



Research article

Forecasting future realized variance paths with depth-weighted ridge and conformal diagnostics

Çağlar Sözen¹ and Fikriye Kabakcı^{2,*}

¹ Görele School of Applied Sciences, Department of Finance and Banking, Giresun University, Giresun, Turkey; ORCID: 0000-0002-3732-5058

² Faculty of Arts and Sciences, Department of Mathematics, Recep Tayyip Erdoğan University, Rize, Turkey; ORCID: 0000-0001-6266-1902

* **Correspondence:** Email: fikriye.kabakci@erdogan.edu.tr.

Abstract: We studied the problem of forecasting full future realized–variance (FRV) paths $y_{t,1:H}$ over $H = 30$ trading days. We proposed a depth–weighted ridge (DW–ridge) estimator that (i) enforces the natural monotonicity of cumulative variance via a pool–adjacent violators post–projection and (ii) adapts to market regimes through observation weights derived from a Wasserstein–based curve depth. At the daily frequency, we took squared returns as a practical realized–variance proxy, so that the FRV path is the cumulative sum of next–day squares. Empirically, we used daily data for two liquid U.S. exchange–traded funds (ETFs; XLE and SLV) and two major cryptocurrencies (BTC–USD and ETH–USD) from January 1, 2020, to December 31, 2024, under a 60%/20%/20% train–calibration–test split. On the ETF benchmarks, DW–ridge improved all–horizon pathwise root mean squared error (RMSE) by about 3.1% (XLE) and 2.8% (SLV) relative to a monotone ridge baseline, with statistically significant short–horizon ($H1-3/H1-5$) mean squared error (MSE) gains under a moving–block bootstrap. On BTC–USD and ETH–USD, all–horizon RMSE reductions were around 6.0% and 6.5%, respectively. A block–conformal diagnostic based on depth–derived nonconformity scores attained near–nominal or conservative coverage on test blocks, so sharper forecasts were not obtained at the expense of reliability. Overall, depth reweighting provided a simple, fast, and empirically effective enhancement to monotone FRV path forecasting across both sector ETFs and major cryptocurrencies.

Keywords: realized variance; variance–path forecasting; isotonic regression; Wasserstein curve depth; conformal prediction; ridge regression; financial econometrics; cryptocurrencies

Mathematics Subject Classification: 62M10, 62P20, 91G70

1. Introduction

Short-horizon risk control, option hedging, and algorithmic execution depend on the entire *path* of near-term variability rather than on a single volatility level. For daily data, widely used realized measures of conditional variability (e.g., realized variance/volatility) are computed from high-frequency returns; under mild regularity, such measures consistently track the latent variance process and underpin much of the modern volatility-forecasting literature [1, 2]. We study the problem of forecasting full future realized-variance (FRV) paths $y_{t,1:H}$ over $H = 30$ trading days. We propose a depth-weighted ridge (DW-ridge) estimator that (i) enforces the natural monotonicity of cumulative variance via a pool-adjacent violators post-projection and (ii) adapts to market regimes through observation weights derived from a Wasserstein-based curve depth. At the daily frequency, we take squared returns as a practical realized-variance proxy, so that the FRV path is simply the cumulative sum of next-day squares. Casting the problem at the level of a curve (the FRV trajectory) rather than a single-horizon variance aligns evaluation with practice, such as one-period-ahead value at risk (VaR) limits, short-term drawdown constraints, and rolling risk budgets, all of which depend on the entire profile of near-term variability.

Two design desiderata are central to FRV-path forecasting. First, the target path is monotone in the horizon by construction; ignoring this structure wastes statistical information and can yield incoherent risk metrics *ex post* (for instance, longer horizons being assigned lower variance than shorter ones). Second, not all historical episodes are equally informative for current market conditions: Training on long windows that mix heterogeneous regimes can degrade stability unless one adapts to the prevailing state. These considerations motivate a learner that (i) enforces monotonicity and (ii) reweights training examples by centrality relative to the present regime.

We propose a simple and fast estimator that meets both goals. Our approach combines (i) a post-projection enforcing monotonicity via the classical pool-adjacent violator (PAV) algorithm from isotonic regression, and (ii) DW-ridge, where observation weights are computed from a Wasserstein-based curve depth in the space of quantile functions. The PAV is a standard tool for order-restricted estimation with provable optimality and linear-time implementations for squared-error loss [3, 4]. We apply it as a post-fit projection rather than embedding the constraint into the regression, because this delivers shape-corrected forecasts in a single linear-time pass without refitting the base learner; by nonexpansiveness, the ℓ_2 error cannot increase under this monotonicity correction. For regime adaptation, we map each FRV path to its quantile representation and use optimal-transport geometry to define distances and robust centers. In one dimension the 2-Wasserstein distance reduces to the L^2 distance between quantile functions [5], endowing the space of volatility paths with an interpretable Euclidean structure. This facilitates the construction of depth notions that rank trajectories by centrality or outlyingness [6–9], yielding weights that up-weight regime-central curves and down-weight outlying episodes. The resulting DW-ridge estimator inherits the stability and interpretability of ridge regression [10] while adapting to market regimes through geometry-aware weights.

Beyond point accuracy, we seek transparent validity diagnostics. We therefore equip DW-ridge with a block-conformal diagnostic based on depth-derived nonconformity scores. On a calibration window, we compute blockwise scores from one-minus-depth, i.e., $1 - \text{depth}$, and derive thresholds that target nominal coverage on future blocks, obtaining distribution-free finite-sample reliability under weak dependence assumptions [11–15]. The resulting thresholds are visualized through block-conformal

control charts that flag when realized FRV paths fall outside the anticipated regime. Our construction is related in spirit to the uniform one-sided conformal bands for FRV curves proposed by [20], which use an isotonic baseline and block-maxima calibration to control the entire trajectory. In contrast, the present paper focuses on improving FRV *forecasts* via depth-adaptive weighting and treats block-conformal bands primarily as an auditing device for calibration.

Empirically, we study a four-asset universe comprising two liquid U.S. exchange-traded funds—XLE (Energy) and SLV (Silver)—and two major cryptocurrencies—BTC-USD and ETH-USD—over 2020–2024. XLE and SLV provide sector/commodity benchmarks with well-studied volatility dynamics, while BTC-USD and ETH-USD offer a high-volatility, heavy-tailed stress test for the methodology. The experimental design emphasizes the short horizons most relevant in practice (H1–3/H1–5). On the ETF benchmarks, depth weighting yields statistically significant short-horizon error reductions relative to a monotone ridge baseline, and block-conformal control charts achieve near-nominal coverage on held-out blocks. On BTC-USD and ETH-USD, we observe qualitatively similar short-horizon RMSE gains together with stable block-conformal diagnostics, indicating that the shape-constrained, depth-adaptive learner remains effective even under pronounced tail risk.

Our contributions are threefold. (i) We introduce a ridge-based FRV curve forecaster that jointly enforces the target's monotone shape (via the PAV) and adapts to regime centrality using Wasserstein-based curve-depth weights; hyperparameters are selected by blocked cross-validation focused on short horizons. (ii) We develop a block-conformal diagnostic tailored to depth scores, yielding interpretable one-sided prediction bands with finite-sample validity under dependence. (iii) A reproducible pipeline on XLE and SLV (2020–2024), complemented by robustness checks on BTC-USD and ETH-USD, documents consistent short-horizon RMSE gains together with calibrated coverage, indicating that combining shape constraints, geometry-aware weighting, and conformal auditing is an effective recipe for variance-path forecasting in risk-sensitive applications. A natural question is whether similar improvements could be obtained from more flexible sequence models, such as long short-term memory (LSTM) approaches or transformers trained on volatility covariates; we view such nonlinear benchmarks as complementary to our simple, auditable learner and highlight a systematic comparison as a direction for future work.

2. Related literature

The realized-variance paradigm shows that integrated variance can be recovered from high-frequency returns, giving rise to an extensive literature on modeling and forecasting realized volatility [1, 2]. Within this framework the HAR-RV model of [16] remains a widely used benchmark, capturing long-memory features through daily/weekly/monthly components. Most contributions, however, focus on one-step-ahead realized variance or conditional variance, whereas risk management often requires a path of short-horizon variance over the next few weeks. Our formulation embraces this curve perspective and evaluates performance horizon-by-horizon, with particular emphasis on H1–3/H1–5, where operational decisions—leverage adjustments, margin calls, and hedging policies—are most sensitive. While much of the empirical volatility literature concentrates on equity indices and sector ETFs, similar realized-volatility tools have increasingly been applied to major cryptocurrencies as these markets have matured; by including BTC-USD and ETH-USD alongside XLE and SLV, our study illustrates that FRV-path methods can be deployed in both traditional and digital-asset settings.

Ridge regression is a classical shrinkage technique that stabilizes noisy, collinear designs and often improves out-of-sample prediction risk [10, 17]. Penalized alternatives such as the LASSO and elastic net induce sparsity and grouped selection [18, 19], but when covariates are low-dimensional and strongly correlated—as in heterogeneous autoregressive (HAR)-style lagged volatility regressors—ridge typically offers a favorable bias-variance trade-off. Shape constraints provide an orthogonal axis of structure: For monotone targets, isotonic regression with the PAV algorithm yields the optimal squared-error fit under monotonicity in linear time [3, 4]. Recent applications in finance and econometrics exploit such constraints to enforce economically plausible response shapes. Our approach fits a multi-horizon ridge model for the FRV curve and then applies a fast isotonic projection to restore monotonicity across horizons, thereby guaranteeing coherent variance paths without re-estimating the underlying regression.

Functional data analysis offers depth notions that rank curves by centrality or outlyingness, enabling robust estimation and outlier detection [6, 7]. In parallel, optimal transport endows probability distributions with a geometric structure via Wasserstein distances and barycenters [5, 8, 9]. In one dimension, the 2-Wasserstein distance reduces to an L^2 distance between quantile functions, which makes computation particularly transparent. We build on these ideas by representing each FRV path through its quantile function and defining a Wasserstein-based curve depth that measures how close a trajectory lies to a barycentric “center” of past paths. Depth-weighted regression then uses these scores as observation weights, up-weighting regime-central curves and down-weighting outlying or off-regime episodes. To our knowledge, such Wasserstein-based depth weighting has not previously been used to adapt volatility forecasts to changing regimes, either for sector ETFs or for cryptocurrencies.

