Which Local Guarantees Compose? A Universal Factorization Theorem for Graph-Indexed E-Certificates on Adaptive Execution Graphs

Manoj Saravanan [1]

Abstract

Adaptive agentic systems generate random execution graphs by routing among modules, joining intermediate certificates, transporting uncertainty across filtrations, and stopping based on observed history. We study which local uncertainty objects remain valid under such adaptive composition and show first that ordinary local marginal coverage and local expected-risk guarantees are not closed under predictable routing, even on finite depth-2 graphs. For bounded-horizon execution graphs whose local null classes are generated by finitely many bounded predictable constraints, we prove a fiberwise local representation theorem: every pointwise-admissible local e-certificate is quasi-surely an affine combination of the local constraints, and along any admissible law the corresponding coefficients can be selected predictably from the observed local state. Under explicit structural regularity—trace-generated normal form, serial precision, admissible same-null joins, and conditionally valid lifts—we then prove a universal factorization theorem: every universally valid graph calculus factors through a game-martingale serial update, weighted arithmetic averaging at same-null joins, and, under a value-based lift assumption, adjusters across filtration changes. Finally, for the canonical full-betting factorization we derive quantitative lower bounds showing an exact log-evidence tax at symmetric joins and an intrinsic lift-dependent loss, and we conclude that rank-based conformal e-predictors and bounded-risk local modules compose globally whenever they enter the workflow through the common local e-certificate interface.

[1]Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA, USA. Correspondence to: Manoj Saravanan <manoj663@vt.edu>.

Accepted at the ICML 2026 Workshop on Statistical Frameworks for Uncertainty in Agentic Systems (AgenticUQ), Seoul, South Korea, 2026. Copyright 2026 by the author(s).

1. Introduction

Adaptive agentic systems do not merely output a single prediction. They grow a *random computation graph*: at each round, the system conditions on an evolving filtration, selects a module or subagent, routes intermediate outputs, may transport uncertainty objects across filtrations, and eventually stops. The statistical question is therefore intrinsically compositional. One must ask not only whether an individual module is valid, but whether local uncertainty certificates survive *adaptive routing, joining, filtration change, and stopping* in the end-to-end execution graph.

Existing theory addresses only fragments of this question. At one end, recent graph-level certification frameworks provide nodewise or executable-graph guarantees, but do not characterize what survives arbitrary adaptive composition (Fayyazi & Akrami, 2026). At the other end, the modern e-value literature gives sharp primitive characterizations for constraint-generated e-variables, precise anytime-valid sequential merging, arbitrary-dependence same-null e-merging, and evidence transport across filtrations (Larsson et al., 2026; Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026). What is missing is a graph-level theorem connecting these two views: *which local uncertainty objects can actually compose on an adaptive execution trace?*

A first answer is negative. Local marginal coverage and local expected-risk guarantees are not closure primitives. Even when two local certificates are individually valid, a predictable router that inspects history may select among them in a way that destroys the advertised guarantee. This is the graph-level counterpart of post-hoc conformal selection failure (Hegazy et al., 2025). In Theorem 3.1, we formalize this obstruction on an explicit finite depth-2 execution graph. The consequence is structural: any universal compositional theory must be built from a primitive stronger than local coverage or local mean-risk control.

Our starting point is therefore the language of *graph-indexed e-certificates*. We model an adaptive workflow as a filtered random DAG generated by serial node additions, while joins and filtration lifts act on certificate objects attached to already-created nodes. Local nulls are given by finitely

many bounded predictable constraints. Under that model, the paper proves two representation theorems, one local and one global. The local theorem (Theorem 4.1) shows that on each history fiber, every pointwise-admissible local e-certificate is quasi-surely equal to an affine combination of the local constraint functions, and that along any fixed admissible law the corresponding coefficients may be chosen predictably. Thus, the locally admissible uncertainty layer collapses to a finite-dimensional predictable cone.

The global theorem (Theorem 5.2) shows that, once one imposes explicit structural regularity—trace-generated normal form, serial precision, primitivewise admissibility of same-null joins, and conditionally valid lifts—a universally valid graph calculus has essentially no remaining freedom. It must factor through three primitives: a game-martingale serial update, weighted arithmetic averaging at same-null joins, and, under the strengthened value-based lift assumption, adjusters across filtration changes (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026). This is a necessity theorem, not merely a construction: the theorem identifies the only primitive algebra compatible with universal composability in the stated regime.

The paper also quantifies the cost of this rigidity. For the canonical full-betting factorization, Theorem 6.1 shows that composability is not free: weighted same-null joins incur an exact log-evidence loss, and value-based lifts incur a lift-dependent lower bound. On symmetric stars, the join tax is exactly $\log k$. Thus the theory is not only structural but quantitative: the primitive restrictions force a measurable evidential price.

A final contribution is an interface result rather than a new primitive. Appendix G shows that two natural module classes fit the local e-certificate abstraction: rank-based conformal e-predictors built from exchangeable ranks (Gauthier et al., 2025), and bounded-risk exponential modules for conditional mean-risk nulls, closely related to conformal risk-control viewpoints (Angelopoulos et al., 2024; Hultberg et al., 2026). Consequently, Theorem I.1 yields a compositional corollary: once such modules are expressed through a common local e-certificate interface and joined only after reduction to a common downstream null, they produce a global e-process on the adaptive workflow.

The scope of the workshop theorem is deliberately narrow. We work with bounded horizon, finitely many bounded predictable constraints, same-null joins, trace-generated evaluation traces, and no envelope theorem for arbitrary certificate families beyond the e-certificate setting. These are not cosmetic restrictions. Each omitted case—intrinsically multi-null joins, unbounded stopping without an infinite-horizon factorization theory, infinite constraint classes without finite active reduction, or arbitrary non-e certificate families without an envelope—would require a genuinely new primitive theory. Within this scope, however, the result is exact: the paper gives a representation theorem for uncertainty on adaptive agent graphs.

The remainder of the paper follows the logical dependency of the theory. Section 2 formalizes adaptive execution graphs and certificate calculi. Section 3 proves the routing obstruction. Section 4 develops the local predictable-constraint representation theorem. Section 5 proves the universal factorization theorem, and Section 6 quantifies the resulting price of universal composability. Section I states the module corollary and records the precise scope of the workshop version.

2. Setup: Adaptive Execution Graphs and Certificate Calculi

We work in discrete time on a deterministic finite horizon T. The full measure-theoretic construction is given in Appendix A; in particular, all measurable spaces are standard Borel, all variables are defined on the whole probability space using cemetery conventions after stopping, and regular conditional laws exist whenever needed. The essential modeling choice is to separate *graph generation* from *certificate manipulation*: the random execution graph is created purely by serial node additions, while same-null joins and filtration lifts act on certificate objects attached to an already-created graph. This separation is the key technical device that later lets us reduce universal composability to three primitive operations from the e-value literature (Larsson et al., 2026; Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026).

Definition 2.1 (Adaptive execution graph)**.** An *adaptive execution graph* is a tuple

$$\mathfrak{E} = \Big(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_t)_{t=0}^T, \tau, (m_t, P_t, X_t, Z_t)_{t=1}^T\Big)$$

such that:

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space and $(\mathcal{F}_t)_{t=0}^T$ is a filtration;
2. τ is an $(\mathcal{F}_t)$-stopping time with $0 \le \tau \le T$;
3. for each $t \in [T]$, the triple (m_t, P_t, X_t) is $\mathcal{F}_{t-1}$-measurable, where
$$m_t \in \mathcal{M}, \qquad P_t \subseteq [t-1], \qquad X_t \in \mathcal{X}_t;$$
4. for each $t \in [T]$, the output $Z_t \in \mathcal{Z}_t$ is $\mathcal{F}_t$-measurable.

The local information sigma-field at round t is

$$\mathcal{H}_t := \sigma(X_t) \subseteq \mathcal{F}_{t-1},$$

and the realized graph at time $u \in [T]_0$ is

$$\begin{aligned} G_u &= (V_u, E_u), \\ V_u &= [u \wedge \tau], \\ E_u &= \{(s,t) : t \le u \wedge \tau,\ s \in P_t\}. \end{aligned}$$

Thus the terminal realized graph is G_τ, and every edge points from an earlier round to a later round.

Definition 2.2 (Local null class)**.** Fix $t \in [T]$. Let $r_t \in \mathbb{N}$, and let

$$\phi_{t,j} : \mathcal{X}_t \times \mathcal{Z}_t \to \mathbb{R}, \qquad j = 1, \dots, r_t,$$

be bounded measurable constraint functions. For a local state $x \in \mathcal{X}_t$, define the local null class

$$\mathcal{P}_t(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \begin{array}{c} \displaystyle\int \phi_{t,j}(x,z)\, Q(dz) \le 0, \\ \text{for all } j = 1, \dots, r_t \end{array} \right\}.$$

A probability law P on $(\Omega, \mathcal{F})$ is called *admissible* if, for every $t \in [T]$, the regular conditional law of Z_t given X_t under P, denoted $K_t^P(\cdot \mid X_t)$, satisfies

$$K_t^P(\cdot \mid X_t) \in \mathcal{P}_t(X_t) \qquad P\text{-a.s.}$$

Equivalently,

$$\begin{aligned} &\mathbb{E}_P\big[\phi_{t,j}(X_t, Z_t) \mid X_t\big] \le 0, \\ &P\text{-a.s. for all } j = 1, \dots, r_t. \end{aligned}$$

Definition 2.3 (Local e-certificate)**.** A *local e-certificate* at round t is a measurable function

$$e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0, \infty)$$

such that

$$\int e_t(x,z)\, Q(dz) \le 1,$$
$$\text{for every } x \in \mathcal{X}_t \text{ and every } Q \in \mathcal{P}_t(x).$$

Its realized value is

$$E_t := e_t(X_t, Z_t).$$

Under every admissible law P, every local e-certificate satisfies the conditional validity relation

$$\mathbb{E}_P[E_t \mid X_t] \le 1 \qquad P\text{-a.s.}$$

For fixed t and $x \in \mathcal{X}_t$, we write

$$\mathcal{E}_t(x) := \left\{ e : \mathcal{Z}_t \to [0, \infty) : \begin{array}{c} \displaystyle\int e(z)\, Q(dz) \le 1, \\ \text{for all } Q \in \mathcal{P}_t(x) \end{array} \right\}$$

for the class of all valid local e-functions over the history fiber x.

Definition 2.4 (Global e-process)**.** Let P be an admissible law. A nonnegative $(\mathcal{F}_t)$-adapted process $M = (M_t)_{t=0}^T$ with $M_0 = 1$ is an *e-process under P* if

$$\mathbb{E}_P[M_t \mid \mathcal{F}_s] \le M_s \qquad P\text{-a.s. for all } 0 \le s \le t \le T.$$

More generally, if $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$ and $\mathcal{G} \subseteq \mathcal{F}$, then a nonnegative $\mathcal{G}$-measurable random variable E is a *$\mathcal{Q}$-e-value on $\mathcal{G}$* if

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[E] \le 1.$$

Definition 2.5 (Same-null join)**.** Let J be a finite index set. A family

$$\big(E_i, \mathcal{G}_i\big)_{i \in J}, \qquad E_i \in L_+^0(\Omega, \mathcal{G}_i),$$

is called a *same-null family* if there exist:

1. a null class $\mathcal{Q}_J \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$,
2. a comparison sigma-field $\mathcal{G}_J \subseteq \mathcal{F}$ with $\mathcal{G}_i \subseteq \mathcal{G}_J$ for every $i \in J$,
3. lift operators

$$A_i : L_+^0(\Omega, \mathcal{G}_i) \to L_+^0(\Omega, \mathcal{G}_J), \qquad i \in J,$$

such that each lifted certificate $A_i(E_i)$ is a $\mathcal{Q}_J$-e-value on $\mathcal{G}_J$. A *same-null join* is any operation that combines such a lifted family into a single certificate object on $\mathcal{G}_J$. At this stage the join rule is left abstract; the main theorem later shows that universal composability forces a unique admissible join primitive.

Definition 2.6 (Certificate calculus)**.** A *certificate calculus* is a rule Γ that maps an adaptive execution graph $\mathfrak{E}$ together with a family of local e-certificates $(e_t)_{t=1}^T$ to a nonnegative $(\mathcal{F}_t)$-adapted process

$$M^\Gamma = (M_t^\Gamma)_{t=0}^T$$

with $M_0^\Gamma = 1$. We call Γ *universally valid* if M^Γ is an e-process under every admissible law.

Assumption 2.7 (Workshop scope)**.** Throughout the workshop version, we impose the following restrictions.

1. **Bounded horizon and bounded stopping:** $T < \infty$ and $\tau \le T$.
2. **Serial graph growth:** each round creates at most one new node, so $V_\tau = [\tau]$.
3. **Constraint-generated local nulls:** every local null class is generated by finitely many bounded predictable constraints as in Theorem 2.2.
4. **Same-null joins only:** every join in the main theorem merges certificates for a common downstream null.

5. **Finite-valued local certificates:** the workshop paper is stated for graph-indexed e-certificates and does not claim an envelope theorem for arbitrary statistical certificates.

6. **Bounded unfolding of recursion:** recursive or self-referential workflows are covered only after finite unfolding into a DAG.

Assumption 2.8 (Axioms for candidate calculi)**.** Any certificate calculus Γ considered in this paper is required to satisfy the following structural axioms.

1. **Universal validity:** for every admissible law P, the output M^Γ is an e-process under P.

2. **Locality:** if two admissible inputs agree on an ancestral subgraph (i.e., a vertex set closed under parenthood), then the restrictions of their outputs to that ancestral subgraph agree.

3. **Monotonicity:** if $E_t \le \widetilde{E}_t$ for every realized node t, then $M_u^\Gamma(E_{1:\tau}) \le M_u^\Gamma(\widetilde{E}_{1:\tau})$ for all $u \le T$.

4. **Graph-isomorphism invariance:** relabelings preserving creation order and parent relations do not change the calculus.

5. **Serial associativity:** composition along chains is associative.

6. **Parallel anonymity:** a same-null join depends on its inputs only through the lifted certificate values and predictable weights.

7. **Filtration naturality:** the rule for lifting certificates from coarser to finer sigma-fields is graph-independent.

Appendix A shows that the realized execution graph is generated purely by serial node additions. The only nontrivial operations are therefore: serial accumulation of evidence along the time axis, same-null joins of certificates already attached to realized nodes, and lifts across filtrations. The later factorization theorem will show that universal composability leaves essentially no freedom beyond these three primitives, matching the chain, same-null merging, and filtration-lift phenomena isolated in (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026).

3. Why Ordinary Local Guarantees Do Not Compose

The first obstacle is structural: local guarantees phrased only in terms of marginal coverage or expected risk are not stable under adaptive routing. This is the graph-theoretic analogue of post-hoc selection failures for conformal sets (Hegazy et al., 2025), but the construction we use is fully finite, exact, and self-contained. The conclusion is not merely that some particular composition rule fails; rather, the class of locally valid set-valued or risk certificates is itself not closed under predictable selection. Hence such objects cannot serve as the primitive layer for a universally composable graph calculus.

Theorem 3.1 (Router non-composability of marginal coverage)**.** *For every $\alpha \in (0,1)$, there exist finite spaces $\mathcal{W} = \mathcal{Y}$, random variables W and Y taking values in $\mathcal{W}$, two set-valued predictors*

$$C_1, C_2 : \mathcal{W} \to \mathcal{P}(\mathcal{Y}),$$

and a measurable router $R : \mathcal{W} \to \{1,2\}$ such that the routed predictor

$$C_R(w) := C_{R(w)}(w)$$

satisfies

$$\mathbb{P}\{Y \in C_i(W)\} = 1-\alpha, \qquad i = 1,2,$$

but

$$\mathbb{P}\{Y \in C_R(W)\} = 1-\alpha-\frac{1}{2}\min\{\alpha, 1-\alpha\} < 1-\alpha.$$

Moreover, the construction embeds into a depth-2 adaptive execution graph with one root node, two child modules, and one router node.

Proof sketch. Fix $\alpha \in (0,1)$, and define

$$p := \frac{1}{2}\min\{\alpha, 1-\alpha\}, \qquad q := \alpha - p, \qquad r := 1-\alpha-p.$$

Let $\mathcal{W} = \mathcal{Y} = \{0,1,2,3\}$, and set

$$\begin{aligned} &\mathbb{P}(W=0) = r, \quad \mathbb{P}(W=1) = p, \\ &\mathbb{P}(W=2) = p, \quad \mathbb{P}(W=3) = q, \qquad Y := W. \end{aligned}$$

Define failure sets

$$A_1 := \{1,3\}, \qquad A_2 := \{2,3\},$$

and predictors

$$C_i(w) := \begin{cases} \varnothing, & w \in A_i, \\ \{w\}, & w \notin A_i. \end{cases} \qquad i = 1,2.$$

Since $Y = W$, predictor C_i fails exactly on A_i, and therefore

$$\mathbb{P}\{Y \notin C_i(W)\} = \mathbb{P}(W \in A_i) = p + q = \alpha,$$

so each C_i has exact marginal coverage $1-\alpha$.

Now define the router by

$$R(w) := \begin{cases} 1, & w \in \{0,1,3\}, \\ 2, & w = 2. \end{cases}$$

Then on every state $w \in \{1, 2, 3\}$, the router selects a predictor that fails at w, while at $w = 0$ both predictors succeed. Hence

$$\{Y \in C_R(W)\} = \{W = 0\},$$

and therefore

$$\begin{aligned}\mathbb{P}\{Y \in C_R(W)\} &= \mathbb{P}(W = 0) = r \\ &= 1 - \alpha - \frac{1}{2}\min\{\alpha, 1 - \alpha\}.\end{aligned}$$

The execution-graph embedding is obtained by letting a root node reveal W, two child modules output $C_1(W)$ and $C_2(W)$, and a final router node output $C_{R(W)}(W)$. The full proof, including the formal selection-as-routing embedding, is given in Appendix B. □

Corollary 3.2 (Local coverage and risk are not compositional primitives)**.** *No universally valid certificate calculus can be based solely on local marginal coverage statements or post-hoc selection among individually valid set-valued predictors. The same obstruction applies to local expected-risk guarantees: if one sets*

$$\ell_i(w, y) := \mathbf{1}\{y \notin C_i(w)\}, \qquad i = 1, 2,$$

then each local predictor has risk $\mathbb{E}[\ell_i(W, Y)] = \alpha$*, whereas the routed predictor has strictly larger risk*

$$\mathbb{E}[\ell_{R(W)}(W, Y)] = \alpha + \frac{1}{2}\min\{\alpha, 1 - \alpha\}.$$

The theorem identifies the core obstruction that drives the rest of the paper. The failure is neither asymptotic nor high-dimensional; it already occurs on a four-point space with two exact local guarantees and a deterministic one-step router. Thus the defect lies in the *primitive object*, not in a particular proof technique. Once a workflow may inspect history and route among locally valid outputs, marginal coverage and expected risk cease to be closed under composition. The positive theory must therefore be built from objects whose validity survives serial accumulation, same-null joining, and filtration refinement. The remainder of the paper shows that graph-indexed e-certificates are precisely such objects.

4. Local Representation Under Predictable Constraints

The local question is the following: once a module is conditioned on its predictable state $X_t = x$, what are the admissible local e-certificates for the fiberwise null class $\mathcal{P}_t(x)$? The answer is unexpectedly rigid. Because the local null is generated by finitely many bounded predictable constraints, every relevant local certificate collapses, up to null sets of the entire fiberwise null class, to a finite-dimensional affine combination of the constraint functions. This is the history-indexed specialization of the finite-constraint e-variable theory of (Larsson et al., 2026), but the version we need here has an additional predictable-selection feature: along any fixed admissible law, the affine coefficients can be chosen as a measurable function of the observed local state.

