

# Adversarial Training for Deep Hedging in Nonstationary Markets

Philipp J. Schneider<sup>\*,1</sup> Lukas Looser<sup>\*,1</sup> Antoine Garin<sup>†,1</sup> Shuhan Liu<sup>†,2</sup> Daniel Kuhn<sup>1</sup>

<sup>1</sup>Risk Analytics and Optimization Chair, EPFL, Lausanne, Switzerland

<sup>2</sup>University of Waterloo, Waterloo, Canada

October 7, 2026

## Abstract

Deep hedging learns trading policies from historical or simulated market trajectories, yet under nonstationarity these training paths may not represent future market conditions. We propose WRAP (Wasserstein–Reweighting Adversarial Perturbation), a drift-aware adversarial training framework derived from a two-budget distributionally robust optimization (DRO) formulation. The formulation is anchored to a weighted empirical reference distribution whose fixed baseline weights are chosen to balance sampling uncertainty against temporal drift. Around this reference distribution, the ambiguity set addresses two complementary forms of distributional misspecification by allowing an adversary to reweight the observed trajectories subject to a  $\phi$ -divergence constraint and perturb their paths subject to an optimal-transport (OT) constraint. We derive a joint first-order expansion in which the leading-order increase over the nominal expected loss decomposes into a reweighting contribution determined by the dispersion of hedging losses across trajectories and a transport contribution determined by the sensitivity of the loss to path perturbations. This expansion yields an explicit finite-dimensional adversarial attack that replaces the distributional inner supremum with a tractable first-order approximation. Across stationary and nonstationary Heston dynamics and a generalized affine diffusion (GAD), the experiments show complementary benefits from reweighting and transport, with joint adversarial training providing the largest gains under nonstationarity.

## 1 Introduction

A hedging strategy dynamically trades financial assets to reduce the risk associated with a contingent claim. Deep hedging parameterizes such strategies with neural networks and learns the resulting sequence of trading decisions from simulated or historical market trajectories (Buehler et al., 2019). The out-of-sample performance of the learned policy therefore depends critically on whether the training distribution represents the market conditions encountered over the hedging horizon. Unlike a static prediction, a hedge consists of a sequence of state-dependent trading decisions, so changes in market dynamics may affect positions throughout the life of the claim (Cont, 2006; Glasserman & Xu, 2014). A policy may consequently achieve a low in-sample hedging loss yet remain vulnerable to market dynamics poorly represented in the training data.

---

<sup>\*</sup>These authors contributed equally to this work.

<sup>†</sup>These authors contributed equally to this work.

Corresponding author: [philipp.schneider@epfl.ch](mailto:philipp.schneider@epfl.ch)

Distributionally robust optimization (DRO) and adversarial training protect a policy against perturbations around a prespecified empirical distribution. Maintaining a representative training distribution is particularly difficult under nonstationarity, where returns, volatility, and dependence structures evolve over time (Stărică & Granger, 2005; Ang & Timmermann, 2012; Capponi et al., 2025). A larger historical sample reduces sampling uncertainty but also includes older observations that may be less representative of the next hedging period, whereas emphasizing recent observations improves temporal relevance at the cost of greater sampling variability. Constructing the reference distribution entails a fundamental trade-off between *statistical precision* and *temporal relevance*. Robust deep hedging under nonstationarity must therefore address two questions: *how should the reference distribution be constructed from temporally heterogeneous observations*, and *how should the hedge be protected against residual misspecification around that distribution*?

We distinguish three channels through which the training distribution may fail to represent the next hedging period. First, historical trajectories differ in their temporal relevance, as older observations may be less informative about future market dynamics. Second, the reference distribution may assign inappropriate probability mass to the sampled trajectories or the regimes they represent. Third, future trajectories may differ from those observed in the training sample. We address these channels through *non-uniform baseline weights*, *adversarial probability reweighting*, and *optimal-transport perturbations*, respectively. The baseline weights are selected independently of the hedging losses and determine the drift-aware center of the ambiguity set. Adversarial reweighting depends on the current policy and reallocates probability mass toward trajectories on which that policy incurs high losses, whereas optimal-transport perturbations move the trajectories themselves in locally adverse directions. These three components constitute our main contributions, summarized as follows.

**Drift-aware reference distributions for deep hedging.** We provide a statistical foundation for constructing reference distributions from nonstationary market data. Our finite-sample guarantees separate sampling uncertainty, governed by the effective sample size induced by the baseline weights, from temporal misspecification, governed by a weighted drift index that captures the influence assigned to older observations. This decomposition formalizes the trade-off between statistical precision and temporal relevance and provides a basis for calibrating the Wasserstein radius to both weight concentration and temporal drift.

**First-order expansion of the robust objective.** We derive a joint first-order expansion of the worst-case hedging loss under probability and trajectory misspecification. The expansion shows that the  $\phi$ -divergence component penalizes the dispersion of hedging losses across trajectories (*variance*), whereas the optimal-transport component penalizes the sensitivity of the hedging loss to perturbations of market paths (*variation*). We further show that a structured empirical attack attains the same first-order value as the full two-budget robust problem.

**WRAP: two-budget distributional adversarial training.** Building on this expansion, we introduce the Wasserstein–Reweighting Adversarial Perturbation (WRAP), an implementable training method for the two-budget robust problem. The two budgets govern complementary deviations from the drift-aware reference distribution: adversarial probability reweighting and trajectory perturbations. The latter are defined through a time- and component-weighted ground metric that permits separate scaling across trading dates and market variables. We derive a one-step attack from the first-order solution and study a multistep variant that iteratively refines the adversarial distribution.

Existing approaches address individual components of this problem. Closest to our setting, He et al. (2025) develop Wasserstein attacks for deep hedging, Keehan et al. (2025) construct drift-aware Wasserstein sets around weighted empirical laws, and Namkoong & Duchi (2016) use  $\phi$ -divergence DRO to adversarially reweight fixed empirical observations.

Several approaches combine likelihood and optimal-transport perturbations, including holistic

robustness, joint OT-divergence constructions, and OT-regularized divergences (Bennouna et al., 2022; Blanchet et al., 2023; Liu et al., 2023; Birrell & Ebrahimi, 2025). We derive a joint first-order expansion for a sequential, path-dependent objective and construct an explicit transport-and-reweighting attack that is first-order accurate for the full DRO problem. Appendix A provides the extended literature discussion.

## 2 Deep Hedging Under Nonstationarity

We consider a discrete-time financial market with trading dates  $t = 0, \dots, T$ . The market contains  $r$  tradable assets with price process  $S = (S_t)_{t=0}^T$ , where  $S_t \in \mathbb{R}^r$ . The observed market process  $X = (X_t)_{t=0}^T$  is a random element of an admissible trajectory space  $\Xi \subseteq \mathbb{R}^{d(T+1)}$ . Here  $X_t \in \mathbb{R}^d$  contains the tradable asset prices together with any exogenous market features used for hedging, such as volatility estimates or macroeconomic indicators; in particular,  $d \geq r$ . We take  $S_t$  to be the first  $r$  components of  $X_t$  and write  $X_{0:t} := (X_0, \dots, X_t)$ . Given  $N$  observed or simulated market trajectories, we denote the  $n$ -th trajectory by  $X_n = (X_{n,0}, \dots, X_{n,T}) \in \Xi$ .

The agent aims to hedge a contingent claim with terminal payoff function  $\Phi : \Xi \rightarrow \mathbb{R}$  at maturity. Let  $\theta = (\theta_0, \dots, \theta_{T-1}) \in \Theta \subseteq \mathbb{R}^{kT}$  parametrize  $T$  neural networks. At time  $t = 0, \dots, T-1$ , the network parametrized by  $\theta_t$  maps  $X_{0:t}$  to holdings  $\delta_t^\theta(X_{0:t}) \in \mathbb{R}^r$  and thus depends on  $\theta$  only through  $\theta_t$ . Since each decision depends only on the observed history  $X_{0:t}$ , the strategy is *non-anticipative* and fully characterized by  $\theta \in \Theta$ . Assuming a self-financing strategy and a fixed initial premium  $p_0 \in \mathbb{R}$  for selling the derivative, the terminal *profit and loss* (PnL) is given by

$$\text{PnL}(\theta, X) = p_0 + \sum_{t=0}^{T-1} \delta_t^\theta(X_{0:t})^\top (S_{t+1} - S_t) - \Phi(X). \quad (1)$$

Following Buehler et al. (2019), the risk of the terminal profit and loss  $Z$  is assessed through an Optimized Certainty Equivalent (OCE) risk measure (Ben-Tal & Teboulle, 1986, 2007),

$$\rho(Z) = \inf_{m \in \mathbb{R}} \{m + \mathbb{E}[\ell(-Z - m)]\}, \quad (2)$$

where  $\ell : \mathbb{R} \rightarrow \mathbb{R}$  is a convex, nondecreasing loss function satisfying  $\ell(0) = 0$  and  $\ell(u) \geq u$  for all  $u \in \mathbb{R}$ . To train a hedging policy, we optimize jointly over the hedging network parameters and the OCE offset by setting  $\vartheta = (\theta, m) \in \Theta \times \mathbb{R}$  and defining the pathwise deep-hedging loss

$$\ell_{\text{DH}}(\vartheta; X) := m + \ell(-\text{PnL}(\theta, X) - m). \quad (3)$$

The numerical experiments consider Conditional Value at Risk (CVaR) and entropic risk (*cf.* Appendix D.1). Let  $\mathcal{P}(\Xi)$  denote the set of Borel probability measures on  $\Xi$ . For a reference measure  $\mu \in \mathcal{P}(\Xi)$ , we train a nominal deep-hedging model by optimizing

$$\min_{\vartheta \in \Theta \times \mathbb{R}} \mathbb{E}_{X \sim \mu} [\ell_{\text{DH}}(\vartheta; X)]. \quad (4)$$

Further details on the market setup and regularity are provided in Appendix B.1.

### 2.1 Trajectory Space and Wasserstein Robust Value

The training problem in (4) evaluates the hedge under a fixed reference law  $\mu$ . To account for distribution shifts and sampling error, we replace this fixed law with a Wasserstein ambiguity set centered at  $\mu$ . Specifying the Wasserstein ambiguity set requires a transportation cost between

market trajectories. We therefore equip the admissible trajectory set  $\Xi \subseteq \mathbb{R}^{d(T+1)}$  with an anisotropic path metric that allows these costs to vary across market variables and trading dates. Let  $p \in (1, \infty)$ , fix time weights  $\kappa_t > 0$  for  $t = 0, \dots, T$ , and component scales  $\lambda_i > 0$  for  $i = 1, \dots, d$ . For an observed trajectory  $X \in \Xi$  and a perturbed trajectory  $\hat{X} \in \Xi$ , define

$$d_{\kappa, \lambda}(X, \hat{X}) := \left( \sum_{i=1}^d \left( \lambda_i \max_{t=0, \dots, T} \kappa_t |X_t^i - \hat{X}_t^i| \right)^p \right)^{1/p}. \quad (5)$$

The inner  $\ell_\infty$  term measures the largest weighted displacement over time within each component, while the outer  $\ell_p$  norm aggregates these component-wise displacements. The scales  $\lambda_i$  and weights  $\kappa_t$  control how perturbations across market variables and trading dates enter the transportation cost. We impose the following condition on the resulting metric space, which ensures finite  $p$ -th moments for all probability laws on  $\Xi$  with respect to  $d_{\kappa, \lambda}$ .

**Assumption 2.1** (Path space). The admissible trajectory set  $\Xi$ , equipped with the metric  $d_{\kappa, \lambda}$  in (5), is a Polish metric space with finite diameter:  $\text{diam}_{\kappa, \lambda}(\Xi) := \sup_{X, \hat{X} \in \Xi} d_{\kappa, \lambda}(X, \hat{X}) < \infty$ .

For  $p \in (1, \infty]$ , let  $W_p$  denote the  $p$ -Wasserstein distance induced by  $d_{\kappa, \lambda}$ . For a reference law  $\mu \in \mathcal{P}(\Xi)$  and radius  $\varepsilon \geq 0$ , define the Wasserstein ambiguity set as the ball  $\mathcal{B}_\varepsilon(\mu) := \{\eta \in \mathcal{P}(\Xi) : W_p(\eta, \mu) \leq \varepsilon\}$ . The Wasserstein robust value  $V_\vartheta$  and the associated deep-hedging problem are

$$\inf_{\vartheta \in \Theta \times \mathbb{R}} \left\{ V_\vartheta(\varepsilon) := \sup_{\eta \in \mathcal{B}_\varepsilon(\mu)} \mathbb{E}_{X \sim \eta} [\ell_{\text{DH}}(\vartheta; X)] \right\}. \quad (6)$$

## 2.2 Drift-Aware Coverage of the Next-Period Market Law

Under temporal nonstationarity, the reference law and Wasserstein radius should reflect the evolution of the data-generating distribution. Let  $X_1, \dots, X_N \in \Xi$  denote historical trajectories ordered from oldest to most recent, with  $X_n \sim P_n \in \mathcal{P}(\Xi)$ , and let  $P_{N+1}$  denote the unknown next-period market law. We model temporal drift through a uniform bound  $\varrho$  on one-step changes in the market law.

**Assumption 2.2** (Nonstationary sampling and bounded drift). The trajectories  $X_1, \dots, X_N$  are independent, with  $X_n \sim P_n$  for  $n = 1, \dots, N$ . Moreover, there exists  $\varrho \geq 0$  such that

$$W_\infty(P_n, P_{n+1}) \leq \varrho, \quad n = 1, \dots, N.$$

In applications,  $\varrho$  may be calibrated using regime-switching models, change-point methods, or rolling estimates of distributional distance. Under this drift model, older observations may be less representative of the next-period law because they are separated from  $P_{N+1}$  by more distributional shifts. The uniform empirical measure  $\mu := \frac{1}{N} \sum_{n=1}^N \delta_{X_n}$  does not account for this difference in recency. We therefore use the weighted empirical reference measure  $\mu_w := \sum_{n=1}^N w_n \delta_{X_n}$ , where  $\delta_x$  denotes the Dirac measure at  $x$  and  $w \in \Delta_N := \{w \in \mathbb{R}_+^N : \sum_{n=1}^N w_n = 1\}$ . The weights  $w_n$  allow older observations to be downweighted, are chosen independently of the hedging losses, and remain fixed during adversarial training. Choosing  $w$  entails a trade-off between recency and statistical efficiency. Sampling uncertainty is governed by the effective sample size  $N_{\text{eff}}(w) := 1 / \sum_{n=1}^N w_n^2$ , whereas exposure to cumulative drift is governed by  $D_p(w) := (\sum_{n=1}^N w_n (N - n + 1)^p)^{1/p}$ . Here  $N - n + 1$  counts the number of one-step distributional shifts from  $P_n$  to  $P_{N+1}$ , so that  $D_p(w)\varrho$  controls the resulting drift bias. Concentrating weight on recent trajectories reduces drift bias, whereas spreading weight across more trajectories increases  $N_{\text{eff}}(w)$  and reduces sampling uncertainty. The following theorem links the baseline weights to a calibrated radius and a next-period risk bound.

**Theorem 2.3** (Drift-aware coverage and next-period risk). *Let  $p \in (1, \infty)$  and suppose Assumptions 2.1 and 2.2 hold. There exist constants  $c_1, c_2 > 0$  and  $\nu \in (0, 1/2)$ , depending only on  $d(T+1)$ ,  $p$ ,  $\text{diam}_{\kappa, \lambda}(\Xi)$ , such that the following holds for every  $w \in \Delta_N$  and  $\beta \in (0, 1)$ . Define*

$$\varepsilon_\varrho(w, \beta) := D_p(w)\varrho + \left( c_2 N_{\text{eff}}(w)^{-\nu} + \sqrt{\frac{\log(1/\beta)}{c_1 N_{\text{eff}}(w)}} \right)^{1/p}.$$

*With probability at least  $1 - \beta$ , we have  $P_{N+1} \in \mathcal{B}_{\varepsilon_\varrho(w, \beta)}(\mu_w)$  and  $\mathbb{E}_{X \sim P_{N+1}}[\ell_{\text{DH}}(\vartheta; X)] \leq V_\vartheta(\varepsilon_\varrho(w, \beta))$  for every  $\vartheta \in \Theta \times \mathbb{R}$ , where  $V_\vartheta$  uses reference law  $\mu_w$ .*

The coverage bound underlying Theorem 2.3 specializes the nonstationary Wasserstein concentration result of Keehan et al. (2025) to the trajectory space  $(\Xi, d_{\kappa, \lambda})$ . The proofs are given in Appendix C.2. For fixed weights, the theorem yields a calibrated radius. Alternatively, for a target radius  $\varepsilon_\varrho$ , we select the weights by minimizing the failure-probability bound by solving

$$w \in \arg \max_{w \in \Delta_N} \left\{ N_{\text{eff}}(w)(\varepsilon_\varrho - D_p(w)\varrho)_+^{2p} \right\}.$$

The optimizer therefore shifts mass toward more recent observations as the drift bound  $\varrho$  increases, whereas smaller values of  $\varrho$  favor a broader allocation across the historical sample. Moreover, Keehan et al. (2025, Theorem 2) show that this generally nonconvex problem admits an optimal solution of the truncated-polynomial form  $w_n = (a_1 - a_2(N - n + 1)^p)_+$ , for some  $a_1, a_2 \geq 0$  such that the  $w_n$  sum to one. The resulting weights define the drift-aware empirical distribution  $\mu_w$  that serves as the fixed center of the two-budget adversarial problem.

### 3 Two-Budget OT- $\phi$ Adversarial Training

The Wasserstein formulation in Section 2 perturbs observed trajectories but leaves their relative probabilities unchanged. To also capture uncertainty in scenario probabilities, we adapt the joint transport-and-reweighting construction of Blanchet et al. (2023, Section 5) to weighted market trajectories, with separate budgets for path perturbations and probability reweighting.

#### 3.1 A Two-Budget Ambiguity Set

We represent both forms of uncertainty through joint laws of reference and perturbed trajectories. For a coupling  $\pi \in \mathcal{P}(\Xi \times \Xi)$ , let  $\pi_1$  and  $\pi_2$  denote its left and right marginals, respectively. Probability reweighting is constrained by  $\phi$ -divergence.

**Definition 3.1** (Entropy function and  $\phi$ -divergence). A function  $\phi : [0, \infty) \rightarrow [0, \infty]$  is an *entropy function* if it is proper, convex, lower semicontinuous, and satisfies  $\phi(1) = 0$ . For  $\pi', \pi \in \mathcal{P}(\Xi \times \Xi)$ , define the  $\phi$ -divergence of  $\pi'$  with respect to  $\pi$  induced by  $\phi$  as

$$D_\phi(\pi' \parallel \pi) := \begin{cases} \int_{\Xi \times \Xi} \phi\left(\frac{d\pi'}{d\pi}\right) d\pi, & \pi' \ll \pi, \\ +\infty, & \pi' \not\ll \pi. \end{cases}$$

Standard choices include  $\phi_{\text{KL}}(u) = u \log u - u + 1$  for the Kullback–Leibler divergence and  $\phi_{\chi^2}(u) = (u - 1)^2$  for the Pearson  $\chi^2$ -divergence. Combining probability reweighting with transportation yields the following ambiguity set.

**Definition 3.2** (Two-budget OT- $\phi$  ambiguity set). For transport budget  $\varepsilon \geq 0$  and reweighting budget  $\tau \geq 0$ , define

$$\mathcal{B}_{\varepsilon, \tau}(\mu_w) := \left\{ \eta \in \mathcal{P}(\Xi) : \begin{array}{l} \text{there exist } \pi, \pi' \in \mathcal{P}(\Xi \times \Xi) \text{ such that} \\ \pi_1 = \mu_w, \quad \pi'_2 = \eta, \quad D_\phi(\pi' \| \pi) \leq \tau, \\ \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi'(x, \hat{x}) \leq \varepsilon^p \end{array} \right\}. \quad (7)$$

The first marginal of  $\pi$  anchors the construction at the reference law  $\mu_w$ , while the second marginal of  $\pi'$  defines the adversarial law  $\eta$ . The divergence budget  $\tau$  limits probability reweighting relative to the reference coupling, whereas the transport budget  $\varepsilon$  limits displacement in the trajectory space.

*Remark 3.3* (Coverage inheritance). For every  $\tau \geq 0$ ,  $\mathcal{B}_\varepsilon(\mu_w) \subseteq \mathcal{B}_{\varepsilon, \tau}(\mu_w)$ , since any Wasserstein-feasible coupling is feasible in (7) by taking  $\pi' = \pi$ . Hence the transport radius calibrated in Section 2.2 retains its next-period coverage guarantee under the two-budget formulation. This inclusion does not justify reducing the calibrated transport radius in exchange for a reweighting budget.

The corresponding two-budget robust value  $V_\vartheta(\varepsilon, \tau)$  and the associated deep-hedging problem are

$$\inf_{\vartheta \in \Theta \times \mathbb{R}} \left\{ V_\vartheta(\varepsilon, \tau) := \sup_{\eta \in \mathcal{B}_{\varepsilon, \tau}(\mu_w)} \mathbb{E}_{X \sim \eta} [\ell_{\text{DH}}(\vartheta; X)] \right\}. \quad (8)$$

Observations assigned zero baseline weight are excluded from the support of  $\mu_w$ , so subsequent sums may be restricted to indices with  $w_n > 0$ . The formulation contains the two corresponding single-budget models as special cases:  $\tau = 0$  reduces to Wasserstein DRO, while  $\varepsilon = 0$  yields pure  $\phi$ -divergence reweighting on the empirical support.

### 3.2 First-Order Approximation

For small transport and reweighting budgets, we derive a tractable local approximation of the inner problem in (8). The transport contribution is determined by the dual geometry of the path metric.

**Proposition 3.4.** (Dual geometry of the weighted path metric) *Let  $p \in (1, \infty)$ ,  $q = p/(p-1)$ , and  $\kappa_t, \lambda_i > 0$ . For  $\Delta X, G \in \mathbb{R}^{d(T+1)}$ , define*

$$\|\Delta X\|_{\kappa, \lambda, p} := \left( \sum_i (\lambda_i \max_t \kappa_t |\Delta X_t^i|)^p \right)^{1/p}, \quad \|G\|_{\kappa, \lambda, p}^* := \left( \sum_i \left( \sum_t \frac{|G_t^i|}{\lambda_i \kappa_t} \right)^q \right)^{1/q}. \quad (9)$$

*The first is a norm inducing the metric in (5) through  $d_{\kappa, \lambda}(X, \hat{X}) = \|X - \hat{X}\|_{\kappa, \lambda, p}$ , and the second is its dual norm under the Euclidean inner product. Moreover, there exists a Borel map  $h : \mathbb{R}^{d(T+1)} \rightarrow \mathbb{R}^{d(T+1)}$ , with  $h(0) = 0$ , such that  $\|h(G)\|_{\kappa, \lambda, p} = 1$  and  $\langle G, h(G) \rangle = \|G\|_{\kappa, \lambda, p}^*$  whenever  $G \neq 0$ .*

For  $n = 1, \dots, N$ , write  $L_n(\vartheta) := \ell_{\text{DH}}(\vartheta; X_n)$  and  $g_n(\vartheta) := \nabla_X \ell_{\text{DH}}(\vartheta; X_n)$ . The local approximation of the robust value  $V_\vartheta(\varepsilon, \tau)$  in (8) requires the loss to admit the following first-order expansion.

**Assumption 3.5** (Loss regularity). The trajectories  $X_n$  lie in the interior of  $\Xi$ , and for every  $\vartheta \in \Theta \times \mathbb{R}$ , the loss  $\ell_{\text{DH}}(\vartheta; \cdot)$  is Borel measurable, bounded on  $\Xi$  and admits at each  $X_n$  the expansion

$$\ell_{\text{DH}}(\vartheta; X_n + \Delta) = L_n(\vartheta) + \langle g_n(\vartheta), \Delta \rangle + o(\|\Delta\|_{\kappa, \lambda, p}), \quad \|\Delta\|_{\kappa, \lambda, p} \rightarrow 0.$$

The local conditions provide the first-order loss expansion and ensure that perturbations remain feasible near the observed trajectories, while boundedness controls losses away from the empirical support. For the reweighting term, we restrict attention to *locally regular* entropy functions:  $\phi$  is twice continuously differentiable near 1, with  $\phi''(1) > 0$ . This local condition is needed only for the small-budget expansion and is satisfied by the Kullback–Leibler and Pearson  $\chi^2$  divergences above. Let  $\bar{L}_w(\vartheta) := \sum_{n=1}^N w_n L_n(\vartheta)$  denote the weighted empirical loss, and define the loss dispersion and transport sensitivity by

$$\sigma_w(\vartheta) := \left( \sum_{n=1}^N w_n (L_n(\vartheta) - \bar{L}_w(\vartheta))^2 \right)^{1/2}, \quad \Upsilon_w(\vartheta) := \left( \sum_{n=1}^N w_n \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} \right)^{1/q}.$$

**Theorem 3.6.** (Joint first-order approximation) *Suppose  $\phi$  is locally regular,  $p \in (1, \infty)$  and Assumption 3.5 holds. For every fixed  $\vartheta \in \Theta \times \mathbb{R}$ , as  $(\varepsilon, \tau) \rightarrow (0, 0)$*

$$V_\vartheta(\varepsilon, \tau) = \bar{L}_w(\vartheta) + \Upsilon_w(\vartheta) \varepsilon + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\varepsilon + \sqrt{\tau}). \quad (10)$$

The expansion separates the leading-order effects of the two uncertainty mechanisms. The coefficient  $\Upsilon_w(\vartheta)$  measures sensitivity to local deformation of trajectories under  $d_{\kappa, \lambda}$ , whereas  $\sqrt{2/\phi''(1)} \sigma_w(\vartheta)$  measures sensitivity to reallocating probability mass across scenarios. Thus, transport captures within-scenario sensitivity and reweighting captures cross-scenario loss heterogeneity. Cross-effects enter only at higher order, although the two mechanisms remain coupled in the exact adversarial problem because the transport allocation depends on the reweighted scenario probabilities. Along  $\tau = 0$  and  $\varepsilon = 0$ , respectively, (10) recovers the weighted Wasserstein and weighted  $\phi$ -divergence first-order limits (Bartl et al., 2021; Duchi et al., 2021).

