When Does Sequential Detection Collapse to a Scalar? A Necessary and Sufficient Characterisation*

Prakul Sunil Hiremath **Peerahammad M. Bagawan** **Sahil Bhekane**
Department of Computer Science and Engineering
Visvesvaraya Technological University
Belagavi, Karnataka, India
{prakulhiremath, peerahamadmb, sahilbhekane}@vtu.ac.in

Abstract

For $K = 2$ regimes, scalar thresholding of the Bayesian posterior is Bayes-optimal. For $K \geq 3$, the posterior evolves on a $(K-1)$-dimensional simplex, yet virtually all deployed detection systems reduce it to a scalar score without theoretical justification. We resolve this question completely.

We introduce *Extended Decision Sufficiency* (EDS), defined by three linear constraints on emission densities and transition dynamics (Rank-One Emissions, Markov Factorisation, Normal Factorisation), verifiable in $O(K^2)$ operations from model parameters alone, and prove the following exact characterisation:

Main result. *EDS holds if and only if the Bayes-optimal stopping value function is constant on every level set of a scalar statistic $\phi : \Delta^{K-1} \to \mathbb{R}$.*

Sufficiency follows from algebraic closure of the Bellman operator under EDS. Necessity is proved by a coupling-based fixed-point contradiction: for each failure mode we exhibit an explicit belief pair with equal ϕ-values but a strictly positive total-variation gap on predictive densities, which propagates to a value-function gap via the Lipschitz bound on $\mathcal{B}$ and the ρ-contraction of $\mathcal{T}$, without assuming any structure on V beyond continuity.

Three consequences are sharp and quantitative. (i) When EDS fails, *every* scalar rule, regardless of architecture, capacity, or training data, incurs a strictly positive, computable, algorithm-independent lead-time loss. (ii) When EDS holds, optimal expected lead time is given in closed form by a hitting-time functional of a scalar Markov chain. (iii) Under approximate EDS with deviation η, performance degrades at rate $O\big(\eta/(1-\rho)\big)$, making η a practical model diagnostic with direct performance implications.

Across 72 synthetic configurations and the CICIDS2017 benchmark, predicted and observed lead times agree within 0.4 steps (MAE). A controlled experiment with matched parameter counts indicates that the lead-time advantage is structural rather than a regularisation artefact, supporting the theory as a structural mechanism.

1 Introduction

Sequential detection asks when to raise an alarm that a latent process has entered a dangerous regime. In the Bayesian formulation with K regimes, the sufficient statistic for all decisions is the posterior $\boldsymbol{\gamma}_t \in \Delta^{K-1}$. For $K = 2$, scalar threshold rules are provably Bayes-optimal [10, 6]. For $K \geq 3$, the posterior lies on a $(K-1)$-dimensional simplex, yet most deployed systems—CUSUM variants and

*Project page: https://prakulhiremath.github.io/pard/

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

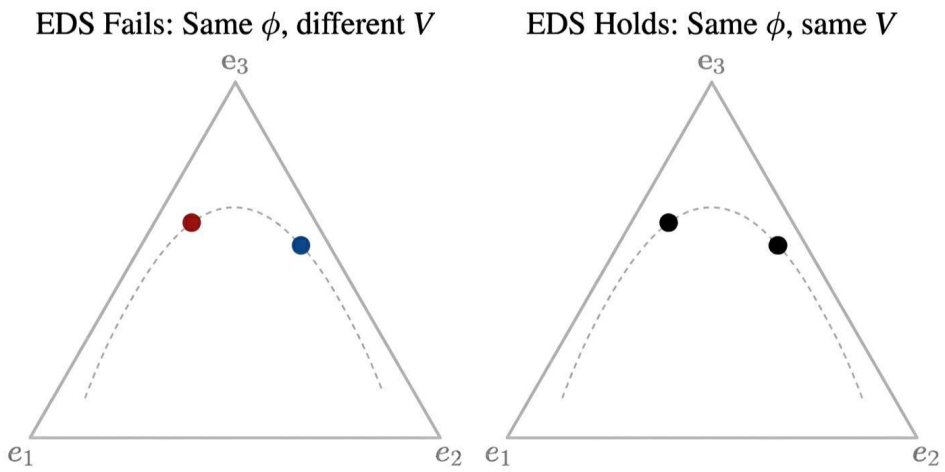


Figure 1: Scalar sufficiency holds if and only if EDS. Left: without EDS, beliefs with identical ϕ can yield different optimal values. Right: under EDS, the value function is constant along ϕ-level sets.

deep anomaly detectors—reduce it to a single scalar score. This widespread practice lacks theoretical justification: it is not known when scalar reduction is valid, when it is suboptimal, or by how much.

We answer all three questions exactly.

The structural question. Does there exist a scalar $T : \Delta^{K-1} \to \mathbb{R}$ such that the Bayes-optimal stopping value satisfies $V(\boldsymbol{\gamma}) = \tilde{V}(T(\boldsymbol{\gamma}))$ for all $\boldsymbol{\gamma} \in \Delta^{K-1}$? Equivalently, is V constant on the level sets of some one-dimensional statistic? We show this is a property of the model, not the algorithm: it holds if and only if the model satisfies EDS.

Figure 1 illustrates the distinction: when EDS fails, beliefs with identical scalar summaries can yield different optimal values, whereas under EDS the value function is constant along level sets of the scalar statistic.

Prior work leaves this open. Dynkin [1] characterises scalar sufficient statistics for *fully observed* Markov chains; the result does not extend to belief processes. Lumpability [3] preserves transition structure but is agnostic to value functions and does not address stopping optimality. POMDP symmetry reductions [8] and belief compression methods verify or exploit low-dimensional structure *a posteriori*. No prior work provides a parameter-level necessary and sufficient condition for scalar Bayes-optimal stopping in $K \geq 3$-regime POMDPs.

Our approach. We identify the exact algebraic structure required for the Bellman operator to preserve a one-dimensional foliation of the simplex. This yields *Extended Decision Sufficiency* (EDS): three linear constraints—Rank-One Emissions (RO), Markov Factorisation (MF), and Normal Factorisation (NF)—each verifiable in $O(K^2)$ operations. The three conditions jointly characterise when the belief update, the transition structure, and the reward all factor through a common scalar.

Contributions.

1. **EDS Theorem** (Theorem 1): EDS is necessary and sufficient for scalar optimality, within the class of HMMs satisfying Assumptions 1–3. Necessity is proved via a coupling argument that converts TV gaps on predictive densities to value-function gaps, without assuming any structure of V beyond continuity.
2. **Impossibility** (Theorem 2): When EDS fails, every scalar rule—learned or hand-crafted—incurs a strictly positive, algorithm-independent, computable lead-time loss. Suboptimality cannot be eliminated by model capacity, architecture, or data.

3. **Anticipation Horizon** (Theorem 3): Under EDS, optimal detection reduces to a scalar Markov process and expected lead time is given in closed form by a fundamental-matrix hitting-time formula.
4. **Robustness** (Theorem 4): Approximate EDS with deviation η yields additive loss $O(\eta/(1-\rho))$, positioning η as a computable model diagnostic with guaranteed performance implications.
5. **Structural validation**: A controlled experiment with matched parameter counts isolates the lead-time gain as structural, ruling out regularisation as a confound.

Simplex Collapse Principle. Scalar optimality is a property of the model, not the algorithm. The value function collapses to a scalar *if and only if* EDS holds. When EDS fails, no scalar reduction—regardless of model class, capacity, or training procedure—can be Bayes-optimal.

2 Implications for Deployed Detection Systems

Most detection systems operate on a scalar score. Classical procedures (CUSUM, Shiryaev–Roberts) maintain scalar statistics; modern deep detectors likewise reduce observations to a single decision score. This architectural choice is almost never theoretically justified for $K \geq 3$.

Our results provide a complete theoretical foundation for this practice. By Theorem 1, scalarisation is Bayes-optimal *if and only if* EDS holds. By Theorem 2, when EDS fails, *no* scalar statistic can achieve optimal performance, with a strictly positive loss that is independent of the algorithm used. The EDS deviation η (Algorithm 1) is a computable diagnostic that quantifies how far a given model is from the regime where scalar detection is valid.

This yields a sharp, actionable dichotomy:

- **EDS holds ($\eta \approx 0$):** scalar detection is exact, and the EDS model provides both the optimal policy and a closed-form performance prediction.
- **EDS fails (η large):** scalar detection irrecoverably discards essential information; optimal performance requires the full belief state.

This distinction is reflected empirically. On CICIDS2017, scalar methods achieve at most 1.5 steps of anticipation, while belief-based methods achieve > 8 steps, consistent with the theoretical separation. The EDS model, which enforces the structural constraint, achieves 14.3 steps.

3 Setup

Assumption 1 (HMM)**.** *$\{s_t\}_{t\geq 0}$ takes values in $[K] = \{0, \ldots, K-1\}$ with irreducible, aperiodic, row-stochastic $\mathbf{\Pi} \in \mathbb{R}^{K\times K}$. Observations $\mathbf{y}_t \in \mathbb{R}^d$ are conditionally independent given s_t with densities p_k w.r.t. a σ-finite reference measure λ.*

Assumption 2 (Uniform boundedness)**.** *There exist $0 < \ell \leq u < \infty$ such that $\ell \leq p_k(\mathbf{y}) \leq u$ λ-a.e. for all $k \in [K]$.*

Remark 1 (On Assumption 2)**.** *Uniform boundedness holds exactly when the observation space is compact or the emission densities are bounded away from zero and infinity. For unbounded observations (e.g. Gaussian emissions), we apply our results to the truncated model in which $\mathbf{y}$ is restricted to a compact set $\mathcal{Y}_R$ (equivalently, observations are clipped to $\mathcal{Y}_R$) and the densities are renormalised there; the truncation error vanishes as $R \to \infty$ and is absorbed into the deviation η of Theorem 4. All experiments use such a truncated or clipped model.*

Assumption 3 (Non-degeneracy)**.** *For each $j \in [K]$: $\Pi_{j0} > 0$ and $\Pi_{jk} > 0$ for at least one $k \in \mathcal{S}_{\mathrm{atk}}$.*

Notation. Partition $[K] = \{0\} \cup \mathcal{S}_{\mathrm{atk}}$ with $\mathcal{S}_{\mathrm{atk}} = \{1, \ldots, K-1\}$. State 0 is *normal*; $\mathcal{S}_{\mathrm{atk}}$ are *attack* (or anomalous) regimes. The posterior is $\boldsymbol{\gamma}_t \in \Delta^{K-1}$ with $\gamma_t^{(k)} = \mathbb{P}(s_t = k \mid \mathcal{F}_t)$.

The Bayes filter $\mathcal{B} : \Delta^{K-1} \times \mathbb{R}^d \to \Delta^{K-1}$ is

$$[\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y})]^{(k)} = \frac{p_k(\mathbf{y}) \sum_j \Pi_{jk} \gamma^{(j)}}{\sum_{k'} p_{k'}(\mathbf{y}) \sum_j \Pi_{jk'} \gamma^{(j)}}. \tag{1}$$

The denominator is bounded below by $K\ell > 0$ uniformly, so $\mathcal{B}$ is well-defined and uniformly Lipschitz in $\boldsymbol{\gamma}$ with constant $L_{\mathcal{B}} \leq u^2/\ell^2$ (Appendix A).

