

State-Dependent Behavioral Hedging: Integrating Herding Regimes into Quantitative Risk Management

Abstract

Purpose – This study proposes a behavioural dynamic hedging framework to address two limitations of the conventional optimal hedge ratio (OHR) literature: the practical friction costs of frequent hedge rebalancing and the limited incorporation of behavioural information into hedging decisions.

Design/methodology/approach – Using VN30 index futures (2019-2025), the study combines a recursive Markov-switching CSAD model for out-of-sample herding detection with a VECM-Robust EWMA volatility system. Two behavioral strategies – Conditional Hedging (OHR_on/off) and Enhanced Behavioral Hedging (OHR_alpha) – are developed and estimated under a strict expanding-window framework to eliminate look-ahead bias.

Findings – The Conditional Hedging Strategy (OHR_on/off) substantially reduces hedge rebalancing frequency while achieving the highest certainty-equivalent return. The Enhanced Behavioral Hedging Strategy (OHR_alpha) improves risk-adjusted performance by adaptively incorporating a behavioural risk premium during periods of elevated market stress.

Originality – This study extends the conventional OHR literature by reframing hedging as an adaptive decision framework rather than solely a variance-minimisation problem. It demonstrates how behavioural state information can be operationalised within quantitative hedge-ratio estimation to adapt hedging intensity to changing market conditions and investor objectives.

Keywords Dynamic hedging; Behavioural finance; Herding; Robust EWMA; Markov-switching; Hedge ratio; Transaction costs.

Paper type Research Paper

1. Introduction

Traditional portfolio management relies heavily on Markowitz's Modern Portfolio Theory (MPT), which optimizes the mean-variance efficient frontier under the assumptions of investor rationality, normally distributed returns, and constant correlations (Fabozzi et al., 2002). However, these static assumptions often fail to capture the dynamics of modern financial markets (Cao et al., 2023). Evidence from behavioral finance reveals that cognitive and emotional biases frequently drive phenomena such as herding and market overreaction (Hirshleifer and Teoh, 2003). Furthermore, empirical studies indicate that asset returns frequently exhibit fat-tailed distributions (Harris and Shen, 2003; Gong et al., 2019) and time-varying dependence (Ang and Chen, 2002). Consequently, modeling risk based on idealized, static assumptions degrades hedging effectiveness, highlighting the need for frameworks that adapt to fluctuating market conditions.

This need for adaptive hedging is particularly acute in frontier markets like Vietnam, where limited derivative instruments (e.g., only stock index futures) exacerbate market inefficiencies and behavioral biases (Chang et al., 2000). While the estimation of the optimal hedge ratio (OHR) – defined economically as the exact proportion of a spot position that must be covered by futures to minimize overall portfolio risk – has seen significant methodological advances – primarily through the continuous development of increasingly sophisticated econometric models (Brooks et al., 2002; Yang and Allen, 2005) – the literature remains heavily fragmented. Individual studies typically address only a single dimension of market behavior, such as fat-tailed return distributions or long-run cointegration relationships. Consequently, they fail to simultaneously incorporate other critical features, notably market asymmetry, into a unified

analytical framework. Because of these fragmented advances, two major practical challenges remain unresolved. First is the long-recognized trade-off between hedging effectiveness and the transaction costs arising from frequent portfolio rebalancing (Park and Switzer, 1995; Hou and Li, 2013). Second is the necessity of asymmetric downside protection, as loss-aversion dictates that losses exert a stronger influence on investors' decisions than gains, consistent with Prospect Theory (Kahneman and Tversky, 1992) and asymmetric market behavior (Ang and Chen, 2002). Furthermore, much of the existing literature lacks the rigorous out-of-sample validation needed to prevent model overfitting and accurately reflect real-world decision-making.

Against this background, this study proposes an Adaptive Hedging Decision Framework that shifts the focus from merely estimating “how much to hedge” to supporting dynamic hedging decisions under changing market conditions. To overcome the aforementioned fragmentation, the proposed framework adopts an integrated architecture combining VECM and Robust EWMA to capture long-run cointegration and fat-tailed distributions, respectively, while a recursive Markov-switching CSAD (MRS-CSAD) model serves as a market-state identification mechanism to capture behavioral asymmetry. Within this framework, two adaptive strategies are developed: the Conditional Hedging Strategy (OHR_on/off) to optimize when to hedge (addressing transaction frictions), and the Enhanced Behavioral Hedging Strategy (OHR_alpha) to determine how much downside protection is required during market distress (addressing Prospect Theory).

Using VN30 index futures 2019-2025, the proposed framework is evaluated under a strictly recursive out-of-sample expanding-window design to eliminate look-ahead bias. The empirical results demonstrate that OHR_on/off substantially reduces rebalancing frequency while delivering superior certainty-equivalent returns. Meanwhile, OHR_alpha effectively mitigates extreme portfolio drawdowns during market distress. These findings confirm that behavioral state information can significantly improve the economic efficiency of hedging decisions.

The remainder of the paper proceeds as follows: Section 2 reviews the literature; Section 3 details the methodology; Section 4 presents empirical results; Section 5 discusses economic implications; and Section 6 concludes.