### 3.3 First-Order Adversarial Construction

To attain the first-order expansion in Theorem 3.6, we construct an explicit first-order attack. Fix  $\vartheta$  and let  $Y_n := L_n - \bar{L}_w$  denote the centered loss of the  $n$ -th trajectory. If  $\sigma_w > 0$ , define

$$\alpha_\tau := \sup \left\{ \alpha \geq 0 : 1 + \alpha Y_n \geq 0 \text{ for all } n, \quad \sum_{n=1}^N w_n \phi(1 + \alpha Y_n) \leq \tau \right\},$$

and set  $a_n := 1 + \alpha_\tau Y_n$ . If  $\sigma_w = 0$ , set  $\alpha_\tau = 0$  and  $a_n = 1$ . Since  $1 + \alpha_\tau Y_n \geq 0$  for all  $n$  and  $\sum_{n=1}^N w_n Y_n = 0$ , the reweighted probabilities  $\hat{w}_n := w_n a_n$  remain nonnegative and normalized. The sign of the centered loss  $Y_n = L_n - \bar{L}_w$  determines whether trajectory  $n$  gains or loses probability mass and  $\alpha_\tau$  sets the largest feasible step. The resulting tilt therefore reallocates mass toward trajectories with above-average loss while preserving normalization and satisfying the budget.

Let  $\Upsilon_{\hat{w}} := (\sum_{n=1}^N \hat{w}_n \|g_n\|_{\kappa, \lambda, p}^{*q})^{1/q}$  denote the transport sensitivity under the reweighted probabilities. If  $\Upsilon_{\hat{w}} > 0$ , define

$$\hat{X}_n := X_n + \varepsilon h(g_n) \|g_n\|_{\kappa, \lambda, p}^{*(q-1)} \Upsilon_{\hat{w}}^{1-q}; \quad (11)$$

otherwise, set  $\hat{X}_n = X_n$ . The direction  $h(g_n)$  maximizes the local loss increase per unit transport distance, while  $\|g_n\|_{\kappa, \lambda, p}^{*(q-1)}$  allocates larger perturbations to trajectories with greater local sensitivity. The normalization by  $\Upsilon_{\hat{w}}$  scales the perturbations to the transport budget under the reweighted probabilities. Thus, the construction first tilts scenario probabilities according to realized loss and then allocates the transport budget according to local pathwise sensitivity. The probability multipliers and perturbed trajectories define the adversarial law  $\hat{\eta}_{\varepsilon, \tau} := \sum_{n=1}^N \hat{w}_n \delta_{\hat{X}_n}$ .

---

**Algorithm 1** Wasserstein Reweighting Adversarial Perturbation (WRAP)

---

**Require:** Weighted sample  $\{(X_n, w_n)\}_{n=1}^N$ , batch  $\mathcal{I} \subseteq \{1, \dots, N\}$  with  $|\mathcal{I}| = M$ , hedge parameter  $\vartheta$ , transport budget  $\varepsilon$ , reweighting budget  $\tau$

**Ensure:** Adversarial weighted batch  $\{(\hat{X}_n, \hat{w}_n)\}_{n \in \mathcal{I}}$

1: Set  $\bar{w}_n \leftarrow w_n / \sum_{j \in \mathcal{I}} w_j$  and compute  $L_n \leftarrow \ell_{\text{DH}}(\vartheta; X_n)$ ,  $g_n \leftarrow \nabla_X \ell_{\text{DH}}(\vartheta; X_n)$  for all  $n \in \mathcal{I}$ .

**$\phi$ -divergence reweighting**

2:   **if**  $\tau = 0$  **then** Set  $a_n \leftarrow 1$  for all  $n \in \mathcal{I}$ .  
3:   **else** Set  $Y_n \leftarrow L_n - \bar{L}_{\mathcal{I}}$  for all  $n \in \mathcal{I}$ , where  $\bar{L}_{\mathcal{I}} = \sum_{n \in \mathcal{I}} \bar{w}_n L_n$ .  
4:       Choose the largest  $\alpha_\tau \geq 0$  such that  $1 + \alpha_\tau Y_n \geq 0$  for all  $n \in \mathcal{I}$  and  $\sum_{n \in \mathcal{I}} \bar{w}_n \phi(1 + \alpha_\tau Y_n) \leq \tau$ .  
5:       Set  $a_n \leftarrow 1 + \alpha_\tau Y_n$  for all  $n \in \mathcal{I}$ .  
6:   Set  $\hat{w}_n \leftarrow \bar{w}_n a_n$  for all  $n \in \mathcal{I}$ .

**Optimal-transport perturbation**

7:   **if**  $\varepsilon = 0$  **then** Set  $\hat{X}_n \leftarrow X_n$  for all  $n \in \mathcal{I}$ .  
8:   **else** Set  $\Upsilon_{\bar{w}}^{\mathcal{I}} \leftarrow (\sum_{n \in \mathcal{I}} \hat{w}_n \|g_n\|_{\kappa, \lambda, p}^{*q})^{1/q}$ .  
9:       Set  $\hat{X}_n \leftarrow X_n + \varepsilon h(g_n) \|g_n\|_{\kappa, \lambda, p}^{*(q-1)} (\Upsilon_{\bar{w}}^{\mathcal{I}})^{1-q}$  for all  $n \in \mathcal{I}$ .

10: **return**  $\{(\hat{X}_n, \hat{w}_n)\}_{n \in \mathcal{I}}$ .

---

**Theorem 3.7** (First-order accuracy of the adversary). *Suppose  $\phi$  is locally regular,  $p \in (1, \infty)$  and Assumption 3.5 holds. For sufficiently small  $(\varepsilon, \tau)$ ,  $\hat{\eta}_{\varepsilon, \tau} \in \mathcal{B}_{\varepsilon, \tau}(\mu_w)$ , and as  $(\varepsilon, \tau) \rightarrow (0, 0)$ ,*

$$V_{\vartheta}(\varepsilon, \tau) - \mathbb{E}_{X \sim \hat{\eta}_{\varepsilon, \tau}}[\ell_{\text{DH}}(\vartheta; X)] = o(\varepsilon + \sqrt{\tau}). \quad (12)$$

Hence the two-budget robust value is attained to first order by an adversary consisting of one probability multiplier and one perturbed trajectory per observation. Its proof is in Appendix C.3.2.

## 4 WRAP: Nonstationary Two-Budget Adversarial Training

The first-order analysis in Section 3 yields a tractable approximation to the two-budget robust objective. Training alternates between an approximate inner maximization over scenario weights and trajectory perturbations and an outer minimization over the hedge parameters. The drift-aware weights from Section 2.2 remain fixed throughout training and define the reference distribution.

**Inner step: Adversarial construction.** For fixed  $\vartheta$ , Algorithm 1 implements the structured first-order attack characterized in Theorem 3.7. Reweighting shifts probability mass toward scenarios with above-average losses, while transport perturbs trajectories along locally adverse directions, yielding a two-budget adversary around the drift-aware reference distribution. The one-step Wasserstein Projection Gradient Descent (WPGD) attack of He et al. (2025, Section 4) is recovered as the special case with zero reweighting budget and uniform reference weights. Appendix D.3 develops the corresponding multistep two-budget variant, which similarly reduces to their iterative transport attack under the same restriction. Appendix D.6.3 compares the one-step and multistep variants.

**Outer step: Loss minimization.** Given the adversarial weighted batch  $\mathcal{I}$ , the outer step minimizes

$$\min_{\vartheta \in \Theta \times \mathbb{R}} \left\{ \mathcal{J}_{\text{adv}}^{\mathcal{I}}(\vartheta) := \sum_{n \in \mathcal{I}} \hat{w}_n \ell_{\text{DH}}(\vartheta; \hat{X}_n) \right\}.$$

The gradient with respect to  $\vartheta$  is evaluated while holding the adversarial weights and perturbed trajectories fixed. Under sufficient regularity, a Danskin argument shows that this update agrees, to leading order in the ambiguity budgets, with the gradient of the corresponding robust value function. After each outer update, the inner adversary is recomputed at the updated hedge parameters, so that the scenario reweighting and transport perturbations adapt to the evolving loss landscape.

## 5 Numerical Experiments

We conduct numerical experiments to evaluate the two central components of the proposed framework. First, we assess whether drift-aware reference weights improve out-of-sample hedging performance under temporal nonstationarity. Second, we isolate the effects of probability reweighting and trajectory perturbations and examine how the corresponding ambiguity budgets affect adversarial paths and hedging performance. We compare uniform reference weights (-U) with drift-aware weights (-W). For each reference distribution, clean training (ERM) applies no adversarial perturbation, “ $\phi$ -” applies probability reweighting, “OT-” applies trajectory perturbations, and “WRAP” combines both mechanisms. Across all experiments, the network architecture and optimization settings are held fixed as reported in Table 2. Ambiguity budgets are selected separately for each experimental setting and training specification using the validation sample, while all reported performance measures are evaluated on the held-out test sample. Appendix D.6 provides additional results, including the stationary Heston benchmark and extensions with proportional transaction costs and a Mean-CVaR (M-CVaR) objective.

### 5.1 Nonstationary Heston Experiments

We evaluate the robustness of a deep-hedging strategy for a European call option trained under the OCE risk measure CVaR in a nonstationary Heston market. Temporal nonstationarity is introduced by varying the stock-price drift and long-run variance linearly across eight consecutive blocks of simulated trajectories. The first five blocks are used for training, the sixth for hyperparameter and model selection, and the final two for evaluating near-future (Near) and far-future (Far) performance under progressively larger distribution shifts. We compare uniform reference weights with the drift-aware weights from Section 2.2, which increasingly emphasize recent observations according to the validation-selected ratio  $(\varrho/\varepsilon_\varrho)^*$ . The reference-weight radius  $\varepsilon_\varrho$  is distinct from the adversarial transport budget  $\varepsilon$  (*cf.* Appendix D.6.1).

Table 1: Results for the nonstationary Heston model.

| Method    | $\varepsilon$ | $\tau$ | Val. CVaR( $\downarrow$ )           | Near CVaR( $\downarrow$ )           | Far CVaR( $\downarrow$ )            | Far Mean PnL( $\uparrow$ ) | Far Turnover( $\downarrow$ ) |
|-----------|---------------|--------|-------------------------------------|-------------------------------------|-------------------------------------|----------------------------|------------------------------|
| ERM-U     | –             | –      | $6.868 \pm 0.028$                   | $7.952 \pm 0.137$                   | $8.992 \pm 0.236$                   | $-4.528 \pm 0.055$         | $10.450 \pm 0.955$           |
| $\phi$ -U | –             | 1.00   | $6.324 \pm 0.034$                   | $7.239 \pm 0.118$                   | $8.128 \pm 0.248$                   | $-4.557 \pm 0.038$         | $10.717 \pm 1.774$           |
| OT-U      | 0.20          | –      | $6.245 \pm 0.093$                   | $7.050 \pm 0.180$                   | $7.811 \pm 0.397$                   | $-4.594 \pm 0.044$         | $10.056 \pm 1.572$           |
| WRAP-U    | 0.30          | 0.60   | $6.005 \pm 0.058$                   | $6.641 \pm 0.136$                   | $7.268 \pm 0.193$                   | $-4.571 \pm 0.029$         | $10.437 \pm 2.500$           |
| ERM-W     | –             | –      | $6.653 \pm 0.044$                   | $7.735 \pm 0.126$                   | $8.707 \pm 0.249$                   | $-4.532 \pm 0.042$         | $7.995 \pm 0.650$            |
| $\phi$ -W | –             | 1.20   | $6.259 \pm 0.036$                   | $7.097 \pm 0.169$                   | $7.919 \pm 0.305$                   | $-4.553 \pm 0.042$         | $8.033 \pm 0.455$            |
| OT-W      | 0.20          | –      | $6.199 \pm 0.093$                   | $6.984 \pm 0.139$                   | $7.772 \pm 0.232$                   | $-4.590 \pm 0.034$         | $7.412 \pm 0.672$            |
| WRAP-W    | 0.40          | 1.00   | <u><math>5.954 \pm 0.029</math></u> | <u><math>6.600 \pm 0.095</math></u> | <u><math>7.221 \pm 0.186</math></u> | $-4.584 \pm 0.034$         | $7.483 \pm 1.125$            |

*Note.* Values are mean  $\pm$  standard deviation over five seeds.

**Results.** As shown in Table 1, under uniform weights, reweighting and transport alone reduce far-future CVaR by 9.6% and 13.1%, respectively, relative to ERM-U, while WRAP-U reduces it by 19.2%. Drift-aware weighting further lowers CVaR for all four methods in every evaluation period. WRAP-W, with  $(\varepsilon, \tau) = (0.40, 1.00)$ , attains the lowest validation, near-future, and far-future CVaR (*cf.* Figure 4 for the validation landscape). The extensions with transaction costs and M-CVaR likewise favor joint training, which attains the lowest risk in both settings (*cf.* Tables 6–7).




Figure 1: Loss landscapes across deep-hedging training schemes from Table 1: (A) ERM-W, (B)  $\phi$ -W, (C) OT-W, and (D) WRAP-W. Each surface depicts the mean hedging loss of a separately trained network on the same reference trajectories  $X_n$ . The axes show terminal moneyness and the root mean square (RMS) of daily log returns in percent. Markers represent reference paths and adversarial trajectories  $\hat{X}_n$ , with sizes proportional to their respective probabilities  $w_n$  and  $\hat{w}_n$ . Arrows indicate selected transport displacements. Within each main panel, subpanels 1–4 report the full feature range, the moneyness marginal, the log-return marginal, and the loss landscape.

Figure 1 shows how the two adversarial mechanisms reshape the training distribution and the resulting loss landscape. Panel (A) shows that under ERM-W, highly volatile paths generate a broad region of high losses. In Panel (B),  $\phi$ -W leaves the trajectories unchanged but reallocates probability toward adverse scenarios, particularly in regions of high realized variation (B.1). This concentration is also visible in the moneyness and log-return marginals (B.2–B.3), where probability mass shifts toward selected groups of adverse paths. Panel (C) isolates the effect of transport. The selected displacements mainly increase realized price variation while inducing comparatively small changes in terminal moneyness (C.1). For a European call, these perturbations stress the sequence of trading gains without requiring a comparable change in the terminal payoff. Accordingly, the

marginals remain relatively close to the reference distribution (C.2–C.3).

Panel (D) combines both mechanisms. At the selected budgets (*cf.* Table 1), WRAP-W reallocates probability toward adverse scenarios while simultaneously transporting trajectories toward locally unfavorable regions (D.1–D.3). The corresponding loss landscape (D.4), together with that of OT-W, exhibits a smaller high-loss region around the strike, although highly volatile paths remain difficult to hedge. Since all loss surfaces are evaluated on the same reference trajectories, these differences reflect the learned hedging policies. The figure thus provides a qualitative counterpart to Theorem 3.6: reweighting acts on loss dispersion across scenarios, whereas transport acts on local sensitivity to perturbations within trajectories.

Figure 2 contrasts reweighting and transport at the level of individual trajectories: the highest-weight path differs from the path most strongly displaced by optimal transport. Reweighting emphasizes high-loss scenarios, whereas transport primarily increases realized variation. Figures 7–8 show the corresponding hedging policies, and Appendix D.7 provides further details.




Figure 2: Pathwise effects in nonstationary Heston: (A) ERM-U, (B)  $\phi$ -U, (C) WRAP-U, and (D) magnified view of the original and transported paths highlighted in (C). Panels (A)–(C) show sampled trajectories and highlight the highest-weight and most-displaced paths.

## 5.2 Market-Data Experiments

We test whether adversarial training improves robustness to changes in market conditions estimated from real equity prices. Following the GAD benchmark of He et al. (2025) and the parameter-uncertainty framework of Lütkebohmert et al. (2022), we hedge a 30-step arithmetic Asian call on AAPL, AMZN, BRK-B, GOOGL, and MSFT under an entropic risk objective. Training, validation, and test trajectories are simulated from GAD models fitted separately to prices from 2008–2019, 2020–2023, and 2024–2025. We consider two training specifications: FIX uses the final parameter estimate from the training period, whereas ROBUST resamples complete rolling estimates at each trading step. Unlike stress-test designs centered on exogenously imposed crash scenarios, we ask whether adversarial training remains beneficial under empirically estimated regime shifts or becomes overly conservative relative to the risk objective. The historical training and validation windows exhibit more severe drawdowns than the test window (*cf.* Appendix D.6.2 and Table 8).

**Results.** WRAP-U attains lower validation and test entropic risk than either single-budget method across all ten stock-training-scheme combinations. Relative to ERM-U, WRAP-U lowers test

entropic risk in six of ten comparisons, while in three additional cases coinciding with ERM-U. Table 9 reports the complete results. This indicates that joint adversarial training can add robustness beyond parameter uncertainty in the training generator, although the gains remain stock dependent.

## 6 Conclusion

We develop a drift-aware distributionally robust framework for deep hedging under potential shifts in market regimes. The framework separates the construction of the reference distribution from the subsequent robustification step. Non-uniform baseline weights balance sampling precision against temporal misspecification from older observations, while separate  $\phi$ -divergence and optimal-transport budgets capture uncertainty in scenario probabilities and market trajectories, respectively.

We establish finite-sample coverage guarantees for the weighted empirical law and derive a joint small-budget expansion of the robust hedging objective. The first-order correction separates scenario-reweighting effects, governed by loss dispersion across trajectories, from transport effects, governed by sensitivity to path perturbations.

**Limitations and Future Research.** The baseline weights considered in this paper depend only on the temporal ordering of the observations. A richer formulation could condition them on observable market states or latent regimes, for example through context-dependent optimization or a hidden Markov model, so that historical trajectories are weighted by economic similarity as well as temporal proximity. An adaptive extension could jointly infer the prevailing regime, update the baseline weights, and estimate the drift bound  $\varrho$  as new observations arrive.

The transport and reweighting budgets are currently selected through validation. A more principled calibration could determine  $\varepsilon$  from the drift-aware coverage guarantee and relate  $\tau$  to finite-sample uncertainty in the weighted scenario probabilities. Finally, a sufficiently large panel of realized market trajectories would permit direct training and evaluation without first fitting a parametric data-generating model.

## AI use statement

In this work, we used generative AI tools to aid and polish the writing. In particular, we used these tools to cross-validate the consistency of mathematical notation throughout the paper and to verify that the paper’s descriptions match the code implementation as closely as possible, ensuring readability. Additionally, upon finalization of the paper, AI was utilized to discover additional related works that were unknown to the authors during development, which were manually verified before inclusion in the related works section. We have reviewed all AI-assisted work and take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

## Acknowledgments

This work was supported as part of NCCR Automation, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 51NF40\_225155).

## References

- Andrew Ang and Allan Timmermann. Regime changes and financial markets. *Annual Review of Financial Economics*, 4(1):313–337, 2012.
- M. Avellaneda, A. Levy, and A. Parás. Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance*, 2(2):73–88, 1995.
- Xingjian Bai, Guangyi He, Yifan Jiang, and Jan Obloj. Wasserstein distributional robustness of neural networks. In *Advances in Neural Information Processing Systems*, volume 36, pp. 26322–26347, 2023.
- Daniel Bartl, Samuel Drapeau, Jan Obłój, and Johannes Wiesel. Sensitivity analysis of Wasserstein distributionally robust optimization problems. *Proceedings of the Royal Society A*, 477(2256): 20210176, 2021.
- Güzin Bayraksan and David K Love. Data-driven stochastic programming using phi-divergences. In *The Operations Research Revolution*, INFORMS TutORials in Operations Research, pp. 1–19. INFORMS, 2015.
- Aharon Ben-Tal and Marc Teboulle. Expected utility, penalty functions, and duality in stochastic nonlinear programming. *Management Science*, 32(11):1445–1466, 1986.
- Aharon Ben-Tal and Marc Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007.
- Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- Amine Bennouna, Bart Van Parys, and Ryan Lucas. Holistic robust data-driven decisions. *arXiv preprint arXiv:2207.09560*, 2022.
- Amine Bennouna, Ryan Lucas, and Bart Van Parys. Certified robust neural networks: Generalization and corruption resistance. In *International Conference on Machine Learning*, pp. 2092–2112. PMLR, 2023.

- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations Research*, 63(5):1227–1244, 2015.
- Jeremiah Birrell and Reza Ebrahimi. Optimal transport regularized divergences: Application to adversarial robustness. *SIAM Journal on Mathematics of Data Science*, 7(4):1801–1827, 2025.
- Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.
- Jose Blanchet, Daniel Kuhn, Jiajin Li, and Bahar Taskesen. Unifying distributionally robust optimization via optimal transport theory. *arXiv preprint arXiv:2308.05414*, 2023.
- Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- Agostino Capponi, Chengpiao Huang, J Antonio Sidaoui, Kaizheng Wang, and Jiacheng Zou. The nonstationarity-complexity tradeoff in return prediction. *arXiv preprint arXiv:2512.23596*, 2025.
- Rama Cont. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16(3):519–547, 2006.
- Rama Cont, Mihai Cucuringu, Renyuan Xu, and Chao Zhang. Tail-GAN: Learning to simulate tail risk scenarios. *Management Science*, 72(4):2917–2936, 2025.
- John Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. Distributionally robust losses for latent covariate mixtures. *Operations Research*, 71(2):649–664, 2023.
- John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46(3):946–969, 2021.
- Darrell Duffie, Damir Filipović, and Walter Schachermayer. Affine processes and applications in finance. *Annals of Applied Probability*, 13(3):984–1053, 2003.
- Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162:707–738, 2015.
- Rui Gao and Anton J. Kleywegt. Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research*, 48(2):603–655, 2023.
- Rui Gao, Xi Chen, and Anton J. Kleywegt. Wasserstein distributionally robust optimization and variation regularization. *Operations Research*, 72(3):1177–1191, 2022.
- Paul Glasserman and Xingbo Xu. Robust risk measurement and model risk. *Quantitative Finance*, 14(1):29–58, 2014.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- Ben Hambly, Renyuan Xu, and Huining Yang. Recent advances in reinforcement learning in finance. *Mathematical Finance*, 33(3):437–503, 2023.
- Guangyi He, Tobias Sutter, and Lukas Gonon. Distributional adversarial attacks and training in deep hedging. In *Advances in Neural Information Processing Systems*, volume 38, pp. 131893–131925, 2025.

- Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.
- Dominic ST Keehan, Edward J Anderson, and Wolfram Wiesemann. Don’t look back in anger: Wasserstein distributionally robust optimization with nonstationary data. *arXiv preprint arXiv:2510.18566*, 2025.
- Daniel Kuhn, Soroosh Shafiee, and Wolfram Wiesemann. Distributionally robust optimization. *Acta Numerica*, 34:579–804, 2025.
- Henry Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.
- Henry Lam. Recovering best statistical guarantees via the empirical divergence-based distributionally robust optimization. *Operations Research*, 67(4):1090–1105, 2019.
- Yannick Limmer and Blanka Horvath. Robust hedging GANs: Towards automated robustification of hedging strategies. *Applied Mathematical Finance*, 31(3):164–201, 2024.
- Zhenyuan Liu, Bart PG Van Parys, and Henry Lam. Smoothed  $f$ -divergence distributionally robust optimization. *arXiv preprint arXiv:2306.14041*, 2023.
- Eva Lütkebohmert, Thorsten Schmidt, and Julian Sester. Robust deep hedging. *Quantitative Finance*, 22(8):1465–1480, 2022.
- Terry J Lyons. Uncertain volatility and the risk-free synthesis of derivatives. *Applied Mathematical Finance*, 2(2):117–133, 1995.
- Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with  $f$ -divergences. In *Advances in Neural Information Processing Systems*, volume 29, pp. 2208–2216, 2016.
- R Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.
- R Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.
- Johannes Ruf and Weiguan Wang. Neural networks for option pricing and hedging: A literature review. *Journal of Computational Finance*, 24(1):1–46, 2020.
- Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.
- Cătălin Stărică and Clive Granger. Nonstationarities in stock returns. *Review of Economics and Statistics*, 87(3):503–522, 2005.
- Bart PG Van Parys, Peyman Mohajerin Esfahani, and Daniel Kuhn. From data to decisions: Distributionally robust optimization is optimal. *Management Science*, 67(6):3387–3402, 2021.

- Jie Wang, Rui Gao, and Yao Xie. Sinkhorn distributionally robust optimization. *Operations Research*, 74(3):1581–1603, 2025.
- Magnus Wiese, Robert Knobloch, Ralf Korn, and Peter Kretschmer. Quant GANs: Deep generation of financial time series. *Quantitative Finance*, 20(9):1419–1440, 2020.
- David Wu and Sebastian Jaimungal. Robust risk-aware option hedging. *Applied Mathematical Finance*, 30(3):153–174, 2023.

# Appendix

|          |                                                                 |           |
|----------|-----------------------------------------------------------------|-----------|
| <b>A</b> | <b>Extended Related Work</b>                                    | <b>18</b> |
| <b>B</b> | <b>Technical Assumptions and Definitions</b>                    | <b>19</b> |
| B.1      | Market Measurability and Payoff Regularity . . . . .            | 19        |
| B.2      | Drift-Aware Coverage Assumptions . . . . .                      | 19        |
| B.3      | Two-Budget Loss Regularity . . . . .                            | 19        |
| <b>C</b> | <b>Proofs</b>                                                   | <b>20</b> |
| C.1      | Weighted Trajectory Metric and Duality . . . . .                | 20        |
| C.1.1    | Wasserstein Distances on Path Space . . . . .                   | 20        |
| C.1.2    | Norm Representation, Dual Norm, and Maximizer Map . . . . .     | 20        |
| C.2      | Drift-Aware Coverage and Calibration . . . . .                  | 23        |
| C.2.1    | Weighted Empirical Concentration . . . . .                      | 23        |
| C.2.2    | Coupling Under Temporal Drift . . . . .                         | 27        |
| C.2.3    | Drift-Aware Coverage Bound . . . . .                            | 29        |
| C.2.4    | Calibration and Next-Period Risk . . . . .                      | 30        |
| C.3      | Two-Budget OT- $\phi$ Adversarial Training . . . . .            | 30        |
| C.3.1    | Exact Reduction and Joint First-order Expansion . . . . .       | 31        |
| C.3.2    | Structured Two-Budget Attack . . . . .                          | 38        |
| <b>D</b> | <b>Numerical Setup and Supplementary Results</b>                | <b>41</b> |
| D.1      | Optimized Certainty Equivalent Risk Measures . . . . .          | 41        |
| D.2      | Contingent Claims . . . . .                                     | 42        |
| D.3      | Multistep WRAP . . . . .                                        | 43        |
| D.4      | Data-Generating Processes . . . . .                             | 43        |
| D.5      | Performance Metrics . . . . .                                   | 44        |
| D.6      | Experiments . . . . .                                           | 45        |
| D.6.1    | Heston Experiments . . . . .                                    | 46        |
| D.6.2    | GAD Experiments . . . . .                                       | 53        |
| D.6.3    | One-step and Multistep WRAP . . . . .                           | 56        |
| D.7      | Pathwise Effects of Reweighting and Optimal Transport . . . . . | 57        |

## A Extended Related Work

*Deep hedging and model uncertainty.* Deep hedging formulates derivative risk management as a model-free, data-driven sequential decision problem (Buehler et al., 2019); related reviews include Ruf & Wang (2020) on neural networks for pricing and Hambly et al. (2023) on reinforcement learning in finance. Classical robust hedging optimizes over families of market dynamics, including uncertain-volatility models (Avellaneda et al., 1995; Lyons, 1995). Recent work studies robustness in deep hedging under parameter uncertainty, alternative data-generating processes, generative adversaries, and distributional perturbations (Lütkebohmert et al., 2022; Wu & Jaimungal, 2023; Limmer & Horvath, 2024; He et al., 2025; Wiese et al., 2020; Cont et al., 2025). We extend this literature by adapting the reference law to temporal drift and distinguishing uncertainty about scenario probabilities from uncertainty about market trajectories.

