WHEN DOES CAUSAL REPAIR ORDERING SURVIVE FEDERATED EVIDENCE FORMATION?

Anonymous authors
Paper under double-blind review

ABSTRACT

Federated causal analysis can have enough data overall to answer a target question while no individual client can form the evidence needed for a particular test. The issue matters when the same target can be settled by several alternative sets of conditional-independence (CI) tests, which we call exact repairs. Federation can reverse this centralised preference: even reallocating the same pooled observations can make a shorter repair less likely to form than a longer one. The gap can be arbitrarily large: a target solvable by one query structurally and after pooled query-statistic aggregation can require arbitrarily many locally formed queries, with the gap persisting through any prescribed bounded aggregation depth. Formation ease and target-decision value can also rank repairs differently. Under random federation, we characterise how route support demands and federation scaling determine how much aggregation is needed, and give a concrete condition that preserves the centralised ordering. Experiments on public Bayesian networks show shallow recovery with limited aggregation, cases where a more expensive repair attains higher local readiness than every cheaper alternative, and strong sensitivity to how an unchanged pooled sample is split across clients.

1 INTRODUCTION

A conclusion can often be supported in more than one way. Consider two hospitals that want to determine whether a treatment has a causal relation with an outcome. One route might require a single statistical test involving several patient groups, while another uses several simpler tests involving different subsets of the population. If all patient records were stored together, the one-test route would appear preferable. Across federated clients, however, the patients needed for that test may be split across hospitals, so neither hospital can form the required evidence on its own. The several-test route may still be usable because each of its tests can be formed locally, possibly at different hospitals. Both routes can fully resolve the same causal target, yet federation can make the longer route easier to realise locally. More surprisingly, the preferred route can change even when the underlying records are exactly the same and only their assignment to hospitals is changed. Thus, having the same information overall does not imply having the same usable causal evidence. *When does a centralised preference for a shorter sufficient evidence route survive federated evidence formation?*

We call each set of CI evidence sufficient to resolve the remaining target ambiguity an *exact repair*, and compare the complete family of inclusion-minimal exact repairs. Federated causal-discovery work studies limited local samples, heterogeneous client distributions, local evidence quality, and whether evidence is formed locally or after statistical aggregation (Wang et al., 2023; Li et al., 2024; Guo et al., 2024b; 2025; Hahn et al., 2026); related CI and query-complexity work studies difficulty of individual tests and central causal procedures (Jamshidi et al., 2024; Shiragur et al., 2024; Zhang et al., 2024). This leaves open whether the preference among alternative sufficient routes for one target survives a change in where their evidence is formed. We therefore compare two matched interfaces. In *local-evidence-first* formation, a whole CI decision must be supported within one client; in exact *aggregate-before-decision*, query-specific sufficient statistics are combined across clients before that decision is formed. Bounded aggregation connects the two. We ask how much cost federation can add and how much aggregation removes it; how federation changes the preferred repair, including disagreement between formability and target value; and which support conditions preserve the centralised preference.

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.

Federation can change the preferred repair without changing pooled information: ordering can reverse with the realised pooled sample fixed, or with the pooled population, client sizes, target task, and complete repair family fixed while only the client decomposition changes. The effect can be arbitrarily severe. With h residual ambiguity modules, costs can be $1{:}h{:}1$: one query structurally, h under client-local formation, and one after exact aggregate-before-decision. This distortion can persist through any prescribed bounded aggregation depth before deeper aggregation restores the structural preference. Under random federation, route support burden and client/sample scaling determine the aggregation depth needed to form the relevant evidence routes. We also show that formation and target value can rank repairs differently: structural cost, readiness, and target-decision order can differ, including a regime of target-decision incomparability. Finally, support-compatible dominance preserves structural ordering, while exact aggregate-before-decision recovers the pooled boundary.

Public Bayesian networks show these mechanisms in fixed causal tasks. Across four harmonised CHILD targets, allowing each query to use evidence from one or two clients makes a structurally cheapest repair match the best repair's readiness. Yet on CHILD/ChestXray a higher-cost repair attains higher local readiness than the complete frontier of cheaper repairs at all five pre-specified sample sizes. On CHILD/GruntingReport, reallocating the same pooled sample sharply changes local readiness while pooled readiness stays fixed. ALARM shows the complementary case, where alternative cheap repairs protect the cheaper frontier; INSURANCE marks a boundary where required evidence remains unsupported even after pooling. We make three contributions:

- **A one-query repair can require arbitrarily many locally formed queries.** We prove a $1{:}h{:}1$ structural/local/pooled separation and, for any prescribed recovery depth b_0, cost remains h for $b < b_0$ before dropping to 1. Recovery depth and distortion are independently selectable, with randomised lower bounds. Under random federation, we derive a formation-depth law for the relevant evidence routes.
- **Federation can change the preferred repair even when pooled information is fixed.** A structurally cheaper exact repair can be less ready with the realised pooled sample fixed, or with the pooled population, client count, and client sample sizes fixed while only the client decomposition changes. We further characterise a sharp target-decision trichotomy: readiness and uniform target advantage can rank repairs differently; at one matched point, the structurally more expensive repair is uniformly more informative for the fixed target with an arbitrarily large cost ratio. Public-network experiments show a higher-cost repair attaining higher local readiness than the complete frontier of cheaper alternatives.
- **A concrete support condition guarantees preservation of the centralised preference.** Under support-compatible dominance, if queries in the longer repair can be matched to no-harder queries in the shorter repair that cover at least the same target disagreements, then forming the longer repair on any declared interface also forms the shorter one. Exact query-specific aggregate-before-decision gives the pooled endpoint, a positive boundary complementary to the reversal results.

2 Related Work

Target-local causal recovery and CI-query complexity. Constraint-based and target-local causal discovery recover global or local structure from CI information, including Markov blankets, causal neighbourhoods, and target relations (Spirtes et al., 2000; Kalisch & Bühlmann, 2007; Aliferis et al., 2010; Gao & Ji, 2015; Loranchet & Assaad, 2026). Related work studies finite-sample CI testing and CI-query counts for causal discovery or equivalence-class testing (Jamshidi et al., 2024; Shiragur et al., 2024; Zhang et al., 2024). Uehara (2026) is closer in analysing adaptive causal certificates and oracle interactions, but its oracle directly supplies causal information, whereas our cost arises from forming valid CI evidence from client-local support. These lines establish central or oracle complexity for causal information. Our object is comparative: we study the full family of alternative exact repairs and ask whether their structural preference survives federated evidence formation.

Federated causal evidence formation. Federated causal-discovery methods combine scores, local structures, CI decisions, parameters, or sufficient statistics under heterogeneous, limited-sample, missing-data, and non-identical-variable settings (Ng & Zhang, 2022; Gao et al., 2023; Huang et al., 2023; Guo et al., 2024a; Li et al., 2024; Guo et al., 2025; Shi et al., 2025; Wang et al., 2025; Chen et al., 2026). FedACD (Guo et al., 2024b) shows that evidence quality can vary across variable spaces. FedC2SL (Wang et al., 2023) aggregates client statistics before global CI decisions, while fedCI (Hahn

et al., 2026) aggregates statistical evidence across heterogeneous data sets and can recover power that is weak locally. We use a bounded aggregate-before-decision interpolation to quantify how much query-specific cross-client aggregation is required before a structural repair preference reappears. This isolates how the evidence-formation interface reorders alternative complete repairs for the target.

Distributed inference, experiment comparison, and combinatorial representations. Distributed testing and estimation characterise how local information, communication budgets, interaction, and shared randomness alter statistical complexity (Acharya et al., 2019; 2020; 2022). Minimal transversals, minimal path sets, and coherent-system representations provide standard tools for coverage and reliability (Barlow & Proschan, 1975; Birnbaum, 1969; Marichal, 2014; Mary, 2024). For the stronger target-level comparison, we use the standard comparison-of-experiments viewpoint (Blackwell, 1953) only after nuisance causal worlds are marginalised under each fixed prior, and compare whole repair-induced experiments through target-dependent decision losses. Together with the complete repair family and support paths, this lets us separate structural repair cost, formation/readiness order, target-decision preorder, and support-compatible preservation under matched pooled information.

3 Problem Formulation

Target-Specific Repair Problem. We study finite discrete observed-variable causal tasks. Fix a target T, finite analysis scope $\mathcal{V}_T$, finite residual admissible DAG class $\mathcal{G}$, and target-summary map $\mathcal{L}_T$. The scope sets the analysis domain and need not equal the true causal neighbourhood of T. A supplied baseline transcript is fixed before repair: surfaced answers may restrict $\mathcal{G}$; abstention adds no restriction. The residual class represents the causal worlds remaining after the supplied baseline evidence and any declared background restrictions; throughout, the complete repair family is complete relative to this fixed residual class and admissible query universe.

Federated Evidence Setting. For a positive integer n, write $[n] := \{1, \dots, n\}$. A realised horizontal federated system has K clients with aligned observed variables and local tables $\mathcal{D}_k$. All admissible queries at client k reuse $\mathcal{D}_k$. Under candidate world G, rows are sampled i.i.d. from client law $P_{k,G}$; client populations and sample sizes may differ. Cross-client independence, homogeneity, and population scaling are imposed only where stated.

To avoid duplicate endpoint orderings, fix a total order $\prec$ on $\mathcal{V}_T$. For maximum conditioning order $r_{\max}$, define the target-relevant CI-query universe as $\mathcal{Q}_T^{(r_{\max})}(\mathcal{V}_T) := \{I = (X_a, X_b, S) : X_a \prec X_b,\ X_a, X_b \in \mathcal{V}_T,\ S \subseteq \mathcal{V}_T \setminus \{X_a, X_b\},\ |S| \le r_{\max}\}$. A query need not contain T. For $G \in \mathcal{G}$, write $d_G(I) := \mathbf{1}[X_a \perp\!\!\!\perp_G X_b \mid S]$; when stochastic client laws realise candidate world G, write $d_{P_{k,G}}(I) := \mathbf{1}[X_a \perp\!\!\!\perp_{P_{k,G}} X_b \mid S]$ for the corresponding population label. Fix the admissible universe $\mathcal{Q}_T \subseteq \mathcal{Q}_T^{(r_{\max})}(\mathcal{V}_T)$ before client support, observation allocation, or query outcomes; baseline queries are excluded from the residual universe. A set of additional admissible queries becomes an *exact repair* only when its correct labels eliminate all residual target disagreements, as formalised below.

Assumption 1 (Shared target-relative admissible-query semantics)**.** *Whenever client population laws are invoked for candidate world $G \in \mathcal{G}$, they preserve the same admissible-query semantics: $d_{P_{k,G}}(I) = d_G(I)$ for every $k \in [K]$ and $I \in \mathcal{Q}_T$.*

Assump. 1 fixes admissible-query labels in $\mathcal{Q}_T$ while allowing client variation in full-system DAGs, CI relations outside $\mathcal{Q}_T$, marginals, conditioning-cell probabilities, sample sizes, and observation allocations. The stochastic results below use strictly positive faithful finite-discrete parameterisations.

Evidence-Formation Interfaces. For each admissible query I, let $\mathcal{P}_I$ denote its finite collection of required conditioning-support cells, indexed by ℓ; let $N_{k,I,\ell}$ be the count contributed by client table $\mathcal{D}_k$ to cell ℓ, and let $m_{I,\ell}$ be the prespecified surfacing threshold for that cell. Define the local and matched pooled readiness indicators by

$$\mathrm{Ready}_I^{\mathrm{loc}} := \mathbf{1}[\exists k \in [K] : N_{k,I,\ell} \ge m_{I,\ell} \text{ for every required } \ell], \tag{1}$$

$$\mathrm{Ready}_I^{\mathrm{pool}} := \mathbf{1}\left[\sum_{k=1}^{K} N_{k,I,\ell} \ge m_{I,\ell} \text{ for every required } \ell\right]. \tag{2}$$

The primary interface is local-evidence-first and non-pooling: each required CI decision must be supportable within one client before it can surface. The matched pooled interface evaluates the same

support requirements after combining the corresponding counts across clients; under Assump. 2, a ready pooled query returns the same admissible-query label $d_G(I)$. Hence pooling changes where a whole CI decision is formed, while the target task and admissible-query label remain fixed. Thresholds $m_{I,\ell}$ govern support-based surfacing; Assump. 2 takes every surfaced admissible-query label to be correct, thereby isolating evidence-formation effects from finite-sample CI-test error. The bounded-aggregation strengthening below interpolates with query-specific aggregation depth b, whose endpoints 1 and K recover these two support tests.

Assumption 2 (Correct-or-abstain evidence)**.** *For either interface $\xi \in \{\mathrm{loc}, \mathrm{pool}\}$, candidate world $G \in \mathcal{G}$, and admissible query $I \in \mathcal{Q}_T$, the evidence response is $A_I^\xi(G) = d_G(I)$ when $\mathrm{Ready}_I^\xi = 1$ and $A_I^\xi(G) = \perp$ otherwise, where $\perp$ denotes abstention.*

Define the residual target disagreements and the disagreements separated by query I as

$$\mathrm{Amb}_T := \{\zeta = \{G, G'\} \subseteq \mathcal{G} : \mathcal{L}_T(G) \neq \mathcal{L}_T(G')\}, \tag{3}$$

$$\mathrm{Cov}_T(I) := \{\zeta = \{G, G'\} \in \mathrm{Amb}_T : d_G(I) \neq d_{G'}(I)\}. \tag{4}$$

Definition 1 (Exact Repairs and Support Geometry)**.** *Let $c_I > 0$ be the cost of query I. Let $\mathcal{R}^{\mathrm{sup}}$ be a finite set of primitive support requirements, each representing one atomic availability requirement of the evidence interface. A realised support state $z \subseteq \mathcal{R}^{\mathrm{sup}}$ is the set of primitive requirements that are available on that realised interface. Let $\mathfrak{P}_I \subseteq 2^{\mathcal{R}^{\mathrm{sup}}}$ be the antichain of minimal support paths sufficient to make I locally ready. The support-route representation abstracts the concrete readiness interface; local and pooled analyses use support states or support-state laws induced from the same observations. A query set $\mathcal{Q} \subseteq \mathcal{Q}_T$ is an exact repair when $\bigcup_{I\in\mathcal{Q}} \mathrm{Cov}_T(I) = \mathrm{Amb}_T$. The inclusion-minimal exact-repair family and repair cost, with $C(\mathcal{Q}) = |\mathcal{Q}|$ under unit query costs, are*

$$\mathfrak{C}_T := \min_{\subseteq} \left\{ \mathcal{Q} \subseteq \mathcal{Q}_T : \bigcup_{I\in\mathcal{Q}} \mathrm{Cov}_T(I) = \mathrm{Amb}_T \right\}, \qquad C(\mathcal{Q}) := \sum_{I\in\mathcal{Q}} c_I. \tag{5}$$

For a support state $z \subseteq \mathcal{R}^{\mathrm{sup}}$, query readiness is $\phi_I(z) := \mathbf{1}[\exists R_I \in \mathfrak{P}_I : R_I \subseteq z]$. Let μ denote a probability law on $2^{\mathcal{R}^{\mathrm{sup}}}$ and $Z \sim \mu$ the corresponding random realised support state. For an exact repair $\mathcal{Q}$, define its composed minimal support-path family by $\mathfrak{P}(\mathcal{Q}) := \min_{\subseteq}\{\bigcup_{I\in\mathcal{Q}} R_I : R_I \in \mathfrak{P}_I \text{ for every } I \in \mathcal{Q}\}$. Its readiness indicator is $\phi_{\mathcal{Q}}(z) := \mathbf{1}[\exists R \in \mathfrak{P}(\mathcal{Q}) : R \subseteq z]$, and its readiness probability under $Z \sim \mu$ is $\rho_\mu(\mathcal{Q}) := \Pr[\phi_{\mathcal{Q}}(Z) = 1]$. For $q \in (0, 1]$, selection minimises $C(\mathcal{Q})$ over repairs with $\rho_\mu(\mathcal{Q}) \geq q$, abstaining if none qualify; without fixed q, we retain the full cost–readiness ordering. A realised client-data/support specification Γ and interface $\xi \in \{\mathrm{loc}, \mathrm{pool}\}$ determine a support state; if Γ is stochastic under $G \in \mathcal{G}$, they induce the law $\mu^\xi_{\Gamma,G}$ and readiness $\rho^\xi_{\Gamma,G}(\mathcal{Q}) := \rho_{\mu^\xi_{\Gamma,G}}(\mathcal{Q})$. When Γ is clear and $G^\star$ is fixed, write $(\mu^{\xi,\star}, \rho^{\xi,\star}) := (\mu^\xi_{\Gamma,G^\star}, \rho^\xi_{\Gamma,G^\star})$.

We compare the complete exact-repair family $\mathfrak{C}_T$ through structural and interface-dependent costs. The minimum structural repair cost is $C_T^{\mathrm{str}} := \min_{\mathcal{Q}\in\mathfrak{C}_T} C(\mathcal{Q})$. A singleton query is *target-sufficient* on the current candidate class when candidates sharing its definite label also share the same target summary, equivalently when its coverage is all of Amb_T. We write $I_0 \downarrow$ for the event that a supplied baseline query I_0 abstains; when it appears as an argument of a cost, it denotes the resulting residual task with I_0 removed from the admissible repair pool. Thus $C_T^{\mathrm{str}}(I_0)$ is the structural optimum for that residual task. After such a baseline abstention, a local adaptive protocol may select additional admissible CI queries sequentially using previous correct labels and abstentions, then stop and output a target summary. It is *exact* if it outputs $\mathcal{L}_T(G)$ for every $G \in \mathcal{G}$ on the realised interface. The quantity $C^{\mathrm{loc,ad}}_{T,\Gamma}(I_0)$ is the minimum, over exact local protocols, of their worst-case cumulative distinct-query cost; $C^{\mathrm{pool,ad}}_{T,\Gamma}(I_0)$ is defined analogously under the pooled control. We omit the argument I_0 when no supplied baseline is being conditioned on. Thm. 1 uses unit query costs, so these quantities equal the corresponding numbers of distinct calls.

4 FEDERATED TRANSFER OF EXACT REPAIRS

4.1 FROM CAUSAL DISAGREEMENT TO ADAPTIVE REPAIR COMPLEXITY

Proposition 1 (Lossless Repair-Support Representation). *For each residual disagreement $\zeta \in \mathrm{Amb}_T$, let $\mathcal{E}_\zeta := \{I \in \mathcal{Q}_T : \zeta \in \mathrm{Cov}_T(I)\}$. The residual-disagreement hypergraph has vertex set $\mathcal{Q}_T$ and hyperedges $\{\mathcal{E}_\zeta : \zeta \in \mathrm{Amb}_T\}$, and $\mathfrak{C}_T$ is exactly its family of minimal transversals. For every exact repair $\mathcal{Q}$, the conjunction of its query-readiness events has exactly the minimal support-path family $\mathfrak{P}(\mathcal{Q})$. Consequently, the query coverages, query costs, minimal query support paths, and support-state law completely determine exact-repair identities, costs, readiness probabilities, and cost–readiness orderings. Any isomorphism preserving these objects preserves these quantities.*

We next characterise the adaptive query cost of resolving this representation on a realised interface. For a realised support state z, let $\mathcal{Q}_z := \{I \in \mathcal{Q}_T : \phi_I(z) = 1\}$ be the queries that surface definite labels. Fix $G_0 \in \mathcal{G}$. A candidate G is *target-wrong relative to* G_0 when $\mathcal{L}_T(G) \neq \mathcal{L}_T(G_0)$. For each such G, define the visible witness set $\mathcal{W}_z(G_0, G) := \{I \in \mathcal{Q}_z : d_{G_0}(I) \neq d_G(I)\}$. Their collection $\mathsf{W}_{T,z}(G_0) := \{\mathcal{W}_z(G_0, G) : \mathcal{L}_T(G) \neq \mathcal{L}_T(G_0)\}$ is the reference-relative witness hypergraph for (T, z, G_0). For the fractional relaxation, $x_I \geq 0$ denotes the non-negative selection weight assigned to visible query I. Its weighted integral and fractional transversal values are

$$\tau_{T,z}(G_0) := \min_{S \subseteq \mathcal{Q}_z} \left\{ \sum_{I \in S} c_I : S \cap \mathcal{W}_z(G_0, G) \neq \varnothing \ \forall G \text{ target-wrong relative to } G_0 \right\}, \tag{6}$$

$$\tau^\star_{T,z}(G_0) := \min_{x_I \geq 0} \left\{ \sum_{I \in \mathcal{Q}_z} c_I x_I : \sum_{I \in \mathcal{W}_z(G_0,G)} x_I \geq 1 \ \forall G \text{ target-wrong relative to } G_0 \right\}. \tag{7}$$

If any witness edge is empty, $\tau_{T,z}(G_0) := +\infty$.

Proposition 2 (Reference-Relative Adaptive Witness Characterisation). *Assume that the visible query labels determine the target summary uniformly over $\mathcal{G}$ on state z. Let $D^{\mathrm{det}}_{T,z}(G_0)$ be the minimum cumulative distinct-query cost incurred on the G_0 transcript among deterministic adaptive protocols that are exact for all $G \in \mathcal{G}$. Fix $0 \leq \delta < 1/2$ and any randomised protocol Π with error probability at most δ for every $G \in \mathcal{G}$ on this realised state. Let C_Π be its cumulative distinct-query cost; let $\mathbb{E}_{G_0}$ denote expectation over the protocol's internal randomness on the G_0 transcript. Then*

$$D^{\mathrm{det}}_{T,z}(G_0) = \tau_{T,z}(G_0), \qquad \mathbb{E}_{G_0}[C_\Pi] \geq (1 - 2\delta)\tau^\star_{T,z}(G_0). \tag{8}$$

The characterisation is reference-relative: it controls the cost of the G_0 transcript after the complete residual task has been fixed. The uniform exact-repair family $\mathfrak{C}_T$ remains the object used for repair comparison.

4.2 HOW CLIENT-LOCAL EVIDENCE FORMATION INFLATES REPAIR COST

Using Props. 1 and 2, proved in Apps. B and C, we now construct a hard family in which client-local formation distorts adaptive difficulty arbitrarily. We classify its complete observational-CI disagreement structure while deriving all query-readiness events from the same client tables. This classification identifies all exact repairs and establishes the no-shortcut property used in Thm. 1.

Theorem 1 (Unbounded Client-Local Repair-Cost Distortion). *For every residual ambiguity-module count $h \geq 2$, client count $K \geq 2$, and integer readiness threshold $m \geq 2$, there exists a connected target-specific causal recovery task on $2h + 5$ observed variables, a designated reference graph $G^\star_h$, and an ordinary horizontal federated system $\Gamma_{h,K,m}$ with a local non-pooling interface with exactly K clients such that, uniformly over every supplied singleton target-sufficient baseline I_0 after $I_0 \downarrow$,*

$$C^{\mathrm{str}}_T(I_0) = 1, \qquad C^{\mathrm{loc,ad}}_{T,\Gamma_{h,K,m}}(I_0) = h, \qquad C^{\mathrm{pool,ad}}_{T,\Gamma_{h,K,m}}(I_0) = 1. \tag{9}$$

Every exact deterministic adaptive local protocol therefore makes at least h distinct additional calls on the reference transcript, and the canonical module protocol attains this bound. For every $0 \leq \delta < 1/2$, every local randomised protocol that is δ-correct for every $G \in \mathcal{G}$ on this realised interface satisfies $\mathbb{E}_{G^\star_h}[C_\Pi] \geq (1 - 2\delta)h$. Consequently, for every fixed $K \geq 2$ the local-to-pooled ratio is unbounded as $h \to \infty$.

Remark 1 (What the $1{:}h{:}1$ separation means in practice)**.** *Imagine that a medical study has h residual ambiguity modules affecting its target conclusion. With all observations together, one CI query may resolve all h modules; across clients, the conditioning groups needed for that query may no longer occur locally, while each module retains a simpler witness. The study can therefore require one query structurally, h locally, and one after the relevant sufficient statistics are pooled. Here h counts ambiguity modules rather than clients: the separation holds for every $K \geq 2$, and the adaptive lower bound applies at protocol level across the complete repair family.*

Proof idea. Every global witness is locally blocked, each module retains a visible witness, and pooling restores a global witness; Prop. 2 then gives the $1{:}h{:}1$ separation and the randomised adaptive lower bound. App. D gives the full construction and proof.

Bounded aggregation. For aggregation depth $b \in [K]$, let $N_{k,I,u}$ be client k's count in required cell $u \in \mathcal{P}_I$, and define

$$\mathrm{Ready}_I^{(b)} := \mathbf{1}\left[\exists C_I \subseteq [K],\ |C_I| \leq b : \sum_{k \in C_I} N_{k,I,u} \geq m_{I,u}\ \forall u \in \mathcal{P}_I\right]. \tag{10}$$

Use the correct-or-abstain semantics of Assump. 2 with $\mathrm{Ready}_I^{(b)}$, and let $C_{T,\Gamma}^{(b),\mathrm{ad}}(I_0)$ be the corresponding adaptive cost. Coalitions are query-specific; $b = 1$ and $b = K$ recover the local and matched pooled interfaces.

Corollary 1 (Persistence under bounded aggregate-before-decision)**.** *For every $h \geq 2$, $m \geq 2$, and $K \geq b_0 \geq 2$, the construction of Thm. 1 admits a legal K-client federation with*

$$C_T^{\mathrm{str}}(I_0) = 1, \qquad C_{T,\Gamma}^{(b),\mathrm{ad}}(I_0) = \begin{cases} h, & 1 \leq b < b_0, \\ 1, & b_0 \leq b \leq K. \end{cases} \tag{11}$$

For $b < b_0$, every δ-correct randomised protocol, $0 \leq \delta < 1/2$, also has expected cost at least $(1 - 2\delta)h$. Thus b_0 and h are independently selectable.

Remark 2 (Worst-case bounded-depth persistence)**.** *A short repair may require support pieces spread across several hospitals. The threshold b_0 is the coalition depth at which it becomes formable and can be chosen independently of the cost gap h. The construction keeps the DAG topology fixed while increasing the state cardinalities of P and R with b_0. App. D gives the bounded-depth construction, while under random support fragmentation the recovery depth obeys the quantitative law below.*

Theorem 2 (Stochastic Bounded-Depth Formation and Recovery)**.** *Fix a support-competitive canonical route I with rare required cells $\mathcal{C}_j$ satisfying $P(\mathcal{C}_j) = \kappa_j p_N\{1 + o(1)\}$, $\kappa_j > 0$, fixed thresholds m_j, rare-support burden $w_I := \sum_j m_j$, and remaining required cells uniformly bounded away from zero. Let $K_N = N^{\beta + o(1)}$ independent near-balanced clients sample from this population with total size N, let $p_N = N^{-\gamma}$, and assume $0 < \beta, \gamma < 1$ and $\beta + \gamma > 1$. Define*

$$e_\ell(\beta, \gamma) := \ell(1 - \gamma) - (\ell - 1)\beta, \qquad \ell_\star := \max\{\ell \geq 1 : e_\ell(\beta, \gamma) > 0\}. \tag{12}$$

Away from critical walls, i.e. $e_\ell(\beta, \gamma) \neq 0$ for $1 \leq \ell \leq w_I$, the formation depth $d_{\Gamma_N}(I) := \min\{b \in [K_N] : \mathrm{Ready}_I^{(b)} = 1\}$ satisfies

$$d_{\Gamma_N}(I) \xrightarrow{p} \left\lceil \frac{w_I}{\ell_\star} \right\rceil. \tag{13}$$

The conclusion holds jointly for every fixed finite family of canonical support-competitive routes and allows arbitrary dependence among those routes.