Conformal prediction furnishes distribution-free prediction sets with finite-sample coverage guarantees under minimal assumptions on the data-generating mechanism [11]. A growing body of work adapts conformal methods to non-exchangeable settings and complex predictors, including conformalized quantile regression, jackknife+ constructions, and procedures tailored to time-series data [12–15]. Within the specific context of FRV, [20] proposed uniform one-sided conformal bands for entire FRV trajectories, using monotone isotonic fits and block-maxima calibration to guarantee that future FRV paths remain below a data-adaptive upper envelope with high probability. Our contribution is complementary: We aim to improve multi-horizon point forecasts via shape-constrained, depth-weighted regression, and we employ block-conformal bands primarily as a calibration diagnostic that assesses whether the regime-adaptive learner delivers reliable short-horizon variance paths.

Relative to HAR-style baselines, standard shrinkage methods, and recent conformal-band approaches, the combination of monotone FRV curve forecasting, Wasserstein-based depth weighting, and block-conformal auditing appears to be new. It delivers a simple, transparent benchmark against which more flexible nonlinear sequence models (e.g., LSTMs or transformers for volatility paths) can be evaluated in future work, both for liquid sector ETFs and for major cryptocurrencies such as BTC-USD and ETH-USD.

3. Data

We work with daily adjusted close prices for the energy sector ETF XLE, the silver ETF SLV, and the two major cryptocurrencies BTC-USD and ETH-USD from January 1, 2020, to December 31, 2024, obtained from Yahoo Finance [21]. For comparability across assets, we align all series to a common

calendar of U.S. equity trading days (New York Stock Exchange (NYSE) and Nasdaq sessions). In particular, for BTC–USD and ETH–USD—traded continuously—we subsample the daily series to the same U.S. trading-day dates used by the ETFs, so that each FRV origin t corresponds to the same set of calendar dates across all four assets.

Log returns are defined as

$$r_t = \log P_t - \log P_{t-1}.$$

At the daily frequency, we follow the realized–variance literature and use *daily squared returns* as a standard realized–variance proxy,

$$RV_t = r_t^2,$$

so that the empirical analysis is conducted in terms of variance levels rather than volatility (no square–root transformation is applied).

The forecasting target is the FRV path

$$y_{t,h} = \sum_{j=1}^h r_{t+j}^2, \quad h = 1, \dots, 30,$$

that is, the cumulative sum of next–day squared returns over a 30–day horizon. After enforcing a full 30–day look–ahead window and the HAR look–backs used as regressors, we obtain $N_{\text{total}} = 1205$ FRV origins per asset, of which $N_{\text{train}} = 723$ are used for training, $N_{\text{cal}} = 241$ for conformal calibration, and $N_{\text{test}} = 241$ for out-of-sample evaluation. The corresponding FRV origin dates range from February 4, 2020, to November 14, 2024, for all four assets. For transparency and reproducibility, the detailed chronological sample splits and associated FRV origin dates are reported in Appendix Table A.1, while Table 1 reports descriptive statistics for the training–sample FRV curves.

For the two ETF series, XLE exhibits a substantially higher scale of terminal FRV than SLV: The mean 30–day FRV level increases from about $\overline{y(H)} \approx 0.014$ for SLV to $\overline{y(H)} \approx 0.023$ for XLE, with corresponding mean one–step increments of approximately 0.00048 vs. 0.00077. Skewness and excess kurtosis of $y(H)$ indicate pronounced right tails for both ETFs, but especially for XLE (skewness ≈ 4.020 , excess kurtosis ≈ 16.370), reflecting occasional high–volatility episodes in the energy sector. The two cryptocurrencies display an even higher FRV scale. For BTC–USD, the mean terminal FRV is around $\overline{y(H)} \approx 0.065$ with mean per–step increments near 0.0022, whereas ETH–USD reaches $\overline{y(H)} \approx 0.119$ and mean increments close to 0.0040. All four assets have clearly right–skewed and leptokurtic FRV distributions, with excess kurtosis ranging from roughly 3.990 (ETH–USD) to 16.370 (XLE), underscoring the relevance of heavy–tailed behavior in short–horizon variance paths. The Wasserstein curve–depth (WCD) statistics in Table 1 show median depths in a relatively narrow band (about 0.220–0.260) across assets. This suggests that, despite large cross–asset differences in overall variance scale, FRV curves remain reasonably concentrated around central shapes, with extreme episodes captured by the upper tails of $y(H)$ and by the spread of the depth distribution. These features motivate the use of a robust, regime–adaptive forecasting pipeline: MAD–type scaling and depth–based reweighting aim to stabilize estimation in the presence of outliers, while the subsequent blockwise conformal step is tailored to serial dependence in FRV sequences.

Table 1. Descriptive statistics of training FRV curves by asset (2020–2024).

Metric	XLE	SLV	BTC–USD	ETH–USD
Train N	723	723	723	723
H (horizons)	30	30	30	30
Mean $y(H)$	0.023	0.014	0.065	0.119
SD $y(H)$	0.033	0.013	0.057	0.100
Min $y(H)$	0.005	0.004	0.006	0.019
Median $y(H)$	0.013	0.009	0.049	0.086
Max $y(H)$	0.205	0.069	0.313	0.499
Mean of avg incr.	0.001	0.000	0.002	0.004
SD of avg incr.	0.001	0.000	0.002	0.003
WCD median	0.250	0.220	0.255	0.254
WCD mean	0.261	0.224	0.266	0.266
Skew. of $y(H)$	4.019	2.340	2.716	2.032
Excess kurt. of $y(H)$	16.370	5.343	8.494	3.993

4. Methodology

We forecast the FRV path $\mathbf{y}_t = (y_{t,1}, \dots, y_{t,H})^\top$ over $H = 30$ trading days using a shape-aware and regime-adaptive linear learner. The design couples four ingredients: (i) HAR-style regressors for parsimony and direct comparability with the realized-variance literature, (ii) ridge shrinkage to stabilize multi-horizon fits under collinearity, (iii) WCD reweighting to adapt to the prevailing volatility regime, and (iv) a monotone post-projection via the PAV operator to enforce the cumulative structure of FRV paths. Alongside ridge regression and DW-ridge, we also estimate a simple heterogeneous autoregressive ordinary least squares (HAR-OLS) benchmark for one-step-ahead realized variance and then iterate it forward to obtain a pseudo-FRV path; this clarifies how much of the performance gains are due to depth reweighting and monotonicity rather than to regularization alone.

4.1. Forecast target and monotone projection

As introduced in Section 3, the FRV path for each origin t is

$$y_{t,h} = \sum_{j=1}^h r_{t+j}^2, \quad h = 1, \dots, H,$$

with $H = 30$. Thus $y_{t,h}$ is the cumulative realized variance over the next h trading days, and the FRV path is nondecreasing in h by construction.

Given a raw forecast $\widehat{\mathbf{y}}_t \in \mathbb{R}^H$, we impose the shape constraint via the isotonic (PAV) projection [3, 4]:

$$\widehat{\mathbf{y}}_t^\uparrow = \arg \min_{\mathbf{z} \in \mathbb{R}^H} \|\mathbf{z} - \widehat{\mathbf{y}}_t\|_2^2 \quad \text{s.t.} \quad z_1 \leq \dots \leq z_H. \quad (4.1)$$

The PAV is the Euclidean projection onto the isotone cone, is nonexpansive in ℓ_2 , and therefore cannot increase the squared error whenever the true FRV path is monotone. We apply the PAV to each predicted

path ex post, rather than embedding monotonicity directly in the fit, which keeps the base learner simple and allows us to interpret DW–ridge as a standard linear smoother followed by a deterministic shape correction.

4.2. Ridge, DW–ridge, and HAR baselines with WCD weights

Let $RV_t = r_t^2$ and define HAR–style regressors

$$X_t = (1, RV_t, \overline{RV}_t^{(w)}, \overline{RV}_t^{(m)}),$$

where $\overline{RV}_t^{(w)}$ and $\overline{RV}_t^{(m)}$ denote trailing weekly and monthly averages of RV_t over $K_w = 5$ and $K_m = 22$ trading days, respectively. Stacking over t gives the design matrix $X \in \mathbb{R}^{n \times p}$ ($p = 4$) and the FRV response matrix $Y \in \mathbb{R}^{n \times H}$.

The multi–horizon ridge estimator solves

$$\widehat{B}_\lambda = \arg \min_{B \in \mathbb{R}^{p \times H}} \sum_{t=1}^n \|y_t - X_t B\|_2^2 + \lambda \|B\|_F^2, \quad \Rightarrow (X^\top X + \lambda I_p) \widehat{B}_\lambda = X^\top Y. \quad (4.2)$$

We favor ridge over unpenalized OLS for the multi–horizon system because the HAR regressors are deliberately overlapping and therefore highly collinear: Daily, weekly, and monthly components are constructed from the same underlying squared returns. In this setting, ridge provides a transparent and computationally trivial way to stabilize the Gram matrix and control the effective degrees of freedom, while preserving the HAR structure that practitioners are familiar with.

As a complementary benchmark, we also estimate a *single–horizon* HAR–OLS model for one–step realized variance,

$$RV_{t+1} = \beta_0 + \beta_1 RV_t + \beta_2 \overline{RV}_t^{(w)} + \beta_3 \overline{RV}_t^{(m)} + \varepsilon_{t+1},$$

and recursively iterate this model forward to generate pseudo–FRV paths via cumulative sums of predicted RV_{t+h} . This HAR–OLS baseline isolates the value added by ridge shrinkage and depth weighting relative to a de facto standard in the realized–variance literature [16].

To construct WCD weights, we associate to each FRV path \mathbf{y} a discrete probability distribution on horizons. We set increments $\Delta_h = y_h - y_{h-1}$ ($\Delta_1 := y_1$) and define discrete masses

$$p_h = \frac{\max(\Delta_h, 0)}{\sum_{j=1}^H \max(\Delta_j, 0)}.$$

On a fixed grid $U = \{u_h\}_{h=1}^H \subset (0, 1)$, we compute the quantile vector $Q(\mathbf{y}) \in \mathbb{R}^H$ associated with (p_h) . In one dimension, the 2–Wasserstein distance between two distributions is just the ℓ_2 distance between their quantile functions on U , so the map $\mathbf{y} \mapsto Q(\mathbf{y})$ embeds FRV paths into a geometrically meaningful and computationally tractable space.