*Distributionally robust optimization and adversarial training.* Distributionally robust optimization protects decisions against perturbations of a reference law. Wasserstein ambiguity sets transport probability mass according to a ground metric and may therefore generate adverse distributions whose support differs from that of the empirical law (Mohajerin Esfahani & Kuhn, 2018; Gao & Kleywegt, 2023; Kuhn et al., 2025). This construction is particularly well suited to trajectory-valued data because the transport cost can encode temporal and component structure (Blanchet & Murthy, 2019). It also provides a distributional interpretation of adversarial training (Sinha et al., 2018). First-order sensitivity analysis yields tractable approximations of worst-case objectives and identifies locally adverse perturbation directions (Bartl et al., 2021; Gao et al., 2022); related methods have been developed for neural networks and deep hedging (Bai et al., 2023; He et al., 2025).

By contrast,  $\phi$ -divergence ambiguity sets preserve the support of the reference law and alter only the relative probability masses assigned to its outcomes (Ben-Tal et al., 2013; Bayraksan & Love, 2015). They are connected to statistical calibration, variance regularization, and robustness to shifts in the prevalence of high-loss observations or subpopulations (Lam, 2019; Duchi et al., 2021; Van Parys et al., 2021; Namkoong & Duchi, 2016; Lam, 2016; Duchi et al., 2023). Wasserstein and  $\phi$ -divergence ambiguity sets therefore address distinct forms of misspecification: the former changes the outcomes themselves according to a ground geometry, whereas the latter changes their relative likelihoods.

Several approaches combine likelihood and optimal-transport perturbations. Holistic robustness protects against statistical error, diffuse data perturbations, and contamination, with a corresponding specialization to neural networks (Bennouna et al., 2022, 2023). Joint OT-divergence constructions include Blanchet et al. (2023) and the smoothed  $\phi$ -divergences of Liu et al. (2023); our ambiguity set can be viewed as a drift-aware, weighted path-space analogue of the latter, with separate transport and divergence budgets. ARMOR instead combines the two costs in a single infimal convolution of OT and an information divergence and develops a global adversarial-training method (Birrell & Ebrahimi, 2025), whereas Sinkhorn DRO entropy-regularizes the optimal-transport plan (Wang et al., 2025). In contrast to these works, we derive a joint small-budget expansion for a sequential, path-dependent objective and construct an explicit transport-and-reweighting attack that is first-order accurate for the full DRO problem.

*Nonstationarity and weighted empirical laws.* Decision-making under nonstationarity must balance the statistical benefit of additional observations against the misspecification introduced by data generated under earlier distributions (Besbes et al., 2015). Keehan et al. (2025) provide the closest statistical precedent by centering Wasserstein ambiguity sets at weighted empirical laws and separating effective sample size from temporal drift. We integrate this principle with deep hedging and a structured path geometry, and extend it with adversarial probability reweighting under a two-budget training framework.

## B Technical Assumptions and Definitions

This appendix records the assumptions and definitions used in Sections 2 and 3. The concentration and sensitivity results built from these conditions are stated and proved in Appendix C.

### B.1 Market Measurability and Payoff Regularity

We use the market setup and hedging strategy from Section 2. All stochastic processes are defined on a filtered probability space  $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t=0}^T, \mathbb{P})$ . The observed process  $X$  is adapted to this filtration, and hence  $\sigma(X_0, \dots, X_t) \subseteq \mathcal{F}_t$  for  $t = 0, \dots, T$ . Since each  $\delta_t^\theta$  is measurable, the holdings  $\delta_t^\theta(X_{0:t})$  are  $\mathcal{F}_t$ -measurable, and the terminal PnL in (1) is  $\mathcal{F}_T$ -measurable. The negative payoff term  $\Phi(X)$  corresponds to a short position in the claim; examples of the payoffs used in the experiments are given in Appendix D.2.

*Remark B.1* (Payoff and loss regularity). Many vanilla payoffs, such as calls or puts, are globally Lipschitz in  $S_T$  and hence in  $X$ . If differentiability everywhere is desired, kinked terms such as  $\max(\cdot, 0)$  or  $|\cdot|$  can be replaced by smooth approximations (*e.g.* softplus for the positive part). The same approach applies to nonsmooth OCE losses, such as the CVaR loss, while preserving convexity, monotonicity, and the OCE normalization. These approximations preserve the structure of the DRO formulation.

### B.2 Drift-Aware Coverage Assumptions

We restate the assumptions from Section 2 and record their interpretation and the coupling form used in the proofs.

**Assumption 2.1** (Path space). The admissible trajectory set  $\Xi$ , equipped with the metric  $d_{\kappa, \lambda}$  in (5), is a Polish metric space with finite diameter:  $\text{diam}_{\kappa, \lambda}(\Xi) := \sup_{X, \hat{X} \in \Xi} d_{\kappa, \lambda}(X, \hat{X}) < \infty$ .

*Remark B.2.* Assumption 2.1 is satisfied when all  $\lambda_i$  and  $\kappa_t$  are positive and  $\Xi$  is compact. In implementations, compactness corresponds to working with normalized and bounded features or clipping extreme inputs. The Polish-space condition ensures that the Wasserstein ambiguity sets and couplings are well defined, while the finite-diameter condition is used in the concentration bounds and sensitivity analysis.

**Assumption 2.2** (Nonstationary sampling and bounded drift). The trajectories  $X_1, \dots, X_N$  are independent, with  $X_n \sim P_n$  for  $n = 1, \dots, N$ . Moreover, there exists  $\varrho \geq 0$  such that

$$W_\infty(P_n, P_{n+1}) \leq \varrho, \quad n = 1, \dots, N.$$

*Remark B.3.* Following the nonstationary Wasserstein DRO model of Keehan et al. (2025), the drift condition in Assumption 2.2 can equivalently be written in coupling form: for each  $n = 1, \dots, N$ , there exists  $\pi_n \in \Pi(P_n, P_{n+1})$  such that  $d_{\kappa, \lambda}(X, \hat{X}) \leq \varrho$   $\pi_n$ -almost surely. Here  $\Pi(P, Q)$  denotes the set of probability measures on  $\Xi \times \Xi$  with marginals  $P$  and  $Q$ ; see Appendix C.1.1. The proof of the drift-aware concentration result uses this coupling form directly.

### B.3 Two-Budget Loss Regularity

This section records the loss-regularity conditions underlying the two-budget formulation and first-order analysis in Section 3.

Let  $\|\cdot\|_{\kappa, \lambda, p}^*$  and the maximizing map  $h$  be as defined in Proposition 3.4. Recall  $L_n(\vartheta) = \ell_{\text{DH}}(\vartheta; X_n)$  and  $g_n(\vartheta) = \nabla_X \ell_{\text{DH}}(\vartheta; X_n)$ . The following assumption specifies the regularity and local feasibility conditions required for the first-order analysis.

**Assumption 3.5** (Loss regularity). The trajectories  $X_n$  lie in the interior of  $\Xi$ , and for every  $\vartheta \in \Theta \times \mathbb{R}$ , the loss  $\ell_{\text{DH}}(\vartheta; \cdot)$  is Borel measurable, bounded on  $\Xi$  and admits at each  $X_n$  the expansion

$$\ell_{\text{DH}}(\vartheta; X_n + \Delta) = L_n(\vartheta) + \langle g_n(\vartheta), \Delta \rangle + o(\|\Delta\|_{\kappa, \lambda, p}), \quad \|\Delta\|_{\kappa, \lambda, p} \rightarrow 0.$$

*Remark B.4.* Assumption 3.5 is local at the observed trajectories: the differentiability and feasibility requirements are imposed only at the empirical support. The boundedness condition is global and is used to control transport mass that may move away from that support. The feasibility condition ensures that the explicit perturbations in (11) remain in  $\Xi$  for sufficiently small transport radii.

## C Proofs

This appendix contains the proofs of the theoretical results stated in the paper. The arguments follow the logical development of the paper: metric and dual-norm properties of the weighted trajectory distance, drift-aware coverage, and first-order sensitivity expansions for the Wasserstein and two-budget ambiguity sets. For ease of reference, each claim is restated before its proof.

### C.1 Weighted Trajectory Metric and Duality

This section develops the norm-duality structure underlying the Wasserstein perturbation analysis. After recalling Wasserstein distances on the trajectory space, we specialize the ground metric to the weighted trajectory distance used in the paper. We then identify its induced norm, dual norm, and associated maximizer for the linearized perturbation problem that determines the first-order adversarial direction.

#### C.1.1 Wasserstein Distances on Path Space

We recall the Wasserstein distance for a generic ground metric before specializing to the weighted trajectory distance. Let  $(\Xi, d)$  be a Polish metric space and let  $\mathcal{P}(\Xi)$  denote the set of Borel probability measures on  $\Xi$ . For  $p \in [1, \infty)$ , the  $p$ -Wasserstein distance between  $\mu, \nu \in \mathcal{P}(\Xi)$  is

$$W_p(\mu, \nu) := \left( \inf_{\pi \in \Pi(\mu, \nu)} \int_{\Xi \times \Xi} d(x, \hat{x})^p \pi(dx, d\hat{x}) \right)^{1/p}, \quad (13)$$

where  $\Pi(\mu, \nu)$  is the set of couplings of  $\mu$  and  $\nu$ . For  $p = \infty$ , the  $\infty$ -Wasserstein distance is

$$W_\infty(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \text{ess sup}_{(x, \hat{x}) \sim \pi} d(x, \hat{x}). \quad (14)$$

**Standing convention.** Unless stated otherwise, all Wasserstein distances in this appendix are computed with the weighted trajectory ground metric  $d_{\kappa, \lambda}$  defined in (5).

#### C.1.2 Norm Representation, Dual Norm, and Maximizer Map

We restate and prove here the results about the dual geometry.

**Proposition 3.4.** (Dual geometry of the weighted path metric) *Let  $p \in (1, \infty)$ ,  $q = p/(p-1)$ , and  $\kappa_t, \lambda_i > 0$ . For  $\Delta X, G \in \mathbb{R}^{d(T+1)}$ , define*

$$\|\Delta X\|_{\kappa, \lambda, p} := \left( \sum_i (\lambda_i \max_t \kappa_t |\Delta X_t^i|)^p \right)^{1/p}, \quad \|G\|_{\kappa, \lambda, p}^* := \left( \sum_i \left( \sum_t \frac{|G_t^i|}{\lambda_i \kappa_t} \right)^q \right)^{1/q}. \quad (9)$$

The first is a norm inducing the metric in (5) through  $d_{\kappa,\lambda}(X, \widehat{X}) = \|X - \widehat{X}\|_{\kappa,\lambda,p}$ , and the second is its dual norm under the Euclidean inner product. Moreover, there exists a Borel map  $h : \mathbb{R}^{d(T+1)} \rightarrow \mathbb{R}^{d(T+1)}$ , with  $h(0) = 0$ , such that  $\|h(G)\|_{\kappa,\lambda,p} = 1$  and  $\langle G, h(G) \rangle = \|G\|_{\kappa,\lambda,p}^*$  whenever  $G \neq 0$ .

*Proof of Proposition 3.4. Primal norm.*

Homogeneity and definiteness follow directly from the positivity of the weights  $\kappa_t$  and  $\lambda_i$ . For the triangle inequality, define, for each component  $i$ ,

$$\|\Delta X^i\|_{\infty,\kappa} := \max_{t=0,\dots,T} \kappa_t |\Delta X_t^i|, \quad \Delta X^i = (\Delta X_0^i, \dots, \Delta X_T^i).$$

The weighted sup-norm satisfies

$$\|\Delta X^i + \Delta Y^i\|_{\infty,\kappa} \leq \|\Delta X^i\|_{\infty,\kappa} + \|\Delta Y^i\|_{\infty,\kappa}.$$

Therefore,

$$\begin{aligned} \|\Delta X + \Delta Y\|_{\kappa,\lambda,p} &= \left( \sum_{i=1}^d (\lambda_i \|\Delta X^i + \Delta Y^i\|_{\infty,\kappa})^p \right)^{1/p} \\ &\leq \left( \sum_{i=1}^d (\lambda_i \|\Delta X^i\|_{\infty,\kappa} + \lambda_i \|\Delta Y^i\|_{\infty,\kappa})^p \right)^{1/p}. \end{aligned}$$

Applying Minkowski's inequality to the outer  $\ell_p$  aggregation, with  $a_i = \lambda_i \|\Delta X^i\|_{\infty,\kappa}$  and  $b_i = \lambda_i \|\Delta Y^i\|_{\infty,\kappa}$ , gives

$$\begin{aligned} \|\Delta X + \Delta Y\|_{\kappa,\lambda,p} &\leq \left( \sum_{i=1}^d (\lambda_i \|\Delta X^i\|_{\infty,\kappa})^p \right)^{1/p} + \left( \sum_{i=1}^d (\lambda_i \|\Delta Y^i\|_{\infty,\kappa})^p \right)^{1/p} \\ &= \|\Delta X\|_{\kappa,\lambda,p} + \|\Delta Y\|_{\kappa,\lambda,p}. \end{aligned}$$

Thus the triangle inequality holds, and  $\|\cdot\|_{\kappa,\lambda,p}$  is a norm.

*Dual norm.* For  $G, \Delta X \in \mathbb{R}^{d(T+1)}$ , write

$$\langle G, \Delta X \rangle := \sum_{i=1}^d \sum_{t=0}^T G_t^i \Delta X_t^i.$$

Let  $q$  be the Hölder conjugate of  $p \in (1, \infty)$ , so that  $1/p + 1/q = 1$ .

For each component  $i$ , define  $b_i := \sum_{t=0}^T |G_t^i|/\kappa_t$ . For any  $\Delta X \in \mathbb{R}^{d(T+1)}$ , the componentwise pairing satisfies

$$\sum_{t=0}^T G_t^i \Delta X_t^i \leq \sum_{t=0}^T |G_t^i| |\Delta X_t^i| \leq \|\Delta X^i\|_{\infty,\kappa} b_i,$$

where  $\|\Delta X^i\|_{\infty,\kappa} := \max_{t=0,\dots,T} \kappa_t |\Delta X_t^i|$ . Summing over  $i$ , we obtain

$$\langle G, \Delta X \rangle \leq \sum_{i=1}^d (\lambda_i \|\Delta X^i\|_{\infty,\kappa}) \frac{b_i}{\lambda_i}.$$

Hölder's inequality in  $\mathbb{R}^d$  gives

$$\langle G, \Delta X \rangle \leq \|\Delta X\|_{\kappa,\lambda,p} \left( \sum_{i=1}^d \left( \frac{b_i}{\lambda_i} \right)^q \right)^{1/q}.$$

Taking the supremum over all  $\Delta X$  with  $\|\Delta X\|_{\kappa,\lambda,p} \leq 1$  yields

$$\|G\|_{\kappa,\lambda,p}^* \leq \left( \sum_{i=1}^d \left( \sum_{t=0}^T \frac{|G_t^i|}{\lambda_i \kappa_t} \right)^q \right)^{1/q}. \quad (15)$$

It remains to show that the upper bound is tight. Since the dual norm is the supremum over all feasible perturbations, it is at least the supremum over any feasible subclass. Choose  $\Delta X_t^i = \text{sgn}(G_t^i) \alpha_i / \kappa_t$ , with  $\alpha_i \geq 0$ . Then  $\|\Delta X^i\|_{\infty,\kappa} \leq \alpha_i$ ,  $\|\Delta X\|_{\kappa,\lambda,p} \leq (\sum_{i=1}^d (\lambda_i \alpha_i)^p)^{1/p}$ , and  $\langle G, \Delta X \rangle = \sum_{i=1}^d \alpha_i b_i$ . Hence

$$\|G\|_{\kappa,\lambda,p}^* \geq \sup_{(\sum_{i=1}^d (\lambda_i \alpha_i)^p)^{1/p} \leq 1} \sum_{i=1}^d \alpha_i b_i.$$

Setting  $c_i = \lambda_i \alpha_i$ , the right-hand side becomes

$$\sup_{\|c\|_p \leq 1} \sum_{i=1}^d c_i \frac{b_i}{\lambda_i} = \left( \sum_{i=1}^d \left( \frac{b_i}{\lambda_i} \right)^q \right)^{1/q},$$

by the standard duality between  $\ell_p$  and  $\ell_q$ . Substituting  $b_i = \sum_{t=0}^T |G_t^i| / \kappa_t$  gives

$$\|G\|_{\kappa,\lambda,p}^* \geq \left( \sum_{i=1}^d \left( \sum_{t=0}^T \frac{|G_t^i|}{\lambda_i \kappa_t} \right)^q \right)^{1/q}. \quad (16)$$

Combining (15) and (16) proves (9).

*Maximizer map.* The maximizer map  $h$  is applied to dual variables, such as pathwise gradients  $G \in \mathbb{R}^{d(T+1)}$ , rather than to trajectories. For  $G \in \mathbb{R}^{d(T+1)}$ , define  $a_i(G) := \sum_{t=0}^T |G_t^i| / (\lambda_i \kappa_t)$  and  $A(G) := (\sum_{j=1}^d a_j(G)^q)^{1/q}$ . By (9),  $A(G) = \|G\|_{\kappa,\lambda,p}^*$ .

For  $G \neq 0$ , define  $r_i(G) := a_i(G)^{q-1} / A(G)^{q-1}$  and set

$$h(G)_t^i := \frac{r_i(G)}{\lambda_i \kappa_t} \text{sign}(G_t^i), \quad i = 1, \dots, d, \quad t = 0, \dots, T. \quad (17)$$

For  $G = 0$ , we set  $h(0) = 0$ .

By construction,  $\kappa_t |h(G)_t^i| \leq r_i(G) / \lambda_i$ . If  $r_i(G) > 0$ , some  $G_t^i$  is nonzero, so equality is attained at that date; if  $r_i(G) = 0$ , the entire component vanishes. Thus  $\lambda_i \max_t \kappa_t |h(G)_t^i| = r_i(G)$ . Hence

$$\|h(G)\|_{\kappa,\lambda,p} = \left( \sum_{i=1}^d r_i(G)^p \right)^{1/p}.$$

From  $r_i(G) = a_i(G)^{q-1} / A(G)^{q-1}$  and the identity  $(q-1)p = q$ , we obtain

$$\sum_{i=1}^d r_i(G)^p = \frac{\sum_{i=1}^d a_i(G)^q}{A(G)^q} = 1,$$

so  $\|h(G)\|_{\kappa,\lambda,p} = 1$ .

The same construction attains the dual norm. Substituting (17) into the pairing gives

$$\langle G, h(G) \rangle = \sum_{i=1}^d r_i(G) \sum_{t=0}^T \frac{|G_t^i|}{\lambda_i \kappa_t} = \sum_{i=1}^d r_i(G) a_i(G).$$

Substituting the definition of  $r_i(G)$ , this becomes

$$\sum_{i=1}^d r_i(G) a_i(G) = \frac{\sum_{i=1}^d a_i(G)^q}{A(G)^{q-1}} = A(G).$$

Since  $A(G) = \|G\|_{\kappa, \lambda, p}^*$  by (9), we conclude that  $\langle G, h(G) \rangle = \|G\|_{\kappa, \lambda, p}^*$ .  $\square$

## C.2 Drift-Aware Coverage and Calibration

This section establishes the drift-aware concentration result underlying Section 2.2 and its calibration consequences. The analysis separates statistical sampling error, controlled by the effective sample size  $N_{\text{eff}}(w)$ , from temporal drift, controlled by the cumulative drift term  $D_p(w)\varrho$ . We first derive the stationary weighted concentration bound, then transfer it to the nonstationary setting through a coupling argument, and finally obtain the calibrated ambiguity radius and the corresponding next-period risk guarantee.

We model nonstationarity as in Assumption 2.2: the observed trajectories are independent,  $X_n \sim P_n$ , and consecutive laws admit couplings with  $d_{\kappa, \lambda}$ -distance at most  $\varrho$  almost surely. In particular, Assumption 2.2 implies  $W_\infty(P_n, P_{n+1}) \leq \varrho$  for  $n = 1, \dots, N$ .

Recall from Section 2.2 the effective sample size and weighted cumulative drift index,

$$N_{\text{eff}}(w) := \frac{1}{\sum_{n=1}^N w_n^2}, \quad D_p(w) := \left( \sum_{n=1}^N w_n (N - n + 1)^p \right)^{1/p}. \quad (18)$$

Here  $N_{\text{eff}}(w)$  captures the statistical information retained by the weighted sample, while  $D_p(w)$  captures the exposure of the weighting scheme to temporal drift toward the next-period law.

### C.2.1 Weighted Empirical Concentration

We first isolate the sampling error from the drift error by considering the stationary benchmark in which  $X_1, \dots, X_N$  are i. i. d. with law  $P$ . The following two results show that the weighted empirical measure concentrates around  $P$  at a rate governed by  $N_{\text{eff}}(w)$ , rather than by the nominal sample size  $N$ . Recall the definition of the simplex  $\Delta_N := \{w \in \mathbb{R}_+^N : \sum_{n=1}^N w_n = 1\}$ .

**Lemma C.1** (Expectation bound with effective sample size). *Suppose Assumption 2.1 holds. For  $p \in (1, \infty)$  and  $P \in \mathcal{P}(\Xi)$ , let  $X_1, \dots, X_N \sim P$  be i.i.d., and define  $\mu_w := \sum_{n=1}^N w_n \delta_{X_n}$  for  $w \in \Delta_N$ . Then there exist constants  $c_0 > 0$  and  $\nu \in (0, 1/2)$ , depending only on the ambient dimension  $d(T+1)$ ,  $p$ , and  $\text{diam}_{\kappa, \lambda}(\Xi)$ , such that*

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq c_0 N_{\text{eff}}(w)^{-\nu}.$$

*Proof of Lemma C.1.* The proof reduces the weighted empirical measure to a convex combination of ordinary empirical measures and then uses the standard Wasserstein convergence rate for i.i.d. samples.

Since  $X_1, \dots, X_N$  are i.i.d., the law of  $\sum_{n=1}^N w_n \delta_{X_n}$  is invariant under permutations of the weights. Since  $N_{\text{eff}}(w)$  is also permutation-invariant, it is enough to prove the claim for weights sorted in nonincreasing order. We therefore write  $w_1 \geq \dots \geq w_N$  and set  $w_{N+1} = 0$ . For  $k = 1, \dots, N$ , define

$$\bar{P}_k := \frac{1}{k} \sum_{i=1}^k \delta_{X_i}, \quad \alpha_k := k(w_k - w_{k+1}).$$

Then  $\mu_w$  admits the decomposition  $\mu_w = \sum_{k=1}^N \alpha_k \bar{P}_k$ . Indeed,  $\alpha_k \geq 0$  and

$$\sum_{k=1}^N \alpha_k = \sum_{k=1}^N k(w_k - w_{k+1}) = \sum_{k=1}^N w_k = 1,$$

so this is a convex combination. Moreover, the coefficient of  $\delta_{X_i}$  in  $\sum_{k=1}^N \alpha_k \bar{P}_k$  is

$$\sum_{k=i}^N \frac{\alpha_k}{k} = \sum_{k=i}^N (w_k - w_{k+1}) = w_i.$$

Hence the decomposition is exact.

The map  $Q \mapsto W_p(Q, P)^p$  is convex in its first argument. Therefore,

$$W_p(\mu_w, P)^p \leq \sum_{k=1}^N \alpha_k W_p(\bar{P}_k, P)^p.$$

Taking expectations gives

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq \sum_{k=1}^N \alpha_k \mathbb{E}[W_p(\bar{P}_k, P)^p].$$

Since  $\Xi$  has finite diameter under  $d_{\kappa, \lambda}$ , standard empirical Wasserstein convergence bounds on compact subsets of finite-dimensional normed spaces, such as Fournier & Guillin (2015), imply that there exist constants  $C > 0$  and  $\nu \in (0, 1/2)$  such that

$$\mathbb{E}[W_p(\bar{P}_k, P)^p] \leq C k^{-\nu}, \quad k = 1, \dots, N.$$

The ground metric  $d_{\kappa, \lambda}$  is induced by a weighted norm on a finite-dimensional space, so the usual Euclidean-rate statement applies up to a change of constants. Thus,

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq C \sum_{k=1}^N k(w_k - w_{k+1}) k^{-\nu} = C \sum_{k=1}^N k^{1-\nu} (w_k - w_{k+1}).$$

By summation by parts,

$$\sum_{k=1}^N k^{1-\nu} (w_k - w_{k+1}) = \sum_{k=1}^N (k^{1-\nu} - (k-1)^{1-\nu}) w_k,$$

where  $0^{1-\nu} = 0$ . In the remainder of the proof,  $C_\nu < \infty$  denotes a constant depending only on  $\nu$ , whose value may change from line to line. Since  $x \mapsto x^{1-\nu}$  is concave and increasing,

$$k^{1-\nu} - (k-1)^{1-\nu} \leq C_\nu k^{-\nu}.$$

Hence

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq C C_\nu \sum_{k=1}^N k^{-\nu} w_k.$$

It remains to control the last sum by the effective sample size. Let  $m := \lfloor N_{\text{eff}}(w) \rfloor$ . Since  $N_{\text{eff}}(w) \geq 1$ , we have  $m \geq 1$ . Split

$$\sum_{k=1}^N k^{-\nu} w_k = \sum_{k=1}^m k^{-\nu} w_k + \sum_{k=m+1}^N k^{-\nu} w_k := H + R.$$

For the head term, Cauchy–Schwarz and the definition of  $N_{\text{eff}}(w)$  give

$$H \leq \left( \sum_{k=1}^m k^{-2\nu} \right)^{1/2} \left( \sum_{k=1}^m w_k^2 \right)^{1/2} \leq N_{\text{eff}}(w)^{-1/2} \left( \sum_{k=1}^m k^{-2\nu} \right)^{1/2}.$$

Because  $\nu < 1/2$ ,

$$\sum_{k=1}^m k^{-2\nu} \leq C_\nu m^{1-2\nu}.$$

Therefore,

$$H \leq C_\nu N_{\text{eff}}(w)^{-1/2} m^{1/2-\nu} \leq C_\nu N_{\text{eff}}(w)^{-\nu},$$

where the last inequality uses  $m \leq N_{\text{eff}}(w)$ .