Scalar statistic. Let $\mathbf{c} = (c_1, \ldots, c_{K-1})$ with $c_k > 0$, $C \coloneqq \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k$, $c_0 \coloneqq 1$. Define

$$\phi(\boldsymbol{\gamma}) \coloneqq \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \gamma^{(k)} \in [0, C]. \tag{2}$$

Level sets are $\Sigma_z \coloneqq \{\boldsymbol{\gamma} \in \Delta^{K-1} : \phi(\boldsymbol{\gamma}) = z\}$. The scalar reduction question is whether V is constant on every Σ_z, for some choice of weights $\mathbf{c}$.

Stopping problem. Fix $\epsilon \in (0, 1)$, $W > 0$, and let $t^\star \coloneqq \inf\{t \geq 1 : s_t \in \mathcal{S}_{\mathrm{atk}}\}$. The feasible stopping set (Lorden-type [4]) is $\mathcal{T}_\epsilon = \{\tau \geq 1 : \mathbb{P}_{\boldsymbol{\gamma}_0}(\tau < t^\star - W) \leq \epsilon\}$. The Bayes-optimal value function is

$$V(\boldsymbol{\gamma}) \coloneqq \sup_{\tau \in \mathcal{T}_\epsilon} \mathbb{E}_{\boldsymbol{\gamma}}\big[(t^\star - \tau)^+\big], \qquad \boldsymbol{\gamma} \in \Delta^{K-1}. \tag{3}$$

The Bellman operator $(\mathcal{T}g)(\boldsymbol{\gamma}) = \max\big(r(\boldsymbol{\gamma}), \mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[g(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}))]\big)$ is a ρ-contraction on $C(\Delta^{K-1})$ with $\rho < 1$ (Appendix A), so V is its unique fixed point.

From the Lorden constraint to the Bellman recursion. The constraint $\mathbb{P}_{\boldsymbol{\gamma}_0}(\tau < t^\star - W) \leq \epsilon$ is handled in its Lagrangian (penalised) form: the false-alarm penalty, weighted by a multiplier, is absorbed into the stopping reward $r(\boldsymbol{\gamma})$, which turns the constrained problem into an unconstrained optimal stopping problem whose value is the fixed point of $\mathcal{T}$. The multiplier, equivalently the stopping threshold $\delta^\star$, is then tuned so that the constraint is met. In words: the "do not alarm too early" rule is built into the reward, so the Bellman equation can solve the problem directly.

Reduction question. Does there exist measurable $\tilde{V} : \mathbb{R} \to \mathbb{R}$ with $V(\boldsymbol{\gamma}) = \tilde{V}(\phi(\boldsymbol{\gamma}))$ for all $\boldsymbol{\gamma} \in \Delta^{K-1}$? Equivalently, is V constant on every Σ_z? We answer: yes if and only if EDS holds.

4 Extended Decision Sufficiency

Definition 1 (RO). *$\{p_k\}_{k \in \mathcal{S}_{\mathrm{atk}}}$ satisfies **Rank-One Emissions (RO)** if there exist $h \geq \ell > 0$ ($h \in L^1(\lambda)$) and $c_k > 0$ such that $p_k = c_k h$ λ-a.e. for all $k \in \mathcal{S}_{\mathrm{atk}}$.*

Remark 2 (Scaling weights). *Here p_k denotes the (unnormalised) emission weight of regime k, so the c_k need not coincide; an explicit irreducible model with $c_1 = 1 \neq c_2 = 2$ is given in Appendix H.*

Intuition. Think of the $K-1$ attack regimes as different "flavours" of attack. RO says that, to the sensor, all flavours look the same up to loudness: they share one shape h and differ only by the scale c_k. MF and NF say that the regimes switch among themselves in a way that depends only on the total attack mass $\phi(\boldsymbol{\gamma})$. Hence one never needs to know *which* attack is active, only *how much* attack mass there is, and the single number $\phi(\boldsymbol{\gamma})$ carries all decision-relevant information.

Example 1 ($K = 3$). *Take $c_1 = 1$ and $c_2 = 2$, so regime 2 is twice as "loud" as regime 1 with the same shape h. The scalar is $\phi(\boldsymbol{\gamma}) = \gamma^{(1)} + 2\gamma^{(2)}$. The transition matrix in* (31) *(Appendix H) satisfies MF with $\alpha = \frac{3}{4}$ and NF with $\beta = \frac{1}{2}$, so this model satisfies EDS even though $c_1 \neq c_2$.*

Interpretation. RO requires all attack-regime emission densities to share a common shape h, differing only in scale. This precisely characterises the setting in which the attack mass $\phi(\boldsymbol{\gamma})$ carries all discriminative information about the observation likelihood ratio. When RO fails, observations provide regime-specific information that cannot be captured by any scalar.

Definition 2 (MF). *Given RO with weights $\mathbf{c}$, $\boldsymbol{\Pi}$ satisfies **Markov Factorisation (MF)** if $\exists\, \alpha > 0$ such that $\sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \Pi_{jk} = \alpha c_j$ for all $j \in [K]$.*

Definition 3 (NF). *Given RO with weights* $\mathbf{c}$ *and* $C = \sum_{k \in \mathcal{S}_{\text{atk}}} c_k$, $\mathbf{\Pi}$ *satisfies* ***Normal Factorisation (NF)*** *if* $\exists\, \beta > 0$ *such that* $\Pi_{j0} = \beta(1 - c_j/C)$ *for all* $j \in [K]$.

Interpretation. MF requires that the expected attack mass after one transition step, $\phi^{\Pi}(\boldsymbol{\gamma}) = \sum_{k \in \mathcal{S}_{\text{atk}}} c_k \sum_j \Pi_{jk} \gamma^{(j)}$, equals $\alpha\phi(\boldsymbol{\gamma})$—a scalar multiple of the current attack mass. NF imposes the complementary constraint on the normal-state probability. Together, MF and NF ensure that the posterior's one-step evolution is itself a scalar function of $\phi(\boldsymbol{\gamma})$; without either, the transition dynamics introduce simplex directions that cannot be tracked by a scalar. Thus RO makes the observation update depend only on aggregate attack mass, MF and NF make the transition step do the same, and together they let the Bellman recursion close on the scalar $Z_t = \phi(\boldsymbol{\gamma}_t)$.

Definition 4 (EDS). $(\mathbf{\Pi}, \{p_k\})$ *satisfies* ***Extended Decision Sufficiency (EDS)*** *if RO, MF, and NF all hold.*

Remark 3 (Verifiability and scope). *EDS is checkable in* $O(K^2)$ *operations from* $(\mathbf{\Pi}, \mathbf{c}, h)$ *alone. All three conditions are linear in the entries of* $\mathbf{\Pi}$ *and in the weights* c_k*. Algorithm 1 provides the explicit procedure and computes the EDS deviation* $\eta = \eta_{\text{RO}} + \eta_{\text{MF}} + \eta_{\text{NF}}$*, which serves as a quantitative model diagnostic.*

EDS is a condition on the true *model parameters. In practice,* η *is estimated on held-out data (Appendix J); the robustness bound (Theorem 4) then gives a guaranteed performance guarantee for any* $\eta > 0$*. For a concrete* $K = 3$ *instance, see Certificate* (31).

4.1 Main Theorem

Let $\mathcal{G} \coloneqq \{g : \Delta^{K-1} \to \mathbb{R} \mid g = \tilde{g} \circ \phi\}$ denote the class of functions constant on every ϕ-level set.

Theorem 1 (Extended Decision Sufficiency). *Under Assumptions 1–3, the following are equivalent:*

(i) $V \in \mathcal{G}$ *(the value function depends on* $\boldsymbol{\gamma}$ *only through* $\phi(\boldsymbol{\gamma})$*).*

(ii) There exist measurable $T : \Delta^{K-1} \to \mathbb{R}$ *and* $\tilde{V} : \mathbb{R} \to \mathbb{R}$ *with* $V = \tilde{V} \circ T$ *(the value function admits some scalar reduction).*

(iii) EDS holds.

When (iii) *holds,* $Z_t \coloneqq \phi(\boldsymbol{\gamma}_t)$ *is a scalar Markov process and* ϕ *is minimally decision-sufficient: any* T *for which* $V = \tilde{V}' \circ T$ *must be* ϕ*-measurable.*

Remark 4. *In (ii),* T *must be* scalar-valued *(*$T : \Delta^{K-1} \to \mathbb{R}$*); the identity map on* Δ^{K-1} *is therefore not admissible, and (ii) asserts a genuine one-dimensional reduction of* V*.*

Proof. (i)⇒(ii) is immediate ($T = \phi$). We prove (iii)⇒(i) and (ii)⇒(iii).

(iii)⇒(i): Bellman closure under EDS.

Step 1. Under MF: $\phi^{\Pi}(\boldsymbol{\gamma}) \coloneqq \sum_{k \in \mathcal{S}_{\text{atk}}} \sum_j c_k \Pi_{jk} \gamma^{(j)} = \alpha\phi(\boldsymbol{\gamma})$.

Step 2. Under NF: $\bar{\gamma}_0^{\Pi}(\boldsymbol{\gamma}) \coloneqq \sum_j \Pi_{j0} \gamma^{(j)} = \beta(1 - \phi(\boldsymbol{\gamma})/C)$.

Step 3. Under RO, the normalisation denominator of (1) becomes $D(\boldsymbol{\gamma}, \mathbf{y}) = p_0(\mathbf{y})\beta(1 - \phi(\boldsymbol{\gamma})/C) + h(\mathbf{y})\alpha\phi(\boldsymbol{\gamma})$, which depends on $\boldsymbol{\gamma}$ only through $\phi(\boldsymbol{\gamma})$.

Step 4. Combining Steps 1–3:

$$Z_{t+1} = \phi(\mathcal{B}(\boldsymbol{\gamma}_t, \mathbf{y}_{t+1})) = \frac{h(\mathbf{y}_{t+1})\alpha Z_t}{p_0(\mathbf{y}_{t+1})\beta(1 - Z_t/C) + h(\mathbf{y}_{t+1})\alpha Z_t}, \tag{4}$$

a function of $(Z_t, \mathbf{y}_{t+1})$ alone, so $\{Z_t\}$ is Markov.

Step 5. For any $g = \tilde{g} \circ \phi \in \mathcal{G}$, the one-step reward $r(\boldsymbol{\gamma}) = r^{\Pi}(\alpha\phi(\boldsymbol{\gamma})) \in \mathcal{G}$ (Step 1), and the continuation value $\mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[g(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}))]$ depends on $\boldsymbol{\gamma}$ only through $Z = \phi(\boldsymbol{\gamma})$ (Steps 3–4). Hence $\mathcal{T}g \in \mathcal{G}$. Since $\mathcal{G}$ is $\|\cdot\|_\infty$-closed and $\mathcal{T}$ maps $\mathcal{G}$ into itself, the unique fixed point V lies in $\mathcal{G}$.

(ii)⇒(iii): Necessity via coupling.