Let \bar{Q} denote the average quantile vector over the training sample and let

$$d(\mathbf{y}) = \|Q(\mathbf{y}) - \bar{Q}\|_2$$

denote its Euclidean displacement. With robust scale $s = \text{MAD}(d(\mathbf{y}))$, we define

$$\text{Depth}(\mathbf{y}) = \frac{1}{1 + d(\mathbf{y})/s} \in (0, 1]. \quad (4.3)$$

Curves with small Wasserstein distance to the center ($d(\mathbf{y})$ small) receive depth values close to 1 and are interpreted as regime–central episodes; curves far from \bar{Q} receive shallow depths. MAD scaling ensures that a fixed fraction of extreme episodes cannot unduly influence the center or the scale, which is particularly important in the presence of volatility spikes and structural breaks. Intuitively, if the future volatility regime resembles a cluster of central historical curves, then these curves should carry more information about the FRV path than atypical stress or calm episodes; the WCD weights implement this idea in a simple geometry–aware way.

For floor $\alpha_0 \in [0, 1)$ and curvature $\gamma > 0$, we set

$$w_t = \alpha_0 + (1 - \alpha_0) \text{Depth}(\mathbf{y}_t)^\gamma, \quad \widehat{B}_{\lambda, \gamma, \alpha_0} = \arg \min_B \sum_{t=1}^n w_t \|\mathbf{y}_t - X_t B\|_2^2 + \lambda \|B\|_F^2, \quad (4.4)$$

with

$$(X^\top W X + \lambda I_p) \widehat{B} = X^\top W Y, \quad W = \text{diag}(w_1, \dots, w_n).$$

The curvature γ controls how aggressively the weighting discriminates between central and outlying curves, while the floor α_0 prevents any episode from being assigned a near–zero weight. In practice, α_0 shrinks DW–ridge continuously back toward standard ridge as it approaches one. Predictions for all models (HAR–OLS, ridge, DW–ridge) are post–processed by the PAV projection horizon–wise to enforce monotonicity of the FRV path.

4.3. Model selection and the block–conformal diagnostic

We adopt K –fold blocked cross–validation (contiguous folds) focused on short horizons (H1–3, optionally H1–5) to select $(\lambda, \gamma, \alpha_0)$ [22]. Folds are formed as consecutive blocks in calendar time to respect temporal dependence; in each CV iteration we fit the model on past blocks and evaluate the mean squared error of PAV–projected predictions on H1–3 in the held–out block. The same blocked CV scheme is used to choose the ridge penalty λ for both the ridge and DW–ridge specifications, ensuring that any subsequent performance differences are attributable to the depth weights rather than to more favorable tuning. The HAR–OLS baseline does not involve a penalty parameter and is estimated once on the training window. In the implementation we use a moderate logarithmic grid for λ (from very mild to strong shrinkage) and low–dimensional grids for (γ, α_0) , which keeps the computational burden negligible.

For finite–sample *validity*, we implement a block–conformal diagnostic: On the calibration window, we compute depth–based nonconformity scores $s_t = 1 - \text{Depth}(\mathbf{y}_t)$; within blocks of length B , we record

$$S_b^{(\text{cal})} = \max_{t \in \text{block } b} s_t$$

and set the threshold $q_{B, \alpha}$ as the empirical $(1 - \alpha)$ quantile (order statistic) of $\{S_b^{(\text{cal})}\}$. On the test window, we form the analogous block maxima $S_b^{(\text{test})}$ and report empirical block coverage

$$\Pr\{S_b^{(\text{test})} \leq q_{B, \alpha}\}$$

and its deviation from the nominal level $1 - \alpha$ [11–13]. Throughout, we view this block–conformal construction as a *diagnostic* rather than as an operational prediction band: It quantifies how often test blocks look at least as “central” as calibration blocks under the depth score, thereby providing an interpretable reliability check that complements RMSE profiles and short–horizon MSE gains.

4.4. Experimental design

The empirical analysis follows a simple, fully time-respecting design that mirrors the R pipeline used for the computations. For each of the four assets in Section 3 (XLE, SLV, BTC-USD, ETH-USD), we take the same adjusted-close series from Yahoo Finance, construct log-returns r_t , and form FRV paths \mathbf{y}_t with horizon $H = 30$ exactly as defined there. After discarding the initial observations needed to compute HAR regressors and the full FRV horizon, we obtain a panel of FRV paths $\{\mathbf{y}_t\}_{t=1}^{n_a}$ for each asset a .

For each asset, the time index is split chronologically into three disjoint windows: A training window (60% of FRV origins), a calibration window (20%), and a test window (20%). All hyperparameters $(\lambda, \gamma, \alpha_0)$ are selected exclusively on the training window via blocked cross-validation, and the calibration window is reserved for the block-conformal diagnostic. The test window is used once for out-of-sample evaluation. This 60/20/20 scheme, implemented asset by asset, avoids any look-ahead or reuse of information across tuning, calibration, and evaluation.

We compare three specifications: (i) The HAR-OLS baseline iterated forward to generate pseudo-FRV paths, (ii) ridge with HAR regressors and PAV post-projection, and (iii) DW-ridge, which augments ridge with WCD weights and the same PAV correction. All models share the same FRV construction, regressors, and train/cal/test splits, ensuring a level comparison.

Performance is summarized along two complementary axes. First, we report full horizon-wise RMSE profiles,

$$\text{RMSE}_h = \left(\frac{1}{n_{\text{te}}} \sum_{t \in \mathcal{T}_{\text{test}}} (y_{t,h} - \widehat{y}_{t,h})^2 \right)^{1/2}, \quad h = 1, \dots, H,$$

for each model and asset. Second, we emphasize short-horizon accuracy via aggregate H1-3 and H1-5 metrics, defined as averages of squared errors over $h \in \{1, 2, 3\}$ and $h \in \{1, \dots, 5\}$, respectively. The tables report relative MSE gains of DW-ridge and ridge over HAR-OLS and of DW-ridge over ridge on these aggregates. To quantify sampling uncertainty, we compute block bootstrap confidence intervals for MSE differences by resampling non-overlapping blocks of FRV origins on the test window; this preserves local dependence while delivering interpretable intervals for short-horizon gains.

To highlight the regime-adaptive aspect of WCD, we also consider a simple calendar-based regime split of the evaluation sample. Specifically, we partition the post-2020 period into two subperiods, R1 (2021–2022) and R2 (2023–2024), and recompute the short-horizon error metrics for all models within each regime (Table A.3). This regime-wise breakdown shows whether DW-ridge improves predominantly in earlier versus more recent market conditions, and is constructed directly from the same R pipeline that underlies the main tables and figures.

4.5. Algorithmic pipeline

Algorithm 1 DW–ridge (HAR regressors, PAV, Wasserstein curve depth) with block–conformal coverage

- Require:** Prices $\{P_t\}$; horizon $H = 30$; windows $(w, m) = (5, 22)$; CV folds K ; block grid \mathcal{B} ; miscoverage grid \mathcal{A} .
- 1: **Returns/RV:** $r_t = \log P_t - \log P_{t-1}$; $RV_t = r_t^2$.
 - 2: **Targets:** $y_{t,h} = \sum_{j=1}^h r_{t+j}^2$; form $Y \in \mathbb{R}^{n \times H}$.
 - 3: **Regressors:** $X_t = (1, RV_t, \overline{RV}_t^{(w)}, \overline{RV}_t^{(m)})$; stack X .
 - 4: **Split:** time–ordered train (60%), calibration (20%), test (20%).
 - 5: **Train depth:** build $Q(\mathbf{y})$ on U ; compute \bar{Q} and s (MAD); set $\text{Depth}(\mathbf{y}_t)$ on train.
 - 6: **Blocked CV:** for $(\lambda, \gamma, \alpha_0)$, fit weighted ridge with w_t on past blocks; predict on holdout; apply PAV; score H1–3 MSE. Pick $(\lambda^*, \gamma^*, \alpha_0^*)$.
 - 7: **Refit:** compute \widehat{B}^* on full train using W^* .
 - 8: **HAR–OLS benchmark:** fit HAR–OLS for RV_{t+1} on train; recursively iterate to obtain FRV paths on calibration/test; apply PAV.
 - 9: **Conformal (calibration):** compute $s_t = 1 - \text{Depth}(\mathbf{y}_t)$; for each (B, α) , set $q_{B,\alpha}$ from block maxima $S_b^{(\text{cal})}$.
 - 10: **Test:** predict $\widehat{Y}_{\text{te}} = X_{\text{te}} \widehat{B}^*$; apply PAV; report RMSE profiles and short–horizon gains vs. ridge and HAR–OLS; compute $S_b^{(\text{test})}$ and empirical coverage vs. $q_{B,\alpha}$.

4.6. Assumptions and auxiliary results

Assumption 4.1. $\{r_t\}$ has finite fourth moments and $RV_t = r_t^2$ has finite second moments; FRV paths are nondecreasing in h .

Assumption 4.2. $\{(X_t, \mathbf{y}_t)\}$ is strictly stationary or, more generally, piecewise locally stationary with weak dependence (e.g., β –mixing), so that blocked cross-validation and block-conformal approximations are valid.

Assumption 4.3. X has full column rank asymptotically; ridge with $\lambda > 0$ ensures well-posedness in finite samples.

Assumption 4.4. The average quantile \bar{Q} and robust scale s on train are well-defined; $\mathbf{y} \mapsto Q(\mathbf{y})$ is bounded and Lipschitz on the fixed grid U in 1D Wasserstein geometry.

Assumption 4.5. Block length B is large enough to weaken within-block dependence and small enough to yield many calibration blocks m_{cal} .

Assumption 4.6. PAV is the ℓ_2 projection onto the isotonic cone and is nonexpansive.

Lemma 4.1. Under Assumption 4.6, for any isotone \mathbf{y} and any vector $\widehat{\mathbf{y}}$,

$$\|\text{PAV}(\widehat{\mathbf{y}}) - \mathbf{y}\|_2 \leq \|\widehat{\mathbf{y}} - \mathbf{y}\|_2 \quad \text{and} \quad \|\text{PAV}(\widehat{\mathbf{y}}) - \text{PAV}(\widehat{\mathbf{z}})\|_2 \leq \|\widehat{\mathbf{y}} - \widehat{\mathbf{z}}\|_2.$$

Lemma 4.2. Under Assumption 4.3, $(X^\top W X + \lambda I_p) > 0$ for any diagonal $W \geq 0$, $\lambda > 0$; thus a unique \widehat{B} exists and depends continuously on (X, Y, W) .

Lemma 4.3. Under Assumption 4.4, $\text{Depth}(\mathbf{y}) \in (0, 1]$ is 1-Lipschitz in $\|\mathbf{Q}(\mathbf{y}) - \bar{\mathbf{Q}}\|_2/s$; MAD scaling grants resistance to a positive fraction of outliers.

