tracking helps selectively under sparse feedback

We finally ask whether the practical router can convert the competence variation established in RQ1 into better deployment decisions. This is strictly harder than demonstrating oracle headroom: the router must infer the current ordering from past audits and decide when the evidence justifies leaving a competitive static expert.

Table 2 reports nCPR (lower is better). Under full auditing, CCR-MS is essentially non-harmful on MIMIC-IV (0.0776 versus 0.0777 for Static) and close to Static on PhysioNet-2019B (0.1517 versus 0.1494), but performs substantially worse on ETTh1 and Weather.

Under sparse auditing ($p = 0.2$), the picture changes. On ETTh1, CCR-MS reduces nCPR from 0.0498 for Static to 0.0280, nearly matching the best fixed-memory result 0.0277. On Weather, it improves from 0.1563 for Static and 0.0180 for Fixed to 0.0151; on exploratory Electricity, the corresponding values are 0.0552, 0.0045, and 0.0039. Multi-scale tracking therefore provides substantial gains in several sparse-feedback settings, although it does not uniformly dominate a

well-chosen fixed horizon. The extended sparse/clinical nCPR panel and the corresponding behavior plot are in Table 8 and Figure 3 (Appendix G.3).

Certification exposes the complementary trade-off. Under full auditing, CCR-Cert returns to 0.031 on ETTh1 and 0.142 on Weather, essentially matching Static, with similar behavior on the clinical datasets. The conservative rule therefore removes much of the degradation of unconstrained switching, but largely by suppressing adaptation rather than by reliably identifying beneficial switches.

Synthesis. RQ1 establishes recurrent competence drift, RQ2 shows a representation-dependent advantage of fresh geometry, and RQ3 shows that sparse feedback can support useful but non-uniform adaptation. The resulting target is *selective adaptation*: switch only when fresh competence evidence is sufficiently informative.

7 Discussion and Limitations

Our results separate two challenges in continual imputation: whether the competent imputer changes and whether online evidence is sufficiently informative to act on that change. The experiments provide clear evidence for the first through competence-oracle headroom and recurrent ordering changes. The second is harder: adaptive headroom does not imply that a router can reliably determine when to leave a strong incumbent.

Freshness and selective adaptation. Historical observations may cover the current context well while carrying stale competence labels. Empirically, fresh geometry is substantially more informative on ETTh1 and Weather but weaker and representation-dependent on the clinical streams; it is therefore a mechanism-dependent diagnostic rather than a universal uncertainty proxy. Likewise, finite-memory tracking can exploit sparse feedback, whereas conservative certification often avoids harmful switches by reducing adaptation. The remaining challenge is selectivity: preserve a strong incumbent under weak evidence while switching when a competence reversal is sufficiently supported.

Scope and limitations. The theory is mechanism-level: the known reset time is an analytical device, and we do not claim an end-to-end regret guarantee for practical scale weighting, certification, and switching. Controlled audits evaluate only originally observed targets; naturally missing clinical values are never treated as ground truth. D1–D6 are controlled rather than exhaustive models of deployment drift. Longer natural streams, richer expert libraries, learned competence representations, and non-stationary audits remain important extensions.

8 Conclusion

We formulated continual time-series imputation as the problem of tracking the changing relative competence of heterogeneous frozen imputers under sparse post-decision feedback. This perspective separates adaptation of the *deployment decision* from adaptation of the underlying imputers: even when every expert remains fixed, changes in missingness structure and data regime can alter which expert is currently preferable.

Our analysis identifies freshness as a fundamental requirement for this decision. The stale–fresh separation shows that arbitrarily strong all-history geometric coverage need not reveal the current expert ordering when its competence labels come from obsolete regimes, whereas sufficiently informative post-change evidence controls fresh competence-estimation error under the stated model. Empirically, frozen expert libraries exhibit meaningful competence-oracle headroom and recurrent reversals across both clinical and general time-series streams. Fresh geometry is markedly informative in some representations, and multi-scale competence tracking can exploit sparse audits in several regimes, but neither mechanism yields uniform gains across datasets.

These results sharpen the deployment question from *which imputer should be selected once?* to *what evidence is sufficient to justify changing the imputer now?* The central challenge is therefore selective adaptation: retaining a strong incumbent when current evidence is weak, while responding promptly when fresh competence evidence supports a change. More broadly, continual imputation requires not merely accumulating evidence, but ensuring that the evidence used for deployment remains relevant to the competence decision being made.

Reproducibility Statement

The submission includes the frozen protocol, preprocessing rules, audit chronology, expert-library specification, implementation details, and experiment definitions in the appendix. Controlled audit masks and intervention realizations are shared across paired methods. Evaluation-only oracle and reversal annotations are excluded from online routing state.

AI Use Statement

Generative AI tools were used for language editing, code assistance, and consistency checking. Scientific claims, mathematical statements, experimental design, implementation decisions, and reported results were reviewed and validated by the authors. No AI-generated empirical result is reported without execution and verification by the authors.

References

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Online least squares estimation with self-normalized processes: An application to bandit problems. *arXiv preprint arXiv:1102.2670*, 2011. URL https://arxiv.org/abs/1102.2670.

Aman Atman and Santosh Nannuru. Mixture-of-experts for time series imputation. In *Proceedings of the 38th Canadian Conference on Artificial Intelligence*, 2025. doi: 10.21428/594757db.e3e5bb44.

Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, and Yitan Li. BRITS: Bidirectional recurrent imputation for time series. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 6775–6785. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/734e6bfcd358e25ac1db0a4241b95651-Paper.pdf.

Nicolò Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Minimizing regret with label efficient prediction. *IEEE Trans. Inf. Theory*, 51(6):2152–2162, 2005. doi: 10.1109/TIT.2005.847729. URL https://doi.org/10.1109/TIT.2005.847729.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8(1):6085, 2018. doi: 10.1038/s41598-018-24271-9.

Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Hedging the drift: Learning to optimize under nonstationarity. *Management Science*, 68(3):1696–1713, 2022. doi: 10.1287/mnsc.2021.4024.

Andrea Cini, Ivan Marisca, and Cesare Alippi. Filling the G_ap_s: Multivariate time series imputation by graph neural networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=kOu3-S3wJ7.

Wenjie Du, David Côté, and Yan Liu. SAITS: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219:119619, 2023. doi: 10.1016/j.eswa.2023.119619.

Andrew Jacobsen and Ashok Cutkosky. Online linear regression in dynamic environments via discounting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 21083–21120, 2024.

Maksims Kazijevs and Manar D. Samad. Deep imputation of missing values in time series health data: A review with benchmarking. *Journal of Biomedical Informatics*, 144:104440, 2023. doi: 10.1016/j.jbi.2023.104440.

Shaoang Li and Jian Li. Constrained feedback learning for non-stationary multi-armed bandits. In *Advances in Neural Information Processing Systems*, volume 38, pages 89517–89550, 2025. doi: 10.52202/085713-2697.

Mingzhe Liu, Han Huang, Hao Feng, Leilei Sun, Bowen Du, and Yanjie Fu. PriSTI: A conditional diffusion framework for spatiotemporal imputation. In *2023 IEEE 39th International Conference on Data Engineering*, pages 1927–1939, 2023. doi: 10.1109/ICDE55515.2023.00150.

Xuefeng Liu, Fangfang Xia, Rick L. Stevens, and Yuxin Chen. Contextual active model selection. In *Advances in Neural Information Processing Systems*, volume 37, pages 49914–49956, 2024. doi: 10.52202/079017-1581.

Ivan Marisca, Andrea Cini, and Cesare Alippi. Learning to reconstruct missing data from spatiotemporal graphs with sparse observations. In *Advances in Neural Information Processing Systems*, volume 35, pages 32069–32082, 2022. doi: 10.52202/068431-2324.

Tong Nie, Guoyang Qin, Wei Ma, Yuewen Mei, and Jian Sun. ImputeFormer: Low rankness-induced transformers for generalizable spatiotemporal imputation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, pages 2260–2271. Association for Computing Machinery, 2024. doi: 10.1145/3637528.3671751.

Michael Poette, Sandrine Mouysset, Daniel Ruiz, Vincent Pey, Jean-Marc Alliot, and Vincent Minville. Benchmarking imputation strategies for missing time-series data in critical care using real-world-inspired scenarios. *Scientific Reports*, 16(1):8116, 2026. doi: 10.1038/s41598-026-39035-z.

Cang Qin, Lina Yang, Ling Peng, Wenyue Zhang, and Meng Wei. STMoE: A structured mixture-of-experts framework for robust spatiotemporal series imputation. *Expert Systems with Applications*, 319:132137, 2026. doi: 10.1016/j.eswa.2026.132137.

Matthew A. Reyna, Chris S. Josef, Salman Seyedi, Russell Jeter, Supreeth P. Shashikumar, M. Brandon Westover, Ashish Sharma, Shamim Nemati, and Gari D. Clifford. Early prediction of sepsis from clinical data: The PhysioNet/computing in cardiology challenge 2019. In *2019 Computing in Cardiology (CinC)*, pages 1–4. IEEE, 2019. doi: 10.23919/CinC49843.2019.9005736.

Yoan Russac, Claire Vernade, and Olivier Cappé. Weighted linear bandits for non-stationary environments. In *Advances in Neural Information Processing Systems*, volume 32, pages 12040–12049, 2019.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems*, volume 34, pages 24804–24816, 2021.

Artur Trindade. ElectricityLoadDiagrams20112014, 2015.

Joel A. Tropp. User-friendly tail bounds for matrix martingales. Technical Report ACM Report 2011-01, California Institute of Technology, 2011.