2. Literature review

2.1. The Variance-Minimization Paradigm and Econometric Fragmentation

The theoretical foundation of the optimal hedge ratio (OHR) originates from the minimum-variance framework independently developed by Johnson (1960) and formalized by Ederington (1979). By taking the first derivative of a hedged portfolio's return variance and setting it to zero, Ederington (1979) demonstrated that the risk-minimizing hedge ratio is mathematically defined as the ratio of the covariance between spot and futures returns (σ_{sf}) to the variance of futures returns (σ_f^2):

$$OHR = \frac{\sigma_{sf}}{\sigma_f^2}$$

Early applications predominantly estimated the OHR using the ordinary least squares (OLS) regression model, which restrictively assumes a constant variance-covariance structure throughout the investment horizon. Recognizing that market volatilities evolve continuously, subsequent literature transitioned away from static OLS toward dynamic time-series models, predominantly utilizing the GARCH family and its multivariate extensions (e.g., MGARCH, BGARCH, DCC-GARCH) (Brooks et al., 2002; Yang and Allen, 2005; Hou and Li, 2013; Qu et al., 2019), as well as copula integrations (Louhichi and Rais, 2019). Yet, despite transitioning from static to dynamic estimation, both approaches remain

fundamentally anchored to the variance-minimization paradigm. By structurally treating risk as symmetric, this paradigm implicitly equates upside gains with downside losses, failing to accommodate asymmetric risk preferences and inadvertently dampening a portfolio's capacity to fully absorb upward market trends.

Furthermore, the econometric treatment of market characteristics within these dynamic models remains heavily fragmented. One prominent research stream emphasizes robustness against non-normal, heavy-tailed return distributions, employing estimators such as the Robust Exponentially Weighted Moving Average (Robust EWMA) (Harris and Shen, 2003) and skewed distributions (Gong et al., 2019). Conversely, a parallel stream underscores the necessity of anchoring short-run price dynamics to long-run cointegration relationships via Error Correction Models (ECM) and Vector Error Correction Models (VECM) (Lien and Luo, 1994; Lien, 1996; Park and Switzer, 1995; Yang and Allen, 2005; Hao and Nhung, 2023). Consequently, the existing literature lacks a hybrid econometric architecture capable of jointly incorporating long-run equilibrium properties and distributional robustness.

2.2. The Validation Deficit and Out-of-Sample Evaluation

Beyond econometric fragmentation, the practical applicability of dynamic hedging strategies is constrained by the design of out-of-sample validation procedures. The majority of existing studies evaluate hedging performance exclusively under in-sample settings, which fails to replicate real-time decision-making environments. Among those employing out-of-sample validation, the split-sample framework effectively freezes estimated parameters, implicitly assuming that market characteristics remain static post-estimation (Park and Switzer, 1995; Hou and Li, 2013). Alternatively, the rolling-window framework continuously updates parameters but introduces delayed responses to structural changes: newly arriving information is averaged with older observations, while historical crisis observations continue to influence estimates until they mechanically exit the window, potentially leading to unnecessarily conservative positions (Harris and Shen, 2003; Qu et al., 2019). Furthermore, the literature lacks a consensus on the appropriate window length, as evidenced by the varied applications in existing studies (Harris and Shen, 2003; Qu et al., 2019).

To address the limitations of parameter freezing and delayed information adjustment, evaluating hedging decisions requires a strictly recursive expanding-window framework combined with exponentially decaying weights. This design preserves the statistical benefits of a large estimation sample while progressively reducing the influence of distant historical observations, allowing the model to adapt to newly emerging information without entirely discarding historical market memory.

2.3. Behavioral State Integration and Hypotheses Development

The continuous parameter updates inherent in dynamic time-series models produce frequently changing hedge ratios, inherently demanding continuous portfolio rebalancing to track the optimal position. This mechanically continuous rebalancing incurs substantial transaction costs that severely deteriorate the net economic utility of hedging (Park and Switzer, 1995; Hou and Li, 2013). Addressing this structural friction requires a shift from schedule-based adjustments to state-dependent decision mechanisms that explicitly determine when hedging is economically justified.

Behavioral finance posits that investor herding – where market participants disregard private information to follow the actions of the majority – amplifies market instability and intensifies downside risk, particularly in emerging and frontier markets (Chang et al., 2000; Nguyen and Vo, 2024). Chang et al. (2000) operationalized this phenomenon through the cross-sectional absolute deviation (CSAD) framework, quantifying the presence of collective irrationality. To enhance the practical usefulness of this behavioral information, Bohl et al. (2016) extended the CSAD framework by incorporating Markov regime-switching dynamics (MRS-CSAD) to identify distinct herding regimes. Nevertheless, their analysis

remained confined to an in-sample setting, rendering the framework a descriptive tool of historical market behavior rather than a predictive mechanism capable of supporting real-time hedging decisions.

To operationalize behavioral information, this study constructs a unified out-of-sample predicting filter utilizing the expanding-window MRS-CSAD framework. Rather than functioning merely as an ex post metric, the predicted probability of a panic-herding regime serves as the common decision signal supporting two independent