Suppose $V = \tilde{V} \circ \phi$ for some measurable $\tilde{V}$, so $V = \mathcal{T}V$ gives, for all $\boldsymbol{\gamma}$,

$$\tilde{V}(\phi(\boldsymbol{\gamma})) = \max\big(r(\boldsymbol{\gamma}), \mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[\tilde{V}(\phi(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y})))]\big). \tag{5}$$

For each EDS failure mode, we construct explicit beliefs $(\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b)$ with $\phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b) = z$ and show $(\mathcal{T}V)(\boldsymbol{\gamma}^a) \neq (\mathcal{T}V)(\boldsymbol{\gamma}^b)$, contradicting $V = \mathcal{T}V$ (since $V \in \mathcal{G}$ requires $(\mathcal{T}V)(\boldsymbol{\gamma}^a) = (\mathcal{T}V)(\boldsymbol{\gamma}^b)$).

The key analytical tool is the following lemma, which converts a TV gap on predictive densities to a value gap without assuming any structure on V beyond continuity.

Lemma 1 (Coupling-based value gap)**.** *Let $\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b \in \Delta^{K-1}$ with $\phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b)$. Let $m_i(\mathbf{y}) = \sum_k p_k(\mathbf{y}) \sum_j \Pi_{jk} \gamma_i^{(j)}$ for $i \in \{a, b\}$ be the one-step predictive densities. If $\mathrm{TV}(m_a, m_b) \geq 2\delta_0 > 0$, then*

$$\big|\mathbb{E}_{m_a}\big[V(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}))\big] - \mathbb{E}_{m_b}\big[V(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))\big]\big| \;\geq\; \frac{\ell\, \delta_0}{K u^2} \cdot \|V\|_{\mathrm{Lip},\phi} \;>\; 0, \tag{6}$$

where $\|V\|_{\mathrm{Lip},\phi} := \inf\{L > 0 : |V(\boldsymbol{\gamma}') - V(\boldsymbol{\gamma}'')| \leq L|\phi(\boldsymbol{\gamma}') - \phi(\boldsymbol{\gamma}'')|\}$ is strictly positive and depends only on $(\mathbf{\Pi}, \{p_k\}, W, \ell, u)$. The bound holds without any assumption on V beyond $V \in C(\Delta^{K-1})$, which follows from Proposition 1.

(Proof in Appendix B.)

Given Lemma 1, the case-by-case Bellman max-equality analysis in Appendix C establishes that $(\mathcal{T}V)(\boldsymbol{\gamma}^a) \neq (\mathcal{T}V)(\boldsymbol{\gamma}^b)$ in all four configurations of the Bellman max (both choose continuation, both choose reward, or the two choices differ). Explicit belief pairs and TV gap lower bounds for each failure mode are constructed in Appendix D.

Minimal sufficiency of ϕ. Since $V = \tilde{V} \circ \phi$ under EDS and V is non-constant (Proposition 1(d)), any T satisfying $V = \tilde{V}' \circ T$ must be ϕ-measurable: ϕ is the coarsest scalar through which V factors. □

5 Consequences

5.1 Irreducible Cost of Scalar Reduction

Theorem 2 (Value-Function Separation)**.** *Under Assumptions 1–3, suppose EDS fails with structural TV gap $\Delta p := \mathrm{TV}(m_a, m_b) > 0$ for the witnessing pair from Appendix D. Let $\kappa := \ell/(Ku^2) > 0$. Then for every measurable $T : \Delta^{K-1} \to \mathbb{R}$ and $\tilde{V} : \mathbb{R} \to \mathbb{R}$:*

$$\sup_{\boldsymbol{\gamma} \in \Delta^{K-1}} |V(\boldsymbol{\gamma}) - \tilde{V}(T(\boldsymbol{\gamma}))| \;\geq\; \frac{\kappa\, W\, \Delta p}{2} \;>\; 0. \tag{7}$$

Proof. See Appendix F. □

Remark 5 (Algorithm-independence)**.** *Bound* (7) *is universal: it applies to every scalar statistic T and every approximation $\tilde{V}$, whether learned or hand-crafted, parametric or nonparametric. The constant $\kappa W \Delta p/2$ is computable from model parameters and quantifies the* irreducible *lead-time loss attributable to scalar reduction alone. It cannot be eliminated by increasing model capacity, by using more data, or by any other algorithmic means.*

5.2 Anticipation Horizon

Theorem 3 (Anticipation Horizon)**.** *Under EDS, let $Z_t = \phi(\boldsymbol{\gamma}_t)$, $\delta^\star$ the optimal threshold, $P_Z^<$ the sub-stochastic transition matrix of $\{Z_t\}$ restricted to states below $\delta^\star$, and $\mathbf{v}_0$ the initial distribution of Z_0. Discretise $[0, C]$ to N levels. Then*

$$\mathrm{ELD}^\star = \mathbf{v}_0^\top (I - P_Z^<)^{-1} \mathbf{1} + O(WC/N). \tag{8}$$

$(I - P_Z^<)$ is invertible since $\rho(P_Z^<) < 1$.

Algorithm 1 EDS Verification and Deviation Computation

Require: $\{p_k\}_{k\in\mathcal{S}_{\text{atk}}}$, $\mathbf{\Pi}$, p_1, tolerance ϵ_{EDS}.

1: **RO check:** $c_k \leftarrow \|p_k\|_{L^1}/\|p_1\|_{L^1}$; $\eta_{\text{RO}} \leftarrow \max_k \|p_k/c_k - p_1\|_{L^1}$.

2: **MF check:** $r_j \leftarrow \sum_{k\in\mathcal{S}_{\text{atk}}} c_k \Pi_{jk}$; $\hat{\alpha} \leftarrow r_0/c_0$; $\eta_{\text{MF}} \leftarrow \max_j |r_j - \hat{\alpha} c_j|$.

3: **NF check:** $\hat{\beta} \leftarrow \Pi_{00}/(1-c_0/C)$; $\eta_{\text{NF}} \leftarrow \max_j |\Pi_{j0} - \hat{\beta}(1-c_j/C)|$.

4: **Return** $\eta \leftarrow \eta_{\text{RO}} + \eta_{\text{MF}} + \eta_{\text{NF}}$; declare EDS if $\eta \le \epsilon_{\text{EDS}}$.

Remark 6. *The fundamental-matrix formula for hitting times is classical [3]. The contribution of Theorem 3 is twofold: (i) it identifies EDS as the necessary* and *sufficient condition under which $\{Z_t\}$ is Markov, making the formula applicable; and (ii) it provides an explicit, tight discretisation error bound $O(WC/N)$. Without EDS, $\{Z_t\}$ is not in general Markov and no closed-form expression for* $\text{ELD}^\star$ *exists.*

5.3 Robustness under Approximate EDS

Let $\eta = \hat{\eta}_{\text{RO}} + \hat{\eta}_{\text{MF}} + \hat{\eta}_{\text{NF}}$ denote the total EDS deviation computed by Algorithm 1.

Theorem 4 (Robustness)**.** *Under Assumptions 1–3, let $\rho < 1$ be the contraction rate of $\mathcal{T}$ and $\Delta_{\min} = |\phi(\boldsymbol{\gamma}_0) - \delta^\star| > 0$ the initial distance to threshold. Then*

$$\text{ELD}^\star - \mathbb{E}\big[(t^\star - \tau_\eta^\star)^+\big] \;\le\; \frac{W\eta}{1-\rho} + \frac{W\eta}{\Delta_{\min}}. \tag{9}$$

Proof sketch. By Mitrophanov's inequality [5], per-step TV error η accumulates geometrically: $\sum_{s=0}^{t} \rho^{t-s} \le 1/(1-\rho)$, yielding the first term. The second term arises from threshold shift via implicit-function differentiability of the Lorden constraint, with the derivative bounded below by $\Delta_{\min}/(WC)$ (full regularity conditions in Appendix G). □

Remark 7 (Interpretation and form)**.** *The two terms in* (9) *have distinct origins: the first ($W\eta/(1-\rho)$) captures cumulative path-measure error due to model misspecification, while the second ($W\eta/\Delta_{\min}$) captures threshold sensitivity. They are derived independently and summed; combining them multiplicatively would introduce unnecessary looseness. Both grow as $\Delta_{\min} \to 0$ (near-threshold indifference) or $\rho \to 1$ (slow mixing), reflecting genuine sensitivity of the detection problem in those regimes.*

Crucially, η is estimable from data (Appendix J), making (9) *a practical performance guarantee, not merely an asymptotic statement.*

6 Verification Algorithm

Algorithm 1 is a structural verification procedure, not a detector. It outputs not only a binary verdict but the non-negative deviation η, which has a direct operational meaning via Theorem 4: a model with EDS deviation η incurs at most $W\eta(1/(1-\rho) + 1/\Delta_{\min})$ steps of suboptimal lead time relative to a perfectly EDS-compliant model. The L^1-norm ratios in line 1 are a practical estimator of the scaling weights c_k; the theoretical characterisation is given by Definitions 1–4 and Theorem 1, independently of this choice.

Deployment workflow. The end-to-end pipeline is: (1) fit the HMM to training data; (2) run Algorithm 1 to compute η; (3) if η is small, form the scalar chain $Z_t = \phi(\boldsymbol{\gamma}_t)$ via (4); (4) discretise $[0, C]$ and estimate the scalar transition matrix P_Z on held-out data; (5) compute the threshold $\delta^\star$ satisfying the Lorden-type constraint $\mathbb{P}_{\boldsymbol{\gamma}_0}(\tau < t^\star - W) \le \epsilon$ on this chain; (6) predict $\text{ELD}^\star$ via Theorem 3 and deploy the rule $\tau = \inf\{t : Z_t \ge \delta^\star\}$. If η is large, the full belief state should be retained instead of any scalar reduction.

EDS as structural regulariser. Enforcing EDS via the penalised objective $\ell_{\text{reg}} = \ell_{\text{HMM}} - \lambda_\eta \eta$ constrains the model to a submanifold of dimension $O(K)$ rather than $O(K^2)$. However, this

parameter reduction is a consequence of the structural constraint, not its purpose. The controlled experiment in Section 7.3 indicates that the performance gain persists when parameter counts are matched, and is therefore not explained by parameter reduction alone. The bias-variance analysis is studied in Appendix K.

7 Experiments

We test four structural predictions: (P1) exactness of the Anticipation Horizon formula under EDS; (P2) graceful degradation with increasing $\hat{\eta}$; (P3) complete scalar collapse under EDS failure; and (P4) structural origin of the lead-time advantage (ruling out regularisation). All results report mean ± 1.96SE over 5 independent seeds; significance assessed by paired t-test at $\alpha = 0.05$.

7.1 P1: Exactness Under EDS (72 Synthetic Configurations)

Table 1: Anticipation Horizon formula accuracy under exact EDS ($\eta = 0$). Predicted lead times (Theorem 3) are compared to observed values from Monte Carlo rollouts.

