PREDICTABLE WASSERSTEIN CRITICS FOR TWO-SAMPLE TESTING

Anonymous authors
Paper under double-blind review

ABSTRACT

A Wasserstein critic learns distributional differences that respect a chosen metric. Reusing past observations to train successive critics preserves out-of-sample evaluation, but leaves a random variance scale that can invalidate Gaussian studentization. We resolve this difficulty for a finite Lipschitz dictionary by letting entropy regularization vanish more slowly than training noise. The learned critic then has vanishing null variance, while retaining the dictionary's population discrepancy under alternatives. Adding independent random signs with positive amplitude gives a studentized Gaussian test under a $(2+\delta)$-moment null. The same fresh-pair scores also admit a closed-form half-normal e-process, valid at arbitrary stopping times for any predictable critic without moment assumptions. Approximate empirical Kantorovich–Rubinstein optimization recovers full or restricted population drift under first moments and fixed-alternative consistency under second moments, without requiring the fitted functions to converge. Simulations show earlier detection than online Newton-step betting in the studied alternatives, and a stratified penguin analysis illustrates learned projection critics.

1 INTRODUCTION

Two-sample inference asks whether two populations have the same law. For example, Sommerfeld and Munk (2018) compare real and synthetic fingerprints using distances and directional differences between ridge features, calibrating the significance of their distributional differences. Wasserstein–1 distance measures the cost of transporting one distribution to the other. Its dual representation turns this comparison into learning a scalar function, or *critic*, whose variation is bounded by the metric. This interpretation underlies Wasserstein generative adversarial networks (Arjovsky et al., 2017); for testing, the additional task is to distinguish genuine population separation from adaptation to sampling noise.

Let P, Q be laws on a metric space $(\mathcal{X}, d)$. We test $H_0 : P = Q$ against $H_1 : P \neq Q$. For finite first moments, Kantorovich–Rubinstein (KR) duality gives

$$W_1(P, Q) = \sup_{\mathrm{Lip}(f)\leq 1} \{Pf - Qf\}, \tag{1}$$

where $Pf = \mathbb{E}_{X\sim P} f(X)$ and $\mathrm{Lip}(f) \leq 1$ means $|f(x) - f(y)| \leq d(x, y)$. A large empirical optimum need not be a large out-of-sample mean difference. Direct empirical transport statistics can also have distribution-dependent, non-Gaussian null limits, and permutation calibration may repeat the transport computation for each relabelling (Ramdas et al., 2017; Sommerfeld and Munk, 2018). These difficulties motivate learning the direction of separation on past observations and evaluating it on a fresh pair.

Learning a witness for a distributional discrepancy is well established. Wasserstein distance and kernel maximum mean discrepancy (MMD) are integral probability metrics (IPMs), specified by their test-function classes (Müller, 1997; Sriperumbudur et al., 2012). Kernel and energy tests supply widely used comparisons (Gretton et al., 2012; Székely and Rizzo, 2013; Sejdinovic et al., 2013), while classifier and deep-kernel tests adapt the representation to the alternative (Kirchler et al., 2020; Liu et al., 2020). The witness test obtains a studentized Gaussian statistic from a learned kernel witness on held-out data (Kübler et al., 2022); cross-MMD combines splitting with studentization and establishes local minimax guarantees on Sobolev classes (Shekhar et al., 2022).

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.

Past-only learning also permits successive evaluation without a permanent training/test split. Testing by betting chooses predictable IPM witnesses and accumulates anytime-valid evidence, including for unbounded function classes through conditional symmetry (Shekhar and Ramdas, 2024). Learned representations can be used in this sequential setting (Pandeva et al., 2024). Most closely on the Gaussian side, the martingale kernel two-sample test uses past kernel witnesses and fresh-pair increments to obtain asymptotic normality (Chatterjee and Ramdas, 2026). For Wasserstein critics, the central issue is how the geometric learning objective interacts with the accumulated variance.

Transport-specific methods address different parts of this problem. Kernel-projected and sliced discrepancies reduce the cost of comparing distributions (Wang et al., 2022). Sliced-Wasserstein testing has permutation-based power and minimax results on specified classes (Tran and Schreuder, 2026). Neural projection-Wasserstein testing uses sample splitting and Gaussian-maximum calibration (Hu and Lin, 2026), and the Radon–Kolmogorov–Smirnov test studies neural witnesses for another IPM (Paik et al., 2025). Here we study an accumulating scalar score from full or restricted one-Lipschitz critics. Our guarantees identify when its learned population target is preserved and when its variance permits studentization. Table 4 summarizes the data-use and calibration distinctions.

Our contributions connect a concrete variance-stabilizing critic learner, population separation, and time-uniform testing for the same scores. Fresh evaluation removes conditional mean bias under the null, while Gaussian studentization requires a stable variance scale. For a fixed Lipschitz dictionary, entropy regularization averages nearly tied empirical witnesses; a temperature that decays more slowly than training noise drives the critic's null variance to zero while sending its optimization error to zero under alternatives. Metric-based perturbations and independent random-sign noise then leave a nonrandom limiting variance, including at point-mass nulls. This extends the use of perturbations for degenerate studentized statistics (Corradi et al., 2024) to a concrete adaptive critic construction.

The resulting dictionary test has a Gaussian null limit under a $(2+\delta)$ moment and retains the dictionary's maximum absolute mean contrast under alternatives. Two further results connect this construction to more general Wasserstein learning. First, approximate empirical KR optimization transfers to population drift under first moments, with second moments giving fixed-alternative consistency; the bound separates training, optimization, and class-approximation errors without requiring convergence of the fitted functions. Second, conditional symmetry yields a closed-form half-normal e-process for the same scores. This monitoring rule has finite-sample validity without moments or variance stabilization, so it also accommodates full KR and restricted projection critics whose Gaussian calibration remains conditional.

An eight-dimensional learned dictionary links the variance theorem to observable behavior and compares two valid sequential tests on identical scores. A tuning study shows how stronger regularization concentrates null variance while delaying detection. Exact one-dimensional transport critics reveal both gains and losses against permutation tests, and a stratified penguin analysis illustrates the symmetry-based route when a Gaussian limit has not been established. Together, these results distinguish learning a separating critic from calibrating the evidence it produces.

Section 2 constructs the scores, Section 3 gives the variance and drift guarantees, Sections 4 and 5 report the experiments and the penguin application, and Section 6 concludes; the appendix gives implementation details, further results, and proofs.

2 Predictable Critic Scores

2.1 Past-only learning and fresh-pair evaluation

Throughout, $(\mathcal{X}, d)$ is Polish (complete and separable), with its Borel σ-field. Fix $x_0 \in \mathcal{X}$. For $r \geq 1$, let $\mathcal{P}_r(\mathcal{X})$ denote laws with finite rth moment of $d(X, x_0)$; this condition is independent of the base point. We use $\mathcal{F}_1 = \{f : f(x_0) = 0,\ \mathrm{Lip}(f) \leq 1\}$, so that $|f(x)| \leq d(x, x_0)$; for $P, Q \in \mathcal{P}_1$, this normalization does not change the supremum in (1). Let $X_1, X_2, \ldots \overset{\text{i.i.d.}}{\sim} P$ and $Y_1, Y_2, \ldots \overset{\text{i.i.d.}}{\sim} Q$, with the two sequences independent. The information $\mathcal{G}_{i-1}$ contains all observations and training and perturbation randomness revealed before pair i.

A *predictable critic* $f_i \in \mathcal{F}_1$ is measurable with respect to $\mathcal{G}_{i-1}$, with the evaluation σ-field on functions. It scores the fresh pair before either observation enters subsequent training. Set $f_1 = 0$. At deterministic update times, train on the first $r_i < i$ observations from each sample and otherwise

retain the last critic. We take r_i nondecreasing, with $r_i \to \infty$, and write $\widehat{P}_i^{\mathrm{tr}} = r_i^{-1} \sum_{j \le r_i} \delta_{X_j}$ and $\widehat{Q}_i^{\mathrm{tr}} = r_i^{-1} \sum_{j \le r_i} \delta_{Y_j}$. Periodic critic updates are covered by $r_i = t_k - 1$ for $t_k \le i < t_{k+1}$.

To supply variance even when the learned critic is nearly constant, we add centered perturbations. The default is the signed pair distance $H_i^{\mathrm{D}} = R_i d(X_i, Y_i)$, where R_i is an independent Rademacher sign (± 1 with equal probability); it costs one distance evaluation and adds variance at every non-point null with finite second moments. A second independent Rademacher sign E_i supplies external jitter, which remains variable even when both observations equal a single point. With deterministic bounded scales $\tau_i, \rho_i \ge 0$, define

$$\xi_i = f_i(X_i) - f_i(Y_i) + \tau_i H_i^s + \rho_i E_i, \qquad S_n = \sum_{i \le n} \xi_i, \quad V_n = \sum_{i \le n} \xi_i^2. \tag{2}$$

Here $s = \mathrm{D}$ by default; the anchor-based alternative $s = \mathrm{A}$ is defined in Section A.2. All current signs and the fresh pair are independent of $\mathcal{G}_{i-1}$. The perturbations preserve the population mean difference and keep the score in the metric's units. Set $T_n = S_n / \sqrt{V_n}$, with $T_n = 0$ when $V_n = 0$.

2.2 A REGULARIZED DICTIONARY AND TRANSPORT TARGETS

Fix deterministic $h_1, \dots, h_K \in \mathcal{F}_1$, and estimate their contrasts by $\widehat{d}_{ij} = (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) h_j$. Entropy regularization over the positive and negative dictionary atoms gives

$$f_i(x) = \frac{\sum_{j=1}^K \sinh(\widehat{d}_{ij}/\vartheta_i) h_j(x)}{\sum_{j=1}^K \cosh(\widehat{d}_{ij}/\vartheta_i)}, \qquad \vartheta_i > 0. \tag{3}$$

This is a convex combination of the $2K$ functions $\{\pm h_j\}$, so the fitted function is globally one-Lipschitz. Its coefficients are learned from past data. As temperature ϑ_i decreases, it concentrates on empirical separation; a larger temperature averages nearly tied directions. The choice $\vartheta_i = r_i^{-1/4}$ makes the temperature vanish more slowly than null training fluctuations. Coefficients and temperature remain fixed until the next update, and $f_i = 0$ during initialization. Proposition 3.2 shows that this choice stabilizes null variance and preserves the dictionary discrepancy.

Full empirical transport learning instead seeks, for deterministic $\varepsilon_i \downarrow 0$,

$$(\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) f_i \ge W_1(\widehat{P}_i^{\mathrm{tr}}, \widehat{Q}_i^{\mathrm{tr}}) - \varepsilon_i. \tag{4}$$

A finite dual linear program on the training points, followed by the McShane extension, gives a measurable exact optimizer (Lemma D.2). For a possibly random, predictable class $\mathcal{H}_i \subseteq \mathcal{F}_1$ containing zero, the corresponding requirement is

$$(\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) f_i \ge \sup_{f \in \mathcal{H}_i} (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) f - \varepsilon_i, \qquad f_i \in \mathcal{H}_i. \tag{5}$$

All random-class suprema and selections are assumed measurable, automatically so for finite classes and under the usual separability and selection conditions for larger classes. The population supremum is an IPM for symmetric $\mathcal{H}_i$ and a directional discrepancy otherwise. In a symmetric class, selecting the best of $f, -f, 0$ on past data ensures a nonnegative training objective; optimization accuracy still requires (5). The dictionary uses the convex hull of $\{\pm h_j\}$ and has gap at most $\vartheta_i \log(2K)$.

2.3 FIXED-TIME AND TIME-UNIFORM TESTS

At a prespecified horizon n and level $\alpha \in (0, 1)$, reject when $T_n > z_{1-\alpha}$, the standard normal quantile. For the regularized dictionary, fixed $\tau \ge 0$ and $\rho > 0$ yield asymptotic level α under the moment conditions in Proposition 3.2. For other learners, Gaussian calibration additionally requires predictable-variance stabilization, as specified in Theorems C.1 and C.2.

For continuous monitoring, use an *e-process*: nonnegative evidence starting at one, whose probability of ever reaching $1/\alpha$ is at most α under the null (Shafer et al., 2011; Howard et al., 2021). Mixing $\exp\{\lambda S_n - \lambda^2 V_n/2\}$ over a half-normal law on $\lambda \ge 0$ with prespecified scale $\eta > 0$ gives

$$E_n^+(\eta) = \frac{2}{\sqrt{1 + \eta^2 V_n}} \exp\left\{ \frac{\eta^2 S_n^2}{2(1 + \eta^2 V_n)} \right\} \Phi\left\{ \frac{\eta S_n}{\sqrt{1 + \eta^2 V_n}} \right\}, \tag{6}$$

where Φ is the standard normal distribution function. Reject at its first crossing of $1/\alpha$, or use an average over prespecified scales. Theorem 3.3 permits every predictable critic, including full KR critics, and allows $\tau_i = \rho_i = 0$. Both tests need only (S_n, V_n) after scoring.

Multiplying the metric, critics, perturbations, and temperature by $c > 0$ leaves T_n unchanged. For monitoring, also set $\eta \mapsto \eta/c$ to preserve the e-process (Proposition C.10); in particular, external jitter ρ_i must scale with the metric. Algorithm 1 summarizes the procedure; Section A and Algorithm 2 give the implementation and numerical evaluation.

Algorithm 1 Predictable Wasserstein critic test

Require: Paired stream, metric, critic learner, deterministic critic update times, stabilizer and scales, level α, and a fixed horizon or a prespecified mixture over η.

Set $f_1 = 0$ and $S_0 = V_0 = 0$.
for each incoming pair $i = 1, 2, \ldots$ **do**
  At an update time with $i \geq 2$, train and orient f_i on past pairs; otherwise retain the last critic.
  Observe (X_i, Y_i), draw independent signs, and form ξ_i by (2).
  Update $S_i = S_{i-1} + \xi_i$ and $V_i = V_{i-1} + \xi_i^2$.
  For monitoring, stop and reject when the mixture of (6) reaches $1/\alpha$; at a fixed horizon n, reject if $T_n > z_{1-\alpha}$.
end for

3 Variance Stabilization and Population Drift

3.1 A learner that makes Gaussian calibration feasible

Under $H_0 : P = Q = \pi$, let $q_i = \mathbb{E}(\xi_i^2 \mid \mathcal{G}_{i-1})$ and $A_n^2 = \sum_{i \leq n} q_i$ denote the conditional and accumulated predictable variances. For $Z, Z' \overset{\text{i.i.d.}}{\sim} \pi$, the default perturbation has variance

$$\kappa_\pi^{\mathrm{D}} = \mathbb{E}\{d(Z, Z')^2\}. \tag{7}$$

For the anchor alternative, κ_π^{A} is given in (19). Either is positive precisely at non-point nulls whenever the moments are finite.

Proposition 3.1. *Under $H_0 : P = Q = \pi$, for either stabilizer $s \in \{\mathrm{A}, \mathrm{D}\}$, the conditional law of ξ_i is symmetric without moment assumptions. If $\pi \in \mathcal{P}_1(\mathcal{X})$, then*

$$\mathbb{E}(\xi_i \mid \mathcal{G}_{i-1}) = 0, \qquad \mathcal{L}(\xi_i \mid \mathcal{G}_{i-1}) = \mathcal{L}(-\xi_i \mid \mathcal{G}_{i-1}). \tag{8}$$

If $\pi \in \mathcal{P}_2(\mathcal{X})$, moreover,

$$q_i = 2 \operatorname{Var}_\pi\{f_i(Z)\} + \tau_i^2 \kappa_\pi^s + \rho_i^2. \tag{9}$$

Consequently,

$$A_n^2 \geq b_{n,\pi,s}^2 := \kappa_\pi^s \sum_{i=1}^n \tau_i^2 + \sum_{i=1}^n \rho_i^2. \tag{10}$$

Equation (9) separates the variation due to the learned critic, the signed distance, and the independent sign. Only the first term depends on the training history. The sum A_n^2 records variance predicted before each observation, whereas V_n records the squared scores actually observed. The martingale central limit argument (Hall and Heyde, 1980, Chapter 3) requires a deterministic reference scale a_n^2 with $A_n^2/a_n^2 \xrightarrow{\mathbb{P}} 1$ and $a_n \to \infty$. Learning can make subsequent variances depend on earlier scores; dividing by $\sqrt{V_n}$ need not remove that dependence (Section C.4). On a bounded metric space, this condition and $b_{n,\pi,s}^2 \to \infty$ suffice. On an unbounded space with $\pi \in \mathcal{P}_{2+\delta}$, we also require $n/(b_{n,\pi,s}^2)^{1+\min(\delta,2)/2} \to 0$ so that accumulated variance dominates large-score fluctuations. Theorems C.1 and C.2 give the complete statements and justify replacing A_n^2 by V_n.

