STRONGER FORCING IS NOT PATTERN CHANGE IN ENVIRONMENTAL STREAMS UNDER EQUAL NORMS

Anonymous authors
Paper under double-blind review

ABSTRACT

We often use field norm to summarise multivariate environmental streams such as ocean temperature, groundwater level, and soil moisture. A changing norm can lead us to mistake stronger forcing of a fixed pattern for a change in the pattern itself. On fully matched equal-norm pairs, both readings share the field-norm and fixed-dictionary residual paths, so those scalar summaries cannot separate them. We treat the active pattern as a subspace and score mean pattern motion on the Grassmannian. On noiseless rank-1 matched streams, we prove that endpoint and block comparisons fail on the stated path classes. Mean pattern motion identifies one-way drift, leave-and-return, and odd-lag equal hopping under the stated lag conditions. For one-way steady drift, we establish trailing-window rates and finite-sample exchangeable-null calibration. CORAL (Construction on Real Atlantic Latents) is our labelled equal-norm suite in North Atlantic EOF coordinates. In this suite, field norm and leave-and-return endpoints exhibit the predicted failures. Mean pattern motion separates drift and leave-and-return but separates equal hopping only partly when we fix the hop lag and window. We release code and artifacts reproducing the experiments and figures in our anonymous repository.

1 INTRODUCTION

We observe environmental systems through multivariate streams and often monitor their field norm (Kim & Wu, 1999; Shields et al., 2026). A large excursion in a field such as ocean temperature can reflect stronger forcing of a fixed pattern or motion of the pattern itself (Fig. 1) (Werb & Rudnick, 2023). Confusing these readings treats amplitude change within a stable pattern as evidence of spatial reorganisation. Our question is diagnostic rather than causal: we ask which monitors distinguish a fixed subspace from a moving one when we match scalar energy. Table 1 answers this question under equal norms for three Grassmannian paths—steady drift, leave-and-return, and equal hopping—with each monitor governed by its stated path and lag conditions.

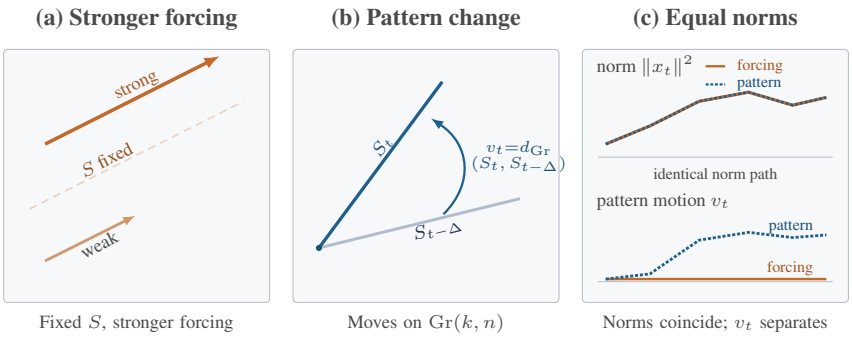

Figure 1: Stronger forcing is not pattern change under equal norms. **(a)** Fixed pattern S under stronger forcing; **(b)** lag-Δ Grassmann motion $v_t = d_{\mathrm{Gr}}(S_t, S_{t-\Delta})$; **(c)** the readings can share one $\|x_t\|^2$ path, while nonzero v_t detects pattern motion at a suitable lag for that path.

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.

Table 1: Population regime map for matched Construction pairs; conditions are in §3–4. **(1)** Matched norm/residual summaries fail on every class. **(2)** Endpoint and block statistics can separate steady drift, but each fails on leave-and-return or equal hopping. **(3)** Mean pattern motion separates one-way drift and leave-and-return at lag 1 (and non-folding paths at general lag); on period-2 hopping, it separates at odd lag and cancels up to boundary terms at even lag in finite records.

Climate monitor	*Path classes specified in §2*		
	Steady drift	**Leave-and-return**	**Equal hopping**
(1) Norm-based summaries (Lorenz, 1956; von Storch & Zwiers, 1999; Shields et al., 2026)	×	×	×
(2) Early/late angle proxy (Xu et al., 2011)	✓	×	△
Fixed-break angle proxy (Koo et al., 2026; Han & Inoue, 2015; Bai et al., 2024)	✓	△	×
Neighbouring-block subspace-angle proxy	✓	△	×
(3) Mean pattern motion (ours)	✓	✓	△

Note. ✓ separates on a subclass; × has risk $1/2$ on the construction; △ depends on lag, window, or schedule.

Related work. Climate science has long organised multivariate records with empirical orthogonal functions (EOFs) and related decompositions (Lorenz, 1956; von Storch & Zwiers, 1999; Kim et al., 2015; Shields et al., 2026); dynamic mode decomposition (DMD) makes an analogous mode–coefficient split (Schmid, 2010; Dylewsky et al., 2022). Neither alone resolves the equal-norm ambiguity: a large fixed-mode projection can reflect stronger forcing or a correlated nearby pattern. Xu, Xie, Kanai, and Koo compare subspaces at selected ends or blocks (Xu et al., 2011; Xie et al., 2018; Kanai et al., 2023; Koo et al., 2026). We use those comparisons as our batch statistics, and we score mean pattern motion over the full record. GROUSE (Balzano et al., 2010) tracks the subspace online and does not decide between stronger forcing and a change in the pattern.

Research gap. Under equal norms, matched streams can share field norm and batch-endpoint angles while S_t still moves on $\mathrm{Gr}(k, n)$. No endpoint or block statistic resolves all three Table 1 classes, including odd-lag hopping. Nor does online tracking provide finite-sample control for the fixed-window, fixed-lag, fixed-rank steady-drift decision problem studied here.

Our contributions. **(1)** We formulate the distinction on $\mathrm{Gr}(k, n)$ (Edelman et al., 1998; Absil et al., 2008), define equal-norm streams, and specify three path classes (§2). **(2)** Under stated conditions, we establish failures and complementarity of norm-based and batch statistics across the three classes, and show that mean pattern motion separates specified subclasses, including odd-lag hopping (§3–4). **(3)** For one-way steady drift, we derive trailing-window tracking and risk transfer under explicit window constraints, with distribution-free false-positive control using held-out exchangeable nulls (§5). **(4)** On our labelled *CORAL* benchmark in frozen ocean-EOF coordinates, we recover the predicted drift and leave-and-return separation with mean pattern motion but separate equal hopping only partly when we fix the hop lag and window in advance (§6).

2 Problem formulation: stronger forcing versus pattern change

We formalise on $\mathrm{Gr}(k, n)$ the distinction between *stronger forcing* of a fixed pattern and *pattern change*. The two readings can share one field-norm path, as in Fig. 1 and Table 1, even though only the latter moves the subspace (Kim et al., 2015; Werb & Rudnick, 2023; Shields et al., 2026).

Definition 1 (Stronger forcing vs. pattern change)**.** Under *stronger forcing*, $S_t \equiv S_1$ and coefficients a_t may vary in a fixed frame; under *pattern change*, S_t is not constant over time.

Generative model and scores. We represent centred $x_t \in \mathbb{R}^n$ through an active pattern $S_t = \mathrm{span}(U_t)$ with orthonormal frame $U_t \in \mathbb{R}^{n\times k}$ and projector P_{S_t} onto that pattern.

$$x_t \;=\; U_t\, a_t \;+\; \varepsilon_t, \tag{1}$$

with amplitudes $a_t \in \mathbb{R}^k$ and observation noise ε_t. For a reference subspace R, $r_t(R) = \|(I - P_R)x_t\|^2$ is the residual outside R. Write $d_{\mathrm{Gr}}(S, S')^2 = \sum_{i=1}^{k} \theta_i^2$ for principal angles θ_i (Edelman

et al., 1998; Chikuse, 2003). Over horizon T and lag $\Delta \geq 1$, we pair the field norm with lagged pattern motion, so both readings enter the following diagnostic score as complementary components:

$$z_t \;=\; \big(e_t,\, v_t\big), \qquad e_t = \|P_{S_t} x_t\|^2, \qquad v_t = d_{\mathrm{Gr}}(S_t, S_{t-\Delta}), \tag{2}$$

Set $S_s := S_1$ for $s \leq 0$, $\eta_t := d_{\mathrm{Gr}}(S_t, S_{t-1})$, $\bar{\eta} := \frac{1}{T}\sum_{t=2}^{T} \eta_t$, and $\bar{\eta}_\Delta := \frac{1}{T}\sum_{t=1}^{T} v_t$. Then $v_t = 0$ under stronger forcing, while pattern change can yield $v_t > 0$ at suitable lags under equal norms.

Equal-norm constraint. To isolate what norm-based summaries miss, we build equal-norm pairs that share $\|x_t\|^2$ at every time yet differ in which pattern S_t is active, as in Fig. 1 and Table 1.

Definition 2 (Norm-based summaries and equal-norm pairs)**.** For fixed $R_\star \in \mathrm{Gr}(k, n)$, a *norm-based summary* is any measurable function of $\big(\|x_s\|^2, r_s(R_\star)\big)_{s \leq T}$. An *equal-norm* (or *matched*) pair shares $\|x_t\|^2$ at every t; a *full matched pair* also shares $r_t(R_\star)$ at every t.

Definition 3 (Path classes). We specify the three path classes of Table 1 on $\mathrm{Gr}(k, n)$—steady drift, leave-and-return, and equal hopping—with parameters $L > 0$, plateau length m, and hop separation $\delta_{\mathrm{sym}} > 0$ on each class. For a binary rule φ, balanced risk is $\mathcal{R}(\varphi) = \frac{1}{2}\Pr_{\mathrm{forc}}(\varphi = \mathrm{patt}) + \frac{1}{2}\Pr_{\mathrm{patt}}(\varphi = \mathrm{forc})$; identical feature laws on the Construction 1 branches imply $\mathcal{R} = 1/2$. Let $\Pi_{\mathrm{drift}}(L)$ collect paths (S_t) with total one-step motion $\sum_{t=2}^{T} \eta_t \geq L$ and endpoint shift $d_{\mathrm{Gr}}(S_T, S_1) > 0$. Let $\Pi_{\mathrm{RT}}(L, m)$ collect leave-and-return paths with $S_1 = S_T$, total motion at least L, and $d_{\mathrm{Gr}}(S_t, S_1) = 0$ on the first and last m observations. Let $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$ collect even-length period-2 paths alternating between S_+ and S_- at separation δ_{sym}, with total motion at least L and equal occupancy of S_+ and S_- in every even-length block. We take $1 \leq m \leq \lfloor T/2 \rfloor$, $\delta_{\mathrm{sym}} \in (0, \pi/2)$, and only nonempty parameter regimes; in particular, Π_{sym} requires even T and $L \leq (T-1)\delta_{\mathrm{sym}}$, since a period-2 path hops at most once per step. We call one-way Π_{drift} schedules *steady drift*, Π_{RT} paths *leave-and-return*, and the Π_{sym} paths *equal hopping*.

Endpoint or block averages may agree while S_t moves. Construction 1 matches energy e_t; under Table 1's path and lag conditions, v_t and $\bar{\eta}_\Delta$ provide the needed separation between the readings.

3 Impossibility under equal norms and path classes

We prove impossibility on equal-norm pairs under Def. 2 and the path classes Π_{drift}, Π_{RT}, and Π_{sym}. Norm-based summaries cannot separate the readings even under rotationally invariant sub-Gaussian noise (Cor. 1). Batch endpoint and block angles fail on leave-and-return and equal hopping (Thm. 1). Table 1 marks the monitor–class combinations with risk $1/2$ on Construction 1 pairs.

Norm-based summaries. Under Def. 2, streams share the field norm and, in the full matched benchmark, the fixed-dictionary residual. We build rank-1 equal-norm pairs at fixed angle ψ to a unit vector r so only the pattern-change branch moves, isolating geometry from scale.

Construction 1 (Constant-angle matched streams)**.** Fix unit $r \in \mathbb{R}^n$, set $R_\star = \mathrm{span}(r)$, and fix angle $\psi \in (0, \pi/2)$. Let $\mathcal{C}_\psi := \{u \in \mathbb{S}^{n-1} : |\langle u, r\rangle| = \cos\psi\}$ (assume $n \geq 3$ for positive-length paths). For equal hopping at line separation δ_{sym}, feasibility additionally requires $\delta_{\mathrm{sym}} \leq 2\psi$ when $\psi < \pi/4$. Pick a rank-1 schedule $(u_t)_{t=1}^{T} \subset \mathcal{C}_\psi$ and amplitudes (α_t) with $A_{\min} \leq |\alpha_t| \leq A_{\max}$ and $A_{\min} > 0$. For smooth one-way drift, take $u_t = \gamma(t/T)$ from a fixed C^1 constant-speed curve $\gamma : [0, 1] \to \mathcal{C}_\psi$ of arc length $\Theta(L)$ whose sampled schedule has total motion at least L; leave-and-return and equal-hopping schedules use their corresponding classes below. Under stronger forcing, hold $u_\star = u_1$ fixed and set $x_t^{\mathrm{forc}} = \alpha_t u_\star$; under pattern change, set $x_t^{\mathrm{patt}} = \alpha_t u_t$. At every t the streams share $\|x_t\|^2 = \alpha_t^2$, while only a nonconstant pattern-change schedule moves through $\{\mathrm{span}(u) : u \in \mathcal{C}_\psi\} \subset \mathrm{Gr}(1, n)$ over the record, giving the matched pair used below.

Proposition 1 (Equal-norm summaries cannot separate)**.** *Under Construction 1, for every t,*

$$\|x_t^{\mathrm{forc}}\|^2 = \|x_t^{\mathrm{patt}}\|^2 = \alpha_t^2, \tag{3}$$

$$r_t^{\mathrm{forc}}(R_\star) = r_t^{\mathrm{patt}}(R_\star) = \alpha_t^2 \sin^2\psi, \tag{4}$$

so every norm-based summary matches on the two streams at every time (App. A.1).

Corollary 1 (Norm-based tests have risk $1/2$)**.** *Under Construction 1, every test measurable only through the norm-based summaries of Def. 2 has $\mathcal{R} = 1/2$; the result remains exact under i.i.d. rotationally invariant noise independent of the signal by subspace-preserving invariance (App. A.1).*

On these pairs, field norm and fixed-reference EOF residuals (Lorenz, 1956; von Storch & Zwiers, 1999; Shields et al., 2026) cannot distinguish the readings because Prop. 1 matches both scalar paths.

Endpoint and block angles. On leave-and-return and equal hopping paths (Π_{RT} and Π_{sym}), explicit early/late, fixed-break, and adjacent-window proxies compare end, break, or adjacent-block subspaces. They follow Xu et al. (2011); Han & Inoue (2015); Koo et al. (2026); Xie et al. (2018); Kanai et al. (2023); Basseville & Nikiforov (1993). Leave-and-return endpoints stay at S_1 despite path length L; equal hopping gives even blocks one leading subspace despite odd-lag motion.

Definition 4 (Even-length batch-endpoint monitors). A test uses *batch-endpoint Grassmann information only* if it is a measurable function of any combination of: (a) norm-based summaries (Def. 2); (b) the angle between leading subspaces of the first and last m observations, with even m; (c) angles between leading pre/post subspaces at finitely many break fractions whose blocks have even length; and (d) angles between leading subspaces of adjacent even-length blocks. We extract each leading subspace from that block's covariance as in App. B.1; write $\Phi_a, \Phi_b, \Phi_c, \Phi_d$ for feature groups (a)–(d). Following Xu et al. (2011); Xie et al. (2018); Kanai et al. (2023); Koo et al. (2026), these proxies are the even-length batch statistics in Table 1; Rem. 1 treats odd early/late windows.

Theorem 1 (Endpoint angles fail on leave-and-return; break and adjacent-window angles fail on equal hopping)**.** *Under Construction 1 with shared (α_t) in noiseless streams, part* (i) *is on Π_{RT} and parts* (ii)–(iii) *on Π_{sym}:* (i) Leave-and-return. *Any test that depends on the data only through early/late subspaces from the first and last m observations (or continuous functionals thereof) has risk $\mathcal{R} = 1/2$ on every path in $\Pi_{\mathrm{RT}}(L, m)$ under the equal-norm pair from Construction 1, since $S_{\mathrm{early}} = S_{\mathrm{late}} = S_1$ on both streams (App. A.4).* (ii) Equal-hopping breaks. *In rank* 1*, suppose $M = \frac{1}{2}(P_{S_+} + P_{S_-})$ has a unique leading line $S_\star$. Under constant amplitude $|\alpha_t| \equiv \alpha > 0$, any even-break pre/post Grassmann angle test has risk $\mathcal{R} = 1/2$ on the matched Construction pair in $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$ (App. A.4).* (iii) Adjacent windows. *For even tracker window w, adjacent length-w block angles vanish on $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$, so any test of those angles has risk $\mathcal{R} = 1/2$ under the mixture leading-space hypothesis of* (ii)*; App. A.4 gives the construction and proof.*

Those statistics use endpoints or blocks, not pattern motion along the record. On Π_{sym}, even-length blocks average to the equal-weight mixture of P_{S_+} and P_{S_-}, whose leading space $S_\star$ is shared across break and adjacent-window comparisons, so those batch angles vanish while odd-lag hops still move. Where endpoints and blocks miss that motion, mean pattern motion on $\mathrm{Gr}(k, n)$ separates the two equal-norm readings throughout the full observation record.

Remark 1 (Reading Theorem 1 in Table 1). Theorem 1 gives risk $1/2$ for the early/late statistic on leave-and-return and for the even-break and adjacent-window statistics on equal hopping. The procedures cited above also use statistics outside these even-length angle proxies. The $\triangle$ on equal hopping marks an odd-window early/late comparison, outside the even-m batch statistics of Def. 4. Thm. 1(i) covers every plateau length m on leave-and-return, with no restriction on the interior excursion between the opening plateau and the final return of the path to the opening subspace S_1.

4 Theoretical identifiability via pattern motion

Mean pattern motion separates stronger forcing from pattern change on Construction 1 equal-norm pairs under one-way drift, non-folding leave-and-return, and odd-lag equal hopping (Thm. 2; App. A.2). Complementarity shows no even-length batch-endpoint monitor (Def. 4) is uniform across steady drift, leave-and-return, and equal hopping on those matched pairs (Thm. 3). Population thresholding of $\bar{\eta}_\Delta$ recovers the Table 1 separation where those batch statistics fail.

Definition 5 (Population identification). Fix lag $\Delta \geq 1$ and threshold $\tau > 0$; given an equal-norm pair, declare *pattern change* when $\bar{\eta}_\Delta > \tau$ and *stronger forcing* otherwise. Mean pattern motion *identifies* a given matched pair when some $\tau = \tau(T)$ satisfies $0 \leq \tau(T) < \bar{\eta}_\Delta^{\mathrm{patt}}$ and yields $\mathcal{R} \to 0$; identification is *uniform* on a class only if one threshold sequence works for every path in that class. The oracle choice $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ is pair-specific, while every norm-based test retains $\mathcal{R} = 1/2$ (Cor. 1). A path *does not fold at lag* Δ when $\sum_{t=\Delta+1}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) \geq c_\Delta \sum_{t=2}^{T} \eta_t$ for some $c_\Delta \in (0, 1]$ (App. A.4). Such drift is *one-way*; such leave-and-return is *non-folding at lag* Δ.