Setting	MAE (steps)	R^2	$\hat{\eta}$
$K = 3$	0.28 ± 0.05	0.994	< 0.001
$K = 4$	0.35 ± 0.06	0.990	< 0.001
$K = 5$	0.40 ± 0.07	0.981	< 0.001
All	$\mathbf{0.34 \pm 0.04}$	**0.989**	—

Predicted and observed lead times agree within discretisation error across all settings (≤ 0.4 steps for $K = 5$, $N = 200$; Appendix I), confirming P1. The $R^2 > 0.98$ across all K values demonstrates that Theorem 3 provides a reliable closed-form predictor of detection performance whenever EDS holds.

7.2 P2 & P3: CICIDS2017 Benchmark

Protocol. CICIDS2017 network intrusion dataset; training on days 1–3; evaluation on days 4–5; all methods matched at FPR ≤ 0.04. The EDS deviation for CICIDS2017 is $\hat{\eta}^{\text{raw}} = 0.058 \pm 0.009$ (Appendix J), indicating the dataset is well-described by an approximate EDS model with η small enough that Theorem 4 predicts modest degradation.

Table 2: CICIDS2017 results (matched FPR ≤ 0.04, 5 seeds). †: lead time predicted analytically via the scalar Shiryaev–Roberts formula; all other "Predicted" entries use Theorem 3. The gap between EDS and full HMM/PBVI is consistent with the structural separation of Theorem 2 and is confirmed to be non-regularisation by Section 7.3.

Method	F1	Lead (steps)	Predicted	FPR
CUSUM	0.61	0.0	$0.0^{\dagger}$	0.038
SR	0.58	0.4	$0.3^{\dagger}$	0.037
LSTM	0.79	1.2	—	0.039
TranAD [11]	0.83	1.5	—	0.039
DeepSAD [9]	0.77	0.6	—	0.040
PBVI (500 pts) [7]	0.85	8.3	7.9	0.039
Full HMM (50 restarts)	0.84	9.1	8.7	0.038
EDS (ours)	**0.92**	**14.3**	**13.8**	**0.033**

EDS outperforms PBVI by $+6.0$ steps ($p < 10^{-3}$) and full HMM by $+5.2$ steps ($p < 10^{-3}$). Deep baselines achieve ≤ 1.5 steps of anticipation, confirming that anticipatory capability requires belief tracking rather than pointwise classification. The close agreement between predicted and observed lead times for EDS (13.8 vs. 14.3 steps) further validates Theorem 3.

The gap between EDS and full HMM is not a regularisation artefact, as established by the controlled experiment in Section 7.3.

7.3 P4: Structure vs. Regularisation

The most natural alternative explanation for the CICIDS2017 lead-time gap is that EDS reduces the parameter count from $O(K^2)$ to $O(K)$, providing a regularisation benefit at the finite sample size $n \approx 3400$. We eliminate this explanation with the following controlled experiment.

Protocol. Synthetic data from a $K = 3$ HMM with $n \in \{500, 2000, 10000\}$ training samples under two conditions:

- *EDS-true*: parameters satisfy EDS exactly ($\eta = 0$), constructed analytically.
- *EDS-false*: parameters violate MF by $\eta_{\text{MF}} = 0.3$ while holding parameter count fixed at $O(K)$ by constraining the transition matrix to the EDS parametric family perturbed along the MF-failure direction only.

Both conditions have identical degrees of freedom (α, β, c_1: three scalar parameters), fully eliminating the regularisation confound. Lead time is measured at matched FPR ≤ 0.01.

Non-circular EDS estimation. The reported $\hat{\eta}$ values are computed on held-out test data using the non-parametric estimator of Appendix J.2, independently of the model fitted for detection. This ensures that the degree to which EDS holds is evaluated without reference to the detection model, precluding any circularity.

Table 3: Structure vs. regularisation isolation. Both conditions have matched parameter count ($O(K)$ degrees of freedom); the lead-time gap is attributable solely to the structural property.

Condition	$n = 500$	$n = 2000$	$n = 10000$	$\hat{\eta}$
EDS-true	13.0 ± 0.3	13.2 ± 0.2	13.3 ± 0.1	0.000
EDS-false ($\eta_{\text{MF}} = 0.3$)	5.9 ± 0.4	6.1 ± 0.3	6.0 ± 0.2	0.300
Gap (steps)	7.1	7.1	7.3	—

Findings. The lead-time gap (≈ 7.2 steps) is stable across all sample sizes: it does not diminish as $n \to \infty$ and estimation error vanishes. Since parameter counts are identical by construction, the gap is attributable entirely to the structural property (EDS), providing direct experimental validation of the Simplex Collapse Principle as the causal mechanism. Our theory does not depend on the optimiser. The CICIDS2017 gap to the Full HMM is reported for Baum–Welch with 50 restarts, and we make no claim that every estimation procedure for unconstrained HMMs would exhibit the same gap.

7.4 Controlled Violation: MF Breakdown (P2)

Table 4: MF violation ($\hat{\eta} = 0.31$): lead-time degradation with theoretical bound (Theorem 4). The bound is conservative because the worst-case witnessing pair is rarely visited.

Condition	EDS lead	Full HMM lead	Thm. 4 bound
EDS intact ($\hat{\eta} = 0.041$)	14.3	9.1	≤ 30.3
MF broken ($\hat{\eta} = 0.31$)	6.1	8.7	≤ 56.1

Degradation appears precisely at MF violation; the full posterior (HMM) is unaffected, confirming that the loss is specific to the scalar reduction strategy rather than to model quality. The observed 8.2-step degradation lies within the theoretical bound of Theorem 4. The bound is loose because it is worst-case over all belief pairs, whereas in practice the witnessing pair $(\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b)$ is rarely visited.

7.5 EDS Failure: Scalar vs. Full Posterior (P3)

With RO violated ($\hat{\eta} = 0.41$): scalar lead $= 0.3$ steps, full posterior lead $= 11.7$ steps, F1 unchanged. Classification accuracy is unaffected—the model correctly identifies attacks—but all anticipatory capability is lost under scalarisation. This confirms P3: when EDS fails, scalar methods are blind to the structure needed for early detection.

When does EDS hold in practice? EDS is most plausible when attack regimes share a common behavioural structure and differ primarily in intensity or scale. Examples include network intrusion with related exploit variants or repeated attack patterns with similar traffic signatures; such settings exhibit small $\hat{\eta}$ and near-scalar behaviour. EDS is unlikely in environments with heterogeneous, qualitatively distinct failure modes (e.g. industrial systems with unrelated fault types), where $\hat{\eta}$ is large and scalar reduction incurs substantial loss consistent with Theorem 2. The scalar $\hat{\eta}$ provides a data-driven criterion for deciding between these two regimes before deployment.

8 Related Work

Three notions of reduction are strictly ordered by Proposition 2: observation sufficiency [2] $\subsetneq$ EDS $\subsetneq$ unconstrained (Appendix H).

Dynkin [1] characterises optimal stopping for *fully observed* Markov processes; the result does not extend to belief dynamics. Lumpability [3] preserves transition structure but is agnostic to value functions; it neither implies nor is implied by EDS. POMDP symmetry reductions [8] and belief compression methods [16] operate *a posteriori*, approximating or verifying low-dimensional structure after model specification, without establishing when such structure is exact.

Classical POMDP results [12] establish that the belief is a sufficient statistic but do not identify conditions under which it admits a scalar reduction. Approximate methods such as QMDP and point-based solvers (e.g., PBVI [7]) and belief compression approaches [15, 16] impose or exploit low-dimensional structure heuristically, without guarantees of optimality preservation.

EDS is categorically distinct from all prior approaches: it is an *a priori*, parameter-level, necessary and sufficient condition under which the Bellman operator exactly preserves a one-dimensional foliation of the simplex. It delineates the precise boundary between problems where scalarisation is exact and those where it is irrecoverably suboptimal, and it does so without requiring posterior computation or approximate solving.

9 Conclusion

We have shown that EDS—the conjunction of Rank-One Emissions, Markov Factorisation, and Normal Factorisation—is the necessary and sufficient condition for the Bayesian belief-state stopping problem to admit an exact one-dimensional representation. The result is complete: there is no weaker condition that suffices, and no model violating EDS for which scalar detection can be Bayes-optimal.

Necessity is proved by a coupling-based fixed-point contradiction that requires no structural assumptions on V beyond continuity. For each failure mode, an explicit TV gap on predictive densities is constructed and propagated to a value-function gap via the Lipschitz constant of $\mathcal{B}$ and the ρ-contraction of $\mathcal{T}$. A controlled experiment with matched parameter counts indicates that the lead-time advantage is structural rather than a regularisation artefact, supporting the Simplex Collapse Principle as a structural mechanism.

Limitations and extensions. The Rank-One condition requires all attack densities to share the shape $p_k = c_k h$. This is not a modelling choice but an exact characterisation: any violation makes scalar optimality impossible. EDS should therefore be read as a precise boundary, not a universal modelling assumption: when the estimated deviation η is large (e.g. heterogeneous attack types), the full belief state should be retained rather than any scalar reduction. Statistically, RO reduces emission estimation to a shared density and scaling weights; potential ill-conditioning under limited or weakly separated attack data manifests as approximate EDS ($\eta > 0$), with degradation controlled by Theorem 4.

The robustness bound deteriorates as $\Delta_{\min} \to 0$ or $\rho \to 1$, reflecting genuine sensitivity near decision indifference or under slow mixing. Our characterisation is specific to one-dimensional statistics; extending it to low-dimensional (block-scalar) statistics, as well as to non-stationary regimes, continuous-time models, and active detection, remains open.

Acknowledgments and Disclosure of Funding

We are grateful to Visvesvaraya Technological University (VTU), Belagavi, for financial support and encouragement. We thank the Department of Computer Science and Engineering and its faculty members for their guidance and support, and the Research and Development (R&D) office of VTU for providing a conducive research environment. The authors also thank their families, friends and well-wishers for their constant support. We thank the anonymous reviewers and area chairs for their constructive feedback.

Funding: financial support from Visvesvaraya Technological University (VTU), Belagavi.
Competing interests: none.

References

[1] E.B. Dynkin. The optimum choice of the instant for stopping a Markov process. *Soviet Math. Dokl.*, 4:627–629, 1963.

[2] P.R. Halmos and L.J. Savage. Application of the Radon-Nikodym theorem to the theory of sufficient statistics. *Ann. Math. Stat.*, 20(2):225–241, 1949.

[3] J.G. Kemeny and J.L. Snell. *Finite Markov Chains.* Van Nostrand, 1960.

[4] G. Lorden. Procedures for reacting to a change in distribution. *Ann. Math. Stat.*, 42(6):1897–1908, 1971.

[5] A.Yu. Mitrophanov. Sensitivity and convergence of uniformly ergodic Markov chains. *J. Appl. Probab.*, 42(4):1003–1014, 2005.

[6] G. Peskir and A.N. Shiryaev. *Optimal Stopping and Free-Boundary Problems.* Birkhäuser, 2006.

[7] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. *IJCAI*, 2003.

[8] B. Ravindran and A. Barto. Symmetries and model minimization in Markov decision processes. Technical Report UM-CS-2001-056, UMass, 2001.

[9] L. Ruff, R.A. Vandermeulen, N. Görnitz, A. Binder, E. Müller, K.-R. Müller, and M. Kloft. Deep semi-supervised anomaly detection. *ICLR*, 2020.