The regularized dictionary verifies this variance condition through its training rule.

Proposition 3.2 (Calibration and drift of the regularized dictionary)**.** *Use* (3) *with deterministic training sizes $r_i < i$, $r_i \to \infty$, $\vartheta_i \to 0$, and $r_i \vartheta_i^2 \to \infty$. Fix $\tau_i = \tau \geq 0$ and $\rho_i = \rho > 0$. Under $P = Q = \pi \in \mathcal{P}_{2+\delta}$ for some $\delta > 0$,*

$$\frac{A_n^2}{n} \xrightarrow{\mathbb{P}} \tau^2 \kappa_\pi^s + \rho^2 > 0, \qquad T_n \xrightarrow{d} N(0,1).$$

For $P, Q \in \mathcal{P}_1$, $(P-Q)f_i \to \max_{j\le K}|(P-Q)h_j|$ almost surely. If this limit is positive and $P, Q \in \mathcal{P}_2$, then $T_n \xrightarrow{\mathbb{P}} +\infty$.

Two bounds explain why this construction works. Write $\|f\|^2_{L^2(\pi)} = \int f^2 \,\mathrm{d}\pi$. Under $P = Q = \pi \in \mathcal{P}_2$, taking expectation over the training samples gives

$$\mathbb{E}_\pi \|f_i\|^2_{L^2(\pi)} \le \frac{2\mathbb{E}_\pi d(Z, x_0)^2 \sum_{j=1}^K \mathrm{Var}_\pi\{h_j(Z)\}}{r_i \vartheta_i^2}. \tag{11}$$

For $P, Q \in \mathcal{P}_1$, the empirical optimization bound gives

$$0 \le \max_{j\le K}|(P-Q)h_j| - (P-Q)f_i \le 2\max_{j\le K}|\widehat{d}_{ij} - (P-Q)h_j| + \vartheta_i \log(2K). \tag{12}$$

Under the null, the empirical contrasts fluctuate at order $r_i^{-1/2}$. A larger temperature gives almost equal weights to h_j and $-h_j$, which cancel; (11) quantifies the resulting small critic variance. Under an alternative, (12) separates estimation error from the entropy penalty $\vartheta_i \log(2K)$, and sending both to zero recovers the largest dictionary contrast. Thus $r_i^{-1/2} \ll \vartheta_i \ll 1$ meets both objectives; for $\vartheta_i = r_i^{-1/4}$, the null bound is of order $r_i^{-1/2}$. Positive independent-sign noise leaves a nonzero limiting variance even at a point mass. Section E.1 proves both bounds.

Fixed $\rho > 0$ gives $b^2_{n,\pi,s} \ge n\rho^2$, verifying the variance-growth requirement for every $\delta > 0$. The Lipschitz envelope transfers the null's $(2+\delta)$ moment to the scores. Finite initialization and periodic critic updates preserve the argument when the most recent training size diverges.

3.2 Anytime-valid inference

The e-process in (6) uses the conditional-symmetry construction of de la Peña et al. (2004, Lemma 1.4). For each $\lambda \ge 0$, the process $L_n(\lambda) = \exp\{\lambda S_n - \lambda^2 V_n/2\}$ has a one-step multiplicative factor with conditional expectation $\mathbb{E}\{\cosh(\lambda\xi_i)\exp(-\lambda^2\xi_i^2/2) \mid \mathcal{G}_{i-1}\} \le 1$. Conditional symmetry pairs positive and negative scores, and $\cosh(u) \le \exp(u^2/2)$ bounds their average factor by one even for arbitrarily large scores. Mixing over $\lambda \ge 0$ therefore gives evidence against positive drift with time-uniform type-I error control, as in testing by betting (Shekhar and Ramdas, 2024).

Theorem 3.3 (Finite-sample anytime-valid test)**.** *Under $H_0 : P = Q$, for every fixed $\eta > 0$, $(E_n^+(\eta))_{n\ge 0}$ is a nonnegative supermartingale with $E_0^+(\eta) = 1$. Therefore*

$$\mathbb{P}_{H_0}\left\{\sup_{n\ge 1} E_n^+(\eta) \ge \frac{1}{\alpha}\right\} \le \alpha. \tag{13}$$

The result requires neither a moment assumption nor divergence of V_n.

Stratified streams. The same e-process proof permits a prespecified sequence of strata: conditional on the past and the design, the fresh X_i and Y_i may be independent draws from a common stratum-specific null law π_i; because the proof uses conditional symmetry one score at a time, the laws need not agree across i. This extension does not automatically carry over to the fixed-horizon Gaussian theorems, whose variance conditions must be checked for the resulting array.

Under an alternative with $S_n/n \xrightarrow{\mathbb{P}} \mu > 0$ and $V_n/n \xrightarrow{\mathbb{P}} \nu > 0$, the evidence grows at exponential rate $\mu^2/(2\nu)$; see Proposition C.9.

3.3 Learning population separation

Under finite first moments, the predictable mean is $m_i = \mathbb{E}(\xi_i \mid \mathcal{G}_{i-1}) = (P-Q)f_i$. If its average tends to $\mu > 0$ and $\sup_i \mathbb{E}\xi_i^2 < \infty$, then $S_n/n \xrightarrow{\mathbb{P}} \mu$ and $T_n \xrightarrow{\mathbb{P}} +\infty$ (Theorem C.8). The moment bound follows from $P, Q \in \mathcal{P}_2$ and bounded stabilizer scales. Thus learning must provide sustained population separation; the following bounds show when empirical optimization achieves it.

Full Wasserstein critics. A large training objective is useful only if it transfers to a large population mean difference. For $P, Q \in \mathcal{P}_1(\mathcal{X})$ and a critic satisfying (4), the loss in drift obeys the pathwise bound

$$0 \le W_1(P,Q) - m_i \le 2\{W_1(\widehat{P}_i^{\mathrm{tr}}, P) + W_1(\widehat{Q}_i^{\mathrm{tr}}, Q)\} + \varepsilon_i.$$

The same transport error bounds empirical-to-population discrepancies for every one-Lipschitz function, including the critic selected from those data. It is paid once for the selected critic and once for the best population value; ε_i is the optimization error. Empirical W_1 convergence makes these training terms vanish under first moments (Villani, 2009, Theorem 6.9); see (45) and (48).

Theorem 3.4. *Let $P, Q \in \mathcal{P}_1(\mathcal{X})$. Suppose $r_i \to \infty$, $\varepsilon_i \to 0$, and f_i satisfies* (4)*. Then $m_i = (P-Q)f_i \to W_1(P,Q)$ almost surely; in particular,*

$$n^{-1}\sum_{i=1}^{n} m_i \longrightarrow W_1(P,Q) \qquad \text{a.s.} \tag{14}$$

If additionally $P, Q \in \mathcal{P}_2(\mathcal{X})$ and $P \neq Q$, then

$$n^{-1}S_n \xrightarrow{\mathbb{P}} W_1(P,Q), \qquad T_n \xrightarrow{\mathbb{P}} +\infty. \tag{15}$$

The critic's population value converges even when the fitted functions keep changing. Quantitative training error still inherits empirical-W_1 rates, which depend on moments and geometry and can be slow in high dimension (Fournier and Guillin, 2015).

Restricted and neural critics. For a possibly random predictable class $\mathcal{H}_i \subseteq \mathcal{F}_1$, define $D_{\mathcal{H}_i}(P,Q) = \sup_{f \in \mathcal{H}_i}(P-Q)f$.

Theorem 3.5. *Let $P, Q \in \mathcal{P}_1(\mathcal{X})$. Suppose* (5) *holds, $r_i \to \infty$, $\varepsilon_i \to 0$, and*

$$D_{\mathcal{H}_i}(P,Q) \longrightarrow D(P,Q) > 0 \qquad \text{a.s.} \tag{16}$$

Then $m_i \to D(P,Q)$ almost surely. If $P, Q \in \mathcal{P}_2(\mathcal{X})$, $T_n \xrightarrow{\mathbb{P}} +\infty$. If instead

$$W_1(P,Q) - D_{\mathcal{H}_i}(P,Q) \longrightarrow 0 \qquad \text{a.s.}, \tag{17}$$

then $m_i \to W_1(P,Q)$ almost surely.

The divergence in Theorem 3.5, like (15), gives fixed-horizon rejection probability tending to one under the stated alternatives. To obtain an asymptotically level-α Gaussian test, one must also verify the null conditions in Theorems C.1 and C.2; Proposition 3.2 does so for the regularized dictionary. For full KR or other restricted critics, finite-sample anytime type-I control follows separately from Theorem 3.3, without a Gaussian variance-stabilization claim.

Restricted optimization has the same training-error bound for $D_{\mathcal{H}_i}(P,Q) - m_i$. Relative to the full Wasserstein target, it incurs the additional approximation gap $W_1(P,Q) - D_{\mathcal{H}_i}(P,Q)$. These terms identify distinct ways to improve drift: growing training samples make empirical error vanish, better optimization reduces ε_i, and a richer class can reduce the approximation gap. A fixed certified neural architecture therefore supplies its class discrepancy, with full-W_1 drift recovered when (17) holds. Proofs of the drift and power results are in Section F.

Vanishing perturbations preserve root-n fluctuations under fixed alternatives, and local triangular arrays link cumulative drift to power under their variance and remainder conditions; see Proposition C.3, corollary C.4, and Theorem C.5.

4 Numerical Experiments

The experiments separate the effects of learning and testing. Section 4.1 checks the dictionary's null variance and compares two sequential tests on identical scores. Section 4.2 varies temperature and perturbation strength to explain the variance–detection tradeoff. Section 4.3 compares full transport learning with permutation tests, where the preferred method depends on the alternative.

Table 1: Matched-score sequential tests in eight dimensions (500 streams per setting, level 0.05). Rejection rates estimate type-I error for null rows and power for alternative rows. Mean stopping time is truncated at 1000 pairs; lower is earlier.

	Rejection by 1000		Mean stopping time	
Setting	Half-normal	ONS	Half-normal	ONS
Gaussian null	0.014	0.006	992	997
Standardized t_5 null	0.030	0.016	977	994
Location shift	0.882	0.790	585	681
Scale change	1.000	1.000	303	389

4.1 Stable variance and matched sequential calibration

We instantiate Proposition 3.2 on $\mathbb{R}^8$ with Euclidean distance and the 16 functions $x_j, |x_j|$, $j = 1, \ldots, 8$. The first 20 scores use the zero critic; thereafter we update the critic every 10 pairs using all past observations, with $\vartheta_i = r_i^{-1/4}$, $\tau = 0.10$, and $\rho = 0.05$. Null coordinates are independent standard Gaussian or standardized t_5. Alternatives shift the first Gaussian coordinate by 0.4 or multiply all standard deviations by 1.3. Each setting has 500 independent streams of 1000 pairs.

The null variance identity is explicit here: independent coordinates of variance one give $\mathbb{E}\|X-Y\|^2 = 16$, so the perturbation contribution is $0.10^2 \times 16 + 0.05^2 = 0.1625$ per pair. Proposition 3.2 predicts $A_n^2/n \xrightarrow{\mathbb{P}} 0.1625$. From $n = 250$ to 1000, its mean decreases from 0.1764 to 0.1705 under the Gaussian null and from 0.1772 to 0.1711 under standardized t_5. Its coefficient of variation (CV; standard deviation divided by mean) decreases from 0.027 and 0.026 to 0.014 under both laws. These results connect the null-variance suppression in (11) to concentration of the accumulated variance. The mean ratio V_n/A_n^2 is 0.997–1.001, and fixed-horizon Gaussian rejection frequencies are 0.040–0.056 at $n = 250, 500, 1000$ (Tables 5 and 6). The feasible tests use the observed V_n.

To isolate calibration from learning, we compare the half-normal rule with online Newton step (ONS) betting (Shekhar and Ramdas, 2024) on *identical realized scores*. Both average four prespecified scales $\eta \in \{0.125, 0.25, 0.5, 1\}$ and monitor times 50–1000 at level 0.05. ONS uses the odd bounded payoff $\tanh(\eta\xi_i)$; the half-normal rule uses the score and its square directly. Section B.1 specifies the ONS update and complete comparison.

Following sequential-testing comparisons (Shekhar and Ramdas, 2024), we report rejection probability and average stopping time: with T_{rej} the first monitored rejection time (∞ if there is no rejection), rejection by 1000 estimates type-I error under each null and power under each alternative, and $\mathbb{E}\min(T_{\mathrm{rej}}, 1000)$ is the mean stopping time truncated at the observation budget, nonrejecting streams contributing 1000. The theoretical error guarantee is uniform over monitoring times, while each reported frequency concerns the stated null law and horizon.

The half-normal rule detects both studied alternatives earlier (Table 1). For location, power rises from 0.790 to 0.882, and mean stopping time falls from 681 to 585 pairs. For scale, both rules reject on all streams by 1000, with mean time falling from 389 to 303 pairs. The paired time reductions before rounding are 97 and 86 pairs. These gains occur at the same nominal level; ONS has lower attained type-I error.

Scale is detected earlier despite its smaller dictionary target, $0.3\sqrt{2/\pi}$ versus 0.4 for location. All eight absolute-coordinate contrasts align under scale change, so averaging them retains drift while reducing evaluation variance through independent coordinates. The location learner must identify one informative coordinate. Detection therefore depends on the learned drift relative to score variation and on how quickly a useful witness is learned. With a fixed dictionary, both sequential tests have work and memory per pair independent of stream length: the learner stores 16 cumulative contrasts, the half-normal rule stores S_n, V_n, and ONS maintains four wealths and update variables.

4.2 Regularization and the variance–detection tradeoff

An independent study evaluates the default and halves or doubles either the temperature or both perturbation scales, using 500 new streams per law paired across five prespecified configurations.

Table 2: Temperature and perturbation sensitivity (500 independent streams per law). CV measures dispersion of A_n^2/n; mean stopping times are truncated at 1000 pairs. Smaller values indicate more concentrated variance or earlier detection. Changes multiply the default temperature $r_i^{-1/4}$ or both scales $(\tau, \rho) = (0.10, 0.05)$.

	Null variance CV		Mean stopping time	
Configuration	Gaussian	t_5	Location	Scale
Default	0.015	0.014	580	308
Temperature $\times 0.5$	0.069	0.062	291	207
Temperature $\times 2$	0.004	0.003	891	596
Perturbations $\times 0.5$	0.053	0.049	368	176
Perturbations $\times 2$	0.004	0.004	834	650

Table 3: Empirical rejection frequencies under alternatives (500 paired replications). Exact one-dimensional fixed KR uses $T_n > z_{0.95}$; permutation tests use nominal level 0.05. Bold marks row maxima; Table 8 gives the full grid.

Alternative	n	Fixed KR	Permutation W_1	Permutation energy
Location shift 0.4	100	0.684	**0.768**	0.766
	250	0.984	0.986	**0.988**
Scale change $1 \to 1.3$	100	**0.302**	0.268	0.184
	250	**0.762**	0.716	0.580

Table 2 shows the tradeoff predicted by (11)–(12). Halving temperature reduces location detection time by 289 pairs, while increasing Gaussian-null variance CV from 0.015 to 0.069. Doubling perturbations lowers this CV to 0.004 but delays location detection by 254 pairs. Perturbations add variance without drift; greater temperature suppresses both noise and learned contrast. Time-uniform validity follows from symmetry for every configuration.

Finite-sample Gaussian calibration varies across batches. Rejection frequencies across the two nulls and five configurations are 0.050–0.074; the default Gaussian cell is 0.072 (95% Wilson interval 0.052–0.098), compared with 0.040 in the initial study. A further 2000 independent Gaussian-null streams with the unchanged default give 0.042 (interval 0.034–0.052). Section B.2 specifies the design and the interpretation of these separate batches.