Remark 3 (Interpreting stochastic formation depth)**.** *A route may require several rare patient subgroups. Here w_I measures its rare-support demand and $\ell_\star$ the largest number of rare threshold units abundant within one client; their ratio determines the required coalition depth.*

Proof idea. At depth b, readiness requires some client to supply at least $\lceil w_I/b \rceil$ rare units. Federation-wide multiplicity of collision order ℓ scales as $N^{e_\ell(\beta,\gamma)+o(1)}$; its exponent crossing and monotonicity in b give Eq. (13). App. E gives the full phase analysis, including critical Poisson-packing.

Corollary 2 (Client-Local Critical Repair-Cost Consequence)**.** *For the sparse causal family underlying Thm. 2, let $h_N \to \infty$ and place the support-best global route on the client-local $b = 1$*

critical wall with limiting exposure $\lambda_G \in [0,\infty]$ (with $e^{-\infty}=0$), while the canonical module route is supercritical and competing global routes are subcritical. Conditional on baseline abstention, $\Pr_N^\star(\mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}) \to 1$ and

$$\Pr_N^\star\big(\mathsf{K}_N^{\mathrm{loc}} = 1 \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}\big) \to 1 - e^{-\lambda_G}, \qquad \Pr_N^\star\big(\mathsf{K}_N^{\mathrm{loc}} = h_N \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}\big) \to e^{-\lambda_G}. \tag{14}$$

Consequently, $\mathbb{E}_N^\star[\mathsf{K}_N^{\mathrm{loc}} \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}]/h_N \to e^{-\lambda_G}$ and $C_N^{\mathrm{loc,det}}/h_N \to e^{-\lambda_G}$; every randomised protocol that is δ-correct on every repairable realised support state satisfies $\liminf_N \mathbb{E}_N^\star[C_\Pi]/h_N \geq (1-2\delta)e^{-\lambda_G}$.

4.3 When Client-Local Evidence Formation Reverses Repair Ordering

Cor. 2, proved via the route-profile transfer in App. F, shows how the client-local $b=1$ global-route critical wall changes repair cost; we now show that the same wall can reverse the ordering of competing exact repairs. On the same residual family after baseline abstention, let $\mathcal{A}_h^G$ be a one-query global repair and $\mathcal{B}_h^M$ the canonical h-query module repair. For client laws $P_{k,N}$, write $\bar{P}_N := \sum_{k=1}^{K_N}(n_{k,N}/N)P_{k,N}$ for the sample-weighted pooled population.

Theorem 3 (Client-Local Repair-Order Reversal)**.** *For every $h \geq 2$, there exist an integer critical global-route rare-support burden $w_G > 1$, a rarity exponent $\gamma \in (0,1)$, and a strictly positive faithful finite-discrete causal population sequence on the connected residual family with rare-cell probability scale $p_N = N^{-\gamma}$. For total sample size N and client count K_N, define $\beta_G := w_G(1-\gamma)/(w_G-1)$, equivalently $e_{w_G}(\beta_G,\gamma)=0$, and the offset $t_N := \log K_N - \beta_G \log N$. Then the following hold.*

(a) Matched-population critical window. *Let all near-balanced clients sample independently from the same population, where near-balanced means that the K_N client sizes sum to N and differ by at most one, and suppose $t_N \to t \in \mathbb{R}$. Conditional on baseline abstention,*

$$\begin{aligned} &\rho^{\mathrm{loc},\star}(\mathcal{A}_h^G) \longrightarrow r_G(t) := 1 - \exp\{-\theta_G e^{-(w_G-1)t}\}, \\ &\rho^{\mathrm{loc},\star}(\mathcal{B}_h^M) \longrightarrow 1, \qquad \rho^{\mathrm{pool},\star}(\mathcal{A}_h^G), \rho^{\mathrm{pool},\star}(\mathcal{B}_h^M) \longrightarrow 1. \end{aligned} \tag{15}$$

Here $\theta_G > 0$ is the limiting exposure coefficient of the support-best global route. Every competing global singleton or refined global singleton is asymptotically unavailable or no easier. Thus the complete local repair ordering exhibits the non-degenerate cost–readiness trade-off $(r_G(t), -1)$ versus $(1, -h)$, whereas pooling retains a one-query repair. The function r_G is continuous, strictly decreasing, and onto $(0,1)$.

(b) Allocation-only same-sample separation. *With probability $1-o(1)$ over the pooled sample, the realised sample admits two admissible observation-to-client allocations that preserve every pooled count and graph-implied CI label but expose different locally minimum-cost exact repairs: one allocation exposes a global singleton, whereas the other exposes only the module repair.*

(c) Threshold-free same-pooled-population separation. *There are two client-population sequences with the same sample-weighted pooled population $\bar{P}_N$, total size N, client count K_N, local sample sizes, target task, and complete repair class, but different client-population decompositions whose effective global-route exposures satisfy $0 < \Theta_G^{\mathrm{hom}} < \Theta_G^{\mathrm{split}}$. The decomposition changes client marginals and conditioning-cell masses, hence route exposure, while every client retains the same admissible-query CI labels required by Assump. 1. At every finite common critical offset t,*

$$r_G^{\mathrm{hom}}(t) := 1 - \exp\{-\Theta_G^{\mathrm{hom}} e^{-(w_G-1)t}\} < r_G^{\mathrm{split}}(t) := 1 - \exp\{-\Theta_G^{\mathrm{split}} e^{-(w_G-1)t}\}. \tag{16}$$

The module repair tends to readiness one in both decompositions, while pooled evidence sees the same pooled population $\bar{P}_N$ and a one-query global repair in both. Client decomposition therefore changes local repair readiness despite the same pooled population.

Remark 4 (How the preferred repair can change)**.** *The ordering can reverse while the cheaper repair still has positive readiness. For example, the same patients can be assigned to hospitals in two different ways: one may place the conditioning groups for a one-query repair together within a client, while another separates them across clients, making a larger repair more ready locally. Part (b) fixes realised pooled sample, while part (c) fixes the pooled population, total and client sample sizes, client count, target task, and complete repair class, changing only the client-population decomposition. Thus aggregate information alone does not determine which repair is preferable under client-local*

evidence formation. Because the limiting readiness curves are already strictly ordered before a readiness threshold q is chosen, every threshold between them yields the same reversal.

Proof idea. The repair classification leaves only global singleton repairs and repairs touching all h modules. At the critical scale, the linked rare-state construction makes the module repair ready while the global route has non-degenerate readiness and competing global shortcuts vanish. Part (b) follows by reallocating a fixed pooled sample; part (c) uses a pooled-law-preserving client split that changes global-route exposure. App. G gives the full proof and phase calculations.

Target-decision ordering. Thm. 3 separates structural and client-local readiness ordering; we next ask whether readiness ordering itself determines value for the fixed target. Let $Y_{\mathcal{Q}}$ be the coordinatewise correct-or-abstain transcript of repair $\mathcal{Q}$, retaining surfaced coordinates even when the repair is not fully ready, and define $\mathcal{R}_\pi(\mathcal{Q}; L) := \inf_\psi \mathbb{E}_\pi[L(\mathcal{L}_T(G), \psi(Y_{\mathcal{Q}}))]$ for prior π and bounded target loss L. Write $\mathcal{A} \succeq_T \mathcal{B}$ when this risk is no larger for $\mathcal{A}$ for every prior, finite action space, and bounded loss depending only on $\mathcal{L}_T(G)$; $\succ_T$ and $\|_T$ denote strict dominance and incomparability. We call $\succeq_T$ the *target-decision preorder*. We write $\mathsf{E}^T_{\mathcal{Q}}$ for the target experiment induced by $Y_{\mathcal{Q}}$, and use the same preorder $\succeq_T$ for these induced experiments. In the route-grid construction below, $\mathcal{A}^{\mathrm{row}}_h$ is the canonical two-query row repair and $\mathcal{B}^{\mathrm{col}}_h$ the canonical h-query column repair; r is the limiting readiness of each row route and c that of the shared column route, and $\mathsf{E}^T_{\mathrm{row}}, \mathsf{E}^T_{\mathrm{col}}$ denote their limiting target experiments.

Theorem 4 (Limiting target-decision trichotomy)**.** *Fix $h \geq 3$, $m \geq 2$, and $r, c \in (0, 1)$. On the route-grid family of App. J, there exists a faithful finite-discrete sequence with a common upstream marginal inducing a finite-N graph-independent readiness-mask law. The row and column repairs have structural costs 2 and h, and their limiting target experiments satisfy*

$$\mathsf{E}^T_{\mathrm{row}} \succeq_T \mathsf{E}^T_{\mathrm{col}} \iff c \leq r^2, \qquad \mathsf{E}^T_{\mathrm{col}} \succeq_T \mathsf{E}^T_{\mathrm{row}} \iff c \geq r. \tag{17}$$

Thus $r^2 < c < r$ gives incomparability; at the matched point $c = r$, the more expensive repair dominates, and since the construction holds for arbitrary h, its structural cost ratio $h/2$ can be made arbitrarily large.

Remark 5 (When Readiness and Target Value Diverge)**.** *The longer hospital repair can become easier to form before becoming uniformly preferable for the target. In the intermediate region, the preferred repair can depend on the target prior or decision loss.*

Proof idea. The causal lift gives every candidate world the same support channel. After nuisance marginalisation, the column repair is a binary erasure experiment with reveal probability c, while the row experiment has target-state total variation r and worst-case target-exclusive reveal mass r^2. Binary garbling criteria yield Eq. (17); App. J.7 gives the full lift and calculations.

4.4 When Repair Ordering Is Preserved

A complementary preservation regime arises under support-compatible dominance. The condition below gives a causal and finite-sample criterion under which a structural preference between exact repairs survives the declared evidence interfaces. Using the conditioning-support cell collections $\mathcal{P}_I$ defined in Sec. 3, write $J \succeq_{\mathrm{cell}} I$ when every J-cell is a union of I-cells and $m_{J,u} \leq \sum_{v \in \mathcal{P}_I : v \subseteq u} m_{I,v}$ for every $u \in \mathcal{P}_J$. By construction, $J \succeq_{\mathrm{cell}} I$ guarantees that readiness of I implies readiness of J on every finite client table, at every bounded aggregation depth b, and after full pooling. Under common per-cell thresholds m_I, m_J, strict conditioning-partition coarsening with $m_J \leq m_I$ is a sufficient special case. Here an *interface-induced* support state is one arising from a finite realised client table under the local, bounded-depth, or matched pooled interface; an interface-induced support-state law is generated by a stochastic version of one of these interfaces.

Corollary 3 (Support-Compatible Coverage and Repair-Order Preservation)**.** *Let $\mathcal{A}, \mathcal{B} \in \mathfrak{C}_T$. Suppose there is a map $\varphi : \mathcal{B} \to \mathcal{A}$ such that $\mathrm{Cov}_T(\varphi(I)) \supseteq \mathrm{Cov}_T(I)$ and $\varphi(I) \succeq_{\mathrm{cell}} I$ for every $I \in \mathcal{B}$. Then $\phi_{\mathcal{A}}(z) \geq \phi_{\mathcal{B}}(z)$ for every interface-induced support state z, and $\rho_\mu(\mathcal{A}) \geq \rho_\mu(\mathcal{B})$ for every interface-induced support-state law μ. If also $C(\mathcal{A}) \leq C(\mathcal{B})$, no readiness threshold can make minimum-cost selection uniquely prefer $\mathcal{B}$ to $\mathcal{A}$.*

Remark 6 (A practical preservation regime)**.** *Suppose every query in repair $\mathcal{B}$ can be matched to one in $\mathcal{A}$ that rules out at least the same target-wrong alternatives with no finer conditioning groups or larger support thresholds. For example, a query in $\mathcal{B}$ may require counts for every age-by-severity*

subgroup whereas its matched query in $\mathcal{A}$ only requires age groups. Any client able to support the former can then support the latter, so every interface-induced support state making $\mathcal{B}$ ready also makes $\mathcal{A}$ ready. If $\mathcal{A}$ is also no more costly, it remains weakly preferred under client-local support.

Together, Thm. 3 and Cor. 3 delimit repair-order reversal and preservation. Alternative low-cost repairs can protect the preferred frontier; Sec. 5 tests this behaviour, while App. H gives the proof, the necessary-and-sufficient abstract path-level boundary, and an ALARM preservation instance.

5 Experiments

Experiments test three theory-facing effects: aggregation-depth recovery, complete-frontier reordering, and allocation sensitivity with the realised pooled sample fixed. For structural cost c, the cost-c readiness frontier maximises over members of the complete repair family with cost c, and the complete lower-cost frontier maximises over all lower-cost members of the complete repair family. The frontier-recovery depth is the first b at which the structural-minimum frontier reaches the best complete-repair readiness and remains there over larger depths. Public-network CI semantics come from the benchmark DAG/CPT populations; full protocols and broader controls are in App. K.

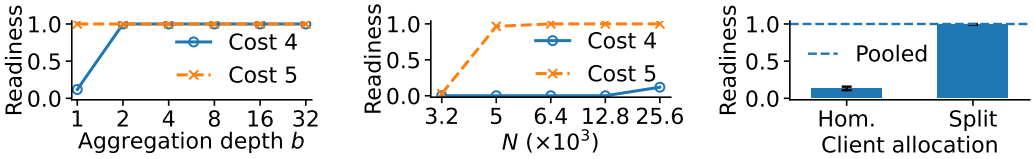


Figure 1: Three public-network results under seeds 42–51. (a) CHILD/ChestXray: the complete cost-4 frontier rises from 0.119 at $b = 1$ to 1.000 at $b = 2$, while the cost-5 frontier is 1.000 throughout. (b) The cost-5 readiness frontier exceeds the complete lower-cost frontier at all five pre-specified sample sizes. (c) CHILD/GruntingReport: reallocating the same pooled sample changes local readiness from 0.133 to 0.994 while pooled readiness remains 1.

On CHILD/ChestXray, this two-client recovery persists through $b = 32$; deeper aggregation provides no further gain on this witness. This localises the observed readiness reversal to shallow cross-client support fragmentation rather than a need for unrestricted pooling. Three additional CHILD targets recover at depth one or two (App. K.4), showing shallow recovery on these tasks alongside the theorem's arbitrary worst-case recovery depth. The same ChestXray task exhibits complete-family reversal: at $K = 32, m = 5, N = 25{,}600$, the cost-5 readiness frontier exceeds the complete lower-cost frontier by 0.881, with a seed-level 95% Monte Carlo precision interval $[0.8672, 0.8948]$; the broader panel contains further reversals, while ALARM shows how alternative low-cost repairs can protect the complete frontier (App. K.5). Finally, on CHILD/GruntingReport at $K = 32, N = 51{,}200$, the realised pooled sample and client sizes are fixed and only allocation changes, producing a paired local-readiness gap of 0.861 (paired seed-level 95% Monte Carlo precision interval $[0.8346, 0.8874]$) while pooled readiness remains 1 (App. K.7.1). Fixed-total/fixed-local controls and the INSURANCE pre-transfer boundary are in Apps. K.6.1–K.7.2.

6 Conclusion

Federated evidence formation can change the structural ordering of alternative exact repairs for the same causal target. We proved an unbounded $1{:}h{:}1$ repair-cost separation that can persist through any prescribed bounded aggregation depth; under random federation, we characterise the aggregation depth required for evidence formation. Client-local formation can reverse repair readiness with pooled information fixed, while readiness can disagree with target-decision order. Support-compatible dominance preserves structural preference, while exact aggregate-before-decision recovers the pooled boundary. Public Bayesian networks exhibit shallow aggregation recovery, complete-family reversal, same-pooled decomposition sensitivity, and pre-transfer infeasibility. Together, the results separate whether an exact repair can be formed under the federation from whether its induced target experiment remains preferable.

AI USE STATEMENT

AI was used to assist with research refinement, theoretical formulation and proof checking, experimental design, analysis and interpretation of experimental results, literature organisation, and manuscript polishing. All AI-assisted mathematical claims, proofs, experimental protocols, reported results, and manuscript content were reviewed and validated by the authors. The authors take full responsibility for the correctness, originality, and integrity of the work.

REPRODUCIBILITY STATEMENT

Complete proofs are provided in the appendix. The supplementary experimental material reports the certified repair panels, full-network probability construction, readiness contract, controlled construction-validity analysis, client-allocation designs, paired controls, and statistics supporting the empirical claims. The accompanying anonymous artifact contains the code, configurations, and final outputs used to reproduce the reported numerical results.

REFERENCES

Jayadev Acharya, Clément L. Canonne, and Himanshu Tyagi. Communication-constrained inference and the role of shared randomness. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 30–39. PMLR, 2019.

Jayadev Acharya, Clément L. Canonne, and Himanshu Tyagi. Distributed signal detection under communication constraints. In *Proceedings of the Thirty-Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 41–63. PMLR, 2020.

Jayadev Acharya, Clément L. Canonne, Ziteng Sun, and Himanshu Tyagi. The role of interactivity in structured estimation. In *Proceedings of the Thirty-Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 1328–1355. PMLR, 2022.

Constantin F. Aliferis, Alexander Statnikov, Ioannis Tsamardinos, Subramani Mani, and Xenofon D. Koutsoukos. Local causal and Markov blanket induction for causal discovery and feature selection for classification Part I: Algorithms and empirical evaluation. *Journal of Machine Learning Research*, 11:171–234, 2010.

Richard E. Barlow and Frank Proschan. *Statistical Theory of Reliability and Life Testing: Probability Models*. Holt, Rinehart and Winston, New York, 1975.

Z. W. Birnbaum. On the importance of different components in a multicomponent system. In P. R. Krishnaiah (ed.), *Multivariate Analysis II*, pp. 581–592. Academic Press, New York, 1969.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24 (2):265–272, 1953. doi: 10.1214/aoms/1177729032.

Wei Chen, Wanyang Gu, Linjun Peng, Ting Yan, Ruichu Cai, Zhifeng Hao, and Kun Zhang. Horizontal and vertical federated causal structure learning via higher-order cumulants. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(24):20280–20288, 2026.

Erdong Gao, Junjia Chen, Li Shen, Tongliang Liu, Mingming Gong, and Howard D. Bondell. FedDAG: Federated DAG structure learning. *Transactions on Machine Learning Research*, 2023.

Tian Gao and Qiang Ji. Local causal discovery of direct causes and effects. In *Advances in Neural Information Processing Systems*, volume 28, pp. 2512–2520, 2015.

Xianjie Guo, Kui Yu, Lin Liu, and Jiuyong Li. FedCSL: A scalable and accurate approach to federated causal structure learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11): 12235–12243, 2024a.

Xianjie Guo, Kui Yu, Hao Wang, Lizhen Cui, Han Yu, and Xiaoxiao Li. Sample quality heterogeneity-aware federated causal discovery through adaptive variable space selection. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pp. 4071–4079, 2024b.

Xianjie Guo, Liping Yi, Xiaohu Wu, Kui Yu, and Gang Wang. Enhancing causal discovery in federated settings with limited local samples. In *Federated Learning in the Age of Foundation Models – FL 2024 International Workshops*, volume 15501 of *Lecture Notes in Computer Science*, pp. 164–179. Springer, 2025.

Maximilian Hahn, Alina Zajak, Dominik Heider, and Adèle Helena Ribeiro. Federated causal discovery across heterogeneous datasets under latent confounding. *arXiv preprint arXiv:2603.05149*, 2026.

Jianli Huang, Xianjie Guo, Kui Yu, Fuyuan Cao, and Jiye Liang. Towards privacy-aware causal structure learning in federated setting. *IEEE Transactions on Big Data*, 9(6):1525–1535, 2023.

Fateme Jamshidi, Luca Ganassali, and Negar Kiyavash. On the sample complexity of conditional independence testing with Von Mises estimator with application to causal discovery. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 21225–21240. PMLR, 2024.

Markus Kalisch and Peter Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–636, 2007.

Loka Li, Ignavier Ng, Gongxu Luo, Biwei Huang, Guangyi Chen, Tongliang Liu, Bin Gu, and Kun Zhang. Federated causal discovery from heterogeneous data. In *International Conference on Learning Representations*, 2024.

Timothée Loranchet and Charles K. Assaad. Local Markov equivalence for PC-style local causal discovery and identification of controlled direct effects. In *Proceedings of the Fifth Conference on Causal Learning and Reasoning*, volume 323 of *Proceedings of Machine Learning Research*, pp. 84–118. PMLR, 2026.

Jean-Luc Marichal. Structure functions and minimal path sets. *arXiv preprint arXiv:1401.0803*, 2014.

Arnaud Mary. Enumeration of minimal transversals of hypergraphs of bounded VC-dimension. *arXiv preprint arXiv:2407.00694*, 2024.

Ignavier Ng and Kun Zhang. Towards federated Bayesian network structure learning with continuous optimization. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 8095–8111. PMLR, 2022.

Jiaqi Shi, Xiaoling Huang, Xianjie Guo, Kui Yu, Chengxiang Hu, and Peng Zhou. Federated causal structure learning with missing data. *Knowledge-Based Systems*, 330:114601, 2025.

Kirankumar Shiragur, Jiaqi Zhang, and Caroline Uhler. Causal discovery with fewer conditional independence tests. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 45060–45078. PMLR, 2024.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2 edition, 2000.

Eichi Uehara. Iterative causal discovery: Per-edge impossibility certificates, tier-aware oracle queries, and the $1 + k$ lower bound. *arXiv preprint arXiv:2605.27477*, 2026.

Yunxia Wang, Fuyuan Cao, Kui Yu, and Jiye Liang. Federated causal structure learning with non-identical variable sets. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 62445–62466. PMLR, 2025.

Zhaoyu Wang, Pingchuan Ma, and Shuai Wang. Towards practical federated causal structure learning. In *Machine Learning and Knowledge Discovery in Databases: Research Track*, volume 14170 of *Lecture Notes in Computer Science*, pp. 351–367. Springer, 2023.

Jiaqi Zhang, Kirankumar Shiragur, and Caroline Uhler. Membership testing in Markov equivalence classes via independence queries. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3925–3933. PMLR, 2024.

Appendix

A NOTATION

Symbol	Meaning
$T, \mathcal{V}_T$	Target variable and finite target-specific analysis scope.
$\mathcal{G}, \mathcal{L}_T(G)$	Residual admissible DAG class and target summary induced by candidate graph G.
$I = (X_a, X_b, S), \mathcal{Q}_T$	An admissible CI query and the fixed residual admissible query universe; supplied baseline queries are excluded from the residual universe.
$d_G(I), d_{P_{k,G}}(I)$	Graph-implied and client-population CI labels for query I.
$\zeta, \mathrm{Amb}_T, \mathrm{Cov}_T(I), \mathcal{E}_\zeta$	A residual target-disagreement pair, the set of all such pairs, the pairs separated by query I, and the queries that separate a fixed pair ζ.
$c_I, C(\mathcal{Q}), \mathfrak{C}_T$	Query cost, total repair cost, and the family of inclusion-minimal exact repairs.
$\mathcal{P}_I, N_{k,I,\ell}, m_{I,\ell}$	Required conditioning-support cell collection for query I, client-k cell count, and prespecified cell threshold. The calligraphic $\mathcal{P}_I$ is distinct from the support-path family $\mathfrak{P}_I$.
$\mathrm{Ready}_I^{\mathrm{loc}}, \mathrm{Ready}_I^{\mathrm{pool}}$	Client-local and matched pooled readiness indicators for query I.
$b, \mathrm{Ready}_I^{(b)}, A_I^{(b)}, d_\Gamma$	Bounded aggregation depth, bounded-depth readiness and response, and query/repair formation depths $d_\Gamma(I)$ and $d_\Gamma(\mathcal{Q})$.
$p_N, K_N, \beta, \gamma, w_I, e_\ell, \ell_\star$	Rare-cell scale, client count, client-count and rarity exponents, rare-support burden of canonical route I, collision exponent $e_\ell(\beta, \gamma)$, and largest supercritical collision order in Thm. 2.
$\mathcal{R}^{\mathrm{sup}}, z$	Primitive support requirements and a realised support state $z \subseteq \mathcal{R}^{\mathrm{sup}}$.
$\mathfrak{P}_I, \phi_I(z)$	Minimal support paths for query I and its readiness indicator on state z.
$\mathfrak{P}(\mathcal{Q}), \phi_{\mathcal{Q}}(z)$	Composed minimal support paths of exact repair $\mathcal{Q}$ and its readiness indicator.
$\mu, \rho_\mu(\mathcal{Q})$	Support-state law and repair readiness probability.
Γ, ξ	Client-data/support specification and evidence interface $\xi \in \{\mathrm{loc}, \mathrm{pool}\}$. A realised Γ and ξ determine a support state; a stochastic Γ induces $\mu_{\Gamma,G}^{\xi}$ and $\rho_{\Gamma,G}^{\xi}$.
$I_0 \downarrow$	Abstention of supplied baseline query I_0; when used in a cost, it denotes the residual task with I_0 removed from the admissible repair pool.
C_T^{str}	Minimum structural repair cost over $\mathfrak{C}_T$ for the current residual task.
$C_{T,\Gamma}^{\mathrm{loc,ad}}, C_{T,\Gamma}^{(b),\mathrm{ad}}, C_{T,\Gamma}^{\mathrm{pool,ad}}$	Minimum worst-case cumulative distinct-query cost of an exact adaptive protocol under client-local, bounded-depth, and matched pooled evidence formation.
$C_{T,\Gamma}^{\mathrm{pool}}$	Minimum cost of a pooled-ready inclusion-minimal exact repair; this static quantity is distinct from $C_{T,\Gamma}^{\mathrm{pool,ad}}$.
$\mathcal{Q}_z$	Admissible queries that surface definite labels on realised support state z.
$\mathcal{W}_z(G_0, G), \mathsf{W}_{T,z}(G_0)$	Visible witness set separating reference graph G_0 from target-wrong candidate G, and the resulting reference-relative witness hypergraph.
$\tau_{T,z}(G_0), \tau_{T,z}^{\star}(G_0)$	Weighted integral and fractional transversal values of the reference-relative witness hypergraph.

Symbol	Meaning
$D^{\mathrm{det}}_{T,z}(G_0), C_\Pi, \delta$	Minimum deterministic reference-transcript cost, cumulative distinct-query cost of protocol Π, and its worst-case error tolerance.
h, K, N, m, q	Residual ambiguity-module count, client count, total sample size, integer readiness threshold used in Thm. 1, and readiness threshold used in thresholded comparisons.
$Q^{\mathrm{call}}_\Pi, Q^{\mathrm{mod}}_\Pi$	Total distinct additional query calls and calls to module-witness classes made by protocol Π in Thm. 1.
G_N, M_N, P_N	Readiness events for the support-best global singleton, common module route, and competing global route in Cor. 2 and App. F.
$\Lambda_{G,N}, \lambda_G$	Finite-N exposure parameter of the support-best global route and its limiting critical value in Cor. 2.
$\mathsf{K}^{\mathrm{loc}}_N, \mathsf{ABS}$	Minimum cost of a locally visible exact repair under the critical route profile of Cor. 2, and the state recorded when no exact repair is locally visible.
$C^{\mathrm{loc,det}}_N$	Minimum expected distinct-call count, under the reference truth and conditional on baseline abstention, of a deterministic repair-or-abstain protocol that is exact on every repairable realised support state.
$\Pr^\star_N, \mathbb{E}^\star_N$	Reference-world probability and expectation conditional on baseline abstention.
$\mathcal{A}^G_h, \mathcal{B}^M_h$	One-query global exact repair and canonical h-query module exact repair in Thm. 3.
$Y_\mathcal{Q}, \mathsf{E}^T_\mathcal{Q}, \mathcal{R}_\pi, \succeq_T, \succ_T, \|_T$	Repair transcript, induced target experiment, target Bayes risk, and target-decision comparison relations.
r, c	Limiting row-route and column-route readiness parameters in Thm. 4.
w_G, β_G, t_N	In Thm. 3, the critical global-route rare-support burden, critical client-scaling exponent defined by $e_{w_G}(\beta_G, \gamma) = 0$, and offset $t_N = \log K_N - \beta_G \log N$.
$\theta_G, r_G(t)$	Limiting exposure coefficient of the support-best global route and its critical-window readiness profile $r_G(t) = 1 - \exp\{-\theta_G e^{-(w_G-1)t}\}$.
$\Theta^{\mathrm{hom}}_G, \Theta^{\mathrm{split}}_G$	Effective global-route exposures for the homogeneous and split client-population decompositions in Thm. 3(c).
$\succeq_{\mathrm{cell}}$	Conditioning-cell domination used by Cor. 3.
$\mathrm{SurfCov}_T(\mathcal{Q}; z), \succeq_{\mathrm{path}}$	Surfaced target-disagreement coverage at state z and the minimal-support-path domination order used in the preservation proof.