Proposition 4.1. Suppose Assumptions 4.1–4.3 hold and the truth \mathbf{y}_t^* is isotone. Let $\widehat{\mathbf{y}}_t = X_t \widehat{\mathbf{B}}_{\lambda, \gamma, \alpha_0}$ and $\widehat{\mathbf{y}}_t^\uparrow = \text{PAV}(\widehat{\mathbf{y}}_t)$. Then

$$\mathbb{E} \|\widehat{\mathbf{y}}_t^\uparrow - \mathbf{y}_t^*\|_2^2 \leq \mathbb{E} \|\widehat{\mathbf{y}}_t - \mathbf{y}_t^*\|_2^2 = \underbrace{\|\mathbb{E} \widehat{\mathbf{y}}_t - \mathbf{y}_t^*\|_2^2}_{\text{bias}^2} + \underbrace{\text{tr Var}(\widehat{\mathbf{y}}_t)}_{\text{variance}}.$$

Moreover, for fixed weights W , the DW-ridge predictor equals ridge on the reweighted dataset $(\tilde{X}, \tilde{Y}) = (W^{1/2}X, W^{1/2}Y)$, yielding the usual shrinkage risk decomposition with effective Gram matrix $\tilde{G} = \tilde{X}^\top \tilde{X}$. In particular, depth reweighting can be viewed as modifying the local geometry of the design in favor of regime-central episodes, while PAV guarantees that the final FRV paths respect the cumulative variance structure.

Lemma 4.4. Under Assumptions 4.2 and 4.5, the calibration quantile $q_{B, \alpha}$ of block maxima $S_b^{(\text{cal})}$ attains empirical test coverage close to nominal $1 - \alpha$, up to $o_p(1)$ deviations decreasing with m_{cal} .

Remark 4.1. Blocked-CV risk minimized on H1–3 (optionally H1–5) aligns with practice; full RMSE profiles are reported for completeness.

Remark 4.2. An expanding-window center at test time can be used to track slow regime drift without look-ahead, consistent with Assumption 4.4.

Remark 4.3. The curvature parameter γ sets reweighting aggressiveness, whereas the floor α_0 guards against vanishing weights and shrinks the procedure toward unweighted ridge as $\alpha_0 \rightarrow 1$.

4.7. Computational complexity and implementation

Let n be the number of curves, $p = 4$, and $H = 30$.

- Feature/target construction: $O(nH)$.
- Depth scores (quantiles + MAD): $O(nH)$.
- Weighted ridge: form $X^\top W X$ in $O(np^2)$; solve the $p \times p$ system in $O(p^3)$.
- Prediction + PAV: $O(n_{\text{te}}(p + H))$.
- Block-conformal per (B, α) : $O(n_{\text{cal}} + n_{\text{te}})$.

Overall cost is linear in nH plus a negligible $O(p^3)$ term; with $H = 30$ and $p = 4$, the procedure is computationally light. This remains true for the HAR-OLS baseline and for regime-wise performance summaries, since these only require reusing the same FRV construction and refitting low-dimensional linear models on subsamples.

5. Results

We compare three learners for each asset (XLE, SLV, BTC-USD, ETH-USD): A HAR-OLS baseline on one-day realized variance, iterated to obtain FRV paths (HAR); a multi-horizon ridge model (ridge); and its depth-weighted counterpart (DW-ridge). Forecast quality is evaluated along three dimensions: (i) Horizon-wise RMSE profiles for the entire 30-day FRV path, (ii) short-horizon accuracy aggregated over H1–3 and H1–5, and (iii) reliability via block-conformal diagnostics (Section 6).

5.1. RMSE profiles over the FRV path

Figure 1 plots RMSE by horizon for all three learners and all four assets. DW-ridge uniformly dominates ridge along the full FRV path: For every asset and every horizon $h = 1, \dots, 30$, the DW-ridge curve lies weakly below the ridge curve. Ridge is broadly comparable to HAR at short horizons and tends to improve on HAR at medium and longer horizons, with the clearest gains for the more volatile cryptocurrency series. For the two ETFs, XLE and SLV, the gains from depth reweighting are moderate but systematic. On XLE, the RMSE gap between DW-ridge and ridge is visually most pronounced at medium and longer horizons ($h \gtrsim 10$), consistent with the idea that depth weights attenuate atypical regimes precisely where extrapolation along the FRV path is strongest. On SLV, the separation is smaller but persistent across h , pointing to nontrivial yet milder efficiency gains. For BTC-USD and ETH-USD, improvements are substantially larger. Both cryptocurrencies display higher FRV scale and heavier tails than the ETFs (Table 1); in this environment DW-ridge delivers clearly sharper RMSE profiles than both ridge and HAR across the entire 30-day horizon. The all-horizon RMSE reductions of DW-ridge relative to ridge, summarized in Table 3, are about 3.1% for XLE and 2.8% for SLV, compared with roughly 6.0% for BTC-USD and 6.5% for ETH-USD. Since all methods are post-processed with the same isotonic PAV projection, these gains arise from depth reweighting rather than from differences in shape enforcement. To condense the most practically relevant horizons, Table 2 reports average RMSE over H1–3 and H1–5 for each learner and asset. For all four series and for both short-horizon aggregates, DW-ridge attains the lowest average RMSE. Relative to HAR, the percentage reductions are again most pronounced for BTC-USD and ETH-USD, reflecting the greater potential for improvement in highly volatile, heavy-tailed variance environments.




Figure 1. RMSE by horizon for XLE, SLV, BTC-USD, and ETH-USD under HAR, ridge, and DW-ridge. DW-ridge uniformly dominates ridge across horizons for all assets, with particularly pronounced gains for the cryptocurrency series.

Table 2. Short-horizon RMSE averages over H1–3 and H1–5 by learner and asset.

Asset	HAR (H1–3)	Ridge (H1–3)	DW–ridge (H1–3)	HAR (H1–5)	Ridge (H1–5)	DW–ridge (H1–5)
XLE	0.0005516096	0.0005562883	0.0005397459	0.0008392604	0.0008409967	0.0008141342
SLV	0.0008910276	0.0008961439	0.0008909808	0.0010760429	0.0010768920	0.0010671359
BTC–USD	0.0039028640	0.0039082656	0.0037147848	0.0051304005	0.0051270675	0.0048072273
ETH–USD	0.0069856133	0.0069531987	0.0066570631	0.0093918490	0.0093035157	0.0088031247

Table 3. Hyperparameters for ridge and DW–ridge, conformal settings, and all-horizon RMSE improvements of DW–ridge relative to ridge.

Metric	XLE	SLV	BTC–USD	ETH–USD
λ_{Ridge}	0.0001	0.0100	0.0100	0.0100
λ_{DW}	0.0001	0.0100	0.0100	0.0100
γ_{DW}	3.000	3.000	2.000	3.000
Weighting	depth	depth	depth	depth
Weight floor α_0	0.250	0.100	0.100	0.100
Block size B	21	10	10	10
Miscoverage α	0.10	0.05	0.20	0.05
Threshold q	0.876	0.894	0.865	0.884
Coverage (block)	0.909	1.000	0.833	0.958
Deviation	0.009	0.050	0.033	0.008
All-horizon improv. [%]	3.061	2.789	5.972	6.507

5.2. Short-horizon gains and bootstrap uncertainty

Table 4 summarizes short-horizon MSE gains of DW–ridge relative to ridge over H1–3 and H1–5, together with 95% moving-block bootstrap confidence intervals based on 499 resamples. A companion appendix table (Table A.2) reports analogous statistics relative to HAR. For XLE, DW–ridge delivers clear and robust short-run improvements. On H1–3, the average MSE gain over ridge is about 2.1×10^{-8} (roughly 6%), with a strictly positive 95% interval $[1.9, 2.3] \times 10^{-8}$. On H1–5, the mean gain rises to approximately 5.6×10^{-8} (about 6.4%), again with a tight, strictly positive interval. Gains relative to HAR are slightly smaller in percentage terms but remain sizable and statistically well supported (Table A.2). For SLV, short-horizon gains are more modest. DW–ridge improves on ridge by about 1.0×10^{-8} on H1–3 ($\approx 1.2\%$) and 2.4×10^{-8} on H1–5 ($\approx 2.0\%$), with 95% intervals that are mostly positive but noticeably wider. Improvements over HAR are correspondingly mild, with intervals that occasionally include zero, in line with a less extreme volatility and depth structure for SLV. For BTC–USD and ETH–USD, DW–ridge yields markedly larger short-horizon gains. On BTC–USD, the H1–3 MSE reduction relative to ridge is around 1.7×10^{-6} (about 10.3%) and remains strictly positive in the 95% interval; H1–5 gains reach roughly 3.8×10^{-6} (about 12.9%). For ETH–USD, the corresponding reductions are approximately 4.6×10^{-6} ($\approx 8.9\%$) on H1–3 and 1.1×10^{-5} ($\approx 11.1\%$) on H1–5, again with strictly positive intervals. Relative to HAR, percentage gains for BTC–USD and ETH–USD exceed 10% on both short-horizon aggregates (Table A.2).

Taken together, Figure 1, Table 2, and Table 4 show that depth reweighting consistently improves short-run FRV forecasts across assets, with particularly pronounced benefits for the two cryptocurrency series where variance paths are more volatile and heavy-tailed.

Table 4. Short-horizon MSE gains of DW-ridge relative to ridge with 95% moving-block bootstrap confidence intervals. Positive values favor DW-ridge.