Theorem 2 (Pattern-motion identifiability)**.** *In noiseless rank-*1 *streams under Construction 1 with shared (α_t), the following statements hold as $T \to \infty$ with $L, m, \delta_{\mathrm{sym}}$, and Δ fixed unless stated*

otherwise: (i) Steady drift. *On one-way schedules in* $\Pi_{\mathrm{drift}}(L)$ *sampled from a fixed smooth curve of arc length* $\Theta(L)$, $e_t^{\mathrm{forc}} = e_t^{\mathrm{patt}} = \alpha_t^2$ *for every* t, $\bar{\eta}_\Delta^{\mathrm{forc}} = 0$, *and* $\bar{\eta}_\Delta^{\mathrm{patt}} = \Theta(L/T)$ *for fixed* Δ *(App. A.2).* (ii) Leave-and-return. *On every path in* $\Pi_{\mathrm{RT}}(L, m)$, *early and late leading subspaces both equal* S_1 *on both streams (Thm. 1*(i)*), yet* $\bar{\eta}_1^{\mathrm{forc}} = 0$ *and* $\bar{\eta}_1^{\mathrm{patt}} \geq L/T$*; if the path does not fold at lag* Δ, *then* $\bar{\eta}_\Delta^{\mathrm{patt}} \geq c_\Delta L/T$ *(App. A.4).* (iii) Equal hopping. *On constant-amplitude* Π_{sym}, $\bar{\eta}_\Delta^{\mathrm{forc}} = 0$ *and break/adjacent risk is* $1/2$*;* $\bar{\eta}_\Delta^{\mathrm{patt}} = \Theta(\delta_{\mathrm{sym}})$ *for odd* Δ *and* $O(\Delta\delta_{\mathrm{sym}}/T)$ *for even* Δ *(App. A.4), so the rate depends on lag parity; the even-lag term comes only from record boundaries.*

Under stronger forcing, $v_t = 0$ and $\bar{\eta}_\Delta = 0$ for every $\Delta \geq 1$, while pattern-change branches carry the $\Theta(L/T)$ and odd-lag $\Theta(\delta_{\mathrm{sym}})$ gaps in Thm. 2; matched field norms and dictionary residuals then leave only mean pattern motion in (2) to separate the two equal-norm readings (Prop. 1).

Theorem 3 (Complementarity under even-length batch statistics)**.** *Under Construction 1 with equal-norm pairs:* (i) *Batch statistics* (b)–(d) *retain risk* $\mathcal{R} = 1/2$ *on the hard classes of Thm. 1 (early/late on* Π_{RT}*; break/adjacent on* Π_{sym} *under constant amplitude).* (ii) *No test that uses batch-endpoint Grassmann information only (Def. 4) has vanishing balanced risk simultaneously on* $\Pi_{\mathrm{drift}}(L)$, $\Pi_{\mathrm{RT}}(L, m)$, *and* $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$*: on* Π_{sym} *under constant amplitude the full feature vector* $\Phi = (\Phi_a, \Phi_b, \Phi_c, \Phi_d)$ *is identical on the matched branches, so every such test has* $\mathcal{R} = 1/2$ *there (App. A.3).* (iii) *On the subclasses of parts* (i)–(ii)*, population identification by mean pattern motion at* $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ *with fixed odd* Δ *separates stronger forcing from pattern change with vanishing risk on steady drift, non-folding leave-and-return, and odd-lag equal hopping (Cor. 2).*

Thms. 2 and 3 together establish the separation in Table 1: on equal hopping the full batch feature vector can match while odd-lag $\bar{\eta}_\Delta$ still moves, and on leave-and-return endpoints can remain at S_1 while $\bar{\eta}_\Delta$ still accumulates (Thm. 1(i)). Mean pattern motion separates at $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ (Cor. 2), and no Def. 4 batch monitor matches this separation across all three Table 1 path classes.

Corollary 2 (Population threshold identification)**.** *For each Theorem 2 pair with* $\bar{\eta}_\Delta^{\mathrm{patt}} > 0$*, the oracle threshold* $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ *gives* $\mathcal{R} = 0$ *in the noiseless model. One threshold is uniform over a class only if it lies below that class's infimum pattern-motion gap uniformly in* T *(App. A.2).*

5 Pattern motion estimation and decision

We estimate mean pattern motion online from (2). Lemma 1 and Theorem 4 bound trailing-window recovery via matrix Bernstein and Davis–Kahan theory (App. A.5). Corollary 3 calibrates $\hat{\tau}$ on stronger-forcing nulls. On one-way Construction 1, the online rule inherits Thm. 2(i) whenever (8) keeps the tracking error ζ_t below the relevant population mean pattern-motion gap.

Online scores. From trailing-window covariance $\hat{\Sigma}_t = \frac{1}{w}\sum_{s=t-w+1}^{t} x_s x_s^\top$ we set $\hat{S}_t = \mathrm{span}(\mathrm{SVD}_{\hat{k}}(\hat{\Sigma}_t))$, $\hat{e}_t = \|P_{\hat{S}_t} x_t\|^2$, and $\hat{v}_t = d_{\mathrm{Gr}}(\hat{S}_t, \hat{S}_{t-\Delta})$. The rate statements take $\hat{k} = k$; the CORAL protocol fixes its empirical rank separately. GROUSE (Balzano et al., 2010) is an empirical short-window substitute; the rates below apply to trailing-window SVD. For $t \geq w + \Delta$, our online decision score is the running mean $\bar{\hat{v}}_t := \frac{1}{t-w-\Delta+1}\sum_{s=w+\Delta}^{t} \hat{v}_s$. It targets population lag-Δ mean pattern motion $\bar{\eta}_\Delta$ and enters the null-calibrated decision rule (App. B). The analysis allows a fixed orthonormal representation $B \in \mathbb{R}^{n\times d}$, chosen independently of the monitored stream, when every $S_t \subseteq \mathrm{span}(B)$: apply the estimator to $B^\top x_t$ and let d denote the resulting dimension. Thus $d = n$ without reduction, whereas CORAL uses its frozen $d = 3$ EOF representation.

Assumption 1 (Standing conditions)**.** Observations follow (1) with $1 \leq k < d \leq n$, i.i.d. isotropic sub-Gaussian noise of coordinate variance σ^2, known rank k, and amplitudes satisfying $A_{\min} \leq \|a_t\| \leq A_{\max}$ with $A_{\max}/A_{\min} = O(1)$. Conditionally on (S_t), the pairs (a_t, ε_t) are independent across t, and each mean-zero ε_t is independent of (a_t, S_t); the signal Gram $\mathbb{E}[a_t a_t^\top \mid S_t]$ does not depend on S_t and is isotropic ($\mathbb{E}[aa^\top] = \lambda I_k$), so the top-$k$ eigengap is $\lambda_{\mathrm{gap}} = \lambda$. On constant-speed steady drift, $\eta_t = \Theta(\bar{\eta})$ uniformly, $A_{\max}^2 w\bar{\eta} = o(\lambda_{\mathrm{gap}})$, and $\bar{\eta}_\Delta = \Theta(\Delta\bar{\eta})$ (Thm. 2). Coordinate-wise isotropy suffices for the rate bounds below; we require rotational invariance only for Cor. 1's exact matched-pair law, not for the tracking and rate bounds in Lemma 1 and Thm. 4.

Lemma 1 (Trailing-window tracking)**.** *Under Assumption 1 with* $A_{\max}^2 w\bar{\eta} = o(\lambda_{\mathrm{gap}})$ *and* $w \gtrsim (A_{\max}^2 + \sigma^2)^2\{d + \log T\}/\lambda_{\mathrm{gap}}^2$*, trailing-window SVD yields* $\hat{S}_t$ *with uniform Grassmann error*

bounded for every $t \geq w$ by

$$\zeta_t := C\Big(\rho_w + \tfrac{A_{\max}^2}{\lambda_{\text{gap}}} \min\big(1, \textstyle\sum_{j=0}^{w-1} \eta_{t-j}\big)\Big), \qquad \rho_w := \tfrac{A_{\max}^2+\sigma^2}{\lambda_{\text{gap}}} \left\{ \sqrt{\tfrac{d+\log T}{w}} + \tfrac{d+\log T}{w} \right\}, \quad (5)$$

with probability at least $1 - T^{-2}$ simultaneously for all $t \geq w$, when using the trailing-window rank-k subspace estimator (matrix Bernstein–Davis–Kahan argument in App. A.5).

Theorem 4 (Steady-drift online rates and risk transfer)**.** *On the event of Lemma 1, the score bound (6) holds for every $t \geq w$, and the motion bound (7) holds for every $t \geq w + \Delta$:*

$$|\hat{e}_t - e_t| \leq C' \|x_t\|^2 \zeta_t, \quad (6)$$

$$|\hat{v}_t - d_{\text{Gr}}(S_t, S_{t-\Delta})| \leq \zeta_t + \zeta_{t-\Delta}. \quad (7)$$

On one-way steady-drift Construction 1 with $|\alpha_t| \equiv A > 0$, suppose $T \to \infty$ along sequences (w_T, Δ_T, d_T) that satisfy the following separation and horizon conditions:

$$\rho_w + \tfrac{A_{\max}^2}{\lambda_{\text{gap}}} w\bar{\eta} = o(\bar{\eta}_\Delta^{\text{patt}}), \qquad w + \Delta = o(T), \quad (8)$$

Then the averaged scores satisfy

$$\tfrac{1}{T} \sum_{t=w+\Delta}^{T} \hat{v}_t^{\text{forc}} = O(\rho_w), \quad (9)$$

$$\tfrac{1}{T} \sum_{t=w+\Delta}^{T} \hat{v}_t^{\text{patt}} \geq \bar{\eta}_\Delta^{\text{patt}} - O\big(\rho_w + \tfrac{A_{\max}^2}{\lambda_{\text{gap}}} w\bar{\eta} + \tfrac{w+\Delta}{T} \bar{\eta}_\Delta^{\text{patt}}\big). \quad (10)$$

Thus the population gap survives in $\bar{\hat{v}}_T$; the oracle threshold $\tau = \bar{\eta}_\Delta^{\text{patt}}/2$ yields vanishing balanced risk, while under rotationally invariant noise norm-based summaries retain $\mathcal{R} = 1/2$ (App. A.6).

Corollary 3 (Finite-sample null calibration)**.** *Let $Z_i = \bar{\hat{v}}_{T,i}$, $i = 1, \ldots, N$, be exchangeable held-out stronger-forcing terminal scores at a common horizon and fixed $(w, \Delta, \hat{k})$, choose level $q \in (0, 1)$, and set $\hat{\tau} = Z_{(\ell)}$ for $\ell = \lceil (N+1)(1-q) \rceil \leq N$. For an exchangeable fresh null Z_{N+1}, the strict rule $Z_{N+1} > \hat{\tau}$ has false-positive probability at most $(N+1-\ell)/(N+1) \leq q$ (Angelopoulos & Bates, 2023). With false-negative probability β, balanced risk is at most $(q+\beta)/2$ (finite-sample form: Prop. 3). Under Theorem 4 and (8), $\beta = o(1)$; taking $q = q_T \to 0$, $N_T \geq q_T^{-1} - 1$, and $N_T = o(T^2)$ gives vanishing risk for the calibrated decision rule (App. A.6).*

Decision rule. At each $t \geq w + \Delta$, declare *pattern change* if $\bar{\hat{v}}_t > \hat{\tau}_t$ and *stronger forcing* otherwise. Independent tuning data fix $(w, \Delta, \hat{k})$, and held-out calibration nulls set $\hat{\tau}_t$ at the same horizon; abstain if $T < w + \Delta$. The rule thresholds $\bar{\hat{v}}_t$ alone (Cor. 1): $\hat{e}_t$ tracks forcing on $\hat{S}_t$ but does not enter $\hat{\tau}$ or the decision. Condition (8) requires $\rho_w \to 0$ and $w\bar{\eta} = o(\bar{\eta}_\Delta)$ so w preserves the $\Theta(L/T)$ gap of Thm. 2(i); with Cor. 3, the online rule inherits steady-drift separation under the tracking conditions of Lemma 1. For fixed-curve drift, this instantaneous-space condition entails $\Delta \gg w$; fixed Δ is a population-identification regime, not an online-rate regime. For constant-amplitude smooth drift with fixed Δ/w, App. A.6 tracks the window-averaged space and removes the explicit $w\bar{\eta}$ remainder. Leave-and-return and equal hopping lie outside this guarantee, so we evaluate both recurrent regimes empirically in the labelled benchmark that follows.

Remark 2 (Null-calibration scope). $\hat{\tau}$ calibrates on stronger-forcing nulls at fixed $(w, \Delta, \hat{k})$ and sets the noise floor for $\bar{\hat{v}}$, not a path-class label on environmental streams. The guarantee assumes $\hat{k} = k$; Cor. 3 gives exact marginal control, while Cor. 4 gives finite-horizon ever-alarm control from stream maxima, both without within-stream independence. If (8) fails, null inflation or signal suppression invalidates the transfer, so we must retune (w, Δ) before applying the null-calibrated transfer.

6 CORAL: Labelled benchmark and SST comparison

This section has two empirical roles. *CORAL* (Construction on Real Atlantic Latents) is a labelled suite in frozen CMEMS North Atlantic SST coordinates (Szekely et al., 2026). We use it to test the Table 1 monitors on Construction 1 drift and leave-and-return signals and on the equal-norm hopping

variant below. The final heatwave comparison has no matched labels and therefore supports only a descriptive comparison of the monitors, not a recovery claim about observed climate.

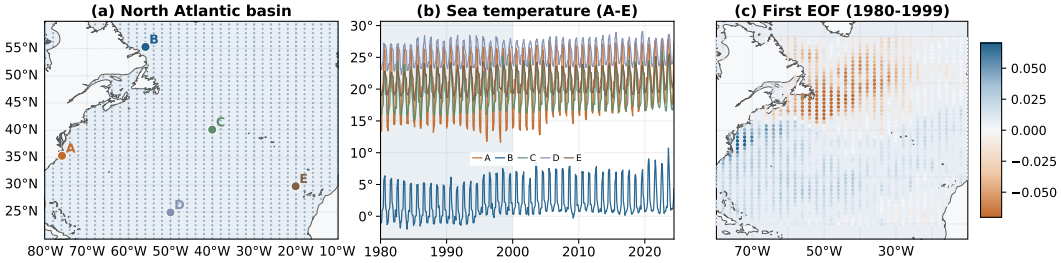

Figure 2: **CORAL**: labelled equal-norm signals in frozen CMEMS North Atlantic EOF frames (Szekely et al., 2026) (Construction 1 for drift and leave-and-return; hop forcing uses the mixture top space). **(a)** Basin/sites; **(b)** site temperatures with atlas window shaded; **(c)** 1980–1999 EOF1. Across three paths, two atlases, and two projection/noise controls, rank-1 monitors score the leading-3 representation. Ambient ρ=0.2 yields $\rho_{\text{score}} \approx 0.0089$ under isotropic noise, while projection removes residual noise; App. C gives the atlas, noise, schedules, and defaults.

Table 2: Paired ranking accuracy (%) for rank-1 monitors in a frozen leading-3 representation (Fig. 2; Construction 1 for drift and leave-and-return; hop forcing uses the mixture top space; CMEMS SST (Szekely et al., 2026)). Cells are mean±s.e. over 50 paired rankings (50% is chance). Field-norm accuracy remains near chance; mean pattern motion separates drift and leave-and-return, while hopping is only partly separated. Defaults are (w, Δ)=$(36, 12)$; hop uses Δ=11, neighbouring-block uses w=120, and we mark fixed residual and projected early/late degenerate on hop rows. At ambient ρ=0.2, $\rho_{\text{score}} \approx 0.0089$ under isotropic noise and vanishes under residual noise; the table tests geometric separation, not robustness to large or correlated in-frame noise.

Climate setting	Path class	Field norm	Fixed residual	Early/late angle	Projected early/late angle	Fixed-break angle	Neighbouring block angle	Mean pattern motion (ours)
CORAL-80s-I	Steady drift	48.0±7.1	58.0±7.1	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-80s-I	Leave-and-return	58.0±7.1	58.0±7.1	54.0±7.1	56.0±7.1	88.0±4.6	100.0±0.0	100.0±0.0
CORAL-80s-I	Equal hopping	50.0±7.1	*deg.*	74.0±6.3	*deg.*	66.0±6.8	80.0±5.7	100.0±0.0
CORAL-80s-R	Steady drift	50.0±0.0	50.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-80s-R	Leave-and-return	50.0±0.0	50.0±0.0	37.0±5.5	64.0±6.1	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-80s-R	Equal hopping	50.0±0.0	*deg.*	57.0±6.5	*deg.*	48.0±5.9	55.0±6.6	53.0±1.7
CORAL-00s-I	Steady drift	44.0±7.1	44.0±7.1	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-00s-I	Leave-and-return	52.0±7.1	52.0±7.1	52.0±7.1	50.0±7.1	86.0±5.0	100.0±0.0	100.0±0.0
CORAL-00s-I	Equal hopping	48.0±7.1	*deg.*	74.0±6.3	*deg.*	68.0±6.7	84.0±5.2	100.0±0.0
CORAL-00s-R	Steady drift	50.0±0.0	50.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-00s-R	Leave-and-return	50.0±0.0	50.0±0.0	43.0±5.9	51.0±6.1	100.0±0.0	100.0±0.0	100.0±0.0
CORAL-00s-R	Equal hopping	50.0±0.0	*deg.*	48.0±5.9	*deg.*	47.0±5.4	53.0±6.9	49.0±2.3

Protocol. For each atlas era (1980–1999 and 2000–2019), we freeze a leading-3 EOF frame on demeaned and detrended CMEMS anomalies over n=1529 North Atlantic grid cells. We then embed Construction 1 (drift and leave-and-return) or its equal-norm hopping variant in that frame (App. C; algorithms in App. B). Each trial uses latent equal-norm signals with shared amplitude $|\alpha_t| \equiv 1$, cone angle ψ=30°, and record length T=480. Pattern motion follows one path class (Π_{drift}, Π_{RT}, or Π_{sym}), and we add independent observation noise to each branch. Four conditions cross atlas era with ambient isotropic noise or frame-orthogonal residual variation, which we remove by projecting onto the frozen frame, at ρ=0.2. The projected isotropic-noise ratio is only $\rho_{\text{score}} = \rho\sqrt{3/1529} \approx 0.0089$, hop rows use Δ=11, and neighbouring-block uses w=120. Rank-1 monitor m yields $s_m(x_{1:T})$ on each branch in the leading-3 representation with default (w, Δ)=$(36, 12)$. Table 2 reports *paired ranking accuracy*: the fraction of 50 trials on which s_m is higher on the pattern-change branch, with chance at 50%. Paired branches share latent geometry but receive independent noise; their ranking uses no threshold, and only Fig. 4 reports null-calibrated decisions in the main text through the held-out steady-drift experiment shown in panel (c).

Equal-norm ranking table on CORAL. Table 2 tests Table 1's separation in two frozen ocean-EOF embeddings and two noise controls at ambient ρ=0.2 (half-turn drift; leave-and-return excursion $=\pi/3$); field norm stays near chance ($\approx 50\%$) across every benchmark cell and condition.

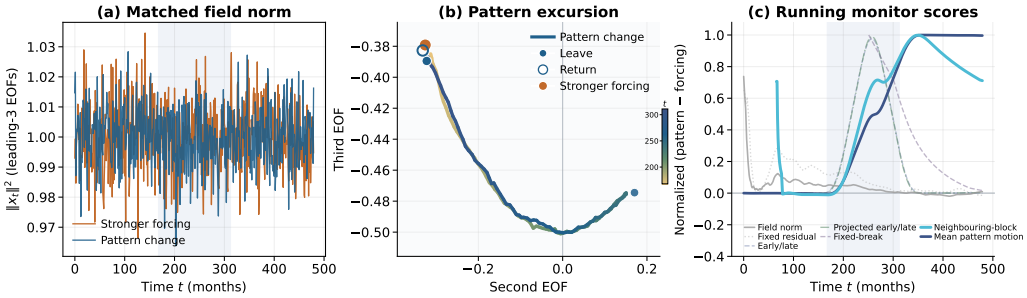

Figure 3: Leave-and-return on CORAL-80s-I (seed-0 pair; atlas as in Fig. 2). **(a)** $\|x_t\|^2$ in the leading-3 EOF subspace on both branches (isotropic noise; leave-and-return interval shaded). **(b)** Pattern-change path in the second–third EOF plane (colour = time t), with leave and return near the stronger-forcing baseline; first EOF is the fixed cone axis. **(c)** Normalised pattern−forcing score differences for the Table 2 monitors over time. Mean pattern motion and neighbouring-block remain separated through the excursion; projected early/late and fixed-break rise then fall as early and late windows realign, while the field-norm difference stays near zero throughout.

