FINITE TEACHING, INFINITE HORIZONS: SHARP GUARANTEES FOR MIN-AGGREGATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We show when errors on a small teaching set certify a learned graph algorithm on unseen graphs and at arbitrary execution lengths. For nonnegative, bias-free min-aggregation networks of arbitrary width, we compute the exact worst-case relative error against Bellman–Ford from the weights alone. Just $K + 2$ labeled endpoints identify exact K-step execution. With normalized mixing and bounded edge coefficients, their teaching error equals the global error. For shared networks with fixed normalized connections, we determine the optimal error amplification over all execution lengths. It is finite exactly when teaching exposes every eventually observable channel. This factor also guides readout design. For irreducible connections with q message channels, an optimized readout achieves one-step amplification within a factor $1 + 1/q$ of optimal. Matching hitting and propagation bounds certify multi-step optimality. In three paired fits, adding one teaching graph reduces global errors initially between 19.92% and 47.99% to below 0.163% in every case.

1 INTRODUCTION

When can errors on a small teaching set certify a learned graph algorithm on unseen graphs and at arbitrary execution lengths? We answer this question for nonnegative, bias-free min-aggregation networks. Each message channel takes its own minimum and may select a different neighbor. Fitting shortest-path examples therefore need not recover the Bellman–Ford (BF) rule. We determine what teaching error guarantees and how to strengthen that guarantee.

Teaching must bridge two gaps: from examples to unseen graphs at the taught length, and from that length to later updates. Exact identification alone does not control either gap for an approximate fit. Even when the examples identify the exact algorithm, a small nonzero residual can hide a large error.

Existing BF guarantees use small teaching sets and near-optimal regularized objectives (Nerem et al., 2026; Wittig et al., 2026). We analyze the weights directly in a model with nonnegative, bias-free linear maps before and after coordinatewise minimum aggregation. Two weight recurrences give the exact worst-case error and connect three contributions:

1. **Certifying unseen graphs.** We compute attainable output-to-target extrema and identify exact BF from $K+2$ fixed labeled endpoints, regardless of width. A sharp coefficient-balance bound determines when approximate fitting also gives a reliable certificate (Sections 3–4).
2. **Certifying longer execution.** For fixed normalized shared connections, we determine when finite teaching controls every execution length. We compute the optimal amplification to any prescribed accuracy, including for periodic connections, and derive sharp teaching-length formulas for structured classes (Section 5).
3. **Improving the guarantee by readout design.** With irreducible connections fixed, we optimize how hidden states are read out. One-step design achieves an additive $1 + \eta$ approximation at computational tolerance $\eta > 0$. Matching hitting and propagation bounds certify multi-step global optimality (Section 6).

Figure 1 maps the main dependencies: Theorem 1 provides the foundation, and Proposition 8 connects long-horizon guarantees to readout design. Section 2 compares prior work; Section 7 illustrates which example to add and how long to teach.

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

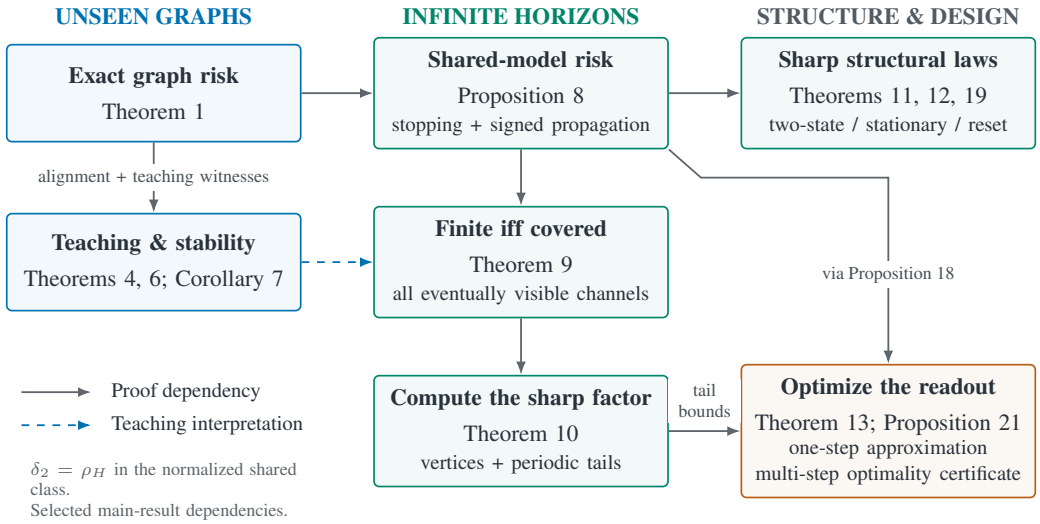

Figure 1: Dependencies among the main results. Solid arrows mark proof dependencies; the dashed arrow gives the teaching interpretation in the normalized shared class.

2 Related Work

From BF learning guarantees to exact risk. Algorithmic alignment motivates networks that follow dynamic programming (Xu et al., 2020; Dudzik & Veličković, 2022; Xu et al., 2021b). The closest BF guarantees use small teaching sets and near-optimal sparsity-regularized objectives (Nerem et al., 2026; Wittig et al., 2026). Vector-min underestimation and teaching inputs that expose it are already identified by Wittig et al. (2026). We derive attainable two-sided risk for arbitrary parameters in the nonnegative, bias-free class. For normalized shared networks, we then quantify amplification across all execution lengths. Retaining vector state between updates introduces channel exposure and cancellation beyond repeated calls to an accurate scalar BF module.

Identification versus stable teaching. Machine teaching studies examples that identify a target within a hypothesis class (Goldman & Kearns, 1995; Zhu et al., 2018). Recent graph-instruction learning and algorithmic-circuit studies also establish generalization in different architectures (Qharabagh et al., 2026; He et al., 2026). Our question is quantitative: how much algorithmic error can a small nonzero teaching residual hide? The sharp stability bound separates exact identification from approximate certification.

Graph expressivity and generalization. Graph-discrimination results (Xu et al., 2019; Morris et al., 2019) and statistical generalization bounds (Garg et al., 2020; Maskey et al., 2026) connect architecture, graph structure, and prediction. Size-generalization analyses identify failures caused by changing local structures (Yehudai et al., 2021). We instead take a deterministic supremum over admissible graphs for a known target algorithm. The resulting channel chain describes internal computation after input graphs have already been eliminated from the risk calculation.

Infinite behavior and readout design. Weighted-automaton metrics study infinite behaviors through seminorms and spectral conditions (Balle et al., 2017). Hitting probabilities, stopping LPs, and polyhedral error bounds supply established tools (Levin & Peres, 2017; Kallenberg, 2013; Peña et al., 2021). Our finite-to-infinite ratio combines stopping with signed propagation, including cancellation. Unlike fastest-mixing chain design (Boyd et al., 2004), we fix the connections and optimize a realizable readout. Positive-system peak optimization provides a related design perspective (Balakrishnan & Boyd, 1992; Ogura et al., 2020; Zhao et al., 2023); our guarantee controls the complete MinAgg risk ratio. One-step design is near-optimal, and matching hitting and propagation bounds certifies multi-step optimality. Appendix K develops the broader literature connections.

3 EXACT ERROR ON EVERY GRAPH

To determine what teaching must observe, we first compute the error it can miss. Let G be any finite nonempty undirected graph with finite nonnegative node values h_v and edge weights w_{uv}, including a zero-weight self-loop at every node. The synchronous BF target is

$$x_v^{(0)} = h_v, \qquad x_v^{(t)} = \min_{u \in N(v)} (x_u^{(t-1)} + w_{uv}), \qquad \Gamma^K(G, v) = x_v^{(K)}. \tag{1}$$

Initial values are arbitrary nonnegative numbers, not just a one-source initialization. Fix $K \geq 1$, scalar endpoint widths $d_0 = d_K = 1$, and arbitrary intermediate state and message widths. Our network is

$$h_v^{(0)} = h_v, \quad z_v^{(t)} = \min_{u \in N(v)} (A_t h_u^{(t-1)} + c_t w_{uv}), \quad h_v^{(t)} = B_t z_v^{(t)}, \qquad F_\theta = h_v^{(K)}, \tag{2}$$

where $A_t \in \mathbb{R}_{\geq 0}^{q_t \times d_{t-1}}$, $B_t \in \mathbb{R}_{\geq 0}^{d_t \times q_t}$ and $c_t \in \mathbb{R}_{\geq 0}^{q_t}$. There are q_t message channels and d_t state channels: A_t forms messages, each message channel takes its own minimum, and B_t mixes the result into the next state. ReLUs leave nonnegative activations unchanged, so nonnegative linear/ReLU stacks reduce to these effective maps. We exclude biases and skip connections. Define the risk by

$$\rho_*(\theta) = \sup_{G, v : \Gamma^K(G,v) > 0} \frac{|F_\theta(G, v) - \Gamma^K(G, v)|}{\Gamma^K(G, v)}. \tag{3}$$

The risk is the largest relative error over all admissible graphs and nodes with positive BF target, rather than an average over a test distribution. All theorems use real arithmetic.

Inactive support. First identify coordinates that can ever carry a nonzero value. The following binary masks track them through the layers:

$$s_0 = 1, \qquad a_t = \mathbf{1}\{A_t s_{t-1} > 0\}, \qquad s_t = \mathbf{1}\{B_t a_t > 0\}. \tag{4}$$

If $a_{t,i} = 0$, that message is always zero because its self-loop supplies a zero candidate. Set its row of A_t, entry of c_t, and column of B_t to zero. Also clear columns that read inactive states. This leaves the network function unchanged. We use these cleaned parameters below and retain every strictly positive weight on the active support (Lemma 14).

We now compute the two extremes of the output-to-target ratio. The lower recurrence tracks the smallest ratio; the upper coefficients track contributions from a walk's initial value and each edge:

$$\ell_0 = 1, \qquad \ell_t = B_t \min(A_t \ell_{t-1}, c_t), \qquad L = \ell_K, \tag{5}$$

$$C_t = B_t A_t, \qquad \gamma_0 = C_K \cdots C_1, \qquad \gamma_j = C_K \cdots C_{j+1} B_j c_j \ (1 \leq j \leq K). \tag{6}$$

Set $U = \max_{0 \leq j \leq K} \gamma_j$. Empty products are identities.

Theorem 1 (Exact graph-uniform endpoints)**.** *For every network* (2)*,*

$$\inf_{\Gamma^K > 0} \frac{F_\theta}{\Gamma^K} = L, \qquad \sup_{\Gamma^K > 0} \frac{F_\theta}{\Gamma^K} = U, \qquad \rho_* = \max\{|L - 1|, |U - 1|\}. \tag{7}$$

Both extrema are attained. One fixed two-node graph attains L for all parameters and depths; each γ_j is attained by a K-edge path with parameter-dependent finite initial values. If $\Gamma^K = 0$, then $F_\theta = 0$.

Two nodes with initial values $(0, 1)$ and connecting weight one attain L. Paths isolate the initial-value and edge coefficients that determine U (Appendix B.1). Both endpoints are computed from the weights in arithmetic time linear in the parameter count, without searching over graphs.

3.1 WHAT EXACTNESS SAYS ABOUT EXTRA CHANNELS

What must the channels compute when the risk is small? Let $p_0 = 1$, $p_t = C_t p_{t-1}$ and $r_t = C_K \cdots C_{t+1} B_t$. The vector p_t tracks the propagated initial-value scale, while r_t measures each message's influence on the output. Define $m_t = r_t |A_t p_{t-1} - c_t|$ to measure the mismatch between state and edge coefficients. A message affects the output exactly when $r_{t,i} > 0$.

Theorem 2 (Quantitative alignment)**.** *For each t,*

$$m_t \le p_K + \gamma_t - 2L \le 2(U - L) \le 4\rho_*, \qquad \rho_* \le |p_K - 1| + \sum_{t=1}^{K} m_t. \tag{8}$$

Corollary 3 (Exact channel alignment)**.** *Zero risk is equivalent to*

$$p_K = 1, \qquad c_{t,i} = (A_t p_{t-1})_i \quad \textit{on every output-relevant message}. \tag{9}$$

Every relevant hidden state is then a scaled BF copy: $h_{v,j}^{(t)} = (p_t)_j x_v^{(t)}$ whenever $(C_K \cdots C_{t+1})_j > 0$.

Exact execution does not require a single active channel. Dense positive matrices can satisfy these conditions: every relevant channel then carries a scaled copy of the same BF state. The proof sums the nonnegative layerwise deficits to obtain the alignment bounds (Appendices B.2–B.3).

4 Finite teaching: exactness versus stability

The paths attaining the upper endpoint depend on the weights. Teaching requires examples chosen in advance. Fix $\beta > 1$ and use K-edge paths $P_j(\beta)$, $j = 0, \dots, K$. On path j, exactly one of the starting value and K edge weights is one: coordinate j. The others are zero, and all remaining node initial values are β. Each final endpoint has BF target one. Let T_β contain these $K + 1$ labeled endpoints and the stationary two-node endpoint. With $\alpha_t = A_t p_{t-1}$, their outputs are

$$\tau_0 = p_K, \qquad \tau_j = r_j \min(c_j, \beta\alpha_j)\ (1 \le j \le K), \qquad F_\theta(\text{stationary}) = L. \tag{10}$$

A fixed path can clip a large edge coefficient at $\beta\alpha_j$, so its output need not reveal the full upper coefficient (Lemma 15). This clipping is why exact identification and reliable approximate fitting require separate arguments.

Theorem 4 (Exact teaching)**.** *For every $\beta > 1$, fitting all $K + 2$ endpoints of T_β exactly is equivalent to exact BF on every admissible graph, without a norm or regularization condition.*

Fitting the stationary and input-path endpoints forces $L = p_K = 1$. The alignment bounds then rule out deficient edge coefficients on relevant messages. Any excess coefficient increases its path output because $\beta > 1$. Fitting those paths removes the remaining mismatches.

Why approximate fitting is different. For $K = 1$, take

$$A = (1, \varepsilon)^\top, \quad B = (1, 1), \quad c = (1, 1)^\top, \qquad 0 \le \varepsilon \le 1. \tag{11}$$

At $\varepsilon = 0$ the second message is inactive; at any $\varepsilon > 0$ it can double the target.

Proposition 5 (Unstable exact identification)**.** *The family* (11) *has $\rho_*(0) = 0$ and $\rho_*(\varepsilon) = 1$ for every $\varepsilon > 0$, while its maximum teaching error equals $\min(\beta\varepsilon, 1)$. Errors on any fixed finite set of labeled nodes tend to zero.*

The unit-bounded witness has initial values $(0, 1)$ and edge weight ε. At the second node, the target is ε and the prediction is 2ε. Here the relative error stays large because the target shrinks with ε. Delayed channel exposure causes a different failure at unit target (Appendix C.3).

The missing quantitative control. We need to control how much the paths can clip. Let κ bound the ratio of each relevant edge coefficient to its state coefficient, and let δ_β be the maximum teaching error:

$$\kappa = \max\{1, \max_{t,i:r_{t,i}>0} c_{t,i}/\alpha_{t,i}\}, \qquad \nu_\beta = \max\{1, (\kappa - 1)/(\beta - 1)\},$$
$$\delta_\beta = \max\{|L - 1|, \max_{0 \le j \le K} |\tau_j - 1|\}. \tag{12}$$

The inner maximum is omitted on empty support; cleanup gives positive denominators elsewhere.

Theorem 6 (Sharp graph-input stability)**.** *For every cleaned network,*

$$\delta_\beta \le \rho_* \le (2\nu_\beta - 1)\delta_\beta. \tag{13}$$

In particular, $\kappa \le \beta$ gives $\rho_* = \delta_\beta$. For every $\kappa > \beta > 1$, strictly positive one-layer, two-message networks attain the factor with arbitrarily small positive teaching error.

The factor measures how much error path clipping can hide. The proof bounds the clipped excess and uses $p_K - L$ to control negative mismatches. The bound has no explicit width or graph-size factor; κ can depend on depth. In (11), $\kappa = 1/\varepsilon$ grows without bound, explaining the instability. A sufficient condition for equality of teaching and global risk is $\kappa \le \beta$.