Fix a round $t \in [T]$ and a local state $x \in \mathcal{X}_t$ with $\mathcal{P}_t(x) \neq \varnothing$. A set $N \in \mathcal{Z}_t$ is $\mathcal{P}_t(x)$*-polar* if $Q(N) = 0$ for every $Q \in \mathcal{P}_t(x)$. For measurable $f, g : \mathcal{Z}_t \to [0, \infty]$, write

$$f \preceq_x g$$

if $f \leq g$ outside a $\mathcal{P}_t(x)$-polar set. For $\lambda = (\lambda_1, \ldots, \lambda_{r_t}) \in \mathbb{R}_+^{r_t}$, define the affine constraint combination

$$H_x^\lambda(z) := 1 + \sum_{j=1}^{r_t} \lambda_j \phi_{t,j}(x, z), \qquad z \in \mathcal{Z}_t,$$

its nonnegative envelope

$$\bar{H}_x^\lambda := H_x^\lambda \vee 0,$$

and the feasible coefficient set

$$\Pi_t(x) := \left\{\lambda \in \mathbb{R}_+^{r_t} : H_x^\lambda \geq 0 \quad \mathcal{P}_t(x)\text{-q.s.}\right\}.$$

For $\lambda \in \Pi_t(x)$, the envelope $\bar{H}_x^\lambda$ belongs to $\mathcal{E}_t(x)$, and $\bar{H}_x^\lambda = H_x^\lambda$ $\mathcal{P}_t(x)$-quasi surely.

Theorem 4.1 (Predictable local representation)**.** *Fix $t \in [T]$ and $x \in \mathcal{X}_t$ with $\mathcal{P}_t(x) \neq \varnothing$. Then the following hold.*

1. ***Fiberwise complete class.*** *For every $e \in \mathcal{E}_t(x)$, there exists $\lambda \in \Pi_t(x)$ such that*

   $$e \preceq_x \bar{H}_x^\lambda.$$

   Hence the family $\{\bar{H}_x^\lambda : \lambda \in \Pi_t(x)\}$ is a complete class for the local e-problem on the history fiber x.

2. ***Representation of pointwise-admissible fibers.*** *If $e \in \mathcal{E}_t(x)$ is pointwise admissible in the sense of Theorem A.14, then there exists $\lambda \in \Pi_t(x)$ such that*

   $$e = H_x^\lambda \qquad \mathcal{P}_t(x)\text{-q.s.}$$

3. ***Predictable measurable selection.*** *Fix an admissible law P and a local e-certificate $e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0, \infty)$. Suppose there exists $B_t \in \mathcal{X}_t$ with $P(X_t \in B_t) = 1$ such that for every $x \in B_t$, the fiber $e_t(x, \cdot)$ is pointwise admissible in $\mathcal{E}_t(x)$. Then there exists a Borel measurable map*

   $$\lambda_{t,P} : \mathcal{X}_t \to \mathbb{R}_+^{r_t}$$

   such that

   $$\begin{aligned} &e_t(X_t, \cdot) = H_{X_t}^{\lambda_{t,P}(X_t)}(\cdot), \\ &K_t^P(\cdot \mid X_t)\text{-a.s.}, \quad P\text{-a.s.}\end{aligned}$$

Equivalently, with $E_t = e_t(X_t, Z_t)$,

$$E_t = 1 + \sum_{j=1}^{r_t} \lambda_{t,P,j}(X_t)\, \phi_{t,j}(X_t, Z_t) \qquad P\text{-a.s.},$$

and the coefficient vector $\lambda_{t,P}(X_t)$ is $\mathcal{H}_t = \sigma(X_t)$-measurable.

Proof sketch. The argument is fiberwise. Fix $x \in \mathcal{X}_t$, and view the measurable space $(\mathcal{Z}_t, \mathcal{Z}_t)$ together with the finite family of bounded constraints $\{\phi_{t,j}(x, \cdot)\}_{j=1}^{r_t}$. This is exactly a finite-constraint e-variable problem, so the complete-class theory of (Larsson et al., 2026) applies on the fiber: every local e-function is quasi-surely dominated by an affine constraint combination H_x^λ with $\lambda \in \Pi_t(x)$. This yields part (1).

For part (2), pointwise admissibility implies quasi-sure admissibility: if a competitor dominates e strictly outside a polar set, then replacing it by its pointwise maximum with e preserves validity because polar sets are null for the entire fiberwise null class, producing a genuine pointwise dominator and contradicting Theorem A.14. The fiberwise complete-class theorem then forces e to coincide quasi surely with some H_x^λ.

For part (3), fix an admissible law P. For each x, consider the correspondence

$$\Lambda_{t,P}^{e_t}(x) := \left\{\lambda \in \mathbb{R}_+^{r_t} : e_t(x, \cdot) = H_x^\lambda(\cdot) \;\; K_t^P(\cdot \mid x)\text{-a.s.}\right\}.$$

By part (2), the sections are nonempty on a P_{X_t}-full set. Appendix C shows that this correspondence has closed values and a Borel graph. Since the parameter space $\mathbb{R}_+^{r_t}$ is Polish, a measurable-selection theorem yields a Borel selector $x \mapsto \lambda_{t,P}(x)$. Substituting $x = X_t$ gives the displayed representation, and $\lambda_{t,P}(X_t)$ is $\sigma(X_t)$-measurable by construction. The full proof appears in Appendix C. □

Theorem 4.1 is the local backbone of the paper. It says that, once the module is conditioned on its predictable history, the apparent infinite-dimensional space of valid local e-certificates collapses to a finite-dimensional cone indexed by nonnegative constraint multipliers. This is the sense in which predictable uncertainty is locally linear. The next section shows that, after this local collapse, the only remaining freedom at the global graph level is how these local certificates may be accumulated, joined, and lifted across filtrations.

5. A Universal Factorization Theorem

Appendices D and E now let us pass from local structure to a global characterization. The reduction lemmas show that every candidate graph calculus hides three primitive restrictions: a sequential constructor on pure chains, an arbitrary-dependence merger on already-lifted same-null stars, and a lift operator on pure filtration refinements. The remaining question is whether these three primitive objects can be recombined into a universally valid graph calculus, and if so, whether there is any structural freedom beyond them. The answer is essentially no: once one imposes the right primitive regularity, the global calculus factors through a game-martingale serial update, weighted arithmetic averaging at same-null joins, and—under a mild strengthening on the lift primitive—adjusters across filtration changes (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026).

Proposition 5.1 (Primitive reduction)**.** *Let Γ be a certificate calculus satisfying Theorems 2.7 and 2.8. Then, by restriction to special execution graphs, Γ induces:*

1. *a bounded sequential e-merging constructor on pure chains;*
2. *a bounded arbitrary-dependence same-null merger on already-lifted stars;*
3. *a bounded lift operator on pure filtration refinements.*

Proof sketch. Appendix D realizes each primitive class inside the execution model of Section 2: chains are encoded by serial node additions with one-step state-based e-values, already-lifted same-null stars are encoded by edgeless graphs on a common comparison sigma-field, and pure refinements are encoded by one-node graphs with a coarse-to-fine filtration change. Universal validity yields valid outputs on each of these special inputs; locality, serial associativity, parallel anonymity, and filtration naturality show that the induced operators depend only on the corresponding primitive data. □

Structural regularity. To turn Theorem 5.1 into a characterization theorem, we impose the exact extra regularity needed to invoke the primitive complete-class results without hidden gaps. Formally, these conditions are stated in Appendix E; informally, they are:

1. *trace-generated normal form*: the output of the calculus is assembled from an ordered sequence of stage certificates, each built from already available local certificates by finitely many lifts followed by one same-null join;
2. *serial precision*: on precise pure chains, the serial restriction is a precise anytime-valid constructor;
3. *primitivewise admissibility of joins*: the pure same-null join restriction is an admissible arbitrary-dependence e-merging function;

4. *conditional lift validity*: every lift used inside the trace preserves one-step conditional e-validity.

When we additionally assume that pure lifts act by increasing deterministic functions that universally lift running maxima of e-processes across filtrations, the lift primitive is forced to be an adjuster (Choe & Ramdas, 2026).

Theorem 5.2 (Universal factorization theorem)**.** *Let Γ be a certificate calculus satisfying Theorems 2.7 and 2.8 together with the structural regularity conditions above. Then for every admissible execution graph $\mathfrak{E}$, every compatible certificate evaluation trace, and every admissible law P, there exist:*

1. *an ordered stage decomposition $u = 1, \dots, m(\mathfrak{E})$;*
2. *stage certificates J_u^Γ and the associated stage filtration*
$$\mathcal{K}_u := \sigma(J_1^\Gamma, \dots, J_u^\Gamma);$$
3. *for each stage u, lifted stage inputs*
$$\widetilde{E}_{u,1}, \dots, \widetilde{E}_{u,k_u} \in L_+^0(\Omega, \mathcal{K}_u),$$
together with nonnegative $\mathcal{K}_{u-1}$-measurable weights $w_{u,1}, \dots, w_{u,k_u}$ satisfying
$$\sum_{i=1}^{k_u} w_{u,i} = 1 \qquad P\text{-a.s.};$$
4. *measurable betting fractions*
$$\beta_u^\Gamma : \mathbb{R}_+^{u-1} \to [0,1], \qquad u = 1, \dots, m(\mathfrak{E}),$$

such that, up to P-a.s. equivalence,

$$J_u^\Gamma = \sum_{i=1}^{k_u} w_{u,i} \widetilde{E}_{u,i}, \qquad u = 1, \dots, m(\mathfrak{E}), \tag{1}$$

and the global process is updated by

$$\begin{aligned} M_0^\Gamma &= 1, \\ M_u^\Gamma &= M_{u-1}^\Gamma \Big(1 + \beta_u^\Gamma(J_1^\Gamma, \dots, J_{u-1}^\Gamma)(J_u^\Gamma - 1)\Big). \end{aligned} \tag{2}$$

Consequently,

$$\begin{aligned} M_u^\Gamma = \prod_{s=1}^{u} \Big(1 &+ \beta_s^\Gamma(J_1^\Gamma, \dots, J_{s-1}^\Gamma) \\ &\cdot \Big(\sum_{i=1}^{k_s} w_{s,i} \widetilde{E}_{s,i} - 1\Big)\Big). \end{aligned} \tag{3}$$

If, in addition, the pure lift restriction is value-based and universally lifts running maxima of e-processes across filtrations, then each lifted input in Equation (1) may be written as

$$\widetilde{E}_{u,i} = a_{u,i}(E_{u,i}),$$

where $a_{u,i}$ is an adjuster for the corresponding filtration pair. Hence, under this strengthened lift assumption, the calculus factors through exactly three primitives:

1. *a game-martingale serial update;*
2. *weighted arithmetic averaging at same-null joins;*
3. *adjusters across filtration lifts.*

Conversely, every trace-generated calculus built from:

1. *lifts preserving one-step conditional e-validity,*
2. *weighted arithmetic averages at same-null joins,*
3. *a game-martingale serial update on the resulting stage sequence,*

is universally valid. Under the strengthened lift assumption, these lifts may be taken to be adjusters.

Proof sketch. The necessity argument is primitive-by-primitive. First, Theorem 5.1 and the serial-precision assumption reduce the chain restriction of Γ to an adapted, anytime-valid, precise sequential constructor; Vovk and Wang's theorem therefore forces it to be a game martingale, yielding Equation (2) (Vovk & Wang, 2024). Second, the pure same-null join restriction is, by assumption, an admissible arbitrary-dependence e-merging function; Wang's complete-class theorem then forces each stage join to be a weighted arithmetic average, giving Equation (1) (Wang, 2025). Third, if pure lifts are value-based and universally lift running maxima across filtrations, then Choe and Ramdas show that they must be adjusters (Choe & Ramdas, 2026). Combining these three identifications along the trace-generated stage decomposition yields the product representation Equation (3).

For the converse, fix an admissible law P. By local validity, each base local certificate is a one-step conditional e-value relative to its native history. By conditional-lift validity, every lifted stage input $\widetilde{E}_{u,i}$ still satisfies

$$\mathbb{E}_P[\widetilde{E}_{u,i} \,|\, \mathcal{K}_{u-1}] \le 1 \qquad P\text{-a.s.}$$

Weighted averaging therefore preserves one-step conditional validity at the stage level, so $(J_u)_{u \ge 1}$ is a bounded sequence of stagewise sequential e-values. Applying a game-martingale serial update to this stage sequence produces an e-process. Since the entire construction is trace-generated from the execution graph, the resulting process is a universally valid graph calculus. Appendix E contains the full proof. □

Theorem 5.2 is the central structural statement of the paper. It says that, once the local uncertainty layer has collapsed to predictable finite-dimensional certificates as in Theorem 4.1, the global graph problem has essentially no additional degrees of freedom. Universal composability does not permit an arbitrary algebra of certificate manipulations: it forces a serial game-martingale update, a same-null weighted average at joins, and—under the natural value-based liftability condition—adjusters across filtration refinements. The next section quantifies the price of this rigidity.

6. The Price of Universal Composability

The structural theorem identifies the only admissible primitive algebra. A different question is quantitative: how much evidence must be lost because a globally valid calculus is forced to merge by weighted averaging and to cross filtration boundaries by value-based lifts? The cleanest answer is obtained for the *canonical full-betting factorization*, namely the special case of Theorem 5.2 with serial update $\beta_u \equiv 1$, so that

$$M_u^\Gamma = \prod_{s=1}^{u} J_s^\Gamma.$$

For general game-martingale serial updates, stagewise losses are attenuated by the map $x \mapsto \log(1+\beta_u(x-1))$, and exact additivity need not hold. Thus the quantity below should be interpreted as the intrinsic primitive tax of the canonical factorized calculus.

Fix a primitive-expanded full-betting trace $\mathfrak{T} = (J_1^\Gamma, \dots, J_m^\Gamma)$. Let $\mathcal{J}$ and $\mathcal{L}$ denote the sets of join and lift stages, respectively. For a join stage $u \in \mathcal{J}$, write

$$J_u^\Gamma = \sum_{i=1}^{k_u} w_{u,i} E_{u,i}, \qquad w_{u,i} \geq 0, \qquad \sum_{i=1}^{k_u} w_{u,i} = 1,$$

and define the corresponding oracle stage certificate

$$J_u^{\mathrm{or}} := \max_{1 \leq i \leq k_u} E_{u,i}.$$

For a lift stage $u \in \mathcal{L}$, write

$$J_u^\Gamma = a_u(S_u), \qquad J_u^{\mathrm{or}} := S_u,$$

where $S_u \geq 1$ is the running maximum of the underlying coarse-filtration e-process and a_u is the value-based lift. Define the stage taxes

$$\Delta_{\mathrm{join}}(u) := \log \frac{1}{w_{u,\min}}, \qquad w_{u,\min} := \min_{1 \leq i \leq k_u} w_{u,i},$$

and

$$\Delta_{\mathrm{lift}}(u) := \sup_{x \geq 1} \frac{1}{x} \log \frac{x}{a_u(x)}.$$

The total primitive tax of the trace is

$$\mathfrak{C}(\mathfrak{T}) := \sum_{u \in \mathcal{J}} \Delta_{\mathrm{join}}(u) + \sum_{u \in \mathcal{L}} \Delta_{\mathrm{lift}}(u).$$

Proposition 6.1 (Additive composition tax for the canonical factorization)**.** *Let Γ be a factorized calculus in the sense of Theorem 5.2, and assume its serial update is the canonical full-betting rule $M_u^\Gamma = \prod_{s=1}^{u} J_s^\Gamma$. Then for every primitive-expanded full-betting trace $\mathfrak{T}$ and every $\varepsilon > 0$, there exist compatible bounded primitive inputs and a product null law Q^ε such that*

$$\mathbb{E}_{Q^\varepsilon}\left[\log \frac{M_m^{\mathrm{or}}}{M_m^\Gamma}\right] \geq \mathfrak{C}(\mathfrak{T}) - \varepsilon,$$

where

$$M_m^{\mathrm{or}} := \prod_{u=1}^{m} J_u^{\mathrm{or}}.$$

If every join stage is symmetric k_u-ary, then

$$\Delta_{\mathrm{join}}(u) = \log k_u.$$

Proof sketch. The join term is exact. Once Theorem 5.2 forces a same-null join to be a weighted arithmetic average, its worst-case log-loss relative to the branchwise oracle $\max_i E_{u,i}$ is exactly $\log(1/w_{u,\min})$, attained by a one-hot family concentrated on the least-weighted branch (Wang, 2025). The lift term is lower bounded by a two-point one-step martingale construction: for every $\varepsilon > 0$, one can choose a bounded coarse-filtration e-process whose running maximum S_u satisfies

$$\mathbb{E}\left[\log \frac{S_u}{a_u(S_u)}\right] \geq \Delta_{\mathrm{lift}}(u) - \varepsilon$$

for the value-based lift a_u, and such lifts are exactly the adjuster-type objects characterized by (Choe & Ramdas, 2026). Taking the product of these stagewise extremizers over the primitive-expanded trace makes the global log-ratio additive:

$$\log \frac{M_m^{\mathrm{or}}}{M_m^\Gamma} = \sum_{u=1}^{m} \log \frac{J_u^{\mathrm{or}}}{J_u^\Gamma}.$$

Summing the stagewise bounds gives the claim. The full proof is in Appendix F. □

Corollary 6.2 (Exact symmetric-star tax)**.** *If a primitive-expanded full-betting trace contains only symmetric join stages and no lift stages, then*

$$\sup_Q \mathbb{E}_Q\left[\log \frac{M_m^{\mathrm{or}}}{M_m^\Gamma}\right] = \sum_{u=1}^{m} \log k_u,$$

where k_u is the arity of stage u. Thus the join contribution to the composition tax is exact, not merely a lower bound.

Proposition 6.1 and Theorem 6.2 are the quantitative counterparts of Theorem 5.2: universal composability leaves only a small admissible primitive algebra, and that algebra carries an intrinsic log-evidence tax.

References

Angelopoulos, A. N., Bates, S., Fisch, A., Lei, L., and Schuster, T. Conformal risk control. In *International Conference on Learning Representations*, 2024.

Choe, Y. J. and Ramdas, A. Combining evidence across filtrations, 2026. Accepted for publication in Journal of the Royal Statistical Society: Series B.

Fayyazi, A. and Akrami, H. Proof-of-perception: Certified tool-using multimodal reasoning with compositional conformal guarantees, 2026.

Gauthier, E., Bach, F., and Jordan, M. I. E-values expand the scope of conformal prediction, 2025.

Hegazy, M., Aolaritei, L., Jordan, M. I., and Dieuleveut, A. Valid selection among conformal sets, 2025.

Hultberg, B., Zachariah, D., and Ribeiro, A. H. Anytime-valid conformal risk control, 2026.

Larsson, M., Ramdas, A., and Ruf, J. Testing hypotheses generated by constraints, 2026.

Vovk, V. and Wang, R. Merging sequential e-values via martingales, 2024.

Wang, R. The only admissible way of merging arbitrary e-values. *Biometrika*, 112(2):asaf020, 2025. doi: 10.1093/biomet/asaf020.

A. Measure-Theoretic Execution Model

This appendix fixes the ambient probabilistic and measurable structure used throughout the paper. The main body only needs a compressed version of these definitions; the appendix records the full setup needed later for conditional null classes, local e-certificates, filtration lifts, and graph-indexed certificate calculi. A technically important point is that *graph growth* and *certificate manipulation* are distinct layers of the model: the realized execution graph is created purely by serial node additions, whereas same-null joins and filtration lifts are operations on certificate objects attached to an already-created graph. Keeping these two layers separate avoids a category error in later factorization arguments.

A.1. Standing Measurable Structure

We work in discrete time on a fixed finite horizon $T \in \mathbb{N}$, and write $[T] := \{1, \dots, T\}$ and $[T]_0 := \{0, 1, \dots, T\}$. For a measurable space S, we write $\mathcal{M}_1(S)$ for the set of all probability measures on S.

Assumption A.1 (Standard Borel ambient spaces and cemetery conventions)**.** For each $t \in [T]$, the following hold.

1. $(\mathcal{X}_t, \mathcal{X}_t)$ and $(\mathcal{Z}_t, \mathcal{Z}_t)$ are standard Borel spaces.
2. The module-label space $(\mathcal{M}, \mathcal{M})$ is standard Borel.
3. Each space $\mathcal{X}_t$, $\mathcal{Z}_t$, and $\mathcal{M}$ contains a distinguished cemetery point, denoted respectively by $\partial_t^X \in \mathcal{X}_t$, $\partial_t^Z \in \mathcal{Z}_t$, and $\partial^{\mathcal{M}} \in \mathcal{M}$.
4. The parent-set space at time t is the finite set $2^{[t-1]}$ endowed with the discrete sigma-field.