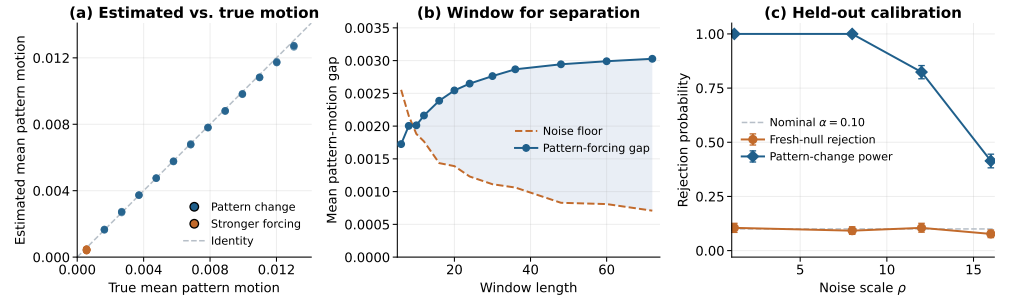

Figure 4: Steady-drift experiments on CORAL-80s-I. **(a)** Estimates track truth while stronger forcing stays near zero; **(b)** larger w clears the noise floor; **(c)** held-out calibration at level 0.10 across noise scales, showing fresh-null rejection and pattern-change power (100 trials, 49 nulls).

On steady drift and leave-and-return, early/late endpoints are near perfect on drift and near chance on isotropic leave-and-return ($\approx$ 50–56%), while mean pattern motion is 100% on every such cell. Neighbouring-block reaches 100% on leave-and-return because mid-record blocks contain the excursion, so Table 1's $\triangle$ is schedule-dependent here. Fixed-break remains partially separating under isotropic noise ($\approx$ 86–88%). Projected early/late removes frozen EOF1; on hop forcing streams this also removes the mixture-axis signal, so the proxy is degenerate. Equal hopping remains only partly separated. Under residual noise, fixed-break, neighbouring-block, and mean-motion accuracies stay near chance ($\approx$ 47–55%) because even w=36 averages period-2 hops toward the mixture (App. A.6). Isotropic hop batch monitors reach $\approx$ 66–84%, and mean pattern motion reaches 100%. Figure 3 shows a noisy CORAL-80s-I leave-and-return pair through latent matched norms, its EOF-plane excursion, and running monitor scores. Appendix D gives drift and hopping examples under the same frozen atlas, rank, and isotropic-noise protocol. These examples illustrate existing benchmark cells; they add no path classes, monitor definitions, scoring rules, or evaluation settings.

Benchmark scope. Drift and leave-and-return keep Construction 1's fixed-S_1 forcing branch and match the hard cases in Table 1. For hopping, the forcing branch is the mixture top space: fixed residual is unmatched and projected early/late is degenerate after removing that axis; Table 2 marks both. At odd Δ, population mean pattern motion follows lag parity, whereas even w=36 can average period-2 states; the hop rows are finite-window stress tests outside Thm. 4's steady-drift guarantee. Under isotropic noise, both atlas embeddings have the same score distribution and are not independent evidence about North Atlantic EOF geometry or CMEMS-era dynamics.

Online rule on steady drift. We test the online rule of §5 on CORAL-80s-I steady drift at elevated noise, where the trailing-window estimate meets the noise floor (Fig. 4). In Apps. C.6–C.8 we test sensitivity, ever-alarm control, and calibration shift under elevated noise while holding the frozen frame fixed. Calibration and test streams share the CORAL-80s-I coordinate frame.

Across $\rho \in \{1.2, 8, 12, 16\}$, fresh-null rejection is 7.7–10.5%, covering 10% except at ρ=16; power is 100%, 100%, 82.4%, and 41.4% as the tracking floor nears the gap (Cor. 3; Lem. 1).

North Atlantic heatwave. On the unlabelled North Atlantic record, the Jul 2023–Jun 2024 heatwave raises field norm ($\times$2.09; OISST confirms warmth) while mean pattern motion falls ($\times$0.77). Cross-era angles exceed within-2010s references ($\approx$ 1.35 rad). Field-norm, residual, and motion ratios fall inside their baseline ranges, including motion at 0.77. In Table 3 and Fig. 5 we report monitor disagreement, not path-class attribution for the 2023–2024 window.

Table 3: Table 1 monitors on demeaned–detrended CMEMS SST (Szekely et al., 2026), 2010s versus Jul 2023–Jun 2024, with OISST (Huang et al., 2021). Angles use a within-2010s reference; continuous rows give heatwave-to-2010s ratios, descriptive month-resampling ranges, and 109 baseline blocks ((w, Δ)=(36, 12); motion in rad per step under the shared basin mask). The final column is a descriptive reading, not a path-class label or causal attribution on this unlabelled record.

Climate monitor	2010s	Heatwave	vs 2010s	Baseline ref rad or $\times$	Descriptive reading
Early/late and fixed-break	across eras		1.35 rad [1.13, 1.56]	0.87 rad	Large cross-era angle
Neighbouring-block	across eras		1.23 rad [1.07, 1.56]	0.87 rad	Large cross-era angle
Field norm	401	838	$\times$2.09 [1.76, 2.45]	$\times$0.97 [0.70, 1.37]	Higher field norm
Prefix-fit residual	328	732	$\times$2.23 [1.87, 2.66]	$\times$0.99 [0.76, 1.23]	Higher residual
Mean pattern motion (ours)	0.0465	0.0359	$\times$0.77 [0.64, 0.91]	$\times$0.92 [0.37, 1.70]	No motion increase
OISST mean anomaly	2023–24		+1.08°C	87% warm	Observed warmth

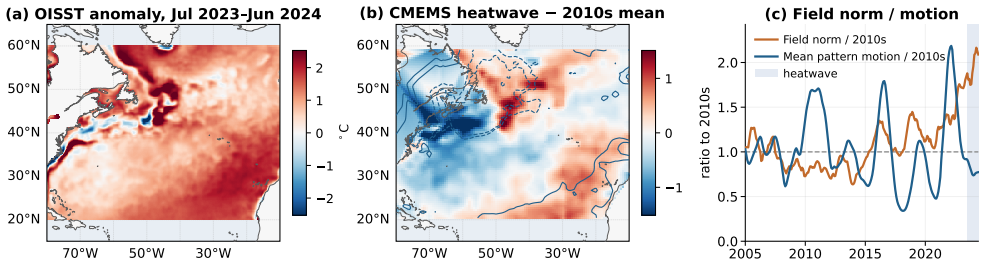

Figure 5: North Atlantic Jul 2023–Jun 2024 heatwave (Table 3): monitors disagree, not one path class. **(a)** OISST anomaly (Huang et al., 2021); **(b)** CMEMS heatwave−2010s mean with leading-EOF contours; **(c)** norm and motion ratios at (w, Δ)=(36, 12) (heatwave window shaded).

7 Discussion and Conclusion

We show that, on matched pairs, field norm and fixed residuals coincide, while endpoint and block statistics each miss a recurrent path. Our mean-pattern-motion statistic identifies drift, non-folding leave-and-return, and odd-lag hopping in noiseless rank 1; we combine trailing-window transfer with held-out calibration to obtain a finite-sample steady-drift rule. At the specified lag and window, CORAL (Szekely et al., 2026) separates drift and leave-and-return but hopping only partly (Table 2); under structured noise, mean motion has the highest hopping point estimate (64–99% power) but loses high-noise return power (App. C.7). We interpret monitor disagreement as path sensitivity, not environmental cause; in the unlabelled heatwave, field norm and OISST (Huang et al., 2021) show warmth while cross-era angles and motion disagree, so we report a descriptive comparison rather than path-class recovery or causal interpretation for this unlabelled event (Table 3; Fig. 5).

Limitations and outlook. We prove exact hopping hardness only in rank 1, and our rates assume known rank, an eigengap, and one-way drift. In CORAL, we fix a leading-3 EOF frame and test only two eras and prescribed paths; structured-noise tests are synthetic, not observational noise models. Our one-basin heatwave analysis depends on preprocessing and reference era and cannot identify cause. Our calibration guarantee requires exchangeability at fixed $(w, \Delta, \hat{k})$; selecting rank, lag, or window from data or facing distribution shift can invalidate false-positive control (App. C.8). We leave higher-rank path families, valid adaptive calibration, dependent-noise theory, time-uniform shift robustness, and minimax rates for environmental streams as central directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools to improve the manuscript's grammar and readability, assist with coding and tests, and provide feedback on experimental design and interpretation of numerical results. The authors reviewed all AI-assisted edits and code, verified the resulting artifacts and numerical outputs against the stated protocol, and take responsibility for the final text, code, and results.

ETHICS STATEMENT

This work studies statistical methods for separating stronger forcing from pattern change in multivariate environmental streams. We use publicly available observational products and no human subjects. The methods are tools for scientific interpretation of environmental change, not automated attribution or operational warning systems.

REPRODUCIBILITY STATEMENT

The anonymous repository contains the implementation, executable tests, frozen CORAL atlases and protocol, experiment scripts, and generated figures and tables. Appendix A gives complete proofs and assumptions, Appendix B specifies the monitor algorithms, and Appendix C documents data processing, path and noise constructions, scoring rules, sample sizes, and calibration settings. The large upstream CMEMS file is not redistributed; the repository's data documentation identifies the public product and acquisition route, while the derived panels and frozen atlases needed by the reported experiments are included.

REFERENCES

P.-A. Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008. doi: 10.1515/9781400830244.

Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023. doi: 10.1561/2200000101.

Jushan Bai, Jiangtao Duan, and Xu Han. The likelihood ratio test for structural changes in factor models. *Journal of Econometrics*, 238(2):105631, 2024. doi: 10.1016/j.jeconom.2023.105631.

Laura Balzano, Robert Nowak, and Benjamin Recht. Online identification and tracking of subspaces from highly incomplete information. In *48th Annual Allerton Conference on Communication, Control, and Computing*, pp. 704–711, 2010. doi: 10.1109/Allerton.2010.5706976.

Michèle Basseville and Igor V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, 1993. ISBN 0-13-126780-9.

Yasuko Chikuse. *Statistics on Special Manifolds*. Number 174 in Lecture Notes in Statistics. Springer, 2003. doi: 10.1007/978-0-387-21540-2.

Daniel Dylewsky, David Barajas-Solano, Tong Ma, Alexandre M. Tartakovsky, and J. Nathan Kutz. Stochastically forced ensemble dynamic mode decomposition for forecasting and analysis of near-periodic systems. *IEEE Access*, 10:33440–33448, 2022. doi: 10.1109/ACCESS.2022.3161438.

Alan Edelman, Tomás A. Arias, and Steven T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998. doi: 10.1137/S0895479895290954.

Xu Han and Atsushi Inoue. Tests for parameter instability in dynamic factor models. *Econometric Theory*, 31(5):1117–1152, 2015. doi: 10.1017/S0266466614000486.

Boyin Huang, Chunying Liu, Viva Banzon, Eric Freeman, Garrett Graham, Bill Hankins, Tom Smith, and Huai-Min Zhang. Improvements of the daily optimum interpolation sea surface temperature (DOISST) version 2.1. *Journal of Climate*, 34(8):2923–2939, 2021. doi: 10.1175/JCLI-D-20-0166.1.

Takumi Kanai, Naoya Sogi, Atsuto Maki, and Kazuhiro Fukui. Time-series anomaly detection based on difference subspace between signal subspaces, 2023. arXiv:2303.17802.

Kwang-Y. Kim and Qigang Wu. A comparison study of EOF techniques: Analysis of nonstationary data with periodic statistics. *Journal of Climate*, 12(1):185–199, 1999. doi: 10.1175/1520-0442(1999)012⟨0185:ACSOET⟩2.0.CO;2.

Kwang-Yul Kim, Benjamin Hamlington, and Hanna Na. Theoretical foundation of cyclostationary EOF analysis for geophysical and climatic variables: Concepts and examples. *Earth-Science Reviews*, 150:201–218, 2015. doi: 10.1016/j.earscirev.2015.06.003.

Bonsoo Koo, Benjamin Wong, and Ze-Yu Zhong. Disentangling structural breaks in factor models for macroeconomic data. *Journal of Business & Economic Statistics*, pp. 1–12, 2026. doi: 10.1080/07350015.2025.2583205.

Edward N. Lorenz. Empirical orthogonal functions and statistical weather prediction. Technical Report Scientific Report No. 1, Statistical Forecasting Project, Department of Meteorology, Massachusetts Institute of Technology, 1956.

Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010. doi: 10.1017/S0022112010001217.

Christine A. Shields, Derek DeSantis, Alexandra Jonko, Jiwoo Lee, Julie M. Caron, Adam S. Phillips, and John Fasullo. Challenges and alternatives to empirical orthogonal functions for earth system data. *Scientific Reports*, 16(1):18794, 2026. doi: 10.1038/s41598-026-49505-z.

Tanguy Szekely, Jerome Gourrion, Sylvie Pouliquen, and Gilles Reverdin. CORA, coriolis ocean dataset for reanalysis. E.U. Copernicus Marine Service Information, doi: 10.48670/mds-00383, 2026. Product INSITU_GLO_PHY_TS_OA_MY_013_052.

Hans von Storch and Francis W. Zwiers. *Statistical Analysis in Climate Research*. Cambridge University Press, 1999. ISBN 978-0-521-45071-3. doi: 10.1017/CBO9780511612336.

Per-Åke Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972. doi: 10.1007/BF01932678.

Benjamin E. Werb and Daniel L. Rudnick. Remarkable changes in the dominant modes of north pacific sea surface temperature. *Geophysical Research Letters*, 50(4):e2022GL101078, 2023. doi: 10.1029/2022GL101078.

Liyan Xie, George V. Moustakides, and Yao Xie. First-order optimal sequential subspace change-point detection. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 111–115, 2018. doi: 10.1109/GlobalSIP.2018.8646377.

Yi Xu, Zhongfei Zhang, Philip S. Yu, and Bo Long. Pattern change discovery between high dimensional data sets. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, pp. 1097–1106, 2011. doi: 10.1145/2063576.2063735.

Yi Yu, Tengyao Wang, and Richard J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015. doi: 10.1093/biomet/asv008.

A Detailed Proofs of Theoretical Claims

We give full proofs of the claims in §2–5: equal-norm confounding and batch-angle impossibility (§3), population identification by $\bar{\eta}_\Delta$ (§4), and online rates with null calibration (§5). Notation follows the main text unless stated otherwise. Throughout, we take d_{Gr} to be the metric with $d_{\mathrm{Gr}}(S, S')^2 = \sum_{i=1}^{k} \theta_i^2$ for principal angles θ_i; any other standard principal-angle metric is equivalent up to constants absorbed in C (cf. §2). In particular $\|P_S - P_{S'}\|_{\mathrm{F}}^2 = 2\sum_{i=1}^{k} \sin^2 \theta_i$ and $\|P_S - P_{S'}\|_{\mathrm{op}} = \sin\theta_{\max}$, so $d_{\mathrm{Gr}}(S, S')$ is comparable to both projector distances up to factors depending only on k. Universal constants may change from line to line; we denote them by C, C', c, c' throughout Apps. A.1–A.6 and absorb them into big-O notation where convenient.

Proof map. The scope remarks prove Prop. 2. Appendix A.1 proves Prop. 1, Cor. 1, and Lem. 4. Appendix A.2 and the three path-specific parts of App. A.4 prove Thm. 2 and Cor. 2. Theorem 3(i) follows from Thm. 1, App. A.3 proves part (ii), and Cor. 2 gives part (iii). Appendix A.5 proves Lem. 1; App. A.6 proves Thm. 4, Cor. 3, Prop. 3, and Cor. 4, then treats overlapping lags and long-window suppression. Shared lemmas appear first; tracking-specific lemmas appear immediately before the online-rate and null-calibration arguments that use them.

Shared lemmas (population geometry and averaging).

Lemma 2 (Eigengap stability under operator perturbation)**.** *Let A, B be real symmetric $n \times n$ matrices with $\|A - B\|_{\mathrm{op}} \le \delta$. Suppose A has a unique top-k eigenspace $S \in \mathrm{Gr}(k, n)$ with eigengap $\lambda_{\mathrm{gap}} := \lambda_k(A) - \lambda_{k+1}(A) > 0$. If $\delta < \lambda_{\mathrm{gap}}/2$, then B has a unique top-k eigenspace $S' \in \mathrm{Gr}(k, n)$ with $\lambda_k(B) - \lambda_{k+1}(B) \ge \lambda_{\mathrm{gap}} - 2\delta$ and $d_{\mathrm{Gr}}(S, S') \le C\,\delta/\lambda_{\mathrm{gap}}$ for a constant $C = C(k)$ depending only on the subspace dimension k (independent of ambient dimension n).*

Proof. Order the eigenvalues $\lambda_1(A) \ge \cdots \ge \lambda_n(A)$. Weyl's inequality gives $|\lambda_j(B) - \lambda_j(A)| \le \delta$ for every j, hence $\lambda_k(B) - \lambda_{k+1}(B) \ge (\lambda_k(A) - \delta) - (\lambda_{k+1}(A) + \delta) = \lambda_{\mathrm{gap}} - 2\delta > 0$. Thus B has a unique top-k space. Davis–Kahan (Yu et al., 2015) applied to the k-dimensional invariant subspaces of A and B yields $d_{\mathrm{Gr}}(S, S') \le C\,\delta/\lambda_{\mathrm{gap}}$ for a constant $C = C(k)$ depending only on k. □

Lemma 3 (Chord-versus-arc comparison on a fixed regular curve)**.** *Fix $k \ge 1$ and sample $(S_t)_{t=1}^{T}$ from one fixed regular C^1 curve on $\mathrm{Gr}(k, n)$ whose speed is bounded away from zero, with one-step increments $\eta_t := d_{\mathrm{Gr}}(S_t, S_{t-1})$ (take $\eta_t = 0$ for $t \le 1$). For $\Delta \ge 1$ and $t \ge \Delta + 1$, write $\ell_t(\Delta) := \sum_{j=0}^{\Delta-1} \eta_{t-j}$ for the path length along the segment ending at t. Then $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) \le \ell_t(\Delta)$ by the triangle inequality. If the sampling mesh tends to zero and the curve is one-way on that segment—it does not reverse in the path parameter between $S_{t-\Delta}$ and S_t—then, uniformly over sampled segments with $\ell_t(\Delta) \to 0$, $d_{\mathrm{Gr}}(S_t, S_{t-\Delta})/\ell_t(\Delta) \to 1$.*

Proof. Parametrise the fixed regular curve by arc length on its compact parameter interval. Its unit tangent is uniformly continuous, so on every subsegment of continuous arc length $L \to 0$ the endpoint displacement in normal coordinates is L times the initial unit tangent plus $o(L)$, uniformly over the segment location. Hence the geodesic endpoint distance is $L\{1 + o(1)\}$ by the standard first-order comparison on Riemannian manifolds (Absil et al., 2008, Ch. 5). The vanishing sampling mesh makes the discrete polygonal length $\ell_t(\Delta) = L\{1 + o(1)\}$ uniformly; combining the two comparisons proves the claim. □

Scope remarks.

Remark 3 (Scope of the scalar class). Definition 2 uses a fixed dictionary $R_\star$ and, in the noisy case, rotationally invariant noise (law invariant under $O(n)$; e.g. isotropic Gaussian). Spatially correlated or only covariance-isotropic ε_t, and adapted dictionaries $\hat{R}_t$, lie outside our class; common fixed-EOF and field-norm monitors fall inside it. Under isotropic sub-Gaussian noise, Eqs. (3)–(4) hold in expectation up to shared $n\sigma^2$ centres and concentrate about those centres (Lem. 4); under rotational invariance the $G_{R_\star}$ argument gives exact risk $\mathcal{R} = 1/2$ in law (App. A.1). Covariance isotropy alone does not imply that law identity: we require rotational invariance of the noise law for the exact risk-$1/2$ conclusion in Cor. 1.