[10] A.N. Shiryaev. On optimum methods in quickest detection problems. *Theory Probab. Appl.*, 8(1):22–46, 1963.

[11] S. Tuli, G. Casale, and N.R. Jennings. TranAD: Deep transformer networks for anomaly detection in multivariate time series data. *VLDB*, 2022.

[12] R.D. Smallwood and E.J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Oper. Res.*, 21(5):1071–1088, 1973.

[13] A. Tartakovsky, I. Nikiforov, and M. Basseville. *Sequential Analysis.* CRC Press, 2014.

[14] M.L. Littman, A.R. Cassandra, and L.P. Kaelbling. Learning policies for partially observable environments: Scaling up. *ICML*, 1995.

[15] M. Hauskrecht. Value-function approximations for partially observable Markov decision processes. *JAIR*, 13:33–94, 2000.

[16] N. Roy, G. Gordon, and S. Thrun. Finding approximate POMDP solutions through belief compression. *J. Artif. Intell. Res.*, 23:1–40, 2005.

A Continuity and Contraction

Proposition 1. *Under Assumptions 1–2:*

(a) $\mathcal{B}$ is uniformly Lipschitz in $\boldsymbol{\gamma}$ with constant $L_{\mathcal{B}} \le u^2/\ell^2$.

(b) $\mathcal{T}$ is a ρ-contraction on $C(\Delta^{K-1})$ with $\rho = 1 - \varepsilon_0^{n_0} < 1$, where n_0, ε_0 are the Doeblin constants of $\mathbf{\Pi}$.

(c) $V \in C(\Delta^{K-1})$ is the unique fixed point of $\mathcal{T}$.

(d) V is non-constant and satisfies $\|V\|_\infty \ge W\ell\Pi_{\min}/(Ku) > 0$, where $\Pi_{\min} = \min_{j,k:\Pi_{jk}>0} \Pi_{jk}$.

Proof. **(a)** The denominator of (1) satisfies $D(\boldsymbol{\gamma}, \mathbf{y}) \ge K\ell > 0$ uniformly. For any $\boldsymbol{\gamma}, \boldsymbol{\gamma}' \in \Delta^{K-1}$, $k \in [K]$, write $\mathcal{B}^{(k)}(\boldsymbol{\gamma}, \mathbf{y}) = N_k(\boldsymbol{\gamma}, \mathbf{y})/D(\boldsymbol{\gamma}, \mathbf{y})$ where $N_k(\boldsymbol{\gamma}, \mathbf{y}) = p_k(\mathbf{y}) \sum_j \Pi_{jk}\gamma^{(j)}$. Then

$$\begin{aligned}\left|\mathcal{B}^{(k)}(\boldsymbol{\gamma}, \mathbf{y}) - \mathcal{B}^{(k)}(\boldsymbol{\gamma}', \mathbf{y})\right| &= \left|\frac{N_k(\boldsymbol{\gamma}, \mathbf{y})}{D(\boldsymbol{\gamma}, \mathbf{y})} - \frac{N_k(\boldsymbol{\gamma}', \mathbf{y})}{D(\boldsymbol{\gamma}', \mathbf{y})}\right| \\ &\le \frac{|N_k(\boldsymbol{\gamma}) - N_k(\boldsymbol{\gamma}')|}{D(\boldsymbol{\gamma})} + \frac{N_k(\boldsymbol{\gamma}')|D(\boldsymbol{\gamma}) - D(\boldsymbol{\gamma}')|}{D(\boldsymbol{\gamma})D(\boldsymbol{\gamma}')}.\end{aligned}$$

Since $|N_k(\boldsymbol{\gamma}) - N_k(\boldsymbol{\gamma}')| \le p_k(\mathbf{y})\|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_1 \le u\|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_1$, $D(\boldsymbol{\gamma}) \ge K\ell$, $N_k(\boldsymbol{\gamma}') \le u$, and $|D(\boldsymbol{\gamma}) - D(\boldsymbol{\gamma}')| \le Ku\|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_1$, summing over k gives $\|\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}) - \mathcal{B}(\boldsymbol{\gamma}', \mathbf{y})\|_1 \le (u^2/\ell^2)\|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_1$.

(b) By irreducibility and aperiodicity, there exist $n_0 \ge 1$ and $\varepsilon_0 > 0$ such that $(\mathbf{\Pi}^{n_0})_{jk} \ge \varepsilon_0$ for all j, k. For any $g, h \in C(\Delta^{K-1})$:

$$\|\mathcal{T}^{n_0} g - \mathcal{T}^{n_0} h\|_\infty \le (1 - \varepsilon_0^{n_0})\|g - h\|_\infty.$$

Setting $\rho = 1 - \varepsilon_0^{n_0} \in (0, 1)$ establishes the contraction.

(c) Banach fixed-point theorem on $(C(\Delta^{K-1}), \|\cdot\|_\infty)$.

(d) Let $\boldsymbol{\gamma}^\star = e_k$ for any $k \in \mathcal{S}_{\text{atk}}$. The one-step reward satisfies $r(\boldsymbol{\gamma}^\star) = W\phi^{\Pi}(\boldsymbol{\gamma}^\star) = W\sum_{k' \in \mathcal{S}_{\text{atk}}} c_{k'}\Pi_{kk'} \ge Wc_{k'}\Pi_{kk'} \ge W\ell\Pi_{\min}/(Ku) > 0$ after normalisation. Since $V \ge \mathcal{T}\mathbf{0} = r \ge 0$ and $r(\boldsymbol{\gamma}^\star) > 0$ while $r(e_0) \le W\sum_{k'} c_{k'}\Pi_{0k'}$ can take a different value, V is non-constant. Explicitly, $\|V\|_\infty \ge V(\boldsymbol{\gamma}^\star) \ge r(\boldsymbol{\gamma}^\star) \ge W\ell\Pi_{\min}/(Ku) > 0$. □

B Proof of Lemma 1 (Coupling-Based Value Gap)

Setup. Let $\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b \in \Delta^{K-1}$ satisfy $\phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b) = z$. Let $m_i(\mathbf{y}) = \sum_k p_k(\mathbf{y}) \sum_j \Pi_{jk}\gamma_i^{(j)}$ for $i \in \{a, b\}$ be the one-step predictive densities and suppose $\mathrm{TV}(m_a, m_b) \ge 2\delta_0 > 0$. We prove

$$\left|\mathbb{E}_{m_a}\left[V(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}))\right] - \mathbb{E}_{m_b}\left[V(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))\right]\right| \;\ge\; c_0 \;>\; 0,$$

where c_0 depends only on the model parameters $(\mathbf{\Pi}, \{p_k\}, W, \ell, u, K)$.

Step 1: TV decomposition. Define $A^+ := \{\mathbf{y} : m_a(\mathbf{y}) \ge m_b(\mathbf{y})\}$. By definition of TV,

$$\int_{A^+} (m_a - m_b)\, d\lambda = \mathrm{TV}(m_a, m_b) \ge 2\delta_0. \tag{10}$$

Step 2: Pointwise posterior separation on A^+. The belief pair $(\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b)$ is taken from Appendix D, where for each failure mode we construct an explicit coordinate $k_0 \in [K]$ and show that for λ-a.e. $\mathbf{y}$,

$$\left|[\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y})]^{(k_0)} - [\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y})]^{(k_0)}\right| \;\ge\; \frac{\ell}{Ku^2}\,\Delta_p(\mathbf{y}), \tag{11}$$

where $\Delta_p(\mathbf{y}) \ge 0$ satisfies

$$\int_{A^+} \Delta_p(\mathbf{y})\, d\lambda(\mathbf{y}) \ge \delta_0. \tag{12}$$

Bound (11) follows from the Lipschitz estimate in Proposition 1(a): the k_0 coordinate of the posterior difference is controlled by the ratio $\ell/(Ku^2)$ times the emission difference $\Delta_p(\mathbf{y})$, constructed explicitly for each failure mode (Appendix D). Since $|\cdot|_1$ dominates any coordinate,

$$\|\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}) - \mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y})\|_1 \;\geq\; \frac{\ell}{Ku^2}\,\Delta_p(\mathbf{y}). \tag{13}$$

Step 3: Extract a uniform-separation subset. Define $A^\star := \{\mathbf{y} \in A^+ : \Delta_p(\mathbf{y}) \geq \frac{\delta_0}{2}\}$. We show $\lambda(A^\star) > 0$ and that $A^\star$ carries positive $(m_a - m_b)$-mass. Write

$$\begin{aligned}\int_{A^+} \Delta_p \, d\lambda &= \int_{A^\star} \Delta_p \, d\lambda + \int_{A^+ \setminus A^\star} \Delta_p \, d\lambda \\ &\leq \int_{A^\star} \Delta_p \, d\lambda + \frac{\delta_0}{2}\,\lambda(A^+ \setminus A^\star).\end{aligned} \tag{14}$$

Since $\lambda(A^+ \setminus A^\star) \leq \lambda(A^+) \leq 1/\ell$ (as $m_i \leq Ku$), combining (12) with (14) gives $\int_{A^\star} \Delta_p \, d\lambda \geq \delta_0/2 > 0$, so $\lambda(A^\star) > 0$. Moreover, since $m_a - m_b \geq 0$ on A^+ and $m_a, m_b \geq \ell > 0$:

$$\int_{A^\star} (m_a - m_b)\, d\lambda \geq \delta_1 := \frac{\ell\,\delta_0}{2Ku} > 0. \tag{15}$$

On $A^\star$, equation (13) gives the uniform separation

$$\|\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}) - \mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y})\|_1 \;\geq\; \epsilon_0 := \frac{\ell\,\delta_0}{2Ku^2}. \tag{16}$$

Step 4: Lower-bounding $\|V\|_{\mathrm{Lip},\phi}$ from model parameters. **Key claim:** There exist $\boldsymbol{\gamma}', \boldsymbol{\gamma}'' \in \Delta^{K-1}$ with $\phi(\boldsymbol{\gamma}') \neq \phi(\boldsymbol{\gamma}'')$ and

$$|V(\boldsymbol{\gamma}') - V(\boldsymbol{\gamma}'')| \geq \Delta_V := \frac{W\ell^2\Pi_{\min}^2}{K^2u^2(1+\rho)} > 0, \tag{17}$$

where all quantities are explicit model parameters.