All random variables appearing below are defined on the whole of Ω. After stopping, the controller is dormant: on $\{\tau < t\}$, we impose

$$m_t = \partial^{\mathcal{M}}, \qquad P_t = \varnothing, \qquad X_t = \partial_t^X, \qquad Z_t = \partial_t^Z.$$

Remark A.2 (Why standard Borel spaces)*.* The standard Borel assumption is used only for the standard measure-theoretic operations required later: existence of regular conditional laws, measurable parameterizations of local null classes, and measurable selections in the proof of the local representation theorem. No topological structure beyond this is used in the workshop version.

A.2. Execution Models and Realized Graphs

Definition A.3 (Adaptive execution graph)**.** An *adaptive execution graph* consists of a tuple

$$\mathfrak{E} = \Big(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_t)_{t=0}^T, \tau, (m_t, P_t, X_t, Z_t)_{t=1}^T\Big)$$

such that:

1. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space and $(\mathcal{F}_t)_{t=0}^T$ is a filtration.
2. τ is an $(\mathcal{F}_t)$-stopping time with $0 \le \tau \le T$.
3. For each $t \in [T]$, the triple (m_t, P_t, X_t) is $\mathcal{F}_{t-1}$-measurable, where
$$m_t \in \mathcal{M}, \qquad P_t \in 2^{[t-1]}, \qquad X_t \in \mathcal{X}_t.$$
4. For each $t \in [T]$, the output $Z_t \in \mathcal{Z}_t$ is $\mathcal{F}_t$-measurable.

The local information sigma-field at round t is

$$\mathcal{H}_t := \sigma(X_t) \subseteq \mathcal{F}_{t-1}.$$

The realized graph at time $u \in [T]_0$ is

$$G_u = (V_u, E_u), \qquad V_u := [u \wedge \tau], \qquad E_u := \{(s,t) : t \le u \wedge \tau,\ s \in P_t\}.$$

In particular, the terminal realized graph is G_τ.

Remark A.4 (Predictable state variables). The state variable X_t is allowed to be arbitrarily rich. In particular, X_t may encode any predictable summary of the past that the controller or the local module wishes to use, including the current module label m_t, the candidate parent set P_t, previously observed outputs, or any predictable graph statistic. Thus the restriction $\mathcal{H}_t = \sigma(X_t)$ is not a loss of generality for the finite-horizon discrete-time setting considered here.

Definition A.5 (Ancestral closure)**.** Let $u \in [T]_0$, and let $S \subseteq V_u$. The *ancestral closure* of S in G_u, denoted $\mathrm{An}_{G_u}(S)$, is the smallest subset $A \subseteq V_u$ such that $S \subseteq A$ and

$$t \in A \quad \Longrightarrow \quad P_t \subseteq A.$$

A subset $A \subseteq V_u$ is called *ancestral* if $\mathrm{An}_{G_u}(A) = A$.

Lemma A.6 (Serial generation of the realized DAG)**.** *Let $\mathfrak{E}$ be an adaptive execution graph. Then the graph process $(G_u)_{u=0}^{T}$ is uniquely generated by the recursion*

$$G_0 = (\varnothing, \varnothing)$$

and, for $t \in [T]$,

$$G_t = \begin{cases} G_{t-1}, & \tau < t, \\ \big(V_{t-1} \cup \{t\}, E_{t-1} \cup \{(s,t) : s \in P_t\}\big), & \tau \geq t. \end{cases}$$

In particular, graph growth itself is purely serial: the terminal DAG G_τ is completely determined by the finite sequence of predictable parent sets $(P_t)_{t \leq \tau}$.

Proof. By construction, $V_t = [t \wedge \tau]$. If $\tau < t$, then no new vertex is present at round t, so $G_t = G_{t-1}$. If $\tau \geq t$, then the only new vertex is t, and all edges incident to t are exactly the edges (s,t) with $s \in P_t \subseteq [t-1]$. Since every edge points from a smaller label to a larger label, the graph is acyclic. The recursion therefore constructs a unique random DAG with the stated vertex and edge sets. Uniqueness is immediate because the vertex set and the parent set of each realized node are explicitly prescribed. □

Remark A.7 (Graph growth versus certificate operations). Lemma A.6 concerns only the *creation* of the random DAG. Later same-null joins and filtration lifts will act on certificate objects attached to already-created nodes; they do *not* modify the vertex set or edge set of the underlying execution graph. This separation is one of the main technical differences between the present framework and existing executable-graph uncertainty formalisms based on stepwise nodewise certification (Fayyazi & Akrami, 2026).

A.3. Local Null Classes and Local E-Certificates

For each $t \in [T]$, fix a finite number $r_t \in \mathbb{N}$ and bounded measurable functions

$$\phi_{t,j} : \mathcal{X}_t \times \mathcal{Z}_t \to \mathbb{R}, \qquad j = 1, \ldots, r_t.$$

These functions define the local null class at round t. We impose the normalization

$$\phi_{t,j}(\partial_t^X, \partial_t^Z) = 0, \qquad j = 1, \ldots, r_t,$$

so that the stopped/cemetery state is automatically feasible.

Definition A.8 (Local null class generated by predictable constraints)**.** For each $t \in [T]$ and $x \in \mathcal{X}_t$, define

$$\mathcal{P}_t(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \int \phi_{t,j}(x,z)\, Q(dz) \leq 0 \text{ for all } j = 1, \ldots, r_t \right\}.$$

For the cemetery state $x = \partial_t^X$, we additionally require

$$\mathcal{P}_t(\partial_t^X) = \{\delta_{\partial_t^Z}\}.$$

Remark A.9 (Constraint-generated local nulls). The map $x \mapsto \mathcal{P}_t(x)$ is the predictable, history-indexed analogue of the constraint-generated hypothesis classes studied in the general theory of e-variables under measurable constraints (Larsson et al., 2026). In the present paper the history variable is the predictable state X_t, and the null class varies fiberwise over the state space.

Definition A.10 (Admissible law)**.** A probability law P on $(\Omega, \mathcal{F})$ is *admissible* for the execution model $\mathfrak{E}$ if, for every $t \in [T]$, the regular conditional distribution of Z_t given X_t under P, denoted $K_t^P(\cdot \mid X_t)$, satisfies

$$K_t^P(\cdot \mid X_t) \in \mathcal{P}_t(X_t) \qquad P\text{-a.s.}$$

Equivalently, every realized round respects the local predictable constraint system encoded by $\mathcal{P}_t(X_t)$.

Lemma A.11 (Kernel and conditional-moment forms of admissibility)**.** *Under Assumption A.1, a law P is admissible if and only if, for every $t \in [T]$ and every $j = 1, \ldots, r_t$,*

$$\mathbb{E}_P\big[\phi_{t,j}(X_t, Z_t) \mid X_t\big] \le 0 \qquad P\text{-a.s.}$$

Proof. Because $\mathcal{Z}_t$ is standard Borel, a regular conditional distribution $K_t^P(\cdot \mid X_t)$ of Z_t given X_t exists. Since each $\phi_{t,j}$ is bounded and measurable, the conditional expectation may be computed by integrating against that kernel:

$$\mathbb{E}_P\big[\phi_{t,j}(X_t, Z_t) \mid X_t\big] = \int \phi_{t,j}(X_t, z)\, K_t^P(dz \mid X_t) \qquad P\text{-a.s.}$$

If P is admissible, then by Definition A.10, $K_t^P(\cdot \mid X_t) \in \mathcal{P}_t(X_t)$ almost surely, and hence the right side is at most 0 almost surely for every j. This proves the forward implication.

Conversely, assume the conditional inequalities hold almost surely. Then for every j,

$$\int \phi_{t,j}(X_t, z)\, K_t^P(dz \mid X_t) = \mathbb{E}_P\big[\phi_{t,j}(X_t, Z_t) \mid X_t\big] \le 0 \qquad P\text{-a.s.}$$

Therefore $K_t^P(\cdot \mid X_t) \in \mathcal{P}_t(X_t)$ almost surely, which is exactly admissibility. □

Definition A.12 (Local e-certificate)**.** A *local e-certificate* at round t is a measurable function

$$e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0, \infty)$$

such that:

1. for every $x \in \mathcal{X}_t$ and every $Q \in \mathcal{P}_t(x)$,
$$\int e_t(x, z)\, Q(dz) \le 1;$$
2. the cemetery value is normalized by
$$e_t(\partial_t^X, \partial_t^Z) = 1.$$

Its realized value is

$$E_t := e_t(X_t, Z_t).$$

Proposition A.13 (Conditional validity of local e-certificates)**.** *Let e_t be a local e-certificate and let $E_t = e_t(X_t, Z_t)$. Then, under every admissible law P,*

$$\mathbb{E}_P[E_t \mid X_t] \le 1 \qquad P\text{-a.s.}$$

In particular, $E_t = 1$ on $\{\tau < t\}$, so stopped rounds contribute the multiplicative identity.

Proof. Fix an admissible law P. By Definition A.10, the regular conditional law $K_t^P(\cdot \mid X_t)$ belongs to $\mathcal{P}_t(X_t)$ almost surely. Therefore,

$$\mathbb{E}_P[E_t \mid X_t] = \int e_t(X_t, z)\, K_t^P(dz \mid X_t) \le 1 \qquad P\text{-a.s.}$$

The cemetery claim follows from the conventions $X_t = \partial_t^X$, $Z_t = \partial_t^Z$ on $\{\tau < t\}$ together with the normalization $e_t(\partial_t^X, \partial_t^Z) = 1$. □

Definition A.14 (Pointwise admissibility of local e-functions)**.** Fix $t \in [T]$ and $x \in \mathcal{X}_t$. Let

$$\mathcal{E}_t(x) := \left\{ e : \mathcal{Z}_t \to [0, \infty) : \int e(z)\, Q(dz) \le 1 \text{ for every } Q \in \mathcal{P}_t(x) \right\}.$$

An element $e \in \mathcal{E}_t(x)$ is *pointwise admissible* if there does not exist $\tilde{e} \in \mathcal{E}_t(x)$ such that

$$\tilde{e}(z) \ge e(z) \quad \text{for all } z \in \mathcal{Z}_t, \qquad \tilde{e}(z_0) > e(z_0) \quad \text{for some } z_0 \in \mathcal{Z}_t.$$

We denote by $\mathrm{Max}(\mathcal{E}_t(x))$ the set of all pointwise admissible elements of $\mathcal{E}_t(x)$.

Remark A.15 (Why pointwise admissibility is frozen now). Definition A.14 is included already at the level of the execution model because the local representation theorem in Appendix C will characterize $\mathrm{Max}(\mathcal{E}_t(x))$ for the predictable constraint-generated null classes above. Freezing this notion here prevents any later ambiguity about what "admissible local certificate" means.

A.4. Global E-Processes, Same-Null Families, and Lift Operators

Definition A.16 (Global e-process)**.** Let P be an admissible law. A nonnegative $(\mathcal{F}_t)$-adapted process $M = (M_t)_{t=0}^T$ with $M_0 = 1$ is an *e-process under* P if

$$\mathbb{E}_P[M_t \,|\, \mathcal{F}_s] \le M_s \qquad P\text{-a.s. for all } 0 \le s \le t \le T.$$

Definition A.17 (Certificate calculus)**.** A *certificate calculus* is a rule Γ that maps:

1. an adaptive execution graph $\mathfrak{E}$, and
2. a family of local e-certificates $(e_t)_{t=1}^T$,

to a nonnegative $(\mathcal{F}_t)$-adapted process

$$M^\Gamma = (M_t^\Gamma)_{t=0}^T$$

with $M_0^\Gamma = 1$. We say that Γ is *universally valid* if, for every admissible law P, the process M^Γ is an e-process under P.

Definition A.18 (Global e-value for a common null)**.** Let $\mathcal{G} \subseteq \mathcal{F}$ be a sub-sigma-field, and let $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$ be a null class. A nonnegative $\mathcal{G}$-measurable random variable E is called a *$\mathcal{Q}$-e-value on $\mathcal{G}$* if

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[E] \le 1.$$

Definition A.19 (Lift operator)**.** Let $\mathcal{G} \subseteq \mathcal{G}' \subseteq \mathcal{F}$, and let $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$ be a null class. A map

$$A : L_+^0(\Omega, \mathcal{G}) \to L_+^0(\Omega, \mathcal{G}')$$

is called a *$(\mathcal{G}, \mathcal{G}'; \mathcal{Q})$-lift operator* if

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[A(E)] \le 1 \qquad \text{whenever} \qquad \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[E] \le 1.$$

Thus a lift operator transports e-values from a coarser sigma-field to a finer one without losing validity for the same null class.

Remark A.20 (Relation to adjusters). Definition A.19 is intentionally abstract. The later factorization theorem will identify the only admissible filtration-lift primitive with an adjuster-type operation in the sense of (Choe & Ramdas, 2026). At the level of the execution model, however, it is important to keep the notion of lift operator independent of any particular formula.

Definition A.21 (Same-null join)**.** Let J be a finite index set. A family

$$\big(E_i, \mathcal{G}_i\big)_{i \in J}, \qquad E_i \in L_+^0(\Omega, \mathcal{G}_i),$$

is called a *same-null family* if there exist:

1. a null class $\mathcal{Q}_J \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$,
2. a common comparison sigma-field $\mathcal{G}_J \subseteq \mathcal{F}$ with $\mathcal{G}_i \subseteq \mathcal{G}_J$ for every $i \in J$,
3. lift operators
$$A_i : L^0_+(\Omega, \mathcal{G}_i) \to L^0_+(\Omega, \mathcal{G}_J), \qquad i \in J,$$

such that every lifted variable $A_i(E_i)$ is a $\mathcal{Q}_J$-e-value on $\mathcal{G}_J$. A *same-null join* is any operation used by a certificate calculus to combine such a lifted family into a single certificate object on $\mathcal{G}_J$.

Remark A.22 (Why the join is kept abstract in Appendix A). At the level of the ambient model, a same-null join is defined only by the common-null requirement. The main theorem later proves that once one imposes universal composability and the structural axioms from the main text, the join primitive is forced to be weighted arithmetic averaging; that conclusion is a theorem, not a definition. For this reason, the present appendix keeps the join operation abstract.

A.5. Certificate Evaluation Traces

Definition A.23 (Certificate evaluation trace)**.** Fix an adaptive execution graph $\mathfrak{E}$ and realized local e-certificates $(E_t)_{t=1}^{\tau}$. A *certificate evaluation trace* is a finite formal expression generated by the following constructors:

1. **Serial insertion:** a realized nodewise local certificate E_t is inserted when node t is created.
2. **Lift:** a certificate E measurable with respect to a sigma-field $\mathcal{G}$ is replaced by $A(E)$, where A is an admissible lift operator from $\mathcal{G}$ to a finer sigma-field $\mathcal{G}'$.
3. **Join:** a finite same-null family of lifted certificates is combined into a single certificate object.

The trace is *compatible* with the execution graph if every serial insertion respects the temporal order of node creation and every lift/join step acts only on certificate objects associated with already-created nodes.

Remark A.24 (What the later factorization theorem will classify). Definition A.23 isolates the algebraic objects that the later factorization theorem will classify. The graph itself is already fixed by serial growth; what remains nontrivial is which lifts and joins can be composed with serial insertions to produce a universally valid global e-process. This formal separation is exactly what allows us to reduce the global theorem to the three primitive characterization results from the e-value literature (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026).

Remark A.25 (Positioning relative to adjacent frameworks). The present appendix is intentionally more austere than a task-specific or benchmark-specific formulation. Local null classes are modeled through predictable constraint systems, as in the abstract e-variable framework of (Larsson et al., 2026); filtration changes are abstracted as lift operators, anticipating the adjuster formalism of (Choe & Ramdas, 2026); and graph execution is separated from certificate manipulation, unlike nodewise executable-graph certification frameworks such as (Fayyazi & Akrami, 2026). This separation is precisely what makes a universal graph-level characterization theorem possible.

B. Full Proof of Router Non-Composability

This appendix proves the negative result from Section 3. We temporarily step outside the graph-indexed e-certificate formalism of Section 2 and consider ordinary local certificates judged only by marginal coverage or expected risk. The goal is to show that such local guarantees are *not closed* under predictable routing, even in the smallest nontrivial adaptive execution graph. The construction below is finite, exact, and completely self-contained. It should be viewed as the routing analogue of the post-hoc selection pathology for conformal sets studied in (Hegazy et al., 2025), but no external result is needed for the proof.

Lemma B.1 (Selection-as-routing embedding)**.** *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let W be an observed random variable taking values in a measurable space $(\mathcal{W}, \mathcal{B}_{\mathcal{W}})$, let $(\mathcal{U}, \mathcal{B}_{\mathcal{U}})$ be a standard Borel space, and let*

$$U_1, \ldots, U_m : \Omega \to \mathcal{U}, \qquad R : \Omega \to [m]$$

be $\sigma(W)$-measurable. Then there exists an adaptive execution graph $\mathfrak{E}$ of horizon $m + 2$ such that:

1. *node* 1 *outputs* W*;*
2. *for each* $i \in [m]$*, node* $i+1$ *outputs* U_i*;*
3. *node* $m+2$ *outputs the selected object* U_R*;*
4. *the realized graph has depth at most* 2*.*

More precisely, one may take

$$P_1 = \varnothing, \qquad P_{i+1} = \{1\} \text{ for } i \in [m], \qquad P_{m+2} = \{1, \dots, m+1\},$$

and $\tau = m+2$ *deterministic.*

Proof. Set the horizon $T := m+2$, and define

$$\mathcal{F}_0 := \{\varnothing, \Omega\}, \qquad \mathcal{F}_t := \sigma(W) \quad \text{for every } t \in [T].$$

Then $(\mathcal{F}_t)_{t=0}^T$ is a filtration. Since $U_1, \dots, U_m$ and R are all $\sigma(W)$-measurable, they are $\mathcal{F}_t$-measurable for every $t \geq 1$. Let $\tau := T$.

We now instantiate the measurable spaces in Theorem 2.1. Take

$$\mathcal{X}_1 := \{\star\}, \qquad \mathcal{Z}_1 := \mathcal{W},$$

and for $i \in [m]$,

$$\mathcal{X}_{i+1} := \mathcal{W}, \qquad \mathcal{Z}_{i+1} := \mathcal{U},$$

while for the router node set

$$\mathcal{X}_{m+2} := \mathcal{W}, \qquad \mathcal{Z}_{m+2} := \mathcal{U}.$$

Choose arbitrary deterministic module labels $m_1, \dots, m_{m+2} \in \mathcal{M}$. Since $\tau = T$ deterministically, the cemetery conventions of Appendix A never become active.

Root node. At round 1, define

$$P_1 := \varnothing, \qquad X_1 := \star, \qquad Z_1 := W.$$

Then P_1 and X_1 are $\mathcal{F}_0$-measurable, while Z_1 is $\mathcal{F}_1$-measurable.

Candidate-output nodes. For each $i \in [m]$, define at round $i+1$,

$$P_{i+1} := \{1\}, \qquad X_{i+1} := W, \qquad Z_{i+1} := U_i.$$

Since W and U_i are $\sigma(W) = \mathcal{F}_i$-measurable and P_{i+1} is deterministic, the required measurability conditions hold.

Router node. At the final round, define

$$P_{m+2} := \{1, \dots, m+1\}, \qquad X_{m+2} := W, \qquad Z_{m+2} := U_R.$$

It remains to verify measurability of U_R. For every $B \in \mathcal{B}_{\mathcal{U}}$,

$$\{U_R \in B\} = \bigcup_{i=1}^{m} \big(\{R = i\} \cap \{U_i \in B\}\big),$$

which belongs to $\sigma(W)$ because R and each U_i are $\sigma(W)$-measurable. Hence U_R is $\mathcal{F}_{m+2}$-measurable.