Corollary 7 (A priori stable class)**.** *The class*

$$A_t\mathbf{1} = B_t\mathbf{1} = \mathbf{1}, \qquad 0 \le c_t \le \beta\mathbf{1} \tag{14}$$

satisfies $\rho_* = \delta_\beta$ for every parameter choice. Conversely, after removing unused coordinates, any model with $p_K = 1$ and $\kappa \le \beta$ has a functionally identical representation of this form, with no greater intermediate widths.

The converse rescales hidden coordinates by positive diagonal matrices without changing the output. Row normalization applies to channel mixing, not to aggregation across graph neighbors, which still uses a minimum. In this class, the teaching residual itself is the exact global risk.

5 From finite teaching to arbitrary execution length

We have controlled unseen-graph error at the taught length. The next question is whether the guarantee survives further updates. Here a *horizon* means a number of execution steps. Later steps can expose channels that short teaching barely reached. To study this effect, retain a vector hidden state and reuse the same weights at each update:

$$h_v^0 = eh_v, \qquad h_v^t = B \min_{u \in N(v)} (Ah_u^{t-1} + cw_{uv}), \qquad F_K(v) = rh_v^K. \tag{15}$$

Let ρ_K be the worst-case relative error over graphs after K updates, and let $\rho_\infty = \sup_{K \ge 1} \rho_K$. This supremum controls every finite execution length; it does not require the state to converge. General nonnegative shared models admit a finite test for exactness at all lengths (Theorem 17). To quantify nonzero error, we fix

$$e = \mathbf{1}, \quad A \in \mathbb{R}_{\ge 0}^{q \times d}, \quad B \in \mathbb{R}_{\ge 0}^{d \times q}, \quad A\mathbf{1} = B\mathbf{1} = \mathbf{1}, \quad r \ge 0, \quad r\mathbf{1} = 1, \tag{16}$$

and vary only the edge coefficients $c \in [0,2]^q$. Expanding H shared updates into H layers gives the stable class of Corollary 7 with $\beta = 2$. Hence $H + 2$ teaching endpoints reveal ρ_H exactly. We seek the smallest factor that bounds all-length error by this observed error:

$$C_H^\star = \inf\{C \ge 1 : \rho_\infty(c) \le C\rho_H(c) \text{ for every } c \in [0,2]^q\}, \tag{17}$$

The constant is infinite if no finite factor works. It applies to every allowed c, so teaching length can be chosen before learning these coefficients. The same risk formulas evaluate learned coefficients afterward (Section 7).

5.1 An exact message-chain representation

The exact endpoint formulas remove the input graph from the calculation. What remains is a Markov chain on message channels: $S = AB$ is its transition matrix and $\mu = rB$ its initial distribution. Set $\xi = c - \mathbf{1}$ to record deviations from the correct edge coefficient. Message time $j = 0$ corresponds to the first update. Define $z_0 = 0$ and

$$z_{k+1}(\xi) = \max\{-\xi, Sz_k(\xi)\}, \qquad p_H(\xi) = \max\{\mu z_H(\xi), \max_{0 \le j < H} \mu S^j \xi\}. \tag{18}$$

Proposition 8 (Exact risk as stopping and propagation)**.** *We have $\rho_H(c) = p_H(\xi)$ and*

$$\rho_\infty(c) = p_\infty(\xi) = \max\{\mu z_\infty(\xi), \sup_{j \ge 0} \mu S^j \xi\}, \qquad z_\infty = \lim_{k \to \infty} z_k. \tag{19}$$

Moreover, $\rho_H = \max_{K \le H} \rho_K$.

The two terms capture different errors. For underestimation, a stopping rule can wait for a channel and collect $-\xi_i$ there; never stopping earns zero. For overestimation, the risk uses the signed average $\mu S^j \xi$ at a chosen time j. Positive and negative deviations can cancel in this average. A hitting calculation alone therefore misses part of the risk. Both p_H and p_∞ extend to convex, positively homogeneous functions on $\mathbb{R}^q$, bounded by $\|\xi\|_\infty$.

5.2 WHICH FINITE TEACHING HORIZONS SUFFICE?

A channel is observed when it has positive probability under μS^j. Compare the channels ever observed with those observed during the H teaching steps:

$$\mathcal{O} = \{i : \exists j \geq 0,\ (\mu S^j)_i > 0\}, \qquad \mathcal{O}_H = \{i : \exists j < H,\ (\mu S^j)_i > 0\}. \tag{20}$$

Theorem 9 (Finite amplification iff channel coverage)**.** *$C_H^\star < \infty$ if and only if $\mathcal{O}_H = \mathcal{O}$. Writing $a_i(H) = \max_{j<H}(\mu S^j)_i$, a finite case satisfies*

$$C_H^\star \leq \frac{2}{\min_{i\in\mathcal{O}} a_i(H)}. \tag{21}$$

The earliest finite horizon is $H_{\mathrm{cov}} = 1 + \max_{i\in\mathcal{O}} \mathrm{dist}(\mathrm{supp}\,\mu, i) \leq \min(|\mathcal{O}|, d)$.

Proof idea. If a channel is missed, $\xi = -e_i$ has zero teaching risk but positive later risk. Conversely, if $p_H(\xi) = \delta$, stopping at a fixed time gives $\mu S^j(-\xi)_+ \leq \delta$. Together with $\mu S^j \xi \leq \delta$, this implies $\mu S^j(\xi)_+ \leq 2\delta$. Positive exposure therefore bounds each negative coordinate by δ/a_i and each positive one by $2\delta/a_i$. The theorem follows from $p_\infty \leq \|\xi\|_\infty$.

Coverage ensures that no relevant channel is completely hidden. It does not ensure a small amplification factor: a channel reached with very low probability can still hide a large error. We therefore need the size of $C_H^\star$, not just whether it is finite.

5.3 THE OPTIMAL CONSTANT, INCLUDING PERIODIC CONNECTIONS

Discard unreachable coordinates and collect error directions with teaching risk at most one in $\mathcal{P}_H = \{\xi : p_H(\xi) \leq 1\}$.

Theorem 10 (Optimal amplification and effective computation)**.** *If coverage holds, $\mathcal{P}_H$ is a compact full-dimensional polytope and*

$$C_H^\star = \max_{\xi\in\mathrm{vertices}(\mathcal{P}_H)} p_\infty(\xi). \tag{22}$$

For rational A, B, r and any rational $\eta > 0$, a finite algorithm returns rational $C^- \leq C_H^\star \leq C^+$ with $C^+ - C^- \leq \eta$. No irreducibility or aperiodicity assumption is required.

How the computation closes the infinite horizon. The computation has two parts. Finite stopping policies describe the polytope of teaching errors, and convexity reduces the search to its vertices. At each vertex, a linear program computes the infinite stopping value exactly:

$$\mu z_\infty(\xi) = \min\{\mu z : z \geq 0,\ z \geq -\xi,\ z \geq Sz\}. \tag{23}$$

The remaining sequence $\mu S^j \xi$ may oscillate. We separate its periodic phases and bound each uncomputed tail. These bounds and a finite prefix enclose the supremum (Appendix F.3). The algorithm computes certified additive approximations; exact threshold equality is outside its guarantee. Vertex enumeration can be exponential, and slow transitions require long prefixes.

5.4 MIXED ERRORS DETERMINE THE SHARP TEACHING LENGTH

The general algorithm computes the factor. A two-channel example shows why it can be large and how longer teaching reduces it. Consider

$$S = \begin{pmatrix} 1-p & p \\ 0 & 1 \end{pmatrix}, \quad \mu = (1-s, s), \quad 0 < p \leq 1, \quad 0 \leq s < 1,$$
$$w = 1 - s, \qquad h = 1 - w(1-p)^{H-1}. \tag{24}$$

Here h is the probability of having reached the absorbing channel by the final teaching decision.

Theorem 11 (Sharp two-state amplification)**.** *For $h > 0$,*

$$C_H^\star = \frac{w+1-h}{wh}, \qquad C_{H,-}^\star = \frac{1}{h}, \qquad \frac{C_H^\star}{C_{H,-}^\star} = 1 + (1-p)^{H-1}, \tag{25}$$

where $C_{H,-}^\star$ restricts to $\xi \leq 0$. The full constant is attained by $\xi = t(-z^\star, 1)$ with $z^\star = h/(w+1-h)$ and $t = 1/\max(1, z^\star)$. For $p > 0, h = 0$, $C_H^\star = \infty$; for $p = 0$ or $s = 1$, $C_H^\star = 1$.

For $\xi = (-a, b)$ with $a, b \geq 0$, cancellation gives

$$\rho_H = \max\{wa, hb - (1-h)a\}, \qquad \rho_\infty = \max\{wa, b\}.$$

In the finite expression, the negative first coordinate partly hides the positive second one. At long lengths, that cancellation disappears. Balancing the two finite-risk terms gives the worst direction. The other sign patterns amplify by at most $1/h$, so a purely underestimating model does not attain the full worst case.

For $s = 0$ and $0 < p < 1$, the formula simplifies to

$$C_H^\star = \frac{1 + (1-p)^{H-1}}{1 - (1-p)^{H-1}},$$
$$H_{\min}(1+\epsilon) = 1 + \left\lceil \frac{\log(\epsilon/(2+\epsilon))}{\log(1-p)} \right\rceil \quad (\epsilon > 0). \tag{26}$$

The constant at $H = 1$ is infinite. At $p = 1/64$, $C_2^\star = 127$, and limiting amplification to 1.1 requires exactly 195 updates. Mixed errors increase the sharp factor and the exact cutoff. Both cases retain the asymptotic $O(p^{-1}\log(1/\epsilon))$ order for small p, ϵ.

Mixing location and periodicity. Let P be the matrix in (24), $r = (1, 0)$ and $0 < p < 1$. The choices $(A, B) = (P, I)$ and (I, P) have the same $BA = AB = P$, but different initial channel distributions $\mu = rB$. Their one-step constants are infinity and $2/p$. Thus identical spectra need not give identical teaching guarantees: the first update observes different channels. Mixing is not necessary either. A deterministic three-cycle with $\mu = (1, 0, 0)$ has infinite amplification for $H < 3$ and factor one for $H \geq 3$, when every channel has been observed.

5.5 Structured multi-channel teaching laws

For more channels, explicit formulas can avoid vertex enumeration. We obtain them for stationary readouts and reset connections. Let $T_i = \inf\{j \geq 0 : X_j = i\}$ be the first visit to channel i, and let $h_i(H) = \Pr_\mu(T_i < H)$ be its probability of being visited during teaching.

Theorem 12 (Stationary readout)**.** *If $\mu S = \mu$, let w_i be the μ-mass of the closed irreducible class containing i, for $\mu_i > 0$. For every $H \geq 1$,*

$$C_H^\star = \max_{i:\mu_i>0} \frac{w_i}{h_i(H)}, \qquad h_i(H) = 1 - \mu_{-i}(S_{-i,-i})^{H-1}\mathbf{1}. \tag{27}$$

No reversibility or aperiodicity assumption is needed.

Stationarity keeps $\mu S^j \xi = \mu\xi$ fixed, so teaching already sees all overestimation. Only the stopping term can grow, at a rate controlled by channel-hitting probabilities. A single negative coordinate attains the ratio, even though the formula allows mixed costs. The condition $rBA = r$ ensures $\mu S = \mu$, including when $q \neq d$.

Nonstationary readout retains a separate cancellation cost. A reset connection holds its current channel with probability $1 - \alpha$ and otherwise draws a channel from a positive probability row π. For sparse readout, Theorem 19 gives the exact constant as the maximum of an inverse hitting probability and a mixed-cancellation term. For $\alpha = 1/4$, $\pi = (1/10, 4/5, 1/10)$, $\mu = (1/2, 0, 1/2)$ and $H = 2$, these terms equal 5 and 7.

A sharp architecture-family rule. Fix $q \geq 2$, $0 < \alpha \leq 1$ and $0 < \beta \leq 1/q$. Over all connections with $\mu = \pi$ satisfying $\pi S = \pi$, $\pi_i \geq \beta$ and $S_{ij} \geq \alpha\pi_j$,

$$\sup_{S,\pi} C_H^\star = \frac{1}{1 - (1-\beta)(1-\alpha\beta)^{H-1}}. \tag{28}$$

A lower bound on each channel's per-step hitting probability proves the bound; reset connections with $\min_i \pi_i = \beta$ attain it. For uniform π, the worst teaching length for factor $1 + \epsilon$ is $\Theta((q/\alpha)\log(1/\epsilon))$. The attaining chain has spectral gap α and starts stationary, yet teaching must still reach rarely observed channels. Appendix I derives these exact MinAgg consequences of classical Markov properties (Levin & Peres, 2017). We next optimize the readout.

6 DESIGNING READOUT FOR RELIABLE TEACHING

Readout determines which hidden channels influence the output. We now choose it to reduce error amplification, while keeping the internal connections fixed. Let $A \in \mathbb{R}_{\geq 0}^{q \times d}$ and $B \in \mathbb{R}_{\geq 0}^{d \times q}$ be row-stochastic, with $S = AB$ irreducible and $q \geq 2$. Choose the state readout $r \in \Delta_d$, the probability simplex. The message readout is then $\mu = rB$; it cannot be chosen independently of B. Write $C_H(r) = C_H^\star(A, B, r)$ and $C_{H,\mathrm{opt}} = \inf_{r \in \Delta_d} C_H(r)$. The chosen readout has a guarantee valid for every $c \in [0, 2]^q$.

Theorem 13 (Near-optimal one-step readout)**.** *For $\mu = rB > 0$, define*

$$\mathcal{U}(r) = \max\left\{ \frac{1}{\min_i \mu_i},\ 2 \sup_{j \geq 0, i} \frac{(\mu S^j)_i}{\mu_i} \right\}. \tag{29}$$

Then $C_1(r) \leq \mathcal{U}(r) \leq C_1(r) + 1$. A minimizing readout $r_{\mathcal{U}}$ exists and satisfies

$$C_1(r_{\mathcal{U}}) \leq C_{1,\mathrm{opt}} + 1, \qquad \frac{C_1(r_{\mathcal{U}})}{C_{1,\mathrm{opt}}} \leq 1 + \frac{1}{q}. \tag{30}$$

For rational connections and any rational $\eta > 0$, finitely many rational LPs with certified periodic tails produce $\widehat{r}$ with $C_1(\widehat{r}) \leq C_{1,\mathrm{opt}} + 1 + \eta$. No reversibility or aperiodicity is required.

The upper bound $\mathcal{U}$ controls both initial channel exposure and later propagation. The exact propagation term is $2 \max_i (\mu S^j)_i / \mu_i - \min_i (\mu S^j)_i / \mu_i$. Dropping the minimum increases the bound by at most one. In return, testing a fixed threshold becomes a set of linear constraints on r. Bisection and bounds on the periodic tails then give the near-optimal design (Appendix J).

Explicit improvements and multi-step optimality. For a non-reset four-channel reversible star, an explicit readout has $C_1 = 4.6$, versus at least 6.1333 for uniform readout and 15 for stationary readout. This design minimizes the envelope and satisfies the approximation guarantee for the true constant.

For $H > 1$, a channel-hitting LP lower-bounds the best amplification. A propagation certificate proves that a candidate attains this bound for all mixed costs by comparing later channel distributions with those observed during teaching. On a non-reset three-channel cycle, the bounds match at $C_{2,\mathrm{opt}} = 11/6$, versus 2 and $29/14$ for uniform and stationary readout. Optimality persists in an open parameter neighborhood. Appendix J develops this sufficient optimality certificate. Both risks in $\rho_\infty \leq C_H(r)\rho_H$ are evaluated at the chosen readout r.

7 TEACHING INTERVENTIONS AND SHARP CONSTRUCTIONS

A paired teaching intervention. Does adding the stationary graph remove error hidden by path-only teaching? We compare the two teaching sets at $K = 2$, with state and message width four and seeds $10, 11, 12$. Each pair starts from the same positive random weights. One fit uses the three coordinate paths; the other adds the stationary endpoint. Targets are 50 and other initial values are 500 ($\beta = 10$). Both fits use nonnegative-bounded TRF/exact least squares without regularization; we evaluate the terminal solver iterates. Table 1 reports exact global errors from the endpoint formulas. Path-only fits have source residuals at most 6.54×10^{-14}, yet large global errors remain. All six models retain 48 strictly positive parameters and satisfy $\kappa < 10$. The stability theorem therefore equates their global error with error on the complete teaching set. The missing stationary endpoint accounts for the path-only discrepancy (Appendix E).

Table 1: Effect of one additional stationary teaching graph on exact graph-uniform error, across all three paired fits.