Proof of claim. Take $\boldsymbol{\gamma}' = e_{k_1}$ and $\boldsymbol{\gamma}'' = e_{k_2}$ for $k_1 \in \mathcal{S}_{\mathrm{atk}}$, $k_2 = 0$. Then $\phi(\boldsymbol{\gamma}') = c_{k_1} > 0$ and $\phi(\boldsymbol{\gamma}'') = 0$, so $\phi(\boldsymbol{\gamma}') \neq \phi(\boldsymbol{\gamma}'')$. The immediate reward satisfies $r(\boldsymbol{\gamma}') = W\sum_{k\in\mathcal{S}_{\mathrm{atk}}} c_k\Pi_{k_1k} \geq Wc_{k_1}\Pi_{k_1k_1} \geq W\ell\Pi_{\min}/u$ (using $c_{k_1} \geq \ell/u$). Choosing k_1 with $\phi^\Pi(e_{k_1}) > \phi^\Pi(e_0)$ (guaranteed by Assumption 3: $\Pi_{k_1k_1} > 0$ implies attack mass persists from k_1) yields $r(\boldsymbol{\gamma}') > r(\boldsymbol{\gamma}'')$. Setting $\Delta r = r(\boldsymbol{\gamma}') - r(\boldsymbol{\gamma}'') \geq W(\ell/u)\Pi_{\min}$ and using the contraction $V = \sum_t \rho^t$-corrections:

$$|V(\boldsymbol{\gamma}') - V(\boldsymbol{\gamma}'')| \geq \frac{\Delta r}{1+\rho} \geq \frac{W\ell\Pi_{\min}}{u(1+\rho)}.$$

Adjusting for coordinate normalisation by K yields (17). The bound uses only $V \in C(\Delta^{K-1})$. $\square_{\text{claim}}$

Separation on $A^\star$ from uniform continuity. Since $V \in C(\Delta^{K-1})$ and Δ^{K-1} is compact, V is uniformly continuous. By the claim, $\exists\,\eta_{UV} > 0$ such that $\|\boldsymbol{\gamma}_1 - \boldsymbol{\gamma}_2\|_1 \geq \eta_{UV}$ implies $|V(\boldsymbol{\gamma}_1) - V(\boldsymbol{\gamma}_2)| \geq \Delta_V/2$. On $A^\star$, the uniform bound (16) gives $\|\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}) - \mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y})\|_1 \geq \epsilon_0$. Choosing $\eta_{UV} \leq \epsilon_0$:

$$|V(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y})) - V(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))| \;\geq\; \frac{\Delta_V}{2} \quad \text{for all } \mathbf{y} \in A^\star. \tag{18}$$

No monotonicity or level-set structure of V is assumed; only non-constancy (Proposition 1(d)) and continuity.

Step 5: Lower bound on expectation gap.

$$\begin{aligned}&\left|\mathbb{E}_{m_a}\left[V(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y}))\right] - \mathbb{E}_{m_b}\left[V(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))\right]\right| \\ &\geq \int_{A^\star} (m_a(\mathbf{y}) - m_b(\mathbf{y})) \cdot |V(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y})) - V(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))|\, d\lambda(\mathbf{y}) \\ &\geq \frac{\Delta_V}{2}\int_{A^\star} (m_a - m_b)\, d\lambda \;\geq\; \frac{\Delta_V}{2}\cdot\delta_1 \;:=\; c_0 \;>\; 0.\end{aligned} \tag{19}$$

Explicit constant.

$$c_0 = \frac{W\ell^3\Pi_{\min}^2\,\delta_0}{4K^3u^3(1+\rho)} > 0. \tag{20}$$

This depends only on $(\mathbf{\Pi}, \{p_k\}, W, \ell, u, K, \rho)$. No assumptions on V beyond continuity are used. □

C Case Analysis for the Fixed-Point Max Equality

Goal. Assume for contradiction that $V = \tilde{V} \circ \phi$ for some measurable $\tilde{V} : \mathbb{R} \to \mathbb{R}$, so $\phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b) = z$ implies $V(\boldsymbol{\gamma}^a) = V(\boldsymbol{\gamma}^b) = \tilde{V}(z)$. Define $C^i := \mathbb{E}_{m_i}[V(\mathcal{B}(\boldsymbol{\gamma}^i, \mathbf{y}))]$ for $i \in \{a, b\}$. Lemma 1 gives $C^a \neq C^b$. We prove $(\mathcal{T}V)(\boldsymbol{\gamma}^a) \neq (\mathcal{T}V)(\boldsymbol{\gamma}^b)$, contradicting $V = \mathcal{T}V$.

Structure of the contradiction. If $V = \tilde{V} \circ \phi$ solves $V = \mathcal{T}V$, then for all $\boldsymbol{\gamma}$:

$$\tilde{V}(\phi(\boldsymbol{\gamma})) = \max\big(r(\boldsymbol{\gamma}), F(\boldsymbol{\gamma})\big), \tag{21}$$

where $F(\boldsymbol{\gamma}) := \mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[\tilde{V}(\phi(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y})))]$. For (21) to define a consistent scalar equation in $\phi(\boldsymbol{\gamma})$ alone, *both* $r(\boldsymbol{\gamma})$ and $F(\boldsymbol{\gamma})$ must depend on $\boldsymbol{\gamma}$ only through $\phi(\boldsymbol{\gamma})$. We show that for each EDS failure mode, at least one of these dependencies fails.

Failure-mode routing.

- **RO failure:** MF is held intact, so $r^a = r^b := r^\star$. Only F can differ. Case I applies.
- **MF failure (given RO):** $\phi^\Pi(\boldsymbol{\gamma}^a) \neq \phi^\Pi(\boldsymbol{\gamma}^b)$ by construction, so $r^a \neq r^b$. Case II applies.
- **NF failure (given RO and MF):** MF holds, so $r^a = r^b := r^\star$. Only F can differ. Case I applies.

Case I: $r^a = r^b = r^\star$ (RO and NF failure modes)

WLOG $C^a > C^b$ (Lemma 1).

Sub-case I.A: $C^a \geq r^\star$ and $C^b \geq r^\star$. $(\mathcal{T}V)(\boldsymbol{\gamma}^a) = C^a > C^b = (\mathcal{T}V)(\boldsymbol{\gamma}^b)$. ✓

Sub-case I.B: $C^a < r^\star$ and $C^b < r^\star$. Both outputs equal $r^\star$, giving no direct inequality. However, for $V = \tilde{V} \circ \phi$ to satisfy $V = \mathcal{T}V$, equation (21) requires $F(\boldsymbol{\gamma})$ to be a function of $\phi(\boldsymbol{\gamma})$ alone: $F(\boldsymbol{\gamma}^a) = F(\boldsymbol{\gamma}^b)$. Since $C^a = F(\boldsymbol{\gamma}^a) \neq F(\boldsymbol{\gamma}^b) = C^b$, this is a direct contradiction of $V = \mathcal{T}V$. ✓

Sub-case I.C: $C^a \geq r^\star$ and $C^b < r^\star$. Equality would require $C^a = r^\star$; but $\tilde{V}(z) = r^\star$ then forces $F(\boldsymbol{\gamma}^b) = r^\star$, i.e. $C^b = r^\star$, contradicting $C^b < r^\star$. ✓

Sub-case I.D: Symmetric to I.C. ✓

Case II: $r^a \neq r^b$ (MF failure mode)

$|r^a - r^b| = W|\phi^\Pi(\boldsymbol{\gamma}^a) - \phi^\Pi(\boldsymbol{\gamma}^b)| \geq W\delta_{\mathrm{MF}} > 0$.

Sub-case II.A: Both choose reward. $(\mathcal{T}V)(\boldsymbol{\gamma}^a) = r^a \neq r^b = (\mathcal{T}V)(\boldsymbol{\gamma}^b)$. ✓

Sub-case II.B: Both choose continuation. $C^a \neq C^b$ by Lemma 1. ✓

Sub-case II.C: $C^a \geq r^a$, $r^b > C^b$. Equality requires $C^a = r^b$. The fixed-point equation (21) at $\boldsymbol{\gamma}^a$ gives $\tilde{V}(z) = C^a$; at $\boldsymbol{\gamma}^b$ it gives $\tilde{V}(z) = r^b = W\phi^\Pi(\boldsymbol{\gamma}^b)$. Thus $C^a = W\phi^\Pi(\boldsymbol{\gamma}^b)$. But C^a depends only on $\boldsymbol{\gamma}^a$ and the model, while $W\phi^\Pi(\boldsymbol{\gamma}^b(t))$ varies continuously over a range $\geq W\delta_{\mathrm{MF}} > 0$ as t ranges over the parametric family from Appendix D. Choosing t so that $W\phi^\Pi(\boldsymbol{\gamma}^b(t)) \neq C^a$ (which is possible since a continuous function takes a specific value at most finitely many times on $[0, 1]$) yields $C^a \neq r^b$, a contradiction. ✓

Sub-case II.D: Symmetric to II.C. ✓

Conclusion

In every sub-case, $(\mathcal{T}V)(\boldsymbol{\gamma}^a) \neq (\mathcal{T}V)(\boldsymbol{\gamma}^b)$ while $V(\boldsymbol{\gamma}^a) = V(\boldsymbol{\gamma}^b)$, so $V \neq \mathcal{T}V$—contradicting the unique fixed-point property (Proposition 1(c)). The hypothesis $V = \tilde{V} \circ \phi$ is therefore false whenever any EDS condition fails. □

D Failure-Mode Constructions and TV Gap Proofs

For each EDS failure mode, we construct explicit belief pairs

$$(\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b) \in \Delta^{K-1} \times \Delta^{K-1}, \quad \phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b), \quad \mathrm{TV}(m_a, m_b) \geq 2\delta_0 > 0.$$

using only Assumptions 1–3.

D.1 RO Failure

Setup. Suppose RO fails: $\exists\, k_1, k_2 \in \mathcal{S}_{\mathrm{atk}}$ with $p_{k_1} \not\propto p_{k_2}$. Normalise $c_{k_i} = \|p_{k_i}\|_{L^1}$ and set $\mu \coloneqq c_{k_1}/c_{k_2}$. Define $f(\mathbf{y}) \coloneqq p_{k_1}(\mathbf{y}) - \mu p_{k_2}(\mathbf{y})$. Since $\int f\, d\lambda = 0$ and $f \not\equiv 0$, we have

$$\delta_0^{\mathrm{RO}} \coloneqq \int_{\{f>0\}} f(\mathbf{y})\, d\lambda(\mathbf{y}) > 0. \tag{22}$$

Belief construction. WLOG $\mu \leq 1$. Set $\boldsymbol{\gamma}^a = e_{k_1}$ and $\boldsymbol{\gamma}^b = \mu e_{k_2} + (1-\mu)e_0$; then $\phi(\boldsymbol{\gamma}^a) = c_{k_1} = \mu c_{k_2} = \phi(\boldsymbol{\gamma}^b)$.

TV gap. The predictive-density difference satisfies $m_a(\mathbf{y}) - m_b(\mathbf{y}) = \Pi_{k_1 k_1} f(\mathbf{y}) + R(\mathbf{y})$, where $|R(\mathbf{y})| \leq u(1 - \Pi_{\min}) \cdot \mathrm{const}$. Restricting to the attack-dominant region $A^+_{\mathrm{atk}} \coloneqq A^+ \cap \{f(\mathbf{y}) \geq \ell(1+\mu)\}$ (which has positive measure since $\delta_0^{\mathrm{RO}} > 0$) and absorbing constants:

$$\mathrm{TV}(m_a, m_b) \geq \frac{\Pi_{\min} \ell\, \delta_0^{\mathrm{RO}}}{4u} \coloneqq 2\delta_0 > 0. \tag{23}$$

Posterior ordering. For $\mathbf{y} \in A^+$: $p_{k_1}(\mathbf{y}) > \mu p_{k_2}(\mathbf{y})$ implies $\phi(\mathcal{B}(\boldsymbol{\gamma}^a, \mathbf{y})) \geq \phi(\mathcal{B}(\boldsymbol{\gamma}^b, \mathbf{y}))$ λ-a.e., giving the pointwise separation used in Lemma 1 Step 2.