Thus $\mathfrak{E}$ is an adaptive execution graph of horizon $m+2$. Its realized DAG has one root node, m children attached to the root, and one final router node attached to all previous nodes. The longest directed path has the form

$$1 \longrightarrow i+1 \longrightarrow m+2,$$

so the graph has depth at most 2. □

Proof of Theorem 3.1. Fix $\alpha \in (0,1)$, and define

$$p := \frac{1}{2}\min\{\alpha, 1-\alpha\}, \qquad q := \alpha - p, \qquad r := 1 - \alpha - p.$$

We first check that p, q, r determine a probability distribution. Because $\alpha \in (0,1)$, we have $p > 0$. Since $p \le \alpha$,

$$q = \alpha - p \ge 0,$$

and since $p \le 1 - \alpha$,

$$r = 1 - \alpha - p \ge 0.$$

Moreover,

$$p + p + q + r = 2p + (\alpha - p) + (1 - \alpha - p) = 1.$$

Now define

$$\mathcal{W} = \mathcal{Y} := \{0, 1, 2, 3\},$$

and specify the law of W by

$$\mathbb{P}(W = 0) = r, \qquad \mathbb{P}(W = 1) = p, \qquad \mathbb{P}(W = 2) = p, \qquad \mathbb{P}(W = 3) = q.$$

Set

$$Y := W.$$

Thus Y is a deterministic function of W, so every coverage event can be computed explicitly.

Next define the two failure regions

$$A_1 := \{1, 3\}, \qquad A_2 := \{2, 3\}.$$

Define set-valued predictors $C_1, C_2 : \mathcal{W} \to \mathcal{P}(\mathcal{Y})$ by

$$C_1(w) := \begin{cases} \varnothing, & w \in A_1, \\ \{w\}, & w \notin A_1, \end{cases} \qquad C_2(w) := \begin{cases} \varnothing, & w \in A_2, \\ \{w\}, & w \notin A_2. \end{cases}$$

Because $\mathcal{W}$ and $\mathcal{Y}$ are finite, no measurability subtlety arises.

We compute the individual coverages. Since $Y = W$,

$$\{Y \in C_1(W)\} = \{W \notin A_1\} = \{W \in \{0, 2\}\},$$

and therefore

$$\mathbb{P}\{Y \in C_1(W)\} = \mathbb{P}(W = 0) + \mathbb{P}(W = 2) = r + p = (1 - \alpha - p) + p = 1 - \alpha.$$

Similarly,

$$\{Y \in C_2(W)\} = \{W \notin A_2\} = \{W \in \{0, 1\}\},$$

so

$$\mathbb{P}\{Y \in C_2(W)\} = \mathbb{P}(W = 0) + \mathbb{P}(W = 1) = r + p = 1 - \alpha.$$

Now define the router $R : \mathcal{W} \to \{1, 2\}$ by

$$R(w) := \begin{cases} 1, & w \in \{0, 1, 3\}, \\ 2, & w = 2. \end{cases}$$

This router is measurable because $\mathcal{W}$ is finite. By construction,

$$W = 1 \implies R(W) = 1 \text{ and } W \in A_1,$$

$$W = 2 \implies R(W) = 2 \text{ and } W \in A_2,$$

$$W = 3 \implies R(W) = 1 \text{ and } W \in A_1.$$

Hence on every state $w \in \{1, 2, 3\}$, the router deliberately chooses a predictor that fails at w. On the remaining state $w = 0$, both predictors succeed and the router chooses 1.

Therefore

$$C_{R(W)}(W) = \begin{cases} \{0\}, & W = 0, \\ \varnothing, & W \in \{1, 2, 3\}. \end{cases}$$

Since $Y = W$, we have

$$\{Y \in C_{R(W)}(W)\} = \{W = 0\},$$

and hence

$$\mathbb{P}\{Y \in C_{R(W)}(W)\} = \mathbb{P}(W = 0) = r = 1 - \alpha - p = 1 - \alpha - \frac{1}{2}\min\{\alpha, 1 - \alpha\}.$$

Because $p > 0$, this is strictly smaller than $1 - \alpha$.

Finally, Theorem B.1 embeds the family $\big(C_1(W), C_2(W), R(W)\big)$ into a depth-2 adaptive execution graph: one root node reveals W, two child modules output $C_1(W)$ and $C_2(W)$, and a final router node outputs $C_{R(W)}(W)$. This proves Theorem 3.1. □

Lemma B.2 (Risk inflation under predictable routing)**.** *Under the construction in Theorem 3.1, define bounded losses*

$$\ell_i(w, y) := \mathbf{1}\{y \notin C_i(w)\}, \qquad i = 1, 2.$$

Then

$$\mathbb{E}[\ell_i(W, Y)] = \alpha, \qquad i = 1, 2,$$

but the routed loss satisfies

$$\mathbb{E}[\ell_{R(W)}(W, Y)] = \alpha + \frac{1}{2}\min\{\alpha, 1 - \alpha\} > \alpha.$$

Proof. Because $Y = W$,

$$\ell_1(W, Y) = \mathbf{1}\{W \in A_1\}, \qquad \ell_2(W, Y) = \mathbf{1}\{W \in A_2\}.$$

Therefore,

$$\mathbb{E}[\ell_1(W, Y)] = \mathbb{P}(W \in A_1) = p + q = \alpha,$$

and similarly,

$$\mathbb{E}[\ell_2(W, Y)] = \mathbb{P}(W \in A_2) = p + q = \alpha.$$

For the routed loss, Theorem 3.1 shows that the routed predictor fails exactly on $\{1, 2, 3\}$. Hence

$$\ell_{R(W)}(W, Y) = \mathbf{1}\{Y \notin C_{R(W)}(W)\} = \mathbf{1}\{W \in \{1, 2, 3\}\},$$

so

$$\mathbb{E}[\ell_{R(W)}(W, Y)] = \mathbb{P}(W \in \{1, 2, 3\}) = 1 - r = \alpha + p = \alpha + \frac{1}{2}\min\{\alpha, 1 - \alpha\}.$$

Since $p > 0$, this exceeds α. □

Proof of Corollary 3.2. The first claim follows directly from Theorem 3.1. Indeed, consider the depth-2 execution graph furnished by Theorem B.1, with two child modules outputting $C_1(W)$ and $C_2(W)$ and a final router node outputting $C_{R(W)}(W)$. Each child module carries an exact local marginal coverage guarantee at level $1 - \alpha$, but the routed output violates that same guarantee. Hence no composition rule whose only local premises are such marginal-coverage statements can be universally valid under predictable routing.

The expected-risk assertion is exactly Theorem B.2. There, each local loss has expectation α, while the routed loss has strictly larger expectation. Therefore local expected-risk guarantees are likewise not closed under predictable routing. This proves Corollary 3.2. □

Remark B.3 (Interpretation). The obstruction proved above is neither asymptotic nor high-dimensional. It already appears on a four-point probability space, with two exact local guarantees and a deterministic one-step router. Thus the failure is *structural*: once a workflow may inspect history and route among locally valid objects, marginal coverage and expected-risk guarantees cease to be compositional primitives. This is precisely why the positive theory in the remainder of the paper is built on e-certificates rather than on local coverage or local risk statements.

Remark B.4 (Relation to conformal-selection pathologies). The present construction is fully self-contained and does not require any external theorem. It is nonetheless consistent with the broader phenomenon identified in (Hegazy et al., 2025): even when several candidate conformal sets are individually valid, selecting among them can invalidate coverage. The contribution of this appendix is to show that the same pathology already manifests itself at the level of adaptive execution graphs and predictable routing.

C. Fiberwise Duality and Predictable Local Representation

This appendix develops the local representation theory used later in the main factorization theorem. Fix a round $t \in [T]$ throughout, and suppress the index t whenever no confusion can arise. The setup of Sections A and 2 associates to each local state $x \in \mathcal{X}_t$ a constraint-generated null class

$$\mathcal{P}_t(x) = \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \int \phi_{t,j}(x,z)\, Q(dz) \le 0, \quad j = 1, \dots, r_t \right\},$$

and a class $\mathcal{E}_t(x)$ of valid local e-functions. The goal of the present appendix is to show that, fiberwise in x, the relevant local certificates collapse to a finite-dimensional cone generated by the predictable constraint functions. This is the precise point at which the abstract theory of constraint-generated e-variables enters: Larsson, Ramdas, and Ruf prove that for finitely many measurable constraints, every e-variable is dominated (in the appropriate quasi-sure order) by a linear combination of the constraints plus one, and they identify the admissible ones within that class (Larsson et al., 2026). Our contribution here is to adapt that finite-constraint description to the history-indexed local model and then prove a predictable measurable-selection statement along any admissible law.

C.1. Fiberwise Quasi-Sure Order and Admissibility

Fix $x \in \mathcal{X}_t$. A set $N \in \mathcal{Z}_t$ is called $\mathcal{P}_t(x)$-*polar* if

$$Q(N) = 0 \qquad \text{for every } Q \in \mathcal{P}_t(x).$$

For measurable functions $f, g : \mathcal{Z}_t \to [0, \infty]$, we write

$$f \preceq_x g$$

if $f \le g$ outside a $\mathcal{P}_t(x)$-polar set, and

$$f = g \quad \mathcal{P}_t(x)\text{-q.s.}$$

if both $f \preceq_x g$ and $g \preceq_x f$. We also write

$$f \prec_x g$$

if $f \preceq_x g$ but $g \not\preceq_x f$, equivalently if the set $\{z : g(z) > f(z)\}$ is not $\mathcal{P}_t(x)$-polar.

Definition C.1 (Fiberwise quasi-sure admissibility)**.** Let $x \in \mathcal{X}_t$ and $e \in \mathcal{E}_t(x)$. We call e $\mathcal{P}_t(x)$-*quasi-surely admissible* if there is no $\tilde{e} \in \mathcal{E}_t(x)$ such that

$$e \prec_x \tilde{e}.$$

Lemma C.2 (Elementary structure of the local e-class)**.** *Fix $x \in \mathcal{X}_t$. Then $\mathcal{E}_t(x)$ is convex and solid: if $e, \tilde{e} \in \mathcal{E}_t(x)$ and $\theta \in [0,1]$, then $\theta e + (1-\theta)\tilde{e} \in \mathcal{E}_t(x)$; and if $0 \le h \le e$ pointwise with $e \in \mathcal{E}_t(x)$, then $h \in \mathcal{E}_t(x)$.*

Proof. Convexity is immediate from linearity of integration: for any $Q \in \mathcal{P}_t(x)$,

$$\int \big(\theta e + (1-\theta)\tilde{e}\big)\, dQ = \theta \int e\, dQ + (1-\theta) \int \tilde{e}\, dQ \le \theta + (1-\theta) = 1.$$

Similarly, if $0 \le h \le e$, then

$$\int h\, dQ \le \int e\, dQ \le 1 \qquad \forall Q \in \mathcal{P}_t(x),$$

so $h \in \mathcal{E}_t(x)$. □

Proposition C.3 (Pointwise admissibility implies quasi-sure admissibility)**.** *Fix $x \in \mathcal{X}_t$. Let $e \in \mathcal{E}_t(x)$. If e is pointwise admissible in the sense of Theorem A.14, then e is $\mathcal{P}_t(x)$-quasi-surely admissible in the sense of Theorem C.1.*

Proof. Assume for contradiction that e is pointwise admissible but not $\mathcal{P}_t(x)$-quasi-surely admissible. Then there exists $\tilde{e} \in \mathcal{E}_t(x)$ such that

$$e \prec_x \tilde{e}.$$

In particular,

$$N := \{z : \tilde{e}(z) < e(z)\}$$

is $\mathcal{P}_t(x)$-polar, while

$$B := \{z : \tilde{e}(z) > e(z)\}$$

is not $\mathcal{P}_t(x)$-polar and therefore nonempty.

Now define

$$\hat{e} := e \vee \tilde{e}.$$

Then $\hat{e} \ge e$ pointwise and $\hat{e}(z) > e(z)$ for every $z \in B$, so e is strictly pointwise dominated by $\hat{e}$. It remains to show that $\hat{e} \in \mathcal{E}_t(x)$. Since $\hat{e} = \tilde{e}$ outside the polar set N, we have

$$\int \hat{e}\, dQ = \int \tilde{e}\, dQ \qquad \forall Q \in \mathcal{P}_t(x),$$

because every $Q \in \mathcal{P}_t(x)$ charges N with probability zero. Hence

$$\int \hat{e}\, dQ \le 1 \qquad \forall Q \in \mathcal{P}_t(x),$$

so $\hat{e} \in \mathcal{E}_t(x)$. This contradicts pointwise admissibility of e. □

Remark C.4 (Why quasi-sure admissibility is the correct local notion). The abstract theory of e-variables is naturally formulated modulo the null sets of the entire null class, not modulo pointwise equality. Proposition C.3 shows that every pointwise-admissible local certificate is automatically quasi-surely admissible, so the quasi-sure theory below already covers the pointwise-admissible objects used in the main text.

C.2. Fiberwise Linear Complete Class

For $x \in \mathcal{X}_t$, write

$$\phi^x(z) := \big(\phi_{t,1}(x,z), \dots, \phi_{t,r_t}(x,z)\big) \in \mathbb{R}^{r_t}.$$

For $\lambda = (\lambda_1, \dots, \lambda_{r_t}) \in \mathbb{R}_+^{r_t}$, define the associated affine constraint combination

$$H_x^\lambda(z) := 1 + \sum_{j=1}^{r_t} \lambda_j \phi_{t,j}(x,z) = 1 + \lambda^\top \phi^x(z).$$

Also define the fiberwise feasible coefficient set

$$\Pi_t(x) := \left\{\lambda \in \mathbb{R}_+^{r_t} : H_x^\lambda \ge 0 \quad \mathcal{P}_t(x)\text{-q.s.}\right\}.$$

If $\lambda \in \Pi_t(x)$, then $H_x^\lambda \in \mathcal{E}_t(x)$, because H_x^λ is nonnegative quasi-surely and

$$\int H_x^\lambda(z)\, Q(dz) = 1 + \sum_{j=1}^{r_t} \lambda_j \int \phi_{t,j}(x,z)\, Q(dz) \le 1 \qquad \forall Q \in \mathcal{P}_t(x).$$

Lemma C.5 (Geometry of the coefficient set)**.** *For each $x \in \mathcal{X}_t$, the set $\Pi_t(x) \subseteq \mathbb{R}_+^{r_t}$ is convex and closed.*

Proof. Convexity is immediate because

$$H_x^{\theta\lambda+(1-\theta)\lambda'} = \theta H_x^{\lambda} + (1-\theta) H_x^{\lambda'}$$

for $\theta \in [0,1]$, and the quasi-sure nonnegativity relation is preserved under convex combinations.

To prove closedness, let $\lambda_n \in \Pi_t(x)$ and suppose $\lambda_n \to \lambda \in \mathbb{R}_+^{r_t}$. Fix $z \in \mathcal{Z}_t$. Since the constraint family is finite,

$$H_x^{\lambda_n}(z) \to H_x^{\lambda}(z).$$

Let

$$A := \{z : H_x^{\lambda}(z) < 0\}.$$

For each $m \in \mathbb{N}$, define

$$A_m := \left\{ z : H_x^{\lambda}(z) \le -\frac{1}{m} \right\}.$$

Then $A = \bigcup_{m=1}^{\infty} A_m$. Fix m. If $z \in A_m$, then for all sufficiently large n,

$$H_x^{\lambda_n}(z) < -\frac{1}{2m} < 0.$$

Thus

$$A_m \subseteq \bigcup_{N=1}^{\infty} \bigcap_{n \ge N} \{z : H_x^{\lambda_n}(z) < 0\}.$$

Each set $\{H_x^{\lambda_n} < 0\}$ is $\mathcal{P}_t(x)$-polar because $\lambda_n \in \Pi_t(x)$, so the countable union on the right is $\mathcal{P}_t(x)$-polar. Therefore every A_m is $\mathcal{P}_t(x)$-polar, and so A is $\mathcal{P}_t(x)$-polar. Hence $\lambda \in \Pi_t(x)$. □

Assumption C.6 (Fiberwise constraint qualification)**.** Fix $x \in \mathcal{X}_t$. We say that the *fiberwise constraint qualification* holds at x if

$$\lambda, \lambda' \in \mathbb{R}_+^{r_t}, \qquad H_x^{\lambda} \preceq_x H_x^{\lambda'} \quad \Longrightarrow \quad H_x^{\lambda} = H_x^{\lambda'} \ \ \mathcal{P}_t(x)\text{-q.s.}$$

Theorem C.7 (Fiberwise linear complete class)**.** *Fix $x \in \mathcal{X}_t$, and assume $\mathcal{P}_t(x) \neq \varnothing$. Then:*

1. *For every $e \in \mathcal{E}_t(x)$, there exists $\lambda \in \Pi_t(x)$ such that*

   $$e \preceq_x H_x^{\lambda}.$$

   In other words, the family

   $$\{H_x^{\lambda} : \lambda \in \Pi_t(x)\}$$

   is a complete class for the fiberwise local e-problem at x.

2. *Every $\mathcal{P}_t(x)$-quasi-surely admissible $e \in \mathcal{E}_t(x)$ is $\mathcal{P}_t(x)$-q.s. equal to H_x^{λ} for some $\lambda \in \Pi_t(x)$.*

3. *If the fiberwise constraint qualification Theorem C.6 holds at x, then every H_x^{λ} with $\lambda \in \Pi_t(x)$ is $\mathcal{P}_t(x)$-quasi-surely admissible.*

Proof. Fix $x \in \mathcal{X}_t$. Consider the measurable space $(\mathcal{Z}_t, \mathcal{Z}_t)$, the finite collection of measurable constraint functions

$$\Phi_x := \{\phi_{t,1}(x,\cdot), \ldots, \phi_{t,r_t}(x,\cdot)\},$$

and the hypothesis class $\mathcal{P}_t(x)$ generated by those constraints. Since Φ_x is finite and each $\phi_{t,j}(x,\cdot)$ is bounded and measurable, the finite-constraint theory of (Larsson et al., 2026) applies directly on the fiber $(\mathcal{Z}_t, \mathcal{Z}_t)$. Their complete-class result implies that every e-variable for $\mathcal{P}_t(x)$ is $\mathcal{P}_t(x)$-q.s. dominated by an e-variable of the form

$$1 + \sum_{j=1}^{r_t} \lambda_j \phi_{t,j}(x,\cdot)$$

for some $\lambda \in \mathbb{R}_+^{r_t}$. The requirement that this affine function be a valid e-variable is exactly the condition $\lambda \in \Pi_t(x)$, which proves part (1).

For part (2), let $e \in \mathcal{E}_t(x)$ be $\mathcal{P}_t(x)$-quasi-surely admissible. By part (1), there exists $\lambda \in \Pi_t(x)$ such that $e \preceq_x H_x^\lambda$. Quasi-sure admissibility of e rules out strict domination, so necessarily

$$e = H_x^\lambda \qquad \mathcal{P}_t(x)\text{-q.s.}$$

For part (3), assume the fiberwise constraint qualification Theorem C.6. The admissibility characterization in (Larsson et al., 2026) then implies that every affine constraint combination H_x^λ with $\lambda \in \Pi_t(x)$ is $\mathcal{P}_t(x)$-quasi-surely admissible. This is exactly the finite-constraint admissibility criterion specialized to the present fiber. □

Corollary C.8 (Representation of pointwise-admissible local certificates)**.** *Fix $x \in \mathcal{X}_t$. If $e \in \mathcal{E}_t(x)$ is pointwise admissible in the sense of Theorem A.14, then there exists $\lambda \in \Pi_t(x)$ such that*

$$e = H_x^\lambda \qquad \mathcal{P}_t(x)\text{-q.s.}$$

If, in addition, the fiberwise constraint qualification Theorem C.6 holds at x, then every H_x^λ, $\lambda \in \Pi_t(x)$, is quasi-surely admissible.

Proof. By Theorem C.3, pointwise admissibility implies quasi-sure admissibility. The conclusion then follows from Theorem C.7. □

C.3. Predictable Selection Along an Admissible Law

We now pass from fiberwise representation to a measurable representation along a fixed admissible law. This is the step needed later in the global graph factorization theorem.

Fix an admissible law P. By Theorem A.10, the regular conditional law of Z_t given X_t under P is a measurable kernel

$$K_t^P(\cdot \,|\, x) \in \mathcal{P}_t(x) \qquad P_{X_t}\text{-a.s.}$$

where P_{X_t} denotes the law of X_t under P. Let $e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0, \infty)$ be a local e-certificate, and define for $x \in \mathcal{X}_t$ and $\lambda \in \mathbb{R}_+^{r_t}$,

$$\Gamma_t^P(x, \lambda) := \int_{\mathcal{Z}_t} \min\Big\{1,\ \big|e_t(x, z) - H_x^\lambda(z)\big|\Big\}\, K_t^P(dz \,|\, x).$$

The quantity $\Gamma_t^P(x, \lambda)$ is zero if and only if $e_t(x, \cdot) = H_x^\lambda(\cdot)$ $K_t^P(\cdot \,|\, x)$-almost surely.