Seed	Paths: global error (%)	Paths + stationary: global error (%)
10	26.7844	1.56×10^{-6}
11	19.9201	0.0937
12	47.9872	0.1627

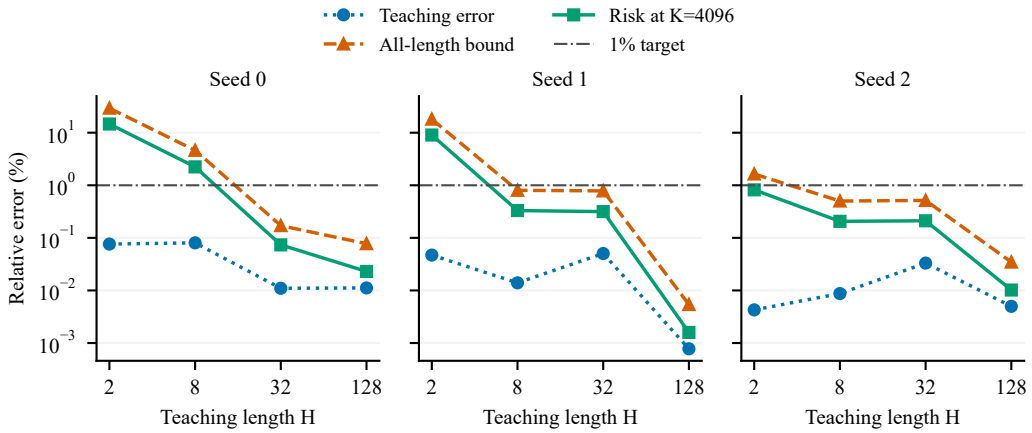

Figure 2: Twelve learned-cost fits at four teaching lengths, with paired initializations within each seed. Lines connect separately trained models; green curves show graph-uniform risk at $K = 4096$.

Table 2: Teaching-length selection at a fixed fitting tolerance of 0.1%. All errors are percentages; exact all-length risk ρ_∞ is evaluated after selection.

Seed	Source-only ρ_∞	Certificate H	Bound	Selected ρ_∞
0	14.6268	32	0.1708	0.0736
1	9.0355	8	0.8019	0.3306
2	0.8192	8	0.5017	0.2062

Choosing a teaching length for learned costs. Does the certificate change how long we teach a model? We fix three-channel connections $A = (63/64)I + \mathbf{1}\mathbf{1}^\top/192$, $B = I$, $r = (1, 0, 0)$, and learn only $c \in [0, 2]^3$. For each seed $0, 1, 2$, we fit all $H+2$ teaching endpoints at each $H \in \{2, 8, 32, 128\}$, restarting from the same random costs. Each fit stops at its first accepted iterate with maximum teaching error $\delta_H \le 0.1\%$. The source-only rule selects the shortest qualifying length in this grid. The certificate rule also requires

$$C_H^+ \delta_H \le 1\%, \qquad C_H^+ = \frac{6}{1 - (63/64)^{H-1}},$$

using the inexpensive coverage bound in Theorem 9.

All twelve fits meet the teaching tolerance, so the source-only rule always selects $H = 2$ (Figure 2). Two selections nevertheless have all-length errors of 14.63% and 9.04% (Table 2). The certificate selects longer teaching and all three errors fall below 1%. The third seed already meets the target at $H = 2$. Appendix H reports the fits and protocol. Appendix I applies the sharper constants from Theorem 19 to certify all twelve models post hoc.

Certified horizon constructions. Rational identities and tail bounds verify the coverage, mixed-error, and periodic constructions. The vertex algorithm encloses constants 127 and 128 within 5×10^{-7} and returns $[1, 1]$ for the covered three-cycle. Appendix G gives the numerical results and verification of the 194/195 cutoff.

8 Conclusion

Exact risk and coefficient balance certify nonnegative MinAgg on unseen graphs. Under normalized shared connections, channel coverage and mixed-error propagation quantify long-horizon error. Readout design strengthens these guarantees, and the experiments guide example choice and teaching length. A general near-optimal readout algorithm for arbitrary teaching lengths would extend the one-step guarantee and sufficient multi-step certificate.

Reproducibility statement

The appendices provide complete proofs and experimental methods. The supplement includes code, initializations and fitted weights for all six paired models and twelve learned-cost fits, optimization trajectories, graph predictions, selected horizons, and rational certificates. The included commands regenerate all tables and figures.

AI use statement

Generative AI tools assisted research ideation and literature retrieval; mathematical formulation, proof development and checking; synthetic graph construction, experiment design, implementation and analysis; and manuscript drafting and revision. The author takes responsibility for the final content, including AI-assisted text, claims, code, and artifacts.

Ethics statement

This work studies mathematical guarantees and synthetic graphs, without human participants or personal data.

References

Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2006.05205.

Dana Angluin. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2):87–106, 1987. ISSN 0890-5401. doi: 10.1016/0890-5401(87)90052-6. URL http://dx.doi.org/10.1016/0890-5401(87)90052-6.

V. Balakrishnan and S. Boyd. On computing the worst-case peak gain of linear systems. *Systems & Control Letters*, 19(4):265–269, 1992. ISSN 0167-6911. doi: 10.1016/0167-6911(92)90064-y. URL http://dx.doi.org/10.1016/0167-6911(92)90064-Y.

Borja Balle, Pascale Gourdeau, and Prakash Panangaden. Bisimulation metrics for weighted automata. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl (eds.), *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 103:1–103:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi: 10.4230/LIPIcs.ICALP.2017.103. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2017.103.

Timon Barlag, Vivian Holzapfel, Laura Strieker, Jonni Virtema, and Heribert Vollmer. Recurrent graph neural networks and arithmetic circuits. In *Proceedings of the 23rd International Conference on Principles of Knowledge Representation and Reasoning*, pp. 826–836, 2026. doi: 10.24963/kr.2026/78. URL https://proceedings.kr.org/2026/78/. Version consulted: arXiv:2603.05140v2.

Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks, 2018. URL https://arxiv.org/abs/1806.01261.

Richard Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958. ISSN 0033-569X. doi: 10.1090/qam/102435. URL http://dx.doi.org/10.1090/qam/102435.

Joost Berkhout and Bernd F. Heidergott. The jump start power method: A new approach for computing the ergodic projector of a finite Markov chain. *Journal of Scientific Computing*, 78(3):

1691–1723, 2019. doi: 10.1007/s10915-018-0828-1. URL https://doi.org/10.1007/s10915-018-0828-1.

Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lió, and Michael Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1026–1037. PMLR, 2021. URL https://proceedings.mlr.press/v139/bodnar21a.html.

Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting, 2021. URL https://arxiv.org/abs/2006.09252.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. ISBN 9780511804441. doi: 10.1017/cbo9780511804441. URL http://dx.doi.org/10.1017/CBO9780511804441.

Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004. ISSN 1095-7200. doi: 10.1137/s0036144503423264. URL http://dx.doi.org/10.1137/S0036144503423264.

Davide Buffelli, Pietro Liò, and Fabio Vandin. Sizeshiftreg: a regularization method for improving size-generalization in graph neural networks. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2207.07888.

Samantha Chen, Sunhyuk Lim, Facundo Memoli, Zhengchao Wan, and Yusu Wang. Weisfeiler-lehman meets gromov-Wasserstein. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 3371–3416. PMLR, 2022. URL https://proceedings.mlr.press/v162/chen22o.html.

Sanjoy Dasgupta, Daniel Hsu, Stefanos Poulis, and Xiaojin Zhu. Teaching a black-box learner. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1547–1555. PMLR, 2019. URL https://proceedings.mlr.press/v97/dasgupta19a.html.

Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, and Michael M. Bronstein. On over-squashing in message passing neural networks: The impact of width, depth, and topology. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 7865–7885. PMLR, 2023. URL https://proceedings.mlr.press/v202/di-giovanni23a.html.

Andrew J Dudzik and Petar Veličković. Graph neural networks are dynamic programmers. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 20635–20647. Curran Associates, Inc., 2022. doi: 10.52202/068431-1500. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8248b1ded388fcdbbd121bcdfea3068c-Paper-Conference.pdf.

Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3419–3430. PMLR, 2020. URL https://proceedings.mlr.press/v119/garg20c.html.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272. PMLR, 2017. URL https://proceedings.mlr.press/v70/gilmer17a.html.

S.A. Goldman and M.J. Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995. ISSN 0022-0000. doi: 10.1006/jcss.1995.1003. URL http://dx.doi.org/10.1006/jcss.1995.1003.

Jesse He, Helen Jenne, Max Vargas, Davis Brown, Gal Mishne, Yusu Wang, and Henry Kvinge. MINAR: Mechanistic interpretability for neural algorithmic reasoning, 2026. URL https://arxiv.org/abs/2602.21442v1. Version 1.

Borja Ibarz, Vitaly Kurin, George Papamakarios, Kyriacos Nikiforou, Mehdi Bennani, Róbert Csordás, Andrew Joseph Dudzik, Matko Bošnjak, Alex Vitvitskyi, Yulia Rubanova, Andreea Deac, Beatrice Bevilacqua, Yaroslav Ganin, Charles Blundell, and Petar Veličković. A generalist neural algorithmic learner. In Bastian Rieck and Razvan Pascanu (eds.), *Proceedings of the First Learning on Graphs Conference*, volume 198 of *Proceedings of Machine Learning Research*, pp. 2:1–2:23. PMLR, 2022. URL https://proceedings.mlr.press/v198/ibarz22a.html.

Haotian Ju, Dongyue Li, Aneesh Sharma, and Hongyang R. Zhang. Generalization in graph neural networks: Improved pac-bayesian bounds on graph diffusion. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 6314–6341. PMLR, 2023. URL https://proceedings.mlr.press/v206/ju23a.html.

Stasys Jukna and Georg Schnitger. On the optimality of Bellman–Ford–Moore shortest path algorithm. *Theoretical Computer Science*, 628:101–109, 2016. doi: 10.1016/j.tcs.2016.03.014. URL https://doi.org/10.1016/j.tcs.2016.03.014.

Lodewijk Kallenberg. Derman's book as inspiration: some results on LP for MDPs. *Annals of Operations Research*, 208(1):63–94, 2013. doi: 10.1007/s10479-011-1047-4. URL https://doi.org/10.1007/s10479-011-1047-4.

Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/ea9268cb43f55d1d12380fb6ea5bf572-Paper.pdf.

Stefan Kiefer. Notes on equivalence and minimization of weighted automata, 2020. URL https://arxiv.org/abs/2009.01217v1. Version 1; exposition of established results.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://arxiv.org/abs/1609.02907.

David Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2 edition, 2017. doi: 10.1090/mbk/107. URL https://doi.org/10.1090/mbk/107.

Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. In *International Conference on Learning Representations*, 2016. URL https://arxiv.org/abs/1511.05493.

Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda B. Smith, James M. Rehg, and Le Song. Iterative machine teaching. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2149–2158. PMLR, 2017. URL https://proceedings.mlr.press/v70/liu17b.html.

Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1907.03199.

Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bb04af0f7ecaee4aae62035497da1387-Paper.pdf.

Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. In *International Conference on Learning Representations*, 2019b. URL https://arxiv.org/abs/1812.09902.

Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the universality of invariant networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4363–4371. PMLR, 2019c. URL https://proceedings.mlr.press/v97/maron19a.html.

Sohir Maskey, Raffaele Paolino, Fabian Jogl, Gitta Kutyniok, and Johannes Lutzeyer. Graph representational learning: When does more expressivity hurt generalization? In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 105345–105392, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/ac29036313fcd6fb84fd48787fe103d0-Paper-Conference.pdf.

Siqi Miao, Mia Liu, and Pan Li. Interpretable and generalizable graph learning via stochastic attention mechanism. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 15524–15543. PMLR, 2022. URL https://proceedings.mlr.press/v162/miao22a.html.

Mehryar Mohri. Weighted automata algorithms. In *Handbook of Weighted Automata*, pp. 213–254. Springer Berlin Heidelberg, 2009. ISBN 9783642014925. doi: 10.1007/978-3-642-01492-5_6. URL http://dx.doi.org/10.1007/978-3-642-01492-5_6.

Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4602–4609, 2019. URL https://arxiv.org/abs/1810.02244.

Christopher Morris, Floris Geerts, Jan Tönshoff, and Martin Grohe. WL meet VC. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 25275–25302. PMLR, 2023. URL https://proceedings.mlr.press/v202/morris23a.html.

Mojmir Mutny, Tadeusz Janik, and Andreas Krause. Active exploration via experiment design in markov chains. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 7349–7374. PMLR, 2023. URL https://proceedings.mlr.press/v206/mutny23a.html.

Robert R. Nerem, Samantha Chen, Sanjoy Dasgupta, and Yusu Wang. Graph neural networks extrapolate out-of-distribution for shortest paths. In Steve Hanneke and Tor Lattimore (eds.), *Proceedings of Thirty Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 5273–5331. PMLR, 2026. URL https://proceedings.mlr.press/v336/nerem26a.html.

Masaki Ogura, Masako Kishida, and James Lam. Geometric programming for optimal positive linear systems. *IEEE Transactions on Automatic Control*, 65(11):4648–4663, 2020. ISSN 2334-3303. doi: 10.1109/tac.2019.2960697. URL http://dx.doi.org/10.1109/TAC.2019.2960697.

Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1905.10947.

Javier Peña, Juan C. Vera, and Luis F. Zuluaga. New characterizations of Hoffman constants for systems of linear constraints. *Mathematical Programming*, 187(1–2):79–109, 2021. doi: 10.1007/s10107-020-01473-6. URL https://doi.org/10.1007/s10107-020-01473-6.

Muhammad Fetrat Qharabagh, Artur Back de Luca, George Giapitzakis, and Kimon Fountoulakis. Learning to execute graph algorithms exactly with graph neural networks, 2026. URL https://arxiv.org/abs/2601.23207v2. Version 2.

F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009. ISSN 1941-0093. doi: 10.1109/tnn.2008.2005605. URL http://dx.doi.org/10.1109/TNN.2008.2005605.

Gernot Schullerus, Volker Krebs, Bart De Schutter, and Ton van den Boom. Input signal design for identification of max-plus-linear systems. *Automatica*, 42(6):937–943, 2006. doi: 10.1016/j.automatica.2006.01.025. URL https://doi.org/10.1016/j.automatica.2006.01.025.

Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2111.14522.

Petar Veličković, Adrià Puigdomènech Badia, David Budden, Razvan Pascanu, Andrea Banino, Misha Dashevskiy, Raia Hadsell, and Charles Blundell. The CLRS algorithmic reasoning benchmark. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 22084–22102. PMLR, 2022. URL https://proceedings.mlr.press/v162/velickovic22a.html.

Petar Veličković, Matko Bošnjak, Thomas Kipf, Alexander Lerchner, Raia Hadsell, Razvan Pascanu, and Charles Blundell. Reasoning-modulated representations. In Bastian Rieck and Razvan Pascanu (eds.), *Proceedings of the First Learning on Graphs Conference*, volume 198 of *Proceedings of Machine Learning Research*, pp. 50:1–50:17. PMLR, 2022. URL https://proceedings.mlr.press/v198/velickovic22a.html.

Petar Veličković and Charles Blundell. Neural algorithmic reasoning. *Patterns*, 2(7):100273, 2021. ISSN 2666-3899. doi: 10.1016/j.patter.2021.100273. URL http://dx.doi.org/10.1016/j.patter.2021.100273.

Petar Veličković, Rex Ying, Matilde Padovano, Raia Hadsell, and Charles Blundell. Neural execution of graph algorithms. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1910.10593.

Solveig Wittig, Antonis Vasileiou, Robert R. Nerem, Timo Stoll, Floris Geerts, Yusu Wang, and Christopher Morris. Which algorithms can graph neural networks learn?, 2026. URL https://arxiv.org/abs/2602.13106v2. Version 2.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/abs/1810.00826.

Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du, Ken ichi Kawarabayashi, and Stefanie Jegelka. What can neural networks reason about? In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1905.13211.

Keyulu Xu, Mozhi Zhang, Stefanie Jegelka, and Kenji Kawaguchi. Optimization of graph neural networks: Implicit acceleration by skip connections and more depth. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11592–11602. PMLR, 2021a. URL https://proceedings.mlr.press/v139/xu21k.html.