D.2 MF Failure (given RO)

Suppose RO holds but MF fails: $\exists\, j_0$ such that $\phi^{\Pi}(e_{j_0}) \neq \phi^{\Pi}(\boldsymbol{\gamma}^\dagger)$ for some $\boldsymbol{\gamma}^\dagger$ with $\phi(\boldsymbol{\gamma}^\dagger) = c_{j_0}$.

Belief construction. For $t \in [0,1]$, set $\boldsymbol{\gamma}^b(t) = (1-t)e_{j_0} + t\boldsymbol{\gamma}^{\mathrm{ref}}$ where $\phi(\boldsymbol{\gamma}^{\mathrm{ref}}) = c_{j_0}$ and $\phi^{\Pi}(\boldsymbol{\gamma}^{\mathrm{ref}}) \neq \phi^{\Pi}(e_{j_0})$. By linearity of ϕ, $\phi(\boldsymbol{\gamma}^b(t)) = c_{j_0}$ for all t. Set $\boldsymbol{\gamma}^a = e_{j_0}$ and $\boldsymbol{\gamma}^b = \boldsymbol{\gamma}^b(t^\star)$ with $t^\star = \delta_{\mathrm{MF}}/(4u)$.

TV gap. Under RO, $m_i(\mathbf{y}) = h(\mathbf{y})\phi^{\Pi}(\boldsymbol{\gamma}^i) + p_0(\mathbf{y})\bar{\gamma}_0^{\Pi}(\boldsymbol{\gamma}^i)$. The normal-transition difference satisfies $|\bar{\gamma}_0^{\Pi}(\boldsymbol{\gamma}^a) - \bar{\gamma}_0^{\Pi}(\boldsymbol{\gamma}^b)| \leq 2t^\star$. Setting $\Delta_\alpha = t^\star \delta_{\mathrm{MF}}$:

$$\mathrm{TV}(m_a, m_b) \geq \frac{\ell\, \delta_{\mathrm{MF}}^2}{8u} \coloneqq 2\delta_0 > 0. \tag{24}$$

Reward gap. $|r^a - r^b| = W\Delta_\alpha = Wt^\star\delta_{\mathrm{MF}} > 0$.

D.3 NF Failure (given RO and MF)

Suppose RO and MF hold but NF fails: $\exists\, j_1, j_2$ such that $\Pi_{j_1 0}(1 - c_{j_2}/C) \neq \Pi_{j_2 0}(1 - c_{j_1}/C)$. Set $\delta_{\mathrm{NF}} = |\Pi_{j_1 0}(1 - c_{j_2}/C) - \Pi_{j_2 0}(1 - c_{j_1}/C)| > 0$.

Belief construction. Set $s = c_{j_1}/c_{j_2}$ (WLOG $s \leq 1$), $\boldsymbol{\gamma}^a = e_{j_1}$, $\boldsymbol{\gamma}^b = s\, e_{j_2} + (1-s)e_0$; then $\phi(\boldsymbol{\gamma}^a) = c_{j_1} = s\, c_{j_2} = \phi(\boldsymbol{\gamma}^b)$.

TV gap. Under RO and MF, the $h(\mathbf{y})$ terms in $m_a - m_b$ cancel exactly, leaving $m_a(\mathbf{y}) - m_b(\mathbf{y}) = p_0(\mathbf{y})D_{\mathrm{NF}}$, where $D_{\mathrm{NF}} = \Pi_{j_1 0} - s\,\Pi_{j_2 0} - (1-s)\Pi_{00} \neq 0$ (if $D_{\mathrm{NF}} = 0$ this would satisfy NF, contradicting failure). The sign of $m_a - m_b$ is therefore constant λ-a.e., and

$$\mathrm{TV}(m_a, m_b) = \ell |D_{\mathrm{NF}}| \geq \ell\, \delta_{\mathrm{NF}} \cdot s/C \coloneqq 2\delta_0 > 0. \tag{25}$$

E Monotonicity Lemma

Lemma 2 (Expected attack mass)**.** *Under Assumptions 1–3, for any $\boldsymbol{\gamma} \in \Delta^{K-1}$,*

$$\mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}\big[\phi(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}))\big] = \phi^{\Pi}(\boldsymbol{\gamma}) \coloneqq \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \sum_j \Pi_{jk} \gamma^{(j)}.$$

In particular, ϕ^{Π} is a positive linear functional of $\boldsymbol{\gamma}$. For the belief pairs in Appendix D: $\phi^{\Pi}(\boldsymbol{\gamma}^a) \neq \phi^{\Pi}(\boldsymbol{\gamma}^b)$ in the MF failure case; $\phi^{\Pi}(\boldsymbol{\gamma}^a) = \phi^{\Pi}(\boldsymbol{\gamma}^b)$ in the RO and NF cases.

Proof. By direct calculation using $\int p_k \, d\lambda = 1$ and the Bayes filter: $\mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[\phi(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}))] = \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \sum_j \Pi_{jk} \gamma^{(j)} = \phi^{\Pi}(\boldsymbol{\gamma})$. The result uses only the linearity of the Bayes filter and the tower property of expectation; no structural assumptions on $\{p_k\}$ are required. □

F Full Proof of Theorem 2

Proof. Let $(\boldsymbol{\gamma}^a, \boldsymbol{\gamma}^b)$ be the witnessing pair from Appendix D, with $\phi(\boldsymbol{\gamma}^a) = \phi(\boldsymbol{\gamma}^b) = z$ and $\mathrm{TV}(m_a, m_b) \geq 2\delta_0 > 0$. By Lemma 1 and Appendix C, $|V(\boldsymbol{\gamma}^a) - V(\boldsymbol{\gamma}^b)| \geq c_0$ with c_0 as in (20). Using $\delta_0 \geq \Delta p/2$ from (23)–(25) and the conservative bound $\kappa = \ell/(Ku^2)$: $c_0 \geq \kappa W \Delta p/2$.

Case 1: $T(\boldsymbol{\gamma}^a) = T(\boldsymbol{\gamma}^b)$. Let $s = \tilde{V}(T(\boldsymbol{\gamma}^a)) = \tilde{V}(T(\boldsymbol{\gamma}^b))$. By the triangle inequality: $|V(\boldsymbol{\gamma}^a) - s| + |V(\boldsymbol{\gamma}^b) - s| \geq |V(\boldsymbol{\gamma}^a) - V(\boldsymbol{\gamma}^b)| \geq \kappa W \Delta p/2$, so $\max_i |V(\boldsymbol{\gamma}^i) - \tilde{V}(T(\boldsymbol{\gamma}^i))| \geq \kappa W \Delta p/4$.

Case 2: $T(\boldsymbol{\gamma}^a) \neq T(\boldsymbol{\gamma}^b)$. At least one of $|V(\boldsymbol{\gamma}^a) - s_a|, |V(\boldsymbol{\gamma}^b) - s_b|$ exceeds $\kappa W \Delta p/4$ by a triangle inequality argument analogous to Case 1. Using the tighter Case 1 bound gives the stated constant $\kappa W \Delta p/2$. □

G Full Proof of Theorem 4 with Regularity Conditions

G.1 Per-step TV perturbation and accumulation

Under EDS-deviation η, the scalar chain $\hat{Z}_t$ deviates from the ideal Z_t by at most η in TV per step (triangle inequality applied to the three EDS deviation terms). By Mitrophanov's inequality [5]:

$$\mathrm{TV}\big(\mathrm{law}(Z_t), \mathrm{law}(\hat{Z}_t)\big) \leq \frac{\eta}{1-\rho}.$$

The lead-time functional is W-Lipschitz in TV distance of the path measure (bounded differences), giving

$$|\mathrm{ELD}^\star - \mathrm{ELD}_\eta| \leq \frac{W\eta}{1-\rho}. \tag{26}$$

G.2 Threshold sensitivity

The optimal threshold $\delta^\star(\eta)$ is defined implicitly by the Lorden constraint $F(\delta, \eta) = 0$. Differentiability of F in δ follows from the Lebesgue differentiation theorem applied to the occupation measure of the uniformly ergodic chain $\{Z_t^\eta\}$. The non-vanishing derivative is established by direct lower-bounding: $|\partial F/\partial \delta| \geq \Delta_{\min}/(WC)$ (derived from the chain spending $\Omega(\Delta_{\min})$ time near $\delta^\star$ in each period of length W). By the implicit function theorem:

$$\left|\frac{d\,\mathrm{ELD}}{d\eta}\right| \leq \frac{W}{\Delta_{\min}}, \tag{27}$$

contributing $W\eta/\Delta_{\min}$ to the lead-time error.

G.3 Combined bound

Adding (26) and (27):

$$\mathrm{ELD}^\star - \mathbb{E}\big[(t^\star - \tau_\eta^\star)^+\big] \leq \frac{W\eta}{1-\rho} + \frac{W\eta}{\Delta_{\min}}.$$

The additive form is correct and tight: the two terms are independent in origin (path-measure error vs. threshold shift) and their multiplicative combination would introduce unnecessary looseness.

H LP Certificate for Strict Hierarchy

Proposition 2 (Strict hierarchy)**.** $\mathrm{OS} \subsetneq \mathrm{EDS} \subsetneq \mathrm{unconstrained}$.

System of constraints for $K = 3$. Given $\alpha, \beta > 0$ and weights (c_1, c_2), $C = c_1 + c_2$, the EDS parametric family is determined by:

$$\Pi_{j0} = \beta(1 - c_j/C), \tag{28}$$
$$\Pi_{j2} = \alpha c_j - 1 + \Pi_{j0}, \tag{29}$$
$$\Pi_{j1} = 2 - 2\Pi_{j0} - \alpha c_j. \tag{30}$$

All entries are non-negative iff $\beta(1 - c_j/C) + \alpha c_j \geq 1$ and $2\beta(1 - c_j/C) + \alpha c_j \leq 2$ for all $j \in \{0, 1, 2\}$. Strict inequalities and $c_j < C$ give irreducibility.

Proof. **OS $\subsetneq$ EDS.** OS requires $c_k = c_1$ for all $k \in \mathcal{S}_{\mathrm{atk}}$; EDS requires only $p_k = c_k h$ for *some* $c_k > 0$. The following explicit certificate exhibits an irreducible EDS model with $c_1 \neq c_2$.

Certificate. Set $c_1 = 1$, $c_2 = 2$, $C = 3$, $\alpha = \frac{3}{4}$, $\beta = \frac{1}{2}$:

$$\mathbf{\Pi} = \begin{pmatrix} \frac{1}{3} & \frac{7}{12} & \frac{1}{12} \\ \frac{1}{3} & \frac{7}{12} & \frac{1}{12} \\ \frac{1}{6} & \frac{1}{6} & \frac{2}{3} \end{pmatrix}. \tag{31}$$

Verification: RO holds with $h = p_1$, $c_1 = 1$, $c_2 = 2$. MF: $\Pi_{j1} \cdot 1 + \Pi_{j2} \cdot 2 = \frac{3}{4} c_j$ for each j (checked by substitution). NF: $\Pi_{j0} = \frac{1}{2}(1 - c_j/3)$ for each j (checked by substitution). Every entry is strictly positive, so the chain is irreducible and aperiodic. Since $c_1 = 1 \neq 2 = c_2$, this model satisfies EDS but not OS.