Lemma C.9 (Measurability of the representation correspondence)**.** *Let $B_t \in \mathcal{X}_t$ be measurable. Assume that for every $x \in B_t$, the fiber $e_t(x, \cdot)$ is pointwise admissible in $\mathcal{E}_t(x)$. Define the set-valued map*

$$\Lambda_{t,P}^{e_t}(x) := \begin{cases} \big\{\lambda \in \mathbb{R}_+^{r_t} : \Gamma_t^P(x, \lambda) = 0\big\}, & x \in B_t, \\ \{0\}, & x \notin B_t. \end{cases}$$

Then:

1. *$\Lambda_{t,P}^{e_t}(x)$ is nonempty for every $x \in \mathcal{X}_t$;*
2. *each $\Lambda_{t,P}^{e_t}(x)$ is closed;*
3. *the graph*

$$\mathrm{Gr}\big(\Lambda_{t,P}^{e_t}\big) := \Big\{(x, \lambda) \in \mathcal{X}_t \times \mathbb{R}_+^{r_t} : \lambda \in \Lambda_{t,P}^{e_t}(x)\Big\}$$

is a Borel subset of $\mathcal{X}_t \times \mathbb{R}_+^{r_t}$.

Proof. We first prove measurability of Γ_t^P. The map

$$(x, \lambda, z) \longmapsto \min\Big\{1,\ \big|e_t(x, z) - 1 - \lambda^\top \phi^x(z)\big|\Big\}$$

is jointly measurable on $\mathcal{X}_t \times \mathbb{R}_+^{r_t} \times \mathcal{Z}_t$, because e_t and the constraint functions are measurable and the remaining operations are continuous. Since the integrand is bounded by 1, the standard kernel-measurability theorem implies that $(x, \lambda) \mapsto \Gamma_t^P(x, \lambda)$ is jointly measurable.

Now fix $x \in B_t$. By Theorem C.8, there exists $\lambda \in \Pi_t(x) \subseteq \mathbb{R}_+^{r_t}$ such that

$$e_t(x, \cdot) = H_x^\lambda(\cdot) \qquad \mathcal{P}_t(x)\text{-q.s.}$$

Since $K_t^P(\cdot \mid x) \in \mathcal{P}_t(x)$, this equality also holds $K_t^P(\cdot \mid x)$-almost surely, which means $\Gamma_t^P(x, \lambda) = 0$. Thus $\Lambda_{t,P}^{e_t}(x) \neq \varnothing$ for every $x \in B_t$. For $x \notin B_t$, nonemptiness is built into the definition by $\Lambda_{t,P}^{e_t}(x) = \{0\}$.

To prove closedness, fix $x \in B_t$ and let $\lambda_n \to \lambda$ with $\lambda_n \in \Lambda_{t,P}^{e_t}(x)$. For each z,

$$H_x^{\lambda_n}(z) \to H_x^\lambda(z),$$

because the constraint family is finite. Hence, by dominated convergence,

$$\Gamma_t^P(x, \lambda_n) \to \Gamma_t^P(x, \lambda).$$

Since each $\Gamma_t^P(x, \lambda_n) = 0$, we obtain $\Gamma_t^P(x, \lambda) = 0$, so $\lambda \in \Lambda_{t,P}^{e_t}(x)$. Thus the section is closed. The case $x \notin B_t$ is trivial.

Finally, the graph is

$$\mathrm{Gr}\big(\Lambda_{t,P}^{e_t}\big) = \Big((B_t \times \mathbb{R}_+^{r_t}) \cap \{(x, \lambda) : \Gamma_t^P(x, \lambda) = 0\}\Big) \cup \big(B_t^c \times \{0\}\big),$$

which is Borel because B_t is measurable and Γ_t^P is measurable. □

Proposition C.10 (Predictable local representation under a fixed law)**.** *Fix an admissible law P, a round $t \in [T]$, and a local e-certificate $e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0, \infty)$. Let $B_t \in \mathcal{X}_t$ be measurable with*

$$P(X_t \in B_t) = 1,$$

and assume that for every $x \in B_t$, the fiber $e_t(x, \cdot)$ is pointwise admissible in $\mathcal{E}_t(x)$. Then there exists a Borel measurable map

$$\lambda_{t,P} : \mathcal{X}_t \to \mathbb{R}_+^{r_t}$$

such that

$$e_t(X_t, \cdot) = H_{X_t}^{\lambda_{t,P}(X_t)}(\cdot),$$
$$K_t^P(\cdot \mid X_t)\text{-a.s.}, \quad P\text{-a.s.}$$

Equivalently, if $E_t = e_t(X_t, Z_t)$, then

$$E_t = 1 + \sum_{j=1}^{r_t} \lambda_{t,P,j}(X_t)\, \phi_{t,j}(X_t, Z_t) \qquad P\text{-a.s.}$$

Moreover, $\lambda_{t,P}(X_t)$ is $\mathcal{H}_t = \sigma(X_t)$-measurable.

Proof. By Theorem C.9, the correspondence $\Lambda_{t,P}^{e_t}$ has a Borel graph and nonempty closed values in the Polish space $\mathbb{R}_+^{r_t}$. Since every closed subset of $\mathbb{R}_+^{r_t}$ is σ-compact, a standard measurable selection theorem for Borel sets with nonempty σ-compact sections yields a Borel selector

$$\lambda_{t,P}(x) \in \Lambda_{t,P}^{e_t}(x) \qquad \text{for every } x \in \mathcal{X}_t.$$

By definition of $\Lambda_{t,P}^{e_t}$, for every $x \in B_t$,

$$\Gamma_t^P\big(x, \lambda_{t,P}(x)\big) = 0,$$

that is,

$$e_t(x,\cdot) = H_x^{\lambda_{t,P}(x)}(\cdot) \qquad K_t^P(\cdot \mid x)\text{-a.s.}$$

Since $P(X_t \in B_t) = 1$, substituting $x = X_t$ yields

$$e_t(X_t,\cdot) = H_{X_t}^{\lambda_{t,P}(X_t)}(\cdot) \qquad K_t^P(\cdot \mid X_t)\text{-a.s.}, \quad P\text{-a.s.}$$

Evaluating at Z_t gives

$$e_t(X_t, Z_t) = H_{X_t}^{\lambda_{t,P}(X_t)}(Z_t) = 1 + \sum_{j=1}^{r_t} \lambda_{t,P,j}(X_t)\, \phi_{t,j}(X_t, Z_t) \qquad P\text{-a.s.}$$

Finally, $\lambda_{t,P}$ is Borel measurable and X_t is $\mathcal{H}_t$-measurable, so $\lambda_{t,P}(X_t)$ is $\mathcal{H}_t$-measurable. □

Remark C.11 (What Appendix C actually proves). The local theorem proved here has two layers. First, fiberwise in the local history x, every relevant local e-certificate is quasi-surely represented by a finite-dimensional affine combination of the constraint functions. Second, along any fixed admissible law, those fiberwise representations can be chosen measurably as a function of the observed local state X_t. This is exactly the form needed later in the graph factorization theorem: local uncertainty objects reduce to predictable finite-dimensional coefficients, and those coefficients can be lifted and merged at the graph level.

Remark C.12 (Scope and deliberate limitations). The present appendix is restricted to finitely many bounded predictable constraints. This is deliberate. The general theory of (Larsson et al., 2026) is broader, but the workshop version of the paper only needs the finite-constraint case. We also formulate predictable selection under a *fixed admissible law* P, because that is the level of generality required for the subsequent global e-process argument.

D. Primitive Reduction Lemmas

This appendix isolates the three primitive settings hidden inside any candidate certificate calculus satisfying the axioms of Theorems 2.7 and 2.8. The key point is that our graph framework is rich enough to realize, within the stated workshop scope, three special classes of inputs:

1. bounded serial chains of state-based one-step e-values;
2. bounded already-lifted same-null families on a common sigma-field;
3. bounded filtration refinements of a single e-value.

The restriction of a universally valid graph calculus to these three special inputs yields, respectively, a sequential e-merging constructor, a same-null merger under arbitrary dependence, and a filtration-lift operator. These are exactly the three primitive settings characterized in (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026). Appendix D proves only the reduction from our graph axioms to those primitive settings; the characterization theorems themselves are used later as black boxes.

D.1. Primitive Trace Classes

We begin by formalizing the three primitive inputs.

Definition D.1 (Bounded chain input)**.** Fix $n \in \mathbb{N}$. A *bounded chain input of length* n consists of:

1. a filtered probability space
$$(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_u)_{u=0}^n);$$
2. for each $u \in [n]$, a standard Borel space $(\mathcal{X}_u, \mathcal{X}_u)$ and an $\mathcal{F}_{u-1}$-measurable state variable $X_u : \Omega \to \mathcal{X}_u$;
3. for each $u \in [n]$, a deterministic bound $B_u < \infty$ and an $\mathcal{F}_u$-measurable variable
$$E_u : \Omega \to [0, B_u]$$
satisfying
$$\mathbb{E}[E_u \mid X_u] \le 1 \qquad \text{a.s.}$$

The associated *pure chain graph* is the deterministic chain

$$1 \to 2 \to \cdots \to n,$$

i.e.

$$P_1 = \varnothing, \qquad P_u = \{u-1\} \text{ for } u = 2, \ldots, n, \qquad \tau = n.$$

Definition D.2 (Bounded same-null family)**.** Fix $k \in \mathbb{N}$. A *bounded same-null family of arity* k consists of:

1. a probability space $(\Omega, \mathcal{F})$;
2. a sigma-field $\mathcal{G} \subseteq \mathcal{F}$;
3. a null class $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$;
4. bounded nonnegative $\mathcal{G}$-measurable variables $E_1, \ldots, E_k$ such that

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[E_i] \le 1, \qquad i = 1, \ldots, k.$$

If additionally $w = (w_1, \ldots, w_k) \in [0,1]^k$ satisfies $\sum_{i=1}^k w_i = 1$, then $(E_1, \ldots, E_k; w)$ is called a *weighted bounded same-null family*.

Definition D.3 (Bounded lift input)**.** A *bounded lift input* consists of:

1. a measurable space $(\Omega, \mathcal{F})$;
2. sigma-fields $\mathcal{G} \subseteq \mathcal{G}' \subseteq \mathcal{F}$;
3. a null class $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$;
4. a bounded nonnegative $\mathcal{G}$-measurable random variable E such that

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[E] \le 1.$$

Remark D.4 (Why the reductions are stated for bounded inputs). The workshop paper is formulated under finitely many bounded predictable constraints and finite-valued local certificates. The primitive reductions therefore need to be realizable inside that scope. This is why Theorems D.1 to D.3 are stated for bounded inputs. This causes no tension with the cited primitive characterization theorems, which are stronger than what we need here.

D.2. Chain Restriction

We now isolate the serial primitive. The only role of the graph here is to realize a chain of one-step local e-certificates.

Lemma D.5 (Chain restriction)**.** *Let Γ be a certificate calculus satisfying Theorems 2.7 and 2.8. Then Γ induces a state-based sequential constructor on bounded chain inputs with the following property.*

For every bounded chain input

$$\big((\mathcal{F}_u)_{u=0}^n, (X_u, E_u)_{u=1}^n\big)$$

in the sense of Theorem D.1, there exists a nonnegative $(\mathcal{F}_u)$-adapted process

$$\mathfrak{S}^\Gamma(E_{1:u} \mid X_{1:u})_{u=0}^n$$

with initial value 1 such that:

1. *$\mathfrak{S}^\Gamma(E_{1:u} \mid X_{1:u})$ is an e-process under $\mathbb{P}$;*
2. *the terminal value is obtained by applying Γ to the pure chain graph associated with the input;*
3. *the constructor is associative along chains in the sense inherited from Axiom 2.8(5).*

In particular, the restriction of Γ *to pure chains is an admissible bounded sequential e-merging constructor.*

Proof. Fix a bounded chain input of length n. We realize it inside the execution model of Sections A and 2 as follows. For each $u \in [n]$, define

$$\mathcal{Z}_u := [0, B_u], \qquad \phi_u(x, z) := z - 1, \qquad e_u(x, z) := z.$$

Since $z \in [0, B_u]$, the constraint function ϕ_u is bounded. For each state $x \in \mathcal{X}_u$, define the local null class

$$\mathcal{P}_u(x) := \left\{ Q \in \mathcal{M}_1([0, B_u]) : \int z\, Q(dz) \le 1 \right\}.$$

This class does not depend on x, but that is harmless.

Now define an adaptive execution graph by

$$\tau := n, \qquad P_1 := \varnothing, \qquad P_u := \{u - 1\} \text{ for } u = 2, \dots, n, \qquad Z_u := E_u.$$

By construction, X_u is $\mathcal{F}_{u-1}$-measurable and $Z_u = E_u$ is $\mathcal{F}_u$-measurable, so this is an admissible execution graph of chain type.

We verify admissibility of the law $\mathbb{P}$. Since $Z_u = E_u$ and $\sigma(X_u) \subseteq \mathcal{F}_{u-1}$,

$$\mathbb{E}[\phi_u(X_u, Z_u) \mid X_u] = \mathbb{E}[E_u - 1 \mid X_u] \le 0 \qquad \text{a.s.}$$

Hence $\mathbb{P}$ is admissible in the sense of Theorem A.10. Also, for every $x \in \mathcal{X}_u$ and every $Q \in \mathcal{P}_u(x)$,

$$\int e_u(x, z)\, Q(dz) = \int z\, Q(dz) \le 1,$$

so e_u is a valid local e-certificate.

Apply Γ to this realized chain input and denote the resulting process by

$$M^\Gamma = (M_u^\Gamma)_{u=0}^n.$$

By Axiom 2.8(1), M^Γ is an e-process under $\mathbb{P}$. Define

$$\mathfrak{S}^\Gamma(E_{1:u} \mid X_{1:u}) := M_u^\Gamma, \qquad u = 0, \dots, n.$$

This proves part (1). Part (2) is immediate from the definition. Part (3) follows from Axiom 2.8(5): if the same chain is grouped or parenthesized in two different serial ways, the output of the calculus is unchanged. Therefore the restriction to pure chains is a well-defined bounded sequential constructor.

The final statement is just terminology: validity comes from the e-process property, and the constructor is sequential because it is obtained from serial composition along a chain. □

Remark D.6 (Relation to the classical sequential setting). The chain reduction is stated with conditioning on the predictable state X_u, because that is the conditioning structure of our local model. When the entire relevant past is encoded in X_u, this is the usual one-step e-variable setting studied in (Vovk & Wang, 2024). No extra assumption is needed for the reduction itself.

D.3. Star Restriction

We next isolate the join primitive. To avoid mixing join and lift, we work with *already-lifted* same-null families on a common sigma-field.

Lemma D.7 (Star restriction)**.** *Let* Γ *satisfy Theorems 2.7 and 2.8. Then for every* $k \in \mathbb{N}$ *and every deterministic weight vector* $w = (w_1, \dots, w_k)$ *with* $w_i \ge 0$ *and* $\sum_{i=1}^k w_i = 1$*, the restriction of* Γ *to a pure* k*-ary same-null join induces an operator*

$$\mathfrak{J}_{k,w}^\Gamma(E_1, \dots, E_k)$$

on weighted bounded same-null families with the following properties.

If $(E_1, \dots, E_k; w)$ *is a weighted bounded same-null family in the sense of Theorem D.2, then:*

1. *$\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)$ is $\mathcal{G}$-measurable;*
2. *$\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)$ is again a $\mathcal{Q}$-e-value on $\mathcal{G}$, i.e.*
$$\sup_{Q\in\mathcal{Q}} \mathbb{E}_Q\left[\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)\right] \le 1;$$
3. *the operator depends on its inputs only through the certificate values and the weight vector w, and is permutation-equivariant: if π is a permutation of $[k]$, then*
$$\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k) = \mathfrak{J}^{\Gamma}_{k,\pi w}(E_{\pi(1)},\dots,E_{\pi(k)}) \qquad \text{a.s.},$$
where $(\pi w)_i := w_{\pi(i)}$.

Hence the restriction of Γ to pure same-null joins is an admissible bounded arbitrary-dependence same-null merger.

Proof. Fix a weighted bounded same-null family
$$(\Omega,\mathcal{F},\mathcal{G},\mathcal{Q},E_1,\dots,E_k;w).$$
Let $B_i<\infty$ be deterministic bounds with $0\le E_i\le B_i$ a.s. for each i.

We realize this family inside the graph framework as follows. Take horizon $T=k$, stopping time $\tau=k$, constant filtration
$$\mathcal{F}_0=\mathcal{F}_1=\cdots=\mathcal{F}_k:=\mathcal{G},$$
and edgeless graph
$$P_i=\varnothing, \qquad i=1,\dots,k.$$
For each i, let $\mathcal{X}_i:=\{\star\}$, $\mathcal{Z}_i:=[0,B_i]$, and define
$$X_i:=\star, \qquad Z_i:=E_i, \qquad \phi_i(\star,z):=z-1, \qquad e_i(\star,z):=z.$$
Then each ϕ_i is bounded, and the local null class is
$$\mathcal{P}_i(\star)=\left\{Q\in\mathcal{M}_1([0,B_i]) : \int z\,Q(dz)\le 1\right\}.$$
Fix $Q\in\mathcal{Q}$. Since E_i is a $\mathcal{Q}$-e-value on $\mathcal{G}$,
$$\mathbb{E}_Q[E_i]\le 1.$$
Thus
$$\mathbb{E}_Q[\phi_i(X_i,Z_i)\mid X_i]=\mathbb{E}_Q[E_i-1]\le 0,$$
so Q is admissible for this special execution model. Also, each e_i is a valid local e-certificate because
$$\int e_i(\star,z)\,Q_i(dz)=\int z\,Q_i(dz)\le 1 \qquad \forall Q_i\in\mathcal{P}_i(\star).$$

Now apply the join primitive of Γ to the family $(E_1,\dots,E_k)$ with deterministic weights w. Denote the resulting merged variable by $\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)$. Because the ambient filtration is constantly equal to $\mathcal{G}$, this merged variable is $\mathcal{G}$-measurable.

To prove validity, fix $Q\in\mathcal{Q}$. By Axiom 2.8(1), the output of the calculus is valid under every admissible law. In the present special execution model, the terminal output of the pure join is exactly $\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)$, so
$$\mathbb{E}_Q\left[\mathfrak{J}^{\Gamma}_{k,w}(E_1,\dots,E_k)\right]\le 1.$$
Since this holds for every $Q\in\mathcal{Q}$, the merged output is a $\mathcal{Q}$-e-value.

Finally, Axiom 2.8(6) says that a same-null join depends on its inputs only through the lifted certificate values and predictable weights. Here all certificates already live on the common sigma-field $\mathcal{G}$, so no lift is involved, and the weights are deterministic. Thus the induced operator depends only on $(E_1,\dots,E_k;w)$. The stated permutation-equivariance follows from Axiom 2.8(4): relabeling the edgeless graph by a permutation π leaves the calculus unchanged, while the weight vector is permuted accordingly. □

Remark D.8 (Why lifts are excluded from the star reduction). The point of Theorem D.7 is to isolate the join primitive itself. Therefore all inputs are assumed already to live on a common comparison sigma-field. Filtration changes are handled separately in the next lemma.

D.4. Lift Restriction

We finally isolate the filtration-lift primitive.

Lemma D.9 (Lift restriction)**.** *Let Γ satisfy Theorems 2.7 and 2.8. Then for every bounded lift input $(\mathcal{G}, \mathcal{G}', \mathcal{Q}, E)$ in the sense of Theorem D.3, the restriction of Γ to a pure single-step refinement induces a random variable*

$$\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)$$

such that:

1. *$\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)$ is $\mathcal{G}'$-measurable;*
2. *$\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)$ is again a $\mathcal{Q}$-e-value, i.e.*
$$\sup_{Q\in\mathcal{Q}} \mathbb{E}_Q\big[\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)\big] \le 1;$$
3. *the operator depends only on the pair $(\mathcal{G}, \mathcal{G}')$ and the input E, not on the surrounding graph.*

Hence the restriction of Γ to pure refinement steps is an admissible bounded lift operator.