B Proof of the Lossless Repair-Support Representation

For each residual disagreement $\zeta \in \mathrm{Amb}_T$, let $\mathcal{E}_\zeta := \{I \in \mathcal{Q}_T : \zeta \in \mathrm{Cov}_T(I)\}$. By the definition of exact repair, $\mathcal{Q}$ is exact exactly when it intersects every $\mathcal{E}_\zeta$, so inclusion-minimal exact repairs are precisely the minimal transversals of this residual-disagreement hypergraph.

For a fixed repair $\mathcal{Q}$, readiness is the conjunction of $\{\phi_I : I \in \mathcal{Q}\}$. Choosing one minimal support path $R_I \in \mathfrak{P}_I$ for each query and taking their union gives a sufficient support-route set for the repair. Conversely, every sufficient route set contains such a union. Removing strict supersets therefore gives exactly the composed path family $\mathfrak{P}(\mathcal{Q})$ and its corresponding monotone readiness indicator $\phi_\mathcal{Q}$ defined in the main text.

Exactness, costs, and the probabilities of these readiness functions under μ determine the repair identities, readiness probabilities, and cost–readiness orderings. An isomorphism preserving coverage, costs, query support paths, and μ preserves each of these quantities. This proves Prop. 1. The

transversal and minimal-path steps are standard; the proposition records their lossless coupling for the target-specific federated repair task.

C Proof of the Adaptive Witness Characterisation

Fix a realised support state z and reference graph G_0. Invisible queries return $\perp$ under every candidate on this fixed interface and hence cannot distinguish graphs. The relevant witness sets are therefore exactly the reference-relative sets $\mathcal{W}_z(G_0, G)$ defined in the main text.

Deterministic equality. Let Π be an exact deterministic adaptive protocol and let S_0 be the set of distinct visible queries asked on its G_0 transcript. If S_0 were not a transversal of $\mathsf{W}_{T,z}(G_0)$, there would be a candidate G with $\mathcal{L}_T(G) \neq \mathcal{L}_T(G_0)$ such that $S_0 \cap \mathcal{W}_z(G_0, G) = \varnothing$. Every visible query on the reference transcript would then return the same label under G_0 and G, while every invisible query returns $\perp$ under both. Induction over the adaptive rounds gives identical transcripts, the same stopping time, and the same output, contradicting exactness. Thus the reference transcript is a transversal and $\sum_{I \in S_0} c_I \geq \tau_{T,z}(G_0)$.

Conversely, let $S^\star$ be a minimum weighted transversal. Query every member of $S^\star$ in a fixed order. If all labels agree with d_{G_0}, output $\mathcal{L}_T(G_0)$; otherwise query every remaining visible query, omitting repeated calls, and output the target summary uniquely determined by the complete visible-label transcript. Every target-wrong candidate differs from G_0 on at least one member of $S^\star$, so only candidates with the same target summary can follow the all-reference-label branch. The all-visible fallback is exact on every remaining branch by the proposition's uniform-determination assumption. On the G_0 transcript the protocol asks exactly $S^\star$, proving the reverse inequality and hence Eq. (8).

Randomised fractional lower bound. For a randomised δ-correct protocol, let $x_I := \Pr_{G_0}(\Pi \text{ queries } I \text{ before stopping}), \quad I \in \mathcal{Q}_z$. Fix a target-wrong candidate G and couple the executions under G_0 and G with the same internal seed. Until a member of $\mathcal{W}_z(G_0, G)$ is queried, the two runs receive identical responses and make identical adaptive choices. Let H_G be the event that the reference execution queries at least one member of this witness set. On H_G^c the complete outputs coincide, so $\Pr_{G_0}(H_G^c) \leq \Pr_{G_0}(\text{error}) + \Pr_G(\text{error}) \leq 2\delta$. Therefore $\sum_{I \in \mathcal{W}_z(G_0,G)} x_I \geq \Pr_{G_0}(H_G) \geq 1 - 2\delta$. Hence $y_I = x_I/(1-2\delta)$ is feasible for the fractional transversal program, and $\mathbb{E}_{G_0}[C_\Pi] \geq \sum_{I \in \mathcal{Q}_z} c_I x_I \geq (1-2\delta)\tau^\star_{T,z}(G_0)$, which proves Eq. (8).

The complete repair family remains a uniform object: an exact repair must distinguish every target-discrepant pair in the residual class. The proposition is reference-relative and characterises the cost of one designated transcript. In general, if $C^{\mathrm{unif}}_{T,z}$ is the minimum visible uniform exact-repair cost, then $\tau_{T,z}(G_0) \leq C^{\mathrm{unif}}_{T,z}$, with equality not required outside the special family below.

D Proof of Client-Local Repair-Cost Distortion

For a graph G and realised support configuration Γ, write

$$\mathsf{O}^G_\Gamma(I) = \begin{cases} d_G(I), & I \text{ is support-ready under } \Gamma, \\ \perp, & \text{otherwise.} \end{cases} \tag{18}$$

The pooled oracle applies the same readiness rule after client counts are summed. Readiness depends on the realised support table and the query, not on G. A deterministic adaptive protocol chooses its next query from the previous query-response transcript and stops after finitely many distinct calls; a randomised protocol also uses an internal seed. It is exact if it returns $\mathcal{L}_T(G)$ for every G in the candidate class, and it is δ-correct if its worst-case error on this fixed realised interface is at most δ.

D.1 Causal Family and Residual Target Ambiguity

Fix $h \geq 2$ and let $\mathcal{V}_h = \{P, R, X, T, Z\} \cup \{U_i, V_i : i \in [h]\}$. The reference DAG $G^\star_h$ contains

$$P \to R \to X, \qquad P \to T, \qquad X \to T, \qquad P \to Z, \qquad Z \to U_i, \quad Z \to V_i \quad (i \in [h]). \tag{19}$$

For every non-empty $A \subseteq [h]$, the target-wrong candidate $G_{h,A}$ replaces $X \to T$ by $T \to X$ and additionally contains $U_i \to V_i$ for every $i \in A$. The target summary is the orientation of the edge

between X and T, and $\mathcal{G}_h = \{G_h^\star\} \cup \{G_{h,A} : \varnothing \neq A \subseteq [h]\}$. For any observational CI query I, define its construction-specific residual coverage by $\mathrm{cov}_h(I) := \{A \subseteq [h] : A \neq \varnothing,\ d_{G_{h,A}}(I) \neq d_{G_h^\star}(I)\}$. A singleton query is target-sufficient exactly when its coverage is every non-empty subset of $[h]$.

D.2 Complete CI Disagreement Structure

Lemma 1 (Exact endpoint-orbit classification)**.** *Let $I = (Y, Y', S)$ be any query in the complete observational CI universe on $\mathcal{V}_h$, and write W_i for either U_i or V_i. Then $\mathrm{cov}_h(I) \neq \varnothing$ if and only if exactly one of the following mutually exclusive cases holds.*

Class	*Endpoints*	*Conditioner condition*	*Coverage*
G1	$\{P, X\}$	$R \in S$	*every* $\varnothing \neq A \subseteq [h]$
G2	$\{R, T\}$	$P \in S$	*every* $\varnothing \neq A \subseteq [h]$
G3	$\{X, Z\}$	$R \in S$, $P \notin S$	*every* $\varnothing \neq A \subseteq [h]$
G4	$\{X, W_i\}$	$R \in S$, $P, Z \notin S$	*every* $\varnothing \neq A \subseteq [h]$
M_i	$\{U_i, V_i\}$	$Z \in S$	*exactly the* A *with* $i \in A$

No other non-zero coverage pattern occurs. Consequently, the singleton target-sufficient queries are precisely G1–G4; every such query conditions on P or R; every non-global informative query belongs to one module; and every exact repair containing no global query has at least h queries. The number of non-zero-coverage queries is

$$2\,2^{2h+2} + 2^{2h+1} + 2h\,2^{2h} + h\,2^{2h+2}, \tag{20}$$

and the number of non-zero coverage patterns is exactly $h + 1$.

Proof. Use the ancestral moral graph criterion for d-separation. The modules are exchangeable, so the complete endpoint set has finitely many orbits independent of h: core-core pairs, core-module pairs, the same-module pair U_i, V_i, and pairs from different modules.

First isolate the effect of adding $U_i \to V_i$. In the ancestral moral graph, this addition can only add the chord $U_i - V_i$ inside the three-node star $U_i - Z - V_i$. If $Z \notin S$, the leaves are already connected through Z. If $Z \in S$, deleting Z isolates the module from the rest of the graph, and the chord changes connectivity only when the endpoints themselves are U_i, V_i. The CI label therefore changes exactly for endpoints U_i, V_i with $Z \in S$, and precisely in candidates with $i \in A$. This gives M_i.

It remains to compare the common module-free backbone of $G_h^\star$ and $G_{h,A}$. Its skeleton is $P - R - X - T - P, \quad P - Z$, with pendant module leaves at Z. The reference has the unshielded collider $P \to T \leftarrow X$, whereas every target-wrong candidate has $R \to X \leftarrow T$. The endpoint orbits reduce as follows.

Endpoint orbit	Graphs differ exactly when	Ancestral-moral connectivity reason
P, X	$R \in S$	If $T \notin S$, only the target-wrong graph retains $P - T - X$; if $T \in S$, only the reference retains the moral edge $P - X$ induced by $P \to T \leftarrow X$.
R, T	$P \in S$	If $X \notin S$, only the reference retains $R - X - T$; if $X \in S$, only the target-wrong graph retains the moral edge $R - T$ induced by $R \to X \leftarrow T$.
X, Z	$R \in S$, $P \notin S$	With R deleted and P retained, the surviving route alternates between $X - T - P - Z$ and the moral edge $X - P$ according as $T \notin S$ or $T \in S$.
X, W_i	$R \in S$, $P, Z \notin S$	The module endpoint is pendant through Z, so this is the X, Z calculation with the additional requirement that Z not be deleted.
P, R; P, T; P, Z; R, X; R, Z; X, T; T, Z	never	A common adjacent edge or a common surviving separator fixes connectivity in both graphs.
P, W_i; R, W_i; T, W_i; Z, W_i	never	The pendant route through Z has the same backbone connectivity in both graphs.
Nodes from different modules	never	Their only common connection is through the unchanged Z-centred module star.

Conditioning on additional module variables cannot alter these statements: after ancestral moralisation they are pendant vertices or chords inside their own shielded module and create no new backbone connection. The table therefore exhausts every endpoint orbit for arbitrary h, proving the five classes and their coverage.

For the count, put $n = 2h + 5$. G1 and G2 each have one mandatory conditioner and $n - 3$ optional variables, giving 2^{n-3} queries each. G3 has one mandatory and one forbidden conditioner, giving 2^{n-4}. Each of the $2h$ G4 endpoint pairs has one mandatory and two forbidden conditioners, giving 2^{n-5}. Each of the h module pairs has one mandatory conditioner and $n - 3$ optional variables, giving 2^{n-3}. Summing gives Eq. (20). □

Corollary 4 (Complete exact-repair class)**.** *Let $q_i = (U_i, V_i, \{Z\}), \quad i \in [h]$. The inclusion-minimal exact repairs are exactly: (i) every global singleton from G1–G4; and (ii) every h-query set containing exactly one M_i query for each $i \in [h]$. In particular, $\mathcal{B}_h^M := \{q_i : i \in [h]\}$ is an exact repair of minimum non-global cost h.*

Proof. A global query covers every target-wrong candidate. Without a global query, candidate $G_{h,\{i\}}$ can be separated from the reference only by an M_i query, so every module must be touched. Conversely, one query from each module covers every non-empty A, because each A contains at least one touched index. Inclusion minimality is immediate. □

D.3 Finite-Client Support Construction

All variables are binary. For this construction, a query is ready at a client when every stratum of its conditioning variables appears at least once. Consider two clients with two observations each and, for every i, the following complete assignments.

Client	Row	P	R	X	T	Z	all U_i	all V_i
0	0	0	0	0	0	0	0	0
0	1	0	0	1	1	1	1	0
1	0	1	1	0	1	0	0	1
1	1	1	1	1	0	1	1	1

Lemma 2 (Legal support table)**.** *Under this table: every global informative query in Lem. 1 abstains at every client; every canonical module query q_i is ready at both clients; and the pooled queries*

$$g_R = (P, X, \{R\}), \qquad g_P = (R, T, \{P\}) \tag{21}$$

are ready. Moreover, for every $G \in \mathcal{G}_h$ there exists a strictly positive faithful binary distribution under which this exact finite sample has positive probability.

Proof. At each client, P and R are constant. Lem. 1 shows that every global informative query conditions on P or R, so at least one binary conditioning stratum is absent and the query abstains. At each client, Z takes both values exactly once, so every q_i is ready. In the pooled table, each state of P and each state of R appears twice, so g_P and g_R are ready.

Every graph in $\mathcal{G}_h$ is a finite DAG. Strictly positive faithful discrete parameterisations exist generically. Under any such binary parameterisation every complete assignment has positive probability, so the four displayed ordered observations have positive joint probability. Thus the same realised support pattern is a legal hard instance for every candidate graph; only the correct CI labels depend on the graph. □

Lemma 3 (K-uniform legal support realisation)**.** *For every $K \geq 2$ and integer threshold $m \geq 2$, there is an ordinary horizontal federated system $\Gamma_{h,K,m}$ with exactly K clients and one local data table per client such that every global informative query in Lem. 1 abstains locally, every canonical module query q_i is locally ready, and g_R, g_P in Eq.* (21) *are pooled-ready.*

Proof. Let a type-A client have realised (P, R, Z) counts $(0, 0, 0) : m, \quad (0, 0, 1) : 1, \quad (1, 1, 1) : m-1$, and let a type-B client have $(0, 0, 0) : m-1, \quad (1, 1, 0) : 1, \quad (1, 1, 1) : m$. Include at least one client of each type and assign each of the remaining $K - 2$ clients either

type. Every client contains $2m$ observations. In every individual client, one state of P and the same state of R has count $m-1$, whereas both states of Z have count exactly m. Every global separating query conditions on P or R, so it misses the threshold in at least one stratum; additional conditioners only refine that deficient stratum. Every $q_i = (U_i, V_i, \{Z\})$ is ready. Since both client types occur, pooling gives at least $2m$ observations in each state of P and R, so both displayed global witnesses are pooled-ready.

The remaining variables can be assigned within the displayed (P, R, Z) groups to form complete binary observations. Strictly positive faithful parameterisations exist generically for every finite candidate DAG, and every complete assignment then has positive probability. Hence the realised tables are legal support instances for every candidate graph. □

Lemma 4 (Uniformity over the supplied baseline)**.** *Let I_0 be any singleton target-sufficient query, supplied as a failed baseline and removed from the repair pool. Under every $\Gamma_{h,K,m}$ in Lem. 3, I_0 is a global query and locally abstains; at least one of g_R, g_P is distinct from I_0 and is a pooled-ready global target-sufficient query; and the residual candidate class after $I_0 \downarrow$ remains $\mathcal{G}_h$.*

Proof. The first and third claims follow from Lem. 1, Lem. 3, and the fact that abstention removes no candidate. The queries g_R and g_P have different endpoint pairs, so at most one equals I_0; both are pooled-ready by Lem. 3. □

D.4 Exact Structural and Adaptive Costs

Proof of Thm. 1. Fix $h \geq 2$, $K \geq 2$, $m \geq 2$, a federated system $\Gamma_{h,K,m}$ from Lem. 3, and any failed singleton baseline I_0. Lem. 4 supplies a global target-sufficient query distinct from I_0. Its graph-implied label separates $G_h^\star$ from every target-wrong candidate, so one additional query is structurally sufficient. Zero additional queries cannot distinguish the two remaining target summaries; hence $C_T^{\mathrm{str}}(I_0) = 1$. The same query is pooled-ready, proving $C_{T,\Gamma_{h,K,m}}^{\mathrm{pool,ad}}(I_0) = 1$.

For the local upper bound, query $q_1, \ldots, q_h$. Under the reference all answers have their reference labels. Under candidate $G_{h,A}$, query q_i has the opposite label exactly when $i \in A$. Since $A \neq \varnothing$, the transcript determines whether the target edge has the reference or reversed orientation. Thus h local calls suffice.

For the deterministic lower bound, fix $i \in [h]$ and write $G_i = G_{h,\{i\}}$. Run an exact deterministic protocol on the reference. Suppose it stops without querying any M_i witness. Every global query abstains under both graphs. Every zero-coverage query has the same label under both graphs, and every M_j query with $j \neq i$ also has the same label. Readiness is graph-independent, so the complete adaptive transcript is identical under $G_h^\star$ and G_i. The protocol would return the same target summary on two graphs with opposite target orientations, contradicting exactness. This holds for every i, so the reference transcript touches all h modules. Lem. 1 assigns each informative non-global query to one module only, hence at least h distinct calls are required. Therefore $C_{T,\Gamma_{h,K,m}}^{\mathrm{loc,ad}}(I_0) = h$.

For the randomised statement, couple the executions under $G_h^\star$ and G_i using the same internal seed. Let E_i be the event that the reference execution queries an M_i witness. Until the first such query, both executions receive identical responses and make identical adaptive choices. On E_i^c, the complete transcripts and outputs coincide, so the total-variation distance between the two output laws is at most $\Pr_{G_h^\star}(E_i)$.

Let $a = \mathcal{L}_T(G_h^\star)$. Worst-case δ-correctness gives

$$\Pr_{G_h^\star}(\widehat{\mathcal{L}}_T = a) \geq 1 - \delta, \qquad \Pr_{G_i}(\widehat{\mathcal{L}}_T = a) \leq \delta. \tag{22}$$

The output laws therefore differ in total variation by at least $1 - 2\delta$, implying $\Pr_{G_h^\star}(E_i) \geq 1 - 2\delta$. By linearity of expectation,

$$\mathbb{E}_{G_h^\star}[Q_\Pi^{\mathrm{mod}}] = \sum_{i=1}^{h} \Pr_{G_h^\star}(E_i) \geq (1 - 2\delta)h. \tag{23}$$

Since one query belongs to at most one module class, $Q_\Pi^{\mathrm{call}} \geq Q_\Pi^{\mathrm{mod}}$. This proves the randomised lower bound in Thm. 1. □

D.5 Bounded Aggregate-Before-Decision Strengthening

Throughout this subsection, bounded-depth query responses use $A_I^{(b)}$ from Sec. 4.2. Define the query and repair formation depths $d_\Gamma(I) := \min\{b : \mathrm{Ready}_I^{(b)} = 1\}$ and $d_\Gamma(\mathcal{Q}) := \max_{I \in \mathcal{Q}} d_\Gamma(I)$, respectively, with $+\infty$ when unavailable.

Cor. 1 keeps the causal topology and the disagreement classification of Lem. 1 fixed while changing state cardinalities and realised support allocation.

Lemma 5 (Formation-depth support realisation)**.** *Fix integers $K \geq b_0 \geq 2$ and $m \geq 2$. Replace the binary variables P and R in the family of App. D by b_0-state variables with state space $[b_0]$, leaving the DAG topology and all other variable state spaces unchanged. There is an ordinary K-client horizontal federation $\Gamma_{h,K,m}^{(b_0)}$ with the following properties.*

(i) Every canonical module query $q_i = (U_i, V_i, \{Z\})$ has $d_\Gamma(q_i) = 1$.

(ii) Every global informative query has formation depth at least b_0.

(iii) The canonical global witnesses $g_R = (P, X, \{R\})$, $g_P = (R, T, \{P\})$ have formation depth exactly b_0.

(iv) For every candidate graph in $\mathcal{G}_h$, the displayed federation is a legal realised support instance under some strictly positive faithful finite-discrete distribution.

Proof. Lem. 1 is a d-separation classification. Changing the number of states of P and R does not alter the DAG, endpoint sets, or conditioning-variable sets, so the five non-zero coverage classes G1–G4 and M_i remain exactly the same. In particular, every global informative query conditions on P or R, every informative non-global query belongs to one module, and every exact repair without a global query must touch all h modules.

For each $j \in [b_0]$, define a type-j client with $2m$ rows. Set $P = R = j$ on every row, place exactly m rows in $Z = 0$ and m rows in $Z = 1$, and complete the remaining coordinates arbitrarily to full finite-discrete observations. Use at least one client of each type; if $K > b_0$, let the remaining clients duplicate arbitrary existing types.

Each q_i conditions only on Z. Every single client contains m observations in each Z stratum, hence q_i is ready already at depth one and $d_\Gamma(q_i) = 1$.

Now take any coalition C with $|C| < b_0$. Since each client carries only one P state and the same single R state, C covers at most $|C|$ states of either variable. Thus some $j^\star \in [b_0]$ is absent for both P and R. Every global informative query conditions on P or R by Lem. 1; its all-strata support partition therefore contains at least one cell with $P = j^\star$ or $R = j^\star$, and that cell has count zero in C. The query cannot be ready. Adding further conditioning variables only refines this empty stratum and cannot create a shortcut. Hence every global informative query has formation depth at least b_0.

Conversely, choose one client of every type. In their union every state of P and every state of R appears exactly $2m$ times. Therefore g_P and g_R satisfy the per-cell threshold m and are ready, so both have formation depth exactly b_0.

Finally, every candidate is a finite DAG with finite variable state spaces. Strictly positive faithful discrete Markov parameterisations exist generically. Under any such parameterisation every complete assignment has positive probability. Hence each displayed finite client table has positive probability under an appropriate faithful distribution for every candidate graph. The construction therefore uses neither population structural zeros nor a query-specific sampling law. □

Proof of Cor. 1. Fix the family and federation from Lem. 5, and let I_0 be any supplied singleton target-sufficient baseline after $I_0 \downarrow$. The structural disagreement classification is unchanged from Lem. 1. The two canonical global witnesses g_R and g_P have different endpoint pairs, so at most one can equal I_0; consequently at least one distinct singleton global witness remains structurally available. Zero additional queries cannot distinguish the two target summaries. Thus $C_T^{\mathrm{str}}(I_0) = 1$.

Suppose first that $1 \leq b < b_0$. Lem. 5 blocks every global informative query, while all canonical module queries $q_1, \ldots, q_h$ are ready. Querying these h witnesses is exact: in $G_h^\star$ all return their

reference labels, whereas under $G_{h,A}$ precisely the modules indexed by $A \neq \varnothing$ change label. Hence h calls suffice.

For the deterministic lower bound, fix $i \in [h]$ and let $G_i = G_{h,\{i\}}$. Run an exact deterministic protocol on the reference world. If the reference execution never queries an M_i witness, then every global informative query abstains under both $G_h^\star$ and G_i, every zero-coverage query has the same label, and every M_j query with $j \neq i$ has the same label. Readiness is fixed by the realised federation and is graph-independent. The complete adaptive transcripts are therefore identical although the target summaries differ, contradicting exactness. This holds for every i, and each informative non-global query belongs to only one module. Hence at least h distinct calls are necessary, giving $C_{T,\Gamma}^{(b),\mathrm{ad}}(I_0) = h, \qquad 1 \leq b < b_0.$

For a worst-case δ-correct randomised protocol, couple the reference execution with G_i using the same internal seed and let E_i be the event that the reference execution queries an M_i witness. On E_i^c the two transcripts and outputs coincide. Exactly as in Eqs. (22)–(23), worst-case δ-correctness gives $\Pr_{G_h^\star}(E_i) \geq 1 - 2\delta$. Summing over modules yields

$$\mathbb{E}_{G_h^\star}[Q_\Pi^{\mathrm{call}}] \geq \mathbb{E}_{G_h^\star}[Q_\Pi^{\mathrm{mod}}] \geq (1 - 2\delta)h. \tag{24}$$

If $b \geq b_0$, Lem. 5 makes both canonical global witnesses ready. At least one is distinct from I_0, so one additional call is sufficient; zero is impossible. Therefore $C_{T,\Gamma}^{(b),\mathrm{ad}}(I_0) = 1, \qquad b_0 \leq b \leq K.$ This proves Eq. (11) and the randomised lower bound. Since b_0 and h were arbitrary subject only to $K \geq b_0 \geq 2$ and $h \geq 2$, cost-recovery depth and cost distortion can be chosen independently. □

Corollary 5 (Local transcript indistinguishability)**.** *On the same causal family, there is a second legal support realisation in which every query with non-zero residual coverage abstains locally. Every adaptive local transcript then has the same distribution under $G_h^\star$ and every $G_{h,A}$ for every query budget, so the minimax target error is at least $1/2$. A pooled global query remains ready and the pooled exact cost is one.*

Proof. Use two clients and let client $k \in \{0, 1\}$ contain only observations with $P = R = Z = k$, while assigning the remaining binary variables so that the pooled table is complete for the displayed single-variable strata. Every global witness conditions on P or R, and every module witness conditions on Z, so Lem. 1 makes every informative query abstain locally. Zero-coverage queries have identical graph labels. The local transcript is therefore graph-independent, and the two-point indistinguishability bound between the reference and any singleton candidate gives minimax error at least $1/2$. Pooling restores both states of P and R, so g_P and g_R are ready. □

Under the same unit-cost convention, the corresponding static costs are also $1 : h : 1$.

E Proof of Stochastic Bounded-Depth Formation and Recovery

This section proves Thm. 2. The main-text result is the off-critical formation-depth consequence; the full bounded-depth lemma below also records the critical Poisson-packing law and its client-local $b = 1$ degeneration.