For the remainder term, since  $k^{-\nu}$  is decreasing and the weights sum to one,

$$R \leq (m+1)^{-\nu} \sum_{k=m+1}^N w_k \leq (m+1)^{-\nu} \leq N_{\text{eff}}(w)^{-\nu}.$$

Combining the bounds on  $H$  and  $R$ , we obtain

$$\sum_{k=1}^N k^{-\nu} w_k \leq C_\nu N_{\text{eff}}(w)^{-\nu}.$$

Absorbing constants into  $c_0$  gives

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq c_0 N_{\text{eff}}(w)^{-\nu},$$

which proves the claim.  $\square$

**Proposition C.2** (Concentration bound with effective sample size). *Let  $p \in (1, \infty)$  and assume  $\text{diam}_{\kappa, \lambda}(\Xi) < \infty$ . Let  $P \in \mathcal{P}(\Xi)$ , and let  $\nu \in (0, 1/2)$  be the exponent from Lemma C.1. Let  $X_1, \dots, X_N \sim P$  be independent and identically distributed, and define  $\mu_w := \sum_{i=1}^N w_i \delta_{X_i}$ . There exist constants  $c_1, c_2 > 0$  such that, for all  $N \in \mathbb{N}$ ,  $w \in \Delta_N$ , and  $\varepsilon > 0$ ,*

$$\mathbb{P}(W_p(\mu_w, P) \geq \varepsilon) \leq \exp \left( -c_1 N_{\text{eff}}(w) \left( \varepsilon^p - c_2 N_{\text{eff}}(w)^{-\nu} \right)_+^2 \right).$$

In particular, whenever  $\varepsilon \geq 2(c_2 N_{\text{eff}}(w)^{-\nu})^{1/p}$ , we have

$$\mathbb{P}(W_p(\mu_w, P) \geq \varepsilon) \leq \exp \left( -\frac{c_1}{4} N_{\text{eff}}(w) \varepsilon^{2p} \right).$$

*Proof of Proposition C.2.* We apply McDiarmid's inequality to the random variable

$$f(X_1, \dots, X_N) := W_p \left( \sum_{i=1}^N w_i \delta_{X_i}, P \right)^p.$$

The key step is to verify a bounded-differences property. Changing one observation affects only the mass  $w_j$  attached to that observation; boundedness of  $\Xi$  therefore yields a bounded-differences constant proportional to  $w_j$ .

Fix  $j \in \{1, \dots, N\}$ , and let  $X'_j \in \Xi$  be an arbitrary replacement of the  $j$ -th trajectory. Define

$$\mu_w^{(j)} := \sum_{i \neq j} w_i \delta_{X_i} + w_j \delta_{X'_j}.$$

We claim that

$$\left| W_p(\mu_w, P)^p - W_p(\mu_w^{(j)}, P)^p \right| \leq w_j \text{diam}_{\kappa, \lambda}(\Xi)^p. \quad (19)$$

To prove the claim, let  $\pi^{(j)} \in \Pi(\mu_w^{(j)}, P)$  be arbitrary. Since  $\mu_w^{(j)}$  is discrete, we may disintegrate  $\pi^{(j)}$  according to the indexed atoms of  $\mu_w^{(j)}$ . Thus, for some conditional laws  $\beta_1, \dots, \beta_N \in \mathcal{P}(\Xi)$ ,

$$\pi^{(j)}(dx, dy) = \sum_{i \neq j} w_i \delta_{X_i}(dx) \beta_i(dy) + w_j \delta_{X'_j}(dx) \beta_j(dy).$$

If some atoms coincide, we split the mass according to this indexed representation; this does not change the marginal measure or the transport cost.

Now construct a coupling  $\pi \in \Pi(\mu_w, P)$  by keeping the same conditional laws and replacing only the source point  $X'_j$  by  $X_j$ :

$$\pi(dx, dy) = \sum_{i \neq j} w_i \delta_{X_i}(dx) \beta_i(dy) + w_j \delta_{X_j}(dx) \beta_j(dy).$$

This coupling is feasible for  $W_p(\mu_w, P)$ . Therefore,

$$\begin{aligned} W_p(\mu_w, P)^p &\leq \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, y)^p d\pi(x, y) \\ &= \sum_{i \neq j} w_i \int_{\Xi} d_{\kappa, \lambda}(X_i, y)^p \beta_i(dy) + w_j \int_{\Xi} d_{\kappa, \lambda}(X_j, y)^p \beta_j(dy). \end{aligned}$$

Since  $X_j, X'_j$ , and  $y$  all belong to  $\Xi$ , we have  $d_{\kappa, \lambda}(X_j, y)^p \leq \text{diam}_{\kappa, \lambda}(\Xi)^p$ , and hence

$$d_{\kappa, \lambda}(X_j, y)^p \leq d_{\kappa, \lambda}(X'_j, y)^p + \text{diam}_{\kappa, \lambda}(\Xi)^p.$$

Substituting this bound gives

$$\begin{aligned} W_p(\mu_w, P)^p &\leq \sum_{i \neq j} w_i \int_{\Xi} d_{\kappa, \lambda}(X_i, y)^p \beta_i(dy) + w_j \int_{\Xi} d_{\kappa, \lambda}(X'_j, y)^p \beta_j(dy) \\ &\quad + w_j \text{diam}_{\kappa, \lambda}(\Xi)^p \\ &= \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, y)^p d\pi^{(j)}(x, y) + w_j \text{diam}_{\kappa, \lambda}(\Xi)^p. \end{aligned}$$

Since  $\pi^{(j)} \in \Pi(\mu_w^{(j)}, P)$  was arbitrary, taking the infimum over all such couplings yields

$$W_p(\mu_w, P)^p - W_p(\mu_w^{(j)}, P)^p \leq w_j \text{diam}_{\kappa, \lambda}(\Xi)^p.$$

Exchanging the roles of  $X_j$  and  $X'_j$  gives the reverse inequality, and therefore proves (19). Thus  $f$  satisfies McDiarmid's bounded-differences condition with constants

$$a_j := w_j \text{diam}_{\kappa, \lambda}(\Xi)^p, \quad j = 1, \dots, N.$$

Hence, for every  $u > 0$ ,

$$\mathbb{P}(f - \mathbb{E}[f] \geq u) \leq \exp\left(-\frac{2u^2}{\sum_{j=1}^N a_j^2}\right).$$

Since

$$\sum_{j=1}^N a_j^2 = \text{diam}_{\kappa, \lambda}(\Xi)^{2p} \sum_{j=1}^N w_j^2 = \text{diam}_{\kappa, \lambda}(\Xi)^{2p} N_{\text{eff}}(w)^{-1},$$

we obtain

$$\mathbb{P}(f - \mathbb{E}[f] \geq u) \leq \exp(-c_1 N_{\text{eff}}(w) u^2), \quad c_1 := 2 \text{diam}_{\kappa, \lambda}(\Xi)^{-2p}.$$

By Lemma C.1, there exists  $c_2 > 0$  such that

$$\mathbb{E}[W_p(\mu_w, P)^p] \leq c_2 N_{\text{eff}}(w)^{-\nu}.$$

Therefore, for any  $\varepsilon > 0$ ,

$$\begin{aligned} \mathbb{P}(W_p(\mu_w, P) \geq \varepsilon) &= \mathbb{P}(W_p(\mu_w, P)^p \geq \varepsilon^p) \\ &\leq \mathbb{P}(W_p(\mu_w, P)^p - \mathbb{E}[W_p(\mu_w, P)^p] \geq \varepsilon^p - c_2 N_{\text{eff}}(w)^{-\nu}) \\ &\leq \exp\left(-c_1 N_{\text{eff}}(w) \left(\varepsilon^p - c_2 N_{\text{eff}}(w)^{-\nu}\right)_+^2\right). \end{aligned}$$

This proves the first claim.

For the simplified bound, suppose  $\varepsilon \geq 2(c_2 N_{\text{eff}}(w)^{-\nu})^{1/p}$ . Then  $c_2 N_{\text{eff}}(w)^{-\nu} \leq 2^{-p} \varepsilon^p$ . Since  $p \in (1, \infty)$ , we have  $1 - 2^{-p} \geq 1/2$ , and therefore

$$\left(\varepsilon^p - c_2 N_{\text{eff}}(w)^{-\nu}\right)_+ \geq \frac{1}{2} \varepsilon^p.$$

Substituting this into the first concentration bound gives

$$\mathbb{P}(W_p(\mu_w, P) \geq \varepsilon) \leq \exp\left(-\frac{c_1}{4} N_{\text{eff}}(w) \varepsilon^{2p}\right),$$

which completes the proof.  $\square$

### C.2.2 Coupling Under Temporal Drift

We first bound the Wasserstein distance between weighted empirical laws by pairing observations with the same index.

**Proposition C.3** (Weighted coupling bound). *Let  $p \in (1, \infty)$ ,  $\mu_w = \sum_{n=1}^N w_n \delta_{X_n}$ , and  $\hat{\mu}_w = \sum_{n=1}^N w_n \delta_{\hat{X}_n}$ . Then,*

$$W_p(\mu_w, \hat{\mu}_w)^p \leq \sum_{n=1}^N w_n d_{\kappa, \lambda}(X_n, \hat{X}_n)^p. \quad (20)$$

*Proof of Proposition C.3.* We construct an explicit coupling by assigning mass  $w_n$  to the pair  $(X_n, \hat{X}_n)$  for each  $n = 1, \dots, N$ , namely  $\pi_w := \sum_{n=1}^N w_n \delta_{(X_n, \hat{X}_n)}$ . The first marginal of  $\pi_w$  is  $\mu_w$ , and its second marginal is  $\hat{\mu}_w$ . Hence  $\pi_w \in \Pi(\mu_w, \hat{\mu}_w)$ . The  $p$ -Wasserstein distance on  $(\Xi, d_{\kappa, \lambda})$  is defined by

$$W_p(\mu_w, \hat{\mu}_w)^p = \inf_{\pi \in \Pi(\mu_w, \hat{\mu}_w)} \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi(x, \hat{x}).$$

Since  $\pi_w$  is an admissible coupling, evaluating the transport cost at  $\pi_w$  gives the upper bound

$$W_p(\mu_w, \hat{\mu}_w)^p \leq \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi_w(x, \hat{x}).$$

Finally, substituting the definition of  $\pi_w$

$$\int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi_w(x, \hat{x}) = \sum_{n=1}^N w_n d_{\kappa, \lambda}(X_n, \hat{X}_n)^p.$$

We therefore obtain the desired bound,

$$W_p(\mu_w, \hat{\mu}_w)^p \leq \sum_{n=1}^N w_n d_{\kappa, \lambda}(X_n, \hat{X}_n)^p.$$

□

We next transfer the stationary concentration estimate to the nonstationary sample. The following coupling argument compares the weighted empirical measure generated by  $X_n \sim P_n$  with a weighted empirical measure drawn from the next-period law  $P_{N+1}$ . The cost of this comparison is the drift term  $D_p(w)\varrho$ . Older observations incur a larger transport cost because they must be coupled through more one-step distributional shifts before reaching  $P_{N+1}$ .

**Lemma C.4** (Drift-aware Wasserstein bound). *Suppose Assumption 2.2 holds, and let  $p \in (1, \infty)$  and  $w \in \Delta_N$ . Then there exist coupled pairs  $(X_n, \tilde{X}_n)$ , independent across  $n$ , with  $X_n \sim P_n$  and  $\tilde{X}_n \sim P_{N+1}$ , such that, for  $\mu_w := \sum_{n=1}^N w_n \delta_{X_n}$  and  $\tilde{\mu}_w := \sum_{n=1}^N w_n \delta_{\tilde{X}_n}$ ,*

$$W_p(\mu_w, P_{N+1}) \leq W_p(\tilde{\mu}_w, P_{N+1}) + D_p(w)\varrho \quad \text{a.s.}$$

*In particular,  $\tilde{X}_1, \dots, \tilde{X}_N$  are i.i.d. with law  $P_{N+1}$ .*

*Proof of Lemma C.4.* By Assumption 2.2, for each  $k = 1, \dots, N$  there is a coupling of  $P_k$  and  $P_{k+1}$  whose  $d_{\kappa, \lambda}$ -cost is at most  $\varrho$  almost surely. Gluing these couplings from time  $n$  to time  $N+1$  yields a coupling of  $P_n$  and  $P_{N+1}$  under which

$$d_{\kappa, \lambda}(X_n, \tilde{X}_n) \leq \sum_{k=n}^N \varrho = (N - n + 1)\varrho \quad \text{a.s.}$$

Choosing these couplings independently across  $n$  gives the coupled pairs in the statement. Since the couplings are independent across  $n$ , and each proxy variable has marginal law  $P_{N+1}$ , the variables  $\tilde{X}_1, \dots, \tilde{X}_N$  are i.i.d. with law  $P_{N+1}$ .

We compare  $\mu_w$  with the proxy measure  $\tilde{\mu}_w$  using the triangle inequality. By the triangle inequality for  $W_p$ ,

$$W_p(\mu_w, P_{N+1}) \leq W_p(\tilde{\mu}_w, P_{N+1}) + W_p(\mu_w, \tilde{\mu}_w).$$

The second term is controlled by the weighted coupling bound in Proposition C.3. Applying it to the paired samples  $(X_n, \tilde{X}_n)$  yields

$$W_p(\mu_w, \tilde{\mu}_w)^p \leq \sum_{n=1}^N w_n d_{\kappa, \lambda}(X_n, \tilde{X}_n)^p.$$

Using the pairwise drift bound and the definition of  $D_p(w)$ , we obtain

$$W_p(\mu_w, \tilde{\mu}_w)^p \leq \varrho^p \sum_{n=1}^N w_n (N - n + 1)^p = \varrho^p D_p(w)^p.$$

Hence  $W_p(\mu_w, \tilde{\mu}_w) \leq D_p(w)\varrho$ . Substituting this into the previous triangle inequality gives

$$W_p(\mu_w, P_{N+1}) \leq W_p(\tilde{\mu}_w, P_{N+1}) + D_p(w)\varrho \quad \text{a.s.}$$

□

### C.2.3 Drift-Aware Coverage Bound

We next combine the weighted empirical concentration bound with the drift coupling. The resulting theorem gives the concentration inequality underlying the calibrated radius in Section 2.2, expressed in terms of the effective sample size  $N_{\text{eff}}(w)$  and cumulative drift index  $D_p(w)$ . It specializes the nonstationary Wasserstein concentration guarantee of Keehan et al. (2025) to the trajectory space  $(\Xi, d_{\kappa, \lambda})$ . Let  $(u)_+ := \max\{0, u\}$ .

**Theorem C.5** (Drift-aware coverage). *Let  $p \in (1, \infty)$  and suppose Assumptions 2.1 and 2.2 hold. Then there exist constants  $c_1, c_2 > 0$  and an exponent  $\nu \in (0, 1/2)$  such that, for every  $w \in \Delta_N$  and every  $\varepsilon > 0$ ,*

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon) \leq \exp\left(-c_1 N_{\text{eff}}(w) \left(\left((\varepsilon - D_p(w)\varrho)_+^p - c_2 N_{\text{eff}}(w)^{-\nu}\right)_+\right)^2\right). \quad (21)$$

In particular, if  $\varepsilon \geq D_p(w)\varrho + 2(c_2 N_{\text{eff}}(w)^{-\nu})^{1/p}$ , then

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon) \leq \exp\left(-\frac{c_1}{4} N_{\text{eff}}(w) (\varepsilon - D_p(w)\varrho)_+^{2p}\right). \quad (22)$$

*Proof of Theorem C.5.* We compare the nonstationary weighted empirical measure  $\mu_w$  to a proxy weighted empirical measure drawn from the target distribution  $P_{N+1}$ . By Lemma C.4, we may construct coupled pairs  $(X_n, \tilde{X}_n)$ , independent across  $n$ , with  $X_n \sim P_n$  and  $\tilde{X}_n \sim P_{N+1}$ , such that, for  $\tilde{\mu}_w := \sum_{n=1}^N w_n \delta_{\tilde{X}_n}$ ,

$$W_p(\mu_w, P_{N+1}) \leq W_p(\tilde{\mu}_w, P_{N+1}) + D_p(w)\varrho \quad \text{a.s.}$$

Consequently, for any  $\varepsilon > 0$ ,

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon) \leq \mathbb{P}(W_p(\tilde{\mu}_w, P_{N+1}) \geq (\varepsilon - D_p(w)\varrho)_+).$$

The proxy variables  $\tilde{X}_1, \dots, \tilde{X}_N$  are i.i.d. with law  $P_{N+1}$ . Hence Proposition C.2, applied with  $P = P_{N+1}$  and with  $\mu_w$  replaced by  $\tilde{\mu}_w$ , gives, for all  $t > 0$ ,

$$\mathbb{P}(W_p(\tilde{\mu}_w, P_{N+1}) \geq t) \leq \exp\left(-c_1 N_{\text{eff}}(w) (t^p - c_2 N_{\text{eff}}(w)^{-\nu})_+^2\right).$$

If  $t = (\varepsilon - D_p(w)\varrho)_+ = 0$ , the desired bound is trivial because the right-hand side of (21) equals one. Otherwise, substituting  $t = (\varepsilon - D_p(w)\varrho)_+$  into the preceding display yields

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon) \leq \exp\left(-c_1 N_{\text{eff}}(w) \left(\left((\varepsilon - D_p(w)\varrho)_+^p - c_2 N_{\text{eff}}(w)^{-\nu}\right)_+\right)^2\right),$$

which proves (21).

The simplified tail bound follows by lower-bounding the residual radius. Suppose that

$$\varepsilon \geq D_p(w)\varrho + 2(c_2 N_{\text{eff}}(w)^{-\nu})^{1/p}.$$

Then

$$c_2 N_{\text{eff}}(w)^{-\nu} \leq 2^{-p}(\varepsilon - D_p(w)\varrho)_+^p.$$

Since  $p > 1$ , we have  $1 - 2^{-p} \geq 1/2$ , and therefore

$$\left(\left(\varepsilon - D_p(w)\varrho\right)_+^p - c_2 N_{\text{eff}}(w)^{-\nu}\right)_+ \geq \frac{1}{2}(\varepsilon - D_p(w)\varrho)_+^p.$$

Substituting this lower bound into (21) gives

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon) \leq \exp\left(-\frac{c_1}{4} N_{\text{eff}}(w) (\varepsilon - D_p(w)\varrho)_+^{2p}\right),$$

which proves (22).  $\square$

*Remark C.6* (Drift-aware truncation weights). Theorem 2 of Keehan et al. (2025) shows that weights optimizing the bias–variance trade-off induced by  $N_{\text{eff}}(w)$  and  $D_p(w)\varrho$  in the concentration bound (22) admit a truncated-polynomial form  $w_n \propto (a_1 - a_2(N - n + 1)^p)_+$  for two scalars  $a_1$  and  $a_2$ .

#### C.2.4 Calibration and Next-Period Risk

**Theorem 2.3** (Drift-aware coverage and next-period risk). *Let  $p \in (1, \infty)$  and suppose Assumptions 2.1 and 2.2 hold. There exist constants  $c_1, c_2 > 0$  and  $\nu \in (0, 1/2)$ , depending only on  $d(T + 1)$ ,  $p$ ,  $\text{diam}_{\kappa, \lambda}(\Xi)$ , such that the following holds for every  $w \in \Delta_N$  and  $\beta \in (0, 1)$ . Define*

$$\varepsilon_\varrho(w, \beta) := D_p(w)\varrho + \left(c_2 N_{\text{eff}}(w)^{-\nu} + \sqrt{\frac{\log(1/\beta)}{c_1 N_{\text{eff}}(w)}}\right)^{1/p}.$$

*With probability at least  $1 - \beta$ , we have  $P_{N+1} \in \mathcal{B}_{\varepsilon_\varrho(w, \beta)}(\mu_w)$  and  $\mathbb{E}_{X \sim P_{N+1}}[\ell_{\text{DH}}(\vartheta; X)] \leq V_\vartheta(\varepsilon_\varrho(w, \beta))$  for every  $\vartheta \in \Theta \times \mathbb{R}$ , where  $V_\vartheta$  uses reference law  $\mu_w$ .*

*Proof of Theorem 2.3.* Choose  $c_1, c_2$  and  $\nu$  as in Theorem C.5. Substituting  $\varepsilon = \varepsilon_\varrho(w, \beta)$  into (21) gives

$$\mathbb{P}(W_p(\mu_w, P_{N+1}) \geq \varepsilon_\varrho(w, \beta)) \leq \exp(-\log(1/\beta)) = \beta.$$

Hence  $E := \{P_{N+1} \in \mathcal{B}_{\varepsilon_\varrho(w, \beta)}(\mu_w)\}$  has probability at least  $1 - \beta$ . On  $E$ , the law  $P_{N+1}$  is feasible in the supremum defining  $V_\vartheta(\varepsilon_\varrho(w, \beta))$ , so  $\mathbb{E}_{X \sim P_{N+1}}[\ell_{\text{DH}}(\vartheta; X)] \leq V_\vartheta(\varepsilon_\varrho(w, \beta))$  for every  $\vartheta \in \Theta \times \mathbb{R}$ . Since  $E$  does not depend on  $\vartheta$ , this bound holds simultaneously for all  $\vartheta$ .  $\square$

### C.3 Two-Budget OT- $\phi$ Adversarial Training

This subsection proves the two-budget sensitivity results. We first derive the joint first-order expansion of the OT- $\phi$  robust value by reducing the ambiguity set to an outer reweighting problem and an inner transport problem. We then prove that the structured empirical attack is feasible and attains the same expansion to first order.

### C.3.1 Exact Reduction and Joint First-order Expansion

The proof of Theorem 3.6 proceeds in three steps. First, Proposition C.7 rewrites the robust problem as an outer optimization over anchor reweightings  $a$  and an inner transport problem  $T_a(\varepsilon)$ . Second, Lemma C.8 provides a first-order expansion of  $T_a(\varepsilon)$ , uniformly for reweightings close to  $\mathbf{1}$ . Third, Lemma C.9 gives the small- $\tau$  expansion of the resulting pure reweighting problem. Since  $Q_\phi(\tau)$  collapses to  $\{\mathbf{1}\}$  as  $\tau \downarrow 0$ , the reweighted transport sensitivity converges uniformly to  $\Upsilon_w(\vartheta)$ , and the interaction between the two budgets is of smaller order. Combining these steps yields the claimed joint expansion.

**Proposition C.7** (Exact empirical reduction for the two-budget OT- $\phi$  set). *For  $p \in (1, \infty)$ , define*

$$Q_\phi(\tau) := \left\{ a \in [0, \infty)^N : \sum_{n=1}^N w_n a_n = 1, \sum_{n=1}^N w_n \phi(a_n) \leq \tau \right\},$$

and for  $a \in Q_\phi(\tau)$  define

$$T_a(\varepsilon) := \sup_{\nu'_1, \dots, \nu'_N \in \mathcal{P}(\Xi)} \left\{ \sum_{n=1}^N w_n a_n \int_{\Xi} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}) : \sum_{n=1}^N w_n a_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}) \leq \varepsilon^p \right\}.$$

The robust value admits the exact reduction

$$V_\vartheta(\varepsilon, \tau) = \sup_{a \in Q_\phi(\tau)} T_a(\varepsilon). \quad (23)$$

*Proof of Proposition C.7.* We prove the claim in two steps. First, we show that any feasible  $\eta \in \mathcal{B}_{\varepsilon, \tau}$  admits a representation  $\eta = \sum_{n=1}^N w_n a_n \nu'_n$ , where  $a$  satisfies the  $\phi$ -divergence constraint and the conditional laws  $\nu'_n$  satisfy the transport constraint. Second, we show that any such representation induces a feasible law in  $\mathcal{B}_{\varepsilon, \tau}$ .

Fix any feasible  $\eta \in \mathcal{B}_{\varepsilon, \tau}$ . By definition, there exist couplings  $\pi, \pi' \in \mathcal{P}(\Xi \times \Xi)$  such that

$$\pi_1 = \mu_w, \quad \pi'_2 = \eta, \quad \int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi'(x, \hat{x}) \leq \varepsilon^p, \quad D_\phi(\pi' \| \pi) \leq \tau.$$

Since  $\pi_1 = \mu_w = \sum_{n=1}^N w_n \delta_{X_n}$ , the first coordinate has support contained in the finite set  $\{X_1, \dots, X_N\}$ . Hence  $\pi$  admits the disintegration

$$\pi(dx, d\hat{x}) = \sum_{n=1}^N w_n \delta_{X_n}(dx) \nu_n(d\hat{x}),$$

where  $\nu_n \in \mathcal{P}(\Xi)$  are the reference conditional laws.