Proposition 2 (Isoclinic rank-k matched lift)**.** *We state Proposition 1 and Theorem 2 for rank 1 under Construction 1. When $n \ge 2k$, let $R, Q_t \in \mathbb{R}^{n \times k}$ be orthonormal with $R^\top Q_t = 0$, set $U_t = R\cos\psi + Q_t \sin\psi$, and let $R_\star = \mathrm{span}(R)$. Then $U_t^\top U_t = I_k$ and, for arbitrary coefficients a_t, $\|U_t a_t\|^2 = \|a_t\|^2$ and $\|(I - P_{R_\star})U_t a_t\|^2 = \sin^2\psi\,\|a_t\|^2$. Thus fixing $Q_\star$ on the forcing branch and varying Q_t on the pattern branch gives a full matched pair; rotationally invariant noise preserves the scalar-law identity because the stabiliser of $R_\star$ acts transitively on the orthogonal k-frames. Mean pattern motion then separates any such higher-rank path satisfying the same lag-gap conditions.*

Proof. The identities follow from $R^\top Q_t = 0$ and $R^\top R = Q_t^\top Q_t = I_k$: $U_t^\top U_t = \cos^2\psi\, I_k + \sin^2\psi\, I_k = I_k$, while $(I - P_{R_\star})U_t = Q_t \sin\psi$. The orthogonal group on $R_\star^\perp$ maps any orthonormal k-frame $Q_\star$ to any other Q_t while fixing $R_\star$. Applying that map to rotationally invariant noise proves

equality in law of the scalar summaries, and the stated lag gap is exactly the separation condition for mean pattern motion. □

Remark 4 (Higher-rank hardness and rates). The exact leave-and-return hardness of Thm. 1(i) extends with d_{Gr} on $\mathrm{Gr}(k, n)$, but its exact finite-block equal-hopping claims (ii)–(iii) are rank 1. Even lags on Π_{sym} yield only the temporal-boundary contribution $O(\Delta\delta_{\mathrm{sym}}/T)$; the full-rank case $k = n$ makes d_{Gr} degenerate and is excluded. Rank-1 hardness on Π_{RT} and Π_{sym} does not use slow drift or an isotropic Gram; extending exact empirical block-angle failure to $k > 1$ would require deterministic blockwise coefficient balance, not merely an isotropic population Gram. Lemma 1 separately permits $k > 1$ under its population-Gram assumptions.

A.1 Matching and confounding of norm-based monitors (Prop. 1)

Restatement (Prop. 1). Under Construction 1, for every t,

$$\|x_t^{\mathrm{forc}}\|^2 = \|x_t^{\mathrm{patt}}\|^2 = \alpha_t^2, \tag{11}$$

$$r_t^{\mathrm{forc}}(R_\star) = r_t^{\mathrm{patt}}(R_\star) = \alpha_t^2 \sin^2 \psi, \tag{12}$$

so every norm-based summary matches on the two streams.

Restatement (Cor. 1). Any test of those statistics alone therefore has risk $\mathcal{R} = 1/2$; rotationally invariant noise preserves that exact risk by $G_{R_\star}$-invariance on $\mathcal{C}_\psi$, while independent-coordinate isotropic sub-Gaussian noise yields concentration about the shared schedule.

Proof. Noiseless matching. Under Construction 1 we use noiseless observations:

$$x_t^{\mathrm{forc}} = \alpha_t\, u_\star, \qquad x_t^{\mathrm{patt}} = \alpha_t\, u_t, \tag{13}$$

with unit vectors $u_\star, u_t \in \mathcal{C}_\psi$ and $R_\star = \mathrm{span}(r)$. By definition of $\mathcal{C}_\psi$,

$$|\langle u_\star, r\rangle| = |\langle u_t, r\rangle| = \cos\psi. \tag{14}$$

Since $\|u_\star\| = \|u_t\| = 1$,

$$\|x_t^{\mathrm{forc}}\|^2 = \alpha_t^2 \|u_\star\|^2 = \alpha_t^2, \tag{15}$$

$$\|x_t^{\mathrm{patt}}\|^2 = \alpha_t^2 \|u_t\|^2 = \alpha_t^2. \tag{16}$$

For the residual to $R_\star$,

$$(I - rr^\top) x_t^{\mathrm{forc}} = \alpha_t \big(u_\star - \langle u_\star, r\rangle\, r\big), \tag{17}$$

so that

$$r_t^{\mathrm{forc}}(R_\star) = \|(I - rr^\top) x_t^{\mathrm{forc}}\|^2 = \alpha_t^2 \big(1 - \langle u_\star, r\rangle^2\big) = \alpha_t^2 \sin^2 \psi, \tag{18}$$

and likewise $r_t^{\mathrm{patt}}(R_\star) = \alpha_t^2 \sin^2 \psi$. Therefore the pairs $\big(\|x_t^{\mathrm{forc}}\|^2,\, r_t^{\mathrm{forc}}(R_\star)\big)$ and $\big(\|x_t^{\mathrm{patt}}\|^2,\, r_t^{\mathrm{patt}}(R_\star)\big)$ coincide for every t. Any norm-based monitor $\Phi_t = f_t\big(\{\|x_s\|^2,\, r_s(R_\star)\}_{s\le t}\big)$ (Def. 2) is a functional of that common sequence, so $\Phi_t^{\mathrm{forc}} = \Phi_t^{\mathrm{patt}}$ identically for every time index t under Construction 1.

Noiseless risk. By the matching above, the sequences $\{\Phi_t^{\mathrm{forc}}\}_{t\le T}$ and $\{\Phi_t^{\mathrm{patt}}\}_{t\le T}$ are identical as deterministic sequences. Hence for any (possibly randomised) measurable test φ that uses only the scalar stream $\{\Phi_t\}_{t\le T}$ and no Grassmann or batch-angle information from the observations,

$$\mathbb{P}_{\mathrm{forc}}(\varphi = \mathrm{patt}) = \mathbb{P}_{\mathrm{patt}}(\varphi = \mathrm{patt}), \qquad \mathbb{P}_{\mathrm{forc}}(\varphi = \mathrm{forc}) = \mathbb{P}_{\mathrm{patt}}(\varphi = \mathrm{forc}). \tag{19}$$

Balanced risk is therefore

$$\mathcal{R} = \tfrac{1}{2}\big(\mathbb{P}_{\mathrm{forc}}(\varphi = \mathrm{patt}) + \mathbb{P}_{\mathrm{patt}}(\varphi = \mathrm{forc})\big) = \tfrac{1}{2}\big(p + (1-p)\big) = \tfrac{1}{2}, \tag{20}$$

where $p := \mathbb{P}_{\mathrm{forc}}(\varphi = \mathrm{patt})$.

Noisy rotationally invariant case. Fix a time index and write $x = \alpha u + \varepsilon$ with unit $u \in \mathcal{C}_\psi$, shared amplitude α, and noise ε independent of u. Assume the noise is centred and *rotationally invariant*: $Q\varepsilon \sim \varepsilon$ for all $Q \in O(n)$ (e.g. isotropic Gaussian, or any spherical sub-Gaussian law; Rem. 3). Let

$G_{R_\star} = \{Q \in O(n) : Qr = \pm r\}$ be the stabiliser of the reference subspace. The group $G_{R_\star}$ acts transitively on $\mathcal{C}_\psi$: for any $u, u' \in \mathcal{C}_\psi$ there exists $Q \in G_{R_\star}$ with $Qu = u'$.

For any $Q \in G_{R_\star}$ one has $Q(I - rr^\top) = (I - rr^\top)Q$, and therefore

$$\|Qx\|^2 = \|x\|^2, \tag{21}$$

$$\|(I - rr^\top)Qx\|^2 = \|(I - rr^\top)x\|^2. \tag{22}$$

If $x = \alpha u + \varepsilon$, then $Qx = \alpha(Qu) + Q\varepsilon$. Rotational invariance gives $Q\varepsilon \sim \varepsilon$, so

$$Qx \overset{d}{=} \alpha(Qu) + \varepsilon. \tag{23}$$

Consequently the joint law of $\big(\|x\|^2,\, r(R_\star)\big)$ depends on u only through its $G_{R_\star}$-orbit. Transitivity on $\mathcal{C}_\psi$ implies that the law is constant for all $u \in \mathcal{C}_\psi$.

Under Construction 1 with a shared schedule $(\alpha_t)_{t=1}^T$ and i.i.d. copies ε_t independent of (u_t), the product measures of $\big\{\|x_t\|^2,\, r_t(R_\star)\big\}_{t \le T}$ therefore coincide on the stronger-forcing stream ($u_t \equiv u_\star$) and any pattern-change schedule $(u_t) \subset \mathcal{C}_\psi$. Any test of that scalar sequence alone therefore has risk $\mathcal{R} = 1/2$ on both the stronger-forcing and pattern-change branches under Construction 1. □

Lemma 4 (Noisy matched scalars concentrate)**.** *Under Construction 1 in rank* 1*, let the noise vectors be independent across time with independent, mean-zero sub-Gaussian coordinates of variance σ^2, and write $x_t = \alpha_t u_t + \varepsilon_t$ with $\|u_t\| = 1$ and $u_t \in \mathcal{C}_\psi$. There exists a constant C depending only on the sub-Gaussian proxy such that, simultaneously for all $t \le T$,*

$$\big|\|x_t\|^2 - \alpha_t^2 - n\sigma^2\big| \le C\big(|\alpha_t|\sigma\sqrt{\log(nT)} + \sigma^2\sqrt{n\log(nT)} + \sigma^2\log(nT)\big), \tag{24}$$

$$\big|r_t(R_\star) - \alpha_t^2\sin^2\psi - (n-1)\sigma^2\big| \le C\big(|\alpha_t|\sigma\sqrt{\log(nT)} + \sigma^2\sqrt{n\log(nT)} + \sigma^2\log(nT)\big), \tag{25}$$

with probability at least $1 - T^{-2}$. In particular both Construction 1 branches share the same (α_t) schedule, so their scalar paths concentrate about a common deterministic centre.

Proof. Fix t and write $x = \alpha u + \varepsilon$ with $\|u\| = 1$, $u \in \mathcal{C}_\psi$, and ε having independent mean-zero sub-Gaussian coordinates of variance σ^2, independent of u. Expand

$$\|x\|^2 = \alpha^2 + 2\alpha\langle u, \varepsilon\rangle + \|\varepsilon\|^2. \tag{26}$$

Isotropy gives $\mathbb{E}[\|\varepsilon\|^2] = n\sigma^2$ and $\mathbb{E}[\langle u, \varepsilon\rangle] = 0$. A standard one-dimensional sub-Gaussian tail yields $|\langle u, \varepsilon\rangle| \le C\sigma\sqrt{\log(nT)}$ with probability at least $1 - T^{-3}$ for a universal constant C (absorbing the proxy). Hanson–Wright applied to the identity matrix yields $\big|\|\varepsilon\|^2 - n\sigma^2\big| \le C\sigma^2\{\sqrt{n\log(nT)} + \log(nT)\}$ on the same event (up to adjusting C). Combining these bounds gives (24) at a fixed t.

For the residual, set $P = rr^\top$ and $Q = I - P$. Then $Qu = \big(u - \langle u, r\rangle\, r\big)$ has $\|Qu\| = \sin\psi$, and

$$r(R_\star) = \|Qx\|^2 = \alpha^2\sin^2\psi + 2\alpha\langle Qu, \varepsilon\rangle + \|Q\varepsilon\|^2. \tag{27}$$

The same one-dimensional tail controls $|\langle Qu, \varepsilon\rangle|$, while Hanson–Wright for the rank-$(n-1)$ projector Q gives $\big|\|Q\varepsilon\|^2 - (n-1)\sigma^2\big| \le C\sigma^2\{\sqrt{n\log(nT)} + \log(nT)\}$ on an event of probability at least $1 - T^{-3}$. This yields (25) at a fixed t.

A union bound over $t = 1, \dots, T$ gives probability at least $1 - T^{-2}$ simultaneously for all t. The centres $\alpha_t^2 + n\sigma^2$ and $\alpha_t^2\sin^2\psi + (n-1)\sigma^2$ depend on the shared amplitude schedule (α_t) and on (n, σ, ψ) alone, not on which $u_t \in \mathcal{C}_\psi$ is active, so both Construction 1 branches concentrate about the same deterministic scalar path. □

A.2 Pattern-motion identifiability (Theorem 2)

Restatement (Thm. 2). In noiseless rank-1 streams under Construction 1 with shared (α_t), the following statements hold as $T \to \infty$ with $L, m, \delta_{\mathrm{sym}}$, and Δ fixed unless stated otherwise: (i) *Steady drift.* On one-way schedules in $\Pi_{\mathrm{drift}}(L)$ sampled from a fixed smooth curve of arc length $\Theta(L)$, $e_t^{\mathrm{forc}} = e_t^{\mathrm{patt}} = \alpha_t^2$ for every t, $\bar\eta_\Delta^{\mathrm{forc}} = 0$, and $\bar\eta_\Delta^{\mathrm{patt}} = \Theta(L/T)$ for fixed Δ. (ii) *Leave-and-return.* On every path in $\Pi_{\mathrm{RT}}(L, m)$, early and late leading subspaces both equal S_1 on both streams

(Thm. 1(i)), yet $\bar{\eta}_1^{\mathrm{forc}} = 0$ and $\bar{\eta}_1^{\mathrm{patt}} \geq L/T$; if the path does not fold at lag Δ, then $\bar{\eta}_\Delta^{\mathrm{patt}} \geq c_\Delta L/T$. (iii) *Equal hopping.* On Π_{sym} at constant nonzero amplitude, $\bar{\eta}_\Delta^{\mathrm{forc}} = 0$; break/adjacent angles have risk $1/2$, while $\bar{\eta}_\Delta^{\mathrm{patt}} = \Theta(\delta_{\mathrm{sym}})$ at odd Δ and $O(\Delta\delta_{\mathrm{sym}}/T)$ at even Δ.

Restatement (Cor. 2). For each stated pair with $\bar{\eta}_\Delta^{\mathrm{patt}} > 0$, setting $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ yields $\mathcal{R} = 0$ in the noiseless model; class-uniform identification requires one threshold sequence below the infimum gap.

Apps. A.4–A.4 prove parts (ii)–(iii) of Theorem 2; part (i) and Corollary 2 follow.

Proof. Under Construction 1 with $\varepsilon_t = 0$ and rank 1, $S_t = \mathrm{span}(u_t)$ and $e_t = \|P_{S_t} x_t\|^2 = \|x_t\|^2 = \alpha_t^2$ on both streams, so $e_t^{\mathrm{forc}} = e_t^{\mathrm{patt}} = \alpha_t^2$.

Stronger forcing. Here $u_t \equiv u_\star$, so S_t is constant. For every lag $\Delta \geq 1$ and every t, $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = 0$, hence $v_t^{\mathrm{forc}} = 0$ and $\bar{\eta}_\Delta^{\mathrm{forc}} = 0$.

Pattern change (one-way, general lag $\Delta \geq 1$). Let $\gamma : [0,1] \to \mathcal{C}_\psi$ be the fixed C^1 constant-speed path of arc length $\Theta(L)$ from Construction 1, set $u_t = \gamma(t/T)$, and write path increment $\eta_t := d_{\mathrm{Gr}}(S_t, S_{t-1})$ and mean increment $\bar{\eta} := \frac{1}{T}\sum_{t=2}^T \eta_t$ as in §3. Constant speed gives $\sum_{t=2}^T \eta_t = \Theta(L)$ because $\eta_t = d_{\mathrm{Gr}}(S_t, S_{t-1})$ is the Riemannian increment of the fixed lift sampled at rate T, hence $\bar{\eta} = \Theta(L/T)$; in particular $\bar{\eta}_1^{\mathrm{patt}} = \Theta(L/T)$ when $\Delta = 1$. For general fixed Δ, write $\ell_t(\Delta) := \sum_{j=0}^{\Delta-1} \eta_{t-j}$ (taking $\eta_t = 0$ for $t \leq 1$). Then $\ell_t(\Delta) = \Theta(\Delta L/T) = \Theta(\Delta\bar{\eta})$ uniformly for $t \in \{\Delta+1, \ldots, T\}$, and $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) \leq \ell_t(\Delta)$ by the triangle inequality. The one-way constant-speed lift satisfies Lemma 3, so there exists $\varepsilon_T \to 0$ with $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) \geq (1-\varepsilon_T)\,\ell_t(\Delta)$ on interior indices. For the first Δ indices, constant speed gives $d_{\mathrm{Gr}}(S_t, S_1) = O(tL/T)$, so their T-normalised contribution is $O(\Delta^2 L/T^2) = o(L/T)$ for fixed Δ. Averaging the interior terms therefore gives

$$\bar{\eta}_\Delta^{\mathrm{patt}} = \frac{1}{T}\sum_{t=1}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = \Theta(\Delta\bar{\eta}) = \Theta(L/T). \tag{28}$$

In particular $\bar{\eta}_\Delta^{\mathrm{patt}} > \bar{\eta}_\Delta^{\mathrm{forc}} = 0$ for all large T.

Population threshold (Cor. 2). On the null, $\bar{\eta}_\Delta^{\mathrm{forc}} = 0$, so the rule $\bar{\eta}_\Delta > \tau$ with $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ never declares pattern change. On the alternative, $\bar{\eta}_\Delta^{\mathrm{patt}} = 2\tau > 0$, so the rule always declares pattern change and has balanced risk zero. A single threshold works uniformly over a class only when it lies below that class's infimum positive gap. Appendices A.4 and A.4 prove the leave-and-return and equal-hopping gaps in Thm. 2(ii)–(iii). □

A.3 Complementarity (Theorem 3)

Restatement (Thm. 3). Under Construction 1 with equal-norm pairs: (i) Batch statistics (b)–(d) retain risk $\mathcal{R} = 1/2$ on the hard classes of Thm. 1 (early/late on Π_{RT}; break/adjacent on Π_{sym} under constant amplitude). (ii) No test that uses batch-endpoint Grassmann information only (Def. 4) has vanishing balanced risk simultaneously on $\Pi_{\mathrm{drift}}(L)$, $\Pi_{\mathrm{RT}}(L, m)$, and $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$: on Π_{sym} under constant amplitude the full feature vector $\Phi = (\Phi_a, \Phi_b, \Phi_c, \Phi_d)$ is identical on the matched branches, so every such test has $\mathcal{R} = 1/2$ there. (iii) On the subclasses of parts (i)–(ii), population identification by mean pattern motion at $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ with fixed odd Δ separates stronger forcing from pattern change with vanishing balanced risk on steady drift, non-folding leave-and-return, and odd-lag equal hopping (Cor. 2).

Proof. Full feature identity on equal hopping. Let φ be any test that uses batch-endpoint Grassmann information only (Def. 4), hence a measurable function of any combination of the features $\Phi = (\Phi_a, \Phi_b, \Phi_c, \Phi_d)$ comprising norm-based summaries (a), early/late subspaces from even-length windows (b), even-break subspaces (c), and adjacent even-length block subspaces (d). Fix a matched Construction 1 pair on $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$ with $|\alpha_t| \equiv \alpha > 0$; its stronger-forcing subspace is fixed at $S_1 = S_+$.

Under equal norms, Φ_a is identical on the stronger-forcing and pattern-change branches (Prop. 1). Apps. A.4 and A.4 show that Φ_c and Φ_d are likewise identical: the within-stream angles are zero

on both branches (pattern-change blocks have leading space $S_\star$; stronger-forcing blocks have S_+). For (b), Def. 4 takes the angle between even early/late windows. Every such pattern-change block has equal occupancy of S_+ and S_- (Def. 3), hence leading space $S_\star$ and early/late angle zero; the fixed stronger-forcing branch at S_+ also has angle zero. Thus Φ_b is identical on both branches, and the full angle-and-scalar vector Φ is identical. Any measurable $\varphi(\Phi)$ therefore has balanced risk $\mathcal{R} = 1/2$ on Π_{sym}.

No uniform coverage. A test with $\mathcal{R} = 1/2$ on Π_{sym} cannot have vanishing balanced risk on all three path classes simultaneously under Def. 4. □

A.4 Endpoint angles fail on leave-and-return; break and adjacent-window angles fail on equal hopping (Theorem 1)

Restatement (Thm. 1). Under Construction 1 with shared (α_t) in noiseless streams, part (i) is on Π_{RT} and parts (ii)–(iii) on Π_{sym}: (i) *Leave-and-return.* Any test that depends on the data only through early/late subspaces from the first and last m observations (or continuous functionals thereof) has risk $\mathcal{R} = 1/2$ on every path in $\Pi_{\mathrm{RT}}(L, m)$ under the equal-norm pair from Construction 1, since $S_{\mathrm{early}} = S_{\mathrm{late}} = S_1$ on both streams. (ii) *Equal-hopping breaks.* In rank 1, suppose $M = \frac{1}{2}(P_{S_+} + P_{S_-})$ has a unique leading line $S_\star$. Under constant amplitude $|\alpha_t| \equiv \alpha > 0$, any even-break pre/post Grassmann angle test has risk $\mathcal{R} = 1/2$ on the matched Construction pair in $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$. (iii) *Adjacent windows.* For even tracker window w, adjacent length-w block angles vanish on $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$, so any test of those angles has risk $\mathcal{R} = 1/2$ under the mixture leading-space hypothesis of (ii).