Jun Wang, Wenjie Du, Yiyuan Yang, Linglong Qian, Wei Cao, Keli Zhang, Wenjia Wang, Yuxuan Liang, and Qingsong Wen. Deep learning for multivariate time series imputation: A survey. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 10696–10704. International Joint Conferences on Artificial Intelligence Organization, 8 2025. doi: 10.24963/ijcai.2025/1187. URL `https://doi.org/10.24963/ijcai.2025/1187`. Survey Track.

Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Tracking the best expert in non-stationary stochastic environments. In *Advances in Neural Information Processing Systems*, volume 29, pages 3972–3980, 2016.

Justin Whitehouse, Zhiwei Steven Wu, and Aaditya Ramdas. Time-uniform self-normalized concentration for vector-valued processes. *The Annals of Applied Probability*, 36(3):1972–2013, 2026. doi: 10.1214/25-AAP2246.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):11106–11115, 2021. doi: 10.1609/aaai.v35i12.17325.

A Formal Setup, Timing, and Audit/Evaluation Interface

This appendix formalizes the information structure used by the theory and online benchmark. All subsequent results inherit the chronology, target-separation, and no-leakage conventions below.

A.1 Information structure and round chronology

All variables are defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$. Let $\mathcal{F}_{t-1}$ contain all information revealed strictly before round t, including past contexts, actions, audit activations and revealed audit losses, previous internal randomization, and the router state. After observing the current context Z_t, define

$$\mathcal{G}_t = \sigma(\mathcal{F}_{t-1}, Z_t). \tag{1}$$

The frozen representation $x_t = \phi(Z_t)$, with $\|x_t\|_2 \leq 1$, is $\mathcal{G}_t$-measurable. With fresh internal randomization U_t, the causal action satisfies

$$I_t = \pi_t(\mathcal{G}_t, U_t) \in \mathcal{M}. \tag{2}$$

The audit is revealed only after I_t is fixed. In the primary model,

$$A_t \mid \mathcal{G}_t, U_t, I_t, L_t^{\text{audit}} \sim \text{Bernoulli}(p), \tag{3}$$

so that

$$A_t \perp\!\!\!\perp (I_t, U_t, L_t^{\text{audit}}) \mid \mathcal{G}_t. \tag{4}$$

When $A_t = 1$, the complete audit-loss vector Y_t^{audit} is revealed and the online state becomes

$$\mathcal{F}_t = \sigma(\mathcal{G}_t, U_t, I_t, A_t, A_t Y_t^{\text{audit}}). \tag{5}$$

Here $A_t Y_t^{\text{audit}}$ denotes selective revelation, not an observed zero vector when $A_t = 0$. Thus current-round audit information cannot affect the current routing decision.

A.2 Audit and evaluation targets

For each frozen expert, let $L_t^{\text{eval}} = (L_{1,t}^{\text{eval}}, \ldots, L_{K,t}^{\text{eval}})$ and $L_t^{\text{audit}} = (L_{1,t}^{\text{audit}}, \ldots, L_{K,t}^{\text{audit}})$ denote the evaluation and audit potential-loss vectors. Their target sets are disjoint,

$$\mathcal{I}_t^{\text{eval}} \cap \mathcal{I}_t^{\text{audit}} = \varnothing. \tag{6}$$

Both target sets are withheld before constructing Z_t, x_t, and the expert predictions. Audit activation controls only whether an already withheld audit loss is revealed after the decision. Evaluation losses, competence oracles, reversal annotations, future regime identities, and evaluation margins are evaluation-only quantities and never enter routing, scale updates, certification, fallback, or online hyperparameter selection.

Define the conditional competences

$$\mu_{m,t}^{\text{eval}} = \mathbb{E}[L_{m,t}^{\text{eval}} \mid \mathcal{G}_t], \tag{7}$$

$$\mu_{m,t}^{\text{audit}} = \mathbb{E}[L_{m,t}^{\text{audit}} \mid \mathcal{G}_t]. \tag{8}$$

The target competence oracle is

$$m_t^{\mu} \in \arg\min_{m \in \mathcal{M}} \mu_{m,t}^{\text{eval}}, \tag{9}$$

whereas the pointwise oracle

$$m_t^{\text{pt}} \in \arg\min_{m \in \mathcal{M}} L_{m,t}^{\text{eval}} \tag{10}$$

is clairvoyant and used only as an offline diagnostic. The evaluation competence excess of a deployed action is

$$r_t^{\text{eval}}(I_t) = \mu_{I_t,t}^{\text{eval}} - \mu_{m_t^{\mu},t}^{\text{eval}}, \tag{11}$$

with $r_t^{\text{audit}}(I_t)$ defined analogously on the audit side.

Because routing depends on relative rather than absolute competence, the natural audit/evaluation interface is invariant to an expert-independent shift. We assume

$$\sup_{1 \leq t \leq T} \inf_{c \in \mathbb{R}} \max_{m \leq K} \left| \mu_{m,t}^{\text{audit}} - \mu_{m,t}^{\text{eval}} - c \right| \leq \delta_{\text{audit}}. \tag{12}$$

The special case $c = 0$ gives direct calibration. Under (12), an audit-side routing bound transfers to evaluation competence with an additive $2\delta_{\text{audit}}$ term; see Appendix D.3. Empirical calibration and ordering agreement are diagnostics of this interface, not proofs of the uniform bound.

A.3 PREDICTABLE RELATIVE-COMPETENCE MODEL

The main geometric results require only relative competence. For each expert pair (i, j), assume

$$\mu_{i,t}^{\text{audit}} - \mu_{j,t}^{\text{audit}} = x_t^\top v_{ij,t} + e_{ij,t}, \qquad \|v_{ij,t}\|_2 \le S_{\text{rel}}, \tag{13}$$

where $v_{ij,t}$ and $e_{ij,t}$ are predictable. The observed contrast noise

$$\zeta_{ij,t} = Y_{i,t}^{\text{audit}} - Y_{j,t}^{\text{audit}} - (\mu_{i,t}^{\text{audit}} - \mu_{j,t}^{\text{audit}}) \tag{14}$$

is conditionally sub-Gaussian:

$$\mathbb{E}[\exp(z\zeta_{ij,t}) \mid \mathcal{G}_t] \le \exp(z^2\sigma_{ij}^2/2), \qquad z \in \mathbb{R}.$$

Exact relative linearity corresponds to $e_{ij,t} = 0$. The model permits an arbitrary predictable component shared by all experts, since such a component cancels from pairwise competence contrasts. No independence across experts or rounds is required.

For results stated through absolute expert risks, we use the stronger model

$$Y_{m,t}^{\text{audit}} = x_t^\top \theta_{m,t} + \xi_{m,t}, \qquad \|\theta_{m,t}\|_2 \le S, \tag{15}$$

with predictable $\theta_{m,t}$ and

$$\mathbb{E}[\exp(z\xi_{m,t}) \mid \mathcal{G}_t] \le \exp(z^2\sigma^2/2). \tag{16}$$

This is the predictable-design setting of Abbasi-Yadkori et al. (2011). The relative model is an alternative to, not an additional requirement on, the absolute model. Under the latter, one may take $S_{\text{rel}} = 2S$ and $\sigma_{ij} = 2\sigma$ without assuming cross-expert independence.

Boundedness. The empirical benchmark reports losses on their declared scales. The theory uses bounded routing excess only where explicitly stated, for example for clipping or low-probability events. Sub-Gaussian regression noise does not require the noisy linear labels themselves to lie almost surely in $[0, 1]$.

Change times and oracle information. A reset time τ may be known inside an oracle diagnostic theorem without being observable to CCR. In particular, the fresh estimator in Theorem 7 is a known-reset analytical object. Empirical reversal times are computed only after the stream is complete and may be used to align diagnostic plots or define adaptation metrics; they are never router inputs or intervention-design variables.

A.4 STANDING ASSUMPTIONS

Unless a result states otherwise, the theory uses the following conventions: (i) the expert library and primary representation are frozen online; (ii) contexts and actions precede current audit activation and feedback; (iii) audits are exogenous Bernoulli events with rate p; (iv) an activated audit reveals losses for all frozen experts on the same controlled target; (v) audit and evaluation targets are disjoint and evaluation-only quantities never enter the online state; and (vi) known reset times are analytical or evaluation-only objects. Linearity, piecewise stationarity, competence margins, directional richness, and bounded normalized losses are imposed only by the results that require them.

A.5 PROOFS OF THE MAIN-TEXT RESULTS

Proof of Theorem 1. Before the reset, present every basis vector $e_1, \dots, e_d$ sufficiently many times and use common noiseless parameters $\theta_1^- = \theta_2^- = e_1/2$. By Bernoulli audit concentration and a finite union bound, with probability at least $1 - \eta$ each direction receives at least $\lceil(\varepsilon^{-1} - \lambda)_+\rceil$ audited observations. Hence $G_t^{\text{all}} \succeq \varepsilon^{-1} I_d$ at the reset and thereafter.