4.3 Full transport learning and permutation comparisons

On $\mathbb{R}$, exact piecewise-linear empirical KR critics remove numerical optimization error and expose the complete procedure's learning–calibration tradeoff. We retain 20 initial zero-critic scores, update the critic every 10 pairs, and compare fixed scales $(0.10, 0.05)$ with vanishing scales $(0.40, 0.20)i^{-1/4}$. Each null or alternative at $n \in \{100, 250, 500\}$ has 500 paired replications. Permutation W_1 and energy use the same samples and 199 relabellings. Nulls include Gaussian, Student t_5, mixture, discrete, and point-mass laws; alternatives change location, scale, mixture weights, or contamination (Section B.3).

Fixed-scale null rejection frequencies are 0.026–0.050, versus 0.030–0.064 for vanishing scales (Table 7). At the point mass, positive external jitter gives fixed-scale frequencies 0.048, 0.040, and 0.040. These are empirical calibration results for exact KR learning: its predictable-variance CV remains about 0.35 at $n = 1000$, so the general Gaussian guarantee still assumes (20).

Table 3 shows the change in ranking across alternatives. At shift 0.4 and $n = 100$, fixed KR power is 0.684 versus 0.768 for permutation W_1; by $n = 250$, all three procedures exceed 0.98. For scale, fixed KR gains 0.034 and 0.046 over permutation W_1 at $n = 100$ and 250. Small-sample mixture and contamination results also favor permutation W_1, with gaps largely disappearing by $n = 500$. Fixed and vanishing scales differ by at most 0.028 across the fixed-alternative grid. The empirical gains are alternative-dependent, consistent with a learned scalar score preserving separation but changing how that separation is weighted by variance.

5 Application: Adélie Penguin Morphology by Sex

We compare adult female and male Adélie morphology within nine expedition-by-island strata using the raw nesting fields of `palmerpenguins` release 0.1.1 (Gorman et al., 2014; Horst et al., 2022). Among 76 sampled nests, 72 have complete measurements and recorded female–male pairs. A metadata-only randomized allocation selects 34 females and 34 males from 68 distinct nests, with one adult per nest; pairing and stream order use identifiers and strata only. The test treats distinct nests as independent. Section B.6 gives the inclusion, allocation, and within-stratum sampling rules.

The four measurements were collected from nesting adults during the 2007/08–2009/10 austral summers and accompanied by molecular sex determinations (Gorman et al., 2014). Within each stratum, $\lfloor N_h/2 \rfloor$ nests are allocated to each sex, where N_h is the number of complete female–male nests; one nest is unused in each of the four odd-sized strata.

Bill length, bill depth, flipper length, and body mass are divided by prespecified units of 10 mm, 5 mm, 20 mm, and 1000 g before using Euclidean distance. Each listed increment contributes one unit of distance when the other coordinates are fixed, making the relative measurement weights interpretable. The critic searches four coordinate directions, which retain contrasts in individual measurements, and 64 fixed random unit directions, which combine measurements; each update compares the 68 exact one-dimensional fits and selects the direction with the largest past projected empirical W_1. After 20 zero-critic pairs, updates at pairs 21, 26, and 31 use all 20, 25, and 30 past pairs and score the next 5, 5, and 4 pairs, respectively. We set $\tau = 0.10$, $\rho = 0.05$, and average the four evidence scales used in the simulations. Pooled training learns the witness, while each evaluation compares independent nests in the same stratum. Equality within every stratum therefore gives scorewise conditional symmetry, as in the stratified-stream extension of Theorem 3.3.

The primary terminal e-value is 49.0, exceeding the level-0.05 threshold of 20 and giving the conservative p-value $\min(1, 1/49.0) \approx 0.020$. The standardized score $T_n = 3.37$ is descriptive: Gaussian variance conditions have not been established for this stratified design. Stratified permutation energy and fixed-direction sliced-W_1 both give $p = 0.002$, the minimum with 499 permutations (Table 9). Across 500 repetitions of allocation, pairing, ordering, and auxiliary signs, 89.4% yield a terminal e-value at least 20 (95% Wilson interval 86.4%–91.8%). This measures randomized-design sensitivity on the observed data.

The finding is an omnibus association between recorded sex and morphology within strata, consistent with the source study's sexual dimorphism (Gorman et al., 2014). Identifying a causal effect or a single responsible coordinate requires a different analysis. Body mass was measured during egg laying, limiting extrapolation to other seasonal phases. The e-process calibrates the learned geometric witness through within-stratum symmetry.

6 Concluding Remarks

A learned Wasserstein critic needs both population separation and a calibrated measure of evidence. Vanishing entropy regularization supplies these jointly for a finite dictionary: it suppresses null critic variance at the training-noise scale while preserving the dictionary's alternative drift. Positive jitter then yields a concrete studentized Gaussian test, including at point-mass nulls. Full and restricted empirical transport optimization have a complementary value-convergence guarantee, and conditional symmetry gives anytime-valid evidence even when Gaussian variance stabilization is unavailable. The experiments show how the choice of regularization and calibration affects detection, with the ranking changing between location and scale alternatives.

Extending the constructive variance argument to growing dictionaries and unrestricted transport critics would connect representation complexity to calibration and power. Quantitative transport-error bounds along a critic update schedule could identify how quickly training samples must grow relative to the dictionary size and the number of testing observations. This would provide a basis for choosing the update schedule together with the regularization strength, balancing the cost of learning a critic against the evidence gained from subsequent observations. For dependent streams, scores based on independent blocks or exchangeable clusters offer a starting point, provided their construction preserves conditional symmetry and controls accumulated variance.

AI USE STATEMENT

Generative AI tools were used to revise wording and improve the clarity and readability of the manuscript; to identify potentially relevant literature, references, and search directions, and to assist with research execution, including drafting, debugging, and refining code. For mathematical proofs, AI was used primarily to review and check the authors' arguments for potential errors, inconsistencies, or gaps. AI-assisted manuscript revisions, literature-search suggestions, feedback on research ideas and proofs, and generated or modified code were reviewed and verified by the authors before incorporation into the final work. The authors independently assessed the accuracy and appropriateness of these outputs and retain full responsibility for the methodology, proofs, implementation, analysis, and conclusions presented in the paper.

REPRODUCIBILITY STATEMENT

The code and an accompanying `README` file are submitted with the manuscript. We provide a dedicated reproduction script for each figure and table, allowing the reported results to be reproduced individually. The `README` provides instructions for running the code.

REFERENCES

C. Anil, J. Lucas, and R. Grosse. Sorting out Lipschitz function approximation. In *Proceedings of the 36th International Conference on Machine Learning*, PMLR 97:291–301, 2019.

M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, PMLR 70:214–223, 2017.

A. Chatterjee and A. Ramdas. A martingale kernel two-sample test. In M. Telgarsky and J. Ullman (editors), *Proceedings of the 37th International Conference on Algorithmic Learning Theory*, PMLR 313:1–44, 2026.

V. Corradi, J. Fosten, and D. Gutknecht. Predictive ability tests with possibly overlapping models. *Journal of Econometrics*, 241(1):105716, 2024.

V. H. de la Peña, M. J. Klass, and T. L. Lai. Self-normalized processes: Exponential inequalities, moment bounds and iterated logarithm laws. *Annals of Probability*, 32(3A):1902–1933, 2004.

N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162:707–738, 2015.

A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.

K. B. Gorman, T. D. Williams, and W. R. Fraser. Ecological sexual dimorphism and environmental variability within a community of Antarctic penguins (genus *Pygoscelis*). *PLOS ONE*, 9(3):e90081, 2014.

P. Hall and C. C. Heyde. *Martingale Limit Theory and Its Application*. Academic Press, 1980.

S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics*, 49(2):1055–1080, 2021.

A. M. Horst, A. P. Hill, and K. B. Gorman. Palmer Archipelago penguins data in the `palmerpenguins` R package: An alternative to Anderson's irises. *The R Journal*, 14(1):244–254, 2022.

X. Hu and Z. Lin. Neural Wasserstein two-sample tests. *arXiv preprint arXiv:2601.21732*, 2026.

M. Kirchler, S. Khorasani, M. Kloft, and C. Lippert. Two-sample testing using deep learning. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, PMLR 108:1387–1398, 2020.

J. M. Kübler, W. Jitkrittum, B. Schölkopf, and K. Muandet. A witness two-sample test. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, PMLR 151:1403–1419, 2022.

F. Liu, W. Xu, J. Lu, G. Zhang, A. Gretton, and D. J. Sutherland. Learning deep kernels for non-parametric two-sample tests. In *Proceedings of the 37th International Conference on Machine Learning*, PMLR 119:6316–6326, 2020.

A. Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.

S. Paik, M. Celentano, A. Green, and R. J. Tibshirani. Integral probability metrics meet neural networks: The Radon–Kolmogorov–Smirnov test. *Journal of Machine Learning Research*, 26(86):1–57, 2025.

T. Pandeva, P. Forré, A. Ramdas, and S. Shekhar. Deep anytime-valid hypothesis testing. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, PMLR 238:622–630, 2024.

A. Ramdas, N. García Trillos, and M. Cuturi. On Wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.

D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Annals of Statistics*, 41(5):2263–2291, 2013.

G. Shafer, A. Shen, N. Vereshchagin, and V. Vovk. Test martingales, Bayes factors and p-values. *Statistical Science*, 26(1):84–101, 2011.

S. Shekhar, I. Kim, and A. Ramdas. A permutation-free kernel two-sample test. In *Advances in Neural Information Processing Systems*, 35:18168–18180, 2022.

S. Shekhar and A. Ramdas. Nonparametric two-sample testing by betting. *IEEE Transactions on Information Theory*, 70(2):1178–1203, 2024.

M. Sommerfeld and A. Munk. Inference for empirical Wasserstein distances on finite spaces. *Journal of the Royal Statistical Society: Series B*, 80(1):219–238, 2018.

B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. G. Lanckriet. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.

G. J. Székely and M. L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8):1249–1272, 2013.

B. T. Tran and N. Schreuder. Minimax-optimal two-sample test with sliced Wasserstein. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, PMLR 300:1927–1935, 2026.

V. S. Varadarajan. On the convergence of sample probability distributions. *Sankhyā*, 19(1/2):23–26, 1958.

C. Villani. *Optimal Transport: Old and New*. Springer, 2009.

J. Wang, R. Gao, and Y. Xie. Two-sample test with kernel projected Wasserstein distance. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, PMLR 151:8022–8055, 2022.

The appendix gives algorithms and implementation (Section A), additional numerical evidence and application design (Section B), further theory (Section C), and supporting lemmas and proofs (Sections D to G).

A Algorithms and Implementation

A.1 Training, scoring, and calibration

Algorithm 2 Predictable Wasserstein critic test: detailed steps

Require: Paired stream (X_i, Y_i); metric d; critic learner; deterministic critic update times $1 = t_1 < t_2 < \cdots$; scales (τ_i, ρ_i); stabilizer $s \in \{\mathrm{A}, \mathrm{D}\}$; level α; either a fixed horizon n or continuous monitoring with a prespecified e-process scale $\eta > 0$.

Set $f_1 \equiv 0$, $S_0 = V_0 = 0$, and $E_0^+(\eta) = 1$.
for $i = 1, 2, \ldots$ **do**
  if $i = t_k$ for some $k \geq 2$ **then**
    Train a certified one-Lipschitz critic using only pairs $1, \ldots, i-1$.
    Orient it using its past empirical objective; use the optimization tolerance required for the intended drift guarantee.
  else
    Keep the most recently trained critic.
  end if
  Observe the fresh pair (X_i, Y_i).
  if $s = \mathrm{D}$ **then**
    Draw independent Rademacher signs R_i, E_i and set $\xi_i = f_i(X_i) - f_i(Y_i) + \tau_i R_i d(X_i, Y_i) + \rho_i E_i$.
  else
    Draw independent Rademacher signs $(R_{i\ell})_{\ell \geq 1}$ (or a finite truncation) and E_i, construct B_i, and set $\xi_i = f_i(X_i) - f_i(Y_i) + \tau_i\{B_i(X_i) - B_i(Y_i)\} + \rho_i E_i$.
  end if
  Update $S_i = S_{i-1} + \xi_i$ and $V_i = V_{i-1} + \xi_i^2$.
  if continuous monitoring **then**
    Update $E_i^+(\eta)$ by (6).
    if $E_i^+(\eta) \geq 1/\alpha$ **then**
      Reject and stop.
    end if
  else if $i = n$ **then**
    Reject if $T_n > z_{1-\alpha}$; stop.
  end if
end for

Each pair is scored before it enters training. Critic hyperparameters may adapt to $\mathcal{G}_{i-1}$, while the perturbation scales and critic update times follow the deterministic schedules in the main results. Gaussian calibration uses Proposition 3.2 for the regularized dictionary or the variance conditions in Theorems C.1 and C.2 for another learner.

A.2 Metric perturbations and critic implementation

Anchor-based perturbations. The anchor alternative uses a dense sequence (a_ℓ) and positive weights w_ℓ summing to one. Set $g_\ell(x) = d(x, a_\ell) - d(x_0, a_\ell)$ and, for independent Rademacher signs $R_{i\ell}$,

$$B_i(x) = \sum_{\ell=1}^{\infty} w_\ell R_{i\ell} g_\ell(x), \qquad H_i^{\mathrm{A}} = B_i(X_i) - B_i(Y_i). \tag{18}$$

The series is absolutely convergent at every x and $B_i \in \mathcal{F}_1$.

For $Z \sim \pi$, its null variance is

$$\kappa_\pi^{\mathrm{A}} = 2 \sum_{\ell=1}^{\infty} w_\ell^2 \operatorname{Var}_\pi\{g_\ell(Z)\}. \tag{19}$$

A dense anchor sequence has $\kappa_\pi^{\mathrm{A}} > 0$ at every non-point null with a finite second moment.

Finite anchor truncation. A finite anchor implementation uses positive weights $w_{\ell,L}$ and

$$B_{i,L}(x) = \sum_{\ell=1}^{L} w_{\ell,L} R_{i\ell} g_\ell(x), \qquad \sum_{\ell=1}^{L} w_{\ell,L} = 1.$$

All null-centering and symmetry arguments remain exact. The corresponding variance is

$$\kappa_{\pi,L}^{\mathrm{A}} = 2 \sum_{\ell=1}^{L} w_{\ell,L}^2 \operatorname{Var}_\pi\{g_\ell(Z)\}.$$

This variance contribution prevents denominator collapse whenever $\kappa_{\pi,L}^{\mathrm{A}} \sum_i \tau_i^2 + \sum_i \rho_i^2 \to \infty$; the fixed-horizon Gaussian theorem also requires (20). Unlike the countably dense construction, a fixed finite dictionary does not guarantee $\kappa_{\pi,L}^{\mathrm{A}} > 0$ for every non-point π. This is why the signed pair-distance version is the default stabilizer.

Periodic critic updates. Suppose the critic is updated at deterministic times t_k and held fixed between updates. Let $r_i = t_k - 1$ for $t_k \le i < t_{k+1}$. Predictability is immediate. The full-W_1 drift proof uses only $r_i \to \infty$ and $\varepsilon_i \to 0$. Thus the gaps $t_{k+1} - t_k$ may increase, provided the most recent training size still diverges. Longer gaps reduce the number of training steps while delaying the critic's response to newly observed separation.

One-sided orientation. The Wasserstein dual is nonnegative because the zero function belongs to $\mathcal{F}_1$. An exact empirical optimizer therefore has a nonnegative training objective. Numerical neural optimization can return a negative value. If the critic class is symmetric, replacing f by $-f$ reverses that value. Selecting the best of f, $-f$, and 0 using past data preserves predictability. This operation aligns the one-sided statistic with the intended alternative but does not guarantee proximity to the empirical supremum.

Lipschitz certification for neural critics. Certified neural architectures, including GroupSort constructions, enforce global Lipschitz control (Anil et al., 2019). A Wasserstein-GAN gradient penalty evaluates gradients only at sampled interpolation points and does not, by itself, prove a global one-Lipschitz bound on an unbounded domain. Such a critic can still be used in a predictable mean-difference test if moments are controlled, because null centering does not depend on the Lipschitz constraint. The Wasserstein drift conclusions in Theorems 3.4 and 3.5 apply to the function class and discrepancy that the implementation certifies.