Proof. Fix a bounded lift input $(\Omega, \mathcal{F}, \mathcal{G}, \mathcal{G}', \mathcal{Q}, E)$, and choose $B < \infty$ such that $0 \le E \le B$ a.s. We realize this inside the execution model by taking horizon $T = 1$, stopping time $\tau = 1$, and filtration

$$\mathcal{F}_0 := \mathcal{G}, \qquad \mathcal{F}_1 := \mathcal{G}'.$$

Let $P_1 = \varnothing$, $\mathcal{X}_1 := \{\star\}$, $\mathcal{Z}_1 := [0, B]$, and define

$$X_1 := \star, \qquad Z_1 := E, \qquad \phi_1(\star, z) := z - 1, \qquad e_1(\star, z) := z.$$

Again ϕ_1 is bounded, and the local null class is

$$\mathcal{P}_1(\star) = \left\{ Q_1 \in \mathcal{M}_1([0, B]) : \int z\, Q_1(dz) \le 1 \right\}.$$

Fix $Q \in \mathcal{Q}$. Since E is a $\mathcal{Q}$-e-value on $\mathcal{G}$,

$$\mathbb{E}_Q[E] \le 1.$$

Therefore

$$\mathbb{E}_Q[\phi_1(X_1, Z_1) \mid X_1] = \mathbb{E}_Q[E - 1] \le 0,$$

so Q is admissible for this one-node execution model. Also, e_1 is a valid local e-certificate.

Now apply the pure lift operation of Γ to this input. Denote the result by $\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)$. Because the post-lift sigma-field is $\mathcal{G}' = \mathcal{F}_1$, the output is $\mathcal{G}'$-measurable.

To prove validity, fix $Q \in \mathcal{Q}$. By universal validity of the calculus,

$$\mathbb{E}_Q\big[\mathfrak{A}^{\Gamma}_{\mathcal{G},\mathcal{G}'}(E)\big] \le 1.$$

Taking the supremum over $Q \in \mathcal{Q}$ proves that the lifted variable is again a $\mathcal{Q}$-e-value.

Finally, Axiom 2.8(7) is exactly the statement that the lift rule is graph-independent. Therefore the induced operator depends only on $(\mathcal{G}, \mathcal{G}', E)$. □

D.5. Consequence for the Main Proposition

The three lemmas above are the only ingredients needed for the reduction statement in the main text. They show that once one freezes the graph-theoretic axioms, every candidate certificate calculus automatically induces:

1. a bounded sequential constructor on pure chains;
2. a bounded same-null merger on already-lifted stars;
3. a bounded lift operator on pure filtrations.

This is exactly the reduction needed later in Section 5: the graph problem decomposes into the three primitive settings that have already been characterized in (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026).

E. Full Proof of the Universal Factorization Theorem

This appendix proves the main structural theorem of the paper. The proof is organized around the three primitive reductions from Appendix D. To make the reduction mathematically airtight, we isolate the exact additional regularity assumptions under which the primitive characterization theorems apply *without hidden loss*:

1. the global calculus is *trace-generated*, so that its output is recursively assembled from serial updates, same-null joins, and filtration lifts;
2. its serial restriction is *precise* on precise chain inputs, so that the process-level characterization of Vovk and Wang applies (Vovk & Wang, 2024);
3. its same-null joins are *primitivewise admissible*, so that Wang's complete-class theorem forces weighted arithmetic averaging (Wang, 2025);
4. when we want a characterization of lifts by adjusters rather than by generic conditional lift operators, we assume the lift acts by an increasing deterministic function of the certificate value and universally lifts running maxima of e-processes across filtrations, so that the characterization theorem of Choe and Ramdas applies (Choe & Ramdas, 2026).

Under these assumptions, the graph calculus has no remaining structural freedom.

E.1. Additional Structural Assumptions

Assumption E.1 (Trace-generated normal form)**.** A certificate calculus Γ is said to be *trace-generated* if, for every admissible execution graph $\mathfrak{E}$ and every compatible certificate evaluation trace, there exists a finite ordered stage sequence

$$u = 1, \dots, m(\mathfrak{E}),$$

together with stage certificates

$$J_u^\Gamma \in L_+^0(\Omega, \mathcal{K}_u), \qquad \mathcal{K}_u := \sigma(J_1^\Gamma, \dots, J_u^\Gamma),$$

such that:

1. each J_u^Γ is obtained from a finite family of already-available local certificates by finitely many lift operations followed by one same-null join;
2. the global process is updated serially by a stagewise constructor

   $$M_u^\Gamma = \mathfrak{S}_u^\Gamma(J_1^\Gamma, \dots, J_u^\Gamma), \qquad u = 0, \dots, m(\mathfrak{E}),$$

   with $M_0^\Gamma = 1$;
3. the map $\mathfrak{S}^\Gamma$ is adapted to the stage filtration $(\mathcal{K}_u)$.

Assumption E.2 (Serial precision on precise chains)**.** The serial restriction of Γ is *precise* on pure chains: whenever $(J_u)_{u=1}^n$ is a bounded chain input such that

$$\mathbb{E}[J_u \mid J_1, \dots, J_{u-1}] = 1 \qquad \text{a.s. for every } u \in [n],$$

the induced chain process

$$\mathfrak{S}_u^\Gamma(J_1, \dots, J_u), \qquad u = 0, \dots, n,$$

is precise in the sense that

$$\mathbb{E}\big[\mathfrak{S}_\sigma^\Gamma(J_1, \dots, J_\sigma)\big] = 1$$

for every stopping time σ with values in $\{0, \dots, n\}$.

Assumption E.3 (Primitivewise admissibility of same-null joins)**.** For every bounded same-null family in the sense of Theorem D.2, the induced join operator of Γ is an admissible e-merging function under arbitrary dependence.

Assumption E.4 (Conditional lift validity)**.** For every stage u, every stage-past sigma-field $\mathcal{K}_{u-1}$, and every nonnegative random variable E satisfying

$$E \in L_+^0(\Omega, \mathcal{K}_u), \qquad \mathbb{E}[E \mid \mathcal{K}_{u-1}] \le 1 \quad \text{a.s.},$$

every lift operator L used by Γ at stage u preserves one-step conditional e-validity:

$$\mathbb{E}[L(E) \mid \mathcal{K}_{u-1}] \le 1 \qquad \text{a.s.}$$

Assumption E.5 (Value-based universal liftability)**.** For every pair of sigma-fields $\mathcal{G} \subseteq \mathcal{G}' \subseteq \mathcal{F}$, the pure lift restriction of Γ is given by an increasing deterministic function

$$a_{\mathcal{G},\mathcal{G}'} : [0, \infty) \to [0, \infty),$$

in the sense that a pure lift of a bounded $\mathcal{G}$-measurable e-value E is $a_{\mathcal{G},\mathcal{G}'}(E)$. Moreover, for every null class $\mathcal{Q}$, every $\mathcal{G}$-adapted e-process $(M_u)_{u \ge 0}$ for $\mathcal{Q}$, and every finer filtration $(\mathcal{G}'_u)_{u \ge 0}$ with $\mathcal{G}_u \subseteq \mathcal{G}'_u$, the process

$$u \longmapsto a_{\mathcal{G},\mathcal{G}'}\left(\max_{0 \le s \le u} M_s\right)$$

is again an e-process for $\mathcal{Q}$ in the finer filtration.

Remark E.6 (Why these assumptions are explicit). The graph axioms from Theorem 2.8 encode invariance, monotonicity, and validity, but they do not by themselves force the calculus to be recursively generated by primitive trace operations, nor do they imply the precision condition required by the game-martingale characterization of (Vovk & Wang, 2024). Likewise, generic graph-independent lift operators are not automatically of the value-based form treated in (Choe & Ramdas, 2026). The present theorem therefore states exactly the strongest result that follows without conflating distinct notions of primitive regularity.

E.2. Auxiliary Lemmas

Lemma E.7 (Stagewise conditional validity is preserved by conditional lifts and weighted joins)**.** *Assume Theorem E.4. Fix a stage u, and let*

$$\widetilde{E}_{u,1}, \dots, \widetilde{E}_{u,k_u}$$

be stage-u certificates such that each $\widetilde{E}_{u,i} \in L_+^0(\Omega, \mathcal{K}_u)$ and

$$\mathbb{E}[\widetilde{E}_{u,i} \mid \mathcal{K}_{u-1}] \le 1 \qquad \textit{a.s. for } i = 1, \dots, k_u.$$

Let $w_{u,1}, \dots, w_{u,k_u}$ be nonnegative $\mathcal{K}_{u-1}$-measurable weights with

$$\sum_{i=1}^{k_u} w_{u,i} = 1 \qquad \textit{a.s.}$$

Then

$$J_u := \sum_{i=1}^{k_u} w_{u,i} \widetilde{E}_{u,i}$$

satisfies

$$\mathbb{E}[J_u \mid \mathcal{K}_{u-1}] \le 1 \qquad a.s.$$

Hence the sequence (J_u) is a bounded sequence of stagewise sequential e-values.

Proof. Since each $w_{u,i}$ is $\mathcal{K}_{u-1}$-measurable, the tower property gives

$$\mathbb{E}[J_u \mid \mathcal{K}_{u-1}] = \sum_{i=1}^{k_u} w_{u,i}\mathbb{E}[\widetilde{E}_{u,i} \mid \mathcal{K}_{u-1}] \le \sum_{i=1}^{k_u} w_{u,i} = 1 \qquad \text{a.s.}$$

The final statement is immediate. □

Lemma E.8 (Game-martingale form of the serial primitive)**.** *Assume Theorems E.1 and E.2. Let Γ be universally valid. Then its serial stage constructor*

$$\mathfrak{S}_u^\Gamma(J_1, \dots, J_u), \qquad u = 0, \dots, n,$$

restricted to any bounded chain input of stagewise sequential e-values, is a game martingale. Consequently, for each $u \in [n]$, there exists a measurable betting fraction

$$\beta_u^\Gamma : \mathbb{R}_+^{u-1} \to [0, 1]$$

such that

$$\mathfrak{S}_0^\Gamma = 1, \qquad \mathfrak{S}_u^\Gamma(J_1, \dots, J_u) = \mathfrak{S}_{u-1}^\Gamma(J_1, \dots, J_{u-1})\Big(1 + \beta_u^\Gamma(J_1, \dots, J_{u-1})(J_u - 1)\Big).$$

Proof. By Theorem D.5, the chain restriction of Γ is adapted and anytime-valid on bounded chain inputs. By Theorem E.2, it is also precise on precise chain inputs. Therefore Theorem 2 of (Vovk & Wang, 2024) applies: any adapted, anytime-valid, and precise sequential constructor is a game martingale. The displayed recursive form is exactly the equivalent definition of a game martingale in that paper. □

Lemma E.9 (Weighted-average form of same-null joins)**.** *Assume Theorem E.3. Let Γ be universally valid. Then for every pure same-null join of arity k, the induced join operator has the form*

$$\mathfrak{J}_{k,w}^\Gamma(E_1, \dots, E_k) = \sum_{i=1}^{k} \bar{w}_i E_i$$

for some weight vector $\bar{w} = (\bar{w}_1, \dots, \bar{w}_k) \in [0, 1]^k$ with $\sum_{i=1}^k \bar{w}_i = 1$. Equivalently, every same-null join primitive of Γ is a weighted arithmetic average.

Proof. By Theorem D.7, the restriction of Γ to pure same-null stars is a valid bounded e-merging function under arbitrary dependence. By Theorem E.3, this e-merging function is primitivewise admissible. Theorem 1(ii) of (Wang, 2025) states that an e-merging function under arbitrary dependence is admissible if and only if it is a weighted arithmetic average. The conclusion follows. □

Lemma E.10 (Adjuster form of value-based universal lifts)**.** *Assume Theorem E.5. Then every pure lift function $a_{\mathcal{G},\mathcal{G}'}$ is an adjuster. In particular, for every bounded $\mathcal{G}$-measurable e-value E,*

$$a_{\mathcal{G},\mathcal{G}'}(E)$$

is a valid lift to $\mathcal{G}'$, and the lift is of adjuster type in the sense of (Choe & Ramdas, 2026).

Proof. The assumption says exactly that the increasing deterministic function $a_{\mathcal{G},\mathcal{G}'}$ lifts the running maximum of *every* e-process from a coarser filtration to a finer filtration while preserving e-process validity. Theorem 3(e)⇒(a) of (Choe & Ramdas, 2026) states that any such function must be an adjuster. Applying that theorem with the current pair $(\mathcal{G}, \mathcal{G}')$ yields the claim. □

E.3. Proof of the Main Theorem

Proof of Theorem 5.2. Fix a universally valid certificate calculus Γ satisfying Theorems E.1 to E.4. We prove the theorem in three steps: identification of the serial primitive, identification of the join primitive, and recursive assembly of the global factorization. We then prove the converse.

Step 1: identification of the serial primitive. By Theorem E.1, the output of Γ on any admissible execution graph can be written stagewise as

$$M_u^\Gamma = \mathfrak{S}_u^\Gamma(J_1^\Gamma, \dots, J_u^\Gamma), \qquad u = 0, \dots, m(\mathfrak{E}),$$

for some stage certificates J_u^Γ. By Theorem E.8, the serial constructor $\mathfrak{S}^\Gamma$ is a game martingale. Therefore there exist measurable betting fractions

$$\beta_u^\Gamma : \mathbb{R}_+^{u-1} \to [0, 1]$$

such that

$$M_0^\Gamma = 1, \qquad M_u^\Gamma = M_{u-1}^\Gamma \Big(1 + \beta_u^\Gamma(J_1^\Gamma, \dots, J_{u-1}^\Gamma)(J_u^\Gamma - 1)\Big). \tag{4}$$

Step 2: identification of the join primitive. Fix a stage u. By Theorem E.1, the stage certificate J_u^Γ is obtained from a finite family of already available certificates by finitely many lifts followed by one same-null join. Write the lifted stage inputs as

$$\widetilde{E}_{u,1}, \dots, \widetilde{E}_{u,k_u}.$$

By Theorem E.9, the final same-null join at stage u must be a weighted arithmetic average. Hence there exist nonnegative weights

$$w_{u,1}, \dots, w_{u,k_u}$$

summing to one such that

$$J_u^\Gamma = \sum_{i=1}^{k_u} w_{u,i} \widetilde{E}_{u,i}. \tag{5}$$

If, in addition, Theorem E.5 holds, then Theorem E.10 implies that each lift in the stage can be written as

$$\widetilde{E}_{u,i} = a_{u,i}(E_{u,i}),$$

where $a_{u,i}$ is an adjuster for the corresponding filtration pair. Substituting into (5) yields

$$J_u^\Gamma = \sum_{i=1}^{k_u} w_{u,i} a_{u,i}(E_{u,i}). \tag{6}$$

Without Theorem E.5, the same formula holds with generic conditional lift operators in place of adjusters.

Step 3: assembly of the factorized form. Combining (4) with (5) gives the recursive factorization

$$M_u^\Gamma = M_{u-1}^\Gamma \left(1 + \beta_u^\Gamma(J_1^\Gamma, \dots, J_{u-1}^\Gamma) \left(\sum_{i=1}^{k_u} w_{u,i} \widetilde{E}_{u,i} - 1\right)\right), \qquad u = 1, \dots, m(\mathfrak{E}). \tag{7}$$

Iterating (7) gives the product form

$$M_u^\Gamma = \prod_{s=1}^{u} \left(1 + \beta_s^\Gamma(J_1^\Gamma, \dots, J_{s-1}^\Gamma) \left(\sum_{i=1}^{k_s} w_{s,i} \widetilde{E}_{s,i} - 1\right)\right). \tag{8}$$

If Theorem E.5 holds, then (8) sharpens to

$$M_u^\Gamma = \prod_{s=1}^{u} \left(1 + \beta_s^\Gamma(J_1^\Gamma, \dots, J_{s-1}^\Gamma) \left(\sum_{i=1}^{k_s} w_{s,i} a_{s,i}(E_{s,i}) - 1\right)\right). \tag{9}$$

This proves the necessity half of the theorem.

Converse: validity of the factorized calculus. Now suppose we are given a trace-generated calculus $\bar{\Gamma}$ constructed as follows. At each stage u,

1. each currently available local certificate $E_{u,i}$ is transformed into a lifted certificate $\widetilde{E}_{u,i}$ by a lift operator that satisfies the conditional validity property of Theorem E.4;
2. the lifted certificates are merged by a weighted arithmetic average
$$J_u := \sum_{i=1}^{k_u} w_{u,i}\widetilde{E}_{u,i}, \qquad w_{u,i} \geq 0, \quad \sum_i w_{u,i} = 1,$$
with weights measurable with respect to the stage-past sigma-field $\mathcal{K}_{u-1}$;
3. the global process is updated by a game martingale
$$M_0^{\bar{\Gamma}} = 1, \qquad M_u^{\bar{\Gamma}} = M_{u-1}^{\bar{\Gamma}}\Big(1 + \beta_u(J_1, \ldots, J_{u-1})(J_u - 1)\Big).$$

We show that $\bar{\Gamma}$ is universally valid.

Fix an admissible law P. By the local validity relation from Theorem 2.3, every base local certificate satisfies the one-step conditional inequality required in Theorem E.4. Hence every lifted certificate $\widetilde{E}_{u,i}$ still satisfies
$$\mathbb{E}_P[\widetilde{E}_{u,i} \mid \mathcal{K}_{u-1}] \leq 1 \qquad \text{a.s.}$$
Therefore, by Theorem E.7,
$$\mathbb{E}_P[J_u \mid \mathcal{K}_{u-1}] \leq 1 \qquad \text{a.s.}$$
for every stage u. Thus the stage sequence (J_u) is a bounded sequence of sequential e-values.

Since the serial update is a game martingale, Lemma 2 of (Vovk & Wang, 2024) implies that $(M_u^{\bar{\Gamma}})_{u=0}^{m(\mathfrak{E})}$ is an e-process with respect to the stage filtration $(\mathcal{K}_u)$. Because the stage filtration is a sub-filtration of the execution filtration and the horizon is bounded, the same process is an e-process relative to the execution history up to the stage resolution used by $\bar{\Gamma}$. Hence $\bar{\Gamma}$ is universally valid.

If, furthermore, the lifts are of the value-based universal form from Theorem E.5, then each lift can be taken to be an adjuster by Theorem E.10; this yields the sharpened factorization with adjusters in (9). This completes the proof. □

Remark E.11 (What is exact and what is conditional). The join characterization in this appendix is exact because Wang's theorem is a complete classification of admissible arbitrary-dependence e-merging functions (Wang, 2025). The serial characterization is exact under the precision assumption because Vovk and Wang's game-martingale theorem is a full process-level characterization of adapted, anytime-valid, precise sequential constructors (Vovk & Wang, 2024). The lift characterization is conditional on the value-based universal-liftability assumption because that is the exact scope of the adjuster characterization proved by Choe and Ramdas (Choe & Ramdas, 2026).

Remark E.12 (How the theorem will be stated in the main text). In the main body, the theorem should be stated in the sharpened form appropriate for the workshop paper: a trace-generated universally valid certificate calculus with precise serial restriction, admissible same-null joins, and value-based universal lifts factors through three primitives:

1. a game-martingale serial update;
2. weighted arithmetic averaging at same-null joins;
3. adjusters across filtration lifts.

The converse also holds for any factorized calculus built from conditional lifts, weighted same-null averages, and a game-martingale serial constructor.

F. Quantitative Lower Bounds

This appendix quantifies the cost of universal composability. The structural theorem in Section 5 shows that a universally composable calculus must factor through three primitives: a serial game-martingale update, weighted arithmetic averaging at same-null joins, and (under the strengthened lift assumption) adjusters across filtration changes. The present appendix asks a different question: *how much evidence must be lost because of these primitive restrictions?*

There is one subtlety. For a general serial game-martingale update, the stagewise losses induced by joins and lifts can be attenuated by the serial betting fractions. Thus the cleanest quantitative statement is obtained for the *full-betting canonical factorization*

$$M_u = \prod_{s=1}^{u} J_s,$$

corresponding to the serial update $\beta_s \equiv 1$. This is the canonical serial rule that preserves all stagewise log-evidence. We therefore proceed in three steps:

1. prove the exact worst-case join loss;
2. prove a general lower bound for value-based lifts;
3. combine them additively along a primitive-expanded factorized trace under the full-betting serial update.