After the reset, set $x_t = e_1$. Under P_+, let $\theta_1^+ = (1/2 - \Delta)e_1$ and $\theta_2^+ = (1/2 + \Delta)e_1$; under P_-, exchange the two. The losses are noiseless and lie in $[0, 1]$. At decision $\tau + j$, let E_j denote the event that no post-reset audit has occurred before the decision. Then $\Pr(E_j) = (1 - p)^j$, and conditional on E_j the observable histories have the same law under P_+ and P_-. Since the two environments have opposite winners, their conditional expected excesses sum to 2Δ. Therefore

$$\mathbb{E}_{P_+} r_{\tau+j}^{\text{audit}}(I_{\tau+j}) + \mathbb{E}_{P_-} r_{\tau+j}^{\text{audit}}(I_{\tau+j}) \ge 2\Delta(1 - p)^j.$$

At least one environment consequently incurs expected excess $\Delta(1 - p)^j$, proving the claim. □

Proof of Theorem 2. Because all experts share the same Gram matrix, differences of their ridge estimates equal the ridge estimate of the corresponding audit-loss contrast. For a fixed pair with stationary post-reset parameter v_{ij}^{+}, exact relative linearity gives

$$\widehat{\mu}_{i,t}^{+} - \widehat{\mu}_{j,t}^{+} - (\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}}) = x_t^{\top}(G_{t,\tau}^{+})^{-1}\sum_{s=\tau}^{t-1} A_s x_s \zeta_{ij,s} - \lambda x_t^{\top}(G_{t,\tau}^{+})^{-1} v_{ij}^{+}.$$

Lemma 3, applied simultaneously to all unordered expert pairs, bounds the first term uniformly over time with probability at least $1-\delta$. Cauchy–Schwarz yields the stochastic term in Theorem 2, while $\|v_{ij}^{+}\|_2 \le S_{\mathrm{rel}}$ gives the ridge-bias term. The estimated contrast between I_t^{+} and a true audit winner is nonpositive, so its true audit excess is bounded by the same simultaneous pairwise error. Finally, (12) transfers the result to evaluation competence with an additional $2\delta_{\mathrm{audit}}$. □

B Geometry-Aware Competence Estimation

This appendix develops the geometric estimation result used by the freshness analysis. Contexts may depend on previously revealed feedback, so the relevant concentration argument must accommodate predictable designs. We use self-normalized martingale concentration (Abbasi-Yadkori et al., 2011) separately for each admissible left endpoint. This avoids treating a sliding-window estimator as a single accumulating martingale process.

B.1 Self-normalized inference over intervals

Fix $\lambda > 0$. Let $\mathcal{A}$ be a deterministic collection of candidate left endpoints and choose deterministic weights $\pi_a > 0$ satisfying

$$\sum_{a\in\mathcal{A}} \pi_a \le 1.$$

For $a \in \mathcal{A}$ and $t \ge a$, define

$$G_{a,t} = \lambda I_d + \sum_{s=a}^{t-1} A_s x_s x_s^{\top}.$$

For a finite horizon, taking $\mathcal{A} = \{1,\dots,T\}$ covers all trailing windows and candidate reset locations; for a single prespecified reset τ, one may instead take $\mathcal{A} = \{\tau\}$ and $\pi_\tau = 1$. The weights π_a are proof devices and are unrelated to the online scale weights of CCR.

Lemma 3 (Self-normalized inference over intervals)**.** *Suppose that, for $k \in \{1,\dots,L\}$, the predictable noise sequence $\varepsilon_{k,s}$ is conditionally σ_k-sub-Gaussian and that audits satisfy the exogeneity condition in Appendix A. Then, for every $\delta \in (0,1)$, with probability at least $1-\delta$, simultaneously for every k, $a \in \mathcal{A}$, and $t \ge a$,*

$$\left\|\sum_{s=a}^{t-1} A_s x_s \varepsilon_{k,s}\right\|_{G_{a,t}^{-1}} \le \sigma_k \sqrt{\log\frac{\det G_{a,t}}{\lambda^d} + 2\log\frac{L}{\delta\pi_a}}. \tag{17}$$

Here $\|z\|_M = (z^{\top}Mz)^{1/2}$.

Proof of Lemma 3. Fix a and k, and write

$$S_t = \sum_{s=a}^{t-1} A_s x_s \varepsilon_{k,s}.$$

For $\sigma_k > 0$ and deterministic $v \in \mathbb{R}^d$,

$$M_t(v) = \exp\left\{\frac{v^{\top}S_t}{\sigma_k} - \frac{1}{2}v^{\top}(G_{a,t} - \lambda I_d)v\right\}$$

is a nonnegative supermartingale. This follows from conditional sub-Gaussianity together with audit exogeneity: when $A_s = 0$ the multiplicative increment is one, while for $A_s = 1$ its conditional expectation is at most one.

Mixing $M_t(v)$ with respect to $N(0, \lambda^{-1} I_d)$ gives

$$\overline{M}_t = \left(\frac{\lambda^d}{\det G_{a,t}} \right)^{1/2} \exp\left\{ \frac{\|S_t\|^2_{G_{a,t}^{-1}}}{2\sigma_k^2} \right\}.$$

The maximal inequality for nonnegative supermartingales implies

$$\Pr\left(\sup_{t \geq a} \overline{M}_t > \frac{L}{\delta \pi_a} \right) \leq \frac{\delta \pi_a}{L}.$$

A union bound over a and k proves (17). If $\sigma_k = 0$, the conditional moment assumption implies $\varepsilon_{k,s} = 0$ almost surely. □

Remark. The concentration event is time-uniform for each fixed left endpoint, with a separate failure budget assigned to each a. It therefore covers the finite family of windows used by the analysis without asserting that a fixed-length moving window itself forms a single accumulating martingale. Sharper time-uniform constructions are possible (Whitehouse et al., 2026), but optimal logarithmic factors are not needed for our results.

B.2 Windowed competence estimation

For a memory horizon h, let

$$a = (t - h) \vee 1, \qquad W_t(h) = \{a, \ldots, t-1\}.$$

The estimator used by the router is

$$G_t(h) = \lambda I_d + \sum_{s \in W_t(h)} A_s x_s x_s^\top, \tag{18}$$

$$\widehat{\theta}^{(h)}_{m,t} = G_t(h)^{-1} \sum_{s \in W_t(h)} A_s x_s Y^{\text{audit}}_{m,s}, \tag{19}$$

$$u_t(h) = x_t^\top G_t(h)^{-1} x_t. \tag{20}$$

The leverage $u_t(h)$ measures how strongly the current context is supported by audited contexts retained at horizon h.

For completeness, under the absolute linear model

$$Y^{\text{audit}}_{m,s} = x_s^\top \theta_{m,s} + \xi_{m,s},$$

the prediction error admits the exact decomposition

$$\begin{aligned} x_t^\top \left(\widehat{\theta}^{(h)}_{m,t} - \theta_{m,t} \right) = {} & x_t^\top G_t(h)^{-1} \sum_{s \in W_t(h)} A_s x_s \xi_{m,s} \\ & + x_t^\top G_t(h)^{-1} b^{\text{dr}}_{m,t}(h) - \lambda x_t^\top G_t(h)^{-1} \theta_{m,t}, \end{aligned} \tag{21}$$

where

$$b^{\text{dr}}_{m,t}(h) = \sum_{s \in W_t(h)} A_s x_s x_s^\top (\theta_{m,s} - \theta_{m,t}).$$

The transported drift contribution satisfies

$$\left| x_t^\top G_t(h)^{-1} b^{\text{dr}}_{m,t}(h) \right| \leq \sqrt{u_t(h)} \left(\sum_{s \in W_t(h)} A_s \left[x_s^\top (\theta_{m,s} - \theta_{m,t}) \right]^2 \right)^{1/2}. \tag{22}$$

This decomposition separates stochastic estimation error, transported competence drift, and ridge bias. The relevant drift is therefore the component transported through the audited design, rather than Euclidean parameter displacement alone.

B.3 RELATIVE COMPETENCE BOUNDS

Routing depends on differences between expert risks rather than their absolute levels. We therefore state the main geometric result directly for relative competence. Assume

$$\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}} = x_t^\top v_{ij,t} + e_{ij,t}, \qquad \|v_{ij,t}\|_2 \le S_{\mathrm{rel}},$$

with conditionally sub-Gaussian contrast noise

$$\zeta_{ij,t} = Y_{i,t}^{\mathrm{audit}} - Y_{j,t}^{\mathrm{audit}} - \big(\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}}\big).$$

Because all experts use the same Gram matrix, the difference between two ridge estimators is already the ridge estimator of their audit-loss contrast. No separate pairwise regression is required.

Let $P = K(K-1)/2$. For $i \neq j$, define

$$D_{ij,t}(h) = \left| x_t^\top G_t(h)^{-1} \sum_{s \in W_t(h)} A_s x_s \left\{ x_s^\top (v_{ij,s} - v_{ij,t}) + e_{ij,s} \right\} - e_{ij,t} \right|, \tag{23}$$

$$B_{ij,t}(h;\delta) = \sigma_{ij} \sqrt{\log \frac{\det G_t(h)}{\lambda^d} + 2 \log \frac{P}{\delta \pi_a}} \sqrt{u_t(h)} + \lambda S_{\mathrm{rel}} \|G_t(h)^{-1} x_t\|_2 + D_{ij,t}(h), \tag{24}$$

where $a = (t-h) \vee 1$.