A.3 Comparison with related procedures

Permutation transport recomputes the discrepancy after relabelling. Once trained, a kernel witness permits linear held-out evaluation (Kübler et al., 2022); cross-MMD uses a quadratic sample-split statistic (Shekhar et al., 2022). Neural projection-Wasserstein aggregates across dimensions and sparsities before Gaussian-maximum calibration (Hu and Lin, 2026). The martingale kernel test likewise uses past witnesses and fresh pairs, with a kernel-based Gaussian construction (Chatterjee and Ramdas, 2026). These computational differences complement the data-use comparison in Table 4.

Here calibration means choosing a reference distribution or evidence threshold to control type-I error. Gaussian and permutation calibration are standard usages in two-sample testing (Chatterjee and Ramdas, 2026); the e-process instead uses a threshold valid uniformly over time.

Table 4: Data use and calibration in representative two-sample procedures. Gaussian calibration is asymptotic; permutation calibration requires exchangeability. For the present method, Theorems C.1 and C.2 give the Gaussian conditions and Proposition 3.2 verifies them for a regularized dictionary.

Method	Data use	Calibration
Permutation W_1 or sliced W_1	Pooled sample; relabelling	Permutation
MMD / energy	Pooled sample	Permutation or asymptotic
Witness test	Train/test split	Studentized Gaussian
Cross-MMD	Sample split	Studentized Gaussian
Martingale kernel test	Past kernel witness, fresh pair	Studentized Gaussian
Testing by betting	Past witness, fresh score	Anytime-valid wealth
Neural projection-Wasserstein	Train/test split	Gaussian maximum
Present method	Past critic, fresh pair	Studentized Gaussian; anytime e-process

B Additional Numerical Evidence and Application Design

B.1 Regularized dictionary and matched sequential calibration

The dictionary study uses $\mathbb{R}^8$ with Euclidean distance and the fixed dictionary $h_j(x) = x_j$, $h_{8+j}(x) = |x_j|$, $j = 1, \dots, 8$. Each function is normalized at zero and globally one-Lipschitz. We fit (3) with $\vartheta_i = r_i^{-1/4}$, 20 initial zero-critic scores, and critic updates every 10 pairs using all past observations. The stabilizer scales are $\tau = 0.10$ and $\rho = 0.05$. The Gaussian and standardized t_5 nulls have independent coordinates of variance one. The alternatives compare $N(0, I_8)$ with $N(0.4e_1, I_8)$ or $N(0, 1.3^2 I_8)$, where I_8 is the identity matrix and e_1 the first coordinate vector. Every setting has 500 independent streams of 1000 pairs; all calibration methods use exactly the same realized score sequence. The Gaussian comparison assesses fixed-horizon studentization; the half-normal and ONS rules compare two ways to accumulate evidence under monitoring.

For numerical evaluation of (3), form the softmax weights on the $2K$ logits $\pm\widehat{d}_{ij}/\vartheta_i$ after subtracting their maximum; the coefficient of h_j is its positive-atom weight minus its negative-atom weight. Evaluate (6) in log scale, using the logarithm of the normal distribution function and log-sum-exp for the mixture.

The half-normal e-process averages over $\eta \in \{0.125, 0.25, 0.5, 1\}$. The betting comparator uses the online Newton step (ONS) rule of Shekhar and Ramdas (2024), Definition 5. For each prespecified η, use the odd bounded payoff $v_i = \tanh(\eta \xi_i)$ as in their Section 2.3, initialize $W_0 = 1$, $\lambda_1 = 0$, and $a_0 = 1$, and update

$$W_i = W_{i-1}(1 + \lambda_i v_i), \qquad z_i = \frac{v_i}{1 + \lambda_i v_i}, \qquad a_i = a_{i-1} + z_i^2,$$
$$\lambda_{i+1} = \Pi_{[-1/2, 1/2]}\left(\lambda_i + \frac{2 z_i}{(2 - \log 3) a_i}\right).$$

Here Π denotes projection onto the indicated interval. Under the null, conditional symmetry makes $\mathbb{E}(v_i \mid \mathcal{G}_{i-1}) = 0$, and the predictable bet keeps each factor nonnegative with conditional expectation one. The equal average of the four wealths is therefore an anytime-valid comparator. Both methods update from the first pair, with crossings checked at every time from 50 through 1000 against the threshold 20. The Gaussian statistic is evaluated only at the three fixed horizons $250, 500, 1000$.

We report crossing probabilities, their paired differences, and the mean first crossing time capped at 1000; a stream with no crossing contributes 1000. Under either null, the feature covariance is diagonal: linear features have variance one and absolute-value features have variance $1 - (\mathbb{E}|U|)^2$,

where U has the common coordinate law. Thus if c_{ij} denotes the fitted coefficient of h_j,

$$q_i = 2\sum_{j=1}^{16} c_{ij}^2 \operatorname{Var}(h_j(Z)) + 0.1625.$$

The constant equals $0.10^2\mathbb{E}\|X - Y\|^2 + 0.05^2$ and is the limit of A_n^2/n in Proposition 3.2. These population variances are used only to examine the variance condition; the feasible statistics use the observed V_n.

Across the two null laws, the mean predictable variance approaches its stabilizer limit and its relative dispersion decreases with sample size (Table 6). The observed quadratic variation tracks this predictable scale. The half-normal rule's higher alternative crossing frequencies accompany higher null crossing frequencies than ONS in these designs. At $n = 500$, its power advantage is 0.138 under both alternatives. At $n = 1000$, both rules detect the scale alternative on all 500 streams; their mean stopping times truncated at 1000 still differ by 86 pairs. Table 5 reports every prespecified horizon and includes the fixed-horizon Gaussian comparison.

Table 5: Crossing probabilities for two anytime-valid tests and fixed-horizon Gaussian rejection frequencies, with 95% Wilson intervals. All methods use the same regularized dictionary scores and 500 streams per setting.

Setting	n	Half-normal crossing	ONS crossing	Gaussian rejection
Gaussian null	250	0.008 [0.003, 0.020]	0.002 [0.000, 0.011]	0.056 [0.039, 0.080]
	500	0.010 [0.004, 0.023]	0.004 [0.001, 0.014]	0.044 [0.029, 0.066]
	1000	0.014 [0.007, 0.029]	0.006 [0.002, 0.017]	0.040 [0.026, 0.061]
Standardized t_5 null	250	0.022 [0.012, 0.039]	0.002 [0.000, 0.011]	0.054 [0.037, 0.077]
	500	0.024 [0.014, 0.041]	0.004 [0.001, 0.014]	0.048 [0.032, 0.070]
	1000	0.030 [0.018, 0.049]	0.016 [0.008, 0.031]	0.056 [0.039, 0.080]
Location shift	250	0.128 [0.102, 0.160]	0.054 [0.037, 0.077]	0.390 [0.348, 0.433]
	500	0.412 [0.370, 0.456]	0.274 [0.237, 0.315]	0.722 [0.681, 0.759]
	1000	0.882 [0.851, 0.907]	0.790 [0.752, 0.823]	0.966 [0.946, 0.979]
Scale change	250	0.380 [0.339, 0.423]	0.206 [0.173, 0.244]	0.790 [0.752, 0.823]
	500	0.924 [0.897, 0.944]	0.786 [0.748, 0.820]	0.996 [0.986, 0.999]
	1000	1.000 [0.992, 1.000]	1.000 [0.992, 1.000]	1.000 [0.992, 1.000]

Table 6: Predictable variance of the regularized dictionary under the Gaussian and standardized t_5 nulls. The limit of A_n^2/n is 0.1625. Entries summarize 500 streams; CV is the standard deviation divided by the mean.

Null	n	Mean A_n^2/n	CV	Mean V_n/A_n^2
Gaussian null	250	0.1764	0.027	0.998
	500	0.1732	0.020	1.000
	1000	0.1705	0.014	1.000
Standardized t_5 null	250	0.1772	0.026	0.997
	500	0.1739	0.018	1.000
	1000	0.1711	0.014	1.001

B.2 Temperature and perturbation sensitivity

The study in Table 2 uses 500 new independent streams per law, with observations and auxiliary signs shared across five fixed configurations. The four laws, dictionary, initialization, critic updates, and evidence mixture are those in Section B.1. Temperature is multiplied by 0.5, 1, or 2 while perturbations stay fixed, or both $(\tau, \rho) = (0.10, 0.05)$ are multiplied by 0.5 or 2 at the default temperature. Gaussian rejection is evaluated at $n = 1000$, while half-normal crossing is monitored from 50 to 1000. No replacement default is selected on these evaluation streams.

Halving temperature shortens location and scale stopping times by 289 and 101 pairs. Halving perturbations shortens them by 212 and 131 pairs. Both changes increase the relative dispersion of predictable variance. Larger perturbations add variance without drift; higher temperature shrinks empirical contrasts. Both delay detection in this study.

The default Gaussian-null rejection frequency is $36/500 = 0.072$ (95% Wilson interval 0.052–0.098) in this batch, compared with 0.040 in Table 5. A further 2000 independent Gaussian-null streams with the unchanged default give $84/2000 = 0.042$ (interval 0.034–0.052). Together, the batches show finite-sample variation around the nominal level while the variance diagnostic supports the proposed stabilization mechanism.

B.3 One-dimensional design and primary comparisons

On the line, the empirical KR critic is piecewise linear between pooled training observations, with slopes determined by the sign of the empirical cumulative-distribution difference and magnitude at most one. It is constant outside the training range, and its training objective equals empirical W_1. We use 20 initial zero-critic scores, update the critic every 10 pairs on all past observations, and compare fixed scales $\tau = 0.10$, $\rho = 0.05$ with $\tau_i = 0.40i^{-1/4}$, $\rho_i = 0.20i^{-1/4}$. Both perturbations use independent Rademacher signs.

The null laws are standard Gaussian, Student t_5 (not standardized in this one-dimensional study), $0.7N(-1, 0.5^2) + 0.3N(1, 1)$, a ten-point discrete law, and a point mass at zero. The discrete support is equally spaced from -2.25 to 2.25, with probabilities proportional to $(1, 2, 3, 5, 8, 8, 5, 3, 2, 1)$. Sample sizes are $n \in \{100, 250, 500\}$ per law. Fixed alternatives use Gaussian shifts 0.2 and 0.4, standard deviations 1 versus 1.3, mixture weights 0.5 versus 0.7 for components $N(-1, 0.6^2)$ and $N(1, 0.6^2)$, or 5% contamination of $N(0, 1)$ by $N(5, 1)$. Every null and fixed-alternative design has 500 independent replications, paired across methods. The 199-draw permutation tests have minimum p-value 0.005.

B.4 Null calibration and stabilization

Table 7 shows that the fixed-stabilizer rejection frequency ranges from 0.026 to 0.050. The most conservative cell is the t_5 null at $n = 250$; its Wilson interval is $(0.015, 0.044)$. The vanishing version ranges from 0.030 to 0.064. Its largest values occur at the point mass, where the effective variance grows only as $\sum_i i^{-1/2}$. These finite-sample values are consistent with slower convergence under a vanishing jitter and do not justify treating the vanishing schedule as uniformly interchangeable with a fixed scale at the studied horizons.

Table 7: Empirical null rejection frequencies (500 replications per design). Fixed and Vanishing use exact one-dimensional KR critics and $T_n > z_{0.95}$; permutation tests use nominal level 0.05.

Null	n	Fixed	Vanishing	Perm. W_1	Perm. energy
Gaussian	100	0.036	0.046	0.052	0.058
t_5	100	0.040	0.050	0.052	0.058
Gaussian mixture	100	0.032	0.042	0.038	0.036
Discrete	100	0.042	0.044	0.036	0.036
Point mass	100	0.048	0.060	0.000	0.000
Gaussian	250	0.050	0.048	0.050	0.058
t_5	250	0.026	0.030	0.040	0.046
Gaussian mixture	250	0.034	0.034	0.030	0.030
Discrete	250	0.038	0.040	0.036	0.028
Point mass	250	0.040	0.052	0.000	0.000
Gaussian	500	0.034	0.032	0.042	0.034
t_5	500	0.042	0.038	0.058	0.058
Gaussian mixture	500	0.034	0.034	0.032	0.034
Discrete	500	0.050	0.056	0.054	0.058
Point mass	500	0.040	0.064	0.000	0.000

Positive external jitter supplies all score variation at a point mass; without it, the zero score cannot reject. The vanishing schedule accumulates perturbation variance at order $\sqrt{n}$. This meets the bounded-space growth requirement, whereas the unbounded-space sufficient condition may require slower decay (Corollary C.4). Fixed scales satisfy the growth requirement for all five null laws. Predictable-variance concentration remains a separate condition for the exact KR learner.

B.5 Power across fixed alternatives

Power increases with sample size under every fixed alternative (Table 8). The comparison depends on how the alternative is expressed through the critic. Under location shifts there is a visible loss at $n = 100$, while under the scale change the fixed version exceeds permutation W_1 by 0.034 and 0.046 at $n = 100$ and 250. Under mixture-weight and contamination alternatives, the gap from permutation W_1 is largest at the smallest sample and largely disappears by $n = 500$. Fixed and vanishing stabilizers have similar power throughout; their largest difference in the fixed-alternative grid is 0.028.

Table 8: Empirical rejection frequencies under fixed alternatives (500 replications per design). Fixed and Vanishing use $T_n > z_{0.95}$; permutation tests use nominal level 0.05. Location shifts are shown in parentheses.

Alternative	n	Fixed	Vanishing	Perm. W_1	Perm. energy
Location (0.2)	100	0.230	0.226	0.286	0.286
Location (0.4)	100	0.684	0.674	0.768	0.766
Location (0.2)	250	0.546	0.518	0.592	0.560
Location (0.4)	250	0.984	0.978	0.986	0.988
Location (0.2)	500	0.818	0.824	0.866	0.858
Location (0.4)	500	1.000	1.000	1.000	1.000
Scale	100	0.302	0.302	0.268	0.184
Scale	250	0.762	0.756	0.716	0.580
Scale	500	0.990	0.988	0.988	0.962
Mixture weight	100	0.642	0.638	0.702	0.720
Mixture weight	250	0.980	0.976	0.982	0.988
Mixture weight	500	1.000	1.000	1.000	1.000
Contamination	100	0.274	0.248	0.320	0.206
Contamination	250	0.680	0.668	0.718	0.494
Contamination	500	0.974	0.966	0.982	0.898

B.6 Application design and metric

We ask whether the joint distribution of four structural measurements differs between adult female and male Adélie penguins in the Palmer Archipelago sample. The measurements—bill length, bill depth, flipper length, and body mass—were collected from nesting adults during the 2007/08–2009/10 austral summers and accompanied by molecular sex determinations (Gorman et al., 2014). The curated `palmerpenguins` release contains 344 penguins and documents both the Environmental Data Initiative records and a CC0 license (Horst et al., 2022). We use the pinned CRAN source release 0.1.1 and its raw nesting fields.

The 152 Adélie records arise from 76 sampled nests, usually with both mates observed. Treating all adults as independent would therefore violate the paper's sampling model. We recover nests by study expedition and the shared prefix of the two documented individual identifiers. Seventy-two nests have one molecularly recorded female, one male, and complete morphometrics for both. Within each of the nine expedition-by-island strata, a recorded seed randomly assigns disjoint nests to the female and male samples, using $\lfloor N_h/2 \rfloor$ nests for each sex, where N_h is the number of complete female–male nests in stratum h. One nest is unused in each of the four odd-sized strata. The primary analysis consequently uses 34 females and 34 males from 68 distinct nests, pairs observations only within strata, and never uses both adults from one nest. The allocation, pairing, and stream order depend only on identifiers and strata, not on morphometrics.

To keep the metric independent of the observed sample, the four coordinates are divided by prespecified measurement units: 10 mm for bill length, 5 mm for bill depth, 20 mm for flipper length, and 1000 g for body mass. These divisors set the relative weights of the measurements: each listed increment contributes one unit of Euclidean distance when the other coordinates are fixed. The weights remain fixed across nest allocations and stream orderings.