Because  $D_\phi(\pi' \| \pi) < \infty$ , we necessarily have  $\pi' \ll \pi$ . Let  $W = d\pi'/d\pi$  denote the Radon-Nikodym derivative, and define its slices as

$$W_n(\hat{x}) := W(X_n, \hat{x}), \quad a_n := \int_{\Xi} W_n(\hat{x}) \nu_n(d\hat{x}).$$

The quantity  $a_n$  represents the total mass assigned by  $\pi'$  to the slice  $\{x = X_n\}$  relative to  $\pi$ . For each  $n$  such that  $a_n > 0$ , define the normalized conditional probability measure  $\nu'_n$  by

$$\nu'_n(d\hat{x}) := \frac{W_n(\hat{x})}{a_n} \nu_n(d\hat{x}),$$

and choose an arbitrary probability measure  $\nu'_n \in \mathcal{P}(\Xi)$  when  $a_n = 0$ . This construction cleanly separates the distorted coupling into a mass component  $a_n$  and a normalized conditional slice  $\nu'_n$ . The distorted coupling therefore factors as:

$$\pi'(dx, d\hat{x}) = \sum_{n=1}^N w_n a_n \delta_{X_n}(dx) \nu'_n(d\hat{x}).$$

Taking the second marginal (integrating out  $dx$ ) directly yields:

$$\pi'_2 = \eta = \sum_{n=1}^N w_n a_n \nu'_n, \quad \text{with} \quad \sum_{n=1}^N w_n a_n = 1,$$

thus any feasible  $\eta$  can be written as a finite mixture with weights  $w_n a_n$ . Consequently, its expected loss equals the objective appearing in  $T_a(\varepsilon)$ :

$$\mathbb{E}_{\hat{X} \sim \eta}[\ell_{\text{DH}}(\vartheta; \hat{X})] = \sum_{n=1}^N w_n a_n \int_{\Xi} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}).$$

Moreover, the transport constraint is preserved under this representation:

$$\int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi'(x, \hat{x}) = \sum_{n=1}^N w_n a_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}) \leq \varepsilon^p.$$

Finally, by Jensen's inequality applied to the convex function  $\phi$ ,

$$\begin{aligned} D_{\phi}(\pi' \parallel \pi) &= \sum_{n=1}^N w_n \int_{\Xi} \phi(W_n(\hat{x})) \nu_n(d\hat{x}) \\ &\geq \sum_{n=1}^N w_n \phi\left(\int_{\Xi} W_n(\hat{x}) \nu_n(d\hat{x})\right) = \sum_{n=1}^N w_n \phi(a_n), \end{aligned}$$

so  $a \in Q_{\phi}(\tau)$  since  $D_{\phi}(\pi' \parallel \pi) \leq \tau$ . Thus, for this specific  $a$ , the conditional slices  $(\nu'_1, \dots, \nu'_N)$  are feasible for the inner supremum  $T_a(\varepsilon)$ , attaining the exact same objective value. This proves

$$V_{\vartheta}(\varepsilon, \tau) \leq \sup_{a \in Q_{\phi}(\tau)} T_a(\varepsilon).$$

Conversely, fix any  $a \in Q_{\phi}(\tau)$  and any  $\nu'_1, \dots, \nu'_N \in \mathcal{P}(\Xi)$  feasible in the definition of  $T_a(\varepsilon)$ . Set

$$\pi(dx, d\hat{x}) := \sum_{n=1}^N w_n \delta_{X_n}(dx) \nu'_n(d\hat{x}), \quad \pi'(dx, d\hat{x}) := \sum_{n=1}^N w_n a_n \delta_{X_n}(dx) \nu'_n(d\hat{x}).$$

Then  $\pi_1 = \mu_w$ ,  $\pi'_2 = \sum_{n=1}^N w_n a_n \nu'_n$ , and the transport cost evaluates to:

$$\int_{\Xi \times \Xi} d_{\kappa, \lambda}(x, \hat{x})^p d\pi'(x, \hat{x}) = \sum_{n=1}^N w_n a_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}) \leq \varepsilon^p.$$

Moreover,  $\pi' \ll \pi$ . Where several components overlap,  $d\pi'/d\pi$  is a weighted average of the corresponding  $a_n$ 's, so applying Jensen's inequality with the convexity of  $\phi$  yields

$$D_{\phi}(\pi' \parallel \pi) \leq \sum_{n=1}^N w_n \phi(a_n) \leq \tau.$$

Hence  $\eta := \sum_{n=1}^N w_n a_n \nu'_n$  belongs to  $\mathcal{B}_{\varepsilon, \tau}$  and attains exactly the same objective value as in  $T_a(\varepsilon)$ . Taking the supremum over all  $a \in Q_\phi(\tau)$  and their corresponding feasible conditional slices  $(\nu'_1, \dots, \nu'_N)$  proves the reverse inequality

$$V_\vartheta(\varepsilon, \tau) \geq \sup_{a \in Q_\phi(\tau)} T_a(\varepsilon).$$

□

Next, we show the *uniform* first-order expansion of the Wasserstein ambiguity.

**Lemma C.8** (Uniform OT expansion under small reweightings). *Let  $p \in (1, \infty)$  and suppose Assumption 3.5 holds, with  $\|\cdot\|_{\kappa, \lambda, p}^*$  and  $h$  as in Proposition 3.4. For  $\delta \in (0, 1)$ , define*

$$A_\delta := \left\{ a \in [0, \infty)^N : \sum_{n=1}^N w_n a_n = 1, \|a - \mathbf{1}\|_\infty \leq \delta \right\}. \quad (24)$$

For each  $a \in A_\delta$ , set  $\hat{w}_n := w_n a_n$  for  $n = 1, \dots, N$  and define the aggregate first-order sensitivity

$$\Upsilon_{\hat{w}}(\vartheta) := \left( \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} \right)^{1/q}.$$

Here  $\hat{w} = \hat{w}(a)$ ; its dependence on  $a$  is implicit in the notation throughout this lemma and its proof. Then, for every fixed  $\delta \in (0, 1)$ , the following expansion holds uniformly over  $a \in A_\delta$ :

$$\sup_{a \in A_\delta} \left| T_a(\varepsilon) - \sum_{n=1}^N \hat{w}_n L_n(\vartheta) - \varepsilon \Upsilon_{\hat{w}}(\vartheta) \right| = o(\varepsilon) \quad \text{as } \varepsilon \downarrow 0. \quad (25)$$

*Proof of Lemma C.8.* We establish matching upper and lower bounds uniformly over  $a \in A_\delta$ . The upper bound follows from a local linearization of  $\ell_{\text{DH}}$  around the nominal anchors; the lower bound uses an explicit deterministic perturbation along the dual maximizing directions.

Let

$$B_\ell := \sup_{\hat{x} \in \Xi} |\ell_{\text{DH}}(\vartheta; \hat{x})| < \infty$$

and fix  $\zeta > 0$ . By the local expansion in Assumption 3.5 and the finiteness of  $N$ , there exists  $r_\zeta > 0$  such that, for every  $n = 1, \dots, N$ ,

$$|\ell_{\text{DH}}(\vartheta; \hat{x}) - L_n(\vartheta) - \langle g_n(\vartheta), \hat{x} - X_n \rangle| \leq \zeta d_{\kappa, \lambda}(X_n, \hat{x}) \quad (26)$$

for all  $\hat{x} \in \Xi$  with  $d_{\kappa, \lambda}(X_n, \hat{x}) \leq r_\zeta$ .

To prove the upper bound, fix  $a \in A_\delta$  and let  $(\nu'_1, \dots, \nu'_N)$  be any feasible family of conditional slices for the inner problem  $T_a(\varepsilon)$ . Define the local neighborhood  $B_n := \{\hat{x} \in \Xi : d_{\kappa, \lambda}(X_n, \hat{x}) \leq r_\zeta\}$ . Since  $\Xi = B_n \cup B_n^c$ , we decompose the expected loss as

$$\int_{\Xi} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}) = \int_{B_n} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}) + \int_{B_n^c} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}).$$

On  $B_n$ , the linearization gives  $\ell_{\text{DH}}(\vartheta; \hat{x}) \leq L_n(\vartheta) + \langle g_n(\vartheta), \hat{x} - X_n \rangle + \zeta d_{\kappa, \lambda}(X_n, \hat{x})$ , while on  $B_n^c$ , global boundedness implies  $\ell_{\text{DH}}(\vartheta; \hat{x}) - L_n(\vartheta) \leq 2B_\ell$ .

Integrating these bounds and utilizing the identity  $L_n(\vartheta)\nu'_n(B_n) + L_n(\vartheta)\nu'_n(B_n^c) = L_n(\vartheta)$ , we obtain

$$\int_{\Xi} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}) \leq L_n(\vartheta) + \int_{B_n} \langle g_n(\vartheta), \hat{x} - X_n \rangle \nu'_n(d\hat{x}) + \zeta \int_{B_n} d_{\kappa, \lambda}(X_n, \hat{x}) \nu'_n(d\hat{x}) + 2B_\ell \nu'_n(B_n^c).$$

By the definition of the dual norm,  $\langle g_n(\vartheta), \hat{x} - X_n \rangle \leq \|g_n(\vartheta)\|_{\kappa, \lambda, p}^* d_{\kappa, \lambda}(X_n, \hat{x})$ . Since this upper bound is non-negative, we can extend the integration domain from  $B_n$  to the entire space  $\Xi$  by monotonicity. Taking the weighted sum with  $\hat{w}_n$  yields the nominal reweighted loss alongside three distinct error terms: a gradient bound ( $I_{\text{grad}}$ ), a linearization error ( $I_{\text{lin}}$ ), and a tail bound ( $I_{\text{tail}}$ ):

$$\begin{aligned} & \sum_{n=1}^N \hat{w}_n \int_{\Xi} \ell_{\text{DH}}(\vartheta; \hat{x}) \nu'_n(d\hat{x}) \\ & \leq \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \underbrace{\sum_{n=1}^N \hat{w}_n \int_{\Xi} \|g_n(\vartheta)\|_{\kappa, \lambda, p}^* d_{\kappa, \lambda}(X_n, \hat{x}) \nu'_n(d\hat{x})}_{I_{\text{grad}}} \\ & \quad + \underbrace{\zeta \sum_{n=1}^N \hat{w}_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x}) \nu'_n(d\hat{x})}_{I_{\text{lin}}} + \underbrace{2B_\ell \sum_{n=1}^N \hat{w}_n \nu'_n(B_n^c)}_{I_{\text{tail}}}. \end{aligned}$$

For the gradient bound  $I_{\text{grad}}$ , Hölder's inequality and the feasibility of  $(\nu'_n)$  imply that  $\sum \hat{w}_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}) \leq \varepsilon^p$ , hence we have

$$\begin{aligned} I_{\text{grad}} & \leq \left( \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} \right)^{1/q} \left( \sum_{n=1}^N \hat{w}_n \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}) \right)^{1/p} \\ & \leq \Upsilon_{\hat{w}}(\vartheta) \varepsilon. \end{aligned}$$

Similarly, the linearization error satisfies  $I_{\text{lin}} \leq \zeta \varepsilon$ . Finally, since  $d_{\kappa, \lambda}(X_n, \hat{x}) > r_\zeta$  on  $B_n^c$ , Markov's inequality gives

$$\nu'_n(B_n^c) \leq r_\zeta^{-p} \int_{\Xi} d_{\kappa, \lambda}(X_n, \hat{x})^p \nu'_n(d\hat{x}),$$

and therefore

$$I_{\text{tail}} \leq 2B_\ell r_\zeta^{-p} \varepsilon^p.$$

Substituting these three bounds and taking the supremum over all feasible slices yields the uniform upper bound:

$$T_a(\varepsilon) \leq \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\hat{w}}(\vartheta) + \zeta \varepsilon + 2B_\ell r_\zeta^{-p} \varepsilon^p. \quad (27)$$

We now construct a matching lower bound by evaluating a specific discrete attack. If  $\Upsilon_w(\vartheta) = 0$ , then  $\|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} = 0$  for every  $n$  (since  $w_n > 0$ ), meaning  $\Upsilon_{\hat{w}}(\vartheta) = 0$  for all  $a \in A_\delta$ . Setting  $\nu'_n = \delta_{X_n}$  yields  $T_a(\varepsilon) \geq \sum_{n=1}^N \hat{w}_n L_n(\vartheta)$ , which is sufficient. We therefore assume  $\Upsilon_w(\vartheta) > 0$ .

For every  $a \in A_\delta$ , the constraint  $\|a - \mathbf{1}\|_\infty \leq \delta$  implies  $a_n \geq 1 - \delta$ . Thus,  $\Upsilon_{\hat{w}}(\vartheta)^q = \sum_{n=1}^N w_n a_n \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} \geq (1 - \delta) \sum_{n=1}^N w_n \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q} = (1 - \delta) \Upsilon_w(\vartheta)^q$ , meaning  $\Upsilon_{\hat{w}}(\vartheta)$  is uniformly bounded away from 0. Define the deterministic adversarial perturbations:

$$\hat{X}_n^a := X_n + \varepsilon h(g_n(\vartheta)) \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q}.$$

The transport distance is  $d_{\kappa,\lambda}(X_n, \hat{X}_n^a) = \varepsilon \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q}$  using that  $h(g_n(\vartheta))$  is a unit vector. Recalling that  $(q-1)p = q$ , the total transport cost evaluates exactly to the budget:

$$\sum_{n=1}^N \hat{w}_n d_{\kappa,\lambda}(X_n, \hat{X}_n^a)^p = \varepsilon^p \Upsilon_{\hat{w}}(\vartheta)^{-q} \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*q} = \varepsilon^p \Upsilon_{\hat{w}}(\vartheta)^{-q} \Upsilon_{\hat{w}}(\vartheta)^q = \varepsilon^p.$$

Hence, the discrete point-mass choice  $\nu'_n = \delta_{\hat{X}_n^a}$  is feasible for  $T_a(\varepsilon)$ . Furthermore,  $d_{\kappa,\lambda}(X_n, \hat{X}_n^a) \rightarrow 0$  uniformly in  $a \in A_\delta$  as  $\varepsilon \downarrow 0$ , and by Assumption 3.5 these perturbations remain in  $\Xi$ .

Using the linearization error (26) at the perturbed points yields:

$$\ell_{\text{DH}}(\vartheta; \hat{X}_n^a) \geq L_n(\vartheta) + \langle g_n(\vartheta), \hat{X}_n^a - X_n \rangle - \zeta d_{\kappa,\lambda}(X_n, \hat{X}_n^a).$$

Multiplying by  $\hat{w}_n$  and summing over  $n$ , we explicitly evaluate the gradient inner product. We proceed in three steps: first, we substitute the definition of the perturbation  $\hat{X}_n^a$ ; second, we apply the dual maximizer property  $\langle g_n(\vartheta), h(g_n(\vartheta)) \rangle = \|g_n(\vartheta)\|_{\kappa,\lambda,p}^*$ ; and third, we recognize the definition of  $\Upsilon_{\hat{w}}(\vartheta)^q = \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*q}$ . This yields:

$$\begin{aligned} \sum_{n=1}^N \hat{w}_n \langle g_n(\vartheta), \hat{X}_n^a - X_n \rangle &= \sum_{n=1}^N \hat{w}_n \left\langle g_n(\vartheta), \varepsilon h(g_n(\vartheta)) \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q} \right\rangle \\ &= \varepsilon \Upsilon_{\hat{w}}(\vartheta)^{1-q} \sum_{n=1}^N \hat{w}_n \langle g_n(\vartheta), h(g_n(\vartheta)) \rangle \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*(q-1)} \\ &= \varepsilon \Upsilon_{\hat{w}}(\vartheta)^{1-q} \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*q} \\ &= \varepsilon \Upsilon_{\hat{w}}(\vartheta)^{1-q} \Upsilon_{\hat{w}}(\vartheta)^q \\ &= \varepsilon \Upsilon_{\hat{w}}(\vartheta). \end{aligned}$$

Substituting this into the aggregated linear bound directly confirms the lower bound:

$$\begin{aligned} T_a(\varepsilon) &\geq \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\hat{w}}(\vartheta) - \zeta \sum_{n=1}^N \hat{w}_n d_{\kappa,\lambda}(X_n, \hat{X}_n^a) \\ &\geq \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\hat{w}}(\vartheta) - \zeta \varepsilon, \end{aligned}$$

uniformly over  $a \in A_\delta$ . Combining this with (27) and noting  $\varepsilon^p = o(\varepsilon)$  as  $\varepsilon \downarrow 0$  proves (25).  $\square$

**Pure reweighting expansion.** We isolate the effect of the  $\phi$ -divergence budget  $\tau$  by analyzing the pure reweighting problem. Recall that  $L_n(\vartheta) := \ell_{\text{DH}}(\vartheta; X_n)$  and

$$\bar{L}_w(\vartheta) := \sum_{n=1}^N w_n L_n(\vartheta), \quad \sigma_w(\vartheta) := \left( \sum_{n=1}^N w_n (L_n(\vartheta) - \bar{L}_w(\vartheta))^2 \right)^{1/2}.$$

Define the maximum expected loss achievable through anchor reweighting under budget  $\tau$  by

$$R_\phi(\tau) := \sup_{a \in Q_\phi(\tau)} \sum_{n=1}^N w_n a_n L_n(\vartheta).$$

With these aggregate statistics defined, we can apply a standard second-order Taylor expansion to the  $\phi$ -divergence constraint to extract the asymptotic rate of the pure reweighting problem.

**Lemma C.9** (Small-budget expansion of the reweighting problem). *Let  $\phi$  be a locally regular entropy function. Then, as  $\tau \downarrow 0$ , the reweighting problem admits the expansion*

$$R_\phi(\tau) = \bar{L}_w(\vartheta) + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\sqrt{\tau}). \quad (28)$$

*Proof of Lemma C.9.* We prove the expansion by establishing matching upper and lower bounds in the small-budget regime. The key idea is that for sufficiently small  $\tau$ , every feasible reweighting stays strictly close to the uniform baseline  $a = \mathbf{1}$ . This forces the divergence constraint into a local regime where  $\phi$  is regular and reduces the optimization to maximizing a linear objective subject to a weighted variance constraint.

If  $\sigma_w(\vartheta) = 0$ , then  $L_n(\vartheta) = \bar{L}_w(\vartheta)$  for all  $n$  (since  $w_n > 0$ ). For any feasible  $a \in Q_\phi(\tau)$ , the expected loss is exactly  $\sum_{n=1}^N w_n a_n L_n(\vartheta) = \bar{L}_w(\vartheta) \sum_{n=1}^N w_n a_n = \bar{L}_w(\vartheta)$ , yielding  $R_\phi(\tau) = \bar{L}_w(\vartheta)$  for all  $\tau \geq 0$ , which trivially satisfies the expansion. We therefore assume  $\sigma_w(\vartheta) > 0$ .

Define the centered losses  $Y_n(\vartheta) := L_n(\vartheta) - \bar{L}_w(\vartheta)$ . By definition of the weighted empirical mean and variance, they satisfy  $\sum_{n=1}^N w_n Y_n(\vartheta) = 0$  and  $\sum_{n=1}^N w_n Y_n(\vartheta)^2 = \sigma_w(\vartheta)^2$ .

Fix  $0 < \zeta < \phi''(1)/2$ . Because  $\phi$  is locally regular, 1 is its unique minimizer and  $\phi'(1) = 0$ . Since  $\phi$  is twice continuously differentiable near 1, a second-order Taylor expansion yields  $\phi(1+s) = \frac{1}{2}\phi''(1)s^2 + o(s^2)$ . Thus, there exists  $\delta_\zeta \in (0, 1)$  such that for all  $|s| \leq \delta_\zeta$ , the divergence is locally bounded by quadratics:

$$\left(\frac{\phi''(1)}{2} - \zeta\right) s^2 \leq \phi(1+s) \leq \left(\frac{\phi''(1)}{2} + \zeta\right) s^2. \quad (29)$$

For this quadratic bound to be useful, we must guarantee that small budgets restrict all feasible weights to this neighborhood. For any feasible  $a \in Q_\phi(\tau)$ , the constraints  $a_n \geq 0$  and  $\sum_{n=1}^N w_n a_n = 1$  strictly bound the weights:  $a_n \leq 1/\min_k w_k$ . Define the compact set  $K_{\delta_\zeta} := [0, 1/\min_k w_k] \setminus (1 - \delta_\zeta, 1 + \delta_\zeta)$ . Since  $1 \notin K_{\delta_\zeta}$  and  $\phi$  is strictly positive away from 1, its minimum on this set is strictly positive:  $m_\zeta := \min_{s \in K_{\delta_\zeta}} \phi(s) > 0$ . If any weight strays such that  $|a_n - 1| \geq \delta_\zeta$ , then  $a_n \in K_{\delta_\zeta}$ , forcing  $\tau \geq \sum_{k=1}^N w_k \phi(a_k) \geq w_n \phi(a_n) \geq (\min_k w_k) m_\zeta$ . Consequently, for all  $\tau < (\min_k w_k) m_\zeta$ , every feasible  $a \in Q_\phi(\tau)$  must satisfy  $|a_n - 1| < \delta_\zeta$  for all  $n$ .

We next bound the objective. For any feasible  $a$ , define the deviations  $u_n := a_n - 1$ . Since  $\sum_{n=1}^N w_n u_n = 0$ , the objective simplifies to  $\sum_{n=1}^N w_n a_n L_n(\vartheta) = \sum_{n=1}^N w_n (1 + u_n) (\bar{L}_w(\vartheta) + Y_n(\vartheta)) = \bar{L}_w(\vartheta) + \sum_{n=1}^N w_n u_n Y_n(\vartheta)$ . For sufficiently small  $\tau$ , we know  $|u_n| < \delta_\zeta$ . Applying the lower bound in (29) with  $s = u_n$  bounds the variance of the weights:

$$\tau \geq \sum_{n=1}^N w_n \phi(1 + u_n) \geq \left(\frac{\phi''(1)}{2} - \zeta\right) \sum_{n=1}^N w_n u_n^2.$$

Applying the Cauchy–Schwarz inequality to the objective deviation yields:

$$\sum_{n=1}^N w_n u_n Y_n(\vartheta) \leq \left(\sum_{n=1}^N w_n u_n^2\right)^{1/2} \left(\sum_{n=1}^N w_n Y_n(\vartheta)^2\right)^{1/2} \leq \sqrt{\frac{\tau}{\phi''(1)/2 - \zeta}} \sigma_w(\vartheta).$$

Taking the supremum over all feasible  $a$  gives the uniform upper bound:

$$R_\phi(\tau) \leq \bar{L}_w(\vartheta) + \frac{\sigma_w(\vartheta)}{\sqrt{\phi''(1)/2 - \zeta}} \sqrt{\tau}. \quad (30)$$

For the matching lower bound, we construct an explicit feasible reweighting. Define the weights  $a_n := 1 + \alpha_\tau Y_n(\vartheta)$ , where  $\alpha_\tau := \sqrt{\tau} / (\sigma_w(\vartheta) \sqrt{\phi''(1)/2 + \zeta})$ . By construction,  $\sum_{n=1}^N w_n a_n = 1$ . For sufficiently small  $\tau$ , we have  $|\alpha_\tau Y_n(\vartheta)| \leq \delta_\zeta$ , allowing us to use the upper bound in (29) to verify feasibility:

$$\sum_{n=1}^N w_n \phi(a_n) \leq \left( \frac{\phi''(1)}{2} + \zeta \right) \alpha_\tau^2 \sum_{n=1}^N w_n Y_n(\vartheta)^2 = \left( \frac{\phi''(1)}{2} + \zeta \right) \alpha_\tau^2 \sigma_w(\vartheta)^2 = \tau,$$

confirming  $a \in Q_\phi(\tau)$ . Evaluating the expected loss under  $a$  yields  $\sum_{n=1}^N w_n a_n L_n(\vartheta) = \bar{L}_w(\vartheta) + \alpha_\tau \sum_{n=1}^N w_n Y_n(\vartheta) L_n(\vartheta)$ . Since  $\sum_{n=1}^N w_n Y_n(\vartheta) \bar{L}_w(\vartheta) = 0$ , the inner product simplifies to  $\sum_{n=1}^N w_n Y_n(\vartheta) (Y_n(\vartheta) + \bar{L}_w(\vartheta)) = \sigma_w(\vartheta)^2$ . Substitution provides the lower bound:

$$R_\phi(\tau) \geq \bar{L}_w(\vartheta) + \frac{\sigma_w(\vartheta)}{\sqrt{\phi''(1)/2 + \zeta}} \sqrt{\tau}.$$

Combining this lower bound with (30) and letting  $\zeta \downarrow 0$  completes the proof of (28).  $\square$

**Joint OT- $\phi$  expansion.** We now combine the exact reduction, the uniform OT expansion, and the pure reweighting expansion.