Keyulu Xu, Mozhi Zhang, Jingling Li, Simon S. Du, Ken ichi Kawarabayashi, and Stefanie Jegelka. How neural networks extrapolate: From feedforward to graph neural networks. In *International Conference on Learning Representations*, 2021b. URL https://arxiv.org/abs/2009.11848.

Gilad Yehudai, Ethan Fetaya, Eli Meirom, Gal Chechik, and Haggai Maron. From local structures to size generalization in graph neural networks. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11975–11986. PMLR, 2021. URL https://proceedings.mlr.press/v139/yehudai21a.html.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander Smola. Deep sets. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf.

Liwen Zhang, Gregory Naitzat, and Lek-Heng Lim. Tropical geometry of deep neural networks. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5824–5832. PMLR, 2018. URL https://proceedings.mlr.press/v80/zhang18i.html.

Chengyan Zhao, Xin Gong, Yoshio Ebihara, and Masaki Ogura. Impulse-to-peak optimization of positive linear systems via dc programming. *IFAC-PapersOnLine*, 56(2):5059–5064, 2023. ISSN 2405-8963. doi: 10.1016/j.ifacol.2023.10.1286. URL http://dx.doi.org/10.1016/j.ifacol.2023.10.1286.

Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An overview of machine teaching, 2018. URL https://arxiv.org/abs/1801.05927.

A SUPPORT CLEANUP AND THE UNSTABLE-TEACHING FAMILY

Lemma 14 (Exact support cleanup)**.** *The cleanup leaves F_θ unchanged on every admissible graph. A state is active according to* (4) *if and only if it is nonzero on some admissible input.*

Proof. Inductively, an inactive message reads only states that are identically zero. Its zero-weight self-loop gives a candidate value of zero, so its minimum is zero regardless of its edge coefficient. Removing that coordinate therefore changes no state. Nonnegativity gives the stated activity rule for B_t. Conversely, set all initial values to one and all edge weights to zero. Every node has state $p_0 = 1$, $p_t = B_t A_t p_{t-1}$, whose positive support is exactly s_t. Thus active coordinates can actually be nonzero. □

A.1 PROOF OF PROPOSITION 5

Proof. The extra message contributes $g_\varepsilon(G, v) = \min_{u \in N(v)}(\varepsilon h_u + w_{uv})$. At $\varepsilon = 0$, the second channel is inactive: its self-loop gives zero. For positive ε, the endpoint formulas give $L = 1 + \varepsilon$ and $U = 2$. More directly, take two nodes with initial values $(0, 1)$ and connecting edge weight ε. At the second node, $\Gamma^1 = \varepsilon$ and $F = 2\varepsilon$. On any fixed finite set, let B_0 bound all initial values and let $a > 0$ be the smallest positive BF target, if any. The self-loop gives $0 \le g_\varepsilon(G, v) \le \varepsilon h_v \le \varepsilon B_0$, so each positive-target relative error is at most $\varepsilon B_0/a$. Zero-target errors are exactly zero. If there are no positive targets, all errors are zero. □

For $K = 1$ on the exact teaching set T_β, the maximum endpoint relative error of this family is exactly $\min(\beta\varepsilon, 1)$: the input coordinate and stationary endpoints have error ε, while the edge coordinate has error $\min(\beta\varepsilon, 1)$. Therefore there is no function $\omega(\delta) \to 0$ that uniformly bounds ρ_* by ω(maximum teaching error) on this class, even at width two with all parameters bounded by one. Exact fitting, as required by Theorem 4, holds only at $\varepsilon = 0$ in this family.

The relative risk is therefore discontinuous even with bounded attributes and parameters. The target in this witness tends to zero, while the additional absolute error is at most εB_0 on a bounded input domain. The delayed-error construction in Appendix C.3 keeps the target equal to one.

B PROOFS FOR FIXED-LENGTH RISK AND TEACHING

B.1 PROOF OF THEOREM 1

Proof. Lower bound. Suppose $h_u^{(t-1)} \ge \ell_{t-1} x_u^{(t-1)}$. Writing $a = A_t \ell_{t-1}$, each message satisfies

$$A_t h_u^{(t-1)} + c_t w_{uv} \ge a x_u^{(t-1)} + c_t w_{uv} \ge \min(a, c_t)(x_u^{(t-1)} + w_{uv}).$$

Taking coordinatewise minima and multiplying by nonnegative B_t proves $h_v^{(t)} \ge \ell_t x_v^{(t)}$. The initial step is equality. On the two-node graph with initial values $(0, 1)$ and connecting edge weight one, the first node stays zero. At the second node, the self-loop and other neighbor give precisely $A_t \ell_{t-1}$ and c_t. Hence its state is ℓ_t, while its BF value remains one.

Upper bound. For any prescribed walk $v_0, \dots, v_K$, the minimum in layer t is at most the candidate from v_{t-1}. Expanding these inequalities using nonnegative matrices gives

$$F_\theta(G, v_K) \le \gamma_0 h_{v_0} + \sum_{t=1}^{K} \gamma_t w_{v_{t-1} v_t}.$$

Select a BF-minimizing walk. The right side is at most $U\Gamma^K(G, v_K)$. This also proves the zero-target assertion.

Upper attainment. Fix $j \in \{0, \dots, K\}$ and put $z = e_j$. Take the path $v_0, \dots, v_K$, with initial value z_0 at v_0, initial value β at every other node, and edge $v_{t-1} v_t$ of weight z_t. Define

$$p_0 = 1, \quad p_t = C_t p_{t-1}, \qquad H_0 = z_0, \quad Q_t = A_t H_{t-1} + c_t z_t, \quad H_t = B_t Q_t. \tag{31}$$

Choose any finite

$$\beta > \max\left\{1, \max_{t,i:\,(A_t p_{t-1})_i > 0} \frac{(Q_t)_i}{(A_t p_{t-1})_i}\right\}. \tag{32}$$

An empty inner maximum contributes nothing. After cleanup, an inactive message has $Q_{t,i} = 0$. First consider a node farther than t edges from v_0. Its state at time t is βp_t: every preceding state is βp_{t-1}, and the zero-weight self-loop supplies the smallest edge contribution. Now consider the frontier node v_t. Its predecessor supplies Q_t, whereas every other active candidate is at least $\beta A_t p_{t-1}$. The choice in (32) makes these alternatives strictly larger than Q_t. Induction gives $h_{v_t}^{(t)} = H_t$, and hence $H_K = \gamma_j$. The BF target at v_K is one. The full path from v_0 costs $\sum_{i=0}^{K} z_i = 1$; any other starting node costs at least $\beta > 1$. Reaching v_K from v_0 in K steps requires the full path. Thus each γ_j is attained. The lower witness completes (7). □

B.2 Proof of Theorem 2

Proof. Put $d_t = p_t - \ell_t \geq 0$. From (5),

$$d_t = C_t d_{t-1} + B_t \left(A_t \ell_{t-1} - c_t\right)_+, \tag{33}$$

$$p_K - L = \sum_{t=1}^{K} r_t \left(A_t \ell_{t-1} - c_t\right)_+. \tag{34}$$

Let $P_t = r_t \left(A_t p_{t-1} - c_t\right)_+$. The positive-part inequality

$$\left(A_t p_{t-1} - c_t\right)_+ \leq A_t d_{t-1} + \left(A_t \ell_{t-1} - c_t\right)_+$$

and the remaining nonnegative terms of (33) give $P_t \leq p_K - L$. Also $r_t A_t p_{t-1} = p_K$ and $r_t c_t = \gamma_t$. Consequently $m_t = 2P_t + \gamma_t - p_K \leq p_K + \gamma_t - 2L$. Since $p_K = \gamma_0 \leq U$ and $|U-1|, |L-1| \leq \rho_*$, the first chain of bounds follows.

For the converse, $\ell_{t-1} \leq p_{t-1}$ in (34) gives $p_K - L \leq \sum_t P_t \leq \sum_t m_t$. Moreover $|\gamma_t - p_K| \leq m_t$. Thus both L and U lie in $[p_K - \sum_t m_t, p_K + \sum_t m_t]$, proving the final inequality. □

B.3 Proof of Corollary 3

Proof. First establish the coefficient condition. The endpoint theorem and (8) make zero risk equivalent to $p_K = 1$ and $m_t = 0$ for all t. Since r_t is nonnegative, $m_t = 0$ is exactly the coordinate condition in (9). Next prove the hidden-state claim by induction. A relevant message can read a preceding state with a positive A_t coefficient only if that state is also relevant. The induction hypothesis therefore gives $A_t h_u^{(t-1)} = (A_t p_{t-1}) x_u^{(t-1)}$ in that coordinate. Alignment assigns the same scale to the edge coefficient. Taking the minimum gives $(A_t p_{t-1})_i x_v^{(t)}$, and multiplying by B_t proves the claim for each relevant state. Coordinates unused by the output remain unconstrained. □

Lemma 15 (A single truncation on each coordinate path)**.** *For every $\beta > 1$, the designated path outputs satisfy $\tau_0 = p_K$ and $\tau_j = r_j \min(c_j, \beta\alpha_j)$ for $1 \leq j \leq K$.*

B.4 Proof of Lemma 15

Proof. As in the upper-attainment proof, nodes beyond the propagation frontier have state βp_t. Fix $j \geq 1$. Before time j, the frontier state is zero. At time j, the unit edge and the zero-weight self-loop give the two candidates, so $H_j = B_j \min(c_j, \beta\alpha_j) \leq \beta p_j$. This is the only truncation. Every later frontier edge has weight zero. If $H_{t-1} \leq \beta p_{t-1}$, then $A_t H_{t-1} \leq \beta \alpha_t$, so neither the high-valued self-loop nor the forward neighbor truncates it again. Hence $H_t = C_t H_{t-1} \leq \beta p_t$ for $t > j$. Expanding to the output proves (10). For $j = 0$, all edges have weight zero. The frontier propagates p_t from initial value one, without truncation because $\beta > 1$. □

B.5 Proof of Theorem 4

Proof. The reverse implication is immediate. For the forward implication, Lemma 15 and the stationary witness give $p_K = L = 1$. Thus $P_t \leq p_K - L = 0$ in the proof of Theorem 2, so

$$(c_t)_i \geq (A_t p_{t-1})_i \quad \text{whenever } r_{t,i} > 0. \tag{35}$$

For every $j \geq 1$, exact path fitting and $r_j \alpha_j = p_K$ give

$$0 = \tau_j - p_K = r_j\big(\min(c_j, \beta\alpha_j) - \alpha_j\big).$$

All contributing terms are nonnegative by (35). Every relevant message has $\alpha_{j,i} > 0$ after cleanup, so $\beta > 1$ makes any strict inequality $c_{j,i} > \alpha_{j,i}$ contribute positively. All relevant coefficients must therefore be equal. Corollary 3 proves exact BF execution. □

B.6 Proof of Theorem 6

Proof. The lower bound follows because all teaching graphs belong to the domain. For $0 \leq \lambda \leq \kappa$, direct separation of the cases $\lambda < 1$, $1 \leq \lambda \leq \beta$, and $\lambda > \beta$ gives

$$\lambda - 1 \leq \nu_\beta(\min(\lambda, \beta) - 1) + (\nu_\beta - 1)(1 - \lambda)_+.$$

For each relevant coordinate in layer j, put $\lambda = c_{j,i}/\alpha_{j,i}$, multiply by $r_{j,i}\alpha_{j,i}$, and sum. Equation (10) and $P_j = r_j(\alpha_j - c_j)_+ \leq p_K - L$ imply

$$\begin{aligned} \gamma_j - p_K &\leq \nu_\beta(\tau_j - p_K) + (\nu_\beta - 1)P_j, \\ \gamma_j &\leq \nu_\beta \tau_j - (\nu_\beta - 1)L \leq 1 + (2\nu_\beta - 1)\delta_\beta. \end{aligned} \tag{36}$$

Also $p_K \leq 1 + \delta_\beta$ and $L \geq 1 - \delta_\beta$. The exact endpoint theorem gives the upper bound in (13). When $\kappa \leq \beta$, $\nu_\beta = 1$ proves equality; equivalently, Lemma 15 then gives $\tau_j = \gamma_j$ for every j.

For sharpness, choose any $0 < \delta < (\beta - 1)/(\beta + 1)$ and set

$$q = \frac{2\delta}{\beta - 1}, \qquad A = (1 - q, q)^\top, \quad c = (1 - \delta - q, \kappa q)^\top, \quad B = (1, 1). \tag{37}$$

All parameters are positive and the actual largest ratio is κ. Direct substitution gives $p_1 = 1$, $L = 1 - \delta$, $\tau_1 = 1 + \delta$, and $U = 1 + (2\nu_\beta - 1)\delta$. Hence $\delta_\beta = \delta$ and (13) is attained. □

B.7 Proof of Corollary 7

Proof. For (14), induction gives $p_t = \mathbf{1}$ and $\alpha_t = A_t p_{t-1} = \mathbf{1}$. Thus $\kappa \leq \beta$, and Theorem 6 gives $\rho_* = \delta_\beta$ for every parameter choice.

For the converse, first clean the support and put $q_t = C_K \cdots C_{t+1}$, with $q_K = 1$. Keep states with $p_{t,j} > 0$ and $q_{t,j} > 0$, and messages with $\alpha_{t,i} > 0$ and $r_{t,i} > 0$. Since $p_K = 1$, some coordinates survive at every stage. We must check that this removal loses no positive contribution. If a kept message reads an active state through $A_{t,ij} > 0$, then $(q_{t-1})_j = (r_t A_t)_j > 0$, so that state is kept. Similarly, if a kept state receives an active message through $B_{t,ji} > 0$, then $(r_t)_i = (q_t B_t)_i > 0$, so that message is kept. Inactive coordinates contribute zero. Recomputing on the retained coordinates therefore leaves p_t and α_t unchanged, with all retained entries positive.

On these reduced coordinates, set $D_t = \mathrm{diag}(p_t)$, $S_t = \mathrm{diag}(\alpha_t)$, and

$$A'_t = S_t^{-1} A_t D_{t-1}, \qquad c'_t = S_t^{-1} c_t, \qquad B'_t = D_t^{-1} B_t S_t. \tag{38}$$

Because $A_t p_{t-1} = \alpha_t$ and $B_t \alpha_t = p_t$, both transformed matrices have row sums one. Relevance and $\kappa \leq \beta$ give $0 \leq c'_t \leq \beta\mathbf{1}$. Positive diagonal scaling commutes with coordinatewise minimum, so $h'^{(t)}_v = D_t^{-1} h^{(t)}_v$ inductively. The scalar endpoints have $D_0 = D_K = 1$, proving equality of output functions. Removing coordinates does not increase widths. The assumption $p_K = 1$ ensures that the final scalar scaling is the identity. □

C A finite criterion for exact shared execution

Here we analyze shared execution with general nonnegative parameters, without the normalization imposed in Section 5. Fix nonnegative $e \in \mathbb{R}^d$, $A \in \mathbb{R}^{q \times d}$, $c \in \mathbb{R}^q$, $B \in \mathbb{R}^{d \times q}$, and a row $r \in \mathbb{R}^{1 \times d}$, and use

$$h^0_v = e h_v, \qquad h^t_v = B \min_{u \in N(v)} (A h^{t-1}_u + c w_{uv}), \qquad F_K(v) = r h^K_v. \tag{39}$$

The graph is fixed throughout execution, and parameters are shared across steps. The admissible graph domain is that of Section 3. Write $C = BA$, $p_n = C^n e$, and $\alpha_i(n) = (A p_n)_i$. Let $\chi_i(n)$ equal one when $\alpha_i(n) > 0$, and zero otherwise.

Proposition 16 (Support-sensitive recurrent endpoints)**.** *For every $K \geq 1$, the exact graph-uniform ratio endpoints are*

$$\begin{aligned} \ell_0 &= e, & \ell_t &= B\min(A\ell_{t-1}, c), & L_K &= r\ell_K, \\ \gamma_{K,0} &= rC^K e, & \gamma_{K,t} &= rC^{K-t}B(c \odot \chi(t-1)) \quad (1 \leq t \leq K), & & \\ U_K &= \max_{0\leq t\leq K} \gamma_{K,t}. & & & & \end{aligned} \tag{40}$$

Consequently $\rho_K = \max(|L_K - 1|, |U_K - 1|)$.

Reduction to fixed length. Unroll K shared updates, absorb e and r at the endpoints, and apply Theorem 1. Support cleanup removes inactive edge injections but preserves every positive suffix path from an active channel. Appendix C.1 proves this support claim.