Asset	Horizon	Gain vs. Ridge	CI lower	CI upper	Percent gain vs. Ridge
XLE	H1-3	0.00000002120715	0.00000001907528	0.00000002330343	6.030303
SLV	H1-3	0.00000001030585	0.00000000063400	0.00000001914498	1.230376
BTC-USD	H1-3	0.00000166508244	0.00000081510260	0.00000237682772	10.284179
ETH-USD	H1-3	0.00000459480104	0.00000315929396	0.00000579319473	8.859073
XLE	H1-5	0.00000005571975	0.00000005055347	0.00000006128373	6.438308
SLV	H1-5	0.00000002418318	0.00000000224923	0.00000004553637	1.963855
BTC-USD	H1-5	0.00000375579297	0.00000209101886	0.00000518519036	12.854812
ETH-USD	H1-5	0.00001086657247	0.00000757520163	0.00001350923578	11.145783

5.3. Regime-wise performance

To examine whether depth reweighting is effective across different market conditions, Table A.3 decomposes short-horizon MSE by calendar regime, distinguishing an earlier period R1 (2021–2022) from a more recent period R2 (2023–2024). For each asset and regime, the table reports H1–3 and H1–5 MSE under ridge, DW-ridge, and HAR, together with absolute and percentage gains. For this regime-wise check, we recompute MSEs over all FRV origins in 2021–2024 (train+cal+test) to gain precision; these regime-specific numbers are descriptive and complement the strictly out-of-sample test results reported in Section 5. For XLE and SLV, DW-ridge improves on ridge and HAR in both regimes, with larger percentage gains in the more recent period R2. For example, short-horizon MSE reductions for XLE relative to ridge are around 2–3% in R1 but increase to about 5–6% in R2, consistent with depth weights becoming more valuable when volatility conditions change and atypical FRV trajectories appear more frequently. SLV shows a similar pattern, albeit with smaller magnitudes. For BTC-USD and ETH-USD, the regime contrast is sharper. In R1, DW-ridge yields modest but positive gains over ridge and HAR. In R2, when both cryptocurrencies experience more frequent volatility spikes, short-horizon MSE reductions relative to ridge exceed 11% for BTC-USD and lie around 12–14% for ETH-USD, with similarly large gains relative to HAR. These findings indicate that depth weights are particularly effective in high-volatility, regime-shift periods, where down-weighting atypical or outdated episodes materially sharpens FRV forecasts.

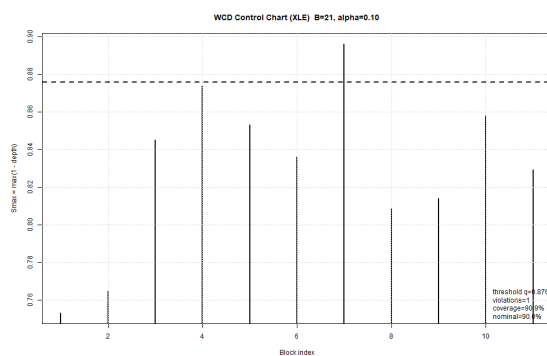
We view these regime-wise results as complementary to the aggregate RMSE and bootstrap analyses: They show that the advantages of DW-ridge are not confined to a single subperiod but persist, and even strengthen, in more stressed market environments.

6. Robustness


We assess the stability of our findings along three complementary dimensions: (i) The conformal grid (B, α) used for reliability diagnostics, (ii) the distribution of WCD on train vs. test splits, and (iii) the sensitivity of short-horizon gains to alternative evaluation windows (H1–3 versus H1–5) and to calendar regimes.

6.1. Conformal grid and coverage behavior


Table 3 reports the selected conformal settings (B, α) and the resulting empirical block coverages for each asset, while Table A.4 provides the full grid over block length B and miscoverage α for all four assets. As expected for block maxima, the calibration threshold q is monotone along the grid: Holding B fixed, q increases as α decreases (more stringent coverage), and holding α fixed, q typically increases with larger B . For XLE, the preferred configuration $(B, \alpha) = (21, 0.10)$ yields a test coverage of 90.9% (deviation $\approx +0.009$ from the nominal 90%), while neighboring grid points typically deliver coverage between roughly 83% and 93%. SLV uses shorter blocks and a stricter nominal level, with $(10, 0.05)$ producing 100% coverage (deviation $+0.05$) and nearby configurations remaining conservative, often at or above 95% coverage. For BTC–USD and ETH–USD, the grid likewise shows that empirical coverage is close to or above the nominal across a broad range of (B, α) pairs. The selected settings $(10, 0.20)$ for BTC–USD and $(10, 0.05)$ for ETH–USD yield test coverages of 83.3% and 95.8%, respectively, corresponding to deviations of about $+0.033$ and $+0.008$ from their nominal levels. Hence, the reliability conclusions drawn in Section 5 do not hinge on a finely tuned choice of block size or miscoverage: Depth-based conformal diagnostics remain near-nominal or conservative for all four assets. Figure 2 visualizes these diagnostics via block-conformal WCD control charts, plotting the test-block maxima S_{\max} against the calibration threshold q for each asset.



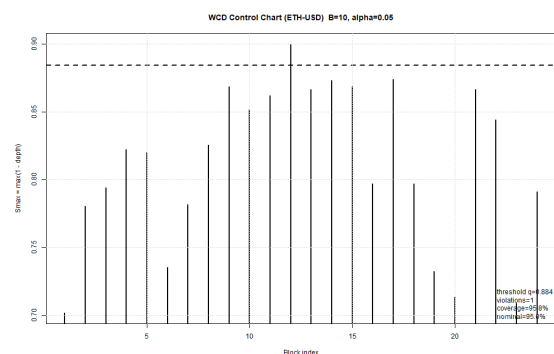
(a) XLE: $B = 21, \alpha = 0.10$



(b) SLV: $B = 10, \alpha = 0.05$



(c) BTC–USD: $B = 10, \alpha = 0.20$



(d) ETH–USD: $B = 10, \alpha = 0.05$

Figure 2. Block-conformal WCD control charts for all four assets. For each test block, the statistic $S_{\max} = \max(1 - \text{Depth}(\mathbf{y}_t))$ is plotted against the calibration threshold q (dashed line). Empirical coverages (Table 3) are close to or above the nominal levels $1 - \alpha$ in all cases.

6.2. Depth distributions and regime stability

Figure 3 compares train and test WCD distributions for all four assets. For XLE and SLV, the train and test histograms exhibit substantial overlap, with mass concentrated in a similar central depth region and only mild differences in the tails. This suggests that, at the FRV–curve level, volatility regimes remain broadly comparable across train and test windows, which is consistent with the near–nominal or conservative coverage obtained in the conformal grid. For BTC–USD and ETH–USD, the central parts of the train and test depth distributions again largely coincide, but upper tails are somewhat heavier on the test window, reflecting more frequent volatility spikes in the later period. Despite this increased tail weight, DW–ridge continues to deliver uniform RMSE improvements and sizeable short–horizon MSE gains (Section 5), together with well–behaved conformal diagnostics. This indicates that the depth weighting remains effective even when the evaluation period is somewhat more volatile than the training sample.

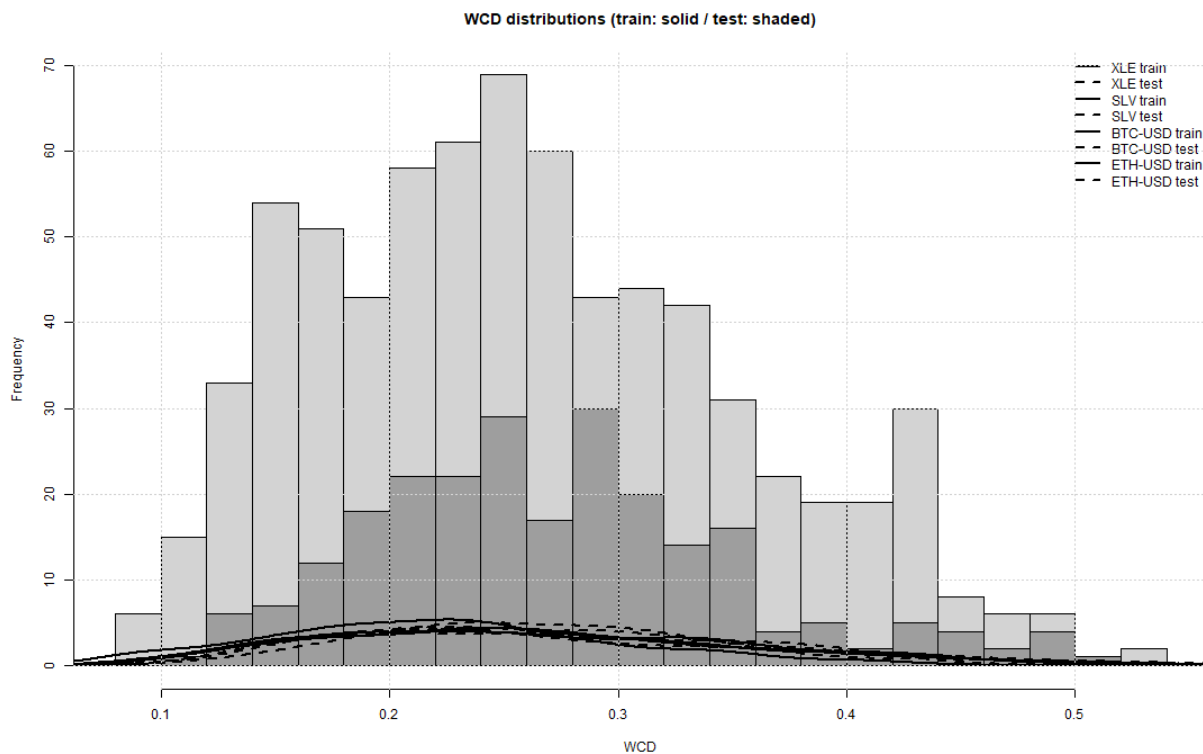


Figure 3. Wasserstein curve–depth distributions for all four assets: Train (solid) versus test (shaded). The strong overlap between train and test densities suggests regime stability at the FRV–curve level, with somewhat heavier upper tails for the cryptocurrency series in the test window.

6.3. Short–horizon targets and regime sensitivity

Short–run improvements are not confined to the first three days. Table 2 shows that DW–ridge attains the lowest average RMSE over both H1–3 and H1–5 for all four assets, and Table 4 confirms that, relative to ridge, the corresponding short–horizon MSE gains are positive with 95% moving–block bootstrap intervals that remain strictly above zero, especially for XLE, BTC–USD, and ETH–USD.

Table A.2 reports analogous gains relative to HAR, with the same qualitative pattern but somewhat weaker and, for SLV, occasionally borderline intervals. Taken together, these results indicate that, for both H1–3 and H1–5, DW–ridge systematically improves upon ridge and HAR, while ridge remains broadly comparable to, and sometimes marginally sharper than, the HAR baseline depending on the asset and horizon.

Regime–wise results, reported in Table A.3, further underscore this stability. For XLE and SLV, DW–ridge improves on ridge and HAR in both regimes, with larger gains in the more volatile recent period R2 (2023–2024). For the two cryptocurrencies, improvements are modest in R1 (2021–2022) but become pronounced in R2, where short–horizon MSE reductions relative to ridge and HAR are on the order of, and often exceed, 10% on both H1–3 and H1–5. These patterns confirm that depth reweighting is particularly advantageous when volatility regimes shift and extreme FRV trajectories become more prevalent.