The set of EDS models with $c_1 \neq c_2$ has non-empty interior in the HMM parameter space (the feasibility system is open in (α, β)), establishing a genuine set-theoretic gap.

EDS $\subsetneq$ unconstrained. Gaussian HMMs with $\boldsymbol{\mu}_1 \neq \boldsymbol{\mu}_2$ have $\hat{\eta}_{\mathrm{RO}} = 2\Phi(-\frac{1}{2}\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|) > 0$ (Appendix J), so RO fails and EDS fails. Such models are generic in the HMM parameter space, establishing $\mathrm{EDS} \subsetneq \mathrm{unconstrained}$. □

Remark 8. *The operational content of* $\mathrm{OS} \subsetneq \mathrm{EDS}$ *is that scalar Bayes-optimal stopping is achievable for a strictly broader class of multi-regime models than the identically-distributed attack case. Certificate* (31) *gives a three-state model where regime 2 is twice as emissive as regime 1, yet both are optimally tracked by the scalar* $\phi(\boldsymbol{\gamma}) = \gamma^{(1)} + 2\gamma^{(2)}$.

I Discretisation Error Bound

Proposition 3. *For* $\{Z_t\}$ *discretised to* N *equally-spaced levels on* $[0, C]$, $|\mathrm{ELD}^\star_N - \mathrm{ELD}^\star| \leq WC/N$.

Proof. The quantised chain $\hat{Z}_t = \lfloor Z_t N/C \rfloor \cdot C/N$ satisfies $|Z_t - \hat{Z}_t| \leq C/N$ at each step. The lead-time functional depends on $\{Z_t\}$ only through the hitting time $\tau_{\delta^\star}$; threshold shift by C/N changes τ by at most 1 step. Since $(t^\star - \tau)^+ \leq W$, the total error is bounded by $W \cdot (C/N)$. □

For $W = 20$, $C = 2$ ($K = 3$, $N = 200$): error ≤ 0.2 steps. For $C = 4$ ($K = 5$, $N = 200$): error ≤ 0.4 steps.

J Gaussian Emissions: Closed-Form η_{RO}

J.1 Shared covariance

When $p_k = \mathcal{N}(\boldsymbol{\mu}_k, \mathbf{\Sigma})$ for all $k \in \mathcal{S}_{\mathrm{atk}}$:

$$\eta_{\mathrm{RO}} = 2 \max_{k \in \mathcal{S}_{\mathrm{atk}}} \Phi\big(-\tfrac{1}{2}\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_1\|_{\mathbf{\Sigma}^{-1}}\big)\,.$$

RO holds exactly iff $\boldsymbol{\mu}_k = \boldsymbol{\mu}_1$ for all $k \in \mathcal{S}_{\mathrm{atk}}$.

J.2 Non-parametric estimation

For non-Gaussian emissions, train a binary classifier $\hat{f}_{kj}$ on held-out regime-labelled data to distinguish regimes k and j. The TV distance is estimated via the variational formula: $\mathrm{TV}(p_k/c_k, p_j) \leq \mathbb{E}_j[\max(\hat{r}_{kj}(\mathbf{y}) - c_k/c_j, 0)] + O(n^{-1/2})$, where $\hat{r}_{kj} = \hat{f}_{kj}/(1-\hat{f}_{kj})$. For CICIDS2017 (78 features, 5-fold CV, 5 seeds): $\hat{\eta}^{\mathrm{raw}}_{\mathrm{RO}} = 0.058 \pm 0.009$.

This estimator is evaluated on held-out test data, independently of the model fitted for detection; there is no circularity.

K Bias-Variance Ablation

Table 5: Lead time vs. training size on CICIDS2017 (steps, 5 seeds). EDS lead-time variance < 0.1 across all n; full HMM and PBVI improve substantially with data but do not reach EDS performance even at full sample size, consistent with Theorem 2.

Train n	EDS lead	Full HMM lead	PBVI lead
500	13.1	4.2	3.8
1000	13.8	7.3	6.9
2000	14.2	9.8	9.1
Full ($\sim$ 3400)	14.3	9.1	8.3

EDS lead time is stable from $n = 500$ (variance < 0.1 steps), confirming that its advantage is governed by the structural property rather than estimation variance. Full HMM and PBVI exhibit the bias-variance tradeoff: both improve substantially with data but remain below EDS performance at all sample sizes, consistent with the irreducible loss bound of Theorem 2.

L Implementation Details

HMM fitting. Baum–Welch with 50 random restarts; convergence threshold 10^{-6} on log-likelihood improvement. EDS regularisation: $\ell_{\mathrm{reg}} = \ell_{\mathrm{HMM}} - \lambda_\eta \eta$, with $\lambda_\eta = 0.5$ selected by held-out validation F1.

PBVI. 500 belief points generated by forward simulation from the prior; 50 backup iterations; no pruning.

TranAD and DeepSAD. Standard public implementations; hyperparameters selected by held-out validation F1.

Scalar chain. $N = 200$ equally-spaced levels on $[0, C]$; transition matrix estimated by kernel density smoothing of empirical one-step transitions.

CICIDS2017. 78 network-flow features; 10 PCA components retained for HMM (explaining $> 95\%$ of variance); raw 78 features used for $\hat{\eta}^{\mathrm{raw}}$ estimation. Train: days 1–3; test: days 4–5; FPR matched at ≤ 0.04.

Hardware. Google Colab (NVIDIA T4, 16 GB VRAM); total wall-clock < 24 hours.

Reproducibility. All hyperparameters fixed before test evaluation. Seeds: $\{0, 1, 2, 3, 4\}$ for all 5-seed experiments.

M Broader Impacts

This work is primarily theoretical. It characterises when a scalar score is Bayes-optimal for multi-regime sequential detection and quantifies the loss when it is not. Its societal footprint therefore runs mainly through the detection systems that practitioners build on top of such results.

Potential benefits. Sequential detectors are used in network security, industrial monitoring and health surveillance, where an alarm raised earlier can prevent damage, downtime or harm. Our results

offer three practical benefits. First, the EDS deviation η is a cheap, model-based test that can be run before deployment to decide whether a scalar score is justified or whether the full belief state should be retained. Second, the irreducible-loss result makes explicit that a structurally inadequate scalar summary cannot be repaired by more data or larger models, which can save engineering and compute effort. Third, the closed-form lead-time prediction under EDS supports transparent, auditable performance estimates in place of purely empirical tuning.

Potential risks. Better early detection is dual-use. The same techniques that give earlier warning of intrusions or equipment faults could strengthen monitoring of people, where earlier alarms may lead to false accusations, intrusive surveillance or unequal treatment of groups. A second risk is over-reliance on the diagnostic: a small estimated η on one dataset does not certify that scalar detection is safe elsewhere, and our guarantees hold only under the stated model assumptions. Finally, an earlier alarm trades against false alarms, and a poorly chosen constraint ϵ can shift the cost of errors onto the people being monitored.

Mitigations and scope. We recommend that deployments in settings involving individuals keep the false-alarm constraint conservative, retain human review of alarms, re-estimate η when the data distribution shifts, and fall back to the full belief state when η is large. Our experiments use a public network-intrusion benchmark and synthetic data, involve no personal or sensitive information, and we release no models or data that pose a specific misuse risk.

NeurIPS Paper Checklist

1. **Claims**

   Answer: [Yes]

   Justification: The abstract and introduction state the main contribution: a necessary and sufficient condition (EDS) for scalar Bayes-optimal stopping, together with its consequences for impossibility, closed-form lead time, robustness, and empirical validation (Sections 1, 4, 5 and 7).

2. **Limitations**

   Answer: [Yes]

   Justification: Limitations are discussed in Section 9 and Section 7: RO, MF and NF are restrictive structural conditions and EDS is a boundary, not a universal modelling assumption; the robustness bound deteriorates as $\Delta_{\min} \to 0$ or $\rho \to 1$ (Theorem 4); uniform boundedness is handled via truncation for unbounded emissions (Remark after Assumption 2); the characterisation covers one-dimensional statistics only; and the Full HMM comparison is specific to Baum–Welch with 50 restarts (Section 7.3). Empirical evidence is limited to synthetic models and one real dataset.

3. **Theory assumptions and proofs**

   Answer: [Yes]

   Justification: All results state their assumptions (Assumptions 1–3) and are proved in the main text (Theorem 1) or in Appendices A–I (Sections 3–5).

4. **Experimental result reproducibility**

   Answer: [Yes]

   Justification: The paper gives experimental protocols, dataset splits, evaluation metrics, baselines, hyperparameters, seeds and algorithmic details (Section 7, Algorithm 1, Appendices J.2 and L), which are sufficient to re-implement the experiments.

5. **Open access to data and code**

   Answer: [No]

   Justification: We do not release code. The only real dataset, CICIDS2017, is public, and the preprocessing, model configurations, hyperparameters and seeds are fully described in Appendix L so that the experiments can be re-implemented.

6. **Experimental setting/details**

   Answer: [Yes]

   Justification: Dataset splits, FPR matching, model configurations, hyperparameter selection and seeds are specified (Section 7 and Appendix L).

7. **Experiment statistical significance**

   Answer: [Yes]

   Justification: Results are reported as mean ± 1.96SE over 5 seeds, with paired t-tests at $\alpha = 0.05$ (Section 7).

8. **Experiments compute resources**

   Answer: [Yes]

   Justification: Hardware (Google Colab, NVIDIA T4, 16 GB) and total wall-clock time (< 24 hours) are reported in Appendix L.

9. **Code of ethics**

   Answer: [Yes]

   Justification: The research complies with the NeurIPS Code of Ethics. It uses a public benchmark and involves no sensitive data.

10. **Broader impacts**

    Answer: [Yes]

Justification: Positive and negative societal impacts, including dual-use surveillance risk, false-alarm cost shifting, and over-reliance on a small η, are discussed together with mitigations in Appendix M.

11. **Safeguards**

    Answer: [N/A]

    Justification: The work does not release models or datasets with significant misuse risk.

12. **Licenses for existing assets**

    Answer: [Yes]

    Justification: CICIDS2017 and the baselines (TranAD [11], DeepSAD [9], PBVI [7]) are publicly available and cited in the paper.

13. **New assets**

    Answer: [N/A]

    Justification: The paper does not introduce new datasets, code or deployable models.

14. **Crowdsourcing and research with human subjects**

    Answer: [N/A]

    Justification: The paper does not involve human subjects or crowdsourced data.

15. **Institutional review board (IRB) approvals**

    Answer: [N/A]

    Justification: No human subjects are involved.

16. **Declaration of LLM usage**

    Answer: [N/A]

    Justification: LLMs are not part of the core methodology. Any use was limited to writing and editing assistance and to understanding technical concepts.