The mask cannot be omitted. Let $A = I_2$, let B swap the coordinates, and take $e = (1,0)^\top$, $c = (1,1)^\top$, $r = (1,1)$. Only one hidden coordinate is active at a time: $h_v^K = B^K e\, \Gamma^K(v)$, so $F_K = \Gamma^K$ at every length. Replacing the masked coefficient by $rC^{K-t}Bc$ would incorrectly give an exact upper endpoint of two rather than one.

Theorem 17 (A finite criterion for all-length exact BF)**.** *Call message channel i observable if $(rC^sB)_i > 0$ for some $0 \leq s \leq d-1$, and put $D = d^2 + d$. Model* (39) *agrees with BF on every admissible graph, at every node and every $K \geq 1$, if and only if*

$$rC^K e = 1 \qquad (1 \leq K \leq d+1), \tag{41}$$

$$\alpha_i(n)(\alpha_i(n) - c_i) = 0 \qquad (0 \leq n < D,\ i \text{ observable}). \tag{42}$$

For rational parameters this is decidable in polynomial bit complexity. In particular, exact agreement on all graphs for $K = 1, \ldots, d^2 + 2d - 1$ implies exact agreement at every length.

Proof outline. The proof separates graph execution from a finite sequence test. Nonnegativity bounds the shortest observable suffix by $d - 1$. Fixed-length alignment then requires scalar normalization and $\alpha_i(n) \in \{0, c_i\}$ for each observable channel. The normalization sequence has a linear representation of dimension $d + 1$. The quadratic alignment sequence has one of dimension $d^2 + d$. Standard Cayley–Hamilton and weighted-automaton constructions give finite zero tests for these sequences (Kiefer, 2020, Propositions 1.1 and 2.2). Accounting for inactive support makes these sequence conditions equivalent to exact graph execution. Appendix C.2 gives the proof and bit complexity.

Using Theorem 4 separately at the specified lengths also gives an $O(d^4)$ finite exact teaching bound. This horizon concerns all-graph equality at each length. Section 5 characterizes approximate teaching for normalized connections. The construction below shows why its bound must depend on the connections.

C.1 Proof of Proposition 16

Proof. Zero self-loops imply $0 \leq h_v^n \leq (\max_u h_u)p_n$ by induction. Thus a channel with $\alpha_i(n) = 0$ is identically zero at that time. Apply Theorem 1 after incorporating e into the first A and r into the last B. The vector $B(c \odot \chi(t-1))$ has support contained in that of p_t. Every positive path starting in this support remains in the support of p_j at each later time j. Cleanup deletes none of these paths; its suffix action is therefore the original rC^{K-t}. The input coefficient is unchanged for the same reason. Since $\ell_n \leq p_n$, inactive channels contribute zero to the lower recurrence without explicit masking. Endpoint attainability follows from the same stationary and coordinate-path witnesses. □

C.2 Proof of Theorem 17

Proof. Nonnegativity identifies a positive coefficient with the existence of a positive path through the d state coordinates. Removing repeated vertices shows that a channel can reach the readout at some length if and only if it can do so with a suffix of length at most $d - 1$. At fixed K, Corollary 3 requires $rC^K e = 1$ and $c_i = \alpha_i(t-1)$ whenever both $\alpha_i(t-1) > 0$ and $(rC^{K-t}B)_i > 0$. The support argument in Proposition 16 justifies using this uncleaned suffix for an active channel. Taking all K gives precisely normalization at all positive lengths and $\alpha_i(n) \in \{0, c_i\}$ at every $n \geq 0$ for each observable i.

The scalar sequence $rC^K e - 1$, starting at $K = 1$, has a linear representation with matrix $\mathrm{diag}(C, 1)$ of dimension $d + 1$. For an observable channel,

$$\alpha_i(n)^2 - c_i\alpha_i(n) = (A_i \otimes A_i)(C \otimes C)^n(e \otimes e) - c_i A_i C^n e \tag{43}$$

has a representation with matrix $\mathrm{diag}(C \otimes C, C)$ of dimension D. Cayley–Hamilton implies that a represented scalar sequence of dimension N vanishes identically when its first N terms vanish. Thus (41)–(42) are sufficient as well as necessary. This finite zero test and the direct-sum/tensor-product construction follow standard weighted-automaton arguments (Kiefer, 2020, Propositions 1.1 and 2.2).

Conditions (41)–(42) can be evaluated by iterating p_n, without forming the tensor product. There are polynomially many rational matrix–vector operations and their intermediate bit lengths are polynomially bounded. Finally, for each $n < D$ and observable i, choose a positive suffix $s \le d - 1$. Exactness at $K = n + 1 + s \le d^2 + 2d - 1$ forces the required alignment whenever $\alpha_i(n) > 0$. The normalization conditions are also within this length range, proving the finite-horizon assertion. □

C.3 A DELAYED-ERROR CONSTRUCTION WITH UNIT TARGET

Fix $0 < p, \delta < 1$ and take $A = I_2$, $e = (1, 1)^\top$, $r = (1, 0)$,

$$B = \begin{pmatrix} 1-p & p \\ p & 1-p \end{pmatrix}, \qquad c = (1, 1-\delta)^\top.$$

We prove

$$\rho_K = \delta\{1 - (1-p)^K\}, \quad \max_{K \le H} \rho_K \le \delta H p, \quad \rho_\infty = \delta. \tag{44}$$

Unlike $C_H^\star$ at fixed connections, this counterexample lets the transition probability tend to zero. Set $b = 1 - \delta$. The lower recurrence is

$$\ell_0 = (1, 1)^\top, \qquad \ell_K = B \min\{\ell_{K-1}, (1, b)^\top\}.$$

We claim that, for $K \ge 1$,

$$(\ell_K)_1 = b + \delta(1-p)^K, \qquad (\ell_K)_2 = b + p\delta(1-p)^{K-1}.$$

At $K = 1$ this follows by multiplying $B(1, b)^\top$. Both coordinates are at least b, and the first is at most one, so the next minimum retains $(\ell_K)_1$ and caps the second coordinate at b. Multiplication by B gives the claimed expressions for $K + 1$. This proves the induction and yields $L_K = b + \delta(1-p)^K$.

All states are active and $B\mathbf{1} = \mathbf{1}$. The input coefficient is $rB^K e = 1$. Since $c \le \mathbf{1}$, every edge coefficient is at most one, hence $U_K = 1$. Proposition 16 now gives

$$\rho_K = 1 - L_K = \delta\{1 - (1-p)^K\}.$$

Bernoulli's inequality gives $\rho_K \le \delta K p$. For any fixed H, the maximum through length H therefore tends to zero as $p \downarrow 0$, while $\sup_{K \ge 1} \rho_K = \delta$ for every $p > 0$. The universal lower witness attains each L_K and has BF target one at every length, proving (44) on the same two-node graph.

D WHY NONNEGATIVITY MATTERS

Consider one message layer with $A = (1, 1)^\top$, $c = (1, 2)^\top$, and signed ReLU readout $F = \mathrm{ReLU}(2z_1 - z_2)$. On the stationary graph with initial values $(0, 1)$ and edge weight one, the second node has $z_1 = z_2 = 1$, so $F = 1$ agrees with BF. Changing the initial values to $(0, 3)$ gives $z_1 = 1$, $z_2 = 2$, and $F = 0$, while the BF value is still one. Thus the stationary graph no longer characterizes the lower ratio once signed cancellation is allowed.

E PAIRED TEACHING EXPERIMENT: METHODS AND RESULTS

The theory identifies a specific missing example: coordinate paths probe the propagation coefficients, while the stationary two-node graph exposes the lower endpoint. We illustrate this distinction by changing only the teaching set and keeping the architecture, initialization, and optimization procedure fixed within each pair.

Table 3: Adding the stationary graph reduces exact graph-uniform error in all three pairs. Source residual is the maximum relative error on the fitted endpoints; global error is computed from the exact ratio endpoints.

Seed	Teaching set	Source residual	Global error (%)	κ
10	Paths	6.54×10^{-14}	26.7844	1.5783
10	Paths + stationary	1.56×10^{-8}	1.56×10^{-6}	1.0000
11	Paths	1.11×10^{-16}	19.9201	1.6323
11	Paths + stationary	9.37×10^{-4}	0.0937	1.0015
12	Paths	1.11×10^{-16}	47.9872	2.9353
12	Paths + stationary	1.63×10^{-3}	0.1627	1.0125

We use $K = 2$, state and message width four, and seeds $10, 11, 12$. Each positive random initialization is fitted on either the three coordinate paths or those paths plus the stationary graph. The designated targets are 50 and the other path initial values are 500, corresponding to unit-target contrast $\beta = 10$. Both conditions use nonnegative-bounded SciPy TRF/exact least squares, an autodifferentiated Jacobian, no regularizer, at most 2000 function evaluations, and 10^{-12} values for `ftol`, `xtol`, and `gtol`, with `x_scale=1`. We evaluate each terminal solver iterate using Theorem 1. Computation uses float64 on CPU with one PyTorch thread.

Path-only fitting reaches near-machine-precision source residuals but leaves global errors between 19.92% and 47.99%. Adding one graph reduces the global error below 0.163% in every pair (Table 3). All learned models have $\kappa < \beta$, so Theorem 6 identifies their global error with maximum error on the complete teaching set. The omitted stationary endpoint therefore accounts for the discrepancy in the path-only fits.

The implementation is `experiments/minagg_fixed_teaching.py`. Solver status and exact duration are missing for one paths-only run; its parameters and numerical outputs are available.

F Proofs for finite-to-infinite teaching

We use the normalized shared class (16) on the same graph domain. Positive homogeneity lets us work with $\xi = c - \mathbf{1} \in \mathbb{R}^q$ in the proofs. Every direction can be scaled back into $[-1, 1]^q$ without changing its amplification ratio.

F.1 Proof of Proposition 8

Proof. The unrolled lower recurrence is $\ell_0 = \mathbf{1}$, $\ell_{k+1} = B \min(A\ell_k, c)$. Write $d_k = \mathbf{1} - \ell_k$. Row stochasticity gives $d_{k+1} = B \max(Ad_k, -\xi)$. Induction yields $d_k = Bz_k$, where $z_0 = 0$ and $z_{k+1} = \max(-\xi, Sz_k)$. Therefore $1 - L_H = \mu z_H$. All input coordinates remain active because their propagated scale is one. The inactive-support mask is therefore identically one. The input coefficient is one, and the edge coefficients are

$$r(BA)^j Bc = \mu S^j(\mathbf{1} + \xi) = 1 + \mu S^j \xi, \qquad 0 \le j < H.$$

The endpoint theorem now gives (18).

Since $S \ge 0$, z_k is coordinatewise nonnegative and nondecreasing; it is bounded above by $\|\xi\|_\infty \mathbf{1}$. Its limit exists, and taking suprema in the finite risk formula gives (19). The increasing stopping value and increasing prefix maximum also give $\rho_H = \max_{K \le H} \rho_K$. Finite stopping policies express μz_H as a maximum of finitely many linear functions of ξ; the upper term has the same property. Thus p_H is convex and positively homogeneous, as is $p_\infty = \sup_H p_H$. Both lie between zero and $\|\xi\|_\infty$. They can be asymmetric and hence are generally not norms.

At a visited channel i, the Bellman recurrence permits stopping with reward $-\xi_i$ or continuing. Terminal reward is zero. Negative rewards can be avoided, giving the same value as reward $(-\xi_i)_+$. The infinite value is the increasing limit of finite-horizon values; a path that never stops earns zero. In contrast, the upper term evaluates a signed expectation at a selected deterministic time. □

F.2 Proof of Theorem 9

Proof. If $i \in \mathcal{O} \setminus \mathcal{O}_H$, choose $\xi = -e_i$. No stopping decision before time H can visit i, and every upper expectation before that time is zero. Hence $p_H(\xi) = 0$. At some later finite time j, $(\mu S^j)_i > 0$. Stopping at j only in state i earns that positive probability. Thus $p_\infty(\xi) > 0$ and no finite amplification constant exists.

Conversely, let $\delta = p_H(\xi)$ and assume coverage. For each $j < H$, the strategy that stops at time j on negative coordinates and otherwise earns zero proves $\mu S^j(-\xi)_+ \le \mu z_H \le \delta$. Also $\mu S^j \xi \le \delta$, so $\mu S^j(\xi)_+ \le 2\delta$. For each observable i, take a time attaining $a_i(H) > 0$ to obtain

$$(-\xi_i)_+ \le \delta / a_i(H), \qquad (\xi_i)_+ \le 2\delta / a_i(H).$$

The observable set is closed under positive transitions; coordinates outside it cannot affect either risk. Therefore $p_\infty(\xi) \le \max_{i \in \mathcal{O}} |\xi_i|$ proves (21), including $\delta = 0$.

A positive entry $(\mu S^j)_i$ is equivalent to a positive-support walk of length j from $\operatorname{supp} \mu$ to i. The first coverage horizon is consequently the longest shortest distance plus one. A shortest walk has no repeated message vertex, giving $H_{\mathrm{cov}} \le |\mathcal{O}|$. For the state-width bound, use $\mu S^j = r(BA)^j B$. If a message is ever observed, a state with a positive coefficient into that message is reachable from $\operatorname{supp} r$. A shortest state walk has at most $d - 1$ edges, so that message is observed at some $j \le d - 1$. This gives $H_{\mathrm{cov}} \le d$. □

F.3 Proof of Theorem 10

Proof. Work only on $\mathcal{O}$, which is closed under S; the restricted matrix and initial row remain stochastic and probabilistic.

The teaching polytope. The constraint $\mu z_H(\xi) \le 1$ is equivalent to the existence of $z_1, \dots, z_H$ satisfying

$$z_0 = 0, \quad z_t \ge -\xi, \quad z_t \ge S z_{t-1} \ (1 \le t \le H), \qquad \mu z_H \le 1.$$

The exact Bellman iterates satisfy these constraints, proving necessity. Conversely, induction shows that any feasible auxiliary sequence dominates the Bellman iterates, proving sufficiency. Add the inequalities $\mu S^j \xi \le 1$ for $0 \le j < H$ to include the upper-risk term. These finitely many linear constraints describe $\mathcal{P}_H$ using auxiliary variables. Eliminating those variables gives a polyhedron. The coordinate bounds from the coverage proof make it bounded. It also contains the unit cube, since $p_H(\xi) \le \|\xi\|_\infty$. Thus it is a compact, full-dimensional polytope, and it is rational when the inputs are rational.

Reduction to vertices. Coverage implies $p_H(\xi) = 0$ only at $\xi = 0$ on the retained coordinates. Positive homogeneity equates (17) with $\max_{\xi \in \mathcal{P}_H} p_\infty(\xi)$. Every point is a convex combination of vertices, and convexity bounds its value by the largest vertex value. This proves (22). The maximum is at least one, since $\xi = \mathbf{1}$ has both risks equal to one. For a vertex outside the original cost box, the direction $\xi / \max(1, \|\xi\|_\infty)$ lies in $[-1, 1]^q$ and has the same ratio.

Exact stopping term. Every feasible z in (23) dominates $z_0 = 0$ and, by induction, every Bellman iterate. It therefore dominates z_∞. Continuity of the recurrence implies $z_\infty = \max(-\xi, S z_\infty)$, making it feasible itself. The constant vector $\|\xi\|_\infty \mathbf{1}$ is also feasible. This proves the LP identity without an absorption assumption. For rational data, rational linear programming gives its exact value. Stopping LPs are classical; see Kallenberg (2013, Section 5.1) for related formulations under different stopping assumptions.