The critic class consists of one-dimensional one-Lipschitz functions composed with a projection from a fixed dictionary: the four coordinate directions and 64 random unit directions generated with a recorded seed before stream evaluation. The coordinate directions retain contrasts in individual measurements; the random directions include contrasts that combine measurements. Each critic update compares 68 exact one-dimensional fits. The method selects the direction with the largest past empirical one-dimensional W_1 objective and fits its exact one-dimensional critic. This gives globally one-Lipschitz critics in a restricted class. Training pools the past pairs across strata, so the score's conditional drift is the current stratum's population mean difference evaluated at the learned critic. Summing these differences gives the cumulative drift; it need not equal the sum of separately optimized stratum discrepancies. The e-process remains valid under equality within every stratum by the stratified-stream extension of Theorem 3.3.

We use the zero critic for the first 20 pairs, then update the critic at pairs 21, 26, and 31. The three training samples contain 20, 25, and 30 past pairs, and the resulting critics score the next 5, 5, and 4 fresh pairs, respectively. This schedule reuses the initial training observations while limiting the projection search to three updates. With $\tau_i = 0.10$ and $\rho_i = 0.05$, the signed distance perturbation has magnitude one tenth of the current pair distance, and the external jitter has magnitude 0.05 in the chosen metric units. We use the same equal mixture over $\eta \in \{0.125, 0.25, 0.5, 1\}$ as in the simulations, distributing weight across four geometrically spaced scales. These fixed choices set the perturbation and evidence scales in the stated metric units; their values affect power while the e-process guarantee follows from conditional symmetry. These settings are held fixed across design randomizations.

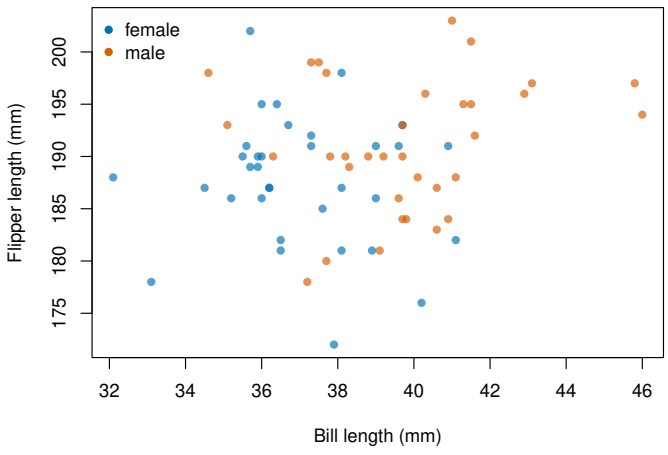

Figure 1: Bill and flipper measurements for the retained Adélie penguins from distinct nests. The test uses these variables together with bill depth and body mass.

B.7 Application results

The prespecified allocation and stream ordering give $T_n = 3.37$. The prespecified equal mixture over $\eta \in \{0.125, 0.25, 0.5, 1\}$ ends at the e-value $E_n^+ = 49.0$ (Table 9); equivalently, the conservative terminal e-to-p calibrator $\min(1, 1/E_n^+)$ is 0.0204. Conditional on the prespecified stratified design, its finite-sample validity is with respect to the joint law of the independent-nest observations and the declared independent auxiliary Rademacher randomization. The recorded seed reproduces one

randomized realization; validity averages over the independent signs. This calibration uses scorewise conditional symmetry and not the predictable-variance stabilization required by the Gaussian theorem. Stratified energy and fixed-direction sliced-Wasserstein permutation tests both attain their minimum p-value of 0.002 with 499 permutations.

Table 9: Stratified analyses of 34 female and 34 male Adélie penguins from distinct nests. The predictable-critic row uses the prespecified nest allocation, stream ordering, and e-process mixture; its p-value is the conservative terminal $\min(1, 1/E)$ calibrator. Permutation rows use 499 within-stratum draws. An e-value of at least 20 rejects at level 0.05. Statistic is T_n for the predictable critic and the corresponding energy or sliced-W_1 discrepancy for each permutation test.

Analysis	Statistic	e-value	p-value
Predictable projected critic	3.37	49.0	0.020
Permutation energy	12.90	–	0.002
Permutation sliced W_1	0.323	–	0.002

C Additional Theoretical Results

This section gives the general Gaussian conditions, perturbation schedules, local limits, and secondary consequences of the main results. Sections D to G contain the supporting lemmas and proofs.

C.1 General fixed-horizon Gaussian calibration

We first treat bounded spaces, where the Lipschitz constraint gives a uniform bound on every score. The Lindeberg condition, which excludes variance from dominant large scores, is then automatic as the variance scale diverges. Proofs of the null results are in Section E.

Theorem C.1 (Bounded-space self-normalized null limit)**.** *Suppose* $\operatorname{diam}(\mathcal{X}) = \Delta < \infty$, $\sup_i(\tau_i + \rho_i) < \infty$, *and*

$$b^2_{n,\pi,s} \longrightarrow \infty.$$

Assume further that there is a deterministic sequence $a_n \to \infty$ *such that*

$$\frac{A_n^2}{a_n^2} \xrightarrow{\mathbb{P}} 1. \tag{20}$$

Under $H_0 : P = Q = \pi$,

$$T_n = \frac{S_n}{\sqrt{V_n}} \xrightarrow{d} N(0,1). \tag{21}$$

Hence $\mathbb{P}_\pi(T_n > z_{1-\alpha}) \to \alpha$.

Condition (20) controls the random variance scale induced by learning: bounded conditionally symmetric increments can otherwise have a non-Gaussian studentized limit even when $V_n = A_n^2$. Section C.4 gives the construction.

On unbounded spaces, an excess moment controls large scores and the fluctuation of V_n around A_n^2. Let $\delta > 0$ denote its order and write $\delta_0 = \min(\delta, 2)$ and

$$r_\delta = 1 + \frac{\delta_0}{2} \in (1, 2]. \tag{22}$$

Theorem C.2 (Unbounded-space self-normalized null limit)**.** *Suppose* $\pi \in \mathcal{P}_{2+\delta}(\mathcal{X})$ *for some* $\delta > 0$, $\sup_i(\tau_i + \rho_i) < \infty$, *and*

$$b^2_{n,\pi,s} \to \infty, \qquad \frac{n}{\{b^2_{n,\pi,s}\}^{r_\delta}} \to 0. \tag{23}$$

Assume also that (20) *holds for a deterministic sequence* $a_n \to \infty$. *Then, under* $H_0 : P = Q = \pi$, (21) *holds and the one-sided test has asymptotic level* α.

C.2 VANISHING STABILIZERS

A stabilizer should regularize the null denominator without dominating power under ordinary alternatives. Let

$$U_i = f_i(X_i) - f_i(Y_i), \qquad Z_i = \tau_i H_i^s + \rho_i E_i, \qquad \xi_i = U_i + Z_i.$$

Proposition C.3. *Suppose $P, Q \in \mathcal{P}_2(\mathcal{X})$, $\tau_i \to 0$, $\rho_i \to 0$, and the scales are bounded. Then*

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n} Z_i \xrightarrow{\mathbb{P}} 0, \tag{24}$$

$$\frac{1}{n} \sum_{i=1}^{n} Z_i^2 \xrightarrow{\mathbb{P}} 0. \tag{25}$$

If $n^{-1} \sum_{i=1}^{n} U_i^2 = O_{\mathbb{P}}(1)$, then

$$\frac{1}{n} \sum_{i=1}^{n} (\xi_i^2 - U_i^2) \xrightarrow{\mathbb{P}} 0. \tag{26}$$

Thus the stabilizer is first-order negligible for root-n alternatives with a nondegenerate unregularized variance, while it can still make V_n diverge under a degenerate null.

The two roles are compatible because null variance accumulation depends on sums such as $\sum_i \tau_i^2$, whereas the perturbation of a root-n score depends on their average $n^{-1} \sum_i \tau_i^2$. For example, $\tau_i, \rho_i \propto i^{-1/4}$ give cumulative stabilizer variance of order $\sqrt{n}$ and average squared scales of order $n^{-1/2}$. Thus perturbations can sustain a degenerate denominator while disappearing on the scale of a nondegenerate alternative. The following corollary states the schedule conditions for the null theorems.

Corollary C.4. *Let $\tau_i = \tau_0 i^{-a}$ and $\rho_i = \rho_0 i^{-\gamma}$.*

(i) On a bounded metric space, the data stabilizer makes $b_{n,\pi,s}^2 \to \infty$ at a non-point null whenever $\tau_0 > 0$ and $0 \le a < 1/2$. External jitter does so at every null, including a point mass, whenever $\rho_0 > 0$ and $0 \le \gamma < 1/2$. In either case, the Gaussian conclusion of Theorem C.1 additionally requires (20)*.*

(ii) On an unbounded space under the $(2+\delta)$ moment condition of Theorem C.2, a sufficient polynomial condition for the variance component that dominates $b_{n,\pi,s}^2$ is

$$s_\star < \frac{\delta_0}{2(2+\delta_0)},$$

where $s_\star$ is the exponent of that dominating component. Fixed scales correspond to $s_\star = 0$. The Gaussian conclusion additionally requires (20)*.*

The broader bounded-space range $s_\star < 1/2$ is enabled by a martingale law of large numbers for observed quadratic variation. The stricter unbounded-space range reflects the finite $(2+\delta)$ moment used to control both quadratic-variation fluctuations and conditional Lindeberg terms. These ranges concern variance accumulation and moment control; neither substitutes for predictable-variance stabilization.

C.3 LOCAL ALTERNATIVES

The preceding negligibility proposition is stated for fixed laws P, Q. Applying the same argument along changing laws requires uniform moment control; the array conditions below specify the requirements for a local limit. The next theorem is stated for a triangular array because both the laws P_n, Q_n and the learned critics may depend on n. Let

$$\xi_{ni} = m_{ni} + e_{ni}, \qquad m_{ni} = \mathbb{E}(\xi_{ni} \mid \mathcal{G}_{n,i-1}), \qquad q_{ni} = \mathbb{E}(e_{ni}^2 \mid \mathcal{G}_{n,i-1}).$$

Here ξ_{ni} is the score from pair i in row n, $\mathcal{G}_{n,i-1}$ is its past information, and e_{ni} is its centered fluctuation. Within each row, set $S_n = \sum_{i=1}^{n} \xi_{ni}$, $V_n = \sum_{i=1}^{n} \xi_{ni}^2$, and $T_n = S_n / \sqrt{V_n}$ with the same zero-denominator convention.

Theorem C.5. *Let $a_n \to \infty$ be deterministic. Suppose*

$$\frac{1}{a_n}\sum_{i=1}^{n} m_{ni} \xrightarrow{\mathbb{P}} \lambda, \tag{27}$$

$$\frac{1}{a_n^2}\sum_{i=1}^{n} q_{ni} \xrightarrow{\mathbb{P}} 1, \tag{28}$$

$$\frac{1}{a_n^2}\sum_{i=1}^{n} \mathbb{E}\{e_{ni}^2 \mathbf{1}(|e_{ni}| > \varepsilon a_n) \mid \mathcal{G}_{n,i-1}\} \xrightarrow{\mathbb{P}} 0 \quad \textit{for every } \varepsilon > 0, \tag{29}$$

$$\frac{1}{a_n^2}\sum_{i=1}^{n} (e_{ni}^2 - q_{ni}) \xrightarrow{\mathbb{P}} 0, \tag{30}$$

$$\frac{1}{a_n^2}\sum_{i=1}^{n} m_{ni}^2 \xrightarrow{\mathbb{P}} 0. \tag{31}$$

Then

$$T_n \xrightarrow{d} N(\lambda, 1).$$

In the standard case $a_n = \sigma\sqrt{n}$, if $n^{-1/2}\sum_i m_{ni} \to \lambda\sigma$ and $n^{-1}\sum_i q_{ni} \to \sigma^2$, the limiting rejection probability is

$$1 - \Phi(z_{1-\alpha} - \lambda).$$

The theorem translates the learned critic's cumulative drift into local power: separation of order a_n shifts the limiting Gaussian mean while the squared conditional means remain negligible in the denominator. Its proof is in Section G. If an empirical full critic has m_{ni} close to $W_1(P_n, Q_n)$ uniformly over most scoring times and $a_n \asymp \sqrt{n}$, the usual root-n scale corresponds heuristically to $W_1(P_n, Q_n) \asymp n^{-1/2}$. A model-specific detection rate would additionally require uniform bounds on this learning error and a matching lower bound.

A uniform conditional $(2+\delta)$ moment bound verifies both the Lindeberg and quadratic-variation conditions when $a_n \asymp \sqrt{n}$. The following corollary gives the sufficient condition for general variance scales; its proof is in Section G.

Corollary C.6. *Suppose, for some $\delta > 0$,*

$$\sup_{n,i} \mathbb{E}(|e_{ni}|^{2+\delta} \mid \mathcal{G}_{n,i-1}) \le C \qquad \text{a.s.},$$

and $a_n = \sigma\sqrt{n}$ with $\sigma > 0$. Then (29) *and* (30) *hold. More generally, they hold for an arbitrary a_n whenever*

$$\frac{n}{a_n^{2+\delta_0}} \to 0, \qquad \delta_0 = \min(\delta, 2). \tag{32}$$

C.4 An adaptive variance counterexample

Let ζ_i be independent Rademacher variables. In an even-length block, take $D_i = \zeta_i$ in the first half and, in the second half, take $D_i = a_+\zeta_i$ if the first-half sum is nonnegative and $D_i = a_-\zeta_i$ otherwise, where $a_+ \neq a_-$ are fixed positive constants. These bounded increments are conditionally symmetric and predictable in scale. At the block endpoint, the self-normalized sum converges to

$$\frac{Z_1 + a(Z_1)Z_2}{\{1 + a(Z_1)^2\}^{1/2}}, \qquad a(z) = a_+\mathbf{1}(z \ge 0) + a_-\mathbf{1}(z < 0),$$

for independent standard normals Z_1, Z_2. Its mean is $(2\pi)^{-1/2}\{(1+a_+^2)^{-1/2} - (1+a_-^2)^{-1/2}\}$, hence it is not standard normal. The observed and predictable quadratic variations are identical and the largest self-normalized increment vanishes. Concatenating blocks whose lengths dominate all preceding lengths embeds the construction in one martingale sequence and produces the same non-Gaussian subsequential limit. Thus conditional symmetry and observed studentization cannot replace (20).

C.5 Degenerate nulls and fixed-alternative power

Proposition C.7. *Let $\pi = \delta_{x_\star}$. Then $f_i(X_i) - f_i(Y_i) = H_i^s = 0$ almost surely and*

$$T_n = \frac{\sum_{i=1}^n \rho_i E_i}{\{\sum_{i=1}^n \rho_i^2\}^{1/2}}.$$

If

$$\sum_{i=1}^n \rho_i^2 \to \infty, \qquad \frac{\max_{1\le i\le n} \rho_i^2}{\sum_{j=1}^n \rho_j^2} \to 0, \tag{33}$$

then $T_n \xrightarrow{d} N(0,1)$. If $\rho_i \equiv 0$, then $V_n = 0$ and the convention $T_n = 0$ yields a conservative test rather than an asymptotically exact level-α test.

At a point-mass null, external jitter supplies the randomness for Gaussian calibration. The e-process remains valid even when all scores are zero.

Theorem C.8. *Suppose*

$$\frac{1}{n}\sum_{i=1}^n m_i \xrightarrow{\mathbb{P}} \mu > 0 \tag{34}$$

and

$$\sup_{i\ge 1} \mathbb{E}\xi_i^2 < \infty. \tag{35}$$

Then

$$\frac{S_n}{n} \xrightarrow{\mathbb{P}} \mu, \qquad T_n \xrightarrow{\mathbb{P}} +\infty.$$

Consequently, T_n crosses every fixed critical value with probability tending to one. Combined with either fixed-horizon null theorem, this gives consistency of the corresponding level-α Gaussian test. If in addition $V_n/n \xrightarrow{\mathbb{P}} \nu \in (0,\infty)$, the e-process satisfies (37)*.*

Positive average drift accumulates linearly in S_n, while the second-moment bound keeps $\sqrt{V_n}$ of order at most $\sqrt{n}$ in probability. This is why any predictable critic with sustained population separation makes T_n diverge. Critic learning enters through m_i, so the theorem also covers learning rules that target a useful witness without recovering a Wasserstein potential. Theorems 3.4 and 3.5 identify when empirical optimization supplies the required drift.