Theorem 4 (Geometry-aware relative competence estimation)**.** *Under the relative-competence and contrast-noise assumptions above, with probability at least $1-\delta$, simultaneously for all covered windows, rounds, and expert pairs,*

$$\big|(\widehat{\mu}_{i,t}^{(h)} - \widehat{\mu}_{j,t}^{(h)}) - (\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}})\big| \le B_{ij,t}(h;\delta). \tag{25}$$

Consequently, for

$$I_t^{(h)} \in \arg\min_m \widehat{\mu}_{m,t}^{(h)},$$

the audit-competence excess satisfies

$$r_t^{\mathrm{audit}}(I_t^{(h)}) \le B_{I_t^{(h)}, m_t^{\mathrm{audit}}, t}(h;\delta) \le \max_{i<j} B_{ij,t}(h;\delta). \tag{26}$$

Moreover,

$$D_{ij,t}(h) \le \sqrt{u_t(h)} \left(\sum_{s \in W_t(h)} A_s \left[x_s^\top (v_{ij,s} - v_{ij,t}) + e_{ij,s} \right]^2 \right)^{1/2} + |e_{ij,t}|. \tag{27}$$

Proof of Theorem 4. Linearity of the shared ridge estimator gives

$$\widehat{\theta}_{i,t}^{(h)} - \widehat{\theta}_{j,t}^{(h)} = G_t(h)^{-1} \sum_{s \in W_t(h)} A_s x_s \big(Y_{i,s}^{\mathrm{audit}} - Y_{j,s}^{\mathrm{audit}}\big).$$

Subtracting the current true contrast decomposes the prediction error into three terms: a contrast-noise projection, the transported drift/approximation term in (23), and the ridge bias

$$-\lambda x_t^\top G_t(h)^{-1} v_{ij,t}.$$

Lemma 3, applied to the P unordered expert pairs, bounds the noise projection by the first term in (24). The parameter norm bound controls the ridge term, proving (25).

Let m_t^{audit} be an audit-competence minimizer. If $I_t^{(h)} \neq m_t^{\mathrm{audit}}$, empirical optimality implies that the estimated contrast of $I_t^{(h)}$ against m_t^{audit} is nonpositive. Subtracting this estimated contrast from the true contrast gives (26); if the two experts coincide, the excess is zero.

Finally, writing the transported residuals as a vector z and the audited contexts as a design matrix X, the inequality

$$X(\lambda I + X^\top X)^{-1} X^\top \preceq I$$

followed by Cauchy–Schwarz yields (27). □

Under exact relative stationarity within the retained window, $v_{ij,s} = v_{ij,t}$ and $e_{ij,s} = e_{ij,t} = 0$, so

$$D_{ij,t}(h) = 0.$$

The remaining uncertainty is then governed by the local geometric factor $\sqrt{u_t(h)}$, the self-normalized complexity term, and ridge bias. Importantly, an arbitrary expert-independent component of the absolute risk cancels from every pairwise contrast. Thus the geometry relevant to routing concerns variation in *relative competence*, not variation in the common risk level.

All audit-side routing bounds transfer to evaluation competence through the audit/evaluation calibration result in Appendix A, contributing at most the corresponding $2\delta_{\mathrm{audit}}$ term.

C Freshness Separation and Post-Drift Recovery

This appendix develops the formal stale-versus-fresh separation underlying the freshness mechanism in Section 4. We show that cumulative geometric coverage can become arbitrarily strong while the current expert ordering remains statistically unidentified after a competence reset. In contrast, under the stated competence model, sufficiently informative fresh post-reset evidence controls competence-estimation error and enables recovery of the current routing decision. The analysis concerns an idealized audit-competence estimation problem. Its results isolate the freshness mechanism and should not be interpreted as end-to-end regret, coverage, or certification guarantees for practical CCR-Full.

C.1 Setup and notation

The context, audit chronology, and filtration are those of Appendix A. In particular,

$$A_t \mid \mathcal{G}_t, U_t, I_t, L_t^{\mathrm{audit}} \sim \mathrm{Bernoulli}(p). \tag{28}$$

The relative model is (13), with conditional contrast noise as in (14). The positive known-reset result requires stationary relative parameters and zero relative approximation residuals after the reset, not stationary common risk levels. When an absolute-risk prediction lemma is used, it separately invokes (15)–(16). The two-expert stale constructions below are special cases with noiseless absolute linear risks and bounded parameter norms.

Define

$$G_t^{\mathrm{all}} = \lambda I_d + \sum_{s<t} A_s x_s x_s^\top, \tag{29}$$

$$G_{t,\tau}^{+} = \lambda I_d + \sum_{s=\tau}^{t-1} A_s x_s x_s^\top, \tag{30}$$

$$u_t^{\mathrm{all}} = x_t^\top (G_t^{\mathrm{all}})^{-1} x_t, \tag{31}$$

$$u_{t,\tau}^{+} = x_t^\top (G_{t,\tau}^{+})^{-1} x_t. \tag{32}$$

Let $m_t^{\mathrm{audit}} \in \arg\min_m \mu_{m,t}^{\mathrm{audit}}$ and write

$$r_t^{\mathrm{audit}}(I_t) = \mu_{I_t,t}^{\mathrm{audit}} - \mu_{m_t^{\mathrm{audit}},t}^{\mathrm{audit}}. \tag{33}$$

A reset used by an oracle diagnostic is not supplied to the deployed router. For a data-selected reset, the interval family in Lemma 3 must cover all candidate left endpoints, rather than assign probability one to a retrospectively chosen endpoint.

C.2 Cumulative coverage can be arbitrarily stale

The next result separates accumulated historical coverage from current competence identifiability.

Theorem 5 (Uniform historical coverage does not identify post-reset competence)**.** *Fix $d \geq 1$, $\varepsilon > 0$, $\Delta \in (0, \frac{1}{2}]$, $p \in (0, 1]$, $\lambda > 0$, and $\eta \in (0, 1)$. There exist a deterministic context sequence, a reset time τ, a pre-reset event $\mathcal{E}^-$ satisfying*

$$\Pr(\mathcal{E}^-) \geq 1 - \eta, \tag{34}$$

and two piecewise-stationary audit-competence environments P_+ and P_- such that the following statements hold.

(i) The two environments induce the same distribution of every router-observable quantity before τ.

(ii) On $\mathcal{E}^-$,

$$G_t^{\mathrm{all}} \succeq \varepsilon^{-1} I_d \tag{35}$$

at every prescribed post-reset decision round $t \geq \tau$ before a post-reset audit is observed. Consequently,

$$\sup_{\|x\|_2 \leq 1} x^\top \left(G_t^{\mathrm{all}}\right)^{-1} x \leq \varepsilon. \tag{36}$$

(iii) At the current context, expert 1 is uniquely optimal under P_+, expert 2 is uniquely optimal under P_-, and selecting the wrong expert incurs audit-competence excess 2Δ.

(iv) Whenever

$$\Pr(A_{\tau:t-1} = 0) > 0,$$

every possibly randomized causal routing rule satisfies

$$\sup_{P \in \{P_+, P_-\}} \mathbb{E}_P\left[r_t^{\mathrm{audit}}(I_t) \,\middle|\, \mathcal{E}^-, A_{\tau:t-1} = 0\right] \geq \Delta. \tag{37}$$

For every $t \geq \tau$, including the cases where the conditioning event in part (iv) has probability zero,

$$\sup_{P \in \{P_+, P_-\}} \mathbb{E}_P[r_t^{\mathrm{audit}}(I_t)] \geq \Delta(1-p)^{t-\tau}. \tag{38}$$

Proof of Theorem 5. Let $e_1, \ldots, e_d$ be the canonical basis of $\mathbb{R}^d$. Before the reset, present each basis vector M times in a deterministic order. Under both environments, use identical pre-reset parameters

$$\theta_1^- = \theta_2^- = \frac{1}{2} e_1. \tag{39}$$

The complete pre-reset observable histories therefore have the same law under P_+ and P_-.

Let

$$N_i = \sum_{s<\tau} A_s \mathbf{1}\{x_s = e_i\},$$

and set

$$n_\varepsilon = \max\left\{0, \left\lceil \varepsilon^{-1} - \lambda \right\rceil\right\}.$$

By choosing M sufficiently large,

$$\mathcal{E}^- = \left\{ \min_{1 \leq i \leq d} N_i \geq n_\varepsilon \right\} \tag{40}$$

has probability at least $1 - \eta$. On this event,

$$G_t^{\mathrm{all}} = \lambda I_d + \sum_{i=1}^{d} N_i e_i e_i^\top \succeq \varepsilon^{-1} I_d$$

before any post-reset audit is incorporated. This proves parts (i) and (ii).

Take every relevant post-reset context to be e_1. Under P_+, define

$$\theta_1^+ = \left(\frac{1}{2} - \Delta\right) e_1, \qquad \theta_2^+ = \left(\frac{1}{2} + \Delta\right) e_1,$$

and under P_- exchange these two parameters. Thus the post-reset competence vectors at e_1 are

$$\left(\frac{1}{2} - \Delta,\ \frac{1}{2} + \Delta\right) \quad \text{under } P_+,$$

and

$$\left(\frac{1}{2} + \Delta,\ \frac{1}{2} - \Delta\right) \quad \text{under } P_-.$$

This proves part (iii).

Now condition on

$$\mathcal{E}^- \cap \{A_{\tau:t-1} = 0\},$$

whenever this event has positive probability. No post-reset competence label has been revealed, so the two environments induce the same conditional law of all information available before I_t is selected. Let q_t be the common conditional probability that the router chooses expert 1. Its conditional expected excess is

$$2\Delta(1 - q_t) \quad \text{under } P_+,$$

and

$$2\Delta q_t \quad \text{under } P_-.$$

Therefore,

$$\max\{2\Delta(1 - q_t), 2\Delta q_t\} \geq \Delta,$$

which proves part (iv).

For the unconditional statement, conditioning on $\mathcal{E}^-$ is unnecessary. On $\{A_{\tau:t-1} = 0\}$, the complete observable-history laws are identical under the two environments, whether or not historical coverage is large. The two conditional expected excesses sum to 2Δ. Since $\Pr(A_{\tau:t-1} = 0) = (1-p)^{t-\tau}$ and excesses are nonnegative,

$$\mathbb{E}_{P_+} r_t^{\text{audit}}(I_t) + \mathbb{E}_{P_-} r_t^{\text{audit}}(I_t) \geq 2\Delta(1-p)^{t-\tau}.$$

Taking half this sum proves the unconditional lower bound. When $p = 1$ and $t > \tau$, the asserted right-hand side is zero. At $t = \tau$, the empty no-audit event has probability one, including when $p = 1$.