Part (i): leave-and-return (Π_{RT}).

Proof. Fix $(S_t) \in \Pi_{\mathrm{RT}}(L, m)$ and the matched stronger-forcing path $S_t \equiv S_1$. By Definition 3,

$$S_t = S_1 \qquad \text{for all } t \le m \text{ and all } t > T - m, \tag{29}$$

and $S_T = S_1$, with path length at least L on the pattern-change stream.

Noiseless early/late subspaces. In the noiseless Construction 1 model, $x_t = \alpha_t u_t$ with $S_t = \mathrm{span}(u_t)$ exactly. The empirical covariance on the first m samples is

$$\hat{\Sigma}_{\mathrm{early}} = \frac{1}{m}\sum_{t=1}^{m} x_t x_t^\top = \frac{1}{m}\sum_{t=1}^{m} \alpha_t^2\, u_t u_t^\top. \tag{30}$$

Since $S_t = S_1$ for all $t \le m$, one has $u_t u_t^\top = P_{S_1}$ throughout the early window, so $\hat{\Sigma}_{\mathrm{early}}$ is a positive multiple of P_{S_1} (using $A_{\min} > 0$) and its leading subspace is exactly $S_{\mathrm{early}} = S_1$. The same argument on $\{T-m+1, \dots, T\}$ yields $S_{\mathrm{late}} = S_1$. Both identities hold under stronger forcing ($S_t \equiv S_1$) and on every leave-and-return pattern-change path. Hence any measurable functional of $(S_{\mathrm{early}}, S_{\mathrm{late}})$—including $d_{\mathrm{Gr}}(S_{\mathrm{early}}, S_{\mathrm{late}})$ and continuous images thereof—takes the same value on both streams, so balanced risk equals $1/2$ on every path in $\Pi_{\mathrm{RT}}(L, m)$.

Mean pattern motion. Under stronger forcing, $\bar{\eta}_\Delta = 0$ for every $\Delta \ge 1$. On $\Pi_{\mathrm{RT}}(L, m)$, path length $\sum_{t=2}^{T} \eta_t \ge L$, so

$$\bar{\eta}_1 = \frac{1}{T}\sum_{t=2}^{T} \eta_t \ge \frac{L}{T}. \tag{31}$$

For a fixed lag $\Delta \ge 1$, the same order need not hold on every path in $\Pi_{\mathrm{RT}}(L, m)$: a period-2 oscillation in the interior yields only $O(\Delta/T)$ boundary motion at even Δ while its one-step path length is $\Theta(T)$. If the path does not fold at lag Δ (Def. 5), then directly

$$\bar{\eta}_\Delta = \frac{1}{T}\sum_{t=\Delta+1}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) \ge \frac{c_\Delta}{T}\sum_{t=2}^{T} \eta_t \ge \frac{c_\Delta L}{T}. \tag{32}$$

□

Part (ii): equal-hopping breaks (Π_{sym}). Fix constant amplitude $|\alpha_t| \equiv \alpha > 0$ and the matched Construction pair whose stronger-forcing stream is fixed at $S_1 = S_+$. In the noiseless model, for any finite grid $\mathcal{G} \subset (0,1)$ of even-length break fractions, writing $S_{\mathrm{pre}}(b)$ and $S_{\mathrm{post}}(b)$ for the leading subspaces before and after time fraction b, any test of $\{d_{\mathrm{Gr}}(S_{\mathrm{pre}}(b), S_{\mathrm{post}}(b))\}_{b\in\mathcal{G}}$ has risk $\mathcal{R} = 1/2$ on $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$.

Proof. In the noiseless model, for rank 1, let $S_\pm = \mathrm{span}(u_\pm)$ with $d_{\mathrm{Gr}}(S_+, S_-) = \delta_{\mathrm{sym}} \in (0, \pi/2)$, and let the path alternate $S_+, S_-, S_+, S_-, \ldots$ for even T, so the path lies in $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$ for $L = \Theta(T\delta_{\mathrm{sym}})$.

Mixture eigenspace (rank 1). The two-point average projector is

$$M := \tfrac{1}{2}\big(P_{S_+} + P_{S_-}\big) = \tfrac{1}{2}\big(u_+u_+^\top + u_-u_-^\top\big). \tag{33}$$

In the plane $\mathrm{span}\{u_+, u_-\}$ the eigenvalues of M are $\frac{1+\cos\delta_{\mathrm{sym}}}{2}$ (multiplicity one, eigenvector along $u_+ + u_-$ up to sign) and $\frac{1-\cos\delta_{\mathrm{sym}}}{2}$ (along $u_+ - u_-$); all other eigenvalues are 0. For $\delta_{\mathrm{sym}} \in (0, \pi/2)$ one has $\frac{1+\cos\delta_{\mathrm{sym}}}{2} > \frac{1-\cos\delta_{\mathrm{sym}}}{2} \geq 0$, so M has a unique top eigenvalue with eigenspace

$$S_\star := \mathrm{span}(u_+ + u_-) \qquad \text{(after normalising)}. \tag{34}$$

Even-length blocks. Under constant amplitude $|\alpha_t| \equiv \alpha$, the population covariance on any time block B is

$$\Sigma_B = \frac{1}{|B|}\sum_{t\in B} \mathbb{E}[x_t x_t^\top] = \alpha^2 \cdot \frac{1}{|B|}\sum_{t\in B} P_{S_t}. \tag{35}$$

By the definition of $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$, every even-length block has equal occupancy of S_+ and S_-, hence

$$\frac{1}{|B|}\sum_{t\in B} P_{S_t} = M \tag{36}$$

and $\Sigma_B = \alpha^2 M$ has leading subspace $S_\star$. In particular, for every $b \in \mathcal{G}$ with even pre/post lengths, $S_{\mathrm{pre}}(b) = S_{\mathrm{post}}(b) = S_\star$, so $d_{\mathrm{Gr}}(S_{\mathrm{pre}}(b), S_{\mathrm{post}}(b)) = 0$.

Matched stronger-forcing stream. The Construction stronger-forcing stream with $S_t \equiv S_+$ also yields the zero break-angle vector on $\mathcal{G}$. Any test of that vector alone therefore has risk $\mathcal{R} = 1/2$.

□

Part (iii): adjacent windows. In the noiseless model with constant amplitude, for even w, writing S_t^{L} and S_t^{R} for the leading subspaces of the length-w blocks ending at t and starting at $t+1$ (whenever both blocks lie in $\{1, \ldots, T\}$), adjacent-block angles $\chi_t(w) = d_{\mathrm{Gr}}(S_t^{\mathrm{L}}, S_t^{\mathrm{R}})$ vanish for every such t, so any test of $\{\chi_t(w)\}_t$ has risk $\mathcal{R} = 1/2$. Odd-lag mean pattern motion separates on the same class: $\bar{\eta}_\Delta = \Theta(\delta_{\mathrm{sym}})$ for odd Δ (even Δ cancels up to temporal-boundary terms).

Proof. We reuse the noiseless constant-amplitude alternating construction of App. A.4 on $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$, with unique mixture top space $S_\star$.

Fix even w and any t for which both the length-w block ending at t and the length-w block starting at $t+1$ lie in $\{1, \ldots, T\}$. Because the window length is even and the alternation has period 2, the length-w block ending at t contains exactly $w/2$ visits to S_+ and $w/2$ visits to S_-, so its population covariance equals $\alpha^2 M$ and has leading subspace $S_t^{\mathrm{L}} = S_\star$. The adjacent block starting at $t+1$ is likewise an even-length period-2 window, hence has the same equal occupancy and leading subspace $S_t^{\mathrm{R}} = S_\star$. Therefore

$$\chi_t(w) = d_{\mathrm{Gr}}(S_t^{\mathrm{L}}, S_t^{\mathrm{R}}) = 0. \tag{37}$$

The matched stronger-forcing stream fixed at S_+ yields the same zero sequence $\{\chi_t(w)\}_t$, so any test of that sequence has risk $\mathcal{R} = 1/2$.

Odd-lag mean pattern motion. On the period-2 path, for interior indices t with $t, t-\Delta \in \{1, \ldots, T\}$, $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = \delta_{\mathrm{sym}}$ when Δ is odd and $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = 0$ when Δ is even (same pattern after an even lag). At most 2Δ indices near the temporal boundary can differ from this rule (using the convention $S_s := S_1$ for $s \leq 0$), contributing at most $2\Delta\delta_{\mathrm{sym}}/T$ to the average when Δ is odd.

Hence $\bar{\eta}^{\mathrm{patt}}_{\Delta} = \frac{1}{T}\sum_{t=1}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = \delta_{\mathrm{sym}} \cdot \frac{T-2\Delta}{T} + O(\Delta/T) = \Theta(\delta_{\mathrm{sym}})$ for odd Δ, and $\bar{\eta}_{\Delta} = O(\Delta\delta_{\mathrm{sym}}/T)$ for even Δ, matching Theorem 2(iii). □

Remark 5 (Trailing-window coverage on Π_{sym}). Trailing-window SVD under Theorem 4 does not cover Π_{sym}; we score the △ cells of Table 1 with the reported batch and trailing-window monitors, outside the steady-drift window of that theorem.

A.5 Subspace tracking (Lemma 1)

Restatement (Lem. 1). Under Assumption 1 with $A^2_{\max} w\bar{\eta} = o(\lambda_{\mathrm{gap}})$ and $w \gtrsim (A^2_{\max} + \sigma^2)^2\{d + \log T\}/\lambda^2_{\mathrm{gap}}$, trailing-window SVD yields $\hat{S}_t$ with uniform Grassmann error bounded for every $t \geq w$ by

$$\zeta_t := C\Big(\rho_w + \tfrac{A^2_{\max}}{\lambda_{\mathrm{gap}}} \min\big(1, \textstyle\sum_{j=0}^{w-1} \eta_{t-j}\big)\Big), \qquad \rho_w := \tfrac{A^2_{\max}+\sigma^2}{\lambda_{\mathrm{gap}}} \left\{ \sqrt{\tfrac{d+\log T}{w}} + \tfrac{d+\log T}{w} \right\}, \quad (38)$$

with probability at least $1 - T^{-2}$ simultaneously for all $t \geq w$, for the trailing-window rank-k subspace estimator.

Proof. We work throughout under Assumption 1. When we use a fixed representation B, replace x_s by $B^\top x_s$ and U_s by $B^\top U_s$. Because B is orthonormal, contains every S_s, and is fixed independently of the monitored stream, this preserves signal norms, principal angles, temporal independence, and isotropic coordinate noise; the proof then takes place in $\mathbb{R}^d$. The unreduced case is $B = I_n$ and $d = n$. Write $x_s = U_s a_s + \varepsilon_s$ as in (1), $S_s = \mathrm{span}(U_s)$, and

$$\Sigma_s := \mathbb{E}[x_s x_s^\top \mid S_s] = U_s\, \mathbb{E}[a_s a_s^\top \mid S_s]\, U_s^\top + \sigma^2 I_n \quad (39)$$

under isotropic noise of coordinate variance σ^2 (Assumption 1). Assumption 1 gives $A_{\min} \leq \|a_s\| \leq A_{\max}$. Together with the Gram hypothesis and Assumption 1, the top-k eigengap of Σ_s is at least $\lambda_{\mathrm{gap}} > 0$; for $k = 1$ the single signal eigenvalue is $\Theta(\|a_s\|^2) \in [\Theta(A^2_{\min}), \Theta(A^2_{\max})]$. Under the isotropic Gram $\mathbb{E}[aa^\top] = \lambda I_k$ in Assumption 1, $\Sigma_s = \lambda P_{S_s} + \sigma^2 I_n$, so each population covariance is a Lipschitz function of the projector P_{S_s} alone.

Step 1: concentration of the windowed covariance. Fix the path $(S_s)_{s \leq T}$ (deterministic under Construction 1, or condition on it in general). Set $Y_s := x_s x_s^\top - \Sigma_s$. Conditionally on the path, the pairs (a_s, ε_s) are independent, so the matrices Y_s are independent and centred. The sub-Gaussian norm of x_s is at most $C(A^2_{\max} + \sigma^2)^{1/2}$ uniformly in s. The standard sub-Gaussian sample-covariance bound (obtainable by truncation, matrix Bernstein, and a sphere net) therefore gives

$$\Big\| \frac{1}{w} \sum_{s=t-w+1}^{t} Y_s \Big\|_{\mathrm{op}} \leq C(A^2_{\max} + \sigma^2) \left\{ \sqrt{\frac{d+\log T}{w}} + \frac{d+\log T}{w} \right\} \quad (40)$$

with conditional probability at least $1 - T^{-3}$ for each fixed endpoint. A union bound over the $T - w + 1 \leq T$ endpoints gives a single event of (conditional) probability at least $1 - T^{-2}$ on which (40) holds simultaneously for all $t \in \{w, \ldots, T\}$. Taking expectations over the path preserves the same probability bound.

Step 2: drift of the population window average. Write $\bar{\Sigma}_t := \frac{1}{w}\sum_{s=t-w+1}^{t} \Sigma_s$ and $\hat{\Sigma}_t := \frac{1}{w}\sum_{s=t-w+1}^{t} x_s x_s^\top$. Under the Gram hypothesis above, $\Sigma(S) = \lambda P_S + \sigma^2 I_n$, so for $d_{\mathrm{Gr}}(S, S') = \theta$ one has $\|P_S - P_{S'}\|_{\mathrm{op}} = \sin\theta_{\max} \leq d_{\mathrm{Gr}}(S, S')$ and

$$\|\Sigma(S) - \Sigma(S')\|_{\mathrm{op}} = \lambda \|P_S - P_{S'}\|_{\mathrm{op}} \leq \lambda\, d_{\mathrm{Gr}}(S, S') \leq C A^2_{\max}\, d_{\mathrm{Gr}}(S, S'), \quad (41)$$

using $\lambda = \Theta(A^2_{\max})$ from Assumption 1. Telescoping along the window therefore gives

$$\|\bar{\Sigma}_t - \Sigma_t\|_{\mathrm{op}} \leq \frac{1}{w} \sum_{j=0}^{w-1} \|\Sigma_{t-j} - \Sigma_t\|_{\mathrm{op}} \quad (42)$$

$$\leq C A^2_{\max} \cdot \frac{1}{w} \sum_{j=0}^{w-1} \sum_{i=0}^{j-1} \eta_{t-i} \quad (43)$$

$$\leq C A^2_{\max} \cdot \min\Big(1, \sum_{j=0}^{w-1} \eta_{t-j}\Big), \quad (44)$$

where the saturation at 1 uses $\|\Sigma - \Sigma'\|_{\mathrm{op}} = O(A_{\max}^2)$ uniformly. Under constant mean increment $\eta_t \equiv \bar{\eta}$ and $w\bar{\eta} = o(1)$, the right-hand side is $O(A_{\max}^2 w\bar{\eta})$.

Step 3: Davis–Kahan. On the high-probability event of Step 1, $\|\hat{\Sigma}_t - \bar{\Sigma}_t\|_{\mathrm{op}}$ is bounded by the right-hand side of (40). Assumption 1 provides eigengap $\lambda_{\mathrm{gap}} = \lambda$ for Σ_t; under $A_{\max}^2 w\bar{\eta} = o(\lambda_{\mathrm{gap}})$, Step 2 gives $\|\bar{\Sigma}_t - \Sigma_t\|_{\mathrm{op}} = O(A_{\max}^2 w\bar{\eta}) = o(\lambda_{\mathrm{gap}})$. Lemma 2 with $A = \Sigma_t$, $B = \bar{\Sigma}_t$, and $\delta = O(A_{\max}^2 w\bar{\eta})$ therefore yields a unique top-k eigenspace of $\bar{\Sigma}_t$ within distance $O(A_{\max}^2 w\bar{\eta}/\lambda_{\mathrm{gap}})$ of S_t whenever $A_{\max}^2 w\bar{\eta} < \lambda_{\mathrm{gap}}/2$. Wedin's theorem and Davis–Kahan (Wedin, 1972; Yu et al., 2015), applied to $(\hat{\Sigma}_t, \bar{\Sigma}_t)$ and $(\bar{\Sigma}_t, \Sigma_t)$, yield

$$\|P_{\hat{S}_t} - P_{S_t}\|_{\mathrm{op}} \le \frac{C}{\lambda_{\mathrm{gap}}}\Big(\|\hat{\Sigma}_t - \bar{\Sigma}_t\|_{\mathrm{op}} + \|\bar{\Sigma}_t - \Sigma_t\|_{\mathrm{op}}\Big). \tag{45}$$

Since $d_{\mathrm{Gr}}(\hat{S}_t, S_t) \le C\|P_{\hat{S}_t} - P_{S_t}\|_{\mathrm{op}}$ for the chosen principal-angle metric, substituting Steps 1–2 produces (5) with

$$\rho_w = \frac{A_{\max}^2 + \sigma^2}{\lambda_{\mathrm{gap}}}\left\{\sqrt{\frac{d + \log T}{w}} + \frac{d + \log T}{w}\right\}. \tag{46}$$

The window lower bound $w \gtrsim (A_{\max}^2 + \sigma^2)^2\{d + \log T\}/\lambda_{\mathrm{gap}}^2$ makes ρ_w small; with a diverging factor it gives $\rho_w = o(1)$ and Davis–Kahan applies uniformly over $t \ge w$. □

A.6 Online rates for $\hat{z}_t$ (Theorem 4)

Restatement (Thm. 4). On the event of Lemma 1, for every $t \ge w$ in the first bound and every $t \ge w + \Delta$ in the second,

$$|\hat{e}_t - e_t| \le C'\|x_t\|^2\zeta_t, \tag{47}$$

$$|\hat{v}_t - d_{\mathrm{Gr}}(S_t, S_{t-\Delta})| \le \zeta_t + \zeta_{t-\Delta}. \tag{48}$$

On one-way steady-drift Construction 1 with $|\alpha_t| \equiv A > 0$, suppose $T \to \infty$ along sequences (w_T, Δ_T, d_T) satisfying

$$\rho_w + \tfrac{A_{\max}^2}{\lambda_{\mathrm{gap}}} w\bar{\eta} = o(\bar{\eta}_\Delta^{\mathrm{patt}}), \qquad w + \Delta = o(T), \tag{49}$$

Then the two displayed mean bounds in Theorem 4 hold, the population gap survives in $\bar{\hat{v}}_T$, and the oracle threshold gives vanishing balanced risk; under rotationally invariant noise, norm-based summaries retain $\mathcal{R} = 1/2$. The regime is nonempty: for fixed d and noise constants, $w = T^{3/5}$ and $\Delta = T^{4/5}$ satisfy the display up to logarithmic factors.

Proof. We work throughout on the high-probability event of Lemma 1, so $d_{\mathrm{Gr}}(\hat{S}_t, S_t) \le \zeta_t$ for all $t \ge w$. The streaming rule waits until $t \ge w + \Delta$, when Lemma 1 covers both estimated subspaces in $\hat{v}_t$.

Projection coordinate.

$$\big|\hat{e}_t - e_t\big| = \big|\|P_{\hat{S}_t}x_t\|^2 - \|P_{S_t}x_t\|^2\big| \tag{50}$$

$$= \big|\langle P_{\hat{S}_t} - P_{S_t},\, x_t x_t^\top\rangle\big| \tag{51}$$

$$\le \|P_{\hat{S}_t} - P_{S_t}\|_{\mathrm{op}}\,\|x_t\|^2. \tag{52}$$

On $\mathrm{Gr}(k, n)$ one has $\|P_{\hat{S}} - P_S\|_{\mathrm{op}} \le C\, d_{\mathrm{Gr}}(\hat{S}, S)$, hence $|\hat{e}_t - e_t| \le C'\|x_t\|^2\zeta_t$, which is (6).

Pattern-motion coordinate. The triangle inequality on $(\mathrm{Gr}(k, n), d_{\mathrm{Gr}})$ gives

$$\big|\hat{v}_t - d_{\mathrm{Gr}}(S_t, S_{t-\Delta})\big| = \big|d_{\mathrm{Gr}}(\hat{S}_t, \hat{S}_{t-\Delta}) - d_{\mathrm{Gr}}(S_t, S_{t-\Delta})\big| \tag{53}$$

$$\le d_{\mathrm{Gr}}(\hat{S}_t, S_t) + d_{\mathrm{Gr}}(\hat{S}_{t-\Delta}, S_{t-\Delta}) \tag{54}$$

$$\le \zeta_t + \zeta_{t-\Delta}. \tag{55}$$

This bound is (7).