For one-Lipschitz critics and bounded stabilizer scales, $P, Q \in \mathcal{P}_2(\mathcal{X})$ imply the second-moment bound in (35); see Corollary F.1 in Section F.

Proposition C.9. *Suppose, under an alternative, that*

$$\frac{S_n}{n} \xrightarrow{\mathbb{P}} \mu > 0, \qquad \frac{V_n}{n} \xrightarrow{\mathbb{P}} \nu \in (0,\infty). \tag{36}$$

Then, for every fixed $\eta > 0$,

$$\frac{1}{n}\log E_n^+(\eta) \xrightarrow{\mathbb{P}} \frac{\mu^2}{2\nu}, \tag{37}$$

so the anytime-valid test eventually rejects with probability tending to one.

The rate in (37) identifies the signal-to-second-moment ratio that governs evidence accumulation. Each fixed mixture scale $\eta > 0$ has the same leading exponential rate under these laws of large numbers, although its finite-sample crossing time can differ.

C.6 Special cases and metric rescaling

Proposition C.10. *The proposed framework has the following reductions.*

(i) If $\mathcal{X}$ is compact, $s = \mathrm{A}$, $\tau_i \equiv \tau > 0$, $\rho_i = 0$, and the null law is not a point mass, the score uses a compact-space empirical KR critic with fixed anchor stabilization. Its fixed-horizon Gaussian limit follows from Theorem C.1 when $n^{-1}\sum_{i\le n}\mathrm{Var}_\pi\{f_i(Z)\} \xrightarrow{\mathbb{P}} v_\pi$ for a deterministic $v_\pi \ge 0$; its full-W_1 drift conclusion does not require this extra null condition.

(ii) If $f_i \equiv f$ for a deterministic square-integrable function and $\tau_i = \rho_i = 0$, then T_n is the ordinary self-normalized paired mean-difference statistic for $f(X) - f(Y)$ whenever its variance is nonzero.

(iii) Predictable witness classes give the same conditional-symmetry mechanism. For classes contained in $\mathcal{F}_1$, empirical maximization under the assumptions of Theorem 3.5 also gives its restricted-IPM conclusion. For other witness classes, the corresponding drift argument requires uniform empirical convergence over the class. Predictable kernel-witness tests share the null mechanism, while cross-MMD uses a quadratic statistic.

(iv) Let $d_c = cd$ for $c > 0$. If $\rho_i = 0$, rescale every critic and data stabilizer by c. Then $\xi_{i,c} = c\xi_i$ and $T_{n,c} = T_n$. Moreover,

$$E^+_{n,c}(\eta/c) = E^+_n(\eta).$$

The same invariance holds with external jitter if ρ_i is also multiplied by c.

Part (iv) explains why the signed pair-distance stabilizer is naturally calibrated in the same units as the Wasserstein critic. A fixed dimensionless external jitter is not scale invariant unless it is rescaled along with the metric. In practice, ρ_i can be tied to a predictable robust scale estimate computed from past distances, but proving validity for such adaptive scales requires keeping the scale predictable and checking the corresponding variance-growth conditions.

D Metric and Stabilizer Lemmas

Lemma D.1. *Let $Z_1, Z_2, \ldots \overset{\text{i.i.d.}}{\sim} \mu$ on a Polish metric space and let $\widehat{\mu}_m = m^{-1}\sum_{j=1}^m \delta_{Z_j}$. If $\mu \in \mathcal{P}_1(\mathcal{X})$, then*

$$W_1(\widehat{\mu}_m, \mu) \longrightarrow 0 \qquad \text{a.s.}$$

Consequently,

$$\sup_{f\in\mathcal{F}_1} |(\widehat{\mu}_m - \mu)f| = W_1(\widehat{\mu}_m, \mu) \longrightarrow 0 \qquad \text{a.s.}$$

Proof. This is the standard empirical-measure convergence theorem for W_1. We give the reduction to established results. First, Varadarajan's theorem yields $\widehat{\mu}_m \Rightarrow \mu$ almost surely on every Polish space (Varadarajan, 1958). Second, the strong law applied to $d(Z_j, x_0)$ gives

$$\int d(x, x_0)\,\widehat{\mu}_m(\,\mathrm{d}x) = \frac{1}{m}\sum_{j=1}^m d(Z_j, x_0) \longrightarrow \int d(x, x_0)\,\mu(\,\mathrm{d}x) \qquad \text{a.s.}$$

Weak convergence together with convergence of first moments is equivalent to convergence in W_1; see, for example, Villani (2009, Theorem 6.9). This proves the first assertion. The second follows from Kantorovich–Rubinstein duality, with the normalization $f(x_0) = 0$ leaving the dual value unchanged. □

Lemma D.2. *For every pair of finitely supported probability measures $\widehat{P}$ and $\widehat{Q}$ on $\mathcal{X}$, an exact maximizer of the empirical dual problem over $\mathcal{F}_1$ exists. It can be constructed by a finite linear program, deterministic tie breaking, and a McShane extension. Therefore the full empirical critic in* (4) *can be chosen predictably and measurably.*

Proof. Write

$$\widehat{P} = \sum_{a=1}^{m} p_a \delta_{x_a}, \qquad \widehat{Q} = \sum_{b=1}^{\ell} q_b \delta_{y_b},$$

using the atoms in their given measurable order and retaining duplicates. Let $(z_1, \dots, z_M) = (x_1, \dots, x_m, y_1, \dots, y_\ell)$ and $(c_1, \dots, c_M) = (p_1, \dots, p_m, -q_1, \dots, -q_\ell)$, so $M = m + \ell$ and $\sum_k c_k = 0$. Consider the finite linear program

$$\max_{u \in \mathbb{R}^M} \sum_{k=1}^{M} c_k u_k \quad \text{subject to} \quad |u_j - u_k| \le d(z_j, z_k), \quad 1 \le j, k \le M, \quad u_1 = 0.$$

The feasible set is closed. The constraint $u_1 = 0$ and the Lipschitz inequalities imply $|u_k| \le d(z_k, z_1)$, so it is bounded. Hence an optimizer exists.

If two listed atoms coincide, their zero-distance constraint forces the corresponding coordinates to agree. Thus retaining duplicates does not alter the objective, while avoiding any need to measurably enumerate distinct support points. For data-dependent empirical measures, the feasible correspondence has nonempty compact values and a Borel graph because it is defined by finitely many inequalities whose coefficients are measurable functions of the observed distance matrix and weights. The objective is continuous in u with measurable coefficients. The measurable maximum theorem, followed by successive coordinate minimization over the current compact optimizer correspondence, therefore makes the lexicographically selected optimizer a Borel measurable function of the input atoms and weights.

Define

$$\widetilde{f}(x) = \min_{1 \le k \le M} \{u_k + d(x, z_k)\}.$$

For any x, y, each function $u_k + d(x, z_k)$ is one-Lipschitz, and the pointwise minimum of functions with a common Lipschitz constant is one-Lipschitz. For $x = z_j$, feasibility gives

$$u_k + d(z_j, z_k) \ge u_j \quad \text{for every } k,$$

while the term $k = j$ equals u_j. Thus $\widetilde{f}(z_j) = u_j$. Finally set

$$f(x) = \widetilde{f}(x) - \widetilde{f}(x_0).$$

Then $f \in \mathcal{F}_1$ and subtracting the constant does not change $(\widehat{P} - \widehat{Q})f$. The finite LP is precisely the Kantorovich dual restricted to the support points, and the extension preserves its value. Thus f is an exact empirical critic. Applying the construction to past empirical measures gives a predictable measurable selector. □

Lemma D.3. *Let $\pi \in \mathcal{P}_2(\mathcal{X})$. Then κ_π^{A} and κ_π^{D} in* (19) *and* (7) *are finite. Moreover,*

$$\kappa_\pi^{\mathrm{A}} > 0 \iff \kappa_\pi^{\mathrm{D}} > 0 \iff \pi \text{ is not a point mass}.$$

Proof. Step 1: finiteness. By the reverse triangle inequality,

$$|g_\ell(x)| = |d(x, a_\ell) - d(x_0, a_\ell)| \le d(x, x_0).$$

Hence $\mathrm{Var}_\pi\{g_\ell(Z)\} \le \mathbb{E}d(Z, x_0)^2$. Since $\sum_\ell w_\ell^2 \le 1$,

$$\kappa_\pi^{\mathrm{A}} \le 2\mathbb{E}d(Z, x_0)^2 < \infty.$$

Also $d(Z, Z')^2 \le 2d(Z, x_0)^2 + 2d(Z', x_0)^2$, so $\kappa_\pi^{\mathrm{D}} < \infty$.

Step 2: the pair-distance stabilizer. If π is a point mass, then $d(Z, Z') = 0$ almost surely. Conversely, if $\mathbb{E}d(Z, Z')^2 = 0$, then $d(Z, Z') = 0$ almost surely, so two independent draws are equal almost surely. This is possible only when π is concentrated at one point.

Step 3: the anchor stabilizer. If π is a point mass, every variance in (19) is zero. Conversely, suppose $\kappa_\pi^{\mathrm{A}} = 0$. Every weight is positive and every summand is nonnegative, hence $\mathrm{Var}_\pi\{g_\ell(Z)\} = 0$ for every ℓ. Thus, for each ℓ, g_ℓ is constant on $\mathrm{supp}(\pi)$. For any x, y in that support,

$$d(x, a_\ell) = d(y, a_\ell) \qquad \text{for every } \ell.$$

The anchors are dense and $a \mapsto d(x, a) - d(y, a)$ is continuous, so the equality holds for all $a \in \mathcal{X}$. Taking $a = x$ gives $d(x, y) = 0$. Therefore the support contains only one point. □

Remark D.4 (Finite anchor dictionaries). If (18) is truncated to $L < \infty$ anchors, its null variance is positive precisely when at least one of $g_1, \ldots, g_L$ is nonconstant under π. A finite arbitrary dictionary need not detect every non-point null. The signed pair-distance stabilizer avoids this coverage issue.

E Proofs of the Null Results

Proof of Proposition 3.1. Fix i and condition on $\mathcal{G}_{i-1}$. Under $\pi \in \mathcal{P}_1$, the Lipschitz bounds make every score term integrable, which justifies Steps 1–3. Steps 4–5 establish symmetry for arbitrary π. Under $\pi \in \mathcal{P}_2$, the same bounds give finite second moments for Steps 6–8.

Step 1: centering of the critic term. The function f_i is fixed under the conditioning. Under H_0, X_i and Y_i are independent with common law π, so

$$\mathbb{E}\{f_i(X_i) - f_i(Y_i) \mid \mathcal{G}_{i-1}\} = \pi f_i - \pi f_i = 0.$$

Step 2: centering of the anchor stabilizer. For fixed x, y, absolute convergence permits termwise expectation:

$$\mathbb{E}\{B_i(x) - B_i(y)\} = \sum_{\ell=1}^{\infty} w_\ell \mathbb{E}(R_{i\ell})\{g_\ell(x) - g_\ell(y)\} = 0.$$

Averaging over (X_i, Y_i) gives $\mathbb{E}(H_i^{\mathrm{A}} \mid \mathcal{G}_{i-1}) = 0$.

Step 3: centering of the pair-distance stabilizer and jitter. Because R_i and E_i are independent centered signs,

$$\mathbb{E}\{R_i d(X_i, Y_i) \mid X_i, Y_i, \mathcal{G}_{i-1}\} = 0, \qquad \mathbb{E}(E_i \mid \mathcal{G}_{i-1}) = 0.$$

Combining Steps 1–3 proves the martingale-difference identity in (8).

Step 4: conditional symmetry for the anchor version. Under the conditional law, the transformation

$$(X_i, Y_i, E_i, (R_{i\ell})_{\ell \geq 1}) \longmapsto (Y_i, X_i, -E_i, (R_{i\ell})_{\ell \geq 1})$$

preserves probability: (X_i, Y_i) is exchangeable, E_i is symmetric, and the anchor signs are unchanged. The critic difference and the anchor difference both change sign when X_i and Y_i are swapped, and the jitter changes sign because E_i is negated. Therefore ξ_i is mapped to $-\xi_i$ under a law-preserving transformation.

Step 5: conditional symmetry for the pair-distance version. Now use

$$(X_i, Y_i, R_i, E_i) \longmapsto (Y_i, X_i, -R_i, -E_i).$$

The critic difference changes sign, $d(X_i, Y_i)$ is unchanged while R_i changes sign, and the jitter changes sign. Again the joint conditional law is invariant and the score is mapped to its negative. This proves the symmetry identity in (8).

Step 6: conditional variance of the critic. Since the two critic evaluations are independent and identically distributed under π,

$$\mathbb{E}\left[\{f_i(X_i) - f_i(Y_i)\}^2 \mid \mathcal{G}_{i-1}\right] = 2\,\mathrm{Var}_\pi\{f_i(Z)\}.$$

Step 7: variance and orthogonality of the anchor stabilizer. For fixed x, y, independence of the signs gives

$$\mathbb{E}_R\{B_i(x) - B_i(y)\}^2 = \sum_{\ell=1}^{\infty} w_\ell^2 \{g_\ell(x) - g_\ell(y)\}^2.$$

For independent $X_i, Y_i \sim \pi$,

$$\mathbb{E}\{g_\ell(X_i) - g_\ell(Y_i)\}^2 = 2\,\mathrm{Var}_\pi\{g_\ell(Z)\}.$$

Tonelli's theorem therefore gives

$$\mathbb{E}\{(H_i^{\mathrm{A}})^2 \mid \mathcal{G}_{i-1}\} = \kappa_\pi^{\mathrm{A}}.$$

Conditioning first on (X_i, Y_i) and using the centered anchor signs shows that the cross product between the critic difference and H_i^{A} has conditional expectation zero. The cross products involving E_i are also zero.

Step 8: variance and orthogonality of the pair-distance stabilizer. Since $R_i^2 = 1$,

$$\mathbb{E}\{(H_i^{\mathrm{D}})^2 \mid \mathcal{G}_{i-1}\} = \mathbb{E}d(X_i, Y_i)^2 = \kappa_\pi^{\mathrm{D}}.$$

The cross product with the critic has conditional expectation zero after conditioning on (X_i, Y_i) and using $\mathbb{E}R_i = 0$. Jitter cross products vanish in the same way.

Adding the mutually orthogonal conditional second moments proves (9); summing its nonnegative terms proves (10). □

We use the following standard martingale central limit theorem with a deterministic variance scale.

Lemma E.1. *Let $(D_i, \mathcal{G}_i)$ be martingale differences and let $a_n \to \infty$ be deterministic. Put $q_i = \mathbb{E}(D_i^2 \mid \mathcal{G}_{i-1})$. If*

$$\frac{1}{a_n^2}\sum_{i=1}^n q_i \xrightarrow{\mathbb{P}} 1,$$

$$\frac{1}{a_n^2}\sum_{i=1}^n \mathbb{E}\{D_i^2 \mathbf{1}(|D_i| > \varepsilon a_n) \mid \mathcal{G}_{i-1}\} \xrightarrow{\mathbb{P}} 0 \quad \textit{for every } \varepsilon > 0,$$

then

$$\frac{\sum_{i=1}^n D_i}{a_n} \xrightarrow{d} N(0,1).$$

Proof. This is the martingale-array central limit theorem; see Hall and Heyde (1980, Corollary 3.1). Applying the array result to D_i/a_n, $1 \le i \le n$, gives the displayed formulation. The array filtrations are nested because the nth row uses the initial segment $(\mathcal{G}_i)_{i \le n}$ of one filtration. □

Lemma E.2. *Let $(D_i, \mathcal{G}_i)$ be martingale differences such that $|\xi_i| \le C$ almost surely, where*

$$D_i = \xi_i^2 - q_i, \qquad q_i = \mathbb{E}(\xi_i^2 \mid \mathcal{G}_{i-1}), \qquad A_n^2 = \sum_{i=1}^n q_i.$$

On the event $A_n^2 \to \infty$,

$$\frac{\sum_{i=1}^n D_i}{A_n^2} \longrightarrow 0 \qquad \text{a.s.} \tag{38}$$

In particular, if $A_n^2 \to \infty$ in probability, then the convergence holds in probability.