Certified signed-orbit tails. Let m be the least common multiple of the periods of all closed irreducible classes of S, and put $Q = S^m$. The recurrent classes of Q are aperiodic; its transient part decays. Finite Markov-chain decomposition thus gives

$$Q^n \longrightarrow L, \qquad L^2 = L, \qquad QL = LQ = L.$$

The decomposition is standard; periodic multichains are treated, for example, in the appendix of Berkhout & Heidergott (2019). Class stationary distributions and transient absorption probabilities solve finite rational linear systems, so L is exactly computable. Enumerate powers until an $n_0 \ge 1$ satisfies $\lambda = \|Q^{n_0} - L\|_\infty < 1$. Convergence ensures termination. For $n = k n_0 + t \ge n_0$, with $k \ge 1$ and $0 \le t < n_0$, the projection identities give

$$Q^n - L = Q^t (Q^{n_0} - L)^k, \qquad \|Q^n - L\|_\infty \le \lambda^k.$$

Here $\|Q^t\|_\infty = 1$. Consequently, for each phase $0 \le a < m$,

$$|\mu S^a Q^n \xi - \mu S^a L\xi| \le \lambda^{\lfloor n/n_0 \rfloor} \|\xi\|_\infty \qquad (n \ge n_0). \tag{45}$$

If $\lambda = 0$, every phase tail is constant. Otherwise choose $N \ge n_0$ so that the tail bound at N is at most the requested tolerance. Compute the prefix $0 \le j < mN$ and all phase limits $\mu S^a L\xi$. Their largest value is a lower bound on the orbit supremum: even an unattained limit lies in the closure of the sequence. For an upper bound, add the tolerance to the largest phase limit and take the maximum with the prefix values. The two bounds differ by at most the tolerance. We compute the initial block directly and use (45) only when $n \ge n_0$.

Finite completion. Enumerate the rational polytope vertices using elimination and full-rank active constraints. For each vertex, combine the exact stopping value with its signed-orbit interval; then take the largest lower and upper endpoints over all vertices. Requesting tolerance η at each vertex yields width at most η for the resulting interval for $C_H^\star$. Each operation terminates on rational input, although vertex enumeration, period least common multiples, and slow tail convergence can all be costly. □

F.4 One-step teaching: an explicit mixed-error formula

Proposition 18 (Full-support one-step constant)**.** *Suppose $\mu_i > 0$ for every observable channel. Let h_i^∞ be the probability, starting from μ, of ever visiting i. Then*

$$C_1^\star = \max\left\{ \max_i \frac{h_i^\infty}{\mu_i},\ \sup_{j \ge 0} \left[2 \max_i \frac{(\mu S^j)_i}{\mu_i} - \min_i \frac{(\mu S^j)_i}{\mu_i} \right] \right\}. \tag{46}$$

Proof. At one step, $p_1(\xi) = \max\{\mu(-\xi)_+, \mu\xi\}$. A stopping policy can stop at i only on a path that visits i, so

$$\mu z_\infty(\xi) \le \sum_i h_i^\infty (-\xi_i)_+ \le \left(\max_i h_i^\infty / \mu_i \right) \mu(-\xi)_+.$$

Concentrating a negative perturbation at a largest-ratio coordinate and stopping at its first visit attains this factor.

For the upper term fix a probability row ν, and let $M = \max_i \nu_i/\mu_i$ and $m_0 = \min_i \nu_i/\mu_i$. On $p_1(\xi) \le 1$, write $N = \mu(-\xi)_+ \le 1$. Then $\mu(\xi)_+ \le 1+N$ and $\nu\xi \le M(1+N) - m_0 N \le 2M - m_0$. If $M > m_0$, choose distinct maximizing and minimizing coordinates i, k and set $\xi_i = 2/\mu_i$, $\xi_k = -1/\mu_k$, and other coordinates zero. This has $p_1 = 1$ and attains the bound. If $M = m_0$, the probability sums imply $\nu = \mu$ and $\xi = \mathbf{1}$ attains one. Putting $\nu = \mu S^j$ and taking the supremum proves the formula. Every direction can be scaled into the cost box. The orbit supremum can be enclosed by the periodic-tail procedure above. □

F.5 Proof of Theorem 11 and its cutoff

Proof. Write $\xi = (x, y)$. The signed expectation starts at $wx + sy$ and moves monotonically toward y. Analyze the four sign patterns.

If $x, y \le 0$, put $a_1 = -x$, $a_2 = -y$. For $a_1 \ge a_2$, stopping immediately is optimal at both horizons, giving $wa_1 + sa_2$. For $a_1 < a_2$, wait for state two or the last decision, whichever occurs first. Finite and infinite values are $(1-h)a_1 + ha_2$ and a_2; the ratio is at most $1/h$, attained at $a_1 = 0$. Thus the purely negative optimum is $1/h$.

If $x, y \ge 0$, the stopping reward is zero. A decreasing expectation achieves its maximum immediately. An increasing expectation gives finite maximum $(1-h)x + hy$ and infinite maximum y, hence ratio at most $1/h$.

If $x \ge 0, y = -a \le 0$, the finite and infinite stopping values are ha and a, respectively. The signed expectation decreases, so its largest positive part is $(wx - sa)_+$ at both horizons. The ratio of the maximum of these terms is at most $1/h$.

Finally, let $x = -a \le 0, y = b \ge 0$. Stopping immediately at state one is optimal for the negative reward, giving wa at both horizons; the upper expectation increases. Thus

$$p_H = \max\{wa, hb - (1-h)a\}, \qquad p_\infty = \max\{wa, b\}.$$

If $wa \geq b$, the ratio is one. Otherwise, put $z = a/b$. Maximizing the ratio is equivalent to minimizing $\max\{wz, h - (1-h)z\}$, whose increasing and decreasing lines intersect at $z^\star = h/(w+1-h)$. Their value is $wh/(w+1-h)$. This proves the full constant and the stated scaled witness. It is at least $1/h$ and therefore dominates the other sign cases.

If $p = 0$, the state never changes and both risks are already fully observed at one step. If $s = 1$, only the absorbing state is observed. Both cases give one. For $p > 0$, $h = 0$ is possible only when $s = 0, H = 1$: $\xi = (0, 1)$ has teaching risk zero and limiting risk one. At $p = 1, H \geq 2$, the formula yields one.

For $s = 0, 0 < p < 1$, solving $C_H^\star \leq 1 + \epsilon$ gives $(1-p)^{H-1} \leq \epsilon/(2+\epsilon)$. Taking logarithms, noting that $\log(1-p) < 0$, and rounding the lower bound on $H - 1$ upward proves (26). In the negative-only case the corresponding threshold is $\epsilon/(1+\epsilon)$. Both give $\Theta(p^{-1}\log(1/\epsilon))$ as p, ϵ become small. □

F.6 Same transition matrix and the periodic boundary

For $0 < p < 1$, let $u = 1-p$. With $r = (1, 0)$, $(A, B) = (P, I)$ gives $\mu = (1, 0)$ and $(A, B) = (I, P)$ gives $\mu = (1-p, p)$, while both products AB and BA are P. Substitution in Theorem 11 gives

$$C_H^\star(P, I) = \frac{1 + u^{H-1}}{1 - u^{H-1}}, \qquad C_H^\star(I, P) = \frac{1 + u^{H-1}}{1 - u^H}.$$

The first expression is infinite at $H = 1$; the second is $2/p$. The decompositions differ in their first-update access to messages while retaining the same eigenvalues. At $p = 1$, the second has constant one at every horizon; the first has infinity at $H = 1$ and one thereafter. At $p = 0$, both are one.

For the three-cycle and $\mu = (1, 0, 0)$, a channel is missed if $H < 3$. For $H \geq 3$, deterministic visits expose every positive and negative coordinate, so $p_H(\xi) = \|\xi\|_\infty = p_\infty(\xi)$. This proves both periodic claims without a mixing assumption.

G Numerical verification of horizon amplification

We evaluate three analytic families using rational recurrences, direct network evaluation on graphs, and the stopping LP, implemented in `minagg_horizon_teaching.py`. All parameters and equalities below are rational.

Table 4: Certified horizon amplification constants. Brackets enclose the exact infinite-horizon constants. The $H = 8$ cases use the closed form to avoid exponential vertex enumeration.

Construction	Horizon	Optimal constant or certified bracket
$p = 1/64,\ (A, B) = (P, I)$	1	∞
Same	2	$127 \in [127, 127.0000005]$
Same	8	$8337027150271/459065871937$
$p = 1/64,\ (A, B) = (I, P)$	1	$128 \in [128, 128.0000005]$
Same	2	64
Same	8	$4201336516672/262355877505$
$p = 1/2,\ s = 3/4$	2	$12/7$
Deterministic three-cycle	1, 2	∞
Same	3	$[1, 1]$

In the two slow-transition cases evaluated by the vertex algorithm, the teaching polytopes have five and four vertices. The phase-tail calculation uses period one, $n_0 = 64$, and $N = 3968$, with final bracket width 5×10^{-7}. When coverage is incomplete, explicit cost directions have zero teaching risk and positive all-length risk. For the three-cycle, eight vertices, period three and a zero tail give exactly one. The mixed-cost example $c = (0, 10/7)$ has teaching risk $1/4$ and all-length risk $3/7$, illustrating the scaling into the cost box when $z^\star > 1$. At $p = 1/64$, exact rational powers establish the cutoff of 195 steps: 194 steps exceed the 1.1 amplification target, whereas 195 satisfy it.

We evaluate the stationary and coefficient-path witnesses in float64 on one CPU thread. The maximum discrepancy between network outputs and the endpoint formulas is 2.23×10^{-16}. Proofs and rational tail bounds certify the infinite horizon.

Total computation took 8.426 seconds of CPU wall time.

The numerical implementation uses exhaustive rational policy and vertex enumeration, with numerical LP candidates verified by exact rational certificates. Candidate reconstruction is heuristic; the termination guarantee in Theorem 10 uses exact rational LP algorithms.

H LEARNED-COST TEACHING-LENGTH CASE STUDY

We compare source-only and certificate-based teaching-length selection.

Model, teaching and fitting. Let $u = 63/64$, $J = \mathbf{1}\mathbf{1}^\top/3$, $A = uI + (1-u)J$, $B = I$, $e = \mathbf{1}$ and $r = (1, 0, 0)$. Only the three entries of $c \in [0, 2]^3$ are trained. For seeds $0, 1, 2$, NumPy's default random generator samples c^0 independently from $U(0.5, 1.5)$; each tested length $H \in \{2, 8, 32, 128\}$ restarts from the same c^0 within a seed. We use the complete contrast-two teaching set of $H + 2$ labeled endpoints with unit targets. Its path outputs are 1 and $rA^{H-j}c$ for $1 \le j \le H$; its stationary output is the first component of ℓ_H, where $\ell_0 = \mathbf{1}$ and $\ell_{k+1} = \min(A\ell_k, c)$. These identities evaluate the finite teaching residuals directly.

Fitting minimizes the sum of squared relative teaching residuals by SciPy TRF/exact least squares with bounds $[0, 2]$, at most 2000 function evaluations, `ftol = xtol = gtol` $= 10^{-12}$, and `x_scale` $= 1$. Jacobians use forward-mode automatic differentiation. The fitted model is the first accepted iterate, including initialization, whose maximum source residual is at most 10^{-3}. All fits reached this source-only stopping threshold in three to five accepted iterates, counting the initial iterate.

Certificate and selection. Here $S = A$, $\mu = r$ and $rA^j = u^j r + (1 - u^j)(1/3, 1/3, 1/3)$. For $H \ge 2$, the smallest exposure in Theorem 9 is $(1 - u^{H-1})/3$, giving $C_H^+ = 6/(1 - u^{H-1})$. The source-only rule chooses the first tested length meeting $\delta_H \le 10^{-3}$. The certificate rule requires this condition and $C_H^+ \delta_H \le 10^{-2}$. Both rules select from the same four-point grid. The reported source risks and bounds use the exact binary-rational representations of the fitted float64 costs and the rational connections. Long-run evaluation follows model selection.

Exact all-length evaluation. For the connections above, writing $\xi = c - \mathbf{1}$ gives

$$\rho_\infty(c) = \max\{0, \max_i(-\xi_i), \xi_1, (\xi_1 + \xi_2 + \xi_3)/3\}. \tag{47}$$

Indeed, this finite chain is irreducible, so a stopping policy can eventually reach a coordinate with maximum negative deviation with probability one. Its stopping value is $\max(0, \max_i(-\xi_i))$. The signed expectation is a convex combination of ξ_1 and the coordinate mean, monotone in j, so its supremum is their maximum. Proposition 8 proves (47).

We also evaluate the finite graph-uniform risks for every $K = 1, \ldots, 4096$ and compute predictions on the stationary two-node graph at $K = 128, 512$ and a 16-node recursive tree and cycle at $K = 128$. The latter graphs use seed 20260913 and initial values/edge weights in $U(0.1, 1)$; each tree node chooses a uniformly random previous node as its parent. All graphs include zero-weight self-loops. All predictions satisfy the certificate and finite-endpoint bounds.

Total experiment time, including checks and evaluation, was 2.237 seconds on one CPU thread, excluding interpreter startup and plotting.

The implementation is `minagg_learned_horizon.py`; its `--report` option regenerates the tables and Figure 2.

I STRUCTURED MULTI-CHANNEL CONSTANTS AND TEACHING LENGTHS

Within the normalized shared model, stationary and reset connections admit explicit optimal constants. The proofs combine the complete MinAgg risk representation with classical hitting and stationary-distribution properties (Levin & Peres, 2017).

Table 5: All twelve fits. Errors and bounds are percentages. Iterates count accepted solver points, including initialization.

Seed	H	Teaching error	All-length bound	ρ_∞	Iterates
0	2	0.07618	29.2537	14.6268	4
0	8	0.08037	4.6199	2.2392	3
0	32	0.01100	0.1708	0.0736	5
0	128	0.01116	0.0774	0.0230	5
1	2	0.04706	18.0709	9.0355	3
1	8	0.01395	0.8019	0.3306	4
1	32	0.05038	0.7825	0.3157	4
1	128	0.00078	0.0054	0.0016	5
2	2	0.00427	1.6384	0.8192	4
2	8	0.00873	0.5017	0.2062	5
2	32	0.03315	0.5149	0.2112	5
2	128	0.00499	0.0346	0.0102	5

I.1 Proof of Theorem 12

Proof. A stationary distribution of a finite chain gives zero mass to transient states. Its positive support is a union of closed irreducible classes. Write these classes as C with masses $w_C = \mu(C) > 0$, and put $M_C = \max_{i \in C}(-\xi_i)_+$. Any stopped reward in class C is at most M_C. Conversely, every state in that class is eventually hit with probability one: from each state a positive finite path reaches the target, and the minimum of these finitely many path probabilities gives a positive success probability over a fixed-length block. Iterating the failure bound proves almost-sure hitting, without aperiodicity. Waiting for a maximizing state therefore gives

$$\mu z_\infty(\xi) = \sum_C w_C M_C.$$

For each class choose a maximizing state i_C. A finite stopping policy waits for i_C through time $H-1$, earning zero if it is not reached. The initial class is observed, so these class-specific policies combine into one admissible policy. Their value yields

$$\mu z_H(\xi) \geq \sum_C M_C h_{i_C}(H).$$

Let $D = \max_{i:\mu_i>0} w_i/h_i(H)$. Each denominator is at least $\mu_i > 0$, and $D \geq 1$ since $h_i(H) \leq w_i$. The two displays give $\mu z_\infty \leq D \mu z_H$. Every propagation value equals $\mu S^j \xi = \mu\xi$, already present in p_H. Consequently $p_\infty \leq D p_H$ for arbitrary mixed ξ.

For a coordinate attaining D, choose $\xi = -e_i$. Its finite and infinite stopping rewards are $h_i(H)$ and w_i; propagation is nonpositive. The ratio is exactly D, with $c = \mathbf{1} - e_i \in [0,2]^q$. Finally, avoiding i requires starting outside it and making $H-1$ transitions within the deleted-state matrix. This gives (27), including $H = 1$. One-state classes and multiple periodic classes require no modification. □

The state-readout condition $rBA = r$ implies $(rB)(AB) = rB$. This also works when $q \neq d$: r acts on the d state channels, whereas rB acts on the q message channels. The theorem holds for every connection satisfying this condition.

I.2 Proof of Theorem 19

Theorem 19 (Reset connections with sparse readout)**.** *Let $S = (1-\alpha)I + \alpha\mathbf{1}\pi$, $q \geq 2$, $0 < \alpha \leq 1$, $\pi_i > 0$, and $\operatorname{supp}\mu \subsetneq \{1,\dots,q\}$. Then $C_1^\star = \infty$. For $H \geq 2$, writing $n = H-1$,*

$$C_H^\star = \max\left\{ \frac{1}{\min_i[1-(1-\mu_i)(1-\alpha\pi_i)^n]},\ \frac{1+(1-\alpha)^n}{1-(1-\alpha)^n} \right\}. \tag{48}$$

Both terms have attaining cost directions in the permitted box.