Separation on steady-drift (one-way) Construction 1. Under stronger forcing, S_t is constant, so $d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) = 0$ and $\zeta_t \le C\rho_w$ (no drift). Averaging (7),

$$\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{forc}} \le \frac{1}{T}\sum_{t=w+\Delta}^{T} (\zeta_t + \zeta_{t-\Delta}) = O(\rho_w). \tag{56}$$

Under one-way pattern change, write

$$\bar{\eta}_\Delta^{\mathrm{patt}} = \frac{1}{T}\sum_{t=1}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}). \tag{57}$$

On the constant-speed path, each lag-Δ increment is $O(\bar{\eta}_\Delta^{\mathrm{patt}})$, so omitting the first $w + \Delta - 1$ terms changes the T-normalised mean by only $O(((w+\Delta)/T)\bar{\eta}_\Delta^{\mathrm{patt}})$. Moreover $\zeta_t \le C\big(\rho_w + (A_{\max}^2/\lambda_{\mathrm{gap}})w\bar{\eta}\big)$ under constant mean increment, so

$$\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{patt}} \ge \frac{1}{T}\sum_{t=w+\Delta}^{T} d_{\mathrm{Gr}}(S_t, S_{t-\Delta}) - \frac{1}{T}\sum_{t=w+\Delta}^{T} (\zeta_t + \zeta_{t-\Delta}) \tag{58}$$

$$\ge \bar{\eta}_\Delta^{\mathrm{patt}} - O\big(\rho_w + \tfrac{A_{\max}^2}{\lambda_{\mathrm{gap}}} w\bar{\eta}\big) - O\big(((w+\Delta)/T)\bar{\eta}_\Delta^{\mathrm{patt}}\big). \tag{59}$$

Whenever $\rho_w + (A_{\max}^2/\lambda_{\mathrm{gap}})w\bar{\eta} = o(\bar{\eta}_\Delta^{\mathrm{patt}})$ and $w = o(T)$, the two population means are separated by $\Theta(\bar{\eta}_\Delta^{\mathrm{patt}})$; under $\bar{\eta}_\Delta^{\mathrm{patt}} = \Theta(\Delta\bar{\eta})$ the separation condition requires $\Delta \gg w$, while scalars remain at risk $\mathcal{R} = 1/2$ by Corollary 1. □

Population threshold. Under Theorem 4, the threshold $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ yields vanishing balanced classification risk $\mathcal{R} = o(1)$ whenever (8) holds.

Proof. On the tracking event of Lemma 1 (probability at least $1 - T^{-2}$), (8) from the proof of Theorem 4 gives $\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{forc}} \le \frac{\bar{\eta}_\Delta^{\mathrm{patt}}}{4}$ and $\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{patt}} \ge \frac{3\bar{\eta}_\Delta^{\mathrm{patt}}}{4}$ for all large T whenever $\rho_w + (A_{\max}^2/\lambda_{\mathrm{gap}})w\bar{\eta} = o(\bar{\eta}_\Delta^{\mathrm{patt}})$. The rule with population threshold $\tau = \bar{\eta}_\Delta^{\mathrm{patt}}/2$ therefore classifies correctly on that event; the complementary event has probability $O(T^{-2})$, hence $\mathcal{R} = o(1)$. If the leading constant in App. A.6 is $c \in (0, 1]$, the same argument uses $\tau = c\bar{\eta}_\Delta^{\mathrm{patt}}/2$. □

Null-calibration lemmas. Lemmas 5–6 bound null-stream pattern motion and give finite-sample false-positive control for a fresh null stream in Cor. 3.

Lemma 5 (Stronger-forcing-stream pattern motion on the tracking event)**.** *Under Assumption 1, suppose $S_t \equiv S$ is constant (stronger forcing). On the high-probability tracking event of Lemma 1 (probability at least $1 - T^{-2}$, uniform in $t \ge w$), for every $t \ge w$ one has $d_{\mathrm{Gr}}(\hat{S}_t, S) \le \zeta_t \le C\rho_w$ because the drift term in (5) vanishes. Consequently, for every $t \ge w + \Delta$,*

$$\hat{v}_t^{\mathrm{forc}} = d_{\mathrm{Gr}}(\hat{S}_t, \hat{S}_{t-\Delta}) \le d_{\mathrm{Gr}}(\hat{S}_t, S) + d_{\mathrm{Gr}}(\hat{S}_{t-\Delta}, S) \le \zeta_t + \zeta_{t-\Delta} \le C'\rho_w. \tag{60}$$

The running mean satisfies $\bar{\hat{v}}_t^{\mathrm{forc}} \le C'\rho_w$ for every $t \ge w + \Delta$.

Proof. When $S_t \equiv S$, $\eta_t = 0$ and the minimum in (5) vanishes, leaving only the sampling term ρ_w. The triangle inequality for d_{Gr} gives the bound on $\hat{v}_t$; averaging preserves it. □

Lemma 6 (Exchangeable upper calibration)**.** *Let $Z_1, \ldots, Z_N, Z_{N+1}$ be exchangeable real-valued null scores, fix $q \in (0, 1)$, and put $\ell = \lceil (N+1)(1-q) \rceil$. If $\ell \le N$, let $\hat{\tau} = Z_{(\ell)}$ be the ℓ-th order statistic of the first N scores; if $\ell = N + 1$, set $\hat{\tau} = +\infty$. Then*

$$\Pr\{Z_{N+1} > \hat{\tau}\} \le \frac{N + 1 - \ell}{N + 1} \le q. \tag{61}$$

The bound does not require independence within the stream that produces each score.

Proof. Break ties among the $N+1$ scores by an auxiliary independent uniform random ordering. Exchangeability makes the resulting rank of Z_{N+1} uniform on $\{1,\dots,N+1\}$. The strict event $Z_{N+1} > Z_{(\ell)}$ can occur only when this rank exceeds ℓ, hence has probability at most $(N+1-\ell)/(N+1)$. Without random tie-breaking the strict inequality can only reduce rejection probability, so the same bound holds for arbitrary ties. □

Proposition 3 (Finite-sample calibrated power and risk)**.** *Let $Z_1,\dots,Z_N$ and a fresh null Z_{N+1} be exchangeable, set $\ell = \lceil (N+1)(1-q) \rceil \le N$ and $\hat\tau = Z_{(\ell)}$, and suppose each calibration score satisfies $\Pr(Z_i \le b_T) \ge 1-\delta_{0,T}$. If a fresh pattern-change score Z^1 satisfies $\Pr(Z^1 \ge g_T) \ge 1-\delta_{1,T}$ for some $g_T > b_T$, then the strict rule has false-negative probability at most $N\delta_{0,T}+\delta_{1,T}$ and balanced risk at most*

$$\mathcal{R} \le \frac{1}{2}\left\{ \frac{N+1-\ell}{N+1} + N\delta_{0,T} + \delta_{1,T} \right\}. \tag{62}$$

For Theorem 4, one may take $b_T = C_0\rho_w$, $g_T = \bar\eta_\Delta^{\mathrm{patt}} - C_1\{\rho_w + (A_{\max}^2/\lambda_{\mathrm{gap}})w\bar\eta + ((w+\Delta)/T)\bar\eta_\Delta^{\mathrm{patt}}\}$, and $\delta_{0,T} = \delta_{1,T} = T^{-2}$; whenever $g_T > b_T$, (62) becomes $\mathcal{R} \le \frac{1}{2}\{(N+1-\ell)/(N+1) + (N+1)T^{-2}\}$.

Proof. If all N calibration scores are at most b_T, then $\hat\tau \le b_T$. On that event and $\{Z^1 \ge g_T\}$, strict rejection follows from $g_T > b_T$, so a union bound gives the false-negative term. Lemma 6 supplies the false-positive term; their average is (62). The final display substitutes the tracking-event bounds from Lemma 5 and Theorem 4. □

Corollary 4 (Finite-horizon ever-alarm calibration)**.** *Fix a finite, data-independent horizon set $\mathcal{H}$ and a score process $(Z_{i,t})_{t\in\mathcal{H}}$ for each stream, with preprocessing and $(w,\Delta,\hat k)$ fixed independently of calibration and test streams. If the N calibration score vectors and one score vector from a fresh null stream are exchangeable, set $Y_i = \max_{t\in\mathcal{H}} Z_{i,t}$ and calibrate $\hat\tau^{\max} = Y_{(\ell)}$ as in Lemma 6. Then*

$$\Pr\{\exists t \in \mathcal{H} : Z_{N+1,t} > \hat\tau^{\max}\} \le \frac{N+1-\ell}{N+1} \le q. \tag{63}$$

The bound does not require independence across horizons within a stream.

Proof. The event in (63) equals $\{Y_{N+1} > \hat\tau^{\max}\}$. Exchangeability of the score vectors implies exchangeability of their maxima, so Lemma 6 applies directly. □

Null calibration. *Restatement (Cor. 3).* For N exchangeable stronger-forcing terminal scores $Z_i = \bar v_{T,i}$ at a common horizon, Lemma 6's threshold controls a fresh null at level q. Proposition 3 bounds finite-sample power loss from calibration and tracking failures; if the false-negative probability on pattern change is β, balanced risk is at most $(q+\beta)/2$. Under Theorem 4 and (8), $\beta = o(1)$; if $q_T \to 0$, $N_T \ge q_T^{-1} - 1$, and $N_T = o(T^2)$, then balanced risk vanishes.

Proof. The false-positive bound is Lemma 6; averaging false-positive and false-negative probabilities gives $\mathcal{R} \le (q+\beta)/2$. For the asymptotic statement, Lemma 5 bounds every calibration score by $C'\rho_w$ whenever its tracking event holds. A union bound over N_T calibration streams fails with probability at most $N_T T^{-2} = o(1)$, so their finite order statistic satisfies $\hat\tau = O(\rho_w) = o(\bar\eta_\Delta^{\mathrm{patt}})$. Theorem 4 gives a fresh pattern-change score at least $\bar\eta_\Delta^{\mathrm{patt}} - o(\bar\eta_\Delta^{\mathrm{patt}})$ except with probability $O(T^{-2})$, hence $\beta = o(1)$. The condition $N_T \ge q_T^{-1} - 1$ ensures $\ell \le N_T$, and therefore $\mathcal{R} \le q_T/2 + o(1) = o(1)$. □

Overlapping lags under constant amplitude. Assume Assumption 1 with $w\bar\eta = o(1)$, a one-way Construction 1 path sampled from a fixed C^2 constant-speed curve of bounded curvature, constant amplitude $|\alpha_t| \equiv \alpha > 0$, and isotropic signal Gram ($\mathbb{E}[aa^\top] = \lambda I_k$). For any fixed ratio $\Delta/w \in (0,1]$, separation still holds when $\rho_w = o(\bar\eta_\Delta)$. Trailing-window SVD tracks the window-averaged top space $\bar S_t$, so the $w\bar\eta$ remainder in (8) does not enter.

Proof. Write $\Sigma_s = \mathbb{E}[x_s x_s^\top \mid S_s]$ and $\bar{\Sigma}_t = \frac{1}{w}\sum_{s=t-w+1}^{t} \Sigma_s$, with top-$k$ eigenspace $\bar{S}_t$. Under constant amplitude $|\alpha_t| \equiv \alpha > 0$ and Construction 1 (rank-1), or more generally when the signal Gram is isotropic on the active subspace ($\mathbb{E}[aa^\top] = \lambda I_k$), one has $\Sigma_s = \lambda P_{S_s} + \sigma^2 I_n$ with $\lambda = \Theta(A_{\max}^2)$ independent of s, so

$$\bar{\Sigma}_t = \lambda \cdot \frac{1}{w} \sum_{s=t-w+1}^{t} P_{S_s} + \sigma^2 I_n, \tag{64}$$

and the top-k space of $\bar{\Sigma}_t$ coincides with that of the unweighted projector average whenever the eigengap persists.

Step 1: short-arc midpoint. Under Assumption 1 with constant mean increment and $w\bar{\eta} = o(1)$, the path segment $\{S_s\}_{s=t-w+1}^{t}$ is a short arc of length $O(w\bar{\eta}) = o(1)$ on $\mathrm{Gr}(k, n)$. Fix the temporal midpoint index $t_{\mathrm{mid}} := t - \lfloor w/2 \rfloor$ and set $S_{\mathrm{mid}} := S_{t_{\mathrm{mid}}}$. On a neighbourhood of radius $o(1)$ the Riemannian exponential $\mathrm{Exp}_{S_{\mathrm{mid}}}$ is a diffeomorphism onto its image and $\mathrm{Log}_{S_{\mathrm{mid}}}$ is bi-Lipschitz (Absil et al., 2008). Write $V_s := \mathrm{Log}_{S_{\mathrm{mid}}}(S_s) \in \mathrm{T}_{S_{\mathrm{mid}}}\mathrm{Gr}(k, n)$. The exponential-chart Taylor expansion along a constant-speed curve, with the curvature remainder controlled on the compact set of projectors, yields

$$V_s = (s - t_{\mathrm{mid}})\, H + R_s, \qquad \|H\| = \Theta(\bar{\eta}), \qquad \|R_s\| \le C(s - t_{\mathrm{mid}})^2 \bar{\eta}^2, \tag{65}$$

whenever $|s - t_{\mathrm{mid}}|\bar{\eta} = o(1)$ (Absil et al., 2008, Ch. 5). Averaging,

$$\frac{1}{w}\sum_{s=t-w+1}^{t} V_s = H \cdot \frac{1}{w}\sum_{s=t-w+1}^{t} (s - t_{\mathrm{mid}}) + O\big((w\bar{\eta})^2\big). \tag{66}$$

The discrete mean of $(s - t_{\mathrm{mid}})$ over a window of length w centred at t_{mid} is $O(1)$ (exactly 0 for odd w; $O(1)$ for even w), hence the linear contribution is $O(\bar{\eta})$. The projector expansion $P_{\mathrm{Exp}(V)} = P_{S_{\mathrm{mid}}} + \mathrm{sym}(V) + O(\|V\|^2)$ and the bound $\|V_s\| = O(w\bar{\eta})$ therefore give

$$\Big\|\frac{1}{w}\sum_{s=t-w+1}^{t} P_{S_s} - P_{S_{\mathrm{mid}}}\Big\|_{\mathrm{op}} \le \Big\|\mathrm{sym}\Big(\tfrac{1}{w}\sum_{s} V_s\Big)\Big\|_{\mathrm{op}} + \frac{1}{w}\sum_{s} O(\|V_s\|^2) = O\big(\bar{\eta} + (w\bar{\eta})^2\big). \tag{67}$$

Under a fixed ratio $r = \Delta/w \in (0, 1]$ with $w \to \infty$ (hence $\Delta \to \infty$), one has $\bar{\eta}_\Delta = \Theta(\Delta\bar{\eta}) = \Theta(w\bar{\eta})$, hence

$$\frac{\bar{\eta} + (w\bar{\eta})^2}{\bar{\eta}_\Delta} = O\Big(\frac{1}{w} + w\bar{\eta}\Big) = o(1). \tag{68}$$

Davis–Kahan about the unique top space of $\bar{\Sigma}_t$ (gap λ_{gap} by Assumption 1 and $w\bar{\eta} = o(1)$) therefore yields

$$d_{\mathrm{Gr}}\big(\bar{S}_t,\, S_{\mathrm{mid}}\big) = O\big(\bar{\eta} + (w\bar{\eta})^2\big). \tag{69}$$

Step 2: pattern motion of window-averaged patterns. Under constant amplitude the window weights are uniform, so the midpoints satisfy $t_{\mathrm{mid}} - (t - \Delta)_{\mathrm{mid}} = \Delta$. The triangle inequality and Step 1 give

$$d_{\mathrm{Gr}}(\bar{S}_t, \bar{S}_{t-\Delta}) \ge d_{\mathrm{Gr}}\big(S_{t_{\mathrm{mid}}},\, S_{t_{\mathrm{mid}}-\Delta}\big) - O\big(\bar{\eta} + (w\bar{\eta})^2\big). \tag{70}$$

Along one-way Construction 1, the midpoint lag equals $\Theta(\Delta\bar{\eta}) = \Theta(\bar{\eta}_\Delta^{\mathrm{patt}})$ by Lemma 3 and the one-way argument of App. A.2, so there exists $c \in (0, 1]$ such that for all large t,

$$d_{\mathrm{Gr}}(\bar{S}_t, \bar{S}_{t-\Delta}) \ge c\,\bar{\eta}_\Delta. \tag{71}$$

Step 3: estimation error relative to $\bar{S}_t$. On the high-probability event of App. A.5 Step 1, $\|\hat{\Sigma}_t - \bar{\Sigma}_t\|_{\mathrm{op}}$ is bounded by the right-hand side of (40). Davis–Kahan about $\bar{\Sigma}_t$ (eigengap λ_{gap}) therefore gives, on the tracking event of App. A.5,

$$d_{\mathrm{Gr}}(\hat{S}_t, \bar{S}_t) \le \frac{C}{\lambda_{\mathrm{gap}}}\|\hat{\Sigma}_t - \bar{\Sigma}_t\|_{\mathrm{op}} = O(\rho_w), \tag{72}$$

without invoking the instantaneous-S_t form of Lemma 1. The triangle inequality yields

$$\big|\hat{v}_t - d_{\mathrm{Gr}}(\bar{S}_t, \bar{S}_{t-\Delta})\big| \le O(\rho_w). \tag{73}$$

Step 4: separation. Under stronger forcing, $\bar{S}_t$ is constant, so $\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{forc}} = O(\rho_w)$. Under one-way pattern change, Steps 2–3 yield

$$\frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{forc}} = O(\rho_w), \qquad \frac{1}{T}\sum_{t=w+\Delta}^{T} \hat{v}_t^{\mathrm{patt}} \geq c\,\bar{\eta}_\Delta - O(\rho_w). \tag{74}$$

Whenever $\rho_w = o(\bar{\eta}_\Delta)$, the two terminal means separate without an explicit $w\bar{\eta}$ remainder. □

Long-window suppression on recurrent paths. Under constant amplitude $|\alpha_t| \equiv \alpha > 0$ and a unique mixture top space with fixed eigengap, consider either strict period-2 Π_{sym} with $\ell := 1$ or Π_{RT} with excursion duration $\ell \ll w$. The estimator of §5 satisfies

$$\mathbb{E}[\hat{\bar{v}}_T^{\mathrm{patt}}] = O(\rho_w + \ell/w). \tag{75}$$

Thus a window much longer than a recurrent excursion can suppress pattern motion even when pointwise motion is nonzero.