6.4. Full RMSE profiles

Finally, Table A.5 reports horizon–wise RMSE values for HAR, ridge, and DW–ridge for each asset. These detailed profiles corroborate the graphical evidence in Figure 1: DW–ridge consistently yields lower RMSE than ridge at every horizon, with the largest absolute gains appearing at medium and long horizons for the ETF series and across the entire FRV path for BTC–USD and ETH–USD. Combined with the conformal and regime–wise analyses above, this indicates that DW–ridge offers a robust and interpretable improvement over standard HAR and ridge baselines across a wide range of volatility environments and evaluation metrics.

7. Conclusions

This paper has developed a shape–aware and regime–adaptive learner for forecasting FRV paths. The proposed DW–ridge estimator couples HAR–style regressors with ridge shrinkage, Wasserstein curve–depth reweighting, and an isotonic post–projection that enforces the natural monotonicity of cumulative variance. On four liquid assets—two sector ETFs (XLE, SLV) and two major cryptocurrencies (BTC–USD, ETH–USD) over 2020–2024—DW–ridge delivers uniformly lower RMSE profiles than unweighted ridge and systematically improves short–horizon errors on H1–3 and H1–5. Across all assets, DW–ridge consistently improves not only upon the monotone ridge specification but also upon a classical HAR–OLS benchmark iterated to FRV paths, with the largest relative gains appearing in the more volatile cryptocurrency series. For the ETFs, the resulting forecasting gains in FRV path accuracy are moderate but persistent; for BTC–USD and ETH–USD they reach the order of 10–13% on short–horizon aggregates. Block–conformal coverage remains close to or above nominal levels, indicating that sharper forecasts are not obtained at the expense of reliability. Taken together, these findings support the core claim that depth–adaptive weighting is an effective and robust mechanism for multi–horizon FRV prediction when combined with simple HAR–type regressors.

Although the exposition emphasizes XLE and SLV for transparency, the empirical results for BTC–USD and ETH–USD show that the same pipeline scales naturally to more volatile and heavy–tailed assets. For larger universes, two practical deployment modes are (i) *asset–separate* fits, where each asset has its own depth center and hyperparameters, and (ii) a *pooled* specification that augments HAR features with asset dummies and uses a common depth center when assets share macro or liquidity

regimes. Both variants preserve the computational lightness of DW–ridge and are compatible with the block–conformal diagnostics used here.

From a modeling perspective, the performance of DW–ridge is primarily driven by its regime–adaptive weighting: All learners share the same monotone PAV post–projection, so improvements over ridge and HAR arise from Wasserstein curve–depth weights rather than from altering the loss or architecture. In this sense, DW–ridge is complementary to more flexible non–linear sequence models, such as LSTMs and transformers trained on volatility covariates. Here we have deliberately focused on a transparent linear benchmark whose shrinkage, depth scores, and post–projection can be inspected and audited, a feature that is particularly attractive in risk–management settings. A systematic comparison against richer non–linear baselines—elastic net, kernel, or random–feature methods, and sequence models including LSTMs and transformers—remains an important direction for future work and would help quantify how much incremental accuracy such models can deliver beyond a regime–adaptive but structurally simple FRV learner.

Several extensions of the empirical evaluation are natural. Diebold–Mariano tests across horizons and loss differentials aggregated over H1–3 and H1–5 would complement the bootstrap intervals, while Mincer–Zarnowitz–type regressions on FRV increments could provide a more formal diagnosis of forecast calibration along the path. Such diagnostics would further quantify the margin by which depth–adaptive weighting improves upon alternative learners and help disentangle horizon–specific improvements from aggregate gains.

Short–horizon FRV improvements translate into actionable gains for volatility targeting and options risk management. Examples include: (i) Rebalancing rules that map FRV paths to target leverage with turnover penalties; (ii) delta–vega hedging where the path forecast informs near–term variance; and (iii) stress testing via depth–based conformal control charts that flag blocks with unusually shallow depth (large S_{\max}). Reporting both statistical accuracy (e.g., RMSE) and reliability (coverage) clarifies when forecasts are not only sharp but also trustworthy, and provides practitioners with a built–in diagnostic for identifying time segments where the FRV system should be treated with greater caution.

The present implementation also has limitations that suggest concrete methodological extensions. First, the choice of block length B and weight floor α_0 trades off variance and robustness; a data–driven selection that combines coverage loss and forecast risk could be developed. Second, our FRV target is based on squared daily returns; incorporating intraday realized measures could reduce measurement noise and sharpen short–run forecasts. Third, the depth computation relies on a fixed quantile grid; extending to learned grids or alternative curve–depth notions may increase sensitivity to shape anomalies. Finally, regime–switching centers and expanding–window recalibration may better accommodate structural change without sacrificing the guarantees afforded by block–conformal calibration.

Overall, a small number of ingredients—HAR features, ridge shrinkage, Wasserstein curve–depth reweighting, PAV monotonicity, and a block–conformal audit—suffice to obtain statistically and practically meaningful gains in FRV path prediction across both ETF and cryptocurrency markets. The approach is fast, transparent, and easily portable to larger cross–sections and richer covariate sets, making it a viable component in risk systems that demand both accuracy and documented reliability, and a useful benchmark against which more complex non–linear FRV forecasters can be evaluated.

Author contributions

Çağlar Sözen: Conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing – original draft, writing – review & editing; Fikriye Kabakcı: Conceptualization, investigation, methodology, project administration, resources, supervision, validation, writing – review & editing. All authors have read and approved the final version of the manuscript for publication.

Use of Generative-AI tools declaration

The authors used ChatGPT (OpenAI) only for language polishing.

Acknowledgments

This study has been supported by the Recep Tayyip Erdogan University Development Foundation (Grant number: 02025011017799).

Conflict of interest

The authors declare no competing interests.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. T. G. Andersen, T. Bollerslev, F. X. Diebold, P. Labys, Modeling and forecasting realized volatility, *Econometrica*, **71** (2003), 579–625. <https://doi.org/10.1111/1468-0262.00418>
2. O. E. Barndorff-Nielsen, N. Shephard, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, *J. Roy. Stat. Soc. Ser. B*, **64** (2002), 253–280. <https://doi.org/10.1111/1467-9868.00336>
3. R. E. Barlow, D. J. Bartholomew, J. M. Bremner, H. D. Brunk, *Statistical inference under order restrictions*, Wiley, 1972.
4. T. Robertson, F. T. Wright, R. L. Dykstra, *Order restricted statistical inference*, Wiley, 1988.
5. C. Villani, *Optimal transport: Old and new*, Springer, 2009. <https://doi.org/10.1007/978-3-540-71050-9>
6. R. Fraiman, G. Muniz, Trimmed means for functional data, *Test*, **10** (2001), 419–440. <https://doi.org/10.1007/BF02595706>
7. S. López-Pintado, J. Romo, On the concept of depth for functional data, *J. Am. Stat. Assoc.*, **104** (2009), 718–734. <https://doi.org/10.1198/jasa.2009.0108>

8. G. Peyré, M. Cuturi, Computational optimal transport, *Found. Trends Mach. Le.*, **11** (2019), 355–607. <https://doi.org/10.1561/22000000073>
9. M. Agueh, G. Carlier, Barycenters in the Wasserstein space, *SIAM J. Math. Anal.*, **43** (2011), 904–924. <https://doi.org/10.1137/100805741>
10. A. E. Hoerl, R. W. Kennard, Ridge regression: Biased estimation for nonorthogonal problems, *Technometrics*, **12** (1970), 55–67. <https://doi.org/10.1080/00401706.1970.10488634>
11. V. Vovk, A. Gammerman, G. Shafer, *Algorithmic learning in a random world*, Springer, 2005. <https://doi.org/10.1007/b106715>
12. Y. Romano, E. Patterson, E. J. Candès, Conformalized quantile regression, *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2019.
13. R. F. Barber, E. J. Candès, A. Ramdas, R. J. Tibshirani, Predictive inference with the jackknife+, *Ann. Stat.*, **49** (2021), 486–507. <https://doi.org/10.1214/20-AOS1965>
14. C. Xu, Y. Xie, Conformal prediction for time series, *IEEE T. Pattern Anal.*, **45** (2023), 11575–11587. <https://doi.org/10.1109/TPAMI.2023.3272339>
15. J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, L. Wasserman, Distribution-free predictive inference for regression, *J. Am. Stat. Assoc.*, **113** (2018), 1094–1111. <https://doi.org/10.1080/01621459.2017.1307116>
16. F. Corsi, A simple approximate long-memory model of realized volatility, *J. Financ. Economet.*, **7** (2009), 174–196. <https://doi.org/10.1093/jjfinec/nbp001>
17. T. Hastie, R. Tibshirani, J. Friedman, *The elements of statistical learning*, 2 Eds., Springer, 2009. <https://doi.org/10.1007/978-0-387-84858-7>
18. R. Tibshirani, Regression shrinkage and selection via the lasso, *J. Roy. Stat. Soc. Ser. B*, **58** (1996), 267–288. <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>
19. H. Zou, T. Hastie, Regularization and variable selection via the elastic net, *J. Roy. Stat. Soc. Ser. B*, **67** (2005), 301–320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>
20. Ç Sözen, Uniform one-sided conformal bands for forward realized volatility curves, *AIMS Math.*, **10** (2025), 27314–27337. <https://doi.org/10.3934/math.20251201>
21. Yahoo Finance, *Historical data for XLE, SLV, BTC–USD and ETH–USD (adjusted close)*, 2020–2024, Retrieved December 31, 2024. Available from: <https://finance.yahoo.com>.
22. C. Bergmeir, J. M. Benítez, On the use of cross-validation for time series predictor evaluation, *Inform. Sciences*, **191** (2012), 192–213. <https://doi.org/10.1016/j.ins.2011.12.028>

Appendix

A. Additional tables and figures

Table A.1. Train/cal/test splits and FRV origin dates by asset.

Asset	n_{total}	n_{train}	n_{cal}	n_{test}	Train start	Train end	Cal start	Cal end	Test start	Test end
XLE	1205	723	241	241	2020-02-04	2022-12-14	2022-12-15	2023-11-30	2023-12-01	2024-11-14
SLV	1205	723	241	241	2020-02-04	2022-12-14	2022-12-15	2023-11-30	2023-12-01	2024-11-14
BTC-USD	1205	723	241	241	2020-02-04	2022-12-14	2022-12-15	2023-11-30	2023-12-01	2024-11-14
ETH-USD	1205	723	241	241	2020-02-04	2022-12-14	2022-12-15	2023-11-30	2023-12-01	2024-11-14

Table A.2. Short-horizon MSE gains of DW-ridge relative to HAR with 95% moving-block bootstrap confidence intervals. Positive values favor DW-ridge.