**Theorem 3.6.** (Joint first-order approximation) *Suppose  $\phi$  is locally regular,  $p \in (1, \infty)$  and Assumption 3.5 holds. For every fixed  $\vartheta \in \Theta \times \mathbb{R}$ , as  $(\varepsilon, \tau) \rightarrow (0, 0)$*

$$V_\vartheta(\varepsilon, \tau) = \bar{L}_w(\vartheta) + \Upsilon_w(\vartheta) \varepsilon + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\varepsilon + \sqrt{\tau}). \quad (10)$$

*Proof of Theorem 3.6.* We first use the exact reduction in Proposition C.7 to rewrite the problem as a supremum over reweightings  $a \in Q_\phi(\tau)$ . We then apply the uniform OT expansion from Lemma C.8 and verify that it applies to all feasible  $a \in Q_\phi(\tau)$  for small  $\tau$ . Finally, we separate the reweighting and OT contributions by an elementary supremum inequality and combine the result with the small- $\tau$  expansion of the pure  $\phi$ -reweighting problem in Lemma C.9. For each  $a \in Q_\phi(\tau)$ , set  $\hat{w}_n := w_n a_n$  for  $n = 1, \dots, N$ ; the dependence  $\hat{w} = \hat{w}(a)$  is implicit throughout this proof. By Proposition C.7,

$$V_\vartheta(\varepsilon, \tau) = \sup_{a \in Q_\phi(\tau)} T_a(\varepsilon).$$

We next verify that the uniform OT expansion applies over the entire feasible set. Fix  $\delta \in (0, 1)$ . We show that  $Q_\phi(\tau) \subseteq A_\delta$  for all sufficiently small  $\tau$ . Define  $w_{\min} = \min_n w_n$ . Note that  $w_{\min} > 0$  and for every  $a \in Q_\phi(\tau)$ , the normalization gives  $0 \leq a_n \leq 1/w_{\min}$ . Since 1 is the unique zero of  $\phi$ , lower semicontinuity implies

$$m_\delta := \inf_{s \in [0, 1/w_{\min}] \setminus (1-\delta, 1+\delta)} \phi(s) > 0.$$

If some coordinate satisfied  $|a_j - 1| \geq \delta$ , then  $\sum_{n=1}^N w_n \phi(a_n) \geq w_j \phi(a_j) \geq w_{\min} m_\delta$ . Hence, for  $\tau < w_{\min} m_\delta$ , no such coordinate exists, and therefore  $Q_\phi(\tau) \subseteq A_\delta$ . Therefore, Lemma C.8 gives

$$\sup_{a \in Q_\phi(\tau)} \left| T_a(\varepsilon) - \sum_{n=1}^N \hat{w}_n L_n(\vartheta) - \varepsilon \Upsilon_{\hat{w}}(\vartheta) \right| = o(\varepsilon), \quad \text{as } \varepsilon \downarrow 0, \quad (31)$$

uniformly for all sufficiently small  $\tau$ . Equivalently, we may write

$$T_a(\varepsilon) = \sum_{n=1}^N \widehat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\widehat{w}}(\vartheta) + r_a(\varepsilon, \tau),$$

where the remainder satisfies

$$\sup_{a \in Q_\phi(\tau)} |r_a(\varepsilon, \tau)| = o(\varepsilon), \quad \text{as } \varepsilon \downarrow 0. \quad (32)$$

We now separate the contribution of the reweighting term and the OT correction. Recall the definition of  $R_\phi(\tau)$ , and compare

$$\sup_{a \in Q_\phi(\tau)} \left( \sum_{n=1}^N \widehat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\widehat{w}}(\vartheta) + r_a(\varepsilon, \tau) \right)$$

with  $R_\phi(\tau) + \varepsilon \Upsilon_w(\vartheta)$ . To control this difference, we use the elementary inequality  $\left| \sup_{a \in Q} (A_a + B_a) - \sup_{a \in Q} A_a - b \right| \leq \sup_{a \in Q} |B_a - b|$ , which follows from  $A_a + B_a \leq A_a + b + |B_a - b|$  together with the corresponding reverse inequality obtained by writing  $A_a = (A_a + B_a) - B_a$ . Applying this inequality with  $A_a := \sum_{n=1}^N \widehat{w}_n L_n(\vartheta)$ ,  $B_a := \varepsilon \Upsilon_{\widehat{w}}(\vartheta) + r_a(\varepsilon, \tau)$  and  $b := \varepsilon \Upsilon_w(\vartheta)$  we obtain

$$\left| V_\vartheta(\varepsilon, \tau) - R_\phi(\tau) - \varepsilon \Upsilon_w(\vartheta) \right| \leq \sup_{a \in Q_\phi(\tau)} |r_a(\varepsilon, \tau)| + \varepsilon \sup_{a \in Q_\phi(\tau)} |\Upsilon_{\widehat{w}}(\vartheta) - \Upsilon_w(\vartheta)|.$$

By (32), the first term is  $o(\varepsilon)$ . The second term is controlled by the collapse of  $Q_\phi(\tau)$  to  $\mathbf{1}$ . Since  $a \mapsto \Upsilon_{\widehat{w}(a)}(\vartheta)$  is continuous, and  $Q_\phi(\tau)$  collapses to  $\{\mathbf{1}\}$  as  $\tau \downarrow 0$ , in the sense that  $\sup_{a \in Q_\phi(\tau)} \|a - \mathbf{1}\|_\infty \rightarrow 0$ , we have  $\sup_{a \in Q_\phi(\tau)} |\Upsilon_{\widehat{w}}(\vartheta) - \Upsilon_w(\vartheta)| \rightarrow 0$ . Hence,

$$\varepsilon \sup_{a \in Q_\phi(\tau)} |\Upsilon_{\widehat{w}}(\vartheta) - \Upsilon_w(\vartheta)| = o(\varepsilon). \quad (33)$$

We conclude that

$$V_\vartheta(\varepsilon, \tau) = R_\phi(\tau) + \Upsilon_w(\vartheta)\varepsilon + o(\varepsilon). \quad (34)$$

Finally, with Lemma C.9 we can substitute the expansion of  $R_\phi(\tau)$  into (34) and get

$$V_\vartheta(\varepsilon, \tau) = \bar{L}_w(\vartheta) + \Upsilon_w(\vartheta)\varepsilon + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\varepsilon + \sqrt{\tau}).$$

The budget axes are included: for  $\tau = 0$ ,  $Q_\phi(0) = \{\mathbf{1}\}$ , so the exact reduction and Lemma C.8 give the pure Wasserstein expansion directly. For  $\varepsilon = 0$ , zero transport cost implies  $T_a(0) = \sum_n \widehat{w}_n L_n(\vartheta)$ , so  $V_\vartheta(0, \tau) = R_\phi(\tau)$  and Lemma C.9 applies.  $\square$

### C.3.2 Structured Two-Budget Attack

The structured attack combines the leading  $\phi$ -reweighting direction with the Wasserstein steepest-ascent perturbation conditional on the reweighted anchors. The proof verifies feasibility of the reweighting and transport budgets and then evaluates the loss expansion along this constructed law. Recall  $Y_n(\vartheta)$ ,  $a_n$ ,  $\widehat{w}_n := w_n a_n$ ,  $\Upsilon_{\widehat{w}}(\vartheta)$ ,  $\widehat{X}_n$ , and  $\widehat{\eta}_{\varepsilon, \tau}$  from Section 3.

**Theorem 3.7** (First-order accuracy of the adversary). *Suppose  $\phi$  is locally regular,  $p \in (1, \infty)$  and Assumption 3.5 holds. For sufficiently small  $(\varepsilon, \tau)$ ,  $\hat{\eta}_{\varepsilon, \tau} \in \mathcal{B}_{\varepsilon, \tau}(\mu_w)$ , and as  $(\varepsilon, \tau) \rightarrow (0, 0)$ ,*

$$V_{\vartheta}(\varepsilon, \tau) - \mathbb{E}_{X \sim \hat{\eta}_{\varepsilon, \tau}}[\ell_{\text{DH}}(\vartheta; X)] = o(\varepsilon + \sqrt{\tau}). \quad (12)$$

*Proof of Theorem 3.7.* We first verify the reweighting constraint. If  $\sigma_w(\vartheta) = 0$ , then  $Y_n(\vartheta) = 0$  for all  $n$ , because  $w_n > 0$ , and the convention  $a = \mathbf{1}$  gives  $a \in Q_{\phi}(\tau)$ . Suppose therefore that  $\sigma_w(\vartheta) > 0$ . Since  $\sum_{n=1}^N w_n Y_n(\vartheta) = 0$ , we have

$$\sum_{n=1}^N w_n a_n = 1 + \alpha_{\tau} \sum_{n=1}^N w_n Y_n(\vartheta) = 1.$$

The nonnegativity and  $\phi$ -budget constraints hold by the definition of  $\alpha_{\tau}$ . Hence  $a \in Q_{\phi}(\tau)$ . We next derive the small- $\tau$  behavior of  $\alpha_{\tau}$ . Since  $\phi$  is  $C^2$  around 1, with  $\phi(1) = \phi'(1) = 0$ , we have

$$\sum_{n=1}^N w_n \phi(1 + \alpha Y_n(\vartheta)) = \frac{\phi''(1)}{2} \alpha^2 \sum_{n=1}^N w_n Y_n(\vartheta)^2 + o(\alpha^2) = \frac{\phi''(1)}{2} \alpha^2 \sigma_w(\vartheta)^2 + o(\alpha^2).$$

Thus the maximal feasible  $\alpha_{\tau}$  satisfies

$$\alpha_{\tau} = \frac{\sqrt{2\tau}}{\sqrt{\phi''(1) \sigma_w(\vartheta)}} + o(\sqrt{\tau}), \quad \tau \downarrow 0. \quad (35)$$

We show that the perturbed trajectories are admissible for all sufficiently small  $(\varepsilon, \tau)$ :  $\hat{X}_n \in \Xi$ . If  $g_n(\vartheta) = 0$ , then  $\hat{X}_n = X_n \in \Xi$ . Otherwise, since  $X_n$  is in the interior of  $\Xi$ ,  $X_n \in \text{int}(\Xi)$ , and  $\|h(g_n(\vartheta))\|_{\kappa, \lambda, p} = 1$ , there exists  $\bar{t}_n > 0$  such that

$$X_n + t h(g_n(\vartheta)) \in \Xi, \quad 0 \leq t \leq \bar{t}_n.$$

The step length in direction  $h(g_n(\vartheta))$  is

$$t_n(\varepsilon, \tau) = \varepsilon \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q}.$$

Since  $\alpha_{\tau} \rightarrow 0$ , the constructed vector  $a = a(\tau)$  satisfies  $a \rightarrow \mathbf{1}$  and  $\hat{w} \rightarrow w$  as  $\tau \downarrow 0$ . Hence, the quantities  $\Upsilon_{\hat{w}}(\vartheta)$  remain bounded away from zero whenever the transport sensitivity  $\Upsilon_w$  is nonzero. Therefore,  $t_n(\varepsilon, \tau) \downarrow 0$  as  $\varepsilon \downarrow 0$ , uniformly over the finitely many samples. Therefore,  $\hat{X}_n \in \Xi$  for all sufficiently small  $(\varepsilon, \tau)$ .

We next verify the transport budget constraint. If  $\Upsilon_{\hat{w}}(\vartheta) = 0$ , then  $\hat{X}_n = X_n$  for all  $n$ , so the transport cost is zero. Assume therefore that  $\Upsilon_{\hat{w}}(\vartheta) > 0$ . By construction,  $\hat{X}_n - X_n = \varepsilon h(g_n(\vartheta)) \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q}$ . Since  $d_{\kappa, \lambda}$  is induced by  $\|\cdot\|_{\kappa, \lambda, p}$ , and since  $\|h(g_n(\vartheta))\|_{\kappa, \lambda, p} = 1$  whenever  $g_n(\vartheta) \neq 0$ , we get

$$d_{\kappa, \lambda}(\hat{X}_n, X_n) = \varepsilon \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*(q-1)} \Upsilon_{\hat{w}}(\vartheta)^{1-q}.$$

For  $g_n(\vartheta) = 0$ , both sides are zero, so the same identity still holds. Taking the  $p$ -th power and using  $(q-1)p = q$  gives

$$d_{\kappa, \lambda}(\hat{X}_n, X_n)^p = \varepsilon^p \Upsilon_{\hat{w}}(\vartheta)^{-q} \|g_n(\vartheta)\|_{\kappa, \lambda, p}^{*q}.$$

Therefore, weighting with  $\hat{w}_n$  and summing over  $n$  shows that the transport budget constraint is satisfied:

$$\begin{aligned} \sum_{n=1}^N \hat{w}_n d_{\kappa,\lambda}(\hat{X}_n, X_n)^p &= \varepsilon^p \Upsilon_{\hat{w}}(\vartheta)^{-q} \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*q} \\ &= \varepsilon^p \Upsilon_{\hat{w}}(\vartheta)^{-q} \Upsilon_{\hat{w}}(\vartheta)^q = \varepsilon^p. \end{aligned}$$

To verify feasibility in the full two-budget ambiguity set, define

$$\pi := \sum_{n=1}^N w_n \delta_{(X_n, \hat{X}_n)}, \quad \pi' := \sum_{n=1}^N \hat{w}_n \delta_{(X_n, \hat{X}_n)}.$$

Then  $\pi_1 = \mu_w$  and  $\pi'_2 = \hat{\eta}_{\varepsilon,\tau}$ . Moreover,  $\pi' \ll \pi$ . If several pairs  $(X_n, \hat{X}_n)$  coincide, the density  $d\pi'/d\pi$  is the corresponding  $w$ -weighted average of the coefficients  $a_n$ ; hence Jensen's inequality and  $a \in Q_\phi(\tau)$  give

$$D_\phi(\pi' \parallel \pi) \leq \sum_{n=1}^N w_n \phi(a_n) \leq \tau.$$

Together with the transport-budget calculation above, this shows that  $\pi, \pi'$  are admissible in Definition 3.2, and therefore

$$\hat{\eta}_{\varepsilon,\tau} \in \mathcal{B}_{\varepsilon,\tau}(\mu_w).$$

It remains to evaluate the expected loss under  $\hat{\eta}_{\varepsilon,\tau}$ . Since  $\alpha_\tau \rightarrow 0$ , the constructed vector  $a = a(\tau)$  satisfies  $a \rightarrow \mathbf{1}$  and  $\hat{w} \rightarrow w$  as  $\tau \downarrow 0$ . Hence, continuity gives  $\Upsilon_{\hat{w}}(\vartheta) = \Upsilon_w(\vartheta) + o(1)$  (see (33)). If  $\Upsilon_w(\vartheta) = 0$ , then all  $g_n(\vartheta) = 0$ , because  $w_n > 0$ , and the transport part contributes nothing at first order. Otherwise,  $\Upsilon_{\hat{w}}(\vartheta)$  is bounded away from zero for all sufficiently small  $\tau$ . Hence the constructed perturbations satisfy  $d_{\kappa,\lambda}(\hat{X}_n, X_n) = O(\varepsilon)$  uniformly in  $n$ . By Assumption 3.5

$$\ell_{\text{DH}}(\vartheta; \hat{X}_n) = L_n(\vartheta) + \langle g_n(\vartheta), \hat{X}_n - X_n \rangle + o(\varepsilon).$$

Multiplying by  $\hat{w}_n$ , summing over  $n$ , and using boundedness of  $a$  for small  $\tau$ , gives

$$\sum_{n=1}^N \hat{w}_n \ell_{\text{DH}}(\vartheta; \hat{X}_n) = \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \sum_{n=1}^N \hat{w}_n \langle g_n(\vartheta), \hat{X}_n - X_n \rangle + o(\varepsilon).$$

The linear transport term is

$$\sum_{n=1}^N \hat{w}_n \langle g_n(\vartheta), \hat{X}_n - X_n \rangle = \varepsilon \Upsilon_{\hat{w}}(\vartheta)^{1-q} \sum_{n=1}^N \hat{w}_n \|g_n(\vartheta)\|_{\kappa,\lambda,p}^{*q} = \varepsilon \Upsilon_{\hat{w}}(\vartheta).$$

Therefore, we conclude that

$$\mathbb{E}_{X \sim \hat{\eta}_{\varepsilon,\tau}} [\ell_{\text{DH}}(\vartheta; X)] = \sum_{n=1}^N \hat{w}_n L_n(\vartheta) + \varepsilon \Upsilon_{\hat{w}}(\vartheta) + o(\varepsilon). \quad (36)$$

If  $\sigma_w(\vartheta) > 0$ , using  $Y_n(\vartheta) = L_n(\vartheta) - \bar{L}_w(\vartheta)$ , the loss evaluates to

$$\sum_{n=1}^N w_n a_n L_n(\vartheta) = \sum_{n=1}^N w_n (1 + \alpha_\tau Y_n(\vartheta)) (\bar{L}_w(\vartheta) + Y_n(\vartheta)) = \bar{L}_w(\vartheta) + \alpha_\tau \sigma_w(\vartheta)^2.$$

Using (35), this becomes

$$\sum_{n=1}^N \widehat{w}_n L_n(\vartheta) = \bar{L}_w(\vartheta) + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\sqrt{\tau}). \quad (37)$$

If  $\sigma_w(\vartheta) = 0$ , the same display holds with zero  $\sqrt{\tau}$ -term. Finally,

$$\varepsilon \Upsilon_{\widehat{w}}(\vartheta) = \varepsilon \Upsilon_w(\vartheta) + o(\varepsilon).$$

Combining (36) and (37) yields

$$\mathbb{E}_{X \sim \widehat{\eta}_{\varepsilon, \tau}} [\ell_{\text{DH}}(\vartheta; X)] = \bar{L}_w(\vartheta) + \Upsilon_w(\vartheta) \varepsilon + \sqrt{\frac{2}{\phi''(1)}} \sigma_w(\vartheta) \sqrt{\tau} + o(\varepsilon + \sqrt{\tau}).$$

Comparing this expansion with Theorem 3.6 gives

$$V_{\vartheta}(\varepsilon, \tau) - \mathbb{E}_{X \sim \widehat{\eta}_{\varepsilon, \tau}} [\ell_{\text{DH}}(\vartheta; X)] = o(\varepsilon + \sqrt{\tau}),$$

which is (12) and completes the proof.  $\square$

## D Numerical Setup and Supplementary Results

This appendix provides further details on the numerical experiments. We first specify the Optimized Certainty Equivalent (OCE) risk measures (D.1), contingent-claim payoffs (D.2), the multistep WRAP algorithm (D.3), and the data-generating processes (D.4). We define the performance metrics (D.5) and present the experimental setup and results (D.6), covering the Heston experiments (D.6.1) and the GAD study (D.6.2), together with the corresponding calibration data and descriptive statistics. This section concludes with a comparison of the one-step and multistep WRAP variants in the stationary Heston setting (D.6.3).

### D.1 Optimized Certainty Equivalent Risk Measures

The numerical experiments adopt two standard Optimized Certainty Equivalent (OCE) risk measures. OCEs were introduced by Ben-Tal & Teboulle (1986) and later developed as a class of convex risk measures by Ben-Tal & Teboulle (2007). Let  $\mathcal{Z}$  denote the domain of real-valued random variables under consideration. A normalized OCE loss is a convex and nondecreasing function  $\ell : \mathbb{R} \rightarrow \mathbb{R}$  satisfying  $\ell(0) = 0$  and  $\ell(u) \geq u$  for all  $u \in \mathbb{R}$ . The associated OCE risk measure is

$$\rho_{\ell}(Z) = \inf_{m \in \mathbb{R}} \{m + \mathbb{E}[\ell(-Z - m)]\}. \quad (38)$$

Here  $Z$  denotes terminal profit and loss, so larger values of  $Z$  are preferred and  $-Z$  is the corresponding loss. For  $Z = \text{PnL}(\theta, X)$ , the objective inside the infimum in (38) gives the samplewise deep-hedging loss  $\ell_{\text{DH}}(\vartheta; X)$  defined in (3).

**Conditional Value at Risk.** Conditional Value at Risk (CVaR), also known as Expected Shortfall, is a coherent risk measure that averages losses beyond a high quantile. The variational representation below follows Rockafellar & Uryasev (2000, 2002). For a confidence level  $\alpha \in (0, 1)$ , define

$$\ell_{\alpha}(u) = \frac{(u)_{+}}{1 - \alpha}, \quad (u)_{+} := \max\{u, 0\}.$$

Substituting this loss into (38) yields

$$\rho_\alpha(Z) = \inf_{m \in \mathbb{R}} \left\{ m + \frac{1}{1-\alpha} \mathbb{E}[(-Z - m)_+] \right\} = \text{CVaR}_\alpha(-Z). \quad (39)$$

Thus, CVaR evaluates the average of the largest  $1 - \alpha$  fraction of losses, with the standard adjustment for probability mass at the quantile. The minimizers in (39) are the  $\alpha$ -quantiles of  $-Z$ ; in particular, one may choose  $m = \text{VaR}_\alpha(-Z)$ .

**Entropic risk.** The entropic risk measure is obtained from the exponential OCE loss (Ben-Tal & Teboulle, 2007). For a risk-aversion parameter  $\lambda > 0$ , define

$$\ell_\lambda(u) = \frac{\exp(\lambda u) - 1}{\lambda}.$$

Whenever  $\mathbb{E}[\exp(-\lambda Z)] < \infty$ ,

$$\begin{aligned} \rho_\lambda(Z) &= \inf_{m \in \mathbb{R}} \left\{ m + \frac{1}{\lambda} \mathbb{E}[\exp(\lambda(-Z - m)) - 1] \right\} \\ &= \frac{1}{\lambda} \log \mathbb{E}[\exp(-\lambda Z)]. \end{aligned} \quad (40)$$

The entropic risk measure is smooth and convex and penalizes large losses exponentially. Unlike CVaR, it is not positively homogeneous.

**Convex combinations with expected loss.** Let  $\rho$  be an OCE risk measure with a convex, nondecreasing loss  $\ell$  satisfying  $\ell(0) = 0$  and  $\ell(u) \geq u$  for all  $u \in \mathbb{R}$ . Then, for  $\omega \in [0, 1]$ , the loss  $\ell_\omega(u) := (1 - \omega)u + \omega \ell(u)$  satisfies the same conditions. Substitution into (38) yields, for integrable  $Z$  with finite  $\rho(Z)$ ,

$$\rho_\omega(Z) = (1 - \omega)\mathbb{E}[-Z] + \omega\rho(Z).$$

In particular, Mean-CVaR (M-CVaR) admits the representation

$$\begin{aligned} \rho_{\omega, \alpha}(Z) &= \inf_{m \in \mathbb{R}} \left\{ m + \mathbb{E} \left[ (1 - \omega)(-Z - m) + \frac{\omega}{1 - \alpha} (-Z - m)_+ \right] \right\} \\ &= (1 - \omega)\mathbb{E}[-Z] + \omega \text{CVaR}_\alpha(-Z). \end{aligned} \quad (41)$$

## D.2 Contingent Claims

We collect the payoff functionals used in the numerical examples. Recall that the agent aims to hedge a contingent claim with payoff function  $\Phi : \Xi \rightarrow \mathbb{R}$  at maturity. Let  $K > 0$  denote the strike price and  $S^1 \in \mathbb{R}^{T+1}$  denote the price process of the considered asset.

The Heston experiments consider a *European call option* on the stock with payoff

$$\Phi^{\text{European}}(X) = (S_T^1 - K)_+.$$

The GAD experiments consider an *arithmetic Asian call* with payoff

$$\Phi^{\text{Asian}}(X) = \left( \frac{1}{T+1} \sum_{t=0}^T S_t^1 - K \right)_+.$$

### D.3 Multistep WRAP

---

**Algorithm 2** Multistep Wasserstein Reweighting Adversarial Perturbation (WRAP)

---

**Require:** Weighted sample  $\{(X_n, w_n)\}_{n=1}^N$ , batch  $\mathcal{I} \subseteq \{1, \dots, N\}$  with  $|\mathcal{I}| = M$ , hedge parameter  $\vartheta$ , transport budget  $\varepsilon$ , reweighting budget  $\tau$ , inner steps  $K$ , stepsize  $\eta$

**Ensure:** Adversarial weighted batch  $\{(\hat{X}_n, \hat{w}_n)\}_{n \in \mathcal{I}}$

```

1: Set  $\bar{w}_n \leftarrow w_n / \sum_{j \in \mathcal{I}} w_j$ ,  $a_n \leftarrow 1$ ,  $\hat{w}_n \leftarrow \bar{w}_n$ ,  $b_n \leftarrow 0$ , and  $d_n \leftarrow 0$  for all  $n \in \mathcal{I}$ .
2: for  $k = 1, \dots, K$  do
3:   Set  $\hat{X}_n \leftarrow X_n + b_n d_n$  for all  $n \in \mathcal{I}$ .
4:   Compute  $L_n \leftarrow \ell_{\text{DH}}(\vartheta; \hat{X}_n)$  for all  $n \in \mathcal{I}$ .
    $\phi$ -divergence reweighting
5:   if  $\tau = 0$  then
6:     Set  $a_n \leftarrow 1$  for all  $n \in \mathcal{I}$ .
7:   else
8:     Set  $\bar{L}_{\mathcal{I}} \leftarrow \sum_{n \in \mathcal{I}} \bar{w}_n L_n$  and  $Y_n \leftarrow L_n - \bar{L}_{\mathcal{I}}$  for all  $n \in \mathcal{I}$ .
9:     Choose the largest  $\alpha_\tau \geq 0$  such that  $1 + \alpha_\tau Y_n \geq 0$  for all  $n \in \mathcal{I}$  and  $\sum_{n \in \mathcal{I}} \bar{w}_n \phi(1 + \alpha_\tau Y_n) \leq \tau$ .
10:    Set  $a_n \leftarrow 1 + \alpha_\tau Y_n$  for all  $n \in \mathcal{I}$ .
11:   end if
12:   Set  $\hat{w}_n \leftarrow \bar{w}_n a_n$  for all  $n \in \mathcal{I}$ .
   Optimal-transport perturbation
13:   Compute  $G_n \leftarrow \nabla_{\hat{X}} \ell_{\text{DH}}(\vartheta; \hat{X}_n)$  for all  $n \in \mathcal{I}$ .
14:   Set  $\tilde{d}_n \leftarrow d_n + (\eta/\varepsilon)h(G_n)$  and  $d_n \leftarrow \tilde{d}_n / \|\tilde{d}_n\|_{\kappa, \lambda, p}$  for all  $n \in \mathcal{I}$ .
15:   Set  $\Upsilon_{\mathcal{I}} \leftarrow (\sum_{n \in \mathcal{I}} \hat{w}_n [\langle G_n, d_n \rangle]_+^q)^{1/q}$  and update  $b_n \leftarrow b_n + \eta [\langle G_n, d_n \rangle]_+^{q-1} \Upsilon_{\mathcal{I}}^{1-q}$  for all  $n \in \mathcal{I}$ .
16:   Rescale  $b_n \leftarrow \varepsilon b_n / (\sum_{j \in \mathcal{I}} \hat{w}_j b_j^p)^{1/p}$  for all  $n \in \mathcal{I}$ .
17: end for
18: Set  $\hat{X}_n \leftarrow X_n + b_n d_n$  and  $\hat{w}_n \leftarrow \bar{w}_n a_n$  for all  $n \in \mathcal{I}$ .
19: return  $\{(\hat{X}_n, \hat{w}_n)\}_{n \in \mathcal{I}}$ .

```

---

The multistep attack refines Algorithm 1 while preserving the same two-budget structure. The reweighting step is unchanged: at each inner iteration, the losses are centered and the adversarial multipliers are updated according to  $a_n = 1 + \alpha_\tau Y_n$ . The transport step differs. Rather than applying the closed-form first-order perturbation once, the algorithm writes  $\hat{X}_n = X_n + b_n d_n$ , where  $d_n$  is a normalized perturbation direction and  $b_n \geq 0$  is the pathwise allocation of the transport budget. Both are updated iteratively and projected back onto the weighted transport constraint.

Algorithm 2 is therefore not the exact first-order construction of Theorem 3.7, but an iterative refinement that builds on He et al. (2025). The  $\phi$ -block continues to tilt probability toward high-loss scenarios through the linearized  $\alpha_\tau$ -step, while the OT block updates the perturbation directions and budget allocations with learning rate  $\eta$ . At each adversarial iterate  $\hat{X}_n$ , the reweighting is recomputed before the transport update proceeds along the local ascent direction. When  $\tau = 0$ , the method reduces to a drift-aware iterative optimal-transport attack; when  $\varepsilon = 0$ , the transport updates vanish and the method reduces to the reweighting-only case of Algorithm 1.

In the numerical experiments, this multistep procedure serves as the stronger iterative benchmark against the one-step attack in Algorithm 1.

### D.4 Data-Generating Processes

We conduct the numerical experiments under two market models: the Heston stochastic-volatility model and the generalized affine diffusion (GAD) model. Both models are sampled at the trading dates  $t_j = j\Delta t$ ,  $j = 0, \dots, T$ , with maturity  $\mathcal{T} = T\Delta t$ . Following the notation introduced in Section 2,  $S_j$  denotes the vector of tradable prices,  $X_j$  the observed market state, and  $X = (X_0, \dots, X_T) \in \Xi$  the resulting trajectory.

**Heston model.** The first specification is the Heston stochastic-volatility model (Heston, 1993),

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{v_t} S_t dW_t^S, \\ dv_t &= \kappa(\theta_H - v_t) dt + \xi \sqrt{v_t} dW_t^v, \end{aligned} \quad d\langle W^S, W^v \rangle_t = \rho dt. \quad (42)$$

Here  $S_0 > 0$  and  $v_0 \geq 0$  are the initial stock price and variance,  $\mu \in \mathbb{R}$  is the stock-price drift,  $\kappa > 0$  the mean-reversion rate,  $\theta_H > 0$  the long-run variance,  $\xi > 0$  the volatility-of-variance parameter, and  $\rho \in [-1, 1]$  the instantaneous correlation between the Brownian motions  $W^S$  and  $W^v$ . We use  $\theta_H$  to distinguish the long-run variance from the hedging-policy parameter  $\theta$ .