Proof. Write $u = 1 - \alpha$ and $\Pi = \mathbf{1}\pi$. The identities $\Pi^2 = \Pi$ and $I\Pi = \Pi I$ give $S^j = u^j I + (1 - u^j)\Pi$ by induction. Since all entries of S are positive, every channel is eventually hit with probability one. The infinite stopping reward is therefore $M = \max_i(-\xi_i)_+$. The finite strategy that waits for a maximizing coordinate i earns $Mh_i(H)$, so the stopping amplification is at most $D_- = 1/\min_i h_i(H)$. The choice $\xi = -e_i$ at a minimizing coordinate attains it. Before i is first hit, the conditional chance of hitting it on the next step is exactly $\alpha\pi_i$, giving

$$h_i(H) = 1 - (1 - \mu_i)(1 - \alpha\pi_i)^{H-1}.$$

For $H \geq 2$ all these probabilities are positive.

For propagation put $a = \mu\xi$, $b = \pi\xi$, and $n = H - 1$. The values $u^j a + (1 - u^j)b$ are monotone in j; hence their finite maximum is the maximum at $j = 0, n$, and their infinite supremum is $\max(a, b)$. Let $\delta = p_H(\xi)$. Stopping immediately on the negative coordinates shows $a \geq -\mu(-\xi)_+ \geq -\delta$, whereas the last teaching propagation value gives $u^n a + (1 - u^n)b \leq \delta$. Therefore

$$b \leq D_+\delta, \qquad D_+ = \frac{1 + u^n}{1 - u^n}.$$

Together with $a \leq \delta$, this proves the upper bound in (48).

To attain D_+ let $J = \operatorname{supp}\mu$ and $s = \pi(J) < 1$. Set $\xi_i = -1$ on J and $\xi_i = b_+ = (D_+ + s)/(1 - s) > 0$ outside J. The chain starts in a negative coordinate with probability one, so its finite and infinite stopping values are both one. Moreover, $a = -1$, $b = D_+$, and the last teaching propagation value is one. Thus $p_H = 1$ and $p_\infty = D_+$. Dividing ξ by $\max(1, b_+)$ places every entry in $[-1, 1]$ without changing the ratio. The two attaining directions prove the maximum of D_- and D_+ is optimal.

At $H = 1$ a coordinate outside J has zero teaching risk but positive infinite risk under $\xi = -e_i$. At $\alpha = 1$ and $H \geq 2$, $u^n = 0$, so $D_+ = 1$ and the proof still applies. The sparse-support assumption ensures that the complement of J used in the attaining construction is nonempty. □

I.3 Proof of the architecture-family rule

In the class of (28), positivity makes S irreducible and $\mu = \pi$ is stationary. Starting from π, a path misses i initially with probability $1 - \pi_i$; conditional on not having hit it, the next step hits it with probability at least $\alpha\pi_i$. It follows that

$$1 - h_i(H) \leq (1 - \pi_i)(1 - \alpha\pi_i)^{H-1} \leq (1 - \beta)(1 - \alpha\beta)^{H-1}.$$

Theorem 12 gives the claimed upper bound. Choose a probability row with minimum entry β and the reset matrix $S = (1 - \alpha)I + \alpha\mathbf{1}\pi$. Such a row exists because $0 < \beta \leq 1/q$. Every displayed bound is then attained at a minimum-mass coordinate; placing a negative cost deviation there attains the constant. This proves the supremum over connections is an equality.

Consequently, for $\epsilon > 0$, the smallest horizon guaranteeing amplification at most $1 + \epsilon$ throughout the class is

$$H_{\min} = 1 + \max\left\{0, \left\lceil \frac{\log[\epsilon/((1 + \epsilon)(1 - \beta))]}{\log(1 - \alpha\beta)} \right\rceil\right\}. \tag{49}$$

Indeed, the bound is equivalent to $(1 - \beta)(1 - \alpha\beta)^{H-1} \leq \epsilon/(1 + \epsilon)$; the denominator in (49) is negative, and the outer maximum handles cases where $H = 1$ already suffices.

When $\beta = 1/q$, stationarity and $\pi_i \geq \beta$ force π to be uniform. The class therefore consists of doubly stochastic matrices with $S_{ij} \geq \alpha/q$, rather than one fixed trained matrix. The attaining reset chain has eigenvalues 1 and $1 - \alpha$, with the latter repeated $q - 1$ times. From any initial law ν, its total variation distance to π at time t is $(1 - \alpha)^t\|\nu - \pi\|_{\mathrm{TV}}$. In our case $\mu = \pi$, so the chain is stationary from the start. Even so, the sharp teaching length scales as $\Theta((q/\alpha)\log(1/\epsilon))$ for small ϵ. What matters is the probability of hitting each rarely observed channel. This is neither marginal mixing nor the time needed to visit all channels on one trajectory.

A readout tradeoff. Stationary readout gives an exact hitting formula but can have a larger constant than other readouts. Take $S = \mathbf{1}(1/100, 99/100)$. Stationary readout gives $C_1^\star = 100$ by Theorem 12. With $\mu = (1/2, 1/2)$, Proposition 18 gives

$$C_1^\star = \max\{2,\ 2(99/50) - 1/50\} = 197/50.$$

Eliminating propagation growth need not minimize the total short-horizon amplification when stationary mass is highly uneven.

I.4 Numerical verification and tighter learned-model bounds

We compare the closed formulas with the general vertex algorithm on two $H = 2$ cases using exact rational arithmetic in `minagg_structured_horizon.py`. The nonreversible stationary chain

$$S = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 0 & 1/3 & 2/3 \\ 3/4 & 0 & 1/4 \end{pmatrix}, \qquad \mu = (12, 9, 8)/29$$

has $C_2^\star = 29/14$ and 16 vertices. The three-channel mixed example after Theorem 19 has 20 vertices and $C_2^\star = 7$; $c = (8/9, 2, 8/9)$ gives risks $1/9$ and $7/9$, whereas the pure underestimation constant is five.

A disjoint two-cycle and three-cycle, with class masses $1/3, 2/3$ and stationary within-class readout, gives constants $3, 3/2, 1$ at $H = 1, 2, 3$. Deleted-state hitting probabilities agree exactly with indicator-reward Bellman recurrences; all 243 directions in $\{-1, 0, 1\}^5$ satisfy the respective bounds. For $\alpha = 1/4$, $\epsilon = 1/10$ and $q = 4, 16, 64$, exact rational powers verify that the sharp uniform cutoffs are $34, 150, 610$: each preceding integer fails and the stated integer passes.

The learned-cost connections in Appendix H are a reset chain with $\alpha = 1/64$, uniform π and $\mu = (1, 0, 0)$. Applying Theorem 19 gives the following sharper constants:

H	Original coverage bound	Optimal $C_H^\star$
2	384	192
8	57.4825	27.8601
32	15.5332	6.6909
128	6.9391	2.0627

All twelve bounds $C_H^\star \delta_H$ are no larger than the coverage bounds and cover the exact risks of the fitted models. For seed zero at $H = 2$, the bound drops from 29.2537% to 14.6268% and equals the actual all-length risk. The complete twelve-row comparison is provided in the supplement. For seed two at $H = 2$, the bound drops from 1.6384% to 0.8192%, certifying the model against the 1% target.

Computation took 5.063 seconds of CPU wall time, excluding interpreter startup. The two $H = 2$ cases use the general vertex algorithm; the other evaluations use the closed formulas and fitted costs.

J Readout design: proofs, constructions, and certificates

Throughout this appendix, $A \in \mathbb{R}_{\geq 0}^{q \times d}$ and $B \in \mathbb{R}_{\geq 0}^{d \times q}$ are row-stochastic, $q \geq 2$, and $S = AB$ is irreducible. The design variable is $r \in \Delta_d$ and $\mu = rB$. All risks use $c \in [0, 2]^q$ and $\xi = c - \mathbf{1}$. Irreducibility ensures that every channel is eventually visible under every realizable readout. All readouts are therefore compared on the same channels.

J.1 Proof of Theorem 13

For $\mu > 0$, write $\nu_j = \mu S^j$, $M_j = \max_i (\nu_j)_i / \mu_i$, $m_j = \min_i (\nu_j)_i / \mu_i$, and $L = 1/\min_i \mu_i$. Proposition 18 simplifies to

$$C_1(r) = \max\{L, \sup_{j \geq 0} (2M_j - m_j)\}. \tag{50}$$

Irreducibility makes every channel eventually reachable with probability one. The infinite stopping value is therefore $\max_i (-\xi_i)_+$. Since μ and ν_j are probability rows, $0 \leq m_j \leq 1 \leq M_j$. Dropping m_j in the exact formula raises the value by at most one, giving $C_1(r) \leq \mathcal{U}(r) \leq C_1(r) + 1$. Minimize $\mathcal{U}$ and compare with the infimum of C_1 to obtain the additive guarantee. The lower bound $C_1(r) \geq L \geq q$ then gives the ratio in (30). Finally, if $\mu_i = 0$, the direction $\xi = -e_i$ has zero one-step risk but unit infinite risk. Such a readout has $C_1(r) = \infty$.

Existence. Every column of B is nonzero, since otherwise AB would have a zero column. A uniform state readout therefore gives $\mu_0 > 0$ and a finite upper bound $U = 2/\min_i (\mu_0)_i$ on $\mathcal{U}$. The sublevel set $\mathcal{U}(r) \leq U$ is contained in the compact set $\{r \in \Delta_d : (rB)_i \geq 1/U\}$. All maps

$\mu \mapsto (\mu S^j)_i/\mu_i$ have a common Lipschitz bound there: the numerators lie in $[0,1]$, their linear coefficients lie in $[0,1]$, and denominators are bounded away from zero. Their supremum is thus continuous on this set. A minimizer exists.

A finite, certified approximation. For a fixed threshold $t > 0$, the condition $\mathcal{U}(r) \le t$ is equivalent to

$$r \ge 0, \quad r\mathbf{1} = 1, \quad (rB)_i \ge 1/t, \quad 2(rBS^j)_i \le t(rB)_i \quad (i = 1, \dots, q,\ j \ge 0). \tag{51}$$

At fixed t, these constraints are linear in r. The dependence on t makes the overall problem quasiconvex. For rational inputs, use the effective periodic decomposition from Appendix F.3. Let m be the period and $L_a = \lim_{n\to\infty} S^{mn+a}$ for $0 \le a < m$. For any rational $\eta > 0$, compute N such that

$$\max_{k,i} |(S^{mn+a} - L_a)_{ki}| \le \eta/(4U) \quad (n \ge N,\ 0 \le a < m).$$

Retain (51) for $j < mN$ and impose the same inequalities with every L_a in place of S^j. This finite LP feasibility problem relaxes the infinite one. Its optimum threshold is at least q and at most U. Exact rational LP feasibility and bisection give a feasible $\widehat{t}$ at most $\eta/2$ above that optimum, with a corresponding $\widehat{r}$. If the lower endpoint q is feasible, it can be returned directly. Otherwise bisection preserves an infeasible lower endpoint and a feasible upper endpoint.

For $\mu = \widehat{r}B$, feasibility gives $\mu_i \ge 1/U$ and, on the tail,

$$2\frac{(\mu S^{mn+a})_i}{\mu_i} \le 2\frac{(\mu L_a)_i}{\mu_i} + \eta/2 \le \widehat{t} + \eta/2.$$

Prefix constraints also control L. Therefore $\mathcal{U}(\widehat{r}) \le \inf_r \mathcal{U}(r) + \eta$ and $C_1(\widehat{r}) \le C_{1,\mathrm{opt}} + 1 + \eta$. Every periodic phase is retained, and the tail estimate is uniform over the candidate readouts. This gives a finite algorithm; its running time depends on the period and tail convergence rate. □

J.2 A NON-RESET FAMILY WITH AN EXPLICIT DESIGN

Proposition 20. *Let $B = I$ and let $S = A$ be irreducible, reversible, and have nonnegative spectrum. Suppose its stationary row is $\pi = (p, b, \dots, b)$, where $1/2 < p < 1$ and $b = (1-p)/(q-1)$. Writing $R = q - 1 + 2p$, the readout*

$$\mu_* = (2p/R, 1/R, \dots, 1/R) \tag{52}$$

satisfies $\mathcal{U}(\mu_*) = \min_\mu \mathcal{U}(\mu) = R$ and $C_1(\mu_*) = R$. Here $\mathcal{U}(\mu)$ denotes the envelope with $B = I$. Its ratio to the true optimal constant is at most R/q.

Proof. Set $a = 1/(Rb)$ and $b_0 = 2/R$. The initial density with respect to π is $a\mathbf{1} - (a - b_0)e_1$, with $a > b_0$. Reversibility gives

$$\frac{(\mu_* S^j)_i}{\pi_i} = a - (a - b_0)(S^j)_{i1}.$$

For $i \ne 1$, this bounds $(\mu_* S^j)_i$ by $ab = 1/R = (\mu_*)_i$. For $i = 1$, the reversible spectral expansion and nonnegative spectrum give $(S^j)_{11} \ge \pi_1 = p$. Since $a(1-p) + b_0 p = 1$, it follows that $(\mu_* S^j)_1 \le p$. Thus $M_j \le R/2$ and $L = R$, proving $\mathcal{U}(\mu_*) = R$. A negative perturbation of any rare channel attains $C_1(\mu_*) = R$.

Conversely, $\mathcal{U}(\mu) \le t$ requires $\mu_i \ge 1/t$ for $i \ne 1$ and $\mu_1 \ge 2p/t$, the latter from the limiting orbit. The chain is aperiodic, since an irreducible reversible chain of period two has eigenvalue -1. Summing the constraints yields $t \ge R$. Finally, $C_{1,\mathrm{opt}} \ge q$ proves the approximation ratio. □

The design is optimal for the envelope and within R/q of optimal for the exact objective. Laziness ($S_{ii} \ge 1/2$) is sufficient for the assumed nonnegative spectrum.

For a concrete non-reset instance take

$$S = \begin{pmatrix} 15/16 & 1/96 & 1/48 & 1/32 \\ 1/8 & 7/8 & 0 & 0 \\ 1/4 & 0 & 3/4 & 0 \\ 3/8 & 0 & 0 & 5/8 \end{pmatrix}, \qquad \pi = (4/5, 1/15, 1/15, 1/15). \tag{53}$$

This reversible lazy star has different exit rates and zero transitions between rare channels. The design is $(8, 5, 5, 5)/23$, with exact $C_1 = 23/5$. Stationary readout has $C_1 = 15$; uniform readout has $C_1 \geq 2qp - qb = 92/15$ from the limiting propagation row. Hence design reduces amplification by at least 25% relative to uniform readout and is within a factor $23/20$ of the true optimum. For fixed q and $p \to 1$, the design constant tends to $q + 1$, the uniform lower bound to $2q$, and the stationary constant diverges. These comparisons hold for the worst-case amplification over the cost box.

J.3 Multi-step optimality from hitting and propagation

Define $(T_H)_{ki} = \Pr_k(T_i < H)$, including time zero, and $V = BT_H$. Consider the hitting LP

$$t_* = \max_{r \in \Delta_d, t} \{t : rV \geq t\mathbf{1}^\top\}. \tag{54}$$

For every readout,

$$C_H(r) \geq \frac{1}{\min_i (rV)_i}, \qquad C_{H,\mathrm{opt}} \geq 1/t_*. \tag{55}$$

To prove the lower bound, take $\xi = -e_i$. Its finite risk is the probability of hitting i during teaching; its infinite risk is one. For a general direction, let $M = \max_i(-\xi_i)_+$. A policy that waits for a maximizing coordinate earns at least $M \min_i (rV)_i$ during teaching, while its infinite reward is M. Thus the inverse of the smallest hitting probability is exactly the amplification of the stopping component. The positive propagation component still needs its own bound. A feasible r_* and a probability column y satisfying

$$r_* V \geq t_* \mathbf{1}^\top, \qquad Vy \leq t_* \mathbf{1}$$

certify the LP optimum, since $\min_i (rV)_i \leq rVy \leq t_*$ for all r.

For probability rows ν and $\sigma > 0$, put $D(\nu, \sigma) = 2 \max_i \nu_i/\sigma_i - \min_i \nu_i/\sigma_i$.