E.1 Bounded-Depth Sparse Route Formation

Lemma 6 (Bounded-depth sparse route formation)**.** *Fix a support-competitive CI route with rare conditioning cells $\mathcal{C}_1, \ldots, \mathcal{C}_s$ satisfying*

$$P(\mathcal{C}_j) = \kappa_j p_N\{1 + o(1)\}, \qquad \kappa_j > 0, \tag{25}$$

fixed thresholds $m_1, \ldots, m_s$, total rare-support burden $w := \sum_{j=1}^s m_j$, and finitely many common conditioning cells whose probabilities are uniformly bounded below. Let K_N independent near-balanced clients have sizes $n_{k,N} = n_N + O(1)$, where $n_N \to \infty$ and $n_N p_N \to 0$. For a fixed aggregation depth $b \geq 1$, set

$$\ell_b := \left\lceil \frac{w}{b} \right\rceil, \qquad r_b := w - b(\ell_b - 1), \qquad \Xi_{\ell_b,N} := K_N (n_N p_N)^{\ell_b}. \tag{26}$$

Then

$$\Xi_{\ell_b,N} \to 0 \Longrightarrow \Pr(\mathrm{Ready}_I^{(b)} = 1) \to 0, \qquad \Xi_{\ell_b,N} \to \infty \Longrightarrow \Pr(\mathrm{Ready}_I^{(b)} = 1) \to 1. \tag{27}$$

If $\Xi_{\ell_b,N} \to \lambda \in (0,\infty)$*, let*

$$\mathcal{A}_{\ell_b}(m) := \left\{ u \in \mathbb{N}^s : 0 \le u_j \le m_j,\ \sum_{j=1}^{s} u_j = \ell_b \right\}, \tag{28}$$

and let Z_u*,* $u \in \mathcal{A}_{\ell_b}(m)$*, be independent Poisson variables with means*

$$\mu_u = \lambda \prod_{j=1}^{s} \frac{\kappa_j^{u_j}}{u_j!}. \tag{29}$$

For a realised vector Z*, define*

$$\pi_m(Z) := \max \left\{ \sum_u x_u : 0 \le x_u \le Z_u,\ x_u \in \mathbb{Z}_+,\ \sum_u x_u u \le m \right\}, \tag{30}$$

where vector inequalities are coordinatewise. Then

$$\Pr(\mathrm{Ready}_I^{(b)} = 1) \longrightarrow \Pr\{\pi_m(Z) \ge r_b\}. \tag{31}$$

At $b = 1$*,* $\ell_1 = w$ *and* $r_1 = 1$*, so the critical law above reduces to a single full-threshold pattern with coefficient* $\theta_I := \prod_j \kappa_j^{m_j} / \prod_j m_j!$*. If additionally* $p_N = N^{-\gamma}$*, writing* $w = w_J$ *and* $r_J(n,N)$ *for one-client route readiness,*

$$r_J(n,N) = \theta_J n^{w_J} N^{-\gamma w_J}\{1 + o(1)\}, \tag{32}$$

and hence

$$R_J(N,K_N) = 1 - \exp\{-\Lambda_J(N,K_N)(1+o(1))\}, \qquad \Lambda_J(N,K_N) = \theta_J N^{w_J(1-\gamma)} K_N^{1-w_J}. \tag{33}$$

Proof. If an at-most-b-client coalition supplies the w required rare observations, at least one selected client supplies $\ell_b = \lceil w/b \rceil$ of them. For a fixed client, the probability of at least ℓ_b query-relevant rare observations is $O((n_N p_N)^{\ell_b})$, uniformly over near-balanced sizes. A union bound over the K_N clients proves the subcritical implication in Eq. (27).

For the supercritical implication, partition the fixed rare-demand multiset into at most b deterministic block patterns, each of size at most ℓ_b. An exact block pattern v of size $u \le \ell_b$ occurs in one client with probability

$$(n_N p_N)^u \prod_{j=1}^{s} \frac{\kappa_j^{v_j}}{v_j!}\{1 + o(1)\}. \tag{34}$$

Its federation-wide expected multiplicity is therefore a positive constant times $K_N(n_N p_N)^u$. This diverges for $u = \ell_b$ by assumption and for every $u < \ell_b$ because $n_N p_N \to 0$. Independence across clients and a finite union bound give enough distinct clients for all finitely required block types with probability tending to one. Conditional on any fixed rare block, the remaining $n_N - O(1)$ observations retain constant-order mass in every common conditioning cell, so failure of any fixed common-cell threshold is exponentially small. This proves the supercritical implication.

Now suppose $\Xi_{\ell_b,N} \to \lambda \in (0,\infty)$. Each exact ℓ_b-pattern $u \in \mathcal{A}_{\ell_b}(m)$ has one-client probability $(\mu_u + o(1))/K_N$, uniformly over near-balanced sizes. These exact pattern categories are mutually exclusive within a client, so the finite multinomial law of small numbers gives joint total-variation convergence of their client counts to the independent Poisson vector $(Z_u)_u$. The expected number of clients with more than ℓ_b query-relevant rare observations is $O(\Xi_{\ell_b,N} n_N p_N) = o(1)$, whereas every exact pattern of size at most $\ell_b - 1$ has diverging multiplicity with probability tending to one.

On this high-probability event, any ready coalition of at most b clients must contain at least $r_b = w - b(\ell_b - 1)$ clients contributing exactly ℓ_b useful rare observations: if only $x < r_b$ selected clients contributed ℓ_b, the coalition would supply at most $b(\ell_b - 1) + x < w$ useful rare observations. Their aggregate pattern must fit coordinatewise inside m, so readiness implies $\pi_m(Z) \ge r_b$. Conversely,

if $\pi_m(Z) \geq r_b$, choose r_b compatible critical-size patterns. The residual demand has total size $w - r_b\ell_b = (b - r_b)(\ell_b - 1)$ and can be partitioned into at most $b - r_b$ blocks of size at most $\ell_b - 1$; every such smaller exact pattern is abundant. Together with the common-cell argument above, these clients form a ready coalition. This proves Eq. (31).

Finally take $b = 1$. Then $\ell_b = w$ and $r_b = 1$, so the only critical-size pattern compatible with the full demand is the threshold vector itself. The exact-pattern expansion gives Eq. (32). Summing the resulting vanishing one-client probabilities over independent near-balanced clients and using $\log(1 - x) = -x + O(x^2)$ gives Eq. (33). Thus the bounded-depth lemma exactly recovers the client-local support engine used below. □

E.2 Off-Critical Formation Depth

Corollary 6 (Off-critical stochastic formation depth)**.** *Under Lem. 6, suppose additionally that* $p_N = N^{-\gamma}$, $K_N = N^{\beta+o(1)}$, $0 < \beta < 1$, *and* $\beta + \gamma > 1$. *Define*

$$e_\ell(\beta, \gamma) := \ell(1 - \gamma) - (\ell - 1)\beta, \qquad \ell_\star := \max\{\ell \geq 1 : e_\ell(\beta, \gamma) > 0\}. \tag{35}$$

If $e_\ell(\beta, \gamma) \neq 0$ *for every integer* $1 \leq \ell \leq w$, *then*

$$d_{\Gamma_N}(I) \xrightarrow{p} \left\lceil \frac{w}{\ell_\star} \right\rceil. \tag{36}$$

The same conclusion holds jointly for every fixed finite family of such routes, without cross-route independence.

Proof. Because $\beta + \gamma > 1$, $e_\ell(\beta, \gamma) = \beta + \ell(1 - \gamma - \beta)$ is strictly decreasing in ℓ, while $e_1 = 1 - \gamma > 0$; hence $\ell_\star$ is finite and positive. Let $b_\star := \lceil w/\ell_\star \rceil$. At depth $b_\star$, $\ell_{b_\star} = \lceil w/b_\star \rceil \leq \ell_\star$, so $\Xi_{\ell_{b_\star},N} = N^{e_{\ell_{b_\star}}(\beta,\gamma)+o(1)} \to \infty$, and Lem. 6 gives readiness probability tending to one. If $b_\star > 1$, then $\ell_{b_\star - 1} > \ell_\star$, so $\Xi_{\ell_{b_\star-1},N} \to 0$ and readiness at depth $b_\star - 1$ tends to zero. Monotonicity of $\mathrm{Ready}_I^{(b)}$ in b proves Eq. (36); the case $b_\star = 1$ is immediate. A finite union bound yields the joint statement for any fixed finite route family. □

Eq. (36) is Eq. (13) in Thm. 2. The general critical law in Eq. (31) is retained here rather than promoted to the main theorem because a reduced critical collision order $\ell \geq 2$ lies in the inherited near-balanced regime only when its corresponding power-law wall has $\beta(\ell) = \ell(1 - \gamma)/(\ell - 1) < 1$, equivalently $\gamma > 1/\ell$.

E.3 Canonical Route Scope

The single-burden law above is intentionally restricted to support-competitive canonical routes. In the complete admissible CI universe, conditioning refinements can create cells on smaller rarity scales such as $O(p_N^2)$, so a single scalar burden need not describe an arbitrary refined query. For the causal family used by Thms. 1 and 3, Lem. 1 gives the complete non-zero coverage classes. Within each informative class, the support-coarsest canonical representative has no larger unit cost, and refinement coherence makes every refined representative no easier at any realised bounded depth. Thus a refined query in an exact repair or reference-witness transversal can be replaced by its canonical representative without losing residual coverage or increasing cost/readiness burden. Consequently the repair frontier and reference-witness minimum used in the headline results are attained within the finite support-competitive canonical classes to which Lem. 6 applies.

F Route-Profile Transfer and Critical Repair Cost

This section separates the route-profile-to-repair implication from the sparse formation mechanism. We first state the general transfer principle and then apply it to the client-local critical profile used in Cor. 2.

F.1 Abstract Route-Profile-to-Repair-Cost Transfer

Proposition 3 (Abstract route-profile-to-repair-cost transfer)**.** *Let $h_N \to \infty$. Under the reference-world support law conditional on baseline abstention, let G_N, M_N, and P_N denote respectively a global-singleton route event, the common module-route event, and the competing-global route event. Suppose*

$$\overset{\star}{\Pr_N}(G_N) \to 1 - e^{-\lambda_G}, \qquad \overset{\star}{\Pr_N}(M_N) \to 1, \qquad \overset{\star}{\Pr_N}(P_N) \to 0, \tag{37}$$

where $\lambda_G \in [0, \infty]$ and $e^{-\infty} := 0$. Assume the complete ex-post local repair classification is

$$\mathsf{K}_N^{\mathrm{loc}} = \begin{cases} 1, & G_N \cup P_N, \\ h_N, & G_N^c \cap P_N^c \cap M_N, \\ \mathsf{ABS}, & G_N^c \cap P_N^c \cap M_N^c, \end{cases} \tag{38}$$

and write $H_N := G_N^c \cap M_N \cap P_N^c$. If every exact deterministic protocol has reference cost at least h_N on H_N, while every support-state-uniformly δ-correct randomised protocol has conditional expected cost at least $(1-2\delta)h_N$ there, then repairability tends to one and

$$\begin{aligned} \overset{\star}{\Pr_N}(\mathsf{K}_N^{\mathrm{loc}} = 1 \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}) &\to 1 - e^{-\lambda_G}, \\ \overset{\star}{\Pr_N}(\mathsf{K}_N^{\mathrm{loc}} = h_N \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}) &\to e^{-\lambda_G}, \end{aligned} \tag{39}$$

$\mathbb{E}_N^\star[\mathsf{K}_N^{\mathrm{loc}} \mid \mathsf{K}_N^{\mathrm{loc}} \neq \mathsf{ABS}]/h_N \to e^{-\lambda_G}$, and

$$\liminf_{N\to\infty} \frac{\mathbb{E}_N^\star[C_\Pi]}{h_N} \geq (1-2\delta)e^{-\lambda_G}. \tag{40}$$

If, in addition, a safe deterministic protocol uses $h_N + O(1)$ calls on H_N, $O(1)$ calls on global-ready states, and total cost $O(h_N)$ elsewhere, then $C_N^{\mathrm{loc,det}}/h_N \to e^{-\lambda_G}$.

Proof. No independence among G_N, M_N, P_N is required. Since $G_N^c \setminus H_N \subseteq M_N^c \cup P_N$, Eq. (37) gives

$$\overset{\star}{\Pr_N}(H_N) = \overset{\star}{\Pr_N}(G_N^c) + o(1) \to e^{-\lambda_G}. \tag{41}$$

Moreover, the abstention event in Eq. (38) is contained in M_N^c, so its probability is $o(1)$. The two conditional repair-cost probabilities in Eq. (39) therefore follow directly from the complete classification and Eq. (41). Dividing the conditional expected ex-post repair cost by h_N removes the cost-one contribution because $h_N \to \infty$.

On every H_N state, the assumed deterministic and randomised lower bounds yield respectively $h_N \Pr_N^\star(H_N)$ and $(1-2\delta)h_N \Pr_N^\star(H_N)$ after integration over the support law, proving the lower bounds and Eq. (40). Under the additional safe-protocol condition, its expected deterministic cost is at most $h_N \Pr_N^\star(H_N) + o(h_N)$; together with the lower bound and Eq. (41), this gives $C_N^{\mathrm{loc,det}}/h_N \to e^{-\lambda_G}$. □

Remark 7 (Average-risk variant)**.** *If correctness is required only after averaging over the random client-data configuration conditional on repairability, the error budget may be concentrated on hard states. The corresponding general lower bound weakens to* $\liminf_{N\to\infty} \mathbb{E}_N^\star[C_\Pi]/h_N \geq (e^{-\lambda_G} - 2\delta)_+$.

F.2 Critical Client-Local Repair-Cost Consequence

Proof of Cor. 2. Let $h_N \to \infty$ in the connected causal family of App. D. The complete endpoint-orbit classification remains valid for every finite h_N, and all h_N canonical module queries share the same Z-support route. Let G_N, M_N, P_N, and B_N denote respectively the support-best global route, common module route, support-coarsest competing-global route, and generated-baseline readiness event. Write their canonical rare-support burdens as

$$w_{\mathrm{base}} > w_G > w_M > 1, \qquad w_P > w_G. \tag{42}$$

For every finite (N, h_N), the existing construction admits strictly positive faithful parameters.

Use the client-local $b = 1$ degeneration in Eq. (33). Define $\beta_G := w_G(1-\gamma)/(w_G-1)$, equivalently $e_{w_G}(\beta_G, \gamma) = 0$, and take $\log K_N = \beta_G \log N + t_N$ with $t_N \to t \in \mathbb{R}$. Then the global exposure converges to $\lambda_G := \theta_G e^{-(w_G-1)t}$, so $\Pr(G_N) \to 1 - e^{-\lambda_G}$. Because $\beta(d) := d(1-\gamma)/(d-1)$ is strictly decreasing for $d > 1$, Eq. (42) and the two-phase part of Lem. 6 give $\Pr(M_N) \to 1$, $\Pr(P_N) \to 0$, and $\Pr(B_N) \to 0$.

Let $\Pr_N^\star(E) := \Pr(E \mid B_N^c)$. Since $\Pr(B_N) = o(1)$,

$$|\Pr(E \mid B_N^c) - \Pr(E)| \le \frac{\Pr(B_N)}{1 - \Pr(B_N)} = o(1) \tag{43}$$

for every event E. Thus these limits give exactly the route-profile assumptions of Prop. 3 under $\Pr_N^\star$.

Cor. 4 gives the complete ex-post classification in Eq. (38): a ready global route gives cost one; if global routes are blocked but the common module route is ready, every exact repair must touch all h_N modules and the canonical bundle attains cost h_N; otherwise the realised residual task is not locally repairable. On the hard states, Prop. 2 gives deterministic reference cost h_N and the support-state-uniform randomised lower bound $(1-2\delta)h_N$. The global-first safe protocol from the deterministic construction has the matching $h_N + O(1)$ versus $O(1)$ cost profile. Applying Prop. 3 proves every conclusion of Cor. 2. □

G Proof of Client-Local Repair-Order Reversal

G.1 Complete Repair Family

Thm. 3 uses the same connected residual DAG family, reference graph, candidate class, and baseline as App. D. Lem. 1 gives the complete all-order non-zero coverage classification. Its four global endpoint orbits all separate every residual target disagreement; hence every query in a global orbit is a singleton exact repair. Every remaining non-zero query belongs to exactly one module orbit M_i, so an exact repair without a global query must touch all h modules. Cor. 4 therefore implies

$$\mathfrak{C}_{T,h} = \big\{\{g\} : g \text{ is global}\big\} \cup \min_{\subseteq}\big\{\mathcal{Q} : \mathcal{Q} \cap M_i \neq \varnothing \ \forall i \in [h]\big\}. \tag{44}$$

In particular, $\mathcal{A}_h^G$ has unit cost, $\mathcal{B}_h^M$ has cost h, and there is no mixed or higher-order shortcut of cost below h.

Among global singleton repairs, the support-coarsest conditioning partitions are the R partition, represented by $g_R = (P, X, \{R\})$, and the P partition, represented by $g_P = (R, T, \{P\})$. Every other global query refines one of these partitions. Each canonical module query $q_i = (U_i, V_i, \{Z\})$ uses the same Z partition; hence all queries in $\mathcal{B}_h^M$ share one readiness event.

G.2 A Single Faithful Matched Population

Let P, R, Z have state spaces $\{0, 1, \dots, \nu_P\}$, $\{0, 1, \dots, \nu_R\}$, and $\{0, 1, \dots, \nu_Z\}$, respectively; state zero is common and positive states are rare. Let T be binary and set $p_N = N^{-\gamma}, \qquad 0 < \gamma < 1.$ Choose constants $a_s, b_j, c_k > 0$ and require

$$\Pr(P = s) = a_s p_N\{1 + o(1)\}, \quad \Pr(R = j) = b_j p_N\{1 + o(1)\}, \quad \Pr(Z = k) = c_k p_N\{1 + o(1)\} \tag{45}$$

for positive states.

For the support-coarsest R-global partition, baseline (R, T) partition, module Z partition, and support-coarsest P-global partition, assign cell thresholds $m_{G,j}$, $m_{0,j,t}$, $m_{M,k}$, and a P-partition threshold vector, with burdens $w_G = \sum_j m_{G,j}, \quad w_{\text{base}} = \sum_{j,t} m_{0,j,t}, \quad w_M = \sum_k m_{M,k}, \quad w_P,$ Choose the burdens and rarity exponent so that

$$\begin{aligned} &w_{\text{base}} > w_G > w_M > 1, \qquad w_P > w_G, \qquad \frac{1}{w_M} < \gamma < 1, \\ &\beta_J := \frac{w_J(1-\gamma)}{w_J - 1}, \qquad J \in \{\text{base}, G, M, P\}. \end{aligned} \tag{46}$$

The threshold contract is refinement coherent: the thresholds of subcells refining a parent cell sum to at least the threshold of that parent. In particular, $\sum_{t=0}^{1} m_{0,j,t} \geq m_{G,j}$. Thus readiness of a refined query implies readiness of its coarser ancestor, and refined R-, P-, and module queries have burden no smaller than w_G, w_P, w_M, respectively.

Lemma 7 (Robust faithful realisation of the linked rare-state population)**.** *For every graph G in the residual family and all sufficiently large N, there is a strictly positive distribution $P_{G,N}$, faithful to G, satisfying Eq.* (45)*. There are constants $q_{j,t} \in (0,1)$, bounded away from zero and one, with $q_{j,0} + q_{j,1} = 1$, such that*

$$\Pr(R = j, T = t) = b_j q_{j,t} p_N \{1 + o(1)\}. \tag{47}$$

All common conditioning cells have probabilities bounded away from zero. Moreover the faithful realisation can be chosen so that every theorem-relevant rare support-cell probability differs from the linked construction below by a relative $o(1)$. Consequently its route coefficients satisfy

$$\theta_{G,N} = \theta_G\{1 + o(1)\}, \qquad \theta_{\mathrm{base},N} = \theta_{\mathrm{base}}\{1 + o(1)\}, \qquad \theta_{M,N} = \theta_M\{1 + o(1)\}, \tag{48}$$

where

$$\theta_G = \frac{\prod_j b_j^{m_{G,j}}}{\prod_j m_{G,j}!}, \quad \theta_{\mathrm{base}} = \frac{\prod_{j,t}(b_j q_{j,t})^{m_{0,j,t}}}{\prod_{j,t} m_{0,j,t}!}, \quad \theta_M = \frac{\prod_k c_k^{m_{M,k}}}{\prod_k m_{M,k}!}. \tag{49}$$

All marginalisation identities of the joint law, including $\Pr(R = j) = \sum_t \Pr(R = j, T = t)$, hold exactly for every N.

Proof. First construct a linked positive Markov centre $\bar{P}_{G,N}$. Put $\bar{P}_{G,N}(P = s) = a_s p_N$ for $s \geq 1$ and the remaining mass on $P = 0$. For every positive state j of R, choose bounded coefficients $\eta_{s,j}$, with at least one varying in s, and set $\bar{P}_{G,N}(R = j \mid P = 0) = b_j p_N, \qquad \bar{P}_{G,N}(R = j \mid P = s) = b_j p_N(1 + \eta_{s,j}),\ s \geq 1$, with coefficients small enough to preserve strict positivity; put the remaining conditional mass on $R = 0$. Since $\bar{P}_{G,N}(P \neq 0) = O(p_N)$, $\bar{P}_{G,N}(R = j) = b_j p_N\{1 + O(p_N)\}$, and variation in s keeps the edge $P \to R$ active. Construct $Z \mid P$ analogously, giving $\bar{P}_{G,N}(Z = k) = c_k p_N\{1 + O(p_N)\}$ and an active $P \to Z$ edge.

Choose every remaining conditional table of the centre in a compact subset of the strictly positive Markov parameter space. In particular, keep all conditional probabilities of T in $[\epsilon, 1 - \epsilon]$ for fixed $\epsilon > 0$. The contribution to $R = j$ from positive P-states is $O(p_N^2)$, so $\bar{P}_{G,N}(T = t \mid R = j) \to q_{j,t} \in [\epsilon, 1 - \epsilon]$. Hence $\bar{P}_{G,N}(R = j, T = t) = b_j q_{j,t} p_N\{1 + o(1)\}$. Summing over t recovers the same R-marginal exactly, so the baseline cells form a genuine refinement of the global cells by construction. The total rare mass is $O(p_N)$, leaving common cells at constant order.

It remains to impose faithfulness without destroying this linking. For fixed (G, N), write Θ_G^+ for the finite-dimensional strictly positive Markov parameter space of G, and let $\bar{\vartheta}_{G,N} \in \Theta_G^+$ parameterise the centre. Collect the finitely many probabilities used by the baseline, global, module, and competing-global support routes into continuous functionals $F_{1,N}, \ldots, F_{M,N}$; include the required common-cell probabilities as well. Choose any deterministic $\delta_N \downarrow 0$. For each sufficiently large fixed N, positivity of the centre and continuity give a radius $r_{G,N} > 0$ such that every parameter within distance $r_{G,N}$ changes each rare-cell functional by relative less than δ_N and preserves all strict common-cell lower bounds. Define the ambient open tube

$$\mathcal{O}_{G,N} := \left\{\vartheta \in \Theta_G^+ : \|\vartheta - \bar{\vartheta}_{G,N}\| < \min\{r_{G,N}, \delta_N\}\right\}. \tag{50}$$

This is a non-empty open subset of the full positive Markov parameter space, not merely of a lower-dimensional equality slice. For a finite discrete DAG, the unfaithful positive Markov parameters lie in a proper algebraic null subset of Θ_G^+. Therefore $\mathcal{O}_{G,N}$ contains a faithful parameter $\vartheta_{G,N}$. Let $P_{G,N}$ be its Markov distribution.

By construction, every theorem-relevant rare support-cell probability of $P_{G,N}$ equals the corresponding centre probability times $1 + o(1)$, while common cells remain bounded away from zero. Because $P_{G,N}$ is itself one joint distribution, all marginalisation identities remain exact after the perturbation; in particular, $\Pr(R = j) = \sum_t \Pr(R = j, T = t)$ for every finite N. The linked asymptotic laws in Eqs. (45) and (47) therefore follow with the same limiting coefficients. Since every threshold vector is fixed, finite products preserve these coefficients; for example, $\prod_j [b_j\{1+o(1)\}]^{m_{G,j}} = \left(\prod_j b_j^{m_{G,j}}\right)\{1+o(1)\}$. This gives Eqs. (48) and (49). The same argument

applies to any displayed competing-global coefficient. The construction is graph-by-graph because the state spaces and the P, R, Z backbone are common to the residual family. □

The perturbation in Lem. 7 is asymptotically invisible to the phase law. Indeed, replacing a route coefficient θ_J by $\theta_{J,N} = \theta_J\{1 + o(1)\}$ leaves $\Lambda_{J,N} = \theta_J N^{w_J(1-\gamma)} K_N^{1-w_J}\{1 + o(1)\}$ unchanged at leading order. Hence the strict burden inequalities in Eq. (46), the disappearance of baseline and competing-global routes, the module-readiness limit, and the critical profile in Eq. (15) all survive the faithful realisation.

For part (a), fix the reference truth and let every client sample from its corresponding P_N. The graph-by-graph clause establishes realisability of the whole causal family but is not an assumption that clients use different graphs.

G.3 Critical Support-Route Regime

The matched-population calculation reuses the client-local $b = 1$ engine of Lem. 6. For the burdens in Eq. (46), the map $\beta(d) = d(1-\gamma)/(d-1)$ is strictly decreasing, hence

$$\beta_{\text{base}} < \beta_G < \beta_M, \qquad \beta_P < \beta_G. \tag{51}$$

Eq. (33) therefore supplies all homogeneous one-client route probabilities and federation-level exposures used below; no separate stochastic support lemma is needed for Thm. 3.

G.4 Matched-Population Reversal

At $\log K_N = \beta_G \log N + t_N$, Eq. (33) gives

$$\log \Lambda_G = \log \theta_G - (w_G - 1)t_N + o(1), \tag{52}$$

so $\Lambda_G \to \theta_G e^{-(w_G-1)t}$ and $R_G \to r_G(t)$. By Eq. (51), $\Lambda_{\text{base}} \to 0$ and every support-coarsest P-global intensity tends to zero. Refinement coherence prevents any refined global query from being easier, while the complete classification in Eq. (44) excludes any other global coverage route.

Since $w_M < w_G$, $\Lambda_M \to \infty$, so the common Z-route becomes ready with probability tending to one. All queries of $\mathcal{B}_h^M$ share that event, and every non-global exact repair must touch every module; hence $\mathcal{B}_h^M$ is support-best in its class.

For any event E_N,

$$|\Pr(E_N \mid I_0 \downarrow) - \Pr(E_N)| \le \frac{\Pr(I_0 \text{ ready})}{\Pr(I_0 \downarrow)} \longrightarrow 0, \tag{53}$$

so generated baseline abstention changes these limits only by $o(1)$. Pooled rare-cell counts have means of order $Np_N = N^{1-\gamma} \to \infty$; fixed thresholds are therefore met with probability tending to one. This proves Eq. (15) and the stated local cost-readiness trade-off. Direct differentiation shows that r_G is continuous and strictly decreasing, with limits one and zero as $t \to -\infty$ and $t \to \infty$, respectively. The same exposure calculation gives the off-critical endpoint limits: taking $t_N \to +\infty$ drives $\Lambda_G \to 0$, while $t_N \to -\infty$ drives $\Lambda_G \to \infty$; choosing these offsets with $|t_N| = o(\log N)$ leaves the route-exponent ordering above unchanged.

G.5 Critical Window and Quantile Width

For $r \in (0, 1)$, define

$$t(r) := \frac{\log \theta_G - \log[-\log(1-r)]}{w_G - 1}. \tag{54}$$

Under a readiness-constrained rule requiring readiness probability $1 - \varepsilon$, the limiting decision boundary is

$$t_\varepsilon = \frac{\log \theta_G - \log[-\log \varepsilon]}{w_G - 1}. \tag{55}$$

For fixed $0 < r_- < r_+ < 1$ and client-count sequences attaining these quantiles,

$$\beta_N(r_-) - \beta_N(r_+) = \frac{t(r_-) - t(r_+)}{\log N} + o\left(\frac{1}{\log N}\right), \qquad \beta_N(r) := \frac{\log K_N(r)}{\log N}. \tag{56}$$

Inverting $r_G(t) = r$ yields Eq. (54). A readiness requirement $1 - \varepsilon$ is crossed at $t_\varepsilon = t(1-\varepsilon)$, which is Eq. (55). For $t < t_\varepsilon$, both repairs qualify and unit cost selects $\mathcal{A}_h^G$; for $t > t_\varepsilon$, the global singleton has limiting readiness below $1 - \varepsilon$ while $\mathcal{B}_h^M$ still qualifies, so the local decision selects the latter. Pooling selects the singleton throughout.

For a sequence attaining fixed $r \in (0,1)$, Eq. (33) requires and suffices that $\Lambda_G \to -\log(1-r)$. Taking logarithms in Eq. (52) gives $\log K_N(r) = \beta_G \log N + t(r) + o(1)$. Subtracting the locations for r_- and r_+ and dividing by $\log N$ proves Eq. (56); strict monotonicity of $t(r)$ gives a positive leading constant.