Asset	Horizon	Gain vs. HAR	CI lower	CI upper	Percent gain vs. HAR
XLE	H1-3	0.0000000154577091	0.0000000102537000	0.00000002067624	4.46849005
SLV	H1-3	0.0000000002215179	-0.0000000146360460	0.00000001517674	0.02676845
BTC-USD	H1-3	0.0000016177084361	0.0000007120876320	0.00000241293405	10.02090055
ETH-USD	H1-3	0.0000051732088576	0.0000037094372400	0.00000638530777	9.86427365
XLE	H1-5	0.0000000550095789	0.0000000391955480	0.00000006942072	6.36146960
SLV	H1-5	0.0000000240681279	-0.0000000022288660	0.00000005167746	1.95469450
BTC-USD	H1-5	0.0000038304967377	0.0000020377776310	0.00000536245092	13.07706190
ETH-USD	H1-5	0.0000130337205488	0.0000095445948490	0.00001558778698	13.07791464

Table A.3. Regime-wise short-horizon MSE for ridge and DW-ridge, and percentage gains of DW-ridge relative to ridge and HAR over H1-3 and H1-5. MSE values are multiplied by 10^6 .

Asset	Regime	n_{origins}	H1-3				H1-5			
			MSE(R)	MSE(DW)	Gain(R) [%]	Gain(H) [%]	MSE(R)	MSE(DW)	Gain(R) [%]	Gain(H) [%]
XLE	R1 (2021-2022)	503	0.991608	0.974199	1.756	0.484	1.753586	1.710683	2.447	0.541
XLE	R2 (2023-2024)	471	0.401967	0.380881	5.246	3.472	0.923081	0.868306	5.934	5.249
SLV	R1 (2021-2022)	503	1.153804	1.137482	1.415	1.126	1.918496	1.876819	2.172	1.574
SLV	R2 (2023-2024)	471	0.687926	0.672099	2.301	0.351	1.046999	1.007708	3.753	2.432
BTC-USD	R1 (2021-2022)	503	46.523453	46.283192	0.516	0.591	73.717249	73.087559	0.854	1.214
BTC-USD	R2 (2023-2024)	471	18.057028	15.963647	11.593	11.516	33.053673	28.372587	14.162	14.629
ETH-USD	R1 (2021-2022)	503	174.353386	173.611286	0.426	1.285	291.046616	289.458685	0.546	1.737
ETH-USD	R2 (2023-2024)	471	44.722355	39.361054	11.988	13.729	90.780317	78.117087	13.949	16.627

Table A.4. Block-conformal coverage over block length B and miscoverage level α for all four assets. The deviation column reports coverage minus the nominal level $1 - \alpha$.

Asset	B	α	Blocks	q	Violations	Coverage (block)	Deviation
XLE	10	0.05	24	0.8757656	2	0.9166667	-0.033333333
XLE	10	0.10	24	0.8643944	4	0.8333333	-0.066666667
XLE	10	0.15	24	0.8461626	6	0.7500000	-0.100000000
XLE	10	0.20	24	0.8326792	10	0.5833333	-0.216666667
XLE	10	0.25	24	0.8305861	10	0.5833333	-0.166666667
XLE	14	0.05	17	0.8757656	2	0.8823529	-0.067647059
XLE	14	0.10	17	0.8757656	2	0.8823529	-0.017647059

Asset	B	α	Blocks	q	Violations	Coverage (block)	Deviation
XLE	14	0.15	17	0.8643944	4	0.7647059	-0.085294118
XLE	14	0.20	17	0.8405691	8	0.5294118	-0.270588235
XLE	14	0.25	17	0.8305861	9	0.4705882	-0.279411765
XLE	21	0.05	11	0.8757656	1	0.9090909	-0.040909091
XLE	21	0.10	11	0.8757656	1	0.9090909	0.009090909
XLE	21	0.15	11	0.8757656	1	0.9090909	0.059090909
XLE	21	0.20	11	0.8643944	2	0.8181818	0.018181818
XLE	21	0.25	11	0.8305861	6	0.4545455	-0.295454545
XLE	28	0.05	8	0.8757656	2	0.7500000	-0.200000000
XLE	28	0.10	8	0.8757656	2	0.7500000	-0.150000000
XLE	28	0.15	8	0.8757656	2	0.7500000	-0.100000000
XLE	28	0.20	8	0.8757656	2	0.7500000	-0.050000000
XLE	28	0.25	8	0.8643944	3	0.6250000	-0.125000000
XLE	35	0.05	6	0.8757656	2	0.6666667	-0.283333333
XLE	35	0.10	6	0.8757656	2	0.6666667	-0.233333333
XLE	35	0.15	6	0.8757656	2	0.6666667	-0.183333333
XLE	35	0.20	6	0.8757656	2	0.6666667	-0.133333333
XLE	35	0.25	6	0.8757656	2	0.6666667	-0.083333333
SLV	10	0.05	24	0.8941828	0	1.0000000	0.050000000
SLV	10	0.10	24	0.8925846	0	1.0000000	0.100000000
SLV	10	0.15	24	0.8820224	0	1.0000000	0.150000000
SLV	10	0.20	24	0.8764456	1	0.9583333	0.158333333
SLV	10	0.25	24	0.8661053	2	0.9166667	0.166666667
SLV	14	0.05	17	0.8941828	0	1.0000000	0.050000000
SLV	14	0.10	17	0.8941828	0	1.0000000	0.100000000
SLV	14	0.15	17	0.8925846	0	1.0000000	0.150000000
SLV	14	0.20	17	0.8783090	1	0.9411765	0.141176471
SLV	14	0.25	17	0.8764456	1	0.9411765	0.191176471
SLV	21	0.05	11	0.8941828	0	1.0000000	0.050000000
SLV	21	0.10	11	0.8941828	0	1.0000000	0.100000000
SLV	21	0.15	11	0.8941828	0	1.0000000	0.150000000
SLV	21	0.20	11	0.8925846	0	1.0000000	0.200000000
SLV	21	0.25	11	0.8783090	1	0.9090909	0.159090909
SLV	28	0.05	8	0.8941828	0	1.0000000	0.050000000
SLV	28	0.10	8	0.8941828	0	1.0000000	0.100000000
SLV	28	0.15	8	0.8941828	0	1.0000000	0.150000000
SLV	28	0.20	8	0.8941828	0	1.0000000	0.200000000
SLV	28	0.25	8	0.8925846	0	1.0000000	0.250000000
SLV	35	0.05	6	0.8941828	0	1.0000000	0.050000000
SLV	35	0.10	6	0.8941828	0	1.0000000	0.100000000
SLV	35	0.15	6	0.8941828	0	1.0000000	0.150000000
SLV	35	0.20	6	0.8941828	0	1.0000000	0.200000000
SLV	35	0.25	6	0.8941828	0	1.0000000	0.250000000
BTC-USD	10	0.05	24	0.8967464	0	1.0000000	0.050000000
BTC-USD	10	0.10	24	0.8855614	1	0.9583333	0.058333333
BTC-USD	10	0.15	24	0.8848417	1	0.9583333	0.108333333
BTC-USD	10	0.20	24	0.8648385	4	0.8333333	0.033333333
BTC-USD	10	0.25	24	0.8645189	4	0.8333333	0.083333333
BTC-USD	14	0.05	17	0.8967464	0	1.0000000	0.050000000
BTC-USD	14	0.10	17	0.8967464	0	1.0000000	0.100000000
BTC-USD	14	0.15	17	0.8855614	1	0.9411765	0.091176471
BTC-USD	14	0.20	17	0.8848417	1	0.9411765	0.141176471
BTC-USD	14	0.25	17	0.8754825	1	0.9411765	0.191176471
BTC-USD	21	0.05	11	0.8967464	0	1.0000000	0.050000000
BTC-USD	21	0.10	11	0.8967464	0	1.0000000	0.100000000
BTC-USD	21	0.15	11	0.8967464	0	1.0000000	0.150000000
BTC-USD	21	0.20	11	0.8855614	1	0.9090909	0.109090909
BTC-USD	21	0.25	11	0.8848417	1	0.9090909	0.159090909
BTC-USD	28	0.05	8	0.8967464	0	1.0000000	0.050000000
BTC-USD	28	0.10	8	0.8967464	0	1.0000000	0.100000000
BTC-USD	28	0.15	8	0.8967464	0	1.0000000	0.150000000
BTC-USD	28	0.20	8	0.8967464	0	1.0000000	0.200000000
BTC-USD	28	0.25	8	0.8855614	1	0.8750000	0.125000000
BTC-USD	35	0.05	6	0.8967464	0	1.0000000	0.050000000
BTC-USD	35	0.10	6	0.8967464	0	1.0000000	0.100000000
BTC-USD	35	0.15	6	0.8967464	0	1.0000000	0.150000000
BTC-USD	35	0.20	6	0.8967464	0	1.0000000	0.200000000

Asset	B	α	Blocks	q	Violations	Coverage (block)	Deviation
BTC-USD	35	0.25	6	0.8967464	0	1.0000000	0.25000000
ETH-USD	10	0.05	24	0.8842006	1	0.9583333	0.00833333
ETH-USD	10	0.10	24	0.8716358	3	0.8750000	-0.02500000
ETH-USD	10	0.15	24	0.8710856	3	0.8750000	0.02500000
ETH-USD	10	0.20	24	0.8569957	8	0.6666667	-0.13333333
ETH-USD	10	0.25	24	0.8563840	8	0.6666667	-0.08333333
ETH-USD	14	0.05	17	0.8842006	1	0.9411765	-0.00882352
ETH-USD	14	0.10	17	0.8842006	1	0.9411765	0.04117647
ETH-USD	14	0.15	17	0.8716358	4	0.7647059	-0.08529411
ETH-USD	14	0.20	17	0.8710856	4	0.7647059	-0.03529411
ETH-USD	14	0.25	17	0.8601574	7	0.5882353	-0.16176470
ETH-USD	21	0.05	11	0.8842006	1	0.9090909	-0.04090909
ETH-USD	21	0.10	11	0.8842006	1	0.9090909	0.00909090
ETH-USD	21	0.15	11	0.8842006	1	0.9090909	0.05909090
ETH-USD	21	0.20	11	0.8716358	3	0.7272727	-0.07272727
ETH-USD	21	0.25	11	0.8710856	3	0.7272727	-0.02727273
ETH-USD	28	0.05	8	0.8842006	1	0.8750000	-0.07500000
ETH-USD	28	0.10	8	0.8842006	1	0.8750000	-0.02500000
ETH-USD	28	0.15	8	0.8842006	1	0.8750000	0.02500000
ETH-USD	28	0.20	8	0.8842006	1	0.8750000	0.07500000
ETH-USD	28	0.25	8	0.8716358	3	0.6250000	-0.12500000
ETH-USD	35	0.05	6	0.8842006	1	0.8333333	-0.11666667
ETH-USD	35	0.10	6	0.8842006	1	0.8333333	-0.06666667
ETH-USD	35	0.15	6	0.8842006	1	0.8333333	-0.01666667
ETH-USD	35	0.20	6	0.8842006	1	0.8333333	0.03333333
ETH-USD	35	0.25	6	0.8842006	1	0.8333333	0.08333333

Table A.5. RMSE by horizon $h = 1, \dots, 30$ for HAR, ridge, and DW-ridge. For each asset, the first row also reports the short-horizon average gain of DW-ridge over ridge (H1-3) in absolute and percentage terms.