This construction also gives the full statement of Theorem 1. On $\mathcal{E}^-$, $G_\tau^{\text{all}} \succeq \varepsilon^{-1} I_d$, and subsequent audit updates are positive semidefinite. Thus the same coverage inequality holds for every $t \geq \tau$, including after post-reset audits. For a segment of length H, summing the preceding two-environment inequality before taking its maximum gives

$$\frac{1}{2} \sum_{P \in \{P_+, P_-\}} \mathbb{E}_P \sum_{j=0}^{H-1} r_{\tau+j}^{\text{audit}}(I_{\tau+j}) \geq \Delta \sum_{j=0}^{H-1} (1-p)^j = \Delta \frac{1-(1-p)^H}{p}.$$

The maximum of the two cumulative expectations is at least their average, which proves the cumulative statement in Theorem 1. In particular, this argument does not interchange a sum with a maximum over environments. □

The same pair of environments gives a cumulative waiting cost. The proof averages the two environments before summing, so it does not interchange a supremum and a time sum.

Corollary 6 (Pre-first-audit cumulative cost)**.** *Under the construction of Theorem 5, for every $H \geq 1$,*

$$\sup_{P \in \{P_+, P_-\}} \sum_{j=0}^{H-1} \mathbb{E}_P[r_{\tau+j}^{\text{audit}}(I_{\tau+j})] \geq \Delta \frac{1-(1-p)^H}{p}. \tag{41}$$

Proof of Corollary 6. For each j such that $\Pr(A_{\tau:\tau+j-1} = 0) > 0$, the conditional action probability is common to P_+ and P_-, and the two conditional excesses sum to 2Δ. If this event has probability zero, the corresponding lower-bound contribution is zero and the same inequality remains true.

Consequently,

$$\begin{aligned} \mathbb{E}_{P_+}[r_{\tau+j}^{\text{audit}}(I_{\tau+j})] + \mathbb{E}_{P_-}[r_{\tau+j}^{\text{audit}}(I_{\tau+j})] \\ \geq 2\Delta(1-p)^j. \end{aligned} \tag{42}$$

Summing over $j = 0, \ldots, H-1$ and using $\max\{a, b\} \geq (a+b)/2$ for the two cumulative excesses gives

$$\sup_{P \in \{P_+, P_-\}} \sum_{j=0}^{H-1} \mathbb{E}_P[r_{\tau+j}^{\text{audit}}(I_{\tau+j})] \geq \Delta \sum_{j=0}^{H-1} (1-p)^j.$$

Evaluating the geometric sum proves (41). □

Interpretation. The obstruction in Theorem 5 is not insufficient historical feature coverage: u_t^{all} can be made arbitrarily small. Rather, the accumulated competence observations characterize the pre-reset state θ_m^-, whereas the current routing decision depends on the post-reset state θ_m^+. Before informative post-reset competence evidence arrives, the two possible current expert orderings remain statistically indistinguishable from the router's observable history.

Thus, arbitrarily small all-history geometric uncertainty does not imply identifiability of the current competence ordering. Historical coverage and current competence identifiability are distinct once competence can change over time.

D Fresh Competence Geometry

This appendix establishes the positive side of the stale–fresh separation. The negative result in the main text shows that arbitrarily strong all-history coverage need not identify the current competence ordering after a reset. We now show that, under stationary post-reset relative competence, geometric coverage supported by observations from the current regime controls routing error.

The reset time is assumed known only for this analysis. The operational CCR router is not given change points; its multiple memory horizons are designed to approximate useful notions of freshness without such information.

D.1 Fresh geometry controls routing error

Fix a prespecified reset time τ. Define the post-reset design matrix and local leverage

$$G_{t,\tau}^{+} = \lambda I_d + \sum_{s=\tau}^{t-1} A_s x_s x_s^\top, \tag{43}$$

$$u_{t,\tau}^{+} = x_t^\top (G_{t,\tau}^{+})^{-1} x_t. \tag{44}$$

For each expert m, the corresponding fresh ridge estimator is

$$\widehat{\theta}_{m,t}^{+} = (G_{t,\tau}^{+})^{-1} \sum_{s=\tau}^{t-1} A_s x_s Y_{m,s}^{\mathrm{audit}}, \tag{45}$$

with predicted competence

$$\widehat{\mu}_{m,t}^{+} = x_t^\top \widehat{\theta}_{m,t}^{+}.$$

The fresh selector is

$$I_t^{+} \in \arg\min_{m \in \mathcal{M}} \widehat{\mu}_{m,t}^{+}. \tag{46}$$

Assume stationary post-reset relative competence:

$$\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}} = x_t^\top v_{ij}^{+}, \qquad t \geq \tau,$$

with

$$\|v_{ij}^{+}\|_2 \leq S_{\mathrm{rel}},$$

and conditionally sub-Gaussian contrast noise with parameters σ_{ij}. Let

$$P = \frac{K(K-1)}{2}, \qquad \sigma_{\mathrm{rel}} = \max_{i<j} \sigma_{ij}.$$

Define

$$b_{\tau,t}^{\mathrm{rel}}(\delta) = \sigma_{\mathrm{rel}} \sqrt{\log \frac{\det G_{t,\tau}^{+}}{\lambda^d} + 2 \log \frac{P}{\delta}} \sqrt{u_{t,\tau}^{+}} + \lambda S_{\mathrm{rel}} \left\| (G_{t,\tau}^{+})^{-1} x_t \right\|_2. \tag{47}$$

Theorem 7 (Fresh geometry controls relative-competence routing)**.** *Under the assumptions above, for every $\delta \in (0,1)$, with probability at least $1-\delta$, simultaneously for all $t \geq \tau$,*

$$r_t^{\mathrm{audit}}(I_t^+) \leq b_{\tau,t}^{\mathrm{rel}}(\delta). \tag{48}$$

If the one-round audit excess is bounded by one, then for every prescribed round,

$$\mathbb{E} r_t^{\mathrm{audit}}(I_t^+) \leq \mathbb{E} \min\{1, b_{\tau,t}^{\mathrm{rel}}(\delta)\} + \delta. \tag{49}$$

Proof of Theorem 7. We specialize Theorem 4 to the single prespecified left endpoint τ with proof weight $\pi_\tau = 1$. Stationary post-reset relative competence and zero relative approximation residual make the transported drift term vanish. Therefore, simultaneously for all $t \geq \tau$ and all expert pairs,

$$\left|(\widehat{\mu}_{i,t}^+ - \widehat{\mu}_{j,t}^+) - (\mu_{i,t}^{\mathrm{audit}} - \mu_{j,t}^{\mathrm{audit}})\right| \leq b_{\tau,t}^{\mathrm{rel}}(\delta).$$

Let m_t^{audit} minimize the current audit competence. Empirical optimality of I_t^+ gives

$$\widehat{\mu}_{I_t^+,t}^+ - \widehat{\mu}_{m_t^{\mathrm{audit}},t}^+ \leq 0.$$

The corresponding true contrast is nonnegative and is bounded above by the simultaneous pairwise estimation error, which proves (48). On the confidence event, unit-bounded excess is at most $\min\{1, b_{\tau,t}^{\mathrm{rel}}(\delta)\}$; on the complement, it is at most one. Taking expectations proves (49). □

Theorem 7 does not require a fixed post-reset winner. It controls the excess of the fresh selector whenever the relative-competence model is stationary. The following consequence makes the decision interpretation explicit.

Corollary 8 (Exact winner recovery under a margin)**.** *On the simultaneous confidence event of Theorem 7, suppose the current audit winner m_t^{audit} satisfies*

$$\Delta_t^{\mathrm{audit}} = \min_{m \neq m_t^{\mathrm{audit}}} \left(\mu_{m,t}^{\mathrm{audit}} - \mu_{m_t^{\mathrm{audit}},t}^{\mathrm{audit}}\right) > b_{\tau,t}^{\mathrm{rel}}(\delta).$$

Then $I_t^+ = m_t^{\mathrm{audit}}$.

Proof of Corollary 8. Suppose, to the contrary, that $I_t^+ \neq m_t^{\mathrm{audit}}$. Empirical optimality gives

$$\widehat{\mu}_{I_t^+,t}^+ - \widehat{\mu}_{m_t^{\mathrm{audit}},t}^+ \leq 0.$$

On the simultaneous confidence event of Theorem 7, the pairwise contrast error is at most $b_{\tau,t}^{\mathrm{rel}}(\delta)$, so

$$\mu_{I_t^+,t}^{\mathrm{audit}} - \mu_{m_t^{\mathrm{audit}},t}^{\mathrm{audit}} \leq b_{\tau,t}^{\mathrm{rel}}(\delta).$$

The definition of $\Delta_t^{\mathrm{audit}}$ gives the opposite strict inequality

$$\mu_{I_t^+,t}^{\mathrm{audit}} - \mu_{m_t^{\mathrm{audit}},t}^{\mathrm{audit}} \geq \Delta_t^{\mathrm{audit}} > b_{\tau,t}^{\mathrm{rel}}(\delta),$$

a contradiction. Therefore $I_t^+ = m_t^{\mathrm{audit}}$. □

Theorem 7 does not require a fixed post-reset winner. It controls the excess of the fresh selector whenever the relative-competence model is stationary. If one expert remains optimal with a positive margin, the same confidence event additionally yields identification of that expert once the radius falls below the margin.