Remark 8 (Why the reversal is not tied to a hard threshold)**.** *Suppose an analyst scores each repair by combining readiness with its other advantages. A cheaper repair $\mathcal{A}$ may be preferred when it is highly ready, while a more expensive repair $\mathcal{B}$ may be preferred once $\mathcal{A}$ becomes insufficiently ready. Let $u_A, u_B : [0,1] \to \mathbb{R}$ be continuous, with u_A strictly increasing and $u_A(1) > u_B(1) > u_A(0)$. Then there is a unique $r_\star \in (0,1)$ such that $u_A(r_\star) = u_B(1)$. Because the critical readiness profile of $\mathcal{A}$ moves continuously through $(0,1)$, the scalar preference changes uniquely at $t(r_\star)$. Thus the repair-order reversal is not an artefact of the hard readiness threshold used in the thresholded comparison above.*

G.6 Same-Sample Allocation Reversal

Under the population sequence of Lem. 7, standard multinomial concentration over the finitely many rare and common support cells places the pooled sample in the typical set below with probability $1 - o(1)$. Fix a pooled sample in this typical set, so the linked rare counts satisfy

$$\begin{aligned} M_{R,j,N} &= b_j N^{1-\gamma}\{1 + o(1)\}, \\ M_{R,T,j,t,N} &= b_j q_{j,t} N^{1-\gamma}\{1 + o(1)\}, \qquad \sum_t M_{R,T,j,t,N} = M_{R,j,N}, \\ M_{Z,k,N} &= c_k N^{1-\gamma}\{1 + o(1)\}, \end{aligned} \tag{57}$$

with analogous consistent counts for the P and finitely many support-competitive refined partitions. Assume these rare counts diverge and are pairwise comparable. Allocate each fixed observation independently and uniformly to K_N clients.

Proposition 4 (Allocation-only critical profile)**.** *If $\log K_N - \beta_G \log N \to t$, then conditional on the fixed pooled sample,*

$$\Pr(I_0 \downarrow) \to 1, \qquad \rho^{\mathrm{loc},\star}(\mathcal{A}_h^G \mid I_0 \downarrow) \to r_G(t), \qquad \rho^{\mathrm{loc},\star}(\mathcal{B}_h^M \mid I_0 \downarrow) \to 1. \tag{58}$$

Every competing P-based or refined global repair remains unavailable or no easier.

Proof. A minimal allocation witness for route J selects exactly $m_{J,\ell}$ pooled observations from every rare cell. Its size is w_J, and the number of choices is $H_J(N) = \prod_\ell \binom{M_{J,\ell,N}}{m_{J,\ell}}$. A selected witness is assigned to one client with probability $K_N^{1-w_J}$, giving expected witness count $\Lambda_J^{\mathrm{alloc}} = K_N^{1-w_J} H_J(N) = \theta_J N^{w_J(1-\gamma)} K_N^{1-w_J}\{1 + o(1)\}$. Indicators whose selected observation sets are disjoint are independent. Ordered witness pairs sharing s observations have order $O(M_N^{2w_J - s})$, while the probability that their union is monochromatic is $K_N^{1-(2w_J-s)}$. The standard dependency-graph moment bounds therefore give probability asymptotic to the intensity when it vanishes, a Poisson limit at finite positive intensity, and readiness tending to one when it diverges. At the w_G critical scale, the w_{base} and w_P intensities vanish, the w_G intensity converges to $\theta_G e^{-(w_G-1)t}$, and the w_M intensity diverges. Common-cell thresholds are met because $N/K_N \to \infty$. The complete repair conclusion follows exactly as in part (a). □

Choose finite t, so $r_G(t) \in (0,1)$. Prop. 4 implies that for all sufficiently large N, the fixed pooled sample has positive probability of an allocation in which a global singleton is locally ready and positive probability of an allocation in which it is not, while the module route is ready with probability arbitrarily close to one. Hence two deterministic allocations with the properties in part (b) exist; they differ only in client assignment and preserve the same pooled sample exactly.

G.7 Same-Pooled Exposure Separation

For part (c), client k has size $n_{k,N}$, population $P_{k,N}$, and independent observations. For route J, let rare cells $\mathcal{C}_{J,\ell}$ have fixed thresholds $m_{J,\ell}$ and probabilities

$$P_{k,N}(\mathcal{C}_{J,\ell}) = c_{k,J,\ell,N} p_N \{1 + \varepsilon_{k,J,\ell,N}\}, \quad 0 < c_- \le c_{k,J,\ell,N} \le c_+, \quad \max_{k,J,\ell} |\varepsilon_{k,J,\ell,N}| \to 0. \tag{59}$$

Assume $\min_k n_{k,N} \to \infty$, $\max_k n_{k,N} p_N \to 0$, and uniform negligibility of common-cell threshold failure. Define

$$\begin{aligned} \theta_{k,J,N} &:= \frac{\prod_\ell c_{k,J,\ell,N}^{m_{J,\ell}}}{\prod_\ell m_{J,\ell}!}, \qquad M_{k,J,N} := \mathbf{1}[\text{client } k \text{ is route-}J \text{ ready}], \\ r_{k,J,N} &:= \Pr(M_{k,J,N} = 1). \end{aligned} \tag{60}$$

Lemma 8 (Heterogeneous one-client expansion)**.** *Uniformly over relevant client-route pairs,*

$$r_{k,J,N} = \theta_{k,J,N} (n_{k,N})_{w_J} p_N^{w_J} \{1 + o(1)\}. \tag{61}$$

Proof. Let $q_{k,J,N}$ be the total rare-cell probability. The configuration with exactly each threshold contributes

$$\frac{(n_{k,N})_{w_J}}{\prod_\ell m_{J,\ell}!} \prod_\ell P_{k,N}(\mathcal{C}_{J,\ell})^{m_{J,\ell}} (1 - q_{k,J,N})^{n_{k,N} - w_J}. \tag{62}$$

Uniformly, $n_{k,N} q_{k,J,N} \to 0$, so the last factor is $1 + o(1)$, and configurations with an extra rare observation have relative order $O(n_{k,N} q_{k,J,N}) = o(1)$. The coefficient bounds keep the threshold allocation probability non-degenerate conditional on w_J rare observations; the common-cell contract contributes only a relative $o(1)$. Substitution gives Eq. (61). □

Define $\Lambda_{J,N} := \sum_{k=1}^{K_N} r_{k,J,N}$, $\qquad R_{J,N}^{\mathrm{loc}} := 1 - \prod_{k=1}^{K_N} (1 - r_{k,J,N})$, and the falling-factorial effective route exposure

$$\Theta_{J,N}^{\mathrm{eff}} := \frac{K_N^{w_J - 1}}{N^{w_J}} \sum_{k=1}^{K_N} \theta_{k,J,N} (n_{k,N})_{w_J}. \tag{63}$$

Proposition 5 (Effective-exposure route law)**.** *If* $\max_k r_{k,J,N} \to 0$*, then*

$$R_{J,N}^{\mathrm{loc}} = 1 - \exp\{-\Lambda_{J,N}(1 + o(1))\}, \tag{64}$$

$$\Lambda_{J,N} = \Theta_{J,N}^{\mathrm{eff}} N^{w_J(1-\gamma)} K_N^{1-w_J} \{1 + o(1)\}. \tag{65}$$

If local sizes diverge, then additionally

$$\Theta_{J,N}^{\mathrm{eff}} = \frac{1}{K_N} \sum_{k=1}^{K_N} \theta_{k,J,N} \left(\frac{K_N n_{k,N}}{N} \right)^{w_J} \{1 + o(1)\}. \tag{66}$$

Proof. The logarithmic expansion of $\prod_k (1 - r_{k,J,N})$ has remainder bounded by $\sum_k r_{k,J,N}^2 \le (\max_k r_{k,J,N}) \Lambda_{J,N} = o(\Lambda_{J,N})$, proving Eq. (64). Substituting Lem. 8 into the sum and rearranging gives Eq. (65); replacing falling factorials by powers gives Eq. (66). □

If $\Theta_{J,N}^{\mathrm{eff}} \to \Theta_J \in (0, \infty)$ and $\log K_N = \beta_J \log N + t + o(1)$, then $R_{J,N}^{\mathrm{loc}} \longrightarrow 1 - \exp\{-\Theta_J e^{-(w_J - 1)t}\}$. Thus bounded route-aligned heterogeneity preserves the burden exponent and shifts the critical location by $\Delta_J^{\mathrm{het}} = \frac{\log \Theta_J - \log \theta_J^{\mathrm{hom}}}{w_J - 1}$.

More generally, if

$$\Lambda_{\mathrm{base},N} \to 0, \quad \Lambda_{G,N} \to \lambda_G \in (0, \infty), \quad \Lambda_{M,N} \to \infty, \quad \Lambda_{P,N} \to 0, \tag{67}$$

then conditioning on baseline abstention is negligible by Eq. (53); Prop. 5 gives global readiness $1 - e^{-\lambda_G}$, module readiness one, and competing P-global readiness zero. The complete-repair classification in Eq. (44) excludes any unmodelled shortcut. If pooled route-cell counts diverge in probability in every fixed route cell, pooling makes all routes ready and selects a one-query global repair. Consequently, at threshold q, the local rule selects the module bundle when $1 - e^{-\lambda_G} < q$ and a global singleton when $1 - e^{-\lambda_G} > q$.

G.8 Same-Pooled Decomposition Reversal

Use the common DAG $P \to R \to X, \quad P \to T, \quad X \to T, \quad P \to Z, \quad Z \to U_i, \quad Z \to V_i$. The root P has two common states and four fixed rare states. The frozen companion centre uses the common-state mixture $\omega = 1/2$; all non-root conditional mechanisms and rare-root masses are shared across client types. Its canonical rare-support burdens are $w_{\text{base}} = 4, \quad w_G = 3, \quad w_M = 2, \quad w_P = 4$, and its limiting route coefficients are

$$\theta_G(\omega) = \tfrac{1}{2}(0.7 + 1.3\omega)^2(0.5 + 0.9\omega), \qquad \theta_M(\omega) = (1.8 - 1.2\omega)(1.5 - 0.7\omega). \tag{68}$$

Both are strictly convex on $[0.25, 0.75]$, since

$$\theta''_G(\omega) = \frac{13(351\omega + 191)}{1000} > 0, \qquad \theta''_M(\omega) = \frac{42}{25} > 0. \tag{69}$$

With $\gamma = 0.42$, $\beta_G = 3(1 - 0.42)/2 = 0.87$; at this scale, $\Lambda_{\text{base}}, \Lambda_P \to 0$ and $\Lambda_M \to \infty$.

Lemma 9 (Semantic preservation under affine same-pooled root mixtures)**.** *For all sufficiently large N, the companion construction can be parameterised by a faithful homogeneous centre $P_{\omega_N^\circ, N}$ with $\omega_N^\circ \to 1/2$, and there is a fixed constant $\varepsilon > 0$, independent of N, such that both $P_{\omega_N^\circ - \varepsilon, N}$ and $P_{\omega_N^\circ + \varepsilon, N}$ are strictly positive and faithful to the same DAG. Hence all three populations satisfy the same admissible-query labels. Moreover,*

$$\frac{1}{2} P_{\omega_N^\circ - \varepsilon, N} + \frac{1}{2} P_{\omega_N^\circ + \varepsilon, N} = P_{\omega_N^\circ, N} \tag{70}$$

pointwise on the full joint state space. If $\theta_{G,N}(\omega)$ is the corresponding finite-N global-route coefficient, then uniformly for $\omega \in [0.25, 0.75]$,

$$\theta_{G,N}(\omega) \longrightarrow \theta_G(\omega), \tag{71}$$

and therefore

$$\Theta_G^{\text{hom}} := \theta_G(1/2) < \Theta_G^{\text{split}} := \frac{\theta_G(1/2 - \varepsilon) + \theta_G(1/2 + \varepsilon)}{2}. \tag{72}$$

Proof. Start from the positive frozen companion centre at common-state mixture $1/2$. Let $\bar{\vartheta}_N$ denote its full Markov parameter vector, and let $\bar{\theta}_{J,N}(\omega)$, $J \in \{\text{base}, G, M, P\}$, denote the corresponding finite-N route coefficients when only the common-root mixture is changed to $\omega \in [0.25, 0.75]$. The explicit frozen construction gives

$$\sup_{\omega \in [0.25, 0.75]} \left| \bar{\theta}_{G,N}(\omega) - \theta_G(\omega) \right| + \sup_{\omega \in [0.25, 0.75]} \left| \bar{\theta}_{M,N}(\omega) - \theta_M(\omega) \right| \longrightarrow 0, \tag{73}$$

while the frozen baseline and competing-P coefficients remain uniformly bounded above on the same compact interval, and the explicit θ_M formula is bounded away from zero there. These properties follow directly from the finite frozen route-cell formulas: the rare-cell probabilities are p_N times coefficient functions that converge uniformly on the interval, and the thresholds are fixed.

Choose any deterministic $\delta_N \downarrow 0$. For fixed N, define $\theta^{\vartheta}_{J,N}(\omega)$ by replacing the common-root mixture coordinate of a nearby Markov parameter vector ϑ by ω while leaving all its other coordinates fixed. The finite collection of maps $(\vartheta, \omega) \mapsto \theta^{\vartheta}_{J,N}(\omega)$ is jointly continuous. Compactness of $[0.25, 0.75]$ therefore gives an ambient radius $\rho_N > 0$ such that

$$\|\vartheta - \bar{\vartheta}_N\| < \rho_N \quad \Longrightarrow \quad \sup_{\omega \in [0.25, 0.75]} \left| \theta^{\vartheta}_{J,N}(\omega) - \bar{\theta}_{J,N}(\omega) \right| < \delta_N \quad \forall J \in \{\text{base}, G, M, P\}. \tag{74}$$

Positivity and continuity of the finitely many theorem-relevant rare support cells likewise give a radius $r_N > 0$ within which their relative changes are below δ_N and all required common-cell lower bounds remain strict. Hence

$$\mathcal{O}_N^{\text{comp}} := \left\{ \vartheta : \|\vartheta - \bar{\vartheta}_N\| < \min\{\rho_N, r_N, \delta_N\} \right\} \tag{75}$$

is a non-empty open subset of the full positive Markov parameter space. The unfaithful parameters form a proper algebraic null subset, so $\mathcal{O}_N^{\text{comp}}$ contains a faithful point ϑ_N°. Its distance from the frozen centre is $o(1)$, and Eqs. (73)–(74) imply the uniform global-route convergence asserted in Eq. (71); they also preserve the positive finite baseline, module, and competing-global coefficient

bounds needed for their burden regimes. Let ω_N° be the proportion of the selected total common-root mass placed on the first common state. Then $\omega_N^\circ \to 1/2$. Freeze, for this N, the selected non-root CPTs and all selected rare-root masses, and vary only the division of the remaining common-root mass between the two common states. This defines a strictly positive one-parameter family $P_{\omega,N}$ for $\omega \in [0.25, 0.75]$.

For every full assignment x, the root factor is either independent of ω (if P is in a rare state) or affine in ω (if P is in a common state), while every conditional factor is fixed. Hence

$$P_{\omega,N}(x) = A_N(x) + \omega B_N(x) \tag{76}$$

for suitable coefficients $A_N(x), B_N(x)$. This immediately gives Eq. (70) for every symmetric split about ω_N°.

It remains to choose one non-vanishing split that preserves the CI semantics for the whole sequence. Fix N. Every CI statement that is d-separated in the common DAG holds under every $P_{\omega,N}$ by the global Markov property. Now take any CI statement $I = (X_a, X_b, S)$ over the finite companion variable set that is d-connected. Strict positivity makes conditional independence equivalent to the vanishing of finitely many conditional 2×2 minors. For any such minor, $f_{I,c,N}(\omega) := P_{\omega,N}(x_a, x_b, s)P_{\omega,N}(s) - P_{\omega,N}(x_a, s)P_{\omega,N}(x_b, s)$ is a polynomial in ω of degree at most two by Eq. (76). Because the centre $P_{\omega_N^\circ,N}$ is faithful, for every d-connected I at least one of its minors is non-zero at ω_N°; the corresponding polynomial is therefore not identically zero. Thus the set $B_{I,N}$ of ω values at which I can become accidentally independent is contained in a finite polynomial-root set. The complete CI universe on the finite companion variable set is finite, so $B_N := \bigcup_{I:d_G(I)=0} B_{I,N}$ is finite.

Choose a fixed open interval $(0, \varepsilon_{\max})$ small enough that, because $\omega_N^\circ \to 1/2$, both $\omega_N^\circ \pm \varepsilon$ stay inside $[0.25, 0.75]$ for every sufficiently large N whenever $0 < \varepsilon < \varepsilon_{\max}$. The forbidden symmetric distances $E_{\text{bad}} := \bigcup_{N \geq N_0} \{|\omega - \omega_N^\circ| : \omega \in B_N\}$ form a countable set. Therefore one may fix $\varepsilon \in (0, \varepsilon_{\max}) \setminus E_{\text{bad}}$. For this single ε, both split populations are faithful to the common DAG for every sufficiently large N, which is stronger than the shared admissible-query semantics required by Assump. 1.

The coefficient-control neighbourhood was chosen precisely so that Eq. (71) holds uniformly on $[0.25, 0.75]$. Consequently

$$\begin{aligned} &\frac{\theta_{G,N}(\omega_N^\circ - \varepsilon) + \theta_{G,N}(\omega_N^\circ + \varepsilon)}{2} - \theta_{G,N}(\omega_N^\circ) \\ &\qquad \longrightarrow \frac{\theta_G(1/2 - \varepsilon) + \theta_G(1/2 + \varepsilon)}{2} - \theta_G(1/2) > 0, \end{aligned} \tag{77}$$

where strict positivity of the limit follows from Eq. (69). This proves Eq. (72) and the claimed fixed non-zero decomposition gap. □

Fix any finite critical offset t. To realise the same-pooled comparison with exactly equal integer local sizes, choose any even sequence $K_m \to \infty$ and set $n_m := \left\lfloor e^{-t/\beta_G} K_m^{(1-\beta_G)/\beta_G} \right\rfloor$, $N_m := K_m n_m$. Then $n_m \to \infty$ and $\log K_m - \beta_G \log N_m = t + o(1)$. Relabel this subsequence by N. Let the homogeneous client configuration use the faithful centre type ω_N° at every client, and let the split client configuration use half of its equal-size clients at $\omega_N^\circ - \varepsilon$ and half at $\omega_N^\circ + \varepsilon$. Eq. (70) then gives exactly the same sample-weighted pooled population $\bar{P}_N$, because the two client types receive equal total sample weight. Prop. 5, Eq. (71), and Eq. (77) give effective exposures converging to Θ_G^{hom} and Θ_G^{split}. Hence $1 - \exp\{-\Theta_G^{\text{hom}} e^{-(w_G-1)t}\} < 1 - \exp\{-\Theta_G^{\text{split}} e^{-(w_G-1)t}\}$, which is exactly Eq. (16). The module route tends to readiness one in both decompositions, and pooled evidence sees the same pooled population $\bar{P}_N$ and a one-query global repair in both. This proves the threshold-free separation in part (c).

Numerical companion anchor. The pre-specified finite diagnostic uses the unperturbed symmetric profile

$$\Pi_{\text{hom}} = \delta_{0.5}, \qquad \Pi_{\text{split}} = \tfrac{1}{2}\delta_{0.25} + \tfrac{1}{2}\delta_{0.75}, \tag{78}$$

for which $\theta_G(0.5) = 0.8656875$, $\frac{\theta_G(0.25)+\theta_G(0.75)}{2} = 1.014578125$. These values illustrate the affine-pooled/nonlinear-exposure mechanism; Lem. 9 supplies the faithful split used in the theorem.

For $q \in (0,1)$ and $a \in \{\mathrm{hom}, \mathrm{split}\}$, define

$$t_a(q) := \frac{\log \Theta_G^a - \log[-\log(1-q)]}{w_G - 1}. \tag{79}$$

Then every $t \in (t_{\mathrm{hom}}(q), t_{\mathrm{split}}(q))$ satisfies

$$r_G^{\mathrm{hom}}(t) < q < r_G^{\mathrm{split}}(t), \tag{80}$$

and the interval width is

$$t_{\mathrm{split}}(q) - t_{\mathrm{hom}}(q) = \frac{\log(\Theta_G^{\mathrm{split}}/\Theta_G^{\mathrm{hom}})}{w_G - 1} > 0. \tag{81}$$

For the threshold consequence, Eq. (79) is the unique solution of $r_G^a(t) = q$. The strict exposure inequality gives $t_{\mathrm{hom}}(q) < t_{\mathrm{split}}(q)$, so every interior offset satisfies Eq. (80). Subtracting the two locations cancels the q term and gives Eq. (81). Thus every $q \in (0,1)$ has a non-empty crossing interval, while no numerical threshold is part of the theorem scope.

Phase Regimes and Scope. For $K_N = \lfloor N^\beta \rfloor$, the homogeneous construction has four coarse regimes: the baseline is ready for $\beta < \beta_{\mathrm{base}}$; both repair classes are ready for $\beta_{\mathrm{base}} < \beta < \beta_G$; global singletons fail while the module repair remains ready for $\beta_G < \beta < \beta_M$; and neither class is locally ready for $\beta > \beta_M$. The critical theorem resolves the second boundary before global readiness reaches zero. The construction is asymptotic, finite-discrete, local non-pooling, and correct-or-abstain, with client-specific populations sharing the admissible CI labels.

H Proof and Natural Instance of Support-Compatible Preservation

H.1 Support-Compatible Preservation

Proof of Cor. 3. For the auxiliary argument, define the surfaced target-disagreement coverage of an arbitrary query set $\mathcal{Q}$ at support state z by $\mathrm{SurfCov}_T(\mathcal{Q}; z) := \bigcup_{I \in \mathcal{Q}:\, \phi_I(z)=1} \mathrm{Cov}_T(I)$. For exact repairs $\mathcal{A}, \mathcal{B} \in \mathfrak{C}_T$, the standard minimal-path order is

$$\mathcal{A} \succeq_{\mathrm{path}} \mathcal{B} \iff \forall R_B \in \mathfrak{P}(\mathcal{B}),\ \exists R_A \in \mathfrak{P}(\mathcal{A}) \text{ with } R_A \subseteq R_B. \tag{82}$$

For monotone readiness functions this abstract order is equivalent to pointwise implication over all support states and therefore to probability ordering under every support-state law. The converse sharpness direction is stated explicitly below. These are standard coherent-system facts; the corollary below requires only the declared interface-induced states and laws.

Assume $J \succeq_{\mathrm{cell}} I$ and that query I is ready on a finite client table or on a coalition table formed at bounded aggregation depth. Every I-cell v then has count at least $m_{I,v}$. A J-cell u is the disjoint union of its constituent I-cells, so

$$N(u) = \sum_{v \subseteq u} N(v) \geq \sum_{v \subseteq u} m_{I,v} \geq m_{J,u}. \tag{83}$$

Thus J is ready on the same table or coalition table. Taking one client gives the local implication, and summing all client counts gives the pooled implication. The deterministic argument also applies to non-identically distributed clients.

Now let z be an interface-induced support state. Take $\zeta \in \mathrm{SurfCov}_T(\mathcal{B}; z)$. Some $I \in \mathcal{B}$ is ready at z and covers ζ. Applying Eq. (83) to the client or query-specific coalition that makes I ready shows that $\varphi(I)$ is ready on the same interface, while the corollary's coverage assumption gives $\zeta \in \mathrm{Cov}_T(I) \subseteq \mathrm{Cov}_T(\varphi(I))$. Hence $\mathrm{SurfCov}_T(\mathcal{A}; z) \supseteq \mathrm{SurfCov}_T(\mathcal{B}; z)$ on every interface-induced state.

If $\phi_{\mathcal{B}}(z) = 1$, every query in $\mathcal{B}$ is ready and therefore $\mathrm{SurfCov}_T(\mathcal{B}; z) = \mathrm{Amb}_T$. The auxiliary inclusion gives $\mathrm{SurfCov}_T(\mathcal{A}; z) = \mathrm{Amb}_T$. Since $\mathcal{A}$ is inclusion-minimal exact, no strict subset of $\mathcal{A}$ covers Amb_T; hence every query in $\mathcal{A}$ must be ready and $\phi_{\mathcal{A}}(z) = 1$. Thus $\phi_{\mathcal{A}}(z) \geq \phi_{\mathcal{B}}(z)$ for every interface-induced state. Integrating under any interface-induced support-state law gives the corollary's readiness inequality; if $C(\mathcal{A}) \leq C(\mathcal{B})$, no readiness threshold can make minimum-cost selection uniquely prefer $\mathcal{B}$. □

Proposition 6 (Abstract path-level sharpness)**.** *For exact repairs $\mathcal{A}, \mathcal{B} \in \mathfrak{C}_T$ in the abstract support system, if $\mathcal{A} \not\succeq_{\mathrm{path}} \mathcal{B}$, then there exists a full-support law μ on support states such that $\rho_\mu(\mathcal{B}) > \rho_\mu(\mathcal{A})$. Consequently, any threshold strictly between the two readiness probabilities yields a strict reversal.*

Proof. Choose $R_B \in \mathfrak{P}(\mathcal{B})$ containing no member of $\mathfrak{P}(\mathcal{A})$. At state $z = R_B$, $\mathcal{B}$ is ready and $\mathcal{A}$ is not. Let U assign positive probability to every support state and put $\mu_\varepsilon = (1-\varepsilon)\delta_z + \varepsilon U$. For sufficiently small $\varepsilon > 0$, $\rho_{\mu_\varepsilon}(\mathcal{B}) > \rho_{\mu_\varepsilon}(\mathcal{A})$; taking $\mu = \mu_\varepsilon$ proves the claim and any threshold strictly between the two values gives the full-support reversal. □

This coherent-system sharpness statement is not needed to enlarge the interface-induced scope of Cor. 3.

H.2 Why Each Preservation Condition Is Needed

For the abstract support system, path-level implication is the necessary-and-sufficient pointwise boundary. The following proposition concerns the three primitive sufficient clauses used by Cor. 3 to certify preservation on the declared evidence interfaces before observing client data.

Proposition 7 (Clause-wise irredundancy of primitive preservation)**.** *Coverage domination, conditioning-cell domination, and threshold domination are each separately non-removable from a uniform causal guarantee. For each clause there is a finite target-specific DAG task, two singleton query sets for which the other two clauses hold, and a finite realised client table on which the proposed operational order fails. All graphs admit strictly positive faithful distributions, and the displayed tables have positive probability.*

Proof. Coverage. Let $G_0 : X \to T \to Y$ and $G_1 : X \to T \leftarrow Y$, with the orientation of the edge between T and Y as target summary. The unconditional query $X \perp\!\!\!\perp Y$ distinguishes the chain from the collider, whereas $X \perp\!\!\!\perp T$ has the same dependence label in both graphs. Their support partitions and thresholds can be identical and both queries can be ready on every non-empty table, but the required surfaced-coverage inclusion fails if the latter query is proposed as the dominating query set.

Conditioning cells. Add an isolated binary Z to the same graphs. The queries $X \perp\!\!\!\perp Y$ and $X \perp\!\!\!\perp Y \mid Z$ have identical residual coverage and equal scalar threshold one. On a one-row client table with $Z = 0$, the unconditional query is ready while the refined query abstains because the $Z = 1$ cell is empty. Thus coverage and threshold clauses hold while conditioning-cell domination is essential.

Threshold. Let G_0 contain $X_1 \to T \leftarrow X_2$ and $T \to Y_1, T \to Y_2$, while G_1 reverses the two edges between T and Y_j. The unconditional queries $(X_1, Y_1, \varnothing)$ and $(X_2, Y_2, \varnothing)$ have identical singleton residual coverage and identical one-cell partitions. Give the first threshold two and the second threshold one. A one-row client table makes only the second ready, so threshold domination is essential.