We sample the variance process from its exact Cox–Ingersoll–Ross transition and approximate the integrated variance entering the stock-price update by the trapezoidal rule. The exact variance transition preserves nonnegativity without requiring the Feller condition.

In the Heston experiments, the claim is written on the stock, while the hedging portfolio may trade both the stock and a variance-linked instrument. At day  $t$ , the hedging network takes  $X_{0:t} = (\log S, v)_{0:t}$  as input, and portfolio gains are computed from changes in both tradable prices.

**Generalized affine diffusion model.** The second specification follows the generalized affine diffusion model used by Lütkebohmert et al. (2022) for robust deep hedging. It extends the classical affine-diffusion case, covered by the affine-process framework of Duffie et al. (2003). In continuous-time notation, we write

$$dS_t = (b_0 + b_1 S_t) dt + (a_0 + a_1 S_t)^\gamma dW_t. \quad (43)$$

Here  $S_0 > 0$  is the initial stock price,  $b_0, b_1 \in \mathbb{R}$  determine the affine drift, and  $a_0, a_1, \gamma$  determine the state-dependent volatility. The stock is the sole tradable asset and the only state variable observed by the hedging network, *i.e.*  $X = S$ .

## D.5 Performance Metrics

We evaluate each trained hedge on a held-out test set  $X_1, \dots, X_M \in \Xi$ . Recall that the terminal profit and loss generated by a hedging strategy  $\theta$  along a path  $X$  is

$$\text{PnL}(\theta, X) = p_0 + \sum_{t=0}^{T-1} \delta_t^\theta (X_{0:t})^\top (S_{t+1} - S_t) - \Phi(X).$$

*Mean PnL.* The average realized profit and loss is

$$\text{MeanPnL} := \frac{1}{M} \sum_{i=1}^M \text{PnL}(\theta, X_i).$$

*Empirical CVaR.* The empirical CVaR of the hedging loss is

$$\text{CVaR} := m + \frac{1}{1 - \alpha_{\text{cvar}}} \frac{1}{M} \sum_{i=1}^M (-\text{PnL}(\theta, X_i) - m)_+.$$

Here  $m$  is the empirical  $\alpha_{\text{cvar}}$ -quantile of the test losses  $-\text{PnL}(\theta, X_i)$ , with  $\alpha_{\text{cvar}} = 0.95$ .

*Empirical M-CVaR.* The empirical combination of mean loss and CVaR is

$$\text{M-CVaR} := -(1 - \omega) \text{MeanPnL} + \omega \text{CVaR}.$$

Here  $\omega \in [0, 1]$  controls the weight on CVaR, with  $\omega = 0.5$  in the M-CVaR experiments.  
*Empirical entropic risk.* The empirical entropic risk is

$$\text{EntR} := \frac{1}{\lambda} \log \left( \frac{1}{M} \sum_{i=1}^M \exp(-\lambda \text{PnL}(\theta, X_i)) \right).$$

Here  $\lambda > 0$  is the risk-aversion parameter.

*Transaction costs.* The average cumulative transaction cost (cf. (45)) per test path is

$$\text{TotalTC} := \frac{1}{M} \sum_{i=1}^M \text{TC}_c(\theta, X_i),$$

where  $\text{TC}_c(\theta, X_i)$  denotes the cost of opening and rebalancing the hedge.

*Turnover.* Additionally, turnover measures how aggressively the learned strategy rebalances along the test paths. Let

$$\text{Turnover} := \frac{1}{MT_r} \sum_{i=1}^M \sum_{t=0}^{T-1} \left\| \delta_t^\theta(X_{i,0:t}) - \delta_{t-1}^\theta(X_{i,0:t-1}) \right\|_1.$$

With  $\delta_{-1}^\theta = 0$ , turnover is the mean absolute position change per path, trading date, and tradable asset, including the opening trade.

*Aggregation across random seeds.* Each metric is first computed on the test paths of one seed. For a scalar metric having seed-level values  $R_1, \dots, R_B$ , we report

$$\frac{1}{B} \sum_{b=1}^B R_b \pm \left( \frac{1}{B} \sum_{b=1}^B \left( R_b - \frac{1}{B} \sum_{c=1}^B R_c \right)^2 \right)^{1/2}.$$

## D.6 Experiments

This section supplements Section 5 with the complete comparisons and budget-sensitivity analysis.<sup>1</sup> We first present the stationary and nonstationary Heston experiments (Appendix D.6.1), including their extensions with transaction costs and M-CVaR, followed by the GAD experiments (Appendix D.6.2). Appendix D.6.3 compares one-step and multistep WRAP.

We use the four training specifications and the uniform and drift-aware weighting schemes introduced in Section 5. Throughout, the probability-reweighting ambiguity set uses the Kullback–Leibler entropy function  $\phi(x) = x \log x - x + 1$ , which satisfies the local regularity conditions of Theorems 3.6 and 3.7. The architecture, optimization settings, and ground-metric parameters are reported in Table 2. Following Goodfellow et al. (2015), we train with a mixture of clean and adversarial losses:

$$\mathcal{J}_{\text{mix}}^{\mathcal{I}}(\vartheta) = (1 - \lambda_{\text{adv}}) \sum_{n \in \mathcal{I}} \bar{w}_n \ell_{\text{DH}}(\vartheta; X_n) + \lambda_{\text{adv}} \mathcal{J}_{\text{adv}}^{\mathcal{I}}(\vartheta). \quad (44)$$

Here  $\lambda_{\text{adv}} \in [0, 1]$  balances the two losses after the Empirical Risk Minimization (ERM) warmup reported in Table 2. Hyperparameters, ambiguity budgets, and model selection are determined using validation performance only.

---

<sup>1</sup>Code and experiment configurations are available at <https://anonymous.4open.science/r/deep-hedging-nonstationary-dro/>.

*Remark D.1.* In Section 3, we construct a first-order adversary using linear tilting of the form  $a_n = 1 + \alpha_\tau(L_n - \bar{L}_w)$ . While this construction is first-order accurate as  $\tau \downarrow 0$ , at larger budgets the nonnegativity constraint can prevent further increases in  $\alpha_\tau$ , leaving part of the divergence budget unused. In the numerical experiments, we therefore replace linear tilting with exponential tilting,  $a_n = \exp(sL_n) / \sum_j w_j \exp(sL_j)$ , where  $s \geq 0$  is chosen by scalar bisection. The two constructions have the same small-budget leading behavior, preserving the first-order expansion while permitting the divergence budget to remain active beyond the range allowed by linear tilting.

Table 2: Hyperparameter settings for the Heston and GAD experiments.

| Parameter                                | Heston                                         | GAD             |
|------------------------------------------|------------------------------------------------|-----------------|
| <i>Architecture and Optimization</i>     |                                                |                 |
| Base architecture                        | Per-date MLP hedger                            |                 |
| Hidden units                             | 20                                             |                 |
| Hidden layers per date                   | 2                                              |                 |
| Activation                               | ReLU                                           |                 |
| Normalization                            | Batch normalization                            |                 |
| Optimizer                                | Adam                                           |                 |
| Learning rate                            | $5 \times 10^{-3}$                             |                 |
| Scheduler                                | Cosine annealing, $\eta_{\min} = 10^{-6}$      |                 |
| Batch size                               | 10,000                                         |                 |
| <i>Adversarial Training and Geometry</i> |                                                |                 |
| ERM warmup epochs                        | 300                                            | 100             |
| Total epochs                             | 600                                            | 300             |
| Attack algorithms                        | WRAP; Multistep WRAP                           |                 |
| Clean/adversarial loss weights           | 1 : 1                                          |                 |
| Ground metric exponent                   | $p = 2$                                        |                 |
| Ground metric scales                     | $(\lambda_S, \lambda_V) = (1, 100)$            | $\lambda_S = 1$ |
| Ground metric time weights               | $\kappa_t = 1$ for all $t \in \{0, \dots, T\}$ |                 |

### D.6.1 Heston Experiments

We first consider the Heston model in (42), which permits controlled comparisons between stationary and nonstationary market environments. The claim is the European call defined in Appendix D.2. The Heston attacks perturb both components of the observed state, corresponding to the stronger SV-Attack in He et al. (2025).

For the stationary and nonstationary benchmarks, CVaR as defined in (39) serves as both the training objective and the model-selection criterion. We also report mean PnL and turnover, as defined in Appendix D.5. CVaR is computed on the hedging loss  $-\text{PnL}$ , so lower values indicate better tail-risk performance. Mean PnL is the average realized profit and loss, while lower turnover indicates less active rebalancing.

**Stationary Heston experiments.** In the stationary benchmark, the training, validation, and test trajectories are generated independently under the same fixed Heston dynamics (42). We split the data equally into training, validation, and test sets. This setting isolates the effects of adversarial probability reweighting and trajectory perturbations in the absence of temporal distribution shift. The stationary benchmark uses uniform reference weights.

We train all adversarial specifications with the one-step WRAP procedure in Algorithm 1. For each method, training epochs and ambiguity budgets are selected by minimizing validation CVaR. The full budget grid is  $\varepsilon, \tau \in \{0, 0.05, 0.10, 0.15, 0.20, 0.25, 0.30\}$ .

*Results.* Table 3 shows that transport drives the improvement in the stationary benchmark. Validation selects  $(\varepsilon, \tau) = (0.10, 0)$ , so WRAP-U coincides with OT-U and reduces test CVaR by 2.4% relative to ERM-U. Reweighting alone provides no improvement at the displayed precision. Mean PnL is nearly unchanged across methods, while turnover remains lowest under ERM-U. Thus, without temporal distribution shift, validation favors robustness to path perturbations but provides no evidence that additional adversarial reweighting improves performance.

Table 3: Results for the stationary Heston model with the CVaR objective.

| Method    | $\varepsilon$ | $\tau$ | Val. CVaR ( $\downarrow$ ) | Test CVaR ( $\downarrow$ ) | Mean PnL ( $\uparrow$ ) | Turnover ( $\downarrow$ ) |
|-----------|---------------|--------|----------------------------|----------------------------|-------------------------|---------------------------|
| ERM-U     | –             | –      | $3.634 \pm 0.013$          | $3.629 \pm 0.018$          | $-2.199 \pm 0.005$      | $3.084 \pm 0.124$         |
| $\phi$ -U | –             | 0.25   | $3.627 \pm 0.014$          | $3.629 \pm 0.025$          | $-2.198 \pm 0.004$      | $3.283 \pm 0.266$         |
| OT-U      | 0.10          | –      | $3.550 \pm 0.011$          | $3.543 \pm 0.016$          | $-2.199 \pm 0.004$      | $3.156 \pm 0.064$         |
| WRAP-U    | 0.10          | –      | $3.550 \pm 0.011$          | $3.543 \pm 0.016$          | $-2.199 \pm 0.004$      | $3.156 \pm 0.064$         |

*Note.* Values are mean  $\pm$  population standard deviation over five seeds.

**Stationary Heston with transaction costs and M-CVaR.** We next extend the stationary Heston benchmark by introducing proportional transaction costs at rate  $c \geq 0$ , which penalize rebalancing in the training objective. Setting  $\delta_{-1}^\theta = 0$ , the cumulative cost of opening and rebalancing the hedge is

$$\text{TC}_c(\theta, X) = c \sum_{t=0}^{T-1} |S_t|^\top |\delta_t^\theta(X_{0:t}) - \delta_{t-1}^\theta(X_{0:t-1})|. \quad (45)$$

The resulting terminal PnL is

$$\text{PnL}_c(\theta, X) = p_0 + \sum_{t=0}^{T-1} \delta_t^\theta(X_{0:t})^\top (S_{t+1} - S_t) - \Phi(X) - \text{TC}_c(\theta, X). \quad (46)$$

We retain the fixed Heston dynamics and European call from the stationary benchmark. We consider both CVaR and M-CVaR objectives, defined in (39) and (41), with transaction-cost rate  $c = 0.005$ , confidence level  $\alpha = 0.95$ , and  $\omega = 0.5$  for M-CVaR. For each objective, we compare ERM-U,  $\phi$ -U, OT-U, and WRAP-U under uniform reference weights. Training epochs and ambiguity budgets are selected separately for each method using validation performance.

*Results.* Table 4 shows that WRAP-U selects  $(\varepsilon, \tau) = (0.05, 0.05)$  and attains the lowest validation and test CVaR. Relative to ERM-U, test CVaR decreases from 4.494 to 4.481 (0.3%), at the cost of higher transaction costs and lower mean PnL. OT-U attains the highest mean PnL and the lowest transaction costs. The additional reduction from joint training is marginal, with test CVaR of 4.481 under WRAP-U versus 4.484 under OT-U.

Figure 3 clarifies the role of the two ambiguity channels. Without transaction costs, transport accounts for nearly all of the improvement, and validation selects  $(\varepsilon, \tau) = (0.10, 0)$ . With transaction costs, gains are smaller and concentrated at low budgets, while larger transport budgets deteriorate validation performance. This is consistent with the rebalancing penalty in (45) already regularizing the hedge, as also suggested by the smoother ERM-U policy in Figure 6. In contrast, reweighting contributes more substantially under nonstationarity (Figure 4). Thus, the two-budget formulation

Table 4: Results for the stationary Heston model with transaction costs and the CVaR objective.

| Method    | $\varepsilon$ | $\tau$ | Val. CVaR ( $\downarrow$ )          | Test CVaR ( $\downarrow$ )          | Mean PnL( $\uparrow$ ) | Total TC( $\downarrow$ )            |
|-----------|---------------|--------|-------------------------------------|-------------------------------------|------------------------|-------------------------------------|
| ERM-U     | —             | —      | $4.500 \pm 0.011$                   | $4.494 \pm 0.016$                   | $-3.003 \pm 0.006$     | $0.803 \pm 0.004$                   |
| $\phi$ -U | —             | 0.05   | $4.491 \pm 0.009$                   | $4.485 \pm 0.015$                   | $-3.019 \pm 0.004$     | $0.819 \pm 0.004$                   |
| OT-U      | 0.05          | —      | $4.488 \pm 0.010$                   | $4.484 \pm 0.013$                   | $-3.002 \pm 0.010$     | <u><math>0.802 \pm 0.008</math></u> |
| WRAP-U    | 0.05          | 0.05   | <u><math>4.487 \pm 0.009</math></u> | <u><math>4.481 \pm 0.012</math></u> | $-3.014 \pm 0.006$     | $0.815 \pm 0.007$                   |

*Note.* Values are mean  $\pm$  population standard deviation over five seeds.

adapts across settings: validation selects transport alone when reweighting adds no value and positive budgets for both channels when their combination improves the objective.




Figure 3: Stationary Heston validation grids with uniform reference weights and the CVaR objective, corresponding to Tables 3 and 4. The surface reports the percentage reduction in validation CVaR relative to the corresponding ERM-U baseline. Crosses indicate the evaluated budget pairs, with contours interpolating between them; the selected WRAP budgets are  $(\varepsilon, \tau) = (0.10, 0)$  in (a) and  $(0.05, 0.05)$  in (b).

Under M-CVaR (Table 5), validation again selects transport alone, with WRAP-U coinciding with OT-U at  $(\varepsilon, \tau) = (0.05, 0)$ . Both attain the lowest validation and test M-CVaR, reducing the test objective from 3.714 to 3.706, while also achieving the highest mean PnL and lowest transaction costs. Reweighting alone attains the lowest test CVaR, but its lower mean PnL yields a higher M-CVaR objective. Hence, the preferred adversarial mechanism depends on the optimization objective, not tail risk alone.

Table 5: Results for the stationary Heston model with transaction costs and the M-CVaR objective.

| Method    | $\varepsilon$ | $\tau$ | Val. M-CVaR ( $\downarrow$ )        | Test M-CVaR ( $\downarrow$ )        | Test CVaR ( $\downarrow$ )          | Mean PnL( $\uparrow$ )               | Total TC( $\downarrow$ )            |
|-----------|---------------|--------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------------------|-------------------------------------|
| ERM-U     | —             | —      | $3.717 \pm 0.010$                   | $3.714 \pm 0.008$                   | $4.500 \pm 0.014$                   | $-2.927 \pm 0.004$                   | $0.727 \pm 0.003$                   |
| $\phi$ -U | —             | 0.05   | $3.716 \pm 0.007$                   | $3.713 \pm 0.009$                   | <u><math>4.486 \pm 0.014</math></u> | $-2.940 \pm 0.004$                   | $0.740 \pm 0.002$                   |
| OT-U      | 0.05          | —      | $3.708 \pm 0.009$                   | <u><math>3.706 \pm 0.007</math></u> | $4.502 \pm 0.013$                   | <u><math>-2.909 \pm 0.004</math></u> | <u><math>0.709 \pm 0.003</math></u> |
| WRAP-U    | 0.05          | —      | <u><math>3.708 \pm 0.009</math></u> | <u><math>3.706 \pm 0.007</math></u> | $4.502 \pm 0.013$                   | <u><math>-2.909 \pm 0.004</math></u> | <u><math>0.709 \pm 0.003</math></u> |

*Note.* Values are mean  $\pm$  population standard deviation over five seeds.

**Nonstationary Heston experiments.** In the nonstationary experiments, the stock-price drift and long-run variance evolve chronologically, so the validation, near-future, and far-future test distributions depart progressively from the training regime. This setting allows us to examine both the drift-aware reference weights from Section 2.2 and the two adversarial ambiguity channels under temporal distribution shift.

We generate  $B = 8$  chronological blocks of 4,000 independent trajectories each. For block  $b = 1, \dots, B$ , the drift and long-run variance are

$$\mu_b = \mu_0 + \Delta\mu \frac{b-1}{B-1}, \quad \theta_{H,b} = \theta_{H,0} + \Delta\theta_H \frac{b-1}{B-1},$$

with  $\mu_0 = 0$ ,  $\Delta\mu = 0.20$ ,  $\theta_{H,0} = 0.04$ , and  $\Delta\theta_H = 0.12$ . Each block starts at  $v_{0,b} = \theta_{H,b}$ , so both initial and long-run variance increase from 0.04 to 0.16, while the drift increases from 0 to 0.20. All remaining Heston parameters retain their baseline values and are fixed across blocks.

The blocks are split chronologically into training (Blocks 1, ..., 5), validation (Block 6), near-future test (Block 7), and far-future test (Block 8). Weight selection precedes adversarial-budget selection. For the drift-aware methods, we train clean hedges over a grid of eleven candidate ratios  $\varrho/\varepsilon_\varrho$  and select the ratio minimizing mean validation CVaR over five seeds. The selected ratio 0.00853 is then held fixed across all weighted adversarial specifications. Here  $\varepsilon_\varrho$  denotes the reference-weighting radius from Section 2.2, distinct from the adversarial transport budget  $\varepsilon$ . The resulting weights are assigned at the block level and distributed uniformly across paths within each block, while  $\varrho/\varepsilon_\varrho = 0$  recovers uniform reference weights.

The nonstationary one-step budget grid is

$$\varepsilon \in \{0, 0.1, \dots, 0.5\}, \quad \tau \in \{0, 0.2, \dots, 1.2\}. \quad (47)$$

*Results.* Table 1 shows that temporal distribution shift materially degrades the clean hedge: the CVaR of ERM-U increases from 6.868 on validation to 7.952 and 8.992 on the near- and far-future samples. In contrast to the stationary benchmark, both ambiguity channels now improve performance under uniform reference weights, and their combination yields the largest reduction. WRAP-U selects  $(\varepsilon, \tau) = (0.30, 0.60)$  and lowers far-future CVaR by 19.2% relative to ERM-U. This contrast with the stationary setting is central: adversarial reweighting becomes useful once the training distribution no longer represents future market conditions, complementing the protection provided by transport.

Drift-aware weighting provides an additional layer of robustness by improving the reference distribution before adversarial perturbations are applied. Across all four training methods, drift-aware weights lower CVaR in every evaluation period, with the largest relative far-future improvement under clean training. WRAP-W selects  $(\varepsilon, \tau) = (0.40, 1.00)$  and attains the lowest CVaR in validation and both test periods. The additional gain over WRAP-U is smaller, consistent with drift-aware weighting and adversarial training addressing partly overlapping sources of temporal misspecification. The improvement is concentrated in tail risk rather than average performance: ERM-U retains the highest mean PnL, while the robust methods achieve substantially lower CVaR.

Figure 4 provides further evidence for the complementarity of the two ambiguity channels. In both panels, the largest validation gains occur in the interior of the budget grid, where both  $\varepsilon$  and  $\tau$  are positive. The selected joint budgets reduce validation CVaR by 12.6% under uniform weights and 10.5% under drift-aware weights, compared with at most 9.1% and 6.8%, respectively, along the single-budget axes. Moreover, sufficiently large transport budgets can worsen performance when applied alone, whereas combining transport with reweighting produces a broader region of improved validation CVaR. The benefit of the two-budget formulation therefore lies not only in

attaining a better selected solution, but also in providing a more robust region of effective adversarial regularization.

The smaller percentage reduction under drift-aware weighting reflects the stronger starting point: weighting alone lowers validation CVaR from 6.868 to 6.653. Panel (b) therefore measures adversarial gains relative to an already improved reference distribution. Despite the smaller relative reduction, WRAP-W attains lower validation CVaR than WRAP-U (5.954 versus 6.005). This pattern shows that drift-aware reference construction and adversarial robustification are complementary rather than substitutes: the former reduces temporal misspecification in the empirical reference, while the latter protects against residual probability and path uncertainty.




Figure 4: Nonstationary Heston validation grids for Table 1. Colours show percentage reductions Val. CVaR relative to ERM-U in (a) and ERM-W in (b); Black crosses mark the budget pairs in (47); contours interpolate between them. Red crosses mark the selected WRAP budgets:  $(\varepsilon, \tau) = (0.30, 0.60)$  in (a) and  $(0.40, 1.00)$  in (b)

**Nonstationary Heston with transaction costs and M-CVaR.** We extend the nonstationary Heston experiment to proportional transaction costs using the European call in Appendix D.2 and the chronological block split described above. Costs follow (45) at rate  $c = 0.005$  and enter the terminal PnL through (46). Reported mean PnL is net of these costs, and “Total TC” denotes the mean cumulative transaction cost per path, including opening trades and excluding terminal liquidation.

We use uniform reference weights and compare ERM-U,  $\phi$ -U, OT-U, and WRAP-U with the one-step attack in Algorithm 1. The first experiment uses the CVaR objective in (39) with  $\alpha = 0.95$ , while the second uses the M-CVaR objective in (41) with the same confidence level and  $\omega = 1/2$ . Epochs and ambiguity budgets are selected on the validation block and subsequently evaluated on the near- and far-future test blocks.

*Results.* Table 6 shows that transaction costs alter the relative contribution of the two ambiguity channels. Reweighting alone substantially improves CVaR across all three periods, whereas transport alone worsens the objective despite producing the lowest transaction costs. Joint training combines the tail-risk benefit of reweighting with a smaller transport component and attains the lowest validation, near-future, and far-future CVaR. In the far-future sample, WRAP-U reduces CVaR by 6.6% relative to ERM-U. The improvement comes with lower mean PnL and higher trading costs

than clean training, but WRAP-U improves both quantities relative to reweighting alone.

Table 6: Results for the nonstationary Heston model with transaction costs and the CVaR objective.

| Method    | $\varepsilon$ | $\tau$ | Val. CVaR( $\downarrow$ )           | Near CVaR( $\downarrow$ )           | Far CVaR( $\downarrow$ )            | Far Mean PnL( $\uparrow$ ) | Far Total TC( $\downarrow$ ) |
|-----------|---------------|--------|-------------------------------------|-------------------------------------|-------------------------------------|----------------------------|------------------------------|
| ERM-U     | —             | —      | $7.592 \pm 0.077$                   | $8.250 \pm 0.155$                   | $8.935 \pm 0.260$                   | $-5.562 \pm 0.078$         | $1.002 \pm 0.073$            |
| $\phi$ -U | —             | 1.20   | $7.122 \pm 0.042$                   | $7.778 \pm 0.131$                   | $8.434 \pm 0.225$                   | $-5.714 \pm 0.082$         | $1.163 \pm 0.076$            |
| OT-U      | 0.10          | —      | $7.752 \pm 0.073$                   | $8.456 \pm 0.134$                   | $9.170 \pm 0.195$                   | $-5.565 \pm 0.101$         | $0.976 \pm 0.095$            |
| WRAP-U    | 0.10          | 0.80   | <u><math>7.082 \pm 0.041</math></u> | <u><math>7.691 \pm 0.111</math></u> | <u><math>8.349 \pm 0.189</math></u> | $-5.695 \pm 0.071$         | $1.130 \pm 0.065$            |

*Note.* Values are mean  $\pm$  population standard deviation over five seeds.

Under M-CVaR, the same qualitative pattern persists (Table 7). Reweighting remains beneficial, transport alone again worsens the risk objective, and WRAP-U attains the lowest validation, near-future, and far-future M-CVaR. The joint method also improves mean PnL and transaction costs relative to reweighting alone, yielding a better balance between average loss and tail risk. Most importantly, the nonstationary experiment selects positive budgets for both perturbation channels under both objectives. This contrasts with the stationary M-CVaR benchmark, where validation selects transport alone, and reinforces the role of reweighting when the market distribution evolves over time.

Table 7: Results for the nonstationary Heston model with transaction costs and the M-CVaR objective.

| Method    | $\varepsilon$ | $\tau$ | Val. M-CVaR( $\downarrow$ )         | Near M-CVaR( $\downarrow$ )         | Far M-CVaR( $\downarrow$ )          | Far Mean PnL( $\uparrow$ ) | Far Total TC( $\downarrow$ )        |
|-----------|---------------|--------|-------------------------------------|-------------------------------------|-------------------------------------|----------------------------|-------------------------------------|
| ERM-U     | —             | —      | $6.281 \pm 0.048$                   | $6.766 \pm 0.077$                   | $7.251 \pm 0.149$                   | $-5.552 \pm 0.084$         | $0.994 \pm 0.077$                   |
| $\phi$ -U | —             | 1.00   | $6.078 \pm 0.011$                   | $6.581 \pm 0.068$                   | $7.090 \pm 0.149$                   | $-5.721 \pm 0.089$         | $1.170 \pm 0.079$                   |
| OT-U      | 0.10          | —      | $6.388 \pm 0.074$                   | $6.898 \pm 0.078$                   | $7.401 \pm 0.151$                   | $-5.562 \pm 0.122$         | <u><math>0.974 \pm 0.112</math></u> |
| WRAP-U    | 0.20          | 0.80   | <u><math>6.045 \pm 0.013</math></u> | <u><math>6.523 \pm 0.062</math></u> | <u><math>7.019 \pm 0.148</math></u> | $-5.681 \pm 0.098$         | $1.112 \pm 0.092$                   |

*Note.* Values are mean  $\pm$  population standard deviation over five seeds.