Proof. Set $M_n = \sum_{i=1}^n D_i$. Since $q_i \le C^2$ and $\xi_i^4 \le C^2 \xi_i^2$,

$$\mathbb{E}(D_i^2 \mid \mathcal{G}_{i-1}) \le \mathbb{E}(\xi_i^4 \mid \mathcal{G}_{i-1}) \le C^2 q_i. \tag{39}$$

We prove the ratio by dyadic stopping.

For an integer $j \ge 0$, define

$$\sigma_j = \inf\{n \ge 1 : A_n^2 \ge 2^j\},$$

with the convention $\inf \varnothing = \infty$. Because each increment $q_i \le C^2$, on $\{\sigma_{j+1} < \infty\}$,

$$A_{\sigma_{j+1}}^2 \le 2^{j+1} + C^2.$$

For any deterministic N, the stopped martingale is square integrable, and (39) gives

$$\mathbb{E}M^2_{\sigma_{j+1}\wedge N} = \mathbb{E}\sum_{i=1}^{\sigma_{j+1}\wedge N}\mathbb{E}(D_i^2 \mid \mathcal{G}_{i-1})$$
$$\leq C^2\mathbb{E}A^2_{\sigma_{j+1}\wedge N} \leq C^2(2^{j+1}+C^2).$$

Doob's maximal inequality therefore implies

$$\mathbb{P}\left(\{\max_{k\leq\sigma_{j+1}\wedge N}|M_k| > \varepsilon 2^j\}\right) \leq \frac{C^2(2^{j+1}+C^2)}{\varepsilon^2 2^{2j}}. \tag{40}$$

Letting $N\to\infty$ by monotone convergence of the events preserves the bound.

Fix J. If $A_n^2 \geq 2^J$ and $|M_n| > \varepsilon A_n^2$, then for some $j \geq J$,

$$2^j \leq A_n^2 < 2^{j+1} \quad \text{and} \quad \max_{k\leq\sigma_{j+1}}|M_k| > \varepsilon 2^j.$$

By a union bound and (40),

$$\mathbb{P}\left(\left\{\sup_{n:A_n^2\geq 2^J}\frac{|M_n|}{A_n^2} > \varepsilon\right\}\right) \leq \sum_{j=J}^{\infty}\frac{C^2(2^{j+1}+C^2)}{\varepsilon^2 2^{2j}} \longrightarrow 0$$

as $J\to\infty$. This proves (38) on $\{A_n^2\to\infty\}$. If $A_n^2\to\infty$ only in probability, choose J large and split according to $\{A_n^2 < 2^J\}$ to obtain convergence in probability. □

Proof of Theorem C.1. Step 1: uniformly bound the scores. For every $f_i \in \mathcal{F}_1$,

$$|f_i(X_i) - f_i(Y_i)| \leq d(X_i, Y_i) \leq \Delta.$$

For the anchor stabilizer, B_i is one-Lipschitz, so $|H_i^{\mathrm{A}}| \leq \Delta$. For the pair-distance stabilizer, $|H_i^{\mathrm{D}}| \leq \Delta$. Therefore, with $\bar\tau = \sup_i \tau_i$ and $\bar\rho = \sup_i \rho_i$,

$$|\xi_i| \leq C_\xi := \Delta(1+\bar\tau) + \bar\rho. \tag{41}$$

Step 2: predictable variance diverges. By Proposition 3.1,

$$A_n^2 \geq b^2_{n,\pi,s} \to \infty.$$

Step 3: observed quadratic variation is equivalent. Apply Lemma E.2 with $C = C_\xi$. Since $V_n - A_n^2 = \sum_{i=1}^n(\xi_i^2 - q_i)$,

$$\frac{V_n}{A_n^2} \xrightarrow{\mathbb{P}} 1.$$

Together with (20), this gives

$$\frac{V_n}{a_n^2} \xrightarrow{\mathbb{P}} 1. \tag{42}$$

Step 4: conditional Lindeberg. For every $\varepsilon > 0$, (41) and $a_n\to\infty$ make $\mathbf{1}(|\xi_i| > \varepsilon a_n) = 0$ for all $i \leq n$ eventually. Thus the conditional Lindeberg condition in Lemma E.1 holds.

Step 5: apply the martingale CLT and studentize. Condition (20) and Step 4 verify Lemma E.1, so $S_n/a_n \xrightarrow{d} N(0,1)$. Combining this with (42) and Slutsky's theorem gives $T_n \xrightarrow{d} N(0,1)$. Continuity of the standard normal distribution at $z_{1-\alpha}$ yields the level statement. □

Lemma E.3. *Let $(D_i, \mathcal{G}_i)$ be martingale differences and let $r \in (1,2]$. There is a constant C_r depending only on r such that*

$$\mathbb{E}\left|\sum_{i=1}^n D_i\right|^r \leq C_r\sum_{i=1}^n\mathbb{E}|D_i|^r.$$

Proof. Burkholder's inequality (Hall and Heyde, 1980, Theorem 2.10) gives

$$\mathbb{E}\left|\sum_{i=1}^{n} D_i\right|^r \le C_r \mathbb{E}\left(\sum_{i=1}^{n} D_i^2\right)^{r/2}.$$

Because $r/2 \le 1$, subadditivity of $x \mapsto x^{r/2}$ on $[0,\infty)$ bounds the right-hand side by $C_r \sum_i \mathbb{E}|D_i|^r$. □

Proof of Theorem C.2. Let $r = r_\delta$ from (22), so $2r = 2 + \delta_0 \le 2 + \delta$.

Step 1: a uniform $(2+\delta)$ score bound. For $f_i \in \mathcal{F}_1$,

$$|f_i(X_i) - f_i(Y_i)| \le d(X_i, x_0) + d(Y_i, x_0).$$

The same bound holds for H_i^{A} because $|B_i(x)| \le d(x, x_0)$, and for H_i^{D} because $d(X_i, Y_i) \le d(X_i, x_0) + d(Y_i, x_0)$. Since the scales are bounded, there is a deterministic C such that

$$|\xi_i| \le C\{1 + d(X_i, x_0) + d(Y_i, x_0)\}.$$

Thus

$$\sup_i \mathbb{E}|\xi_i|^{2+\delta} < \infty. \tag{43}$$

Step 2: observed versus predictable quadratic variation. Set $D_i = \xi_i^2 - q_i$. The inequality $|a-b|^r \le 2^{r-1}(|a|^r + |b|^r)$, conditional Jensen, and $q_i = \mathbb{E}(\xi_i^2 \mid \mathcal{G}_{i-1})$ give

$$\begin{aligned}\mathbb{E}|D_i|^r &\le 2^{r-1}\left\{\mathbb{E}|\xi_i|^{2r} + \mathbb{E}q_i^r\right\} \\ &\le 2^r \mathbb{E}|\xi_i|^{2r} \le C.\end{aligned}$$

By Lemma E.3,

$$\mathbb{E}\left|\sum_{i=1}^{n} D_i\right|^r \le Cn.$$

Since $A_n^2 \ge b_{n,\pi,s}^2$, Markov's inequality yields, for every $\varepsilon > 0$,

$$\begin{aligned}\mathbb{P}\left(\frac{|V_n - A_n^2|}{A_n^2} > \varepsilon\right) &\le \mathbb{P}\left(\left|\sum_{i=1}^{n} D_i\right| > \varepsilon b_{n,\pi,s}^2\right) \\ &\le \frac{Cn}{\varepsilon^r \{b_{n,\pi,s}^2\}^r} \to 0.\end{aligned}$$

Therefore

$$V_n / A_n^2 \xrightarrow{\mathbb{P}} 1.$$

Because $b_{n,\pi,s}^2 \to \infty$, both A_n^2 and V_n diverge in probability.

Step 3: conditional Lindeberg. Because $A_n^2/a_n^2 \xrightarrow{\mathbb{P}} 1$ and $A_n^2 \ge b_{n,\pi,s}^2$ deterministically, $b_{n,\pi,s}^2/a_n^2$ is eventually bounded above. Hence (23) implies

$$\frac{n}{a_n^{2+\delta_0}} \to 0.$$

For every $\varepsilon > 0$, conditional Markov's inequality gives

$$\begin{aligned}&\frac{1}{a_n^2}\sum_{i=1}^{n} \mathbb{E}\{\xi_i^2 \mathbf{1}(|\xi_i| > \varepsilon a_n) \mid \mathcal{G}_{i-1}\} \\ &\quad\le \frac{1}{\varepsilon^{\delta_0} a_n^{2+\delta_0}} \sum_{i=1}^{n} \mathbb{E}(|\xi_i|^{2+\delta_0} \mid \mathcal{G}_{i-1}).\end{aligned}$$

By (43), the expectation of the right-hand side is $O(n/a_n^{2+\delta_0})$, so it converges to zero in probability by Markov's inequality.

Step 4: conclude. Condition (20) and Step 3 verify Lemma E.1, so $S_n/a_n \xrightarrow{d} N(0,1)$. Step 2 and (20) give $V_n/a_n^2 \xrightarrow{\mathbb{P}} 1$. Slutsky's theorem now yields $T_n \xrightarrow{d} N(0,1)$. □

Proof of Proposition C.7. If $X_i = Y_i = x_\star$ almost surely, every function difference and both data stabilizers vanish. Hence $\xi_i = \rho_i E_i$ and $V_n = \sum_i \rho_i^2$ deterministically. The variables $\rho_i E_i$ are independent, centered, and have total variance $\sum_i \rho_i^2$. Condition (33) is the Lindeberg condition for this bounded triangular array because, for every $\varepsilon > 0$, eventually

$$\max_{i \le n} \rho_i \le \varepsilon \left(\sum_{j=1}^{n} \rho_j^2 \right)^{1/2}.$$

The Lindeberg–Feller central limit theorem gives the result. When every $\rho_i = 0$, both numerator and denominator are zero by construction. □

E.1 Regularized finite dictionaries

Proof of Proposition 3.2. Write $\mathcal{H} = \mathrm{conv}\{\pm h_1, \dots, \pm h_K\}$. The coefficients in (3) are the signed sums of the softmax weights

$$w_{ij}^{\pm} = \frac{\exp(\pm \widehat{d}_{ij}/\vartheta_i)}{\sum_{k=1}^{K} \{\exp(\widehat{d}_{ik}/\vartheta_i) + \exp(-\widehat{d}_{ik}/\vartheta_i)\}}.$$

These nonnegative weights sum to one. Hence $f_i \in \mathcal{H} \subseteq \mathcal{F}_1$, and each evaluation is a measurable function of the past observations. The empirical objective is nonnegative because $u \sinh(u) \ge 0$. The weights maximize the empirical objective minus ϑ_i times relative entropy with respect to the uniform distribution on the $2K$ signed functions. Comparing with a vertex of the simplex and using nonnegativity of relative entropy gives

$$0 \le \max_{j \le K} |\widehat{d}_{ij}| - (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) f_i \le \vartheta_i \log(2K).$$

Under $P = Q = \pi \in \mathcal{P}_2$, the weights proportional to $\cosh(\widehat{d}_{ij}/\vartheta_i)$ sum to one across j. Since $|\tanh u| \le \min(1, |u|)$ and $|h_j(x)| \le d(x, x_0)$,

$$|f_i(x)| \le d(x, x_0) \min\{1, \max_{j \le K} |\widehat{d}_{ij}|/\vartheta_i\}.$$

The training means have expectation zero and $\mathbb{E}_\pi \widehat{d}_{ij}^2 = 2 \operatorname{Var}_\pi(h_j)/r_i$. Thus

$$\mathbb{E}_\pi \|f_i\|_{L^2(\pi)}^2 \le \frac{2\mathbb{E}_\pi d(Z, x_0)^2 \sum_{j=1}^{K} \operatorname{Var}_\pi\{h_j(Z)\}}{r_i \vartheta_i^2} \longrightarrow 0.$$

The same bound holds with periodic critic updates because r_i and ϑ_i refer to the most recent training sample. Cesàro averaging gives $n^{-1} \sum_{i \le n} \operatorname{Var}_\pi(f_i) \to 0$ in L^1. Equation (9) then yields $A_n^2/n \xrightarrow{\mathbb{P}} \tau^2 \kappa_\pi^s + \rho^2$. This constant is positive since $\rho > 0$. Under the stated $(2+\delta)$ moment, Theorem C.2 applies with $a_n^2 = n(\tau^2 \kappa_\pi^s + \rho^2)$; its variance-growth condition holds because $b_{n,\pi,s}^2 \asymp n$.

For $P, Q \in \mathcal{P}_1$, the strong law for the finite set of integrable h_j gives $\max_{j \le K} |\widehat{d}_{ij} - (P-Q)h_j| \to 0$ almost surely. The empirical optimization bound implies

$$0 \le \max_{j \le K} |(P-Q)h_j| - (P-Q)f_i \le 2 \max_{j \le K} |\widehat{d}_{ij} - (P-Q)h_j| + \vartheta_i \log(2K) \longrightarrow 0.$$

This proves drift recovery. For a positive limiting discrepancy and second moments, Corollary F.1 and Theorem C.8 give consistency. □

F Proofs for Anytime Validity and Power

Proof of Theorem 3.3. Fix $\lambda \ge 0$. By Proposition 3.1, the conditional law of ξ_i is symmetric. Therefore, conditioning on $\mathcal{G}_{i-1}$,

$$\mathbb{E}\left[\exp\left\{ \lambda \xi_i - \frac{\lambda^2}{2} \xi_i^2 \right\} \middle| \mathcal{G}_{i-1} \right] = \mathbb{E}\left[\cosh(\lambda \xi_i) \exp\left\{ -\frac{\lambda^2}{2} \xi_i^2 \right\} \middle| \mathcal{G}_{i-1} \right]$$

$$\leq 1,$$

where the equality uses conditional symmetry and the inequality uses $\cosh(u) \leq \exp(u^2/2)$. Since

$$L_n(\lambda) = L_{n-1}(\lambda) \exp\left\{\lambda \xi_n - \frac{\lambda^2}{2}\xi_n^2\right\},$$

$(L_n(\lambda))_{n\geq 0}$ is a nonnegative supermartingale.

The half-normal mixing density is

$$h_\eta(\lambda) = \frac{\sqrt{2}}{\eta\sqrt{\pi}} \exp\left(-\frac{\lambda^2}{2\eta^2}\right) \mathbf{1}(\lambda \geq 0).$$

Tonelli's theorem shows that

$$\widetilde{E}_n(\eta) := \int_0^\infty L_n(\lambda) h_\eta(\lambda)\ \mathrm{d}\lambda$$

is also a nonnegative supermartingale with initial value one. To evaluate the integral, set

$$a = V_n + \eta^{-2}.$$

Completing the square gives

$$\lambda S_n - \frac{a}{2}\lambda^2 = -\frac{a}{2}\left(\lambda - \frac{S_n}{a}\right)^2 + \frac{S_n^2}{2a}.$$

Thus

$$\begin{aligned}\widetilde{E}_n(\eta) &= \frac{\sqrt{2}}{\eta\sqrt{\pi}} \exp\left(\frac{S_n^2}{2a}\right) \int_0^\infty \exp\left\{-\frac{a}{2}\left(\lambda - \frac{S_n}{a}\right)^2\right\} \mathrm{d}\lambda \\ &= \frac{2}{\eta\sqrt{a}} \exp\left(\frac{S_n^2}{2a}\right) \Phi\left(\frac{S_n}{\sqrt{a}}\right).\end{aligned}$$

Because $\eta\sqrt{a} = \sqrt{1+\eta^2 V_n}$ and $S_n/\sqrt{a} = \eta S_n/\sqrt{1+\eta^2 V_n}$, this equals (6). Ville's inequality for nonnegative supermartingales yields (13). □

Proof of Proposition C.9. Write

$$\begin{aligned}\frac{1}{n}\log E_n^+(\eta) = &-\frac{1}{2n}\log(1+\eta^2 V_n) + \frac{\eta^2 S_n^2}{2n(1+\eta^2 V_n)} \\ &+ \frac{1}{n}\log 2 + \frac{1}{n}\log\Phi\left(\frac{\eta S_n}{\sqrt{1+\eta^2 V_n}}\right).\end{aligned}$$

The first and third terms converge to zero in probability. Under (36),

$$\frac{\eta S_n}{\sqrt{1+\eta^2 V_n}} = \sqrt{n}\,\frac{\eta(S_n/n)}{\{n^{-1}+\eta^2(V_n/n)\}^{1/2}} \xrightarrow{\mathbb{P}} +\infty.$$

Hence the last logarithm converges to zero and, divided by n, remains negligible. Finally,

$$\frac{\eta^2 S_n^2}{2n(1+\eta^2 V_n)} = \frac{(S_n/n)^2}{2\{n^{-1}\eta^{-2} + V_n/n\}} \xrightarrow{\mathbb{P}} \frac{\mu^2}{2\nu}.$$

Combining the terms proves (37). □

F.1 Fixed-alternative consistency and drift

Proof of Theorem C.8. Let

$$e_i = \xi_i - m_i.$$

Then $(e_i, \mathcal{G}_i)$ is a martingale-difference sequence. Conditional Jensen gives

$$m_i^2 = \{\mathbb{E}(\xi_i \mid \mathcal{G}_{i-1})\}^2 \leq \mathbb{E}(\xi_i^2 \mid \mathcal{G}_{i-1}),$$

and therefore

$$\mathbb{E}e_i^2 = \mathbb{E}\,\mathrm{Var}(\xi_i \mid \mathcal{G}_{i-1}) \le \mathbb{E}\xi_i^2 \le C,$$

where $C = \sup_i \mathbb{E}\xi_i^2 < \infty$. Orthogonality of martingale differences gives

$$\mathbb{E}\left(\frac{1}{n}\sum_{i=1}^n e_i\right)^2 = \frac{1}{n^2}\sum_{i=1}^n \mathbb{E}e_i^2 \le \frac{C}{n} \to 0.$$

Hence

$$\frac{1}{n}\sum_{i=1}^n e_i \xrightarrow{\mathbb{P}} 0.$$

Together with (34), this proves $S_n/n \xrightarrow{\mathbb{P}} \mu$.