In particular, if the current audit winner m_t^{audit} satisfies

$$\Delta_t^{\mathrm{audit}} = \min_{m \neq m_t^{\mathrm{audit}}} \left(\mu_{m,t}^{\mathrm{audit}} - \mu_{m_t^{\mathrm{audit}},t}^{\mathrm{audit}}\right) > b_{\tau,t}^{\mathrm{rel}}(\delta),$$

then $I_t^+ = m_t^{\mathrm{audit}}$ on the simultaneous confidence event. Thus, fresh geometry controls not only excess competence risk but also recovery of the current ordering whenever the remaining uncertainty is smaller than the decision margin.

D.2 RANDOM AUDITS AND DIRECTIONAL COVERAGE

The preceding theorem is expressed directly through realized fresh leverage. To make the effect of sparse auditing explicit, we now give a sufficient condition under which this leverage decreases with elapsed post-reset time.

Suppose the post-reset contexts lie in a fixed subspace $E \subseteq \mathbb{R}^d$ of dimension r, with orthogonal projection Π_E, and satisfy the conditional excitation condition

$$\mathbb{E}\left[x_s x_s^\top \mid \mathcal{F}_{s-1}\right] \succeq \kappa \Pi_E, \qquad s \geq \tau, \tag{50}$$

for $0 < \kappa \leq 1/r$. This condition is used only to derive an explicit elapsed-time bound; the geometric result above applies directly to the realized design without it.

Lemma 9 (Audited directional coverage)**.** *Under* (50) *and exogenous Bernoulli auditing, for every prescribed $n \geq 1$,*

$$\Pr\left(\sum_{s=\tau}^{\tau+n-1} A_s x_s x_s^\top \Big|_E \not\succeq \frac{1}{2} p\kappa n I_E\right) \leq r \exp\Big(-\frac{p\kappa n}{8}\Big). \tag{51}$$

Proof of Lemma 9. Let

$$M_s = (A_s x_s x_s^\top)|_E.$$

Because $\|x_s\|_2 \leq 1$, we have $0 \preceq M_s \preceq I_E$. Audit exogeneity and (50) imply

$$\mathbb{E}[M_s \mid \mathcal{F}_{s-1}] \succeq p\kappa I_E,$$

and hence

$$\lambda_{\min}\left(\sum_{s=\tau}^{\tau+n-1} \mathbb{E}[M_s \mid \mathcal{F}_{s-1}]\right) \geq p\kappa n.$$

The matrix Chernoff inequality for adapted positive-semidefinite sequences (Tropp, 2011, Theorem 3.1), applied with matrix dimension r, upper bound $R = 1$, lower-mean parameter $\mu = p\kappa n$, and deviation $1/2$, gives

$$\begin{aligned}\Pr\left(\lambda_{\min}\left(\sum_{s=\tau}^{\tau+n-1} M_s\right) < \frac{1}{2} p\kappa n\right) \\ \leq r\left(\frac{e^{-1/2}}{(1/2)^{1/2}}\right)^{p\kappa n} \leq r e^{-p\kappa n/8},\end{aligned}$$

where the last inequality follows from $\frac{1}{2}(1 - \log 2) \geq \frac{1}{8}$. This is (51). □

The important quantity here is directional information, not the number of audits alone. A large audit count concentrated in irrelevant context directions need not reduce uncertainty at the current context.

Corollary 10 (Recovery under random audits)**.** *Under Theorem 7 and* (50)*, let*

$$n = t - \tau, \qquad v_n = \lambda + \frac{1}{2} p\kappa n,$$

and define

$$B_n(\delta) = \sigma_{\mathrm{rel}} \sqrt{r \log\left(1 + \frac{n}{\lambda r}\right) + 2 \log \frac{P}{\delta}}.$$

Then, with probability at least

$$1 - \delta - r e^{-p\kappa n/8},$$

$$r_t^{\mathrm{audit}}(I_t^+) \leq \frac{B_n(\delta)}{\sqrt{v_n}} + \frac{\lambda S_{\mathrm{rel}}}{v_n}. \tag{52}$$

Proof of Corollary 10. On the event of Lemma 9,

$$G_{t,\tau}^{+}\big|_{E} \succeq v_n I_E.$$

Since $x_t \in E$,

$$u_{t,\tau}^{+} \le v_n^{-1}, \qquad \left\|(G_{t,\tau}^{+})^{-1} x_t\right\|_2 \le v_n^{-1}.$$

The unregularized design has rank at most r and trace at most n, so

$$\log \frac{\det G_{t,\tau}^{+}}{\lambda^d} \le r \log\left(1 + \frac{n}{\lambda r}\right).$$

Substitution into (47), followed by a union bound with the directional-coverage event, proves the result. □

The corollary makes the sparse-feedback penalty explicit. Up to logarithmic and problem-dependent factors, the information accumulated after n elapsed rounds scales with pn. Consequently, reducing the audit probability increases the elapsed time required to obtain comparable fresh directional coverage by the characteristic $1/p$ factor.

This statement should not be read as saying that audit count alone determines recovery. The factor p controls how often information is revealed, while κ controls whether the revealed contexts span directions relevant to the current competence decision.

D.3 AUDIT-TO-EVALUATION TRANSFER

The theoretical bounds above concern audit competence, whereas the benchmark is evaluated using disjoint evaluation targets. Because routing depends on relative expert competence, the two risk systems need only agree up to an expert-independent offset.

Assume

$$\sup_{1\le t\le T} \inf_{c\in\mathbb{R}} \max_{m\le K} \left|\mu_{m,t}^{\mathrm{audit}} - \mu_{m,t}^{\mathrm{eval}} - c\right| \le \delta_{\mathrm{audit}}. \tag{53}$$

Proposition 11 (Audit-to-evaluation transfer)**.** *For*

$$a_{m,t} = \mu_{m,t}^{\mathrm{eval}} - \mu_{m,t}^{\mathrm{audit}}, \qquad \omega_t = \max_m a_{m,t} - \min_m a_{m,t},$$

every action I_t satisfies

$$r_t^{\mathrm{eval}}(I_t) \le r_t^{\mathrm{audit}}(I_t) + \omega_t. \tag{54}$$

Moreover,

$$\omega_t = 2 \inf_{c\in\mathbb{R}} \max_m |a_{m,t} - c|. \tag{55}$$

Hence (53) *implies*

$$r_t^{\mathrm{eval}}(I_t) \le r_t^{\mathrm{audit}}(I_t) + 2\delta_{\mathrm{audit}}.$$

Proof of Proposition 11. For every action i,

$$\mu_{i,t}^{\mathrm{eval}} \le \mu_{i,t}^{\mathrm{audit}} + \max_m a_{m,t},$$

whereas

$$\min_m \mu_{m,t}^{\mathrm{eval}} \ge \min_m \mu_{m,t}^{\mathrm{audit}} + \min_m a_{m,t}.$$

Subtracting gives (54).

For any c,

$$\max_m |a_{m,t} - c| \ge \frac{1}{2}\left(\max_m a_{m,t} - \min_m a_{m,t}\right),$$

with equality when c is the midpoint of the largest and smallest $a_{m,t}$. This proves (55) and the stated $2\delta_{\mathrm{audit}}$ consequence. □

The centered formulation is weaker than requiring absolute agreement between audit and evaluation risks. An expert-independent discrepancy between the two loss systems has no effect on their relative ordering and therefore need not be penalized by a competence-routing guarantee. The sharper cumulative transfer penalty is $\sum_t \omega_t$.

D.4 Scope of the freshness guarantee

The results above establish a deliberately focused statement. Historical coverage can be strong while the associated competence labels are stale; under stationary post-reset relative competence, fresh directional coverage controls the prediction and routing error of a known-reset estimator. Conditional excitation then makes explicit how sparse random audits slow the accumulation of such information.

These guarantees do not assume that the practical router knows the reset time. Nor do they imply that leverage itself detects a competence change, that the operational certification rule is a formal confidence procedure, or that the complete CCR deployment policy satisfies a dynamic-regret guarantee under arbitrary drift. The practical method instead maintains several finite competence memories and uses their observed audit performance to approximate useful freshness without access to the analytical reset oracle.

The theory therefore supports the mechanism tested empirically in the main paper: accumulated geometric coverage and currently relevant competence evidence are distinct objects, and routing after competence drift depends on the latter.

E Dataset Preprocessing and Frozen Splits

This appendix specifies the preprocessing and data-splitting protocol used throughout the experiments. All partitions, normalization statistics, controlled reconstruction targets, and audit realizations are fixed before online evaluation. Test evaluation targets are never used for expert fitting, representation learning, router selection, certification calibration, drift construction, or audit generation.

E.1 Common protocol

Split integrity. General time-series datasets use fixed chronological train/validation/test partitions. Clinical datasets are partitioned by the relevant patient or stay identifier so that identities do not cross partitions. Whenever temporal ordering is available, future test observations cannot affect preprocessing, expert fitting, representation parameters, or router configuration.

Training-only preprocessing. Channel definitions, normalization statistics, and trainable preprocessing components are estimated from the training partition only. Validation data are used for checkpoint selection and for hyperparameters that must be fixed before online test evaluation. No test evaluation loss is used for model or router selection.

Paired evaluation streams. Within each dataset, intervention family, and random seed, competing methods receive the same source stream, controlled reconstruction targets, frozen expert predictions, audit activations, and revealed audit-loss vectors. Evaluation-only quantities, including competence-oracle identities, reversal annotations, and competence-effect labels, are computed only after the corresponding online decisions and are never available to the router.