Asset	h	HAR	Ridge	DW-ridge	ShortAvg gain (abs)	ShortAvg gain (%)
XLE	1	0.0003068296	0.0003068295	0.0002992463	0.00001654241	2.8578044
XLE	2	0.0005415510	0.0005519549	0.0005361600		
XLE	3	0.0008064483	0.0008100805	0.0007838314		
XLE	4	0.0011029127	0.0011003883	0.0010637980		
XLE	5	0.0014385605	0.0014357301	0.0013876353		
XLE	6	0.0018223455	0.0018202959	0.0017571818		
XLE	7	0.0022227647	0.0021893533	0.0021121953		
XLE	8	0.0026524289	0.0026272850	0.0025355403		
XLE	9	0.0031064353	0.0030991265	0.0029930395		
XLE	10	0.0035843338	0.0035724927	0.0034558455		
XLE	11	0.0040818235	0.0040589275	0.0039344153		
XLE	12	0.0045947202	0.0046081332	0.0044708563		
XLE	13	0.0051225244	0.0051607114	0.0050109846		
XLE	14	0.0056659053	0.0057202979	0.0055548926		
XLE	15	0.0062213929	0.0062875107	0.0061069674		
XLE	16	0.0067889273	0.0068687986	0.0066729270		
XLE	17	0.0073681994	0.0074320696	0.0072211325		
XLE	18	0.0079553087	0.0080058582	0.0077796998		
XLE	19	0.0085508417	0.0086011038	0.0083580842		
XLE	20	0.0091540929	0.0091959660	0.0089334850		
XLE	21	0.0097634127	0.0097903427	0.0095066250		
XLE	22	0.0103808471	0.0103941737	0.0100893103		
XLE	23	0.0110026106	0.0109956382	0.0106672646		
XLE	24	0.0116285845	0.0115112151	0.0111640052		
XLE	25	0.0122584629	0.0120195408	0.0116531427		
XLE	26	0.0128930219	0.0125342053	0.0121500795		
XLE	27	0.0135314649	0.0130186884	0.0126157365		
XLE	28	0.0141741231	0.0134886319	0.0130675456		
XLE	29	0.0148202760	0.0139460618	0.0135065425		
XLE	30	0.0154736433	0.0144143295	0.0139543985		
SLV	1	0.0006540103	0.0006540094	0.0006523310	0.00000516315	0.5297094

Asset	h	HAR	Ridge	DW-ridge	ShortAvg gain (abs)	ShortAvg gain (%)
SLV	2	0.0009185868	0.0009286469	0.0009238057	0.00019348083	4.5520258
SLV	3	0.0011004857	0.0011057756	0.0010968056		
SLV	4	0.0012842403	0.0012775431	0.0012637910		
SLV	5	0.0014228916	0.0014184849	0.0013989459		
SLV	6	0.0015235331	0.0015313230	0.0015049014		
SLV	7	0.0016442855	0.0016664124	0.0016334753		
SLV	8	0.0017596620	0.0017834514	0.0017433860		
SLV	9	0.0019332854	0.0019521365	0.0019051774		
SLV	10	0.0021158913	0.0021258736	0.0020718077		
SLV	11	0.0022834780	0.0023034066	0.0022419715		
SLV	12	0.0024422020	0.0024679191	0.0023987416		
SLV	13	0.0025758793	0.0026072186	0.0025301258		
SLV	14	0.0027509175	0.0027877392	0.0027038761		
SLV	15	0.0029501841	0.0029936946	0.0029035505		
SLV	16	0.0031247443	0.0031774211	0.0030810530		
SLV	17	0.0032824928	0.0033428330	0.0032396198		
SLV	18	0.0034386110	0.0034982190	0.0033878746		
SLV	19	0.0035914364	0.0036490892	0.0035319871		
SLV	20	0.0037474878	0.0038070320	0.0036832166		
SLV	21	0.0038984931	0.0039557949	0.0038242384		
SLV	22	0.0040481370	0.0040998404	0.0039592096		
SLV	23	0.0041949237	0.0042366795	0.0040867379		
SLV	24	0.0043421297	0.0043810451	0.0042215009		
SLV	25	0.0044761232	0.0045109831	0.0043402396		
SLV	26	0.0046128345	0.0046439828	0.0044616426		
SLV	27	0.0047526463	0.0047721217	0.0045776727		
SLV	28	0.0049008691	0.0049087742	0.0047038858		
SLV	29	0.0050367237	0.0050175420	0.0048031816		
SLV	30	0.0051561689	0.0051167850	0.0048927710		
BTC-USD	1	0.0026960195	0.0026960098	0.0026286759	0.00029613560	3.9278887
BTC-USD	2	0.0039827023	0.0039927679	0.0038058551		
BTC-USD	3	0.0050298701	0.0050360192	0.0047098234		
BTC-USD	4	0.0063316937	0.0063382730	0.0058845525		
BTC-USD	5	0.0076117172	0.0075722677	0.0070072296		
BTC-USD	6	0.0088222184	0.0087656114	0.0080785054		
BTC-USD	7	0.0099413678	0.0098537118	0.0090585415		
BTC-USD	8	0.0110109914	0.0108991878	0.0100081945		
BTC-USD	9	0.0120425692	0.0119219258	0.0109537138		
BTC-USD	10	0.0130694744	0.0129369219	0.0119158550		
BTC-USD	11	0.0141066188	0.0139582765	0.0128973064		
BTC-USD	12	0.0151465699	0.0149759774	0.0139201640		
BTC-USD	13	0.0161719682	0.0160004223	0.0149594354		
BTC-USD	14	0.0171717932	0.0169980917	0.0159901079		
BTC-USD	15	0.0182012089	0.0180240755	0.0170348388		
BTC-USD	16	0.0192754588	0.0191006298	0.0181464010		
BTC-USD	17	0.0203285436	0.0201804705	0.0192509335		
BTC-USD	18	0.0213486943	0.0212411953	0.0203144191		
BTC-USD	19	0.0223800778	0.0223053755	0.0213726715		
BTC-USD	20	0.0233846674	0.0233237686	0.0223311384		
BTC-USD	21	0.0244001416	0.0243592200	0.0232954850		
BTC-USD	22	0.0254119423	0.0253962804	0.0242397123		
BTC-USD	23	0.0263874646	0.0263970523	0.0250998475		
BTC-USD	24	0.0273485667	0.0273559885	0.0259213875		
BTC-USD	25	0.0283261678	0.0283626743	0.0267770266		
BTC-USD	26	0.0292867667	0.0293389044	0.0275966316		
BTC-USD	27	0.0302391919	0.0299878320	0.0281746121		
BTC-USD	28	0.0311861032	0.0306174985	0.0287205347		
BTC-USD	29	0.0321463282	0.0312573803	0.0292684227		
BTC-USD	30	0.0331314439	0.0319203273	0.0298155966		
ETH-USD	1	0.0045961143	0.0045960922	0.0044856315	0.00029613560	3.9278887
ETH-USD	2	0.0070916173	0.0070778060	0.0067967598		
ETH-USD	3	0.0092691083	0.0091856979	0.0086887980		
ETH-USD	4	0.0117936963	0.0116644891	0.0109586617		
ETH-USD	5	0.0142087089	0.0139934935	0.0130857724		
ETH-USD	6	0.0165931443	0.0162911315	0.0151753260		
ETH-USD	7	0.0189241902	0.0184831877	0.0171810987		
ETH-USD	8	0.0212360475	0.0206462764	0.0191866843		

Asset	h	HAR	Ridge	DW-ridge	ShortAvg gain (abs)	ShortAvg gain (%)
ETH-USD	9	0.0235280552	0.0227837400	0.0211979573		
ETH-USD	10	0.0257954350	0.0248939259	0.0231912427		
ETH-USD	11	0.0280371820	0.0270057995	0.0251904688		
ETH-USD	12	0.0302688191	0.0291099088	0.0272176442		
ETH-USD	13	0.0324748440	0.0312523459	0.0292744871		
ETH-USD	14	0.0346856093	0.0334223497	0.0313532483		
ETH-USD	15	0.0368998870	0.0355902570	0.0334100902		
ETH-USD	16	0.0391628863	0.0377921352	0.0355096031		
ETH-USD	17	0.0414511103	0.0400442805	0.0376414995		
ETH-USD	18	0.0437298154	0.0423744875	0.0398257299		
ETH-USD	19	0.0459904870	0.0446450200	0.0419265797		
ETH-USD	20	0.0482202133	0.0468631084	0.0439299256		
ETH-USD	21	0.0504607257	0.0491742865	0.0460078752		
ETH-USD	22	0.0527322516	0.0515289175	0.0481310597		
ETH-USD	23	0.0549695006	0.0538864407	0.0502145045		
ETH-USD	24	0.0571925674	0.0562662155	0.0523098285		
ETH-USD	25	0.0594275769	0.0587424706	0.0544567340		
ETH-USD	26	0.0616430826	0.0611849280	0.0565585448		
ETH-USD	27	0.0639043805	0.0631594294	0.0582782865		
ETH-USD	28	0.0661349140	0.0650892661	0.0599424282		
ETH-USD	29	0.0683505653	0.0669586169	0.0615217794		
ETH-USD	30	0.0705971455	0.0688659463	0.0631116007		



AIMS Press

©2025 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)