Next, Markov's inequality gives, for every $K > 0$,

$$\mathbb{P}(V_n/n > K) \le \frac{1}{Kn}\sum_{i=1}^n \mathbb{E}\xi_i^2 \le \frac{C}{K}.$$

Thus $V_n/n = O_{\mathbb{P}}(1)$. To prove divergence of T_n, fix $M > 0$ and $\zeta > 0$. Choose K so that $\limsup_n \mathbb{P}(V_n/n > K) < \zeta$. On the event

$$S_n/n > \mu/2, \qquad V_n/n \le K,$$

we have

$$T_n = \frac{S_n}{\sqrt{V_n}} \ge \frac{\mu\sqrt{n}}{2\sqrt{K}} > M$$

for all sufficiently large n. Since the first event has probability tending to one, $\limsup_n \mathbb{P}(T_n \le M) \le \zeta$. Letting $\zeta \downarrow 0$ proves $T_n \xrightarrow{\mathbb{P}} +\infty$. The final assertion follows from Proposition C.9. □

Corollary F.1. *Suppose $P, Q \in \mathcal{P}_2(\mathcal{X})$, the critics lie in $\mathcal{F}_1$, and $\sup_i(\tau_i + \rho_i) < \infty$. Then* (35) *holds for either stabilizer.*

Proof of Corollary F.1. Let $r(x) = d(x, x_0)$. The critic difference satisfies

$$|f_i(X_i) - f_i(Y_i)| \le r(X_i) + r(Y_i).$$

For the anchor stabilizer, $|B_i(x)| \le r(x)$, so $|H_i^{\mathrm{A}}| \le r(X_i) + r(Y_i)$. For the pair-distance stabilizer, $|H_i^{\mathrm{D}}| = d(X_i, Y_i) \le r(X_i) + r(Y_i)$. With bounded scales, there is $C_0 < \infty$ such that

$$|\xi_i| \le C_0\{1 + r(X_i) + r(Y_i)\}.$$

Therefore

$$\sup_i \mathbb{E}\xi_i^2 \le C\{1 + \mathbb{E}_P r(X)^2 + \mathbb{E}_Q r(Y)^2\} < \infty.$$

□

Proof of Theorem 3.4. Define

$$\Delta_i = \sup_{f\in\mathcal{F}_1} \left|\{(\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) - (P - Q)\}f\right|. \tag{44}$$

By the triangle inequality and Lemma D.1,

$$\begin{aligned}\Delta_i &\le \sup_{f\in\mathcal{F}_1} |(\widehat{P}_i^{\mathrm{tr}} - P)f| + \sup_{f\in\mathcal{F}_1} |(\widehat{Q}_i^{\mathrm{tr}} - Q)f| \\ &= W_1(\widehat{P}_i^{\mathrm{tr}}, P) + W_1(\widehat{Q}_i^{\mathrm{tr}}, Q) \longrightarrow 0 \qquad \text{a.s.}\end{aligned} \tag{45}$$

The population drift is $m_i = (P - Q)f_i$ because the stabilizer is centered. The upper bound follows directly from duality:

$$m_i \le \sup_{f\in\mathcal{F}_1} (P - Q)f = W_1(P, Q). \tag{46}$$

For the lower bound, approximate empirical optimality gives

$$\begin{aligned} m_i &\geq (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}})f_i - \Delta_i \\ &\geq W_1(\widehat{P}_i^{\mathrm{tr}}, \widehat{Q}_i^{\mathrm{tr}}) - \varepsilon_i - \Delta_i. \end{aligned}$$

The elementary inequality $|\sup_f a_f - \sup_f b_f| \leq \sup_f |a_f - b_f|$ gives directly

$$\left| W_1(\widehat{P}_i^{\mathrm{tr}}, \widehat{Q}_i^{\mathrm{tr}}) - W_1(P, Q) \right| \leq \Delta_i. \tag{47}$$

Consequently,

$$m_i \geq W_1(P, Q) - 2\Delta_i - \varepsilon_i. \tag{48}$$

Equations (47), (46), (48), and (45) prove $m_i \to W_1(P, Q)$ almost surely. Cesàro's lemma proves (14). When $P, Q \in \mathcal{P}_2$ and $P \neq Q$, Corollary F.1 verifies the second-moment condition and Theorem C.8 gives (15). □

Proof of Theorem 3.5. Use the same Δ_i as in (44); it converges to zero almost surely by (45). Since $\mathcal{H}_i \subseteq \mathcal{F}_1$, the pathwise deviation bound

$$\sup_{f \in \mathcal{H}_i} \left| \{(\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) - (P - Q)\} f \right| \leq \Delta_i \tag{49}$$

holds even though $\mathcal{H}_i$ may be random and trained on the same past data.

Let

$$\widehat{D}_i = \sup_{f \in \mathcal{H}_i} (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}}) f.$$

By (49),

$$|\widehat{D}_i - D_{\mathcal{H}_i}(P, Q)| \leq \Delta_i.$$

Approximate optimization and another application of (49) give

$$\begin{aligned} m_i = (P - Q)f_i &\geq (\widehat{P}_i^{\mathrm{tr}} - \widehat{Q}_i^{\mathrm{tr}})f_i - \Delta_i \\ &\geq \widehat{D}_i - \varepsilon_i - \Delta_i \\ &\geq D_{\mathcal{H}_i}(P, Q) - 2\Delta_i - \varepsilon_i. \end{aligned}$$

Because $f_i \in \mathcal{H}_i$,

$$m_i \leq D_{\mathcal{H}_i}(P, Q).$$

Under (16), the upper and lower bounds imply $m_i \to D(P, Q)$ almost surely. If $P, Q \in \mathcal{P}_2$, Theorem C.8 gives consistency. If (17) holds, then $D_{\mathcal{H}_i}(P, Q) \to W_1(P, Q)$ and the same squeeze proves the final assertion. □

Proof of Proposition C.3. The stabilizer Z_i is conditionally centered under every P, Q. As in the proof of Corollary F.1, there is a constant C such that

$$\mathbb{E} Z_i^2 \leq C(\tau_i^2 + \rho_i^2). \tag{50}$$

Orthogonality of martingale differences gives

$$\begin{aligned} \mathbb{E}\left| \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_i \right|^2 &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} Z_i^2 \\ &\leq \frac{C}{n} \sum_{i=1}^n (\tau_i^2 + \rho_i^2) \to 0 \end{aligned}$$

by Cesàro's lemma, proving (24). Also, by Markov's inequality and (50),

$$\mathbb{E}\left(\frac{1}{n} \sum_{i=1}^n Z_i^2 \right) \leq \frac{C}{n} \sum_{i=1}^n (\tau_i^2 + \rho_i^2) \to 0,$$

which proves (25).

Finally,

$$\left|\frac{1}{n}\sum_{i=1}^{n}(\xi_i^2 - U_i^2)\right| \le \frac{2}{n}\left|\sum_{i=1}^{n} U_i Z_i\right| + \frac{1}{n}\sum_{i=1}^{n} Z_i^2$$
$$\le 2\left(\frac{1}{n}\sum_{i=1}^{n} U_i^2\right)^{1/2}\left(\frac{1}{n}\sum_{i=1}^{n} Z_i^2\right)^{1/2} + \frac{1}{n}\sum_{i=1}^{n} Z_i^2.$$

The first factor is $O_{\mathbb{P}}(1)$ by assumption and the second converges to zero, proving (26). □

Proof of Corollary C.4. For $a < 1/2$,

$$\sum_{i=1}^{n} i^{-2a} \asymp n^{1-2a},$$

and the analogous relation holds for γ. Equation (10) proves the asserted divergence; Theorem C.1 gives Gaussian calibration when (20) is also imposed.

For the unbounded statement, suppose the dominating variance component has exponent $s_\star$, so $b_{n,\pi,s}^2 \asymp n^{1-2s_\star}$. The second condition in (23) becomes

$$n^{1-r_\delta(1-2s_\star)} \to 0,$$

which is equivalent to

$$r_\delta(1-2s_\star) > 1.$$

Since $r_\delta = (2+\delta_0)/2$, rearrangement gives

$$s_\star < \frac{\delta_0}{2(2+\delta_0)}.$$

□

G Proofs for Local Alternatives and Reductions

Proof of Theorem C.5. Let

$$M_n = \sum_{i=1}^{n} e_{ni}.$$

By the martingale triangular-array central limit theorem of Hall and Heyde (1980), conditions (28) and (29) imply

$$\frac{M_n}{a_n} \xrightarrow{d} N(0,1).$$

Combining this with (27) gives

$$\frac{S_n}{a_n} = \frac{M_n}{a_n} + \frac{1}{a_n}\sum_{i=1}^{n} m_{ni} \xrightarrow{d} N(\lambda, 1). \tag{51}$$

It remains to show $V_n/a_n^2 \xrightarrow{\mathbb{P}} 1$. Expand

$$\frac{V_n}{a_n^2} = \frac{1}{a_n^2}\sum_{i=1}^{n} e_{ni}^2 + \frac{2}{a_n^2}\sum_{i=1}^{n} m_{ni}e_{ni} + \frac{1}{a_n^2}\sum_{i=1}^{n} m_{ni}^2. \tag{52}$$

By (28) and (30),

$$\frac{1}{a_n^2}\sum_{i=1}^{n} e_{ni}^2 \xrightarrow{\mathbb{P}} 1.$$

The final term in (52) converges to zero by (31). For the cross term, Cauchy–Schwarz gives

$$\left|\frac{1}{a_n^2}\sum_{i=1}^{n} m_{ni}e_{ni}\right| \le \left(\frac{1}{a_n^2}\sum_{i=1}^{n} m_{ni}^2\right)^{1/2}\left(\frac{1}{a_n^2}\sum_{i=1}^{n} e_{ni}^2\right)^{1/2} \xrightarrow{\mathbb{P}} 0.$$

Therefore $V_n/a_n^2 \xrightarrow{\mathbb{P}} 1$. Slutsky's theorem applied to (51) yields

$$T_n = \frac{S_n/a_n}{\sqrt{V_n/a_n^2}} \xrightarrow{d} N(\lambda, 1).$$

The power formula follows by evaluating the upper tail at $z_{1-\alpha}$. □

Proof of Corollary C.6. We first verify conditional Lindeberg. For every $\varepsilon > 0$,

$$\begin{aligned}
&\frac{1}{a_n^2}\sum_{i=1}^{n} \mathbb{E}\{e_{ni}^2 \mathbf{1}(|e_{ni}| > \varepsilon a_n) \mid \mathcal{G}_{n,i-1}\} \\
&\quad \le \frac{1}{\varepsilon^\delta a_n^{2+\delta}} \sum_{i=1}^{n} \mathbb{E}(|e_{ni}|^{2+\delta} \mid \mathcal{G}_{n,i-1}) \\
&\quad \le \frac{Cn}{\varepsilon^\delta a_n^{2+\delta}}.
\end{aligned}$$

For $a_n = \sigma\sqrt{n}$, the bound is $O(n^{-\delta/2})$. For the more general statement, replacing δ by $\delta_0 = \min(\delta, 2)$ and using conditional Jensen if necessary gives convergence under (32).

Next define

$$D_{ni} = e_{ni}^2 - q_{ni}.$$

Then $(D_{ni}, \mathcal{G}_{n,i})$ is a martingale-difference array. Put

$$r = 1 + \frac{\delta_0}{2} \in (1, 2].$$

The same argument used in the proof of Theorem C.2 gives

$$\sup_{n,i} \mathbb{E}|D_{ni}|^r < \infty.$$

By Lemma E.3,

$$\mathbb{E}\left|\sum_{i=1}^{n} D_{ni}\right|^r \le Cn.$$

Therefore, for every $\varepsilon > 0$,

$$\mathbb{P}\left(\left|\frac{1}{a_n^2}\sum_{i=1}^{n} D_{ni}\right| > \varepsilon\right) \le \frac{Cn}{\varepsilon^r a_n^{2r}} = \frac{Cn}{\varepsilon^r a_n^{2+\delta_0}} \to 0.$$

This proves (30). In the root-n case, the last bound is $O(n^{-\delta_0/2})$. □

Proof of Proposition C.10. Part (i) follows by substituting the stated choices into (2). Compactness implies bounded diameter and finite moments, while Lemma D.3 gives $\kappa_\pi^{\mathrm{A}} > 0$ for a non-point null. Thus $b_{n,\pi,\mathrm{A}}^2 = n\tau^2\kappa_\pi^{\mathrm{A}}$ up to the immaterial first index. If the displayed average critic variance converges to v_π, then

$$\frac{A_n^2}{n} \xrightarrow{\mathbb{P}} 2v_\pi + \tau^2\kappa_\pi^{\mathrm{A}} > 0,$$

so (20) holds with the corresponding deterministic a_n^2. The null conclusion follows from Theorem C.1; the drift and consistency conclusions follow separately from Theorem 3.4.

For part (ii), the score is $\xi_i = f(X_i) - f(Y_i)$. Thus S_n is the sum of paired scalar differences and V_n is their uncentered observed sum of squares. Under the null the differences have mean zero, so this is the usual self-normalized mean statistic. If their variance is positive, the standard limit follows from the classical self-normalized CLT.

For part (iii), swapping the fresh observations negates any real-valued witness difference and proves conditional symmetry. Integrable witnesses also give null centering. The proof of Theorem 3.5 bounds the population optimization gap by twice the uniform empirical deviation plus the training

error. For subclasses of $\mathcal{F}_1$, empirical W_1 convergence supplies that deviation bound; for other classes it must be verified for the chosen witnesses.

For part (iv), if $d_c = cd$, then cf_i is one-Lipschitz relative to d_c whenever f_i is one-Lipschitz relative to d. Both the anchor distance functions and the pair distance are multiplied by c. With $\rho_i = 0$, every score becomes $\xi_{i,c} = c\xi_i$, so

$$T_{n,c} = \frac{cS_n}{\sqrt{c^2 V_n}} = T_n.$$

For the e-process,

$$E^+_{n,c}(\eta/c) = \frac{2}{\sqrt{1 + (\eta/c)^2 c^2 V_n}} \exp\left\{ \frac{(\eta/c)^2 c^2 S_n^2}{2\{1 + (\eta/c)^2 c^2 V_n\}} \right\} \times \Phi\left\{ \frac{(\eta/c) c S_n}{\sqrt{1 + (\eta/c)^2 c^2 V_n}} \right\} = E^+_n(\eta).$$

If the external jitter scale is also multiplied by c, the same score identity holds. □