E.2 Clinical streams

MIMIC-IV. We use the first 48 hours of each eligible adult stay at hourly resolution and retain four physiological variables: heart rate (HR), oxygen saturation (SpO_2), respiratory rate (RR), and mean arterial pressure (MAP). The reported cohort contains $1{,}640$ stays.

We consider two routing granularities. In the default stay-level protocol, one routing decision is associated with each stay. In the hour-level protocol, routing proceeds over hourly observations within a stay and the router state is reset between stays. These protocols prevent competence state from being carried across unrelated patients.

The competence representation is evaluated with both a base context and a predeclared enriched context. The enriched representation augments the base features with observable information describing coverage and gap structure, expert disagreement, and audit dynamics. It contains no competence-oracle, evaluation-loss, reversal, or other evaluation-only features.

Naturally missing clinical measurements are preserved as part of the observed stream and are never treated as ground-truth reconstruction targets. Controlled audit and evaluation targets are created only at locations where a source measurement is observed. Thus, natural missingness defines the clinical observation regime, whereas controlled hiding provides evaluable reconstruction targets within that regime.

PhysioNet/CinC 2019. We use Hospital B from the PhysioNet/CinC 2019 release (Reyna et al., 2019) as the second clinical environment. The reported evaluation cohort contains $2{,}000$ stays. Results summarized in the main experiments are aggregated over three predefined shards, with the same frozen preprocessing, intervention, and audit protocol applied across competing methods.

E.3 General time-series streams

ETTh1 and Weather. ETTh1 (Zhou et al., 2021) and Weather are transformed into multivariate imputation windows of length $L = 48$ using fixed chronological partitions. The reported evaluation cohorts contain 253 ETTh1 windows and 767 Weather windows.

A window contributes evaluation loss to a partition only according to the deterministic boundary-handling rule implemented by the corresponding dataset adapter. In particular, no future-partition target or observation is exposed to an earlier partition. The same realized windows and controlled interventions are shared across routing methods.

Electricity. UCI ElectricityLoadDiagrams (Trindade, 2015) is transformed using the same chronological windowing protocol. Electricity is used as an exploratory extension of the primary general time-series panel and is reported separately from the primary ETTh1 and Weather results.

E.4 Expert fitting and online freeze

All candidate imputers are fitted offline using only their designated training data, with checkpoint and hyperparameter selection restricted to validation data. Once selected, expert parameters remain fixed throughout the online test stream. The primary competence representation is likewise frozen during the reported routing experiments.

Consequently, online adaptation is confined to the router's competence memories, scale weights, and deployment state. Changes in expert performance therefore arise from interaction between the frozen imputers and the evolving stream rather than from online expert retraining.

The released experiment manifests record the split identifiers, random seeds, expert checkpoints, and other configuration information required to reproduce each reported evaluation. Additional expert and implementation details are given in Appendix F.1.

F Implementation and Hyperparameters

This appendix specifies the implementation and hyperparameter configuration used in the reported experiments. All expert parameters, representation parameters, and routing hyperparameters are fixed before online test evaluation. Unless explicitly stated otherwise, only competence memories, scale weights, and deployment state are updated online.

F.1 Candidate expert library

The reported experiments use the frozen expert library

$$\mathcal{M} = \{\mathrm{Linear}, \mathrm{SAITS}, \mathrm{T1}\}. \tag{56}$$

The Linear expert provides a deterministic low-complexity reference. SAITS and T1 are trained offline with sequence length $L = 48$, batch size 8, and 20 epochs. Checkpoint selection uses validation data only, after which the selected expert parameters remain fixed throughout the corresponding test stream. The released experiment artifacts record the training configuration, random seed, and checkpoint checksum associated with each expert.

F.2 ROUTING CONFIGURATION

The default multi-scale routing configuration is

Parameter	Value
Memory horizons $\mathcal{H}$	$\{8, 16, 32, 64, 128\}$
Ridge regularization λ	0.1
Scale-learning rate η	0.05

For the reported clinical routing experiments, the remaining parameters are fixed to

$$\mathrm{sh} = 2, \qquad \kappa = 2.5, \qquad \alpha = 0.5, \qquad \gamma_{\mathrm{sw}} = 0.05,$$

where sh controls scale-weight sharpening, κ scales the geometric uncertainty contribution, α controls the certification threshold, and γ_{sw} is the additional margin required to authorize a change from the currently deployed expert. The same configuration is retained across the reported clinical audit probabilities.

Audit probabilities. The primary experiments use $p \in \{1.0, 0.2\}$. Clinical sparse-audit experiments additionally include $p = 0.5$, while $p = 0.1$ is reserved as a stress setting. Audit activations and controlled reconstruction targets are generated before routing and shared across competing methods within each paired experimental realization.

F.3 ONLINE AND FROZEN STATE

At round t, the online state consists of the per-horizon ridge competence estimates, scale weights $q_t(h)$, aggregate competence and uncertainty scores, and the certification and deployment state. These quantities may change only through information revealed according to the causal audit protocol.

All expert checkpoints, channel definitions, data splits, representation parameters used to construct $x_t = \phi(Z_t)$, and controlled audit and evaluation masks are frozen before test evaluation. In particular, test evaluation losses, competence-oracle identities, reversal annotations, and other evaluation-only quantities never enter the online state.

Certification and switch authorization are implemented as separate decisions. A candidate that satisfies the certification criterion replaces the incumbent only when it also clears the prescribed switch-margin condition. A candidate whose switch is blocked does not become the fallback certified deployment state indirectly.

F.4 REPRODUCIBILITY

Each reported run is uniquely specified by its dataset, intervention family, audit probability, stream granularity, competence-context configuration when applicable, random seed, expert checkpoints, and router configuration. Competing methods within a paired comparison use the same source stream, intervention realization, expert predictions, audit activations, and controlled evaluation targets.

The released manifests record these configurations together with the corresponding split identifiers and checkpoint checksums. Reproducing a reported result therefore uses the associated frozen configuration rather than a hyperparameter search performed after observing test performance.

G EXPERIMENT DETAILS

This appendix provides additional protocol details and full numeric panels for the experiments in Section 6. All comparisons use frozen expert outputs and paired streams: within each dataset, intervention, audit probability, and seed, competing routers receive identical contexts, expert predictions, audit activations, and evaluation targets.

G.1 RQ1: COMPETENCE DRIFT

RQ1 asks whether the identity of the most competent frozen imputer changes under controlled non-stationarity. For each dataset and intervention family, we compare the best static expert with the

Table 3: **RQ1 competence-drift panel (detail).** Primary G1/G2 gate across clinical + TS. PhysioNet entries are means over 3 shards. n: audit-cohort size. Companion to the main-text summary in Table 1.

Dataset	Family	Static	Headroom (%)	Strong	n	Gate
Clinical (G1/G2)						
MIMIC-IV	none	T1	7.68	53	1640	PASS
MIMIC-IV	D1	T1	10.43	60	1640	PASS
MIMIC-IV	D2	T1	10.80	68	1640	PASS
MIMIC-IV	D3	T1	9.46	53	1640	PASS
MIMIC-IV	D6	T1	11.67	55	1640	PASS
PhysioNet-2019B	none	T1	13.48	84	2000	PASS
PhysioNet-2019B	D1	T1	15.82	93	2000	PASS
PhysioNet-2019B	D2	T1	16.61	93	2000	PASS
PhysioNet-2019B	D3	T1	15.29	88	2000	PASS
PhysioNet-2019B	D6	T1	17.44	89	2000	PASS
Time-series (G1/G2)						
ETTh1	none	T1	9.08	3	253	PASS
ETTh1	D1	T1	16.29	12	253	PASS
ETTh1	D2	SAITS	17.99	9	253	PASS
ETTh1	D3	T1	15.22	11	253	PASS
ETTh1	D6	T1	16.41	6	253	PASS
Weather	none	Linear	18.97	10	767	PASS
Weather	D1	Linear	17.21	7	767	PASS
Weather	D2	Linear	26.71	13	767	PASS
Weather	D3	Linear	22.16	8	767	PASS
Weather	D6	Linear	17.41	11	767	PASS

Table 4: **RQ1 Electricity exploratory pilot.** Smaller eligible-boundary pilot than ETTh1/Weather; not part of the primary G1/G2 panel in Table 3.

Dataset	Family	Static	Headroom (%)	Strong	*n/elig.*	Gate
Electricity*	D1	T1	10.40	5	15	PASS
Electricity*	D2	T1	7.03	10	15	PASS
Electricity*	D3	T1	0.003	0	80	FAIL
Electricity*	D6	T1	4.87	15	20	soft PASS‡

*Smaller eligible-boundary pilot than ETTh1/Weather. ‡Headroom 4.87% just below the nominal 5% cut; G2 clears, hence soft PASS.

evaluation-only competence oracle. Relative oracle headroom is computed as the improvement of the competence oracle over the best static expert under the same evaluation targets.

A competence reversal is recorded only from evaluation-side expert competences after the stream has been generated. These annotations are never available to the router. In addition to aggregate headroom, we report recurrent strong reversals to distinguish persistent changes in relative competence from isolated changes caused by small or nearly tied loss differences.

The intervention families D1–D6 are fixed independently of the resulting expert ordering. Consequently, an intervention is not assumed to induce a competence reversal. Episodes in which the ordering is preserved are retained as negative controls rather than removed from the evaluation. The D7 preserving/reversing taxonomy is assigned only after evaluation and is not an additional stream intervention.

Electricity sensitivity. Electricity is treated as an exploratory extension because its competence effects are less uniform than those of the primary ETTh1 and Weather panel. The reported results therefore retain both reversal-producing and ordering-preserving interventions rather than selecting only conditions that favor adaptive routing.