Generic strictly positive discrete Bayesian-network parameters are faithful outside a proper algebraic exceptional set, and strict positivity makes every complete finite assignment possible. Hence all three realised support witnesses have legal causal realisations. The three constructions establish the stated clause-wise irredundancy. □

H.3 A Natural ALARM Instance of Support-Compatible Coverage Preservation

Use the parent-closed scope

$$\mathcal{V}_T = \{\texttt{HYPOVOLEMIA}, \texttt{LVFAILURE}, \texttt{LVEDVOLUME}, \texttt{CVP}, \texttt{PCWP}\}, \qquad T = \texttt{LVEDVOLUME}, \tag{84}$$

with the complete order-zero baseline. Let $\mathcal{W}_T$ be the 229 residual target-wrong labelled DAGs and define the reference-relative coverage $\mathrm{Cov}^\star_T(I) := \{G \in \mathcal{W}_T : d_G(I) \neq d_{G^\star}(I)\}$. Consider the singleton query sets

$$\begin{aligned} \mathcal{A}_{\mathrm{AL}} &= \{\texttt{LVFAILURE} \perp\!\!\!\perp \texttt{CVP} \mid \texttt{LVEDVOLUME}\}, \\ \mathcal{B}_{\mathrm{AL}} &= \{\texttt{LVFAILURE} \perp\!\!\!\perp \texttt{CVP} \mid \texttt{HYPOVOLEMIA}, \texttt{LVEDVOLUME}\}. \end{aligned} \tag{85}$$

Exact enumeration over all 29,281 labelled five-node DAGs gives

$$|\mathrm{Cov}^\star_T(\mathcal{A}_{\mathrm{AL}})| = 219, \qquad |\mathrm{Cov}^\star_T(\mathcal{B}_{\mathrm{AL}})| = 207, \qquad \mathrm{Cov}^\star_T(\mathcal{B}_{\mathrm{AL}}) \subsetneq \mathrm{Cov}^\star_T(\mathcal{A}_{\mathrm{AL}}). \tag{86}$$

For the reference-relative family above, define $\mathrm{SurfCov}^\star_T(\mathcal{Q}; z) := \bigcup_{I \in \mathcal{Q}:\, \phi_I(z)=1} \mathrm{Cov}^\star_T(I)$. The partition induced by LVEDVOLUME is a strict coarsening of the partition induced by (HYPOVOLEMIA, LVEDVOLUME). Under $m_A = m_B = 5$, readiness of $\mathcal{B}_{\mathrm{AL}}$ therefore implies readiness of $\mathcal{A}_{\mathrm{AL}}$ on every finite client table and after pooling. Together with the strict reference-relative coverage inclusion in Eq. (86), the same surfaced-coverage mapping argument used to prove Cor. 3 gives

$$\mathrm{SurfCov}^\star_T(\mathcal{A}_{\mathrm{AL}}; z) \supseteq \mathrm{SurfCov}^\star_T(\mathcal{B}_{\mathrm{AL}}; z) \qquad \text{for every interface-induced support state } z. \tag{87}$$

Thus every set-monotone value of surfaced reference-relative target-disagreement coverage is pointwise no smaller for $\mathcal{A}_{\mathrm{AL}}$.

The two singleton query sets instantiate the reference-relative coverage-preservation clause. Their $219/229$ and $207/229$ values refer to $\mathrm{Cov}^\star_T$, while the exact-repair specialization of Cor. 3 applies to query families covering the full residual target-disagreement class. The finite grid below shows that this strict coverage-and-support ordering remains non-degenerate across high, intermediate, and low refined-query readiness regimes.

Table 2: Finite ALARM support-compatible coverage anchor. Reference-relative coverage fractions are $219/229 = 0.9563$ and $207/229 = 0.9039$.

N	K	$\rho(A)$	$\rho(B)$	$V_{\mathrm{soft}}(A)$	$V_{\mathrm{soft}}(B)$
500	2	1.0000	0.1237	0.9563	0.1118
1000	2	1.0000	0.8558	0.9563	0.7735
1000	4	1.0000	0.2320	0.9563	0.2097
2000	2	1.0000	0.9998	0.9563	0.9038
2000	4	1.0000	0.9792	0.9563	0.8851
2000	8	1.0000	0.4102	0.9563	0.3708
5000	8	1.0000	1.0000	0.9563	0.9039
5000	16	1.0000	0.9517	0.9563	0.8602
5000	32	1.0000	0.0877	0.9563	0.0793

The complete order-one/order-two singleton enumeration found ten strict candidates in this ALARM scope, none in the pre-specified parent-closed CHILD scope, and none in the non-parent-closed INSURANCE scope. The calculation provides a natural non-degenerate preservation anchor for the evaluated ALARM scope.

I Extensions: Weighted Costs, Residual Tasks, and Pooled Ordering

Corollary 7 (Weighted distortion and residual-system transfer)**.** *Assign every query a positive cost. In the Thm. 1 family, let $\mathcal{Q}_G$ be the global witness class G1–G4 and let $\mathcal{Q}_{M_i}$ be the module-i witness class from Lem. 1. Define*

$$\begin{aligned} c_G^{\mathrm{str}} &:= \min_{I \in \mathcal{Q}_G} c_I, & c_i^{\mathrm{str}} &:= \min_{I \in \mathcal{Q}_{M_i}} c_I, \\ c_i^{\mathrm{loc}} &:= \min_{I \in \mathcal{Q}_{M_i}:\, \mathrm{Ready}_I^{\mathrm{loc}}=1} c_I, & c_G^{\mathrm{pool}} &:= \min_{I \in \mathcal{Q}_G:\, \mathrm{Ready}_I^{\mathrm{pool}}=1} c_I, \\ c_i^{\mathrm{pool}} &:= \min_{I \in \mathcal{Q}_{M_i}:\, \mathrm{Ready}_I^{\mathrm{pool}}=1} c_I. \end{aligned} \tag{88}$$

The local and pooled minima are well defined on the hard support configuration in Thm. 1: every module has a locally ready canonical witness, and pooling exposes at least one global witness. Then

$$C_T^{\mathrm{str}} = \min\left\{ c_G^{\mathrm{str}}, \sum_{i=1}^h c_i^{\mathrm{str}} \right\}, \qquad C_{T,\Gamma}^{\mathrm{loc,ad}} = \sum_{i=1}^h c_i^{\mathrm{loc}}, \qquad C_{T,\Gamma}^{\mathrm{pool,ad}} = \min\left\{ c_G^{\mathrm{pool}}, \sum_{i=1}^h c_i^{\mathrm{pool}} \right\}. \tag{89}$$

In particular, if constants $0 < \underline{c} \le \bar{c} < \infty$ independent of h satisfy $c_i^{\mathrm{loc}} \ge \underline{c}$ for every module and $c_G^{\mathrm{str}}, c_G^{\mathrm{pool}} \le \bar{c}$, then

$$\frac{C_{T,\Gamma}^{\mathrm{loc,ad}}}{C_T^{\mathrm{str}}} \ge \frac{h\underline{c}}{\bar{c}}, \qquad \frac{C_{T,\Gamma}^{\mathrm{loc,ad}}}{C_{T,\Gamma}^{\mathrm{pool,ad}}} \ge \frac{h\underline{c}}{\bar{c}}, \tag{90}$$

so both distortions remain unbounded under uniformly non-degenerate query costs. Any supplied baseline transcript whose residual disagreement incidence, query costs, query support-path families, and support-state law are isomorphic to this residual system inherits the same repair families, costs, readiness probabilities, and cost–readiness orderings. If the isomorphism additionally maps the designated reference graph to its counterpart and preserves the corresponding reference-relative witness hypergraph, then it also preserves the reference-transcript lower bounds.

Proof. Cor. 4 gives the complete inclusion-minimal exact-repair classification: a repair is either a global singleton or contains exactly one module witness from every module. Minimising cost over these two structural classes gives the first equality in Eq. (89). In the realised hard local-support configuration every global witness abstains. Prop. 2 then forces the reference transcript to contain at least one locally visible witness from every module, and querying one cheapest such witness per module attains $\sum_i c_i^{\mathrm{loc}}$, giving the second equality. After pooling, global witnesses and the canonical module witnesses are available. The same reference-witness argument says that an exact pooled protocol must either query a pooled-ready global witness or hit every module, while either corresponding static route attains the bound; this gives the third equality. The inequalities in Eq. (90) follow from $C_T^{\mathrm{str}} \le c_G^{\mathrm{str}}$, $C_{T,\Gamma}^{\mathrm{pool,ad}} \le c_G^{\mathrm{pool}}$, and $C_{T,\Gamma}^{\mathrm{loc,ad}} = \sum_i c_i^{\mathrm{loc}} \ge h\underline{c}$. The repair-family, cost, readiness, and cost–readiness-ordering claims under residual-system transfer follow from Prop. 1 and support-system isomorphism. When the isomorphism also preserves the designated reference and its reference-relative witness hypergraph, Prop. 2 transfers the corresponding reference-transcript bounds. □

For a realised client-data configuration Γ, define the minimum cost among pooled-ready inclusion-minimal exact repairs by

$$C_{T,\Gamma}^{\mathrm{pool}} := \min\left\{C(\mathcal{Q}) : \mathcal{Q} \in \mathfrak{C}_T,\ \mathrm{Ready}_I^{\mathrm{pool}} = 1 \text{ for every } I \in \mathcal{Q}\right\}, \tag{91}$$

with the value $+\infty$ when no pooled-ready exact repair exists. This is a static repair cost and is distinct from the pooled adaptive protocol cost $C_{T,\Gamma}^{\mathrm{pool,ad}}$.

Proposition 8 (Pooled recovery of structural ordering)**.** *If some structural minimum-cost exact repair is pooled-ready, then $C_{T,\Gamma}^{\mathrm{pool}} = C_T^{\mathrm{str}}$. If every admissible query is pooled-ready, every exact repair has readiness one after pooling and the complete pooled cost ordering equals the structural cost ordering.*

Proof. Restricting $\mathfrak{C}_T$ to pooled-ready repairs cannot produce a repair cheaper than the unrestricted structural minimum, while pooled readiness of a structural minimiser attains that minimum. If every admissible query is pooled-ready, the restriction leaves all of $\mathfrak{C}_T$ available, so the complete structural cost ordering is preserved after pooling. □

J A Route-Grid Causal Witness for Repair Ordering and Target Decisions

J.1 Causal Construction

A second causal family isolates repair-order reversal after matching the canonical coarsest representatives in query order, threshold, and limiting query-wise marginal readiness. The reversal is generated by repair-level support-route composition. Fix $h \ge 3$ and let $\mathcal{V}_h^{\mathrm{grid}} = \{X, Y, T, Z\} \cup \{U_i, V_i : i \in [h]\}$. The reference DAG $G_h^{\mathrm{grid},\star}$ contains $X \to Y, \quad X \to T, \quad Y \to T, \quad Z \to U_i, \quad Z \to V_i \quad (i \in [h])$. The target summary for this auxiliary witness is the binary target-parent property $\mathcal{L}_T(G) := \mathbf{1}\{X \to T \in G \text{ and } Y \to T \in G\}$. Thus the task is to recover whether the

designated two-parent target relation is intact; the framework permits such a target-specific summary map. For each $i \in [h]$, define

$$\begin{aligned} G_{A_i} &= G_h^{\mathrm{grid},\star} - \{X \to T\} + \{U_i \to V_i\}, \\ G_{B_i} &= G_h^{\mathrm{grid},\star} - \{Y \to T\} + \{U_i \to V_i\}. \end{aligned} \tag{92}$$

Then $\mathcal{L}_T(G_h^{\mathrm{grid},\star}) = 1$ and every G_{A_i}, G_{B_i} has target summary 0. Take $\mathcal{G}_h^{\mathrm{grid}} = \{G_h^{\mathrm{grid},\star}\} \cup \{G_{A_i}, G_{B_i} : i \in [h]\}$ and choose a baseline containing no non-zero discrepancy query (the baseline may be empty for this residual witness). The target-disagreement set is therefore exactly

$$\mathrm{Amb}_T^{\mathrm{grid}} = \{\zeta_{A_i}, \zeta_{B_i} : i \in [h]\}, \qquad \zeta_{A_i} := \{G_h^{\mathrm{grid},\star}, G_{A_i}\}, \quad \zeta_{B_i} := \{G_h^{\mathrm{grid},\star}, G_{B_i}\}. \tag{93}$$

The residual disagreement set $\mathrm{Amb}_T^{\mathrm{grid}}$ contains pairs with different target summaries; target-wrong candidates share the same target summary and are therefore excluded. The result concerns the complete residual observational CI query universe on this candidate class.

For this stochastic construction, the designated reference world is $G_h^{\mathrm{grid},\star}$. When the client-configuration sequence is clear, write $\rho^{\xi,\mathrm{grid},\star}$ for $\rho^{\xi}_{\Gamma,G_h^{\mathrm{grid},\star}}$.

Define two row queries and h column queries: $a = (X, T, \{Y\}), \qquad b = (Y, T, \{X\}), \qquad q_i = (U_i, V_i, \{Z\}), \quad i \in [h]$.

J.2 Complete Disagreement Structure

Lemma 10 (Exact $2 \times h$ discrepancy grid)**.** *Across the complete all-order observational CI universe, every non-zero target-discrepancy mask is exactly one of*

$$\{\zeta_{A_i} : i \in [h]\}, \qquad \{\zeta_{B_i} : i \in [h]\}, \qquad \{\zeta_{A_i}, \zeta_{B_i}\} \quad (i \in [h]). \tag{94}$$

No query covers all $2h$ target disagreements, and no other row subset, mixed subset, interval, or cross-module subset occurs. Moreover, every query carrying the first row mask is exactly a conditioning refinement of a, every query carrying the second row mask is exactly a conditioning refinement of b, and every query carrying column mask i is exactly a conditioning refinement of q_i: the endpoints are unchanged and the conditioning set is the representative conditioning set plus an arbitrary subset of variables outside those endpoints.

Proof. The target core $\{X, Y, T\}$ and the module component $\{Z, U_i, V_i : i \in [h]\}$ are disconnected. In the complete three-node target core, deleting $X \to T$ changes exactly the relation $X \not\perp\!\!\!\perp T \mid Y$ to $X \perp\!\!\!\perp T \mid Y$, whereas deleting $Y \to T$ changes exactly $Y \not\perp\!\!\!\perp T \mid X$ to $Y \perp\!\!\!\perp T \mid X$. No core CI label changes under both deletions. Conditioning additionally on variables in the disconnected module component only refines the same row label.

In module i, adding $U_i \to V_i$ changes the fork-blocking relation $U_i \perp\!\!\!\perp V_i \mid Z$ and its refinements by variables outside that module. The change occurs in G_{A_i} and G_{B_i} and in no candidate with another module index. Cross-component endpoint queries remain d-separated in every graph. Thus the only non-zero masks are the two rows and the h columns in Eq. (94). □

The exhaustive enumeration covers the complete all-order universe for $h = 2, 3, 4$:

Table 3: Complete-query enumeration of the route-grid family. The observed masks are exactly two rows plus h columns.

h	variables	target-wrong candidates	complete CI queries	non-zero queries
2	8	4	1,792	128
3	10	6	11,520	640
4	12	8	67,584	3,072

J.3 Exact Repair Classes

Let $\mathcal{A}_h^{\mathrm{row}} = \{a, b\}, \qquad \mathcal{B}_h^{\mathrm{col}} = \{q_1, \ldots, q_h\}$ be canonical representatives. Because the complete all-order universe contains conditioning refinements with the same coverage masks, the exact-repair statement is naturally a statement about two *classes*, not two literal query sets. Define the coverage-equivalence classes

$$\begin{aligned} \mathcal{E}_A &:= \{I \in \mathcal{Q}_T : \mathrm{Cov}_T(I) = \{\zeta_{A_i} : i \in [h]\}\}, \\ \mathcal{E}_B &:= \{I \in \mathcal{Q}_T : \mathrm{Cov}_T(I) = \{\zeta_{B_i} : i \in [h]\}\}, \\ \mathcal{E}_i &:= \{I \in \mathcal{Q}_T : \mathrm{Cov}_T(I) = \{\zeta_{A_i}, \zeta_{B_i}\}\}, \qquad i \in [h]. \end{aligned} \tag{95}$$

The corresponding row- and column-repair classes are

$$\begin{aligned} \mathfrak{A}_h^{\mathrm{row}} &:= \big\{\{I_A, I_B\} : I_A \in \mathcal{E}_A,\ I_B \in \mathcal{E}_B\big\}, \\ \mathfrak{B}_h^{\mathrm{col}} &:= \big\{\{I_i : i \in [h]\} : I_i \in \mathcal{E}_i\ \forall i \in [h]\big\}. \end{aligned} \tag{96}$$

Lemma 11 (Exactly two inclusion-minimal repair classes)**.** *Across the complete all-order observational CI universe, $\mathfrak{C}_{T,h}^{\mathrm{grid}} = \mathfrak{A}_h^{\mathrm{row}} \cup \mathfrak{B}_h^{\mathrm{col}}$. Under unit query costs, every member of $\mathfrak{A}_h^{\mathrm{row}}$ has cost 2 and every member of $\mathfrak{B}_h^{\mathrm{col}}$ has cost h. Hence for $h \geq 3$ the structural minimum cost is 2, attained by the canonical row repair $\mathcal{A}_h^{\mathrm{row}}$ and its coverage-equivalent refinements.*

Proof. Lem. 10 says that every non-zero query belongs to exactly one of the two row coverage classes or one of the h column coverage classes. A minimal cover that uses both row masks covers all target disagreements and, because duplicate queries from the same coverage class are removable, contains exactly one query from each row class. If either row mask is absent, every target disagreement in that row must be covered by its own column mask, forcing exactly one query from every column class. A set containing one row and all columns strictly contains an all-column repair, while a set containing both rows and any column strictly contains a two-row repair. Thus the inclusion-minimal repairs are exactly the two classes in Eq. (96). Unit costs give the stated costs. □

J.4 A Faithful Realisation of the Matched-Marginal Centre

Unlike Thm. 3, neither canonical repair is a global singleton: the structurally cheaper row class requires two distinct evidence routes, whereas the canonical member of the structurally more expensive column class uses h queries sharing one route. Let all variables be binary and set $p_N = a_0 N^{-\gamma}, \qquad \frac{1}{m} < \gamma < 1, \qquad m \geq 2$, with $a_0 > 0$. First define a positive Markov *matched-marginal centre* $\bar{P}_{h,N}$ by $\bar{P}_{h,N}(X=1) = \bar{P}_{h,N}(Y=1) = \bar{P}_{h,N}(Z=1) = p_N$, and, for a fixed $\kappa \in (0,1)$,

$$\bar{P}_{h,N}(Y=1 \mid X=1) = \kappa p_N, \qquad \bar{P}_{h,N}(Y=1 \mid X=0) = \frac{p_N - \kappa p_N^2}{1 - p_N}. \tag{97}$$

Thus $\bar{P}_{h,N}(X=1, Y=1) = \kappa p_N^2$ while the X- and Y-marginals are exactly p_N. Complete the centre with any strictly positive CPT $\bar{P}(T \mid X, Y)$ having non-zero direct effects of both parents and, independently across modules conditional on Z, strictly positive mechanisms $\bar{P}(U_i \mid Z)\bar{P}(V_i \mid Z)$. This defines a strictly positive distribution Markov to $G_h^{\mathrm{grid},\star}$; the following ambient-space realisation step establishes faithfulness while preserving the exact-equality slice.

Lemma 12 (Faithful realisation of the matched-marginal centre)**.** *For every fixed $h \geq 3$, the matched-marginal centres above admit a sequence of strictly positive distributions $P_{h,N}$ faithful to $G_h^{\mathrm{grid},\star}$ such that*

$$\begin{aligned} P_{h,N}(X=1) &= p_N\{1 + o(1)\}, & P_{h,N}(Y=1) &= p_N\{1 + o(1)\}, \\ P_{h,N}(Z=1) &= p_N\{1 + o(1)\}, & P_{h,N}(X=1, Y=1) &= \kappa p_N^2\{1 + o(1)\}. \end{aligned} \tag{98}$$

Hence the three canonical support routes have the same leading rare-state probability, while the overlap between X and Y remains of order p_N^2.

Proof. Fix h and N, and let Θ_h^+ be the finite-dimensional strictly positive Markov parameter space of $G_h^{\mathrm{grid},\star}$. Let $\bar{\vartheta}_{h,N} \in \Theta_h^+$ denote the centre above. Choose any deterministic $\delta_N \downarrow 0$. The four

maps $\vartheta \mapsto P_\vartheta(X=1), P_\vartheta(Y=1), P_\vartheta(Z=1), P_\vartheta(X=1,Y=1)$ are continuous and strictly positive at $\bar{\vartheta}_{h,N}$. Therefore the set of $\vartheta \in \Theta_h^+$ for which each of the first three probabilities is within relative error δ_N of p_N and the fourth is within relative error δ_N of κp_N^2 is a non-empty open neighbourhood of $\bar{\vartheta}_{h,N}$ in the *ambient* positive Markov parameter space. For a fixed finite discrete DAG, the unfaithful positive Markov parameters form a proper algebraic null subset of this ambient space. Hence the neighbourhood contains a faithful parameter $\vartheta_{h,N}$. Taking $P_{h,N} = P_{\vartheta_{h,N}}$ gives Eq. (98). No exact-equality constraint is imposed on the selected faithful point; the exact equal marginals serve only as the centre whose leading rare-state laws are preserved. □

An explicit numerical witness for $h = 2$ at three rare probabilities has complete-query conditional-mutual-information statuses matching the graph-implied labels; it is a finite diagnostic and is not used in the existence proof above.

J.5 Refinement-Coherent Support Contract

For this witness, every admissible query uses the ordinary all-strata conditioning-support contract: each state of its conditioning set is a required cell with the same fixed threshold m. For the canonical queries a, b, and q_i, the common conditioning stratum has probability $1 - p_N\{1 + o(1)\}$ and is ready with probability $1 - o(K_N^{-1})$ in the critical regime below; the rare stratum is therefore the unique asymptotically critical cell. The canonical queries have conditioning order one and the same threshold.

The refinement description in Lem. 10 now rules out an operational refinement shortcut as well as a structural one. If $I_A \in \mathcal{E}_A$ is obtained from a by adding conditioning variables, its conditioning partition refines that of a and uses the same per-cell threshold. Hence readiness of I_A implies readiness of a on every realised client table and after pooling. The same implication holds from every $I_B \in \mathcal{E}_B$ to b and from every $I_i \in \mathcal{E}_i$ to q_i. Equivalently, the canonical representative is support-easiest within each coverage class. Consequently every row-class repair has readiness at most that of $\mathcal{A}_h^{\mathrm{row}}$, while $\mathcal{B}_h^{\mathrm{col}}$ is an available canonical member of the column class.

J.6 Critical Route Reversal

Work along an integer sequence with equal client size $n_N = N/K_N$, $n_N \to \infty$, and $K_N \sim \chi N^\beta$, $\beta = \frac{m(1-\gamma)}{m-1} < 1$. Such equal-size sequences exist: for integers $n \to \infty$, take $K = \lfloor c^{1/(1-\beta)} n^{\beta/(1-\beta)} \rfloor$ and $N = Kn$. Then $n_N p_N \to 0$. Let $p_{R,N}$ denote the rare-state probability of route $R \in \{X, Y, Z\}$. By Lem. 12, $p_{R,N} = p_N\{1 + o(1)\}$ for every route.

Let $E_{R,k}$ denote full canonical readiness at client k, including both conditioning strata, and let $\widetilde{E}_{R,k}$ be the event that the rare stratum alone contains at least m observations. Since the common stratum has probability $1 - p_{R,N} = 1 - o(1)$ and $n_N \to \infty$,

$$K_N \Pr(E_{R,k} \triangle \widetilde{E}_{R,k}) \to 0. \tag{99}$$

For example, the difference is bounded by the probability that a $\mathrm{Bin}(n_N, 1 - p_{R,N})$ count is below m, which is $O(n_N^{m-1} p_{R,N}^{n_N - m + 1})$ and hence smaller than any polynomial in N^{-1}. Therefore the one-client readiness probability satisfies $\pi_{R,N} := \Pr(E_{R,k}) = \frac{(n_N p_N)^m}{m!}\{1 + o(1)\} =: \pi_N\{1 + o(1)\}$. Given any desired $r \in (0,1)$, set $\lambda = -\log(1-r)$ and choose $\chi = \left(\frac{a_0^m}{m!\lambda}\right)^{1/(m-1)}$, so that $K_N \pi_N \to \lambda$. Every canonical representative query then has system-level marginal readiness converging to $r = 1 - e^{-\lambda}$.

Lemma 13 (Asymptotic separation of the two canonical row routes)**.** *Let $E_{X,k}$ and $E_{Y,k}$ denote client-k readiness for the canonical X- and Y-conditioned row routes. Then $K_N \Pr(E_{X,k} \cap E_{Y,k}) \to 0$, and the numbers of X-ready and Y-ready clients converge jointly to independent Poisson variables with mean λ. Consequently, $\rho^{\mathrm{loc,grid},\star}(\mathcal{A}_h^{\mathrm{row}}) \to r^2$.*

Proof. Write $p_{11,N} = P_{h,N}(X=1, Y=1)$, $p_{10,N} = P_{h,N}(X=1, Y=0)$, and $p_{01,N} = P_{h,N}(X=0, Y=1)$. Lem. 12 gives $p_{11,N} = O(p_N^2)$, $p_{10,N} = p_N\{1 + o(1)\}$, $p_{01,N} = p_N\{1 + o(1)\}$. First ignore the asymptotically negligible common-stratum constraints in Eq. (99).

Expanding the joint rare-ready event under the multinomial counts of the $(1,1)$, $(1,0)$, and $(0,1)$ cells, the leading contribution uses exactly m $(1,0)$ observations and m $(0,1)$ observations and no shared $(1,1)$ observation; it has order $n_N^{2m}p_N^{2m}$. A configuration using $s \geq 1$ shared $(1,1)$ observations needs only $2m-s$ rare observations but has probability order $n_N^{2m-s}p_N^{2m} = o(n_N^{2m}p_N^{2m})$, because $n_N \to \infty$. Configurations with additional rare observations are lower order because $n_N p_N \to 0$. Hence $\Pr(\widetilde{E}_{X,k} \cap \widetilde{E}_{Y,k}) = \pi_{X,N}\pi_{Y,N}\{1+o(1)\}$. The bound in Eq. (99) transfers the same asymptotic to the full readiness events. Since $K_N\pi_N \to \lambda$ and $K_N \to \infty$, $K_N\pi_N^2 \to 0$. The multivariate law of small numbers then gives independent Poisson limits with common mean λ, and the probability that both canonical row-route types occur converges to $(1-e^{-\lambda})^2 = r^2$. □

All canonical column queries have exactly the same conditioning partition $\{Z=0, Z=1\}$ and the same threshold, so they share one readiness event exactly. Hence $\rho^{\mathrm{loc,grid},\star}(\mathcal{B}_h^{\mathrm{col}}) \to r$. Because $Np_{R,N} \to \infty$ for each $R \in \{X,Y,Z\}$ and the common strata are abundant, the pooled data make every canonical query ready with probability tending to one; therefore both canonical repairs have pooled readiness tending to one.