F.1. Exact Loss at a Weighted Same-Null Join

We begin with the join primitive. Let

$$w = (w_1, \dots, w_k) \in [0,1]^k, \qquad \sum_{i=1}^{k} w_i = 1,$$

and define the weighted arithmetic average

$$A_w(E_1, \dots, E_k) := \sum_{i=1}^{k} w_i E_i.$$

The natural branchwise oracle at such a join is the maximal branch

$$O(E_1, \dots, E_k) := \max_{1 \le i \le k} E_i.$$

The next proposition identifies the exact worst-case log-loss of the admissible join rule A_w relative to O.

Proposition F.1 (Exact weighted-join tax)**.** *Fix $k \in \mathbb{N}$ and a deterministic weight vector $w = (w_1, \dots, w_k) \in [0,1]^k$ with $\sum_{i=1}^{k} w_i = 1$. Let*

$$w_{\min} := \min_{1 \le i \le k} w_i.$$

Then

$$\sup_{(\mathcal{Q}, E_1, \dots, E_k)} \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q \left[\log \frac{\max_{1 \le i \le k} E_i}{\sum_{i=1}^{k} w_i E_i} \right] = \log \frac{1}{w_{\min}}, \tag{10}$$

where the supremum is taken over all bounded same-null families $(E_1, \dots, E_k)$ for the null class $\mathcal{Q}$ in the sense of Theorem D.2, with the convention $0/0 := 1$ inside the logarithm. In particular, if $w_i = 1/k$ for all i, then the exact join tax is $\log k$.

Proof. Fix any bounded same-null family $(E_1, \dots, E_k)$. Choose an index $i^\star \in \operatorname{argmax}_i E_i$. Since $w_i \ge w_{\min}$,

$$\sum_{i=1}^{k} w_i E_i \ge w_{i^\star} E_{i^\star} \ge w_{\min} \max_{1 \le i \le k} E_i.$$

Therefore, pathwise,

$$1 \le \frac{\max_i E_i}{\sum_i w_i E_i} \le \frac{1}{w_{\min}},$$

with the convention $0/0 := 1$ on the event $\{E_1 = \cdots = E_k = 0\}$. Taking logarithms and expectations yields

$$\mathbb{E}_Q\left[\log \frac{\max_i E_i}{\sum_i w_i E_i}\right] \le \log \frac{1}{w_{\min}} \qquad \forall Q \in \mathcal{Q}.$$

To prove equality, choose an index $i_\star \in \operatorname{argmin}_i w_i$, and consider the singleton null class $\mathcal{Q} = \{Q\}$ on the one-point probability space. Define deterministic bounded e-values by

$$E_{i_\star} \equiv 1, \qquad E_i \equiv 0 \text{ for } i \ne i_\star.$$

These are all valid e-values under Q, because their expectations are at most 1. For this family,

$$\max_i E_i = 1, \qquad \sum_{i=1}^k w_i E_i = w_{i_\star} = w_{\min},$$

hence

$$\log \frac{\max_i E_i}{\sum_i w_i E_i} = \log \frac{1}{w_{\min}} \qquad Q\text{-a.s.}$$

Taking expectations gives equality in (10). □

Corollary F.2 (Symmetric star tax)**.** *Under the symmetry enforced by graph-isomorphism invariance and parallel anonymity, a k-ary same-null join has equal weights $w_i = 1/k$. Therefore its exact worst-case log-loss is*

$$\log k.$$

Proof. By symmetry, every child of a pure k-ary star plays the same role, so the weights must be equal. The conclusion follows from Theorem F.1. □

F.2. A Lower Bound for Value-Based Lift Operators

We next isolate the filtration-lift primitive. In the strengthened part of Theorem 5.2, a pure lift is given by an increasing deterministic function

$$a : [0, \infty) \to [0, \infty)$$

that universally lifts running maxima of e-processes across filtrations. Such a function is an adjuster by (Choe & Ramdas, 2026). The relevant question is how much evidence is lost when one replaces the raw running maximum S by its lifted version $a(S)$.

Lemma F.3 (Normalization at one)**.** *Let a be a value-based universal lift in the sense of Theorem E.5. Then*

$$a(1) \le 1.$$

Proof. Apply the lift rule to the constant e-process $M_u \equiv 1$. Its running maximum is identically equal to 1. By universal liftability, $a(1)$ must itself be an e-value under every null class, hence in particular under any singleton null. Therefore $a(1) \le 1$. □

Proposition F.4 (Lift tax lower bound)**.** *Let $a : [0, \infty) \to [0, \infty)$ be a value-based universal lift, and interpret $\log(x/a(x)) = +\infty$ whenever $a(x) = 0 < x$. Define*

$$\Delta_{\text{lift}}(a) := \sup_{x \ge 1} \frac{1}{x} \log \frac{x}{a(x)} \in [0, \infty].$$

Then for every $\varepsilon > 0$, there exists a bounded one-step e-process $M = (M_0, M_1)$ under a singleton null law Q such that, with running maximum

$$S := \max\{M_0, M_1\},$$

one has

$$\mathbb{E}_Q\left[\log \frac{S}{a(S)}\right] \geq \Delta_{\text{lift}}(a) - \varepsilon.$$

In particular,

$$\sup_{Q,M} \mathbb{E}_Q\left[\log \frac{S}{a(S)}\right] \geq \Delta_{\text{lift}}(a),$$

where the supremum ranges over all bounded e-processes M and all singleton null laws Q.

Proof. Fix $\varepsilon > 0$. By definition of $\Delta_{\text{lift}}(a)$, choose $x_\varepsilon \geq 1$ such that

$$\frac{1}{x_\varepsilon} \log \frac{x_\varepsilon}{a(x_\varepsilon)} \geq \Delta_{\text{lift}}(a) - \varepsilon.$$

Let Q be a singleton null law on a two-point probability space $\Omega = \{\omega_1, \omega_0\}$ with

$$Q(\omega_1) = \frac{1}{x_\varepsilon}, \qquad Q(\omega_0) = 1 - \frac{1}{x_\varepsilon}.$$

Define a one-step process by

$$M_0 \equiv 1, \qquad M_1(\omega_1) = x_\varepsilon, \qquad M_1(\omega_0) = 0.$$

Then M is a bounded nonnegative martingale under Q, since

$$\mathbb{E}_Q[M_1] = \frac{1}{x_\varepsilon} \cdot x_\varepsilon + \left(1 - \frac{1}{x_\varepsilon}\right) \cdot 0 = 1 = M_0.$$

Hence M is a bounded e-process.

Its running maximum is

$$S(\omega_1) = x_\varepsilon, \qquad S(\omega_0) = 1.$$

Therefore

$$\mathbb{E}_Q\left[\log \frac{S}{a(S)}\right] = \frac{1}{x_\varepsilon} \log \frac{x_\varepsilon}{a(x_\varepsilon)} + \left(1 - \frac{1}{x_\varepsilon}\right) \log \frac{1}{a(1)}.$$

By Theorem F.3, $a(1) \leq 1$, so the second term is nonnegative. Hence

$$\mathbb{E}_Q\left[\log \frac{S}{a(S)}\right] \geq \frac{1}{x_\varepsilon} \log \frac{x_\varepsilon}{a(x_\varepsilon)} \geq \Delta_{\text{lift}}(a) - \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, the result follows. □

Remark F.5 (Interpretation). The quantity $\Delta_{\text{lift}}(a)$ is a *worst-case expected log-evidence loss* generated by the lift function a, tested against the simplest possible coarse-filtration e-process. The lower bound is intentionally primitive: it does not attempt to optimize over all filtration pairs or all null classes. Its purpose is to isolate an intrinsic cost attached to the lift function itself.

F.3. Additive Graph Tax for Primitive-Expanded Full-Betting Traces

We now combine the primitive losses additively. Because arbitrary serial game-martingale updates may downweight stagewise evidence, the clean additive result is stated for the canonical full-betting factorization.

Definition F.6 (Primitive-expanded full-betting trace)**.** A *primitive-expanded full-betting trace* consists of a finite ordered sequence of stages $u = 1, \ldots, m$, where each stage is one of the following:

1. a *join stage*, specified by a weighted bounded same-null family

$$(E_{u,1}, \ldots, E_{u,k_u}; w_{u,1}, \ldots, w_{u,k_u});$$

2. a *lift stage*, specified by a bounded running maximum $S_u \geq 1$ of a coarse-filtration e-process together with a value-based lift a_u.

The factorized stage certificate is

$$J_u^{\Gamma} = \begin{cases} \sum_{i=1}^{k_u} w_{u,i} E_{u,i}, & \text{if } u \text{ is a join stage,} \\ a_u(S_u), & \text{if } u \text{ is a lift stage,} \end{cases}$$

and the corresponding *oracle stage certificate* is

$$J_u^{\mathrm{or}} = \begin{cases} \max_{1 \leq i \leq k_u} E_{u,i}, & \text{if } u \text{ is a join stage,} \\ S_u, & \text{if } u \text{ is a lift stage.} \end{cases}$$

The associated full-betting global processes are

$$M_u^{\Gamma} := \prod_{s=1}^{u} J_s^{\Gamma}, \qquad M_u^{\mathrm{or}} := \prod_{s=1}^{u} J_s^{\mathrm{or}}, \qquad u = 0, \ldots, m,$$

with the convention $M_0^{\Gamma} = M_0^{\mathrm{or}} = 1$.

For convenience, recall the stage taxes from Section 6. For a join stage u,

$$\Delta_{\mathrm{join}}(u) := \log \frac{1}{w_{u,\min}}, \qquad w_{u,\min} := \min_{1 \leq i \leq k_u} w_{u,i},$$

and for a lift stage u,

$$\Delta_{\mathrm{lift}}(u) := \Delta_{\mathrm{lift}}(a_u) = \sup_{x \geq 1} \frac{1}{x} \log \frac{x}{a_u(x)}.$$

The total primitive tax is

$$\mathfrak{C}(\mathfrak{T}) := \sum_{u \in \mathcal{J}} \Delta_{\mathrm{join}}(u) + \sum_{u \in \mathcal{L}} \Delta_{\mathrm{lift}}(u).$$

Proof of Proposition 6.1. Fix a primitive-expanded full-betting trace $\mathfrak{T}$ and $\varepsilon > 0$. We distribute the error budget over the lift stages: for each $u \in \mathcal{L}$, choose $\varepsilon_u > 0$ such that

$$\sum_{u \in \mathcal{L}} \varepsilon_u \leq \varepsilon.$$

Join stages. For each $u \in \mathcal{J}$, choose an index $i_u^{\star} \in \operatorname{argmin}_i w_{u,i}$. Let Q_u be a singleton null law on a one-point space, and choose deterministic inputs

$$E_{u,i_u^{\star}} \equiv 1, \qquad E_{u,i} \equiv 0 \ \text{ for } i \neq i_u^{\star}.$$

Then, by the attainment argument in the proof of Theorem F.1,

$$\log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}} = \log \frac{1}{w_{u,\min}} = \Delta_{\mathrm{join}}(u) \qquad Q_u\text{-a.s.}$$

Lift stages. For each $u \in \mathcal{L}$, apply Theorem F.4 to the lift function a_u with tolerance ε_u. This yields a singleton null law Q_u and a bounded one-step e-process whose running maximum S_u satisfies

$$\mathbb{E}_{Q_u}\left[\log \frac{S_u}{a_u(S_u)}\right] \geq \Delta_{\mathrm{lift}}(u) - \varepsilon_u.$$

Since $J_u^{\mathrm{or}} = S_u$ and $J_u^{\Gamma} = a_u(S_u)$, this is exactly

$$\mathbb{E}_{Q_u}\left[\log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}}\right] \geq \Delta_{\mathrm{lift}}(u) - \varepsilon_u.$$

Product construction. Take the product probability space over all stages and let

$$Q^{\varepsilon} := \bigotimes_{u=1}^{m} Q_u.$$

By construction, the primitive inputs are stagewise independent, and every stage is compatible with the primitive-expanded trace. Since

$$M_m^{\mathrm{or}} = \prod_{u=1}^{m} J_u^{\mathrm{or}}, \qquad M_m^{\Gamma} = \prod_{u=1}^{m} J_u^{\Gamma},$$

we have the exact pathwise decomposition

$$\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}} = \sum_{u=1}^{m} \log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}}.$$

Taking Q^{ε}-expectations and using linearity,

$$\mathbb{E}_{Q^{\varepsilon}}\left[\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}}\right] = \sum_{u\in\mathcal{J}} \mathbb{E}_{Q_u}\left[\log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}}\right] + \sum_{u\in\mathcal{L}} \mathbb{E}_{Q_u}\left[\log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}}\right].$$

Substituting the stagewise bounds above gives

$$\mathbb{E}_{Q^{\varepsilon}}\left[\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}}\right] \geq \sum_{u\in\mathcal{J}} \Delta_{\mathrm{join}}(u) + \sum_{u\in\mathcal{L}} \big(\Delta_{\mathrm{lift}}(u) - \varepsilon_u\big).$$

Since $\sum_{u\in\mathcal{L}} \varepsilon_u \leq \varepsilon$, this proves the claim. □

Proof of Corollary 6.2. Assume every stage is a symmetric k_u-ary join stage and that there are no lift stages. Then $\Delta_{\mathrm{join}}(u) = \log k_u$ by Theorem F.2. The lower bound

$$\sup_Q \mathbb{E}_Q\left[\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}}\right] \geq \sum_{u=1}^{m} \log k_u$$

follows from Proposition 6.1 with $\varepsilon \downarrow 0$, and is attained by the product of one-hot star instances constructed stagewise as in the proof of Proposition 6.1.

For the upper bound, Theorem F.1 yields the pathwise inequality

$$\frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}} \leq k_u \qquad \text{a.s. for every stage } u.$$

Therefore

$$\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}} = \sum_{u=1}^{m} \log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}} \leq \sum_{u=1}^{m} \log k_u \qquad \text{a.s.}$$

Taking expectations and then the supremum over Q proves the reverse inequality. Hence

$$\sup_Q \mathbb{E}_Q\left[\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}}\right] = \sum_{u=1}^{m} \log k_u.$$

□

Remark F.7 (Why the theorem is stated for the full-betting serial rule). The additive identity

$$\log \frac{M_m^{\mathrm{or}}}{M_m^{\Gamma}} = \sum_{u=1}^{m} \log \frac{J_u^{\mathrm{or}}}{J_u^{\Gamma}}$$

holds because the full-betting serial update is multiplicative: $M_u = \prod_{s\leq u} J_s$. For a general game-martingale serial update with betting fractions $\beta_u \in [0,1]$, the stagewise losses are attenuated by the map

$$x \mapsto \log\big(1 + \beta_u(x-1)\big),$$

and exact additivity is no longer available. Thus Proposition 6.1 should be interpreted as the intrinsic primitive tax revealed by the canonical full-betting factorization.

G. Embedding Conformal / Risk-Control Modules

This appendix shows how two natural classes of local uncertainty modules fit the graph-indexed e-certificate formalism of the main paper. The first class is rank-based and covers local conformal e-predictors built from exchangeable ranks. The second class is risk-based and covers modules whose local guarantee is a bounded conditional mean-risk constraint. The point is not to claim that every conformal or risk-control algorithm already arrives in the exact form used here; rather, the propositions below identify clean *interfaces* by which such modules can enter the factorized graph calculus.

G.1. A General Plug-in Theorem

We begin with an abstract embedding statement. It is a direct corollary of the converse direction of Theorem 5.2, but it is useful to state it explicitly because it turns local module design into a plug-in problem.

Theorem G.1 (Graph embedding of local e-certifying modules)**.** *Assume the hypotheses of Theorem 5.2, and let Γ be a trace-generated factorized calculus satisfying the converse conditions stated there. Fix an admissible execution graph $\mathfrak{E}$, and suppose that at every realized node $t \le \tau$, the module emits a local e-certificate E_t in the sense of Theorem 2.3. Assume further that, at every join stage u,*

1. *the lifted stage inputs $\widetilde{E}_{u,1}, \dots, \widetilde{E}_{u,k_u}$ are obtained from previously available local certificates by lifts preserving one-step conditional e-validity; and*
2. *these lifted inputs certify a common downstream null, so that the stage join is same-null.*

Then the global process M^Γ produced by the factorized calculus is an e-process under every admissible law.

In particular, any family of local modules that outputs valid local e-certificates can be composed serially, merged stagewise under same-null weighted averaging, and transported across filtrations by conditionally valid lifts without losing global e-validity.

Proof. Fix an admissible law P. By local validity,

$$\mathbb{E}_P[E_t \mid X_t] \le 1 \qquad P\text{-a.s.}$$

for every realized node t. By assumption, every lift used at stage u preserves one-step conditional e-validity, so each lifted stage input $\widetilde{E}_{u,i}$ satisfies

$$\mathbb{E}_P[\widetilde{E}_{u,i} \mid \mathcal{K}_{u-1}] \le 1 \qquad P\text{-a.s.}$$

with respect to the stage-past sigma-field $\mathcal{K}_{u-1}$. Since the join is same-null and weighted averaging is used at the stage level, the stage certificate

$$J_u^\Gamma = \sum_{i=1}^{k_u} w_{u,i} \widetilde{E}_{u,i}, \qquad \sum_{i=1}^{k_u} w_{u,i} = 1,$$

again satisfies

$$\mathbb{E}_P[J_u^\Gamma \mid \mathcal{K}_{u-1}] \le 1 \qquad P\text{-a.s.}$$

Thus the stage sequence (J_u^Γ) is a bounded sequence of stagewise sequential e-values. By the converse part of Theorem 5.2, the serial game-martingale update applied to this stage sequence yields a global e-process. Hence M^Γ is an e-process under P. Since P was arbitrary admissible, the claim follows. □

G.2. Rank-Based Local Modules: A Conformal Interface

We now formalize a rank-based interface that captures local conformal e-predictors. The proposition itself is elementary: uniform ranks imply local e-certificates by averaging over the rank support. The conformal interpretation comes from the fact that exchangeability of the test score with the calibration scores implies the required uniform-rank null, which is the basic mechanism of conformal e-prediction (Gauthier et al., 2025).

Proposition G.2 (Uniform-rank modules yield local e-certificates)**.** *Fix a round $t \in [T]$, and let $n_t \in \mathbb{N}$. Suppose there exists a measurable map*

$$r_t : \mathcal{X}_t \times \mathcal{Z}_t \to \{1, \dots, n_t + 1\}.$$

Define the local rank-null class

$$\mathcal{P}_t^{\mathrm{rank}}(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : Q \circ r_t(x,\cdot)^{-1} = \mathrm{Unif}\{1,\dots,n_t+1\} \right\}.$$

Let

$$h_t : \mathcal{X}_t \times \{1,\dots,n_t+1\} \to [0,\infty)$$

be measurable and satisfy

$$\frac{1}{n_t+1} \sum_{r=1}^{n_t+1} h_t(x,r) \le 1 \qquad \textit{for every } x \in \mathcal{X}_t.$$

Define

$$e_t^{\mathrm{rank}}(x,z) := h_t(x, r_t(x,z)).$$

Then e_t^{rank} is a valid local e-certificate for the null class $\mathcal{P}_t^{\mathrm{rank}}$. Equivalently, if

$$E_t^{\mathrm{rank}} := e_t^{\mathrm{rank}}(X_t, Z_t),$$

then under any admissible law P satisfying $K_t^P(\cdot \mid X_t) \in \mathcal{P}_t^{\mathrm{rank}}(X_t)$ almost surely,

$$\mathbb{E}_P[E_t^{\mathrm{rank}} \mid X_t] \le 1 \qquad P\textit{-a.s.}$$

Proof. Fix $x \in \mathcal{X}_t$ and $Q \in \mathcal{P}_t^{\mathrm{rank}}(x)$. By definition of $\mathcal{P}_t^{\mathrm{rank}}(x)$, the pushforward of Q under $r_t(x,\cdot)$ is the uniform distribution on $\{1,\dots,n_t+1\}$. Hence

$$\int e_t^{\mathrm{rank}}(x,z)\, Q(dz) = \int h_t(x, r_t(x,z))\, Q(dz) = \frac{1}{n_t+1} \sum_{r=1}^{n_t+1} h_t(x,r) \le 1.$$

Since this holds for every $Q \in \mathcal{P}_t^{\mathrm{rank}}(x)$, the function e_t^{rank} is a valid local e-certificate. The conditional form is then immediate from Theorem 2.3. □

Remark G.3 (How split conformal e-predictors fit). Suppose a local module computes calibration scores $(S_{t,1},\dots,S_{t,n_t})$ from predictable data already encoded in X_t, and computes a candidate test score S_{t,n_t+1} from (X_t, Z_t). If, under the local null, the tie-broken rank of S_{t,n_t+1} among the n_t+1 scores is conditionally uniform given X_t, then Theorem G.2 applies. This is exactly the exchangeable-rank mechanism behind conformal e-prediction (Gauthier et al., 2025). The workshop theorem does not need any additional details of the conformal construction; only the uniform-rank property matters.