Proof. We use the dimension-aware ρ_w from Lemma 1. Assume constant amplitude $|\alpha_t| \equiv \alpha > 0$ (as in Theorem 1). Suppose first that the pattern-change path lies in strict period-2 $\Pi_{\mathrm{sym}}(\delta_{\mathrm{sym}}, L)$ with odd Δ, and set $\ell = 1$. A trailing window of length w then contains alternating visits to the two states. Writing $n_\pm$ for the number of visits to $S_\pm$ in the window,

$$\big|n_+ - n_-\big| \leq \ell \qquad \text{and} \qquad n_+ + n_- = w, \tag{76}$$

so

$$\frac{1}{w}\sum_{s=t-w+1}^{t} P_{S_s} = \frac{n_+}{w}P_{S_+} + \frac{n_-}{w}P_{S_-} = M + O\Big(\frac{\ell}{w}\Big) \tag{77}$$

in operator norm, where $M = \frac{1}{2}(P_{S_+} + P_{S_-})$ as in App. A.4. For rank 1, the eigengap of M between its top eigenvalue $\frac{1+\cos\delta_{\mathrm{sym}}}{2}$ and the next $\frac{1-\cos\delta_{\mathrm{sym}}}{2}$ equals $\cos\delta_{\mathrm{sym}} > 0$ for $\delta_{\mathrm{sym}} \in (0, \pi/2)$; for $k > 1$ we use the unique top-k gap of M from Rem. 4. Davis–Kahan about the unique top space $S_\star$ of M therefore yields $d_{\mathrm{Gr}}(\bar{S}_t, S_\star) = O(\ell/w)$ (with implied constant depending on that gap). The same bound holds at time $t - \Delta$, so $d_{\mathrm{Gr}}(\bar{S}_t, \bar{S}_{t-\Delta}) = O(\ell/w)$. On the high-probability event of App. A.5 Step 1, $\|\hat{\Sigma}_t - \bar{\Sigma}_t\|_{\mathrm{op}}$ is bounded by (40); Davis–Kahan about $\bar{\Sigma}_t$ (which remains gapped about $S_\star$ for $\ell \ll w$) gives $d_{\mathrm{Gr}}(\hat{S}_t, \bar{S}_t) = O(\rho_w)$ without invoking the slow-drift hypothesis of Lemma 1. Hence on that event

$$\hat{v}_t = d_{\mathrm{Gr}}(\hat{S}_t, \hat{S}_{t-\Delta}) \leq O(\rho_w + \ell/w), \tag{78}$$

and taking expectations (using $d_{\mathrm{Gr}} = O(1)$ on the complementary $O(T^{-2})$-probability event) yields

$$\mathbb{E}\big[\hat{\bar{v}}_T^{\mathrm{patt}}\big] = O(\rho_w + \ell/w). \tag{79}$$

For a leave-and-return path in $\Pi_{\mathrm{RT}}(L, m)$ whose nontrivial excursion has duration $\ell \ll w$, write $E \subset \{1, \ldots, T\}$ for the (at most ℓ-long) set of times at which $S_t \neq S_1$. For any window ending at t that is disjoint from E, one has $\bar{\Sigma}_t = \alpha^2 P_{S_1} + \sigma^2 I_n$ and $\bar{S}_t = S_1$. If the window intersects E, at most ℓ of the w summands differ from P_{S_1}, so

$$\Big\|\frac{1}{w}\sum_{s=t-w+1}^{t} P_{S_s} - P_{S_1}\Big\|_{\mathrm{op}} = O\Big(\frac{\ell}{w}\Big) \tag{80}$$

and Davis–Kahan gives $d_{\mathrm{Gr}}(\bar{S}_t, S_1) = O(\ell/w)$. The same holds at $t - \Delta$, hence $d_{\mathrm{Gr}}(\bar{S}_t, \bar{S}_{t-\Delta}) = O(\ell/w)$ for every t. Combining with the Bernstein–Davis–Kahan bound $d_{\mathrm{Gr}}(\hat{S}_t, \bar{S}_t) = O(\rho_w)$ as above yields $\mathbb{E}[\hat{\bar{v}}_T^{\mathrm{patt}}] = O(\rho_w + \ell/w)$. □

B MONITOR ALGORITHMS

Under equal norms, stronger forcing and pattern change share one field-norm path, so norm-based summaries cannot separate them, and even-length batch endpoints can fail when the pattern leaves and returns or hops between two states (Table 1, §3–4). We first state the norm-based and batch statistics from Defs. 2 and 4. We then give the trailing-window estimator and null-calibrated mean pattern-motion rule from §5. Throughout, $X = (x_t)_{t=1}^{T}$ with $x_t \in \mathbb{R}^n$, rank parameter $\hat{k} \geq 1$, and Grassmann metric d_{Gr} as in §2. Experimental defaults and the CORAL protocol are in App. C.

B.1 SHARED PRIMITIVES

Windowed leading subspace. For a contiguous index set $B \subset \{1, \dots, T\}$, form the sample covariance $\hat{\Sigma}_B := \frac{1}{|B|} \sum_{t \in B} x_t x_t^\top$ (or, equivalently, take the leading right singular vectors of the data matrix with rows $\{x_t : t \in B\}$). Write $\mathrm{SVD}_{\hat{k}}(\hat{\Sigma}_B)$ for an orthonormal basis of the top-$\hat{k}$ eigenspace and $\hat{S}_B := \mathrm{span}\big(\mathrm{SVD}_{\hat{k}}(\hat{\Sigma}_B)\big) \in \mathrm{Gr}(\hat{k}, n)$. For a trailing window of length w ending at time t, write $\hat{\Sigma}_t := \frac{1}{w} \sum_{s=t-w+1}^{t} x_s x_s^\top$ and $\hat{S}_t := \mathrm{span}\big(\mathrm{SVD}_{\hat{k}}(\hat{\Sigma}_t)\big)$ whenever $t \geq w$.

Grassmann distance. If $\hat{U}, \hat{V} \in \mathbb{R}^{n \times \hat{k}}$ are orthonormal bases for subspaces S, S', the principal angles θ_i satisfy $\cos \theta_i = \sigma_i(\hat{U}^\top \hat{V})$ with σ_i the singular values of $\hat{U}^\top \hat{V}$, and $d_{\mathrm{Gr}}(S, S')^2 = \sum_{i=1}^{\hat{k}} \theta_i^2$ as in the main text. For $\hat{k} = 1$ this reduces to $d_{\mathrm{Gr}}(\mathrm{span}(u), \mathrm{span}(v)) = \arccos |\langle u, v \rangle|$ with unit vectors u, v.

Monitor score. Each monitor m returns a scalar $s_m(X) \in \mathbb{R}_{\geq 0}$. On an equal-norm pair we treat larger values as evidence of pattern change relative to stronger forcing; the online rule of §5 thresholds a running mean instead of a single batch scalar.

B.2 NORM-BASED MONITORS

These are the norm-based summaries of Def. 2 (batch statistic (a) in Def. 4). Under Construction 1 they match across the two equal-norm readings (Prop. 1), so these monitors do not separate the two readings.

Field norm. Set

$$s_{\mathrm{fn}}(X) := \frac{1}{T} \sum_{t=1}^{T} \|x_t\|^2. \tag{81}$$

Any measurable functional of the path $\big(\|x_t\|^2\big)_{t \leq T}$ alone is likewise a field-norm monitor; the mean is the batch score used for ranking.

Fixed-dictionary residual. Fix a reference dictionary $R_\star = \mathrm{span}(r)$ with unit $r \in \mathbb{R}^n$ (e.g. a frozen leading EOF direction). The residual energy outside $R_\star$ is $r_t(R_\star) = \|(I - rr^\top)x_t\|^2$, and the batch score is

$$s_{\mathrm{res}}(X; R_\star) := \frac{1}{T} \sum_{t=1}^{T} r_t(R_\star). \tag{82}$$

This is the fixed-dictionary residual in Def. 2 and Prop. 1. Table 2 labels this monitor *Fixed residual*: under Construction 1 it is residual energy outside the frozen EOF axis r. The heatwave table's *prefix-fit residual* is the related prefix-fit score: estimate $\hat{R}$ on an early window, then take mean suffix energy outside $\hat{R}$. Because $\hat{R}$ is data-dependent, that score falls outside Def. 2 and Prop. 1. It is not the projected early/late angle.

B.3 BATCH-ENDPOINT AND BLOCK MONITORS

These explicit early/late, break, and adjacent-window proxies implement the even-length batch statistics (b)–(d) of Def. 4, motivated by the procedures of Xu et al. (2011); Koo et al. (2026); Xie et al. (2018); Kanai et al. (2023). Theorem 1 identifies their hard cases: the proxies separate under steady drift but fail on the stated leave-and-return or equal-hopping classes, as in Table 1.

Early/late angle. Fix a window length $m \geq 2$ (Def. 4 takes even m). Estimate $\hat{S}_{\mathrm{early}}$ from $\{x_1, \dots, x_m\}$ and $\hat{S}_{\mathrm{late}}$ from $\{x_{T-m+1}, \dots, x_T\}$, and set

$$s_{\mathrm{el}}(X; m) := d_{\mathrm{Gr}}\big(\hat{S}_{\mathrm{early}}, \hat{S}_{\mathrm{late}}\big). \tag{83}$$

On $\Pi_{\mathrm{RT}}(L, m)$, Thm. 1(i) gives risk $1/2$ for any such early/late functional. The $\triangle$ on equal hopping in Table 1 marks odd early/late windows outside the even-m batch statistics of Def. 4 (Rem. 1).

Projected early/late angle. Fix a reference axis r and project $x_t^{\perp} := (I - rr^{\top})x_t$. Apply the early/late procedure of App. B.3 to the projected stream $(x_t^{\perp})_{t\le T}$ with even window m:

$$s_{\mathrm{lel}}(X; r, m) := s_{\mathrm{el}}\big((x_t^{\perp})_{t\le T};\, m\big). \tag{84}$$

This is the early/late angle in the orthogonal complement of a fixed dictionary axis. It is undefined when either projected block covariance lacks a leading eigengap. In particular, CORAL's hop stronger-forcing branch lies on the removed mixture axis, so Table 2 marks this proxy degenerate on hop rows rather than assigning an arbitrary SVD direction.

Fixed-break angle. Fix a break fraction $b \in (0,1)$ (or a finite grid $\mathcal{G} \subset (0,1)$). Split the record at $t_b := \lfloor bT \rfloor$ into pre and post blocks with even lengths in the equal-hopping hard case of Thm. 1(ii). Estimate $\hat{S}_{\mathrm{pre}}(b)$ and $\hat{S}_{\mathrm{post}}(b)$ from those blocks and set

$$s_{\mathrm{br}}(X; b) := d_{\mathrm{Gr}}\big(\hat{S}_{\mathrm{pre}}(b),\, \hat{S}_{\mathrm{post}}(b)\big). \tag{85}$$

On a grid, the batch score is $s_{\mathrm{br}}(X; \mathcal{G}) := \max_{b\in\mathcal{G}} s_{\mathrm{br}}(X; b)$ (or the mid-break $b = \frac{1}{2}$ alone).

Neighbouring-block angle. Fix an even block length $w_{\mathrm{nb}} \ge 2$, kept distinct from the trailing-window length w of App. B.4. Thm. 1(iii) and App. A.4 analyse adjacent length-w_{nb} blocks ending at t and starting at $t{+}1$ (whenever both lie in $\{1,\dots,T\}$), with score $\chi_t(w_{\mathrm{nb}}) = d_{\mathrm{Gr}}(S_t^{\mathrm{L}}, S_t^{\mathrm{R}})$ and batch functional any measurable map of $\{\chi_t(w_{\mathrm{nb}})\}_t$. A contiguous partition of $\{1,\dots,T\}$ forms non-overlapping length-w_{nb} blocks $B_0, B_1, \dots$; we discard a trailing remainder shorter than w_{nb}. The monitor averages adjacent angles

$$s_{\mathrm{nb}}(X; w_{\mathrm{nb}}) := \frac{1}{J}\sum_{j=0}^{J-1} d_{\mathrm{Gr}}\big(\hat{S}_{B_j},\, \hat{S}_{B_{j+1}}\big), \tag{86}$$

where J is the number of adjacent pairs and $\hat{S}_{B_j}$ is the leading subspace of block B_j. On Π_{sym} with even w_{nb}, both the sliding adjacent windows of App. A.4 and this contiguous average vanish when the leading space is the equal-weight mixture.

B.4 Mean pattern motion

This is the online estimator of §5: we track $\hat{S}_t$, form lag-Δ Grassmann increments, average them over the record, and (optionally) threshold against stronger-forcing nulls. Where the batch statistics above compare subspaces at ends or blocks, this score records pattern motion along the full record (row (3) of Table 1). Corollary 3 gives vanishing risk for the null-calibrated rule on steady drift; for leave-and-return and partial equal hopping, we use the same score outside that guarantee, as in §5.

Trailing-window tracking. For $t \ge w$, set $\hat{S}_t = \mathrm{span}\big(\mathrm{SVD}_{\hat{k}}(\hat{\Sigma}_t)\big)$ with trailing covariance $\hat{\Sigma}_t$ as in App. B.1. For $t < w$, leave $\hat{S}_t$ undefined (or copy $\hat{S}_w$). Under Assumption 1, Lemma 1 bounds $d_{\mathrm{Gr}}(\hat{S}_t, S_t)$ for this trailing-window SVD tracker.

Lag-Δ motion and running mean. Fix lag $\Delta \ge 1$. For every $t \ge w + \Delta$ with both $\hat{S}_t$ and $\hat{S}_{t-\Delta}$ defined, set

$$\hat{v}_t := d_{\mathrm{Gr}}\big(\hat{S}_t,\, \hat{S}_{t-\Delta}\big), \qquad \bar{\hat{v}}_t := \frac{1}{t - w - \Delta + 1}\sum_{s=w+\Delta}^{t} \hat{v}_s. \tag{87}$$

As in §5, we do not divide population v_t in (2) by Δ. For equal-norm ranking, we use the terminal mean $s_{\mathrm{mpm}}(X; w, \Delta) := \hat{v}_T$. A ranking table may report a per-step speed $\hat{v}_t/\Delta$ instead; ranking and null calibration must use one consistent scale (App. C.5). Optionally record $\hat{e}_t = \|P_{\hat{S}_t} x_t\|^2$; it equals the squared projection energy on $\hat{S}_t$ but does not enter the decision threshold of §5.

Computational cost. The reference implementation recomputes a truncated SVD of each $w \times d$ window, costing $O\{Twd\min(w,d)\}$ time and $O(wd)$ working memory; in CORAL, $d = 3 \le w$, so this is $O(Twd^2)$. Each Grassmann comparison additionally takes an SVD of a $\hat{k} \times \hat{k}$ cross-product. Rank-one covariance updates or an incremental tracker can avoid rebuilding every window, but our reported experiments use the direct SVD implementation.

Null-calibrated decision rule. Fix $(w, \Delta, \hat{k})$ on independent tuning data, then hold out a disjoint pool $\mathcal{N}$ of stronger-forcing calibration streams ($S_t \equiv S$). For level q, set $\ell = \lceil (|\mathcal{N}| + 1)(1 - q) \rceil$ and, at each horizon t, let $\hat{\tau}_t$ be the ℓ-th order statistic of the null running means $\{\bar{v}_t^{\mathrm{forc}}\}_{\mathcal{N}}$ (or $+\infty$ if $\ell > |\mathcal{N}|$). Declare *pattern change* when $\bar{v}_t > \hat{\tau}_t$ and *stronger forcing* otherwise; abstain if $t < w + \Delta$. Cor. 3 gives finite-sample false-positive control for a fresh null stream under stream-level exchangeability and vanishing balanced risk under its steady-drift asymptotics. The procedure is well-defined beyond steady drift. Vanishing risk holds only on steady drift.

Short-window option (GROUSE). When w is too short for a stable trailing SVD, replace $(\hat{S}_t)$ by a GROUSE (Balzano et al., 2010) tracker. Maintain an orthonormal frame $U_t \in \mathbb{R}^{n \times \hat{k}}$, take the residual of x_t outside $\mathrm{span}(U_t)$, and apply a geodesic step of stepsize $\gamma_{\mathrm{G}} > 0$ toward that residual to obtain U_{t+1}. Set $\hat{S}_t = \mathrm{span}(U_t)$ and reuse $\hat{v}_t$ and $\bar{v}_t$ from App. B.4. This substitutes the subspace path only; the decision score remains mean lag-Δ pattern motion. Lemma 1 and Cor. 3 analyse trailing-window SVD, not GROUSE, so §5 does not give online rates or null-calibrated risk transfer for the GROUSE path.

C CORAL CONSTRUCTION AND PROTOCOL

CORAL (Construction on Real Atlantic Latents) is our labelled suite on CMEMS North Atlantic SST (Szekely et al., 2026). We freeze a leading EOF frame on demeaned and detrended anomalies, embed Construction 1 drift and leave-and-return signals plus an equal-norm hopping variant whose forcing branch is the mixture top space, and score the App. B monitors on forcing and pattern-change branches. We specify the data, atlas, Construction embedding, path schedules, noise, and scores used in Fig. 2, Table 2, and the unlabelled heatwave comparison.

C.1 DATA AND ANOMALY ATLAS

Source and domain. We use monthly CMEMS CORA surface temperature on a fixed North Atlantic mask (Szekely et al., 2026) with n=1529 ocean grid cells (basin extent roughly $80°$W–$10°$W, $20°$N–$60°$N). We work with a seasonal-anomaly field on that mask (climatological monthly means removed before freeze-window processing).

Atlas eras. We build the atlases from two non-overlapping freeze windows: 80s = Jan 1980–Dec 1999 and 00s = Jan 2000–Dec 2019, each with T_{era}=240 months. Within each window we (i) column-demean, (ii) remove the spatial mean at each time, and (iii) remove a linear trend in time at each grid cell. The processed panel $X^{\mathrm{era}} \in \mathbb{R}^{T_{\mathrm{era}} \times n}$ is the sole input to the EOF freeze for that era.

Frozen EOF frame. On X^{era} we take a thin SVD and freeze a rank-3 frame. The leading right singular vector r is EOF1, sign-normalised so that its largest absolute entry is positive. We Gram–Schmidt orthonormalise EOF2 and EOF3 against r and each other to obtain (q_1, q_2). Write $B := [r\ q_1\ q_2] \in \mathbb{R}^{n \times 3}$. The residual bank is the freeze-window field orthogonal to $\mathrm{span}(B)$,

$$R^{\mathrm{era}} := X^{\mathrm{era}} - X^{\mathrm{era}} B B^{\top}. \tag{88}$$

We compute these residuals row-wise and store each row for CORAL-R bootstrap noise (App. C.4). We freeze the frame (r, q_1, q_2) and residual bank before drawing any Construction path and never re-estimate them on synthetic trials.

C.2 EQUAL-NORM LIFT ONTO THE FRAME

Cone geometry. Fix cone angle ψ=30° and unit amplitude $|\alpha_t| \equiv \alpha$=1. A direction on the constant-angle cone about r is

$$u(\phi) = \cos\psi\, r + \sin\psi \big(\cos\phi\, q_1 + \sin\phi\, q_2\big), \qquad \phi \in \mathbb{R}, \tag{89}$$

so $|\langle u(\phi), r \rangle| = \cos\psi$ for every ϕ and Construction 1 residual matching to $R_\star = \mathrm{span}(r)$ holds on the noiseless signal.

Matched pair. On each trial we draw $\phi_0 \sim \mathrm{Unif}[0, 2\pi)$ and a path class $\Pi \in \{\Pi_{\mathrm{drift}}, \Pi_{\mathrm{RT}}, \Pi_{\mathrm{sym}}\}$ (App. C.3). On the pattern-change branch we set $u_t = u(\phi_t)$ along the scheduled angles (ϕ_t) and $x_t^{\mathrm{patt}} = \alpha_t u_t$ (we keep the sign continuous along smooth paths, and we do not enforce that continuity on period-2 hops). On the stronger-forcing branch we fix a unit $u_\star$: for drift and leave-and-return, $u_\star = u(\phi_0)$; for equal hopping, $u_\star$ is a unit leading direction of $\frac{1}{2}\big(P_{\mathrm{span}(u(\phi_0))} + P_{\mathrm{span}(u(\phi_0+\pi))}\big)$ (Thm. 1(ii)–(iii)). On this cone geometry the hop mixture top space is $\mathrm{span}(r)$, so the hop forcing branch is aligned with EOF1 rather than at angle ψ. Noiseless fields are vectors in $\mathbb{R}^n$ on the full grid and form an equal-norm pair: $\|x_t^{\mathrm{forc}}\|^2 = \|x_t^{\mathrm{patt}}\|^2 = \alpha_t^2$ for every t before noise. We then add independent noise to each branch (App. C.4), so the *observed* streams are no longer pathwise equal-norm; the Construction labels and Prop. 1 identities refer to the noiseless signals.

Record length. Synthetic records use $T{=}480$ (twice one freeze-era length), so early/late and break windows fit inside the record.

C.3 Path schedules

The three schedules below are the Table 1 path classes of Def. 3 on the frozen cone: steady drift, leave-and-return, and equal hopping. We express all angles through the cone parameter ϕ and measure pattern motion by $d_{\mathrm{Gr}}(\mathrm{span}(u(\phi)), \mathrm{span}(u(\phi')))$.

Steady drift (Π_{drift}). For times $t = 0, \ldots, T{-}1$, we set

$$\phi_t = \phi_0 + 2\pi\,\nu\,\frac{t}{T}, \qquad \nu{=}0.5 \tag{90}$$

(half a turn over the record; the grid excludes the endpoint so the path does not close).