Proposition 21 (A finite certificate for full-risk optimality)**.** *Let r_* solve* (54)*, and let $\mathcal{K} = \{k : 0 \leq k < H,\ \sigma_k = \mu S^k > 0\}$ be nonempty, where $\mu = r_* B$. If some $N \geq 1$ satisfies*

$$\min_{k \in \mathcal{K}} D(\mu S^j, \sigma_k) \leq 1/t_* \qquad (0 \leq j < N), \tag{56}$$

$$\min_{k \in \mathcal{K}} D(e_\ell^\top S^N, \sigma_k) \leq 1/t_* \qquad (1 \leq \ell \leq q), \tag{57}$$

then $C_H(r_*) = C_{H,\mathrm{opt}} = 1/t_*$.

Proof. At every teaching time k, waiting until k and then stopping only on negative coordinates yields $\sigma_k(-\xi)_+ \leq p_H(\xi)$. Teaching also observes $\sigma_k \xi \leq p_H(\xi)$. The one-step propagation bound from Proposition 18 therefore implies

$$\nu\xi \leq D(\nu, \sigma_k) p_H(\xi) \quad \text{for every probability row } \nu.$$

Condition (56) bounds the finite prefix. To handle every later time, choose a qualifying teaching anchor for each row $e_\ell^\top S^N$ in (57). Each row then evaluates to at most $p_H(\xi)/t_*$ for every ξ. Any probability mixture of the rows obeys the same bound. For $j \geq N$, the identity $\mu S^j = (\mu S^{j-N}) S^N$ expresses the later distribution as precisely such a mixture. This bounds propagation at all times; no convexity of the minimum over anchors is needed. Combining this with the stopping bound gives the upper bound. Equation (55) gives the matching lower bound. □

These conditions provide a sufficient test for global optimality.

A non-reset optimal design and an open neighborhood. For $B = I$, $H = 2$, and

$$S = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 0 & 1/3 & 2/3 \\ 3/4 & 0 & 1/4 \end{pmatrix}, \qquad r_* = (18, 21, 16)/55, \tag{58}$$

the hitting matrix has diagonal one and the same off-diagonal entries as S. The dual $y = (4, 4, 3)^\top/11$ and $t_* = 6/11$ satisfy the primal and dual constraints with equality. Using both teaching anchors and $N = 4$, the maximum prefix and tail expressions are $1091/756$ and $7535/4608$, respectively,

both below $11/6$. Thus $C_2(r_*) = C_{2,\mathrm{opt}} = 11/6$. Uniform readout has hitting lower bound 2 and matching propagation upper bound: its prefix and tail expressions at $N = 4$ are $11299/7920$ and $1463/864$. Stationary readout is $(12, 9, 8)/29$, with exact constant $29/14$.

This example belongs to an open family. Set $S_{ii} = 1 - p_i$, $S_{i,i+1} = p_i$ cyclically, and all other entries to zero, with $0 < p_i < 1$. Define $W = 3 - \sum_i p_i + \sum_{i<j} p_i p_j$. The positive rows

$$r_* = (1 - p_3 + p_2 p_3,\ 1 - p_1 + p_3 p_1,\ 1 - p_2 + p_1 p_2)/W,$$
$$y^\top = (1 - p_1 + p_1 p_2,\ 1 - p_2 + p_2 p_3,\ 1 - p_3 + p_3 p_1)/W$$

sum to one and satisfy $r_* T_2 = t_* \mathbf{1}^\top$, $T_2 y = t_* \mathbf{1}$, where $t_* = (1 + p_1 p_2 p_3)/W$. At $(p_1, p_2, p_3) = (1/2, 2/3, 3/4)$, the finite propagation certificate has strict margin $913/4608$. Its finitely many expressions are continuous near this point, so the complete optimality conclusion persists on a nonempty open neighborhood of this point.

J.4 Realizability and limits of simpler design rules

Rectangular interfaces. Let $q = 3, d = 2, H = 2$ and

$$A = \begin{pmatrix} 1 & 0 \\ 1/2 & 1/2 \\ 0 & 1 \end{pmatrix}, \qquad B = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \end{pmatrix}.$$

Every realizable μ has $\mu_2 = 1/2$. At $r_* = (1/2, 1/2)$, $\mu_* = (1/4, 1/2, 1/4)$ is stationary and $\mu_* T_2 = (3/8, 3/4, 3/8)$. The hitting dual $y = (1/2, 0, 1/2)^\top$ gives $BT_2 y = (3/8)\mathbf{1}$. Theorem 12 therefore certifies $C_{2,\mathrm{opt}} = 8/3$. The row $(1/2, 0, 1/2)$ would have minimum hitting probability $1/2$, but violates the constraint $\mu_2 = 1/2$. The design must therefore retain the realizability constraint $\mu = rB$.

Hitting-only design can be suboptimal. Take $S = \mathbf{1}(a, 1 - a)$, $B = I$, $H = 1$, and $0 < a \le 1/3$. The hitting LP selects uniform readout, whose full constant is $4 - 6a$. However,

$$\mu_{\mathrm{opt}} = \left(\frac{1 + a}{3 - a}, \frac{2(1 - a)}{3 - a} \right), \qquad C_{1,\mathrm{opt}} = \frac{3 - a}{1 + a}. \tag{59}$$

For any $\mu = (x, 1 - x)$, (50) gives the lower bound

$$C_1(\mu) \ge \max \left\{ \frac{1}{x}, \frac{2(1 - a)}{1 - x} - \frac{a}{x} \right\}.$$

The first term decreases and the second increases in x. They intersect at $x = (1 + a)/(3 - a)$, where $a \le x \le 1/2$. Substitution in the complete formula attains this lower bound and proves (59). For $a = 1/100$, hitting design has constant 3.94, while the optimum is $299/101$.

When uniform readout is already optimal. If S is doubly stochastic and uniform message readout is realizable, then $C_1 = q = C_{1,\mathrm{opt}}$ at uniform readout. More generally, suppose S is invariant under a transitive permutation group. Then T_H has equal row and column sums, say b_H. For every μ, $\sum_i (\mu T_H)_i = b_H$, so $C_H(\mu) \ge q/b_H$. Uniform readout is stationary and has all hitting probabilities b_H/q; Theorem 12 attains the bound. In this symmetric class, uniform readout is already optimal at every horizon.

J.5 Numerical verification

We evaluate the non-reset multi-step optimum, rectangular realizability, hitting-only design, and one-step envelope design using `experiments/minagg_readout_design.py`. For (53), 15 numerical feasibility LPs generate the rational candidate $(3277, 2048, 2048, 2048)/9421$. A prefix of length 160 and an exact rational tail certify an envelope upper bound of 4.60009765625, within 0.000098 of the analytic optimum 4.6. Together with the analytic lower bound in Proposition 20, this certifies the approximation gap.

Using time zero alone as a teaching anchor gives an upper bound of $241/126$ at $j = 2$, above the optimum in (58). Anchoring at all observed teaching times yields the matching certificate. The prefix and rational tail bounds cover all execution lengths.

K ADDITIONAL LITERATURE CONTEXT

Section 2 compares the closest guarantees and states our contribution relative to them. This appendix develops the broader connections to graph representation, teaching, algebraic computation, and positive-system design.

K.1 NEURAL ALGORITHMIC REASONING AND THE CLOSEST BF GUARANTEES

Neural algorithmic reasoning seeks neural computations that implement algorithmic structure and generalize beyond the training regime (Veličković & Blundell, 2021). Step-wise graph-algorithm execution explicitly supervises intermediate operations, including BF and reachability (Veličković et al., 2020). Algorithmic alignment connects a network's computational organization to learning complexity (Xu et al., 2020), and the algebraic relation between message passing and dynamic programming has been developed beyond individual shortest-path examples (Dudzik & Veličković, 2022). Extrapolation analyses further emphasize encoding the appropriate task-specific nonlinearities in the architecture or features (Xu et al., 2021b). We adopt this algorithmic bias through coordinatewise minimum aggregation.

The CLRS benchmark provides a common collection of algorithmic tasks (Veličković et al., 2022). Generalist neural algorithmic learners study shared processors across many such tasks (Ibarz et al., 2022), while reasoning-modulated representations incorporate pretrained reasoning modules into representation learning (Veličković et al., 2022). These works motivate algorithms as reusable computations and evaluate learning across tasks. Our analysis complements that setting with exact guarantees for a specified BF model, over every admissible graph and execution length.

The closest teaching guarantees for BF are the near-optimal L_0 and weighted-L_1 results of Nerem et al. (2026); Wittig et al. (2026). Wittig et al. (2026) already identify vector-min underestimation and teaching inputs that expose it; positive walk coefficients are also prior tools. Our fixed-length results give attainable two-sided risk for arbitrary parameters in the nonnegative class. They also characterize channel alignment and the sharp stability of approximate teaching. For normalized shared networks, we control every future length, uniformly over bounded costs at fixed connections. Repeated scalar BF calls admit guarantees from an accurate scalar module (Nerem et al., 2026). Our model retains vector-valued hidden state between updates, which introduces delayed channel exposure and cancellation.

Exact graph-instruction learning with infinite-width MLP ensembles and binary representations provides guarantees for a different architecture (Qharabagh et al., 2026). MINAR examines generalizing algorithmic circuits and already observes them before final sparsification (He et al., 2026). These results establish that generalizing circuits can be dense and that finite teaching can identify algorithms. We quantify how much error finite teaching can hide and choose readouts to reduce that amplification.

K.2 GRAPH REPRESENTATION, EXPRESSIVITY, AND DEPTH

The original GNN model (Scarselli et al., 2009), gated graph-sequence networks (Li et al., 2016), graph convolutional networks (Kipf & Welling, 2017), and the message-passing framework (Gilmer et al., 2017) establish different ways to organize graph computations. Relational inductive biases motivate structured computation for combinatorial generalization (Battaglia et al., 2018). Permutation-invariant set architectures (Zaheer et al., 2017) and characterizations of invariant/equivariant graph layers (Maron et al., 2019b) describe architectural symmetries. Our readout acts on channels through a probability row. Its design preserves node-permutation equivariance.

Weisfeiler–Leman analyses characterize the graph-discrimination power of message-passing GNNs and motivate higher-order models (Xu et al., 2019; Morris et al., 2019). Higher-order invariant networks provide stronger graph distinctions (Maron et al., 2019a); invariant and equivariant universality results describe which functions suitable architectures can approximate (Maron et al., 2019c; Keriven & Peyré, 2019). Subgraph counts (Bouritsas et al., 2021) and simplicial message passing (Bodnar et al., 2021) offer additional routes beyond standard message-passing expressivity. These are representational results or architectural proposals. We allow arbitrary channel width and study whether teaching observes the computations that can affect BF error. This concerns the accuracy of a specified algorithm across graphs.

Depth and width also impose computational constraints independent of optimization (Loukas, 2020). Oversmoothing analyses describe loss of node information under repeated graph propagation (Oono & Suzuki, 2020). Oversquashing concerns compression of distant-node information (Alon & Yahav, 2021); curvature and topology help characterize its bottlenecks (Topping et al., 2022), with width, depth, and graph commute times having distinct effects (Di Giovanni et al., 2023). These mechanisms concern propagation between graph vertices. Our Markov chain instead acts on internal message channels. The exact-risk representation has already taken the worst case over input graphs.

K.3 Generalization and distribution shift

Representational limits and data-dependent GNN generalization bounds are studied together by Garg et al. (2020). VC-dimension analyses relate distinguishability and statistical capacity (Morris et al., 2023); PAC-Bayesian bounds connect graph diffusion and noise stability to generalization (Ju et al., 2023). The WL distance compares labeled measure Markov chains and provides structural distances compatible with graph discrimination (Chen et al., 2022). Most directly at the expressivity–generalization intersection, Maskey et al. (2026) derive bounds involving structural train–test distance, model complexity, and sample size. Their results explain how additional expressivity can help or hurt statistical generalization. We quantify a different gap: the ratio between global algorithmic error and the residual on a finite teaching set.

Size generalization can fail when local structures change with graph size, even at good training minima (Yehudai et al., 2021). SizeShiftReg simulates size shifts through coarsening to regularize predictions (Buffelli et al., 2022). Stochastic graph attention uses an information-bottleneck principle to select task-relevant subgraphs under stated assumptions (Miao et al., 2022). These works study statistical prediction and structural distribution shift. Our setting is deterministic: the target algorithm is known, and risk is a supremum over admissible graph sizes, attributes, and execution lengths. The readout design uses the network connections without a test-graph distribution.

K.4 Machine teaching, exact identification, and optimization

Teaching dimension measures how many examples identify a target concept within a hypothesis class (Goldman & Kearns, 1995). The wider machine-teaching framework separates assumptions about the learner, teacher, and teaching objective (Zhu et al., 2018). Iterative machine teaching explicitly aims to accelerate learner convergence through sequential examples (Liu et al., 2017), while black-box teaching studies interaction when the learner's representation is unknown (Dasgupta et al., 2019). Query-based automata learning gives another classical example of obtaining exact behavior from a finite interactive process (Angluin, 1987).

Our $K + 2$ endpoints constitute a fixed, noninteractive teaching set for a specified real-valued model. The central distinction is between identification at zero residual and stability at nonzero residual. We give a sufficient teaching set and quantify its stability for each parameter vector. Convergence analyses instead study optimization, for example through depth and skip connections in linearized GNNs (Xu et al., 2021a). Our paired experiment isolates the effect of one additional observation. Our design objective is the resulting risk-amplification constant.

K.5 Algebraic paths and weighted automata

BF belongs to the classical shortest-path dynamic-programming tradition (Bellman, 1958). Its optimality in restricted circuit models has been studied independently of neural learning (Jukna & Schnitger, 2016). Weighted automata provide algebraic representations and algorithms for weighted path computations (Mohri, 2009), and tropical geometry analyzes piecewise-linear neural representations (Zhang et al., 2018). Max-plus system identification also studies how input design reveals unknown system parameters (Schullerus et al., 2006). In our model, each coordinatewise minimum can choose a different neighboring vertex. Nonnegative channel mixing combines these choices, producing the additional risk terms we analyze.

Finite equivalence and minimization methods for weighted automata (Kiefer, 2020), together with arithmetic-circuit views of recurrent GNNs (Barlag et al., 2026), inform our exact shared-execution criterion. Quantitative weighted-automaton metrics compare infinite behaviors using seminorms and spectral conditions (Balle et al., 2017). Our infinite-horizon objective is an undiscounted ratio of

all-length BF risk to observed finite-horizon risk. The ratio is uniform over a box of edge coefficients. Its value depends on both optimal stopping and cancellation between positive and negative deviations.

K.6 Markov chains, error bounds, and positive-system design

Finite-chain hitting, stationary structure, and spectral mixing are classical (Levin & Peres, 2017). Periodic ergodic projectors can be computed without assuming aperiodicity (Berkhout & Heidergott, 2019), and stopping/MDP problems admit linear-programming formulations (Kallenberg, 2013). Hoffman-type bounds quantify polyhedral residuals through constants depending on the constraint geometry (Peña et al., 2021). We use these established mechanisms after deriving the exact MinAgg risk. Coverage gives finiteness, while the geometry of stopping and propagation gives the sharp constant, including their cancellation.

Optimizing a Markov chain is itself well established: fastest-mixing design chooses transition probabilities to improve a spectral objective (Boyd et al., 2004). Experiment design in Markov chains optimizes policies that control where measurements are collected (Mutny et al., 2023). Our design keeps the transition matrices fixed and chooses a realizable readout to minimize risk amplification. The objective combines channel exposure with signed propagation. The hitting-only counterexample in Appendix J shows why the second term can change the optimal design.

Control theory also studies peak behavior quantitatively. Balakrishnan & Boyd (1992) compute worst-case peak gains using bounds on impulse-response tails. Parameter tuning for positive linear systems can be reduced to geometric programming for several system norms (Ogura et al., 2020). Impulse-to-peak optimization has also been formulated via difference-of-convex programming (Zhao et al., 2023). These works provide precedents for peak optimization, tail certification, and positive-system design. We also use standard fixed-threshold feasibility and quasiconvex bisection (Boyd & Vandenberghe, 2004).

The readout contribution is the bound $C_1(r) \le \mathcal{U}(r) \le C_1(r) + 1$ for the complete MinAgg mixed risk. It gives a cost-independent, realizable design within an additive $1 + \eta$ of the true optimum. For multi-step teaching, a hitting LP supplies a lower bound, while observed teaching times provide anchors for bounding later propagation. Matching these bounds proves exact optimality. The non-reset examples demonstrate how this objective differs from uniform, stationary, and hitting-only readout choices.