**Hedging surfaces.** Figure 5 shows the stock position  $\delta_t^\theta$  as a function of moneyness and trading date. Across methods, the learned hedge has the expected structure: the stock position is close to zero when the call is far out of the money, approaches one when the option is well in the money, and transitions more sharply around the strike as maturity approaches. The main difference across training schemes is therefore not the overall shape of the hedge, but its local regularity. In particular, OT-U produces a smoother policy than ERM-U, with fewer local fluctuations across nearby moneyness levels and trading dates. This coincides with the regularization effect of Wasserstein described by Gao et al. (2022).

Transaction costs reinforce this regularization effect. Relative to Figure 5, the policies in Figure 6 vary more gradually across states and time, including under ERM-U. The transaction-cost term in (45) directly penalizes changes in position, whereas OT regularizes the hedge indirectly through sensitivity of the loss to path perturbations, as characterized by Theorem 3.6. The smoother OT-U and WRAP-U surfaces therefore reflect a distinct robustness mechanism that remains visible even when rebalancing is already penalized.

Under nonstationarity, Figures 7 and 8 show the corresponding hedges under uniform and drift-aware reference weights. The later Heston blocks generate a wider range of moneyness states, but the qualitative structure of the hedge remains stable. The main effect of robust training is again local:  $\phi$ -W and OT-W reduce irregularities around the strike relative to ERM-W, while WRAP-W

produces the smoothest transition among the weighted specifications. The same pattern appears under uniform weights, where WRAP-U also reduces local fluctuations relative to ERM-U. These surfaces suggest that the robustness gains documented above are accompanied by more regular hedging policies rather than a qualitatively different trading rule.




Figure 5: Stock-hedging policies for stationary Heston with uniform reference weights and the CVaR objective, using the selections in Table 3: (A) ERM-U, (B)  $\phi$ -U, (C) OT-U, and (D) WRAP-U. Colours show stock holdings averaged over validation paths and five seeds within each date–moneyness bin. Dashed curves mark the 5th and 95th moneyness percentiles; the solid line marks the strike.




Figure 6: Stock-hedging policies for stationary Heston with uniform reference weights, the CVaR objective, and proportional transaction costs  $c = 0.005$ , using the selections in Table 4. Panels, averaging, and plotting conventions follow Figure 5, with the same colour scale.




Figure 7: Stock-hedging policies for nonstationary Heston with uniform reference weights and the CVaR objective, using the selections in Table 1: (A) ERM-U, (B)  $\phi$ -U, (C) OT-U, and (D) WRAP-U. Averaging and plotting conventions follow Figure 5, with the same colour scale.




Figure 8: Stock-hedging policies for nonstationary Heston with drift-aware reference weights and the CVaR objective, using the selections in Table 1: (A) ERM-W, (B)  $\phi$ -W, (C) OT-W, and (D) WRAP-W. The displayed ratio  $r = \varrho/\varepsilon_\varrho \approx 0.00853$  controls reference weighting. Averaging and plotting conventions follow Figure 5, with the same colour scale.

### D.6.2 GAD Experiments

We complement the controlled Heston experiments with generalized affine diffusion (GAD) models calibrated to historical equity prices, using the dynamics introduced in Appendix D.4. Following the market-data benchmark of He et al. (2025, Section 5.4) and the parameter-uncertainty approach of Lütkebohmert et al. (2022), we compare training under a fixed calibration with training that samples parameters from rolling estimates. Evaluation on simulated paths from models fitted to later, held-out periods assesses whether adversarial probability reweighting and path perturbations improve hedging performance under changes in market conditions estimated from real prices.

**Data and descriptive statistics.** We use daily auto-adjusted closing prices from Yahoo Finance for AAPL, MSFT, AMZN, GOOGL, and BRK-B. Training calibration uses prices from 26 September 2008 through 31 December 2019, while validation and test models are calibrated separately to 2020–2023 and 2024–2025 prices. Each period is divided into overlapping windows of 250 prices, shifted by 100 observations, with no window crossing a period boundary. Each window is rescaled to start at  $S_0 = 10$ . Calibration settings are selected using the first 200 prices for fitting and the following 50 returns for validation, after which the coefficients are refitted on the full window. Test-period calibration is performed only after epoch and budget selection has been frozen. The resulting experiment is therefore a retrospective simulation benchmark conditional on period-specific fitted market models.

Table 8: Descriptive statistics of daily simple adjusted-close returns across the training, validation, and test periods. Mean, median, volatility, and maximum drawdown are reported in percentages. The  $p_{KS}$  column reports two-sample Kolmogorov–Smirnov  $p$ -values relative to the corresponding training-period returns.

| Stock | Split | $\mu$ (%) | Median (%) | $\sigma$ (%) | MaxDD (%) | AC(1)  | $p_{KS}$ |
|-------|-------|-----------|------------|--------------|-----------|--------|----------|
| AAPL  | Train | 30.24     | 25.39      | 29.33        | −43.80    | −0.000 |          |
|       | Val   | 29.90     | 24.90      | 33.55        | −31.43    | −0.133 | 0.016    |
|       | Test  | 23.51     | 29.92      | 27.87        | −33.36    | 0.045  | 0.046    |
| MSFT  | Train | 21.56     | 16.11      | 26.86        | −43.95    | −0.084 |          |
|       | Val   | 27.58     | 24.65      | 32.60        | −37.15    | −0.197 | < 0.001  |
|       | Test  | 16.54     | 28.00      | 22.16        | −23.73    | −0.018 | 0.682    |
| AMZN  | Train | 35.49     | 23.10      | 36.18        | −51.86    | −0.018 |          |
|       | Val   | 18.90     | 19.69      | 37.67        | −56.15    | −0.039 | 0.007    |
|       | Test  | 26.62     | 16.34      | 31.37        | −30.88    | −0.030 | 0.918    |
| GOOGL | Train | 19.92     | 16.17      | 27.21        | −40.27    | −0.010 |          |
|       | Val   | 23.52     | 26.61      | 33.52        | −44.32    | −0.108 | < 0.001  |
|       | Test  | 46.09     | 73.66      | 30.23        | −29.81    | −0.014 | < 0.001  |
| BRK-B | Train | 10.84     | 5.87       | 22.36        | −50.54    | −0.032 |          |
|       | Val   | 13.87     | 15.01      | 23.19        | −29.57    | −0.179 | 0.201    |
|       | Test  | 17.85     | 21.29      | 16.71        | −14.95    | 0.016  | 0.256    |

*Note.* Daily mean and median returns are annualized by a factor of 252, and volatility by  $\sqrt{252}$ . MaxDD reports the maximum decline from a previous peak, and AC(1) the lag-one autocorrelation. KS  $p$ -values are nominal and are not adjusted for serial dependence or multiple testing.

Table 8 reports annualized mean and median returns, annualized volatility, maximum drawdown, lag-one autocorrelation, and two-sample Kolmogorov–Smirnov  $p$ -values relative to the training period. The three calibration periods exhibit different market characteristics. In particular, test-period volatility is lower than validation-period volatility for all five stocks, with the largest declines for MSFT, AMZN, and BRK-B. Maximum drawdowns are also generally smaller in the test period. Hence, the test period represents a different market regime, but not simply a more severe version of the validation period.

The KS statistics provide additional evidence that distributional changes differ across stocks. AAPL and GOOGL exhibit small test-period  $p$ -values relative to the training sample, while the corresponding values for MSFT, AMZN, and BRK-B are larger. The market-data experiment therefore covers heterogeneous changes in return distributions rather than a single common stress scenario. This provides a useful setting for testing whether adversarial training remains beneficial across different

forms of market change.

**Calibration.** The full five-parameter GAD model in (43) can be weakly identified on short calibration windows: nearly flat likelihood directions allow substantially different coefficients to fit the data similarly. To reduce this instability and obtain a positive-state model, we set  $a_0 = 0$  and impose  $a_1, b_0 > 0$  and  $\gamma \in [1/2, 1]$ , with  $2b_0 > a_1$  when  $\gamma = 1/2$ . We fit this restricted model using regularized Gaussian quasi-likelihood for log returns, pooling the volatility exponent across stocks and windows within each period and regularizing the drift. For each candidate exponent and drift penalty, we fit the coefficients on the first 200 prices of each window and evaluate the negative log-likelihood on the following 50 returns. We select the pair with the lowest average validation negative log-likelihood across all five stocks and windows in that period. The exponent grid is  $\gamma \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$  and the drift-penalty grid is  $\{0, 0.01, 0.1, 1, 10, 100\}$ . With the selected pair fixed, we refit  $(b_0, b_1, a_1)$  separately for each stock and window using all 250 prices.

We repeat this procedure independently for the three historical datasets: 2008–2019 prices determine the training parameters, 2020–2023 prices determine the validation parameters, and 2024–2025 prices determine the test parameters. Each period has a separately selected exponent and drift penalty and a corresponding collection of fitted coefficient vectors. These collections generate the respective simulated samples; validation and test parameters do not enter the training generator. Test calibration takes place only after model and adversarial-budget selection is complete.

**Simulation: FIX and ROBUST.** Following Buehler et al. (2019); He et al. (2025), we train on simulated trajectories. The FIX specification uses the coefficient vector from the final training-period calibration window throughout every training path. The ROBUST specification draws one complete coefficient vector uniformly from the stock’s rolling estimates, independently for each path and trading step. This differs from the implementation of He et al. (2025). Sampling the coefficients jointly preserves their dependence within each fitted vector and follows the empirical-parameter alternative discussed by Lütkebohmert et al. (2022, Remark 4.2). The pooled exponent remains fixed within each period. Both specifications use the log-Euler scheme, consistent with the calibration, and produce positive stock prices.

For each stock, each model replicate is trained on 20,000 simulated paths, which are shared across methods and budget candidates. Validation and testing each use a separate sample of 20,000 independent ROBUST paths per stock, generated from the corresponding period’s calibration and shared across methods, training specifications, and model seeds. At each trading step, the evaluation generators sample coefficient vectors only from the corresponding period’s rolling estimates, following the same sampling rule as ROBUST training. The large evaluation samples reduce Monte Carlo variance, since a small number of extreme losses can otherwise dominate the entropic-risk estimate and make comparisons unstable. This improves precision conditional on the fitted models, while calibration uncertainty remains.

**Hedging setup and budget selection.** We hedge the arithmetic Asian call defined in Appendix D.2, with  $S_0 = K = 10$ ,  $T = 30$  trading steps, and  $\Delta t = 1/250$ , using the entropic OCE in Appendix D.1. The stock is the sole tradable asset and observed state, so adversarial transport acts only on the stock-price path. Under both FIX and ROBUST training, we compare ERM-U,  $\phi$ -U, OT-U, and WRAP-U using uniform reference weights and the one-step attack in Algorithm 1. Architecture and optimization settings follow Table 2, with 300 training epochs and a 100-epoch clean warmup for adversarial methods.

Training and evaluation use the entropic OCE in (40) with  $\lambda = 13$  and PnL floored at  $-1$ ; reported mean PnL is computed without this floor. Clean and adversarial losses receive equal weight after warmup. For each candidate, we retain the epoch with the lowest validation entropic risk and select ambiguity budgets separately for each stock, training distribution, and adversarial specification using mean validation performance over five model seeds. We consider

$$\tau \in \{0, 0.05, 0.10, 0.15, 0.20, 0.30, 0.40\}, \quad \varepsilon \in \{0, 0.025, 0.05, 0.10, 0.15, 0.20\}.$$

The selected models are then evaluated on the shared held-out ROBUST test paths. Reported standard deviations therefore reflect training variability across model seeds rather than market-sampling or calibration uncertainty.

**Results.** Table 9 shows that joint adversarial training is consistently stronger than either ambiguity channel in isolation. WRAP-U attains lower validation and test entropic risk than both  $\phi$ -U and OT-U in all ten stock-training comparisons. Relative to ERM-U, WRAP-U lowers test entropic risk in six of ten comparisons. In three additional cases, validation selects the clean origin, so WRAP-U recovers ERM-U. Only FIX BRK-B selects positive adversarial budgets on validation without improving test risk. The results therefore suggest that joint adversarial training can add robustness beyond parameter uncertainty in the training generator while reverting to empirical training when additional robustness is not supported by validation.

The gains vary substantially across stocks and market periods. The largest improvements under ROBUST occur for AAPL and GOOGL, whereas several comparisons show little or no benefit over ERM-U. This heterogeneity is consistent with the descriptive statistics in Table 8: validation and test periods differ in both the magnitude and direction of distribution shift across stocks. In particular, lower test-period volatility for MSFT, AMZN, and BRK-B reduces entropic risk for all methods, so lower test scores alone do not indicate a better hedge. Performance should instead be assessed relative to ERM-U on the same test sample.

The remaining performance measures reveal a clear trade-off. WRAP-U reduces turnover relative to ERM-U in all ten comparisons and attains the lowest turnover in nine. In the six cases with lower test entropic risk, mean PnL is also lower. The improvements therefore reflect better control of risk rather than uniformly better average performance, accompanied by less active rebalancing.

### D.6.3 One-step and Multistep WRAP

We compare one-step and multistep WRAP (Algorithms 1 and 2) using the stationary Heston setup above and the Heston training settings in Table 2. Budgets are selected separately for each attack and specification according to validation CVaR.

*Results.* Table 10 shows that both attacks select the same budgets within each specification. Both OT-U and WRAP-U select  $(\varepsilon, \tau) = (0.10, 0)$ , so the selected WRAP configuration uses transport alone for both attacks. One-step training achieves slightly lower test CVaR than multistep training (3.543 versus 3.545), with a difference smaller than the reported seed standard deviations. The additional inner iterations therefore provide no observed improvement at the selected budgets, supporting the use of the cheaper one-step attack in this setting.

Table 9: Hedging under GAD dynamics calibrated to successive periods of real equity prices.

| Training | Stock | Method    | $\varepsilon$ | $\tau$ | Val. OCE ( $\downarrow$ ) | Test OCE ( $\downarrow$ ) | Mean PnL ( $\uparrow$ ) | Turnover ( $\downarrow$ ) |
|----------|-------|-----------|---------------|--------|---------------------------|---------------------------|-------------------------|---------------------------|
| Fix      | AAPL  | ERM-U     | –             | –      | $0.287 \pm 0.002$         | $0.258 \pm 0.001$         | $-0.227 \pm 0.000$      | $0.072 \pm 0.002$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.308 \pm 0.012$         | $0.272 \pm 0.008$         | $-0.227 \pm 0.002$      | $0.089 \pm 0.005$         |
|          |       | OT-U      | 0.025         | –      | $0.299 \pm 0.003$         | $0.267 \pm 0.002$         | $-0.226 \pm 0.001$      | $0.072 \pm 0.005$         |
|          |       | WRAP-U    | 0.05          | 0.15   | $0.289 \pm 0.002$         | $0.261 \pm 0.001$         | $-0.226 \pm 0.001$      | $0.066 \pm 0.001$         |
| Fix      | AMZN  | ERM-U     | –             | –      | $0.310 \pm 0.001$         | $0.289 \pm 0.001$         | $-0.255 \pm 0.000$      | $0.069 \pm 0.001$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.328 \pm 0.004$         | $0.301 \pm 0.003$         | $-0.252 \pm 0.002$      | $0.083 \pm 0.003$         |
|          |       | OT-U      | 0.025         | –      | $0.322 \pm 0.002$         | $0.300 \pm 0.002$         | $-0.254 \pm 0.001$      | $0.068 \pm 0.005$         |
|          |       | WRAP-U    | 0.05          | 0.30   | $0.314 \pm 0.002$         | $0.294 \pm 0.002$         | $-0.255 \pm 0.001$      | $0.067 \pm 0.001$         |
| Fix      | BRK-B | ERM-U     | –             | –      | $0.204 \pm 0.005$         | $0.170 \pm 0.003$         | $-0.142 \pm 0.004$      | $0.100 \pm 0.004$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.263 \pm 0.044$         | $0.203 \pm 0.025$         | $-0.131 \pm 0.013$      | $0.148 \pm 0.037$         |
|          |       | OT-U      | 0.025         | –      | $0.219 \pm 0.011$         | $0.186 \pm 0.007$         | $-0.151 \pm 0.004$      | $0.086 \pm 0.019$         |
|          |       | WRAP-U    | 0.025         | 0.10   | $0.197 \pm 0.005$         | $0.172 \pm 0.006$         | $-0.147 \pm 0.004$      | $0.085 \pm 0.007$         |
| Fix      | GOOGL | ERM-U     | –             | –      | $0.292 \pm 0.008$         | $0.301 \pm 0.009$         | $-0.243 \pm 0.003$      | $0.086 \pm 0.006$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.327 \pm 0.015$         | $0.340 \pm 0.016$         | $-0.238 \pm 0.008$      | $0.114 \pm 0.010$         |
|          |       | OT-U      | 0.025         | –      | $0.303 \pm 0.008$         | $0.303 \pm 0.008$         | $-0.248 \pm 0.002$      | $0.074 \pm 0.009$         |
|          |       | WRAP-U    | 0.05          | 0.10   | $0.286 \pm 0.002$         | $0.292 \pm 0.003$         | $-0.248 \pm 0.002$      | $0.071 \pm 0.004$         |
| Fix      | MSFT  | ERM-U     | –             | –      | $0.318 \pm 0.019$         | $0.213 \pm 0.004$         | $-0.175 \pm 0.001$      | $0.088 \pm 0.003$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.363 \pm 0.029$         | $0.226 \pm 0.006$         | $-0.181 \pm 0.005$      | $0.108 \pm 0.009$         |
|          |       | OT-U      | 0.025         | –      | $0.323 \pm 0.011$         | $0.217 \pm 0.004$         | $-0.174 \pm 0.001$      | $0.076 \pm 0.009$         |
|          |       | WRAP-U    | 0.025         | 0.05   | $0.280 \pm 0.002$         | $0.205 \pm 0.000$         | $-0.176 \pm 0.001$      | $0.070 \pm 0.004$         |
| Robust   | AAPL  | ERM-U     | –             | –      | $0.328 \pm 0.019$         | $0.288 \pm 0.012$         | $-0.223 \pm 0.001$      | $0.091 \pm 0.005$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.352 \pm 0.023$         | $0.296 \pm 0.008$         | $-0.227 \pm 0.002$      | $0.108 \pm 0.005$         |
|          |       | OT-U      | 0.025         | –      | $0.332 \pm 0.015$         | $0.289 \pm 0.011$         | $-0.222 \pm 0.001$      | $0.077 \pm 0.008$         |
|          |       | WRAP-U    | 0.025         | 0.10   | $0.297 \pm 0.002$         | $0.267 \pm 0.002$         | $-0.224 \pm 0.001$      | $0.071 \pm 0.002$         |
| Robust   | AMZN  | ERM-U     | –             | –      | $0.329 \pm 0.005$         | $0.296 \pm 0.003$         | $-0.243 \pm 0.001$      | $0.075 \pm 0.002$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.336 \pm 0.011$         | $0.302 \pm 0.006$         | $-0.245 \pm 0.003$      | $0.086 \pm 0.006$         |
|          |       | OT-U      | 0.025         | –      | $0.336 \pm 0.006$         | $0.303 \pm 0.003$         | $-0.248 \pm 0.001$      | $0.065 \pm 0.004$         |
|          |       | WRAP-U    | 0.05          | 0.15   | $0.321 \pm 0.001$         | $0.294 \pm 0.002$         | $-0.249 \pm 0.001$      | $0.067 \pm 0.002$         |
| Robust   | BRK-B | ERM-U     | –             | –      | $0.190 \pm 0.003$         | $0.161 \pm 0.002$         | $-0.137 \pm 0.002$      | $0.095 \pm 0.002$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.223 \pm 0.013$         | $0.183 \pm 0.008$         | $-0.139 \pm 0.007$      | $0.142 \pm 0.019$         |
|          |       | OT-U      | 0.025         | –      | $0.206 \pm 0.005$         | $0.173 \pm 0.006$         | $-0.142 \pm 0.004$      | $0.094 \pm 0.014$         |
|          |       | WRAP-U    | 0.025         | 0.05   | $0.193 \pm 0.004$         | $0.167 \pm 0.004$         | $-0.143 \pm 0.003$      | $0.087 \pm 0.009$         |
| Robust   | GOOGL | ERM-U     | –             | –      | $0.325 \pm 0.023$         | $0.339 \pm 0.030$         | $-0.233 \pm 0.006$      | $0.105 \pm 0.010$         |
|          |       | $\phi$ -U | –             | 0.05   | $0.392 \pm 0.028$         | $0.420 \pm 0.033$         | $-0.246 \pm 0.019$      | $0.137 \pm 0.013$         |
|          |       | OT-U      | 0.025         | –      | $0.332 \pm 0.024$         | $0.333 \pm 0.018$         | $-0.242 \pm 0.004$      | $0.092 \pm 0.016$         |
|          |       | WRAP-U    | 0.10          | 0.40   | $0.297 \pm 0.007$         | $0.302 \pm 0.005$         | $-0.245 \pm 0.003$      | $0.077 \pm 0.005$         |
| Robust   | MSFT  | ERM-U     | –             | –      | $0.310 \pm 0.017$         | $0.212 \pm 0.005$         | $-0.176 \pm 0.001$      | $0.089 \pm 0.005$         |
|          |       | $\phi$ -U | –             | 0.10   | $0.369 \pm 0.012$         | $0.235 \pm 0.004$         | $-0.178 \pm 0.003$      | $0.124 \pm 0.008$         |
|          |       | OT-U      | 0.025         | –      | $0.316 \pm 0.008$         | $0.213 \pm 0.003$         | $-0.176 \pm 0.001$      | $0.080 \pm 0.005$         |
|          |       | WRAP-U    | 0.05          | 0.15   | $0.285 \pm 0.004$         | $0.208 \pm 0.003$         | $-0.178 \pm 0.003$      | $0.075 \pm 0.004$         |

*Note.* Values are mean  $\pm$  sample standard deviation over five model seeds. The WRAP-U rows exclude the clean origin  $(\varepsilon, \tau) = (0, 0)$ , reported separately as ERM-U.

## D.7 Pathwise Effects of Reweighting and Optimal Transport

Figure 9, reproduced from Figure 2, illustrates how the two budgets in Section 3 affect scenario probabilities and market trajectories in the nonstationary Heston experiment with joint drift. Panels (A)–(C) show a random sample of reference trajectories, together with the highest-weight and most-displaced trajectories highlighted separately; panel (C) shows their transported versions. Stock prices are normalized by the strike price. Panel (A) shows the original paths and the pointwise 5th and 95th percentiles under uniform reference probabilities  $w_n = 1/N$ . The quantiles in panels (A)–(C) are computed from all  $N = 4000$  validation trajectories using the corresponding scenario probabilities.

Table 10: Results for one-step and multistep WRAP in the stationary Heston setting.

| Attack         | Method    | $\varepsilon$ | $\tau$ | Val. CVaR ( $\downarrow$ ) | Test CVaR ( $\downarrow$ ) | Mean PnL ( $\uparrow$ ) | Turnover ( $\downarrow$ ) |
|----------------|-----------|---------------|--------|----------------------------|----------------------------|-------------------------|---------------------------|
| Baseline       | ERM-U     | –             | –      | $3.634 \pm 0.013$          | $3.629 \pm 0.018$          | $-2.199 \pm 0.005$      | $3.084 \pm 0.124$         |
| One-step WRAP  | $\phi$ -U | –             | 0.25   | $3.627 \pm 0.014$          | $3.629 \pm 0.025$          | $-2.198 \pm 0.004$      | $3.283 \pm 0.266$         |
|                | OT-U      | 0.10          | –      | $3.550 \pm 0.011$          | $3.543 \pm 0.016$          | $-2.199 \pm 0.004$      | $3.156 \pm 0.064$         |
|                | WRAP-U    | 0.10          | –      | $3.550 \pm 0.011$          | $3.543 \pm 0.016$          | $-2.199 \pm 0.004$      | $3.156 \pm 0.064$         |
| Multistep-WRAP | $\phi$ -U | –             | 0.25   | $3.627 \pm 0.014$          | $3.629 \pm 0.025$          | $-2.198 \pm 0.004$      | $3.283 \pm 0.266$         |
|                | OT-U      | 0.10          | –      | $3.552 \pm 0.011$          | $3.545 \pm 0.015$          | $-2.200 \pm 0.004$      | $3.167 \pm 0.060$         |
|                | WRAP-U    | 0.10          | –      | $3.552 \pm 0.011$          | $3.545 \pm 0.015$          | $-2.200 \pm 0.004$      | $3.167 \pm 0.060$         |

*Note.* Values are mean  $\pm$  standard deviation over five seeds.

Panel (B) isolates reweighting with divergence budget  $\tau = 0.6$  and transport budget  $\varepsilon = 0$ . The adversary replaces  $w_n$  by  $\hat{w}_n$ , allocating greater probability to trajectories with larger loss integrands  $L_n(\vartheta)$ . This changes the weighted quantiles while leaving the trajectories unchanged. The path labeled “Highest-weight path” has the largest adversarial probability  $\hat{w}_n$ . The resulting probability concentration is summarized by the effective sample size,

$$N_{\text{eff}}(\hat{w}) = \frac{1}{\sum_{n=1}^N (\hat{w}_n)^2} \approx 691,$$

compared with  $N_{\text{eff}}(w) = 4000$  before reweighting. The reduced effective sample size reflects the concentration of adversarial probability on a subset of scenarios. The same construction accommodates nonuniform reference probabilities  $w_n$  in the -W specification.

Panel (C) adds optimal transport with budget  $\varepsilon = 0.3$ , holding the reweighted probabilities fixed. The attack perturbs trajectories in directions that increase the linearized loss, with perturbations scaled according to local loss sensitivity and the transport geometry. Both stock prices and variances are perturbed; the figure displays the stock-price component. Panel (D) provides a magnified view of the region highlighted in (C), comparing the stock-price components of the original and transported versions of the most displaced trajectory. This trajectory has the largest displacement among all  $N$  validation trajectories with respect to the ground metric  $d_{\kappa,\lambda}$  in (5).




Figure 9: Pathwise effects in nonstationary Heston: (A) ERM-U, (B)  $\phi$ -U, (C) WRAP-U, and (D) magnified view of the original and transported paths highlighted in (C). Panels (A)–(C) show sampled trajectories and highlight the highest-weight and most-displaced paths.