Proposition 9 (Matched-marginal route-diversity reversal)**.** *For every $h \geq 3$ and $r \in (0,1)$, the family above has exactly two combinatorial classes of inclusion-minimal exact repairs. Every row-class repair has unit-cost total 2, every column-class repair has unit-cost total h, and the canonical representatives have conditioning order one, equal per-cell threshold, and the same limiting local marginal readiness r. Moreover,*

$$\rho^{\mathrm{loc,grid},\star}(\mathcal{A}_h^{\mathrm{row}}) \to r^2, \qquad \rho^{\mathrm{loc,grid},\star}(\mathcal{B}_h^{\mathrm{col}}) \to r, \qquad \rho^{\mathrm{pool,grid},\star}(\mathcal{A}_h^{\mathrm{row}}), \rho^{\mathrm{pool,grid},\star}(\mathcal{B}_h^{\mathrm{col}}) \to 1. \tag{100}$$

For every readiness threshold $q \in (r^2, r)$, no cost-2 row-class repair qualifies locally because every such repair is support-dominated by the canonical row repair, whereas the canonical cost-h column repair qualifies. Any exact repair contains an inclusion-minimal exact subset; because repair readiness is conjunctive, a qualifying exact repair of cost below h would therefore contain a qualifying row-class subset, which is impossible. Hence the minimum local cost among repairs with readiness at least q is h. After pooling, the canonical cost-2 row repair qualifies, so the pooled minimum is 2.

Thus the reversal survives the complete all-order query universe: the cheap repair must combine two distinct critical routes, whereas the expensive class contains a canonical repair that reuses one shared route.

J.7 All-World Support Channel and Target-Decision Order

Thm. 4 additionally requires the canonical support-mask law to be shared by every candidate world. We first construct this all-world lift and then identify the limiting repair experiments.

Lemma 14 (All-world common-support faithful lift)**.** *Fix $h \geq 3$. Let $p_{r,N}, p_{z,N} \downarrow 0$ and fix $\kappa \in (0,1)$. For all sufficiently large N, one can choose, simultaneously for every $G \in \mathcal{G}_h^{\mathrm{grid}}$, a strictly positive finite-discrete distribution $P_{G,N}$ faithful to G such that all candidates have exactly the same marginal law $P_{G,N}(X,Y,Z) = \mu_N(X,Y,Z), \qquad G \in \mathcal{G}_h^{\mathrm{grid}}$, and this common marginal satisfies*

$$\begin{aligned} \mu_N(X=1) &= p_{r,N}\{1+o(1)\}, \quad \mu_N(Y=1) = p_{r,N}\{1+o(1)\}, \\ \mu_N(X=1,Y=1) &= \kappa p_{r,N}^2\{1+o(1)\}, \quad \mu_N(Z=1) = p_{z,N}\{1+o(1)\}. \end{aligned} \tag{101}$$

Consequently, at every sufficiently large finite N the readiness-mask law of the canonical row queries a, b and all canonical column queries q_i is exactly graph-independent.

Proof. Every graph in $\mathcal{G}_h^{\mathrm{grid}}$ contains the same upstream factors $P(X)$, $P(Y \mid X)$, and $P(Z)$. Candidate changes affect only the mechanism of T and, in one module, the descendants U_i, V_i. Hence the strictly positive Markov parameter space of each candidate can be written in product coordinates $\Theta_G^+ = \mathcal{X}^+ \times \mathcal{Y}_G^+$, where $\mathcal{X}^+$ parameterises the common upstream block $(P(X), P(Y \mid X), P(Z))$ and $\mathcal{Y}_G^+$ contains the candidate-specific descendant CPTs.

For each fixed finite DAG G, the unfaithful positive Markov parameters form a proper algebraic null subset $\mathcal{U}_G \subset \Theta_G^+$. By Fubini's theorem, for almost every $x \in \mathcal{X}^+$ the fibre $\mathcal{U}_{G,x} := \{y \in \mathcal{Y}_G^+ :$

$(x, y) \in \mathcal{U}_G\}$ has measure zero in $\mathcal{Y}_G^+$. The candidate class is finite, so the set $\mathcal{X}_{\text{good}}$ of upstream parameters having null unfaithful fibres for every candidate still has full measure in $\mathcal{X}^+$.

For each sufficiently large N, construct a positive upstream centre $\bar{x}_N$ with $P(X = 1) = P(Y = 1) = p_{r,N}, \quad P(X = 1, Y = 1) = \kappa p_{r,N}^2, \quad P(Z = 1) = p_{z,N}$; for example use the same $X \to Y$ parameterisation as in Eq. (97), with p_N replaced by $p_{r,N}$, and choose the Z root marginal separately. Choose a deterministic sequence $\varepsilon_N \downarrow 0$ and let $\mathcal{B}_N$ be a non-empty open neighbourhood of $\bar{x}_N$ in which the four displayed probabilities have relative error at most ε_N. Since $\mathcal{B}_N$ has positive measure and $\mathcal{X}_{\text{good}}$ has full measure, choose $x_N \in \mathcal{B}_N \cap \mathcal{X}_{\text{good}}$. For each candidate graph G, select $y_{G,N} \in \mathcal{Y}_G^+ \setminus \mathcal{U}_{G,x_N}$. Then $(x_N, y_{G,N})$ is strictly positive and faithful to G, while the common upstream coordinate x_N makes the (X, Y, Z) marginal exactly identical across all candidates. Membership in $\mathcal{B}_N$ gives Eq. (101).

The support event for $a = (X, T \mid Y)$ depends only on the counts of the Y strata; that for $b = (Y, T \mid X)$ depends only on the X strata; and every $q_i = (U_i, V_i \mid Z)$ uses the same Z partition and threshold. Their readiness masks are therefore functions only of the sampled (X, Y, Z) coordinates. Since those coordinates have exactly the same law in every candidate world, the mask law is graph-independent for each finite N covered by the construction. □

Lemma 15 (Independent tuning of row and column route limits)**.** *Fix $m \geq 2$, $r, c \in (0, 1)$, and $1/m < \gamma < 1$. There is a choice of common marginal in Lem. 14 and an equal-size client sequence such that the two canonical row routes each have limiting system-level readiness r, their exposures are asymptotically independent, and the shared canonical column route has limiting readiness c.*

Proof. Let $\lambda_r := -\log(1 - r), \quad \lambda_c := -\log(1 - c)$, and choose $p_{r,N} = a_r N^{-\gamma}, \quad p_{z,N} = a_z N^{-\gamma}, \quad \beta = \frac{m(1-\gamma)}{m-1}$. Write the client-count scaling as $K_N \sim \chi N^\beta, \quad \chi = \left(\frac{a_r^m}{m!\lambda_r}\right)^{1/(m-1)}$. For a rare conditioning state of probability $aN^{-\gamma}\{1 + o(1)\}$ and equal client size $n_N = N/K_N$, the one-client threshold probability is $\pi_N(a) = \frac{(n_N a N^{-\gamma})^m}{m!}\{1 + o(1)\}$, and therefore $K_N \pi_N(a) \longrightarrow \frac{a^m}{m!\chi^{m-1}}$. The choice of χ makes the X- and Y-route intensities converge to λ_r. Set $a_z = a_r \left(\frac{\lambda_c}{\lambda_r}\right)^{1/m}$. Then the Z-route intensity converges to λ_c. Standard rare-event Poisson limits give system-level readiness $1 - e^{-\lambda_r} = r$ for each row route and $1 - e^{-\lambda_c} = c$ for the shared column route.

Finally, Lem. 14 gives $P(X = 1, Y = 1) = O(p_{r,N}^2)$, while $n_N p_{r,N} \to 0$. The proof of Lem. 13 therefore applies unchanged and makes the two row-route exposure counts asymptotically independent. The all-world common-support lift makes these limits identical in every candidate causal world. □

For a repair $\mathcal{Q} = \{I_1, \ldots, I_q\}$, write its reference-normalised causal codeword as $\mathbf{c}_\mathcal{Q}(G) := \big(d_G(I_1), \ldots, d_G(I_q)\big)$, with the reference label encoded as zero in each coordinate. If $M_\mathcal{Q} \in \{0, 1\}^q$ is the readiness mask, correct-or-abstain gives the exact factorisation

$$Y_\mathcal{Q} = \text{Erase}(\mathbf{c}_\mathcal{Q}(G), M_\mathcal{Q})\,. \tag{102}$$

Thus the target-labelled causal codebook and the joint support-mask channel determine the complete repair-induced experiment.

For the canonical row and column repairs, the codewords are

$$\begin{aligned} &G_h^{\text{grid},\star} \mapsto (0, 0), \qquad G_{A_i} \mapsto (1, 0), \quad G_{B_i} \mapsto (0, 1) \quad \text{for } \mathcal{A}_h^{\text{row}}, \\ &G_h^{\text{grid},\star} \mapsto 0^h, \qquad G_{A_i}, G_{B_i} \mapsto e_i \qquad\qquad\qquad\quad \text{for } \mathcal{B}_h^{\text{col}}. \end{aligned} \tag{103}$$

The row code distinguishes A- from B-type candidates but not module index, whereas the column code distinguishes module index but not A from B; hence the comparison is restricted to losses on the fixed target summary. In the next two lemmas, $\succeq_B$ denotes ordinary Blackwell dominance between finite binary experiments.

Lemma 16 (Binary-erasure upper envelope)**.** *Let $E = (P_0, P_1)$ be a finite binary statistical experiment and let $\mathrm{BEC}(c)$ denote the binary erasure experiment that reveals the state with probability c and otherwise emits a common erasure symbol. Then*

$$\mathrm{BEC}(c) \succeq_B E \quad \Longleftrightarrow \quad c \geq \mathrm{TV}(P_0, P_1). \tag{104}$$

Proof. Necessity follows from contraction of total variation under stochastic post-processing. For sufficiency, let $M(y) := \min\{P_0(y), P_1(y)\}, \qquad A := P_0 - M, \qquad B := P_1 - M$. The non-negative measures A and B have disjoint supports and equal total mass $t := \mathrm{TV}(P_0, P_1)$, while M has mass $1 - t$. From $\mathrm{BEC}(t)$, map the state-0 reveal branch according to A/t, the state-1 reveal branch according to B/t, and the erasure branch according to $M/(1-t)$, with the obvious conventions at $t = 0$ or 1. This stochastic garbling produces exactly (P_0, P_1). If $c > t$, first erase an additional fraction of the revealed symbols to obtain $\mathrm{BEC}(t)$. □

Lemma 17 (Binary-erasure extraction)**.** *For the same binary experiment define the state-exclusive masses $s_0 := P_0\{y : P_1(y) = 0\}, \qquad s_1 := P_1\{y : P_0(y) = 0\}$. Then*

$$E \succeq_B \mathrm{BEC}(c) \quad \Longleftrightarrow \quad c \leq \min\{s_0, s_1\}. \tag{105}$$

Proof. If E is garbled to $\mathrm{BEC}(c)$, the state-0 reveal symbol must have zero probability under state 1. It can therefore receive mass only from outcomes on which P_1 is zero, and its state-0 probability is at most s_0. The symmetric argument gives $c \leq s_1$. Conversely, from each exclusive support retain, by random thinning if necessary, total mass exactly c and map it to its corresponding reveal symbol; map all remaining outcomes to erasure. This yields $\mathrm{BEC}(c)$. □

Proof of Thm. 4. Use Lem. 14 with the route scales supplied by Lem. 15. The complete discrepancy and repair classification in Lems. 10 and 11 is unchanged, so the canonical row and column repairs have structural costs 2 and h. The readiness-mask channel is exactly graph-independent for every sufficiently large finite N. Lem. 15 gives convergence of that channel to independent row masks with marginal readiness r and one shared column mask with readiness c. Since the output alphabets and candidate class are finite, the induced repair channels converge in total variation to the limiting experiments determined by Eq. (102).

Fix an arbitrary prior on the causal worlds and let $H := \mathcal{L}_T(G) \in \{0, 1\}$. If the prior assigns probability one to either target value, every target-only decision problem is already degenerate and both repairs attain the same Bayes envelope. It therefore suffices to consider priors giving both target values positive mass. Conditional on $H = 0$, write $\alpha := \Pr(G \text{ is of A-type} \mid H = 0)$. The module-index mixture is irrelevant to the row repair. With $s := 1 - r$ and outcome order $(\perp\perp, 0\perp, 1\perp, \perp 0, \perp 1, 00, 10, 01)$, the limiting row target experiment has

$$\begin{aligned} P_1 &= (s^2, rs, 0, rs, 0, r^2, 0, 0), \\ P_0 &= \big(s^2, (1-\alpha)rs, \alpha rs, \alpha rs, (1-\alpha)rs, 0, \alpha r^2, (1-\alpha)r^2\big). \end{aligned} \tag{106}$$

A direct calculation gives

$$\mathrm{TV}(P_0, P_1) = r \qquad \text{for every } \alpha \in [0, 1]. \tag{107}$$

For $0 < \alpha < 1$, the $H = 1$ exclusive mass is exactly r^2 (the outcome 00), while the total $H = 0$ exclusive mass is r. At the endpoints $\alpha \in \{0, 1\}$ both exclusive masses equal r. Hence the worst-case, over all nuisance mixtures induced by priors, of the smaller exclusive mass is

$$\inf_{\alpha \in [0,1]} \min\{s_0(\alpha), s_1(\alpha)\} = r^2. \tag{108}$$

For the column repair, failure of the shared route produces the common transcript $\perp^h$. When the route is ready, $H = 1$ produces 0^h and an $H = 0$ world with module index i produces e_i. Under any fixed prior, map every e_i to the $H = 0$ reveal symbol; conversely, from that reveal symbol sample i according to the prior's conditional module-index distribution. Thus the reduced column target experiment is Blackwell-equivalent to $\mathrm{BEC}(c)$ for every fixed prior.

Lem. 17 with Eq. (108) gives row dominance iff $c \leq r^2$, while Lem. 16 with Eq. (107) gives column dominance iff $c \geq r$, proving Eq. (17). At $c = r$, the column experiment strictly dominates with structural cost ratio $h/2 \to \infty$. The common-support equality is exact for every sufficiently large finite N; the Poisson route limits and decision ordering are limiting-experiment statements. □

K TRANSFER EXPERIMENTS AND EMPIRICAL BOUNDARIES

K.1 CONTROLLED TRANSFER EXPERIMENTS

The controlled diagnostics use exact finite route calculations where available, with Monte Carlo reserved for joint-route validation.

Critical-regime calculation.

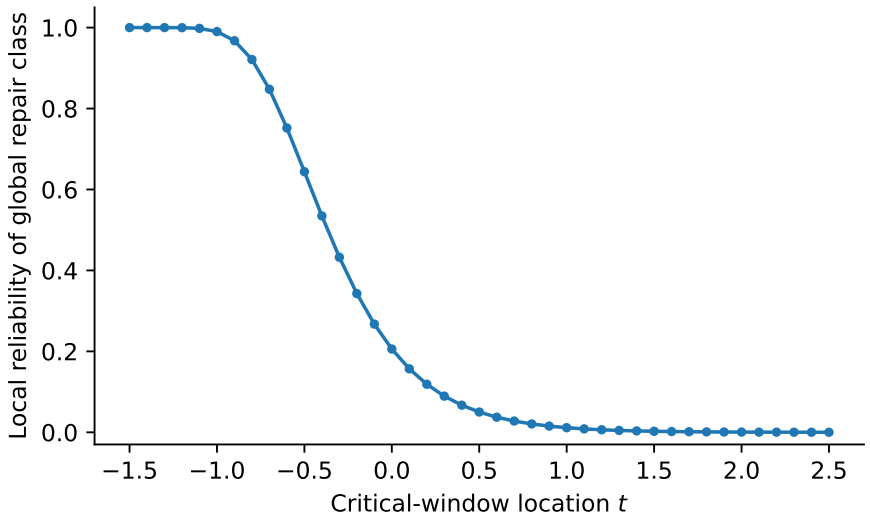


Figure 2: Controlled critical profile for Thm. 3(a). The finite curves converge toward the non-degenerate global-route limit across the pre-specified critical window.

The finite critical profile is evaluated by exact numerical calculation at 41 critical-profile locations and 30 fixed-quantile locations, together with the corresponding joint-population evaluations. The 0.1–0.9 transition width decreases from 0.1488 at $N = 10^3$ to 0.0558 at $N = 10^8$. The complete admissible competitor class is also checked: no allowed global refinement, mixed repair, or higher-order query bypasses the stated global-versus-module comparison. These calculations have no Monte Carlo uncertainty.

Matched-decomposition calculation.

Primary finite qualification and repair-selection decisions are computed from exact route marginals. A separate Monte Carlo validation grid checks joint-route behaviour and finite-sample diagnostics. Across the 63 validation cells, the maximum absolute Monte Carlo error is 0.00942 for the baseline route, 0.01525 for the global-R route, 0.01477 for the module route, and 0.00509 for the pooled/global-P route. These errors are small relative to the exact decision margins used for the primary ordering comparisons.

Finite-to-limit support-channel convergence. This finite support-channel diagnostic tracks convergence to the limiting parameters in Thm. 4. Fix $r = 0.6$, $r^2 = 0.36$, and $c \in \{0.25, 0.48, 0.72\}$ for the three limiting regions; exact calculations over $N = 10^3, \ldots, 10^8$ evaluate single-row readiness r_N, joint-row readiness j_N, and shared-column readiness c_N.

At $N = 10^8$, $r_N = 0.59552$, $j_N = 0.35465$, and the three column probabilities are 0.24851, 0.47649, and 0.71487; the evaluated grid remains in the predicted three regions. All 90 finite-channel garbling LP checks across five nuisance mixtures agree with the analytic finite thresholds.

The pre-specified value $q = 0.61$ in the repair-readiness construction is applied only after exact finite readiness probabilities are computed; Thm. 3(c) is the threshold-free same-pooled readiness separation, while the arbitrary-threshold crossing is its subsequent consequence.

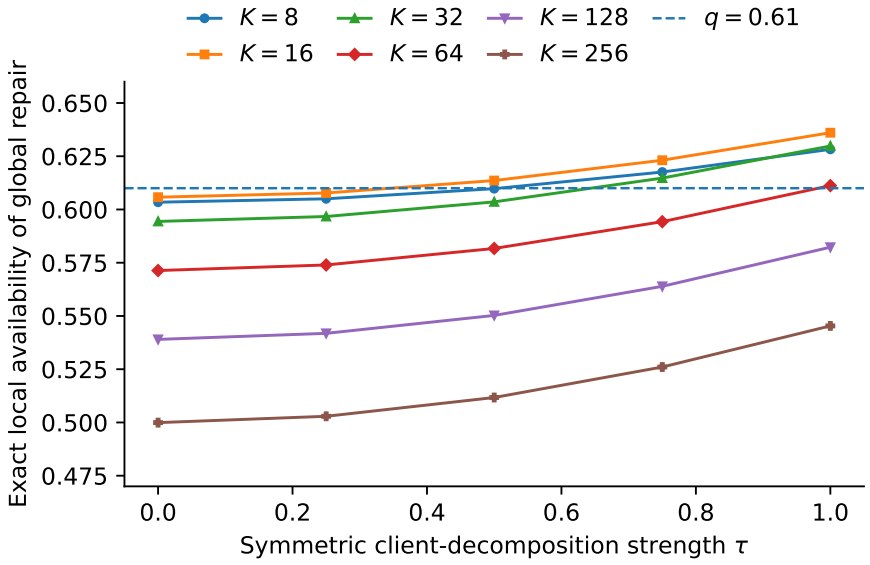


Figure 3: Exact global-route availability in the controlled matched-decomposition construction. The pooled availability is invariant across the compared client decompositions up to numerical precision.

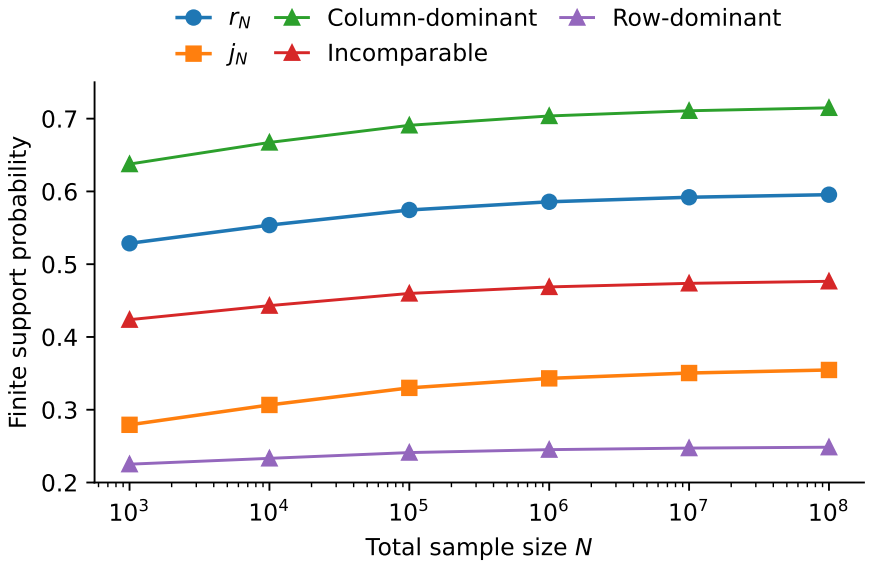


Figure 4: Finite support-channel convergence for Thm. 4. The single-row and joint-row probabilities approach $r = 0.6$ and $r^2 = 0.36$, while the three shared-column sequences approach the pre-specified limits $c = 0.25, 0.48, 0.72$.

K.2 Construction Validity and Semantic Stability

The construction-validity analysis checks CI-semantic preservation, faithful near-centre realisations, candidate roots for accidental independences, and same-pooled exposure separation. Here *pre-specified* means fixed before evaluation. **Pre-specified semantic validation.** For $h = 2$, the validation uses all 16 pre-specified total-sample sizes

$$N \in \{800, 1600, 2400, 3200, 4800, 6400, 9600, 12800,\\ 19200, 25600, 38400, 51200, 76800, 102400, 153600, 204800\},$$

while $h \in \{3, 4\}$ uses the pre-specified sentinel sample sizes $N \in \{800, 19200, 204800\}$. Each population is evaluated at $\omega \in \{0.25, 0.5, 0.75\}$, giving 66 semantic-validation cells and 1,912,320 complete CI-query evaluations. All 66 cells have zero semantic mismatches. Six large-N cells contain d-connected queries whose ordinary-precision conditional mutual information is below the pre-specified high-precision trigger; all six are recomputed at 100-digit precision and remain strictly positive, with zero high-precision failures. The six triggered near-zero connected cases remain strictly dependent at 100-digit precision.

Faithful near-centre realisations. The pre-specified constructive validation evaluates nonzero perturbations of the linked finite centre while recomputing the full joint distribution and all theorem-relevant route quantities. The four pre-specified witness settings, at $(h, N) = (2, 800), (2, 19200), (2, 204800), (3, 19200)$, all retain zero semantic mismatches at perturbation radius 10^{-4}; the same holds for all 12 pre-specified perturbations. Across the selected settings, the maximum relative route-cell change is 4.12×10^{-5}, the maximum relative change in a route coefficient θ is 6.27×10^{-6}, and the maximum baseline-to-global marginalisation error is 6.94×10^{-18}.

Candidate-root geometry. For each d-connected query, the validation selects a conditional-independence minor that is nonzero at the centre and writes the resulting necessary condition for accidental independence as a polynomial in the root-mixture parameter ω. The resulting roots form a conservative candidate superset. An exact catalogue of bad mixtures would require simultaneous vanishing of all relevant minors. At $N = 19{,}200$, all 2,512 d-connected queries for $h = 2$ and all 14,592 for $h = 3$ admit a resolved nonzero minor. The feasible interval contains 64 candidate-root occurrences for $h = 2$ (34 distinct after rounding to 12 decimal places) and 256 for $h = 3$ (130 distinct), and none of the displayed anchors $\omega \in \{0.25, 0.5, 0.75\}$ hits this candidate set.

Same-pooled identity and exposure separation. For each $h \in \{2, 3, 4\}$, the validation evaluates 41 pre-specified symmetric mixture points $\omega_-(\tau) = 0.5 - 0.25\tau, \quad \omega_+(\tau) = 0.5 + 0.25\tau, \quad \tau \in \{0, 0.025, \ldots, 1\}$, for 123 path cells in total. The maximum atomwise discrepancy in $\frac{1}{2}P_{\omega_-(\tau)} + \frac{1}{2}P_{\omega_+(\tau)} = P_{0.5}$ is 1.39×10^{-17} over the full pre-specified path (maximum ℓ_1 discrepancy 2.20×10^{-16}). At the same time, every pre-specified $\tau > 0$ has a positive global-route exposure gap. The smallest positive gap is 9.31×10^{-5} at the weakest pre-specified nonzero split; at $\tau = 1$, the $0.25/0.75$ split gives $\theta_G^{\mathrm{hom}} = 0.8656875$ and $\theta_G^{\mathrm{split}} = 1.014578125$, a gap of 0.148890625.

Table 4: Finite validation of construction validity and semantic stability. The table checks the finite witnesses used by the controlled construction; the asymptotic results are supplied by the analytic proofs.

Validation component	Pre-specified coverage	Result
Shared CI semantics	66 population cells; 1,912,320 complete CI-query evaluations	66/66 valid; 0 mismatches
High-precision safeguard	6 near-zero d-connected cases	6/6 remain strictly dependent
Faithful near-centre realisation	4 pre-specified witness settings; 12 pre-specified perturbations	4/4 settings and 12/12 perturbations valid
Route-law stability	selected faithful witnesses	max cell rel. error 4.12×10^{-5}; max θ rel. error 6.27×10^{-6}
Candidate-root geometry	$h = 2, 3$ at $N = 19{,}200$	0 unresolved minors; anchors hit no candidate root
Same-pooled identity	123 symmetric-path cells	max atomwise error 1.39×10^{-17}
Exposure separation	all pre-specified $\tau > 0$	positive throughout; minimum 9.31×10^{-5}

This finite validation complements, rather than replaces, the analytic lemmas establishing asymptotic faithfulness and countable avoidance. Candidate-root enumeration is reported for $h = 2, 3$, while the $h = 4$ sentinel settings check the remaining semantic and same-pooled conditions.

K.3 PUBLIC-BN EXPERIMENTAL DESIGN

K.3.1 CERTIFIED REPAIRS

Tab. 5 lists the repairs fixed before client-support realisations are generated, together with their structural certification levels.

Table 5: Fixed primary repair panel used in the full-network experiments. Certification is completed before any client-support realisation is generated.

Network	Target	Scope size	Certification	Queries
ALARM	BP	5	exact minimum	4
		6	exact minimum	5
		7	exact minimum	4
		8	exact minimum	4
CHILD	GruntingReport	5	exact minimum	3
		6	exact minimum	3
		7	exact minimum	4
		8	certified exact repair	12
INSURANCE	DrivHist	5	exact minimum	4
		6	exact minimum	5
		7	exact minimum	6
		8	inclusion-minimal exact repair	9

K.3.2 FULL-NETWORK SAMPLING AND READINESS

For each primary target, probabilities are induced by the complete public Bayesian network and its original BIF/CPT parameterisation. Let $\mathcal{V}_T^{(s)}$ denote the declared scope of size $s \in \{5, 6, 7, 8\}$. A master variable set contains the union of these scopes together with the root ancestor used by the matched-decomposition design, and exact marginalisation of the full network determines the master population. Each Monte Carlo trial draws one realised pooled master table; every scope, query, conditioning route, and repair in that trial is obtained by marginalising the same table, so all within-target comparisons share one realised full-network sample. For a query $I = (X_a, X_b, S)$, with $n_{k,S}(\omega)$ the count of state $\omega \in \Omega_S$ at client k, define

$$M_{k,I} := \min_{\omega \in \Omega_S} n_{k,S}(\omega),$$
$$\max_{k \in [K]} M_{k,I} \geq m. \tag{109}$$

for local query readiness, and $M_I^{\text{pool}} := \min_{\omega \in \Omega_S} \sum_{k=1}^{K} n_{k,S}(\omega)$ for the pooled control. Queries with the same conditioning set share one support route, and a fixed repair is ready when every distinct route it requires is ready. Different local routes may be supplied by different clients. The primary contract uses $m = 5$, with $m = 10$ used for sensitivity analysis. Throughout the public-network experiments, all reported ρ quantities denote readiness under the fixed full-network benchmark population for the reported target–scope cell; the candidate-world index is not varied. Subscripts such as loc, pool, hom, and split identify the evidence interface or client decomposition, not causal worlds. CI labels follow the correct-or-abstain evidence model in Assump. 2, so the empirical variable is the support readiness of each fixed repair under the realised client tables.