Leave-and-return (Π_{RT}). Hold $\phi_t \equiv \phi_0$ outside a central excursion window $[t_0, t_1)$ with $t_0 = \lfloor 0.35\,T \rfloor$ and $t_1 = \lfloor 0.65\,T \rfloor$. Inside the window, ϕ rises linearly by $\pi/3$ then returns to ϕ_0. With the scoring window $w{=}36$ of App. C.5, the first and last w observations lie in the plateaus, so early/late subspaces both equal $S_1 = \mathrm{span}(u(\phi_0))$ on both branches and Thm. 1(i) applies.

Equal hopping (Π_{sym}). Period-2 alternation on the same $t = 0, \ldots, T{-}1$ index,

$$\phi_t = \begin{cases} \phi_0, & t \text{ even}, \\ \phi_0 + \pi, & t \text{ odd}, \end{cases} \tag{91}$$

so the two cone states are antipodal in the (q_1, q_2)-plane and every even-length block has equal occupancy (Def. 3).

C.4 Noise protocols

We set protocol noise level $\rho{=}0.2$ for a typical noise-to-signal ratio $\|\varepsilon_t\|/|\alpha| \approx \rho$. Four conditions cross era $\times$ noise type—CORAL-80s-I, CORAL-80s-R, CORAL-00s-I, CORAL-00s-R. The isotropic condition tests in-frame isotropic noise; the residual condition verifies that projection onto the frozen frame rejects frame-orthogonal residual variation and is not a test of in-frame correlated-noise robustness.

Isotropic (I). Draw $\varepsilon_t \sim \mathcal{N}(0, \sigma^2 I_n)$ independently with $\sigma = \rho|\alpha|/\sqrt{n}$, so $\mathbb{E}[\|\varepsilon_t\|^2] = \rho^2\alpha^2$. After projection into the frozen $d = 3$ scoring frame, the expected noise-to-signal norm ratio is $\rho_{\mathrm{score}} = \rho\sqrt{d/n}$. At the ranking table's $\rho = 0.2$, this ratio is approximately 0.0089, so the table tests geometric separation at low projected noise rather than robustness.

Frame-orthogonal residual (R). Draw rows i.i.d. with replacement from the frozen residual bank R^{era}, then apply a global scale c so that the mean squared norm of scaled draws equals $\rho^2\alpha^2$ (energy-matched to the isotropic law). Every residual draw is already orthogonal to $\mathrm{span}(B)$, so scaled residual draws remain orthogonal to $\mathrm{span}(B)$. We use independent draws on the two matched branches. Because scoring projects onto $\mathrm{span}(B)$, residual noise vanishes exactly from the scored streams; this condition tests rejection of variation outside the scoring frame.

Elevated noise for rates. For Fig. 4, we use ambient $\rho \in \{1.2, 8, 12, 16\}$ on the same CORAL-80s-I drift construction. These values correspond to projected noise ratios approximately $\{0.053, 0.354, 0.532, 0.709\}$, so the trailing-window estimate meets the noise floor. The ranking table in Table 2 remains at ambient ρ=0.2.

C.5 Scoring protocol

Coordinate system. We evaluate monitors after projecting each branch onto the frozen leading-3 subspace $\mathrm{span}(B)$ (coefficients in the (r, q_1, q_2) basis, then lifted back for App. B routines). We keep scores in the same frozen atlas as the constructed signals, rather than re-estimating the EOF frame on each trial. All subspace monitors in the ranking table use $\hat{k} = 1$ within this $d = 3$ representation. Dictionary residuals and projected early/late use the frozen EOF1 axis r as $R_\star$.

Windows and lags. We set the default trailing and batch window to w=36 and the default lag to Δ=12. On hop rows only, we set the lag to Δ=11 (the odd lag adjacent to the annual lag) and use even neighbouring-block length w_{nb}=120. This contiguous non-overlapping average from App. B.3 is the mixture hard case of Thm. 1(iii) and App. A.4. Early/late and fixed-break use the even lengths specified in Def. 4 and App. B.

Monitor map. Table 2 lists the App. B scores: field norm; fixed-dictionary residual; early/late and projected early/late angles; fixed mid-break (b=1/2); neighbouring-block average; terminal mean pattern motion s_{mpm}. On hop rows, we mark fixed residual and projected early/late as degenerate. The stronger-forcing branch lies in $\mathrm{span}(r)$, the mixture top space, so fixed residual is not matched across branches; removing that same axis leaves its projected covariance without a population eigengap, so projected early/late has no defined population direction.

Mean pattern motion scale. Section 5 and App. B.4 define $\hat{v}_t = d_{\mathrm{Gr}}(\hat{S}_t, \hat{S}_{t-\Delta})$ and $s_{\mathrm{mpm}} = \bar{\hat{v}}_T$ without dividing by Δ. On the CORAL ranking table we record a per-step speed $\hat{v}_t/\Delta$ and average that series after a burn-in of order $w + \Delta$; that average equals s_{mpm}/Δ on the same path, so paired ranking is unchanged. Ranking and any null calibration stay on this one consistent scale on both branches. Population plots and the theorems use d_{Gr} without dividing by Δ.

Paired ranking accuracy. This is the labelled test of Table 1 recovery in frozen ocean-EOF coordinates in §6; orthogonal invariance makes the coordinate score distribution era-invariant. For each monitor m and each of N=50 independent trials we obtain scores s_m^{forc} and s_m^{patt} on the matched pair. Paired ranking accuracy is the fraction of trials with $s_m^{\mathrm{patt}} > s_m^{\mathrm{forc}}$ (numerical ties count as half); chance is 50%, and 0% means s_m is always larger on the stronger-forcing branch. We report each Table 2 cell as mean±Monte Carlo s.e. of this bounded paired-ranking indicator across the 50 trials; zero s.e. means no variation occurred in those trials, not zero population uncertainty.

Held-out calibration experiment. Figure 4(c) uses the exact leading-3 coordinate law of CORAL-80s-I, with calibration and test draws kept disjoint. At each noise level, 100 independent repetitions use $N = 49$ stronger-forcing streams to set the $\ell = \lceil 50(1 - 0.10) \rceil = 45$ order-statistic threshold of Cor. 3. In each repetition, 10 fresh stronger-forcing and 10 fresh steady-drift streams estimate marginal false-positive rate and power; we form percentile-bootstrap intervals over repetitions. We divide the plotted score by Δ on every stream, which leaves decisions and the order-statistic guarantee unchanged.

C.6 Window, lag, and rank sensitivity

Figure 6 varies one estimator parameter at a time on 580 shared CORAL-80s-I steady-drift pairs at elevated noise ρ=12 from Fig. 4(c). We fixed the protocol defaults before varying the parameters. Twenty disjoint repetitions each use 19 calibration nulls (permitting exact level-0.10 calibration) and 10 fresh streams per class. At correct rank, rejection on fresh null streams stays 4.5–10.5% as the window and the lag vary, while power rises from 12.5 to 100% with w and from 36 to 97% with Δ. At default $(w, \Delta, \hat{k})$=(36, 12, 1), paired accuracy is 96.7% and power 79.5%; $\hat{k} = 2$ reduces them to 54.8% and 10.0%, while the degenerate full-coordinate rank gives chance ranking and zero power. This degradation illustrates the known-rank assumption in Lemma 1.

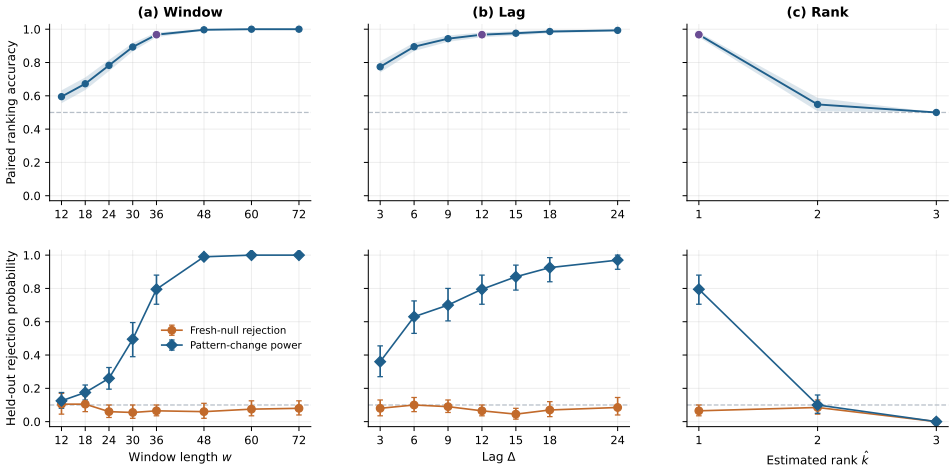

Figure 6: Sensitivity of mean pattern motion on CORAL-80s-I steady drift at ρ=12 (580 shared paired trials; one parameter varied at a time). Top row: paired ranking accuracy with paired-trial bootstrap 95% bands (dashed line = chance; purple points = protocol defaults). Bottom row: rejection on held-out null streams and pattern-change power over 20 disjoint repetitions (bootstrap 95% intervals; dashed line = $q = 0.10$). **(a)** Longer windows clear the tracking floor. **(b)** Separation is stable across nearby lags. **(c)** Rank overestimation removes the rank-1 signal; $\hat{k} = 3$ spans the full coordinate space and is therefore degenerate.

C.7 Recurrent paths under structured in-frame noise

Figure 7 removes the low projected-noise advantage of the main ranking table by adding noise directly in the frozen leading-3 coordinates, where projection cannot discard it. We draw anisotropic AR(1) noise $z_t = 0.85 z_{t-1} + \sqrt{1 - 0.85^2} L \xi_t$ with $LL^\top = \Sigma$, $\mathrm{tr}(\Sigma) = 1$, and eigenvalues $(0.70, 0.22, 0.08)$ in a fixed non-axis-aligned basis. The stationary condition uses $\epsilon_t = \rho_{\text{score}} z_t$; the shifted condition multiplies the first and second record halves by 0.65 and 1.35, respectively, then normalises the scale sequence to unit root-mean-square. Thus ρ_{score} is the expected root-mean-square in-frame noise-to-signal ratio in both conditions.

For each path, noise law, and $\rho_{\text{score}} \in \{0.10, 0.25, 0.50, 0.75\}$, 20 disjoint repetitions use 19 calibration nulls and 10 fresh streams per class at level $q = 0.10$; calibration and test draws share the stated nuisance law but are independent. Besides mean pattern motion, we test a Grassmann-jump CUSUM heuristic inspired by Subspace-CUSUM (Xie et al., 2018), defined as the maximum cumulative sum of adjacent-window Grassmann jumps; we also test a maximum sliding two-window angle and the fixed mid-break angle, all at $w = 36$ (mean motion uses $\Delta = 12$, or 11 for hopping). Mean pattern motion has the highest point estimate in all eight hopping cells, with power 64–99%; the Grassmann-jump CUSUM gives 63.5–95%, sliding two-window 34.5–77.5%, and fixed break 22–36.5%. For leave-and-return, the three sequential statistics have full power at $\rho_{\text{score}} = 0.10$, but none retains useful power once $\rho_{\text{score}} \geq 0.50$; at 0.25, the best method depends on the noise condition. Across all 64 method–cell combinations, realised fresh-null rejection is 5–17%, and 59 bootstrap intervals cover the nominal 10% level.

C.8 Finite-horizon alarms and calibration shift

For the fixed post-burn-in horizons $\mathcal{H} = \{w + \Delta, \ldots, T\}$, we also score each stream by the maximum of its running mean and apply Cor. 4 to the $N = 49$ calibration-null maxima. Across the four matched-noise settings of Fig. 4(c), ever-alarm rejection on fresh null streams is 9.1–11.0%. Power is 100%, 45.7%, 23.2%, and 14.9%. Simultaneous control lowers power relative to the terminal decisions.

Figure 8 deliberately breaks exchangeability by applying each repetition's threshold calibrated at the row's noise level to fresh nulls at the column's noise level. The diagonal reproduces matched calibration; low-noise terminal thresholds applied to noisier nulls reject 98.3–100%, while the ever-alarm maximum rejects 41.9–100%. Reverse shifts are conservative (0–0.6%). This is a stress test,

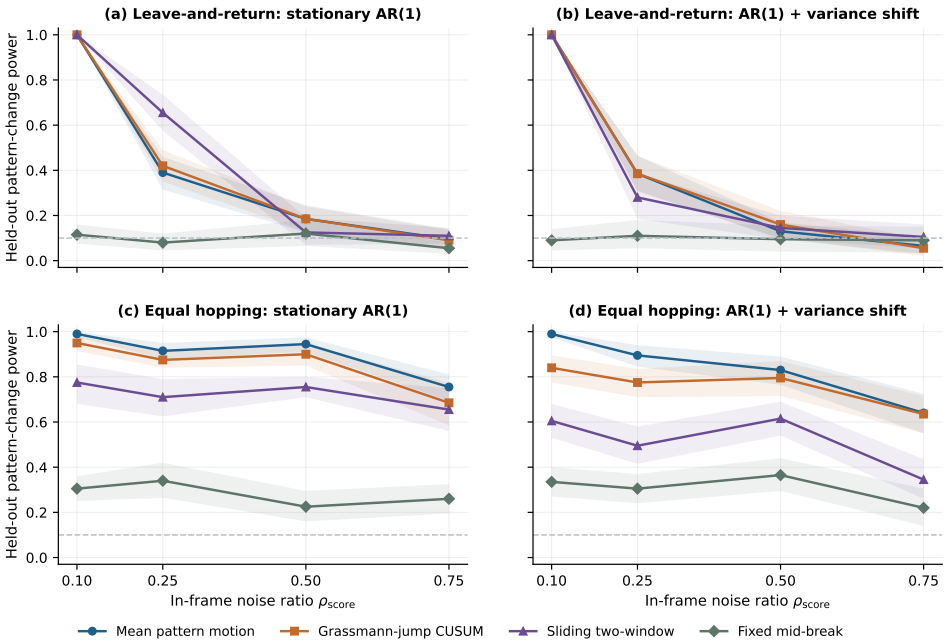

Figure 7: Held-out recurrent-path power under structured noise in the frozen leading-3 CORAL coordinates (20 disjoint calibration/test repetitions; bands are repetition-bootstrap 95% intervals; dashed line $= q = 0.10$). **(a,b)** Leave-and-return loses detectability as in-frame noise obscures the finite excursion. **(c,d)** Mean pattern motion retains the highest point estimate on equal hopping, including under the prescribed mid-record variance shift. The left column uses stationary anisotropic AR(1) noise; the right adds the variance shift.

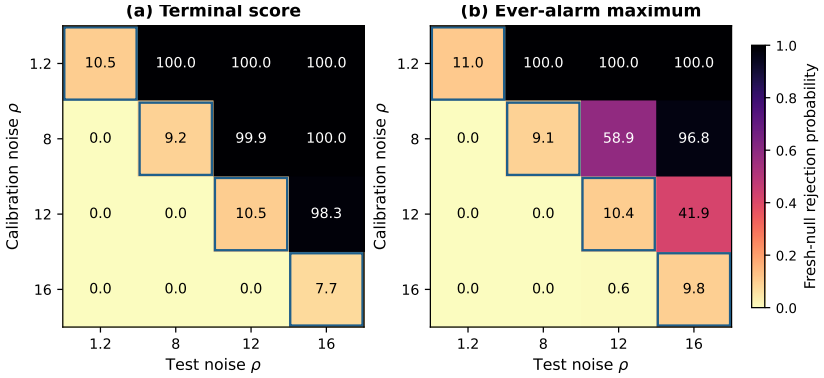

Figure 8: Calibration under null-law shift (100 repetitions; entries are rejection percentages on fresh null streams). Rows give the calibration noise and columns the test noise; blue boxes mark exchangeable matched-noise cells. **(a)** Terminal running-mean score. **(b)** Maximum running mean over the fixed post-burn-in horizon set. Off-diagonal cells intentionally violate exchangeability and quantify failure rather than guaranteed coverage.

not a coverage claim: Corollaries 3 and 4 do not apply off the diagonal. Cross-era CORAL-I would not produce this shift because, conditional on each frozen atlas, its constructed leading-3 coordinate law is era-invariant; changing noise scale produces a null-law shift.

C.9 Unlabelled heatwave comparison

The Jul 2023–Jun 2024 North Atlantic marine heatwave analysis in §6 is not a Construction trial: there are no matched labels. We apply the same demean–detrend preprocessing and App. B monitors to the observed CMEMS record in the CORAL basin, with shared $(w, \Delta)=(36, 12)$. We report continuous scores as heatwave-to-2010s ratios, with descriptive month-resampling ranges rather than dependence-valid coverage intervals, and length-matched twelve-month 2010s baseline blocks

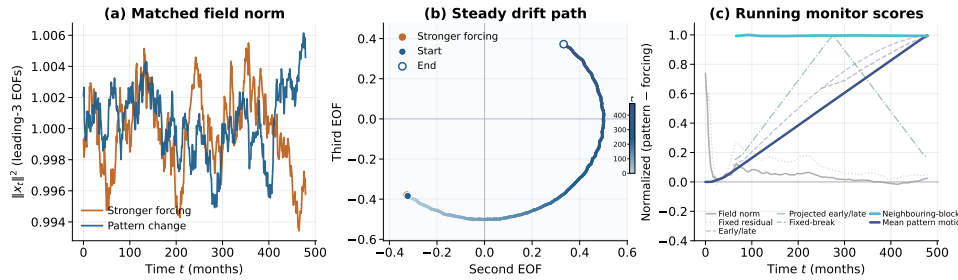

Figure 9: Steady drift on CORAL-80s-I (one matched pair; same frozen atlas as Fig. 2). **(a)** $\|x_t\|^2$ in the leading-3 EOF subspace on both branches (isotropic noise). **(b)** Pattern-change path in the second–third EOF plane (colour = time t); first EOF is the fixed cone axis. **(c)** Normalised pattern−forcing score differences for the Table 2 monitors over time. Mean pattern motion and batch angles separate through the record; the field-norm score difference stays near zero.

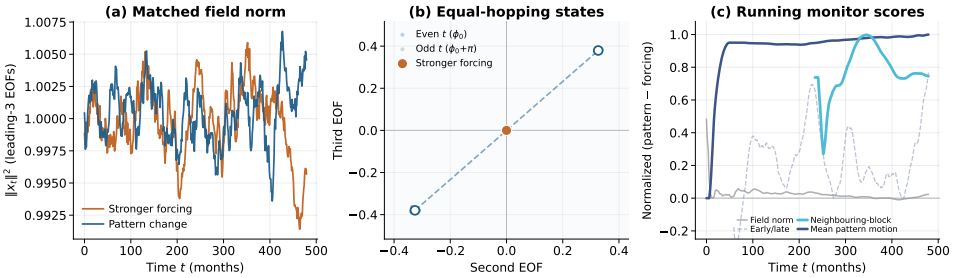

Figure 10: Isotropic equal hopping on CORAL-80s-I (one matched pair; odd lag Δ=11; same frozen atlas as Fig. 2). **(a)** $\|x_t\|^2$ in the leading-3 EOF subspace on both branches (isotropic noise). **(b)** Even/odd samples in the second–third EOF plane: pattern change hops between antipodal cone states, while stronger forcing lies on the EOF1 axis and therefore at the origin of this plane. **(c)** Running scores for field norm, early/late, neighbouring-block (w=120), and mean pattern motion. On this isotropic pair, the hop contrast is the antipodal pair in **(b)**, and mean pattern motion separates in **(c)** while the field-norm score difference stays near zero. Table 2 shows that hopping is only partly separated across conditions.

(n=109). Field norm is the batch norm. For the prefix-fit residual of App. B.2, we fit the leading mode on the 2010s prefix and measure suffix energy outside it; this score is distinct from the projected early/late angle. We measure mean pattern motion in rad per step. We measure angle monitors as cross-era principal angles (Def. 4) against within-2010s references. OISST (Huang et al., 2021) confirms basin warmth only; it does not enter the CORAL latent construction. Table 3 therefore records monitor disagreement under a shared protocol, not labelled recovery of Table 1.

D Additional CORAL figures

Figure 3 in §6 shows one labelled leave-and-return pair on CORAL-80s-I. Figures 9 and 10 show steady drift and isotropic equal hopping under the same frozen atlas, construction, and isotropic-noise condition. Each display contains matched field norms, cone geometry in the EOF2–EOF3 plane, and running monitor gaps under the §6 protocol.