G.3. Bounded-Risk Local Modules: An Exponential E-Certificate Interface

We next give a concrete local e-certificate for bounded-risk guarantees. This provides a direct interface between our graph calculus and local mean-risk modules. It is generic and does not depend on conformal machinery, but it is aligned with the bounded-risk viewpoint underlying conformal risk control and its anytime-valid variants (Angelopoulos et al., 2024; Hultberg et al., 2026).

Proposition G.4 (Exponential e-certificate for bounded mean risk)**.** *Fix a round $t \in [T]$. Let*

$$\ell_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0,1]$$

be a measurable local loss, let

$$\alpha_t : \mathcal{X}_t \to [0,1]$$

be a measurable predictable risk budget, and let

$$\lambda_t : \mathcal{X}_t \to [0,\infty)$$

be a measurable tuning parameter. Define the local risk-null class

$$\mathcal{P}_t^{\mathrm{risk}}(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \int \ell_t(x,z)\, Q(dz) \le \alpha_t(x) \right\}.$$

Define

$$e_t^{\mathrm{risk}}(x,z) := \exp\Big(\lambda_t(x)\,\ell_t(x,z) - \log\big(1-\alpha_t(x)+\alpha_t(x)e^{\lambda_t(x)}\big)\Big). \tag{11}$$

Then e_t^{risk} is a valid local e-certificate for the null class $\mathcal{P}_t^{\mathrm{risk}}$. Equivalently, with

$$E_t^{\mathrm{risk}} := e_t^{\mathrm{risk}}(X_t, Z_t),$$

every admissible law P satisfying $K_t^P(\cdot \mid X_t) \in \mathcal{P}_t^{\mathrm{risk}}(X_t)$ almost surely obeys

$$\mathbb{E}_P[E_t^{\mathrm{risk}} \mid X_t] \le 1 \qquad P\text{-a.s.}$$

Proof. Fix $x \in \mathcal{X}_t$ and $Q \in \mathcal{P}_t^{\mathrm{risk}}(x)$, and write

$$L := \ell_t(x, Z), \qquad Z \sim Q.$$

Then $L \in [0,1]$ and $\mathbb{E}_Q[L] \le \alpha_t(x)$. For every $\lambda \ge 0$ and every $u \in [0,1]$, convexity of the exponential implies

$$e^{\lambda u} \le (1-u)e^0 + ue^{\lambda} = 1 - u + ue^{\lambda}.$$

Applying this pointwise with $u = L$ and taking Q-expectations gives

$$\mathbb{E}_Q[e^{\lambda_t(x)L}] \le 1 - \mathbb{E}_Q[L] + \mathbb{E}_Q[L]e^{\lambda_t(x)} \le 1 - \alpha_t(x) + \alpha_t(x)e^{\lambda_t(x)},$$

because the right-hand side is increasing in $\mathbb{E}_Q[L]$ when $\lambda_t(x) \ge 0$. Therefore

$$\int e_t^{\mathrm{risk}}(x,z)\,Q(dz) = \frac{\mathbb{E}_Q[e^{\lambda_t(x)L}]}{1-\alpha_t(x)+\alpha_t(x)e^{\lambda_t(x)}} \le 1.$$

Since this holds for every $Q \in \mathcal{P}_t^{\mathrm{risk}}(x)$, the function e_t^{risk} is a valid local e-certificate. □

Remark G.5 (Interpretation of the exponential interface). Proposition G.4 is a local e-certificate for the mean-risk null $\mathbb{E}[\ell_t(X_t, Z_t) \mid X_t] \le \alpha_t(X_t)$. The construction is exact and self-contained: it uses only boundedness of the loss and a one-step exponential supermartingale inequality. It therefore provides a clean interface for any local risk-control module whose target guarantee can be expressed as a bounded conditional mean-risk constraint.

G.4. Compositional Corollary

We can now combine the abstract embedding theorem with the two module propositions above.

Proof of Corollary I.1. By Theorem G.2 and Theorem G.4, every realized node emits a valid local e-certificate for its designated local null. Applying Theorem G.1 to these local certificates yields the claim. □

Remark G.6 (Same-null joins remain essential). The current workshop theorem does *not* allow arbitrary mixing of coverage and risk semantics inside a single join stage. A same-null join may combine several conformal modules for the same downstream claim, or several bounded-risk modules for the same bounded-risk null, or more generally any collection of modules that have already been lifted to a common null. What is excluded in the present version is joining heterogeneous certificates that target different nulls without first passing to a common downstream hypothesis.

H. Scope Boundaries and Extensions

This appendix records the precise boundaries of the workshop theorem and several rigorous extension principles that do *not* require a new primitive theory. The theme is simple: whenever a seemingly broader setting can be reduced to the same-null, finite-constraint, finite-trace e-certificate setting of the main paper, the factorization theorem applies unchanged. What is left open are the settings for which no such reduction is presently available.

H.1. Multiple-Null Joins and Common-Null Reducibility

The main theorem only treats *same-null* joins. The sharp boundary is not "single semantic type" versus "multiple semantic types", but rather whether the certificates entering a join can be lifted to a *common downstream null*.

Definition H.1 (Common-null reducibility)**.** Let

$$(E_i, \mathcal{G}_i, \mathcal{Q}_i)_{i \in J}$$

be a finite family, where each E_i is a $\mathcal{Q}_i$-e-value on $\mathcal{G}_i$. We say that the family is *common-null reducible* if there exist:

1. a sigma-field $\mathcal{G} \subseteq \mathcal{F}$ with $\mathcal{G}_i \subseteq \mathcal{G}$ for every $i \in J$;
2. a null class $\mathcal{Q} \subseteq \mathcal{M}_1(\Omega, \mathcal{F})$;
3. lift operators
$$A_i : L^0_+(\Omega, \mathcal{G}_i) \to L^0_+(\Omega, \mathcal{G}), \qquad i \in J,$$

such that each lifted variable $A_i(E_i)$ is a $\mathcal{Q}$-e-value on $\mathcal{G}$.

Proposition H.2 (Reduction of common-null-reducible joins)**.** *Let Γ satisfy the hypotheses of Theorem 5.2. Suppose an execution trace contains a join stage whose incoming certificates are not same-null in their native form, but are common-null reducible in the sense of Theorem H.1. Then, after replacing that join stage by the corresponding lifted common-null family, the factorization theorem applies without modification.*

Proof. By common-null reducibility, the incoming certificates can be transported to a common sigma-field $\mathcal{G}$ and a common null class $\mathcal{Q}$. After this reduction, the join stage is literally a same-null join in the sense of Theorem 2.5. Theorem 5.2 can therefore be applied stagewise to the lifted family. Since the reduction occurs before the join is executed, no other part of the trace is altered. □

Remark H.3 (What remains open). The genuinely open problem is to characterize *intrinsic* multi-null joins: those for which no common-null reduction is available a priori. Such joins do not fit the current primitive algebra, and their treatment would require a new merging theory beyond same-null e-values.

H.2. Unbounded Stopping by Infinite-Horizon Extension

The workshop theorem is stated for bounded horizon and bounded stopping. At the level of nonnegative supermartingales, however, optional stopping is not the real obstruction. The real obstruction is the absence of an infinite-horizon primitive theory for the local and stagewise construction. The next proposition records the exact extension principle once an infinite-horizon factorized process has been built.

Proposition H.4 (Extension from bounded truncations to unbounded stopping)**.** *Let $(\mathcal{F}_t)_{t \geq 0}$ be a filtration, and let*

$$M = (M_t)_{t \geq 0}$$

be a nonnegative adapted process such that, for every $N \in \mathbb{N}$, the truncated process

$$M_t^{(N)} := M_{t \wedge N}, \qquad t \geq 0,$$

is an e-process. Then:

1. *M itself is a nonnegative supermartingale;*
2. *for every stopping time $\tau : \Omega \to \mathbb{N} \cup \{\infty\}$, the random variable*
$$M_\tau := \liminf_{N \to \infty} M_{\tau \wedge N}$$
satisfies
$$\mathbb{E}[M_\tau] \leq 1.$$

Proof. Fix $0 \le s \le t < \infty$, and choose $N \ge t$. Since $M^{(N)}$ is an e-process,

$$\mathbb{E}[M_t \mid \mathcal{F}_s] = \mathbb{E}[M_t^{(N)} \mid \mathcal{F}_s] \le M_s^{(N)} = M_s \qquad \text{a.s.}$$

Hence M is a nonnegative supermartingale.

Now let τ be any stopping time with values in $\mathbb{N} \cup \{\infty\}$. For each N, $\tau \wedge N$ is bounded, so optional stopping for the bounded e-process $M^{(N)}$ gives

$$\mathbb{E}[M_{\tau\wedge N}] \le M_0 = 1.$$

Because $M_{\tau\wedge N} \ge 0$, Fatou's lemma yields

$$\mathbb{E}[M_\tau] = \mathbb{E}\Big[\liminf_{N\to\infty} M_{\tau\wedge N}\Big] \le \liminf_{N\to\infty} \mathbb{E}[M_{\tau\wedge N}] \le 1.$$

□

Remark H.5 (Interpretation). Proposition H.4 shows that bounded stopping is not a conceptual limitation of e-processes. It is only a limitation of the current finite-horizon construction. Once an infinite-horizon factorized process exists, arbitrary stopping times are handled automatically by nonnegative supermartingale theory.

H.3. Infinite Constraint Families via Local Finite Activity

Appendix C is stated for finitely many bounded predictable constraints. The next proposition records a precise reduction from certain infinite-constraint models to the current theorem. The key idea is *local finite activity*: although the ambient constraint family may be infinite, the particular local certificate under study only interacts with a finite measurable subfamily on the realized history fiber.

Proposition H.6 (Finite-activity reduction for countable constraint classes)**.** *Fix $t \in [T]$. Let $\mathcal{J}_t$ be a countable index set, and let*

$$\phi_{t,j} : \mathcal{X}_t \times \mathcal{Z}_t \to \mathbb{R}, \qquad j \in \mathcal{J}_t,$$

be bounded measurable constraint functions. Define the infinite local null class

$$\mathcal{P}_t^\infty(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \int \phi_{t,j}(x,z)\, Q(dz) \le 0 \text{ for all } j \in \mathcal{J}_t \right\}.$$

Let P be an admissible law for this infinite-constraint model, and let $e_t : \mathcal{X}_t \times \mathcal{Z}_t \to [0,\infty)$ be a local e-certificate for $\mathcal{P}_t^\infty$. Assume there exists a measurable finite-valued selector

$$I_t : \mathcal{X}_t \to 2^{\mathcal{J}_t}$$

such that for P_{X_t}-almost every x,

1. *$e_t(x,\cdot)$ is pointwise admissible for the reduced finite null class*

$$\mathcal{P}_t^{I_t(x)}(x) := \left\{ Q \in \mathcal{M}_1(\mathcal{Z}_t) : \int \phi_{t,j}(x,z)\, Q(dz) \le 0 \text{ for all } j \in I_t(x) \right\};$$

2. *$e_t(x,\cdot)$ is valid for $\mathcal{P}_t^{I_t(x)}(x)$.*

Then there exists a measurable coefficient map

$$\lambda_{t,P} : \mathcal{X}_t \to [0,\infty)^{\mathcal{J}_t}$$

with finite support contained in $I_t(X_t)$ almost surely such that

$$e_t(X_t, Z_t) = 1 + \sum_{j\in\mathcal{J}_t} \lambda_{t,P,j}(X_t)\, \phi_{t,j}(X_t, Z_t) \qquad P\text{-a.s.}$$

Proof. For P_{X_t}-almost every x, the selector $I_t(x)$ is finite. By assumption, the reduced finite constraint family indexed by $I_t(x)$ satisfies the hypotheses of Appendix C, and $e_t(x, \cdot)$ is pointwise admissible and valid for the corresponding reduced null class. Therefore Appendix C yields a finite-dimensional representation on that fiber. The measurable-selection step used in Proposition C.10 applies unchanged because the active parameter space is finite-dimensional on each selected fiber. Extending the selected coefficient vector by zero outside $I_t(x)$ yields the claimed countable-indexed representation with finite support. □

Remark H.7 (What this does and does not solve). The proposition does not solve the full infinite-constraint problem. It shows only that whenever the relevant local certificate is controlled by a measurable finite active subset, the finite-constraint theorem already suffices. The truly new case is where no such finite active reduction exists.

H.4. Recursive Graphs Beyond Bounded Unfolding

The workshop theorem treats recursive workflows only after bounded unfolding into a finite DAG. The next proposition gives a precise extension principle by truncation: if deeper unfoldings produce consistent e-processes and those processes converge pointwise, then the limit inherits e-validity.

Proposition H.8 (Recursive truncation limit)**.** *Let $(\mathcal{F}_t)_{t\geq 0}$ be a filtration, and let*

$$M^{(d)} = (M_t^{(d)})_{t\geq 0}, \qquad d \in \mathbb{N},$$

be nonnegative adapted processes such that each $M^{(d)}$ is an e-process. Assume that for every fixed $t \geq 0$,

$$M_t^{(d)} \longrightarrow M_t \qquad \text{a.s. as } d \to \infty$$

for some nonnegative $\mathcal{F}_t$-measurable random variable M_t. Then the limit process $M = (M_t)_{t\geq 0}$ is also an e-process.

Proof. Fix $0 \leq s \leq t$. Since each $M^{(d)}$ is an e-process,

$$\mathbb{E}[M_t^{(d)} \mid \mathcal{F}_s] \leq M_s^{(d)} \qquad \text{a.s. for every } d.$$

Because the sequence $(M_t^{(d)})_{d\geq 1}$ is nonnegative, conditional Fatou's lemma gives

$$\mathbb{E}[M_t \mid \mathcal{F}_s] = \mathbb{E}\left[\liminf_{d\to\infty} M_t^{(d)} \middle| \mathcal{F}_s\right] \leq \liminf_{d\to\infty} \mathbb{E}[M_t^{(d)} \mid \mathcal{F}_s] \leq \liminf_{d\to\infty} M_s^{(d)} = M_s \qquad \text{a.s.}$$

Hence M is a nonnegative supermartingale, i.e. an e-process. □

Remark H.9 (Interpretation). Proposition H.8 shows that recursive systems can be handled by finite-depth approximation whenever the unfolded factorized processes converge. What remains open is an intrinsic factorization theorem for recursion *without* passing through a convergent sequence of finite DAGs.

H.5. Beyond E-Certificates: Threshold-Generated Families

The main paper is formally stated for graph-indexed e-certificates and e-processes. A natural question is how to extend the theory to broader decision objects such as set-valued predictors, abstention rules, or nested acceptance regions. The correct boundary is whether the family can be generated by thresholding an e-process.

Definition H.10 (Threshold-generated certificate family and e-envelope)**.** Let $(D_\alpha)_{\alpha\in(0,1)}$ be a family of measurable events on $(\Omega, \mathcal{F})$. A nonnegative random variable M is called an *e-envelope* for (D_α) if, for every $\alpha \in (0, 1)$,

$$D_\alpha = \left\{ M < \frac{1}{\alpha} \right\}.$$

More generally, a family of decision objects is called *threshold-generated by M* if each member of the family is obtained by applying a measurable monotone threshold rule to M.

Proposition H.11 (Envelope reduction)**.** *Let M be a global e-process produced by a universally valid factorized calculus. Let $(D_\alpha)_{\alpha\in(0,1)}$ be threshold-generated by M, and in particular suppose*

$$D_\alpha = \left\{ M_\tau < \frac{1}{\alpha} \right\}$$

for some stopping time τ. Then the threshold tests

$$\varphi_\alpha := \mathbf{1}\left\{ M_\tau \geq \frac{1}{\alpha} \right\}$$

satisfy

$$\mathbb{E}[\varphi_\alpha] \leq \alpha$$

under every admissible law. Consequently, any certificate family admitting an e-envelope inherits the compositional validity of that envelope.

Proof. Since M is an e-process, M_τ is an e-value at every bounded stopping time τ, and more generally at every stopping time covered by Theorem H.4. Hence

$$\mathbb{E}[M_\tau] \leq 1.$$

By Markov's inequality,

$$\mathbb{P}\left(M_\tau \geq \frac{1}{\alpha} \right) \leq \alpha\mathbb{E}[M_\tau] \leq \alpha.$$

This is exactly

$$\mathbb{E}[\varphi_\alpha] \leq \alpha.$$

The final statement follows because every member of a threshold-generated family is obtained from M_τ by such a threshold rule. □

Remark H.12 (What remains open). The open problem is not how to use an e-envelope once it exists, but how to construct one for broad classes of non-e certificate families. In particular, an exact envelope theorem for arbitrary graph-indexed confidence sets or abstention families is outside the scope of the present paper.

Remark H.13 (Summary of what is genuinely outside scope). The main theorem does *not* currently solve:

1. intrinsically multi-null joins without common-null reduction;
2. infinite-horizon graph construction without an already defined limiting e-process;
3. infinite constraint systems lacking a measurable finite active reduction;
4. recursive workflows without convergent finite truncations;
5. arbitrary certificate families lacking an e-envelope.

Each item represents a genuinely new primitive problem rather than a small technical modification of the current argument.

I. Additional Discussion, Scope, and Conformal Corollary

Theorems 4.1 and 5.2 together identify the right primitive for uncertainty in adaptive workflows. The first theorem shows that, once a module is conditioned on its predictable history, admissible local uncertainty objects collapse to finite-dimensional e-certificates generated by the local constraint system. The second theorem shows that, once those local certificates are attached to an adaptive execution graph, universal composability leaves only a small primitive algebra: a serial game-martingale update, weighted arithmetic averaging at same-null joins, and, under the strengthened lift assumption, adjusters across filtration changes (Vovk & Wang, 2024; Wang, 2025; Choe & Ramdas, 2026). In this sense, the paper is a representation theorem for graph-indexed uncertainty, not merely a sufficiency result.

The workshop version is deliberately narrow. We assume bounded horizon, bounded predictable constraints, trace-generated normal form, same-null joins, and, for the sharpened lift statement, value-based universal lifts. We do *not* treat joins across multiple downstream nulls, unbounded stopping, infinite constraint classes, or an envelope theorem for arbitrary certificate families. These are not cosmetic omissions: each would require a genuinely new primitive theory. The negative result of Theorem 3.1 already shows why one cannot hope to start from local marginal coverage or local mean-risk statements alone.

Corollary I.1 (Conformal and bounded-risk modules compose through the e-certificate interface)**.** *Assume the hypotheses of Theorem 5.2. Suppose every realized node emits either:*

1. *a rank-based local conformal e-certificate obtained through the exchangeable-rank interface of Appendix G; or*
2. *a bounded-risk local exponential e-certificate for a conditional mean-risk null as in Appendix G.*

Assume further that, within each join stage, the lifted stage inputs certify a common downstream null. Then the factorized graph calculus yields a global e-process under every admissible law. In particular, local conformal e-predictors and bounded-risk modules can be composed inside an adaptive workflow once they are expressed through the common local e-certificate interface (Gauthier et al., 2025; Angelopoulos et al., 2024; Hultberg et al., 2026).

Proof sketch. Appendix G proves that both module families produce valid local e-certificates. Theorem G.1 then implies that conditionally valid lifts preserve stagewise e-validity, same-null weighted averaging preserves the common-null guarantee at each join, and the serial game-martingale update turns the resulting stage sequence into a global e-process. □

The corollary should be read as an interface theorem, not as a claim that arbitrary conformal or risk-control procedures may be joined indiscriminately. The same-null requirement remains essential: what is excluded in the present version is the direct joining of heterogeneous certificates that target different downstream nulls without first passing to a common hypothesis.