Table 5: **RQ1 sensitivity (stricter frozen / eligible-boundary).** Not used for the primary G1/G2 verdict in Table 3.

Dataset	Family	Static	Headroom (%)	Strong	*n/elig.*	Gate
ETTh1	D1	SAITS	4.13	7	15	FAIL
ETTh1	D2	SAITS	2.27	5	15	FAIL
ETTh1	D3	SAITS	0.73	3	35	FAIL
ETTh1	D6	SAITS	1.69	8	20	FAIL
Weather	D1	SAITS	25.45	0	15	FAIL
Weather	D2	SAITS	13.67	9	15	PASS
Weather	D3	SAITS	0.60	8	105	PASS
Weather	D6	SAITS	7.61	15	20	FAIL

G.2 RQ2: FRESH VERSUS ACCUMULATED GEOMETRY

RQ2 tests whether geometric support from recent audited observations is more informative about current competence than coverage accumulated over the full history. We compare the all-history leverage

$$u_t^{\mathrm{all}}$$

with the deployed multi-scale surrogate

$$u_t^{\mathrm{dep}} = \sum_{h \in \mathcal{H}} q_t(h) u_t(h).$$

The primary diagnostic is

$$\Delta\mathrm{AUC} = \mathrm{AUC}_{\mathrm{fresh}} - \mathrm{AUC}_{\mathrm{all}},$$

where the binary target indicates mismatch between the currently deployed expert and the evaluation-side competence oracle. Positive values indicate that the finite-memory geometric score better discriminates rounds on which the deployed expert is no longer competence-optimal.

The oracle-reset leverage $u_{t,\tau}^{+}$ used in the theoretical analysis is never exposed to the router and is not substituted for the deployed multi-scale quantity in the primary empirical comparison.

Clinical geometry variants. Clinical experiments consider two stream granularities and two causal competence representations. In the stay-level protocol, one routing decision is associated with each stay. In the hour-level protocol, routing proceeds within a stay and router state is reset before the next stay.

The base representation uses the standard causal router context. The enriched representation additionally contains predeclared observable features describing coverage and gap structure, expert disagreement, and audit dynamics. It contains no evaluation losses, competence-oracle identities, reversal annotations, or future information.

The stay-level base representation at $p = 1$ serves as the reference configuration. The additional variants are used to assess whether the freshness diagnostic is sensitive to temporal resolution and to how strongly the representation exposes competence-relevant geometry.

G.3 RQ3: ROUTING UNDER SPARSE FEEDBACK

RQ3 evaluates whether finite and multi-scale competence memories can exploit sparse post-decision audits without modifying the frozen imputers. We compare the best static expert, fixed-memory routing, multi-scale routing (CCR-MS), and the certification variant (CCR-Cert) where available. All adaptive methods use the same audit realizations within each paired run.

The multi-scale router combines horizon-specific competence estimates through the online weights $q_t(h)$. Scale weights are updated only after audit feedback becomes available. The certification variant uses the same competence estimates and scale weights as CCR-MS; it differs only in the decision layer governing whether a proposed change of deployed expert is authorized.

Table 6: **Failure-discrimination geometry.** ΔAUC is fresh minus all-history (positive favors fresh). Numbers underlying the main-text RQ2 discussion.

Dataset / protocol	Audit p	Units	Fresh	All-hist.	ΔAUC	Gate / 95% CI
Time-series panel (pooled)	1.0	8	0.6915	0.6590	0.0325	$[-0.0432,\ 0.1140]$
Time-series panel (pooled)	0.2	8	0.6296	0.6195	0.0100	$[-0.0802,\ 0.0981]$
ETTh1 (MS)	1.0	2	0.6213	0.5246	+0.0967	soft+strong PASS
Weather (MS)	1.0	2	0.5416	0.4364	+0.1052	soft+strong PASS
MIMIC-IV (stay, base)	1.0	1	0.4830	0.5009	-0.0179	FAIL
MIMIC-IV (stay, enriched)	1.0	1	0.5305	0.5210	+0.0095	soft PASS
MIMIC-IV (stay, base)	0.2	1	0.5030	0.4857	+0.0173	soft PASS
PhysioNet-2019B	1.0	6	0.5115	0.5060	+0.0054	soft PASS
PhysioNet-2019B	0.2	3	0.5060	0.4912	+0.0147	soft PASS

Primary clinical soft PASS at $p = 1$ is stay+enriched.

Table 7: **Additional MIMIC-IV geometry protocol cells.** Not used for the primary clinical soft-PASS call in Table 6.

Dataset / protocol	Audit p	Units	Fresh	All-hist.	ΔAUC	Gate
MIMIC-IV (hour, base)	1.0	1	0.4962	0.4961	+0.0001	soft PASS†

†Near-zero hour-level Δ; primary clinical soft PASS at $p = 1$ remains stay+enriched.

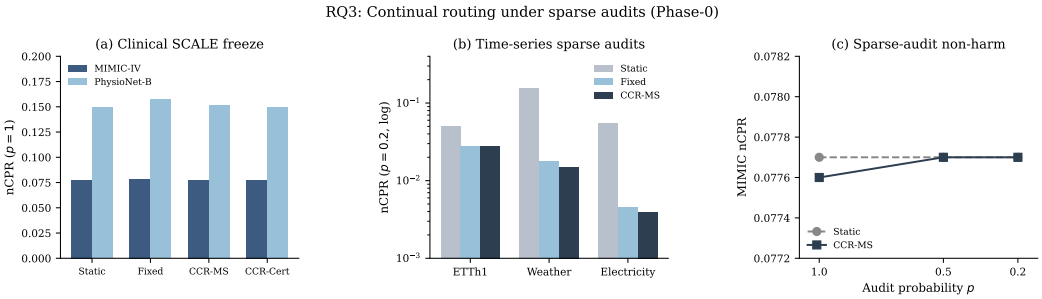


Figure 3: **Online adaptation of CCR under sparse audits.** (a) Clinical nCPR at $p = 1$ (Table 2). (b) Time-series nCPR at $p = 0.2$ on a log scale (Table 8). (c) MIMIC-IV sparse-audit non-harm across $p \in \{1.0, 0.5, 0.2\}$.

Table 8: **Extended nCPR panel (lower better).** TS panel at $p = 0.2$; clinical freeze at $p = 1$ and sparse $p \in \{0.5, 0.2\}$. Discounted baseline: TS protocol only. Companion to the main-text Table 2.

Dataset	Static	Fixed	Discounted	CCR-MS	CCR-Cert
Time-series panel @ $p = 0.2$					
ETTh1	0.0498	**0.0277**	0.0304	0.0280	—
Weather	0.1563	0.0180	0.0174	**0.0151**	—
Electricity	0.0552	0.0045	0.0043	**0.0039**	—
Clinical @ $p = 1$ (nCPR; OMA in parentheses)					
MIMIC-IV	0.0777 (0.327)	0.0784 (0.323)	—‡	**0.0776** (0.327)	0.0777 (0.327)
PhysioNet-2019B	0.1494 (0.240)	0.1574 (0.239)	—‡	0.1517 (0.239)	0.1497 (0.241)
Clinical sparse (nCPR)					
MIMIC-IV @ $p = 0.5$	0.0777	0.0785	—‡	0.0777	0.0777
MIMIC-IV @ $p = 0.2$	0.0777	0.0781	—‡	0.0777	0.0777
PhysioNet-2019B @ $p = 0.2$	0.1494	0.1505	—‡	0.1496	0.1494

‡Discounted baseline not run on clinical.

Certification and switching are distinct operations. A proposed expert must first satisfy the competence-separation criterion and, when it differs from the incumbent, must additionally exceed the prescribed switch margin γ_{sw}. A blocked candidate does not become the certified fallback state.

Sparse-audit protocol. The primary audit probabilities are $p = 1$ and $p = 0.2$, with $p = 0.5$ additionally evaluated for the clinical sparse-audit analysis. Audit masks are generated before routing and shared across competing methods. Thus changes in routing performance across audit probabilities reflect reduced feedback availability rather than different underlying streams or evaluation targets.

Gap-conditioned analysis. To separate adaptation near meaningful competence changes from behavior on the full stream, we additionally evaluate routing on a gap-conditioned subset defined entirely from evaluation-side quantities. This subset is used only for retrospective analysis and never determines whether the router switches.

On the full-audit clinical streams, the gap-conditioned results are nearly tied on MIMIC-IV ($\mathrm{nCPR}_{\mathrm{gap}} = 0.1172$ for CCR-MS versus 0.1173 for Static) and slightly favor Static on PhysioNet/CinC 2019 (0.1994 versus 0.1982). These results indicate that the existence of competence-changing regions does not by itself guarantee that the current routing rule can exploit them.

Certification behavior. Certification is evaluated as a conservative decision layer rather than as a separate competence estimator. On MIMIC-IV with full audits, CCR-Cert makes no certified switches and therefore matches the static deployment. This behavior illustrates the tradeoff observed in the current router: stricter switching can suppress harmful adaptation, but can also remove the adaptive benefit when evidence for a change is insufficient.

G.4 Paired comparisons and reporting

All method comparisons are paired at the stream level. A paired experimental cell fixes the dataset split, intervention realization, frozen expert outputs, audit targets, audit activations, evaluation targets, and random seed before the competing routing policies are evaluated. Router-specific randomness, if present, is recorded separately.

Primary metrics are computed over the complete declared evaluation stream. Oracle quantities, reversal labels, gap-conditioned subsets, and preserving/reversing annotations are used only for retrospective diagnosis. No result is selected or excluded according to whether adaptation improves over the static baseline.