K.3.3 PAIRED CLIENT-ALLOCATION DESIGNS

Fixed-total: one realised pooled master table has size $N \in \{3{,}200, 12{,}800, 51{,}200, 204{,}800\}$ and is partitioned into $K \in \{2, 4, 8, 16, 32, 64, 128, 256\}$ clients whenever $25 \leq N/K \leq 2000$. Increasing K recursively splits the same observations, so the pooled table remains identical along the paired trajectory. **Fixed-local:** the local sample size is held at $n \in \{25, 50, 100, 200, 250, 400, 500, 1000, 2000\}$ with the same client-count grid and $Kn \leq 204{,}800$; existing clients are retained and new ones appended as K grows. **Same-pooled:** two client allocations are constructed from the same realised pooled master table. A result-blind root ancestor

defines the pre-specified split, with heterogeneity strength $\alpha \in \{0.25, 0.50, 0.75\}$; balanced integer transportation preserves every client size and pooled root-state count exactly, and observations are allocated within root states without conditioning on descendants. The complete pooled joint table, pooled route counts, and pooled repair readiness are therefore identical between the two allocations, with only the observation-to-client assignment changing.

All repairs, scope sizes, and paired design conditions for a given target share random draws so that repair, scope, and decomposition contrasts are paired.

K.3.4 Paired Monte Carlo Statistics

For a paired contrast between conditions A and B, let $Y^A_{s,t}, Y^B_{s,t} \in \{0,1\}$ denote the two readiness indicators on shared trial t of seed s, and define the seed-level paired mean $d_s := \frac{1}{T}\sum_{t=1}^{T}\big(Y^B_{s,t} - Y^A_{s,t}\big)$. With S independent seeds, we report $\widehat{\Delta} = S^{-1}\sum_{s=1}^{S} d_s$ and the normal-approximation paired 95% interval

$$\widehat{\Delta} \;\pm\; 1.96\,\frac{s_d}{\sqrt{S}}, \qquad s_d^2 := \frac{1}{S-1}\sum_{s=1}^{S}(d_s - \widehat{\Delta})^2. \tag{110}$$

Pairing is performed within each shared trial and the independent seed is the sampling unit for the interval; identical seed-level means give a point interval, reflecting zero observed seed-level variation rather than exact identification of the population contrast. Primary cells use $S = 10$ seeds and $T = 100$ trials per seed. Fig. 1 uses the harmonised seeds 42–51; broader evaluations use their pre-specified precision protocols unless stated otherwise. For the scope-ladder and same-pooled support-stress analyses below, a triggered precision cell is re-evaluated using 20 dedicated precision seeds ($S = 20$, $T = 100$), and the reported precision summary uses those 20 seeds alone rather than pooling them with the 10 primary seeds. Re-evaluation is triggered when the primary paired interval crosses zero, its width exceeds 0.02, readiness lies within 0.02 of a declared qualification threshold, or the across-seed standard deviation exceeds 0.05. The fixed-task complete-family analysis uses a separate precision protocol, described in App. K.5.1.

K.4 Aggregation-Depth Extension

For the aggregation-depth extension, each query route is independently assigned its exact minimum client-coalition size by solving the binary covering problem induced by its client-by-conditioning-cell count matrix; coalitions are therefore query-specific. By construction, $b = 1$ coincides with the client-local readiness mask and $b = K$ with the pooled readiness mask. The headline CHILD/ChestXray result in Fig. 1(a) uses $K = 32$, $N = 25{,}600$, $m = 5$, seeds 42–51, and 100 trials per seed.

To assess breadth without expanding the main figure, we apply the same harmonised cell to three additional CHILD targets. Tab. 6 reports the structural-minimum frontier at $b = 1$ and the first aggregation depth at which it reaches the best complete-repair readiness. Recovery is shallow but target dependent: Grunting and GruntingReport are already frontier-optimal locally, whereas ChestXray and XrayReport recover at depth two. This natural-network breadth check is descriptive rather than a universal bound; Cor. 1 supplies the worst-case statement.

Table 6: Harmonised CHILD aggregation-depth breadth at $K = 32, N = 25{,}600, m = 5$ with seeds 42–51 and 100 trials per seed. Frontier-recovery depth is the smallest b at which the structural-minimum frontier reaches the best complete-repair readiness and remains there over the evaluated larger depths.

CHILD target	C_T^{str}	$U^{(1)}_{C_T^{\mathrm{str}}}$	best higher-cost $U^{(1)}$	Recovery depth
ChestXray	4	0.119	1.000	2
Grunting	4	1.000	1.000	1
GruntingReport	3	1.000	1.000	1
XrayReport	3	0.130	1.000	2

K.5 Fixed-Task Complete Repair-Family Analysis

K.5.1 Complete Repair Families and Persistence Criterion

The fixed-task panel contains 17 eligible scopes, including 15 with multiple structural cost levels. Repairs with identical support-route masks are grouped into readiness-equivalent route classes; the panel contains 2,308,632 exact repairs and 476,420 certified route classes.

Empirical pairwise and frontier construction. For a fixed design cell and structural cost c, define the empirical cost-level extrema

$$\widehat{L}_c := \min_{\mathcal{A}:\, C(\mathcal{A})=c} \widehat{\rho}_{\mathrm{loc}}(\mathcal{A}), \qquad \widehat{U}_c := \max_{\mathcal{A}:\, C(\mathcal{A})=c} \widehat{\rho}_{\mathrm{loc}}(\mathcal{A}). \tag{111}$$

For each cost pair $c < c'$, the pairwise analysis compares a least-ready route-class representative at cost c with a most-ready representative at cost c'. This gives the strongest empirical pairwise contrast $\widehat{U}_{c'} - \widehat{L}_c$ for that cost pair. For the cost frontier, define

$$\widehat{F}_{<c} := \max_{c'<c} \widehat{U}_{c'}, \qquad \widehat{\Delta}_c^{\mathrm{front}} := \widehat{U}_c - \widehat{F}_{<c}. \tag{112}$$

Each displayed extremum is represented by one attaining route class; ties among attaining classes do not change the corresponding cost-level or frontier value.

Selected pairwise and frontier contrasts use seed-level 95% Monte Carlo precision intervals from the same primary seeds; these intervals summarise Monte Carlo precision of the selected empirical contrasts rather than simultaneous post-selection uncertainty over the full repair family. The persistence criterion below determines the reported outcomes.

A pairwise local-only reversal is declared for a cost pair when its strongest empirical pairwise contrast is greater than 0.01, the lower endpoint of the selected-contrast 95% Monte Carlo precision interval is positive, at least eight of the ten primary seeds have the same direction, the effect persists in at least two adjacent pre-specified cells, and the matched pooled advantage of the costlier representative is at most 0.02. A frontier reversal applies the same empirical gate to $\widehat{\Delta}_c^{\mathrm{front}}$. This distinction tests whether pairwise reordering changes the preferred cost level or is absorbed by another low-cost repair.

Table 7: Persistent fixed-task outcomes in the 15 pre-specified multi-cost scopes at $m = 5$. "Pairwise" denotes a strongest empirical pairwise contrast that satisfies the persistence criterion for at least one cost pair; "Frontier" denotes a persistent reversal against the complete empirical lower-cost frontier. HYPOVOLEMIA contains one isolated local-only candidate but fails the persistence requirement. These counts summarise persistent outcomes within the finite pre-specified multi-cost panel.

Network	Target	Structural costs	Pairwise	Frontier
ALARM	BP	4–13	yes	no
ALARM	CVP	2–3	no	no
ALARM	ERRCAUTER	1–2	no	no
ALARM	EXPCO2	4–10	no	no
ALARM	HYPOVOLEMIA	2–5	no	no
ALARM	KINKEDTUBE	1–2	no	no
ALARM	LVEDVOLUME	3–6	yes	no
ALARM	PCWP	2–3	no	no
ALARM	STROKEVOLUME	2–4	yes	no
CHILD	ChestXray	4–9	yes	yes
CHILD	Grunting	4–9	yes	yes
CHILD	GruntingReport	3–9	yes	yes
CHILD	RUQO2	3–9	yes	yes
CHILD	XrayReport	3–9	yes	yes
INSURANCE	DrivHist	4–10	no	no

A fixed-task lower-cost frontier reversal. CHILD/ChestXray is the strongest frontier example. The five-variable target task and complete conditioning-order-≤ 3 repair family are fixed; the best cost-four and cost-five classes use four and three support routes, respectively. Tab. 8 reports their

$K = 32$ frontier across five adjacent pre-specified sample sizes; ties among maximising route classes do not change the frontier values.

Table 8: Harmonised-seed fixed-task lower-cost frontier reversal for CHILD/ChestXray at $K = 32$ and $m = 5$, using seeds 42–51 and 100 trials per seed. Readiness values are empirical maxima over exact-repair route classes at each structural cost; the frontier gap compares the cost-five maximum with the complete empirical lower-cost frontier. Each displayed interval summarises seed-level 95% Monte Carlo precision for the selected contrast.

N	Best cost-4	Best cost-5	Frontier gap	95% MC precision interval	Pooled gap
3,200	0.000	0.028	0.028	$[0.0193, 0.0367]$	0.002
5,000	0.000	0.966	0.966	$[0.9617, 0.9703]$	0.000
6,400	0.000	1.000	1.000	$[1.0000, 1.0000]$	0.000
12,800	0.000	1.000	1.000	$[1.0000, 1.0000]$	0.000
25,600	0.119	1.000	0.881	$[0.8672, 0.8948]$	0.000

Aggregation-depth extension. For the harmonised headline cell $K = 32$, $N = 25{,}600$, and $m = 5$, we additionally compute the minimum number of clients whose query-specific conditioning-cell counts jointly meet the threshold for each support route. The minimum is solved exactly as a binary covering problem, and a repair's depth is the maximum route depth among its required routes. At $b = 1$, the complete cost-4 and cost-5 frontiers are 0.119 and 1.000. Across seeds 42–51, the cost-4 frontier at $b = 1$ has standard deviation 0.0281 and a 95% interval $[0.1016, 0.1364]$, and the cost-5 frontier exceeds it in all 10 seeds. At $b = 2$, both frontiers equal 1.000 in every seed and remain so for every larger $b \leq 32$, giving an absolute cost-4 recovery of 0.881. Thus this natural witness is limited by a one-client support bottleneck that is completely removed once each required route may aggregate over at most two clients; deeper aggregation provides no further gain. This is the empirical curve shown in Fig. 1(a). Accordingly, the two endpoint depths coincide exactly with the client-local and pooled readiness definitions.

At $N = 51{,}200$, both cost-four and cost-five frontiers reach readiness one, so the ChestXray reversal disappears once the support bottleneck is resolved. Across the panel, Tab. 7 records eight persistent pairwise reversals and five persistent frontier reversals.

K.5.2 Repair-Family Redundancy and Support Controls

ALARM supplies the complementary family-level outcome. BP, LVEDVOLUME, and STROKEVOLUME have persistent pairwise reversals but no persistent frontier reversal, so another lower-cost repair remains competitive after an individual low-cost comparison fails. Same-cost comparisons reinforce this interpretation: 13 of the 17 eligible scopes have non-zero readiness spread among repairs with equal structural cost, and several contain both zero- and one-readiness route classes. Structural query count therefore leaves substantial operational variation within a fixed cost layer.

The two client-scaling trajectories separate fragmentation from total-information growth. All 361 comparable homogeneous fixed-total frontier trajectories are non-increasing in K, whereas all 480 homogeneous fixed-local trajectories are non-decreasing. Pooling also provides a direct support control: 15 of 17 eligible scopes have positive maximum pooling rescue. KINKEDTUBE and DrivHist remain locally and pooled infeasible throughout the pre-specified grid, while EXPCO2 includes cells with zero local repair readiness and complete pooled rescue. These cases separate fragmentation from tasks whose support bottleneck is already present before local evidence formation.

Precision and same-pooled sensitivity. For panel-wide precision assessment, blocks selected by the pre-specified precision-trigger rule are re-evaluated without replacing the harmonised 42–51 headline values above. The 11 selected blocks use 30 total seeds (16001–16030), consisting of the 10 primary seeds and 20 additional precision seeds, with 100 trials per seed; all certified route classes in each selected block are included. Cost-level extrema and frontiers are recomputed from these 30-seed route-class estimates. After this higher-precision evaluation, none of the 11 blocks satisfies the trigger rule, providing a panel-wide Monte Carlo stability check for the complete-family results.

The fixed-task analysis also includes a mild same-pooled root-prior sensitivity; its maximum absolute frontier change is approximately 0.035. The stronger natural decomposition result in the main text therefore comes from the separate same-realised-pooled CHILD allocation control in App. K.7.1. The route-preservation check finds zero violations within matched conditioning-route comparisons, providing a finite consistency check alongside the analytic preservation guarantee.

K.6 Scope-Ladder Fragmentation and Support Geometry

K.6.1 Fixed-Total Fragmentation

For the five-variable ALARM/BP four-query exact minimum at $N = 51{,}200$, local readiness collapses as the same realised pooled sample is recursively divided among more clients, while pooled readiness remains one; Tab. 9 reports the high-K transition.

Table 9: Fixed-total fragmentation for the five-variable ALARM/BP repair at $N = 51{,}200$ and $m = 5$. Each larger-K setting further splits the same realised pooled observations.

Clients K	Local size N/K	ρ_{loc}	ρ_{pool}
32	1600	1.0000	1.0000
64	800	0.7020	1.0000
128	400	0.0040	1.0000
256	200	0.0000	1.0000

The target, full-network population, repair, total sample size, and realised pooled observations are fixed, so the transition isolates within-client conditioning support. Recursive splitting also supplies a trialwise structural check because it cannot create a locally ready route that was absent before the split, while pooled route counts remain invariant.

K.6.2 Fixed-Local Sample-Size Control

The complementary fixed-local design holds per-client sample size fixed and appends clients as K increases, retaining earlier local support opportunities while total information grows as $N = Kn$. Together, the fixed-total and fixed-local designs separate fragmentation of a fixed information budget from growth in total information.

K.6.3 Scope-Ladder Cost–Readiness Ordering

Scope ladders compare structural query count with local support readiness. In ALARM/BP, the five- and seven-variable repairs are both four-query certified exact minima yet can have sharply different local readiness. The five-versus-six comparison reverses the structural ordering: although $C(S5) = 4 < C(S6) = 5$, at $K = 128$ and $N = 51{,}200$ the local-readiness difference $\rho_{\mathrm{loc}}(S6) - \rho_{\mathrm{loc}}(S5)$ is 0.5775 with paired 95% interval $[0.5582, 0.5968]$. Enlarging the target scope changes the residual ambiguity class and admissible query universe, so the result identifies an ordering reversal at the evaluated scopes.

K.6.4 Conditioning-Support Mechanism in CHILD

CHILD/GruntingReport makes the mechanism explicit because structural query count and support burden move differently along the evaluated scope ladder. Tab. 10 summarises the fixed repairs and the minimum full-network conditioning-state probability across the distinct support routes used by each repair. The first three rows are exact minima and the eight-variable row is a certified exact repair.

The five- and six-variable repairs both contain three queries, yet the weakest route becomes nearly an order of magnitude rarer at six variables; the seven-variable repair has only four queries but an even weaker bottleneck whose conditioning partition reaches 36 cells. By contrast, the twelve queries in the eight-variable repair recombine onto only three distinct support routes, and its weakest route has substantially larger population mass than the smaller-scope bottlenecks. Several queries can therefore reuse a small number of well-supported routes while a shorter repair depends on one rare or highly

Table 10: Conditioning-support mechanism for the fixed CHILD/GruntingReport repair ladder. For a repair $\mathcal{Q}$, let $\mathcal{R}(\mathcal{Q})$ denote its distinct conditioning routes and Ω_R the state space of route R. The weakest-route probability is $p_{\mathrm{weak}}(\mathcal{Q}) := \min_{R\in\mathcal{R}(\mathcal{Q})} \min_{r\in\Omega_R} P_{\mathrm{BN}}(R = r)$, computed from the exact full-network BN marginal. Query count is a structural property of the repair, whereas $p_{\mathrm{weak}}(\mathcal{Q})$ and route composition describe the support burden induced by its conditioning requirements.

Scope	Queries	Weakest-route probability	Support summary
$S5$	3	1.47×10^{-2}	compact repair; moderate weakest route
$S6$	3	1.89×10^{-3}	same query count, substantially rarer bottleneck
$S7$	4	8.65×10^{-4}	weakest bottleneck; up to 36 conditioning cells
$S8$	12	2.40×10^{-2}	three distinct conditioning-support routes

stratified route. The composed support structure $\mathfrak{P}(\mathcal{Q})$ captures this repair-specific support geometry beyond raw query cardinality.

K.7 Same-Pooled Decomposition and Pre-Transfer Feasibility

K.7.1 Same-Pooled Decomposition in CHILD

For the six-variable CHILD/GruntingReport repair, the reported matched cell uses $N = 51{,}200$, $K = 32$, and $\alpha = 0.75$. The two client allocations have the same target, full Bayesian-network population, exact repair, total sample size, client count and sizes, realised pooled master table, pooled conditioning-route counts, and pooled repair readiness; only the allocation of observations across clients changes, as summarised in Tab. 11.

Table 11: Matched quantities in the CHILD same-pooled decomposition experiment. The comparison changes only how an already realised pooled sample is assigned to clients.

Quantity	Matched-control status
Target and fixed repair	identical
Full-network model and realised pooled observations	identical
Total sample size N and client count K	identical
Client-size vector	identical
Pooled conditioning-route counts	identical
Pooled repair-readiness mask	identical
Observation-to-client allocation	varied

The homogeneous allocation balances observations without using the pre-specified root-state contrast, whereas the split allocation redistributes the same observations according to the result-blind rule in Sec. K.3.3. The support-readiness probabilities are

$$\rho_{\mathrm{hom}} = 0.133, \qquad \rho_{\mathrm{split}} = 0.994, \qquad \rho_{\mathrm{split}} - \rho_{\mathrm{hom}} = 0.861, \qquad \rho_{\mathrm{pool}} = 1, \tag{113}$$

with seed-level paired 95% Monte Carlo interval $[0.8346, 0.8874]$ under the harmonised seeds 42–51. Matching the pooled sample fixes population, sample size, repair selection, client count, and pooled evidence. The remaining difference is the client allocation and whether the repair's conditioning support co-occurs within individual clients. The matched control therefore isolates decomposition sensitivity of repair readiness.

K.7.2 Pre-Transfer Support Infeasibility in INSURANCE

INSURANCE/DrivHist marks a different boundary. Under the primary all-strata readiness contract, some required conditioning routes contain nominal states with exact full-network probability zero. If $P(S = s) = 0$ for a required state, its realised count is zero for every finite pooled sample,

giving $\rho_{\text{pool}}(\mathcal{Q}) = \rho_{\text{loc}}(\mathcal{Q}) = 0$ for any repair that requires that route. The resulting pooled infeasibility identifies a pre-transfer support boundary under the common all-nominal-strata contract used throughout the primary analysis.

K.8 Robustness and Applicability

K.8.1 Threshold Sensitivity

For the fixed-task complete-family analysis, the stricter $m = 10$ sensitivity matches 574 frontier points to the primary $m = 5$ analysis; no matched $m = 10$ frontier value exceeds its $m = 5$ counterpart, and pre-specified reversal remains present in a subset of scopes.

K.8.2 Secondary-Target Robustness

The separate scope-ladder protocol evaluates reduced $m = 10$ grids and secondary targets ALARM/CVP, ALARM/VENTALV, and CHILD/XrayReport. Target-scope cases that fail the structural prerequisites for a fixed certified repair remain structural boundary information.

K.8.3 Paired-Design Consistency

Paired-design consistency is checked directly against the claims it supports. Fixed-total recursive splitting cannot create a locally ready route that was absent in its parent partition; fixed-local nested addition retains existing clients and therefore cannot destroy an earlier route; and the same-pooled construction enforces exact equality of pooled observations, pooled route counts, and pooled repair masks across the two allocations. Every evaluated preservation comparison obeys the order predicted by Cor. 3. These checks validate the paired contrasts and the intended separation among fragmentation, total-information growth, and decomposition-only effects.

K.8.4 Target–Scope Applicability Boundaries

Broader structural diagnostics delimit where the repair-transfer question is well posed. Fig. 5 separates tasks requiring no residual repair from those admitting a low-cost exact repair, while Fig. 6 illustrates a scope-admissibility boundary.

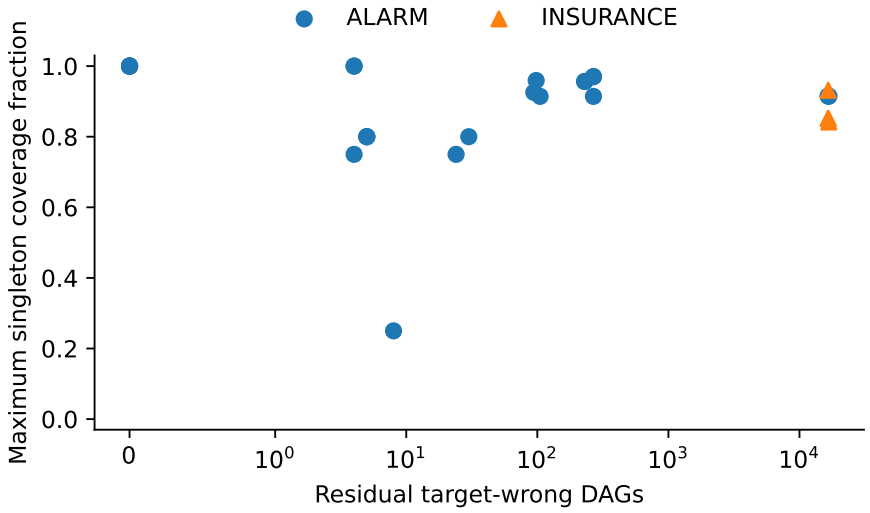


Figure 5: Structural repair boundary across 30 evaluated target-scope instances. In six instances the supplied baseline already identifies the target; sixteen admit no exact repair of cost at most three; the remaining eight admit a singleton, pair, or triple repair.

A target whose supplied baseline already resolves the residual ambiguity requires no repair. A target-scope pair that fails the declared structural scope contract remains outside the corresponding

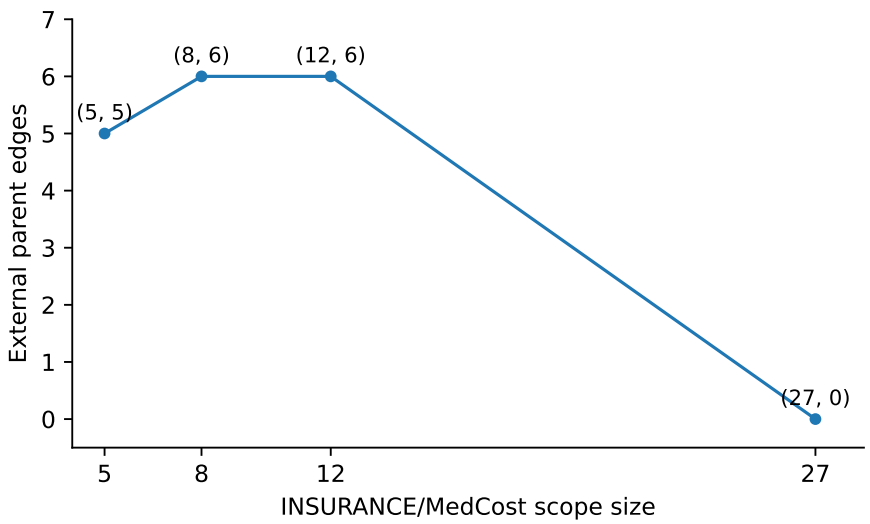


Figure 6: Scope-admissibility boundary for INSURANCE/MedCost. The evaluated five-node scope is not parent-closed, while parent closure expands to 16 variables.

federated readiness analysis, and a failed search below a stated cost cap is reported under that finite cap.

L Evidence-Interface Extensions and Controls

The following results characterise two interface extensions of the main evidence model: graph-independent garbling and exact aggregate-before-decision.

Corollary 8 (Graph-independent garbling preserves the module-touch lower bound)**.** *Let O_Γ^G be the correct-ready local oracle in Thm. 1. For each query I, let W_I be a Markov kernel that may depend on I and the realised transcript history but not on the hidden graph except through $\mathsf{O}_\Gamma^G(I)$. If $\widetilde{Y}_I \sim W_I\big(\cdot \mid \mathsf{O}_\Gamma^G(I)\big)$ reveals no graph identity or additional side information and preserves the readiness pattern of the hard family, then every protocol with worst-case target error at most $\delta < 1/2$ satisfies*

$$\mathbb{E}_{G_h^\star}[Q_\Pi^{\mathrm{mod}}] \geq (1-2\delta)h. \tag{114}$$

Proof. Run the garbled-oracle protocol using the more informative correct-ready oracle. Whenever the protocol requests I, obtain $\mathsf{O}_\Gamma^G(I)$, apply the same graph-independent kernel W_I, and return the simulated response. The transcript law and requested modules are exactly those under the garbled oracle. A violation of Eq. (114) would therefore induce a protocol violating Thm. 1 under the correct-ready oracle. □

Cor. 8 applies to graph-independent answer erasure and fixed label-noise channels under the declared oracle response. The bound therefore transfers to these garbled evidence channels while preserving the same module-touching argument.

Proposition 10 (Exact sufficient-statistic aggregation reproduces pooled evidence)**.** *For a discrete query $I = (X_a, X_b, S)$, let $N_{k,I}(x_a, x_b, s)$ be client k's complete local contingency-table count. Suppose an aggregation interface returns*

$$N_I^{\mathrm{agg}}(x_a, x_b, s) = \sum_{k=1}^{K} N_{k,I}(x_a, x_b, s) \tag{115}$$

for every state and applies the same readiness thresholds and correct-or-abstain response rule as the pooled interface. Then its response equals $A_I^{\mathrm{pool}}(G)$ for every allowed query. Consequently every

exact repair has identical readiness probabilities and thresholded ordering under exact aggregation and pooling, and the corresponding adaptive costs satisfy

$$C_{T,\Gamma}^{(K),\mathrm{ad}} = C_{T,\Gamma}^{\mathrm{pool,ad}}, \qquad C_T^{\mathrm{str}} : C_{T,\Gamma}^{\mathrm{loc,ad}} : C_{T,\Gamma}^{(K),\mathrm{ad}} : C_{T,\Gamma}^{\mathrm{pool,ad}} = 1 : h : 1 : 1 \tag{116}$$

for the Thm. 1 family. The local reversal in Thm. 3 is therefore absent under the matched exact-aggregation control.

Proof. For every cell, Eq. (115) equals the count in the contingency table formed from the union of the same client observations. Thus the query-specific empirical object is identical to the pooled table. The matched readiness rule therefore returns the same support status query by query; under Assump. 2, the surfaced label is then the same graph-implied label. This proves equality of every repair-level operational quantity. □

Prop. 10 identifies exact aggregate-before-decision as the matched positive interface. Under this interface, repair readiness and ordering coincide with the pooled control, isolating within-client evidence formation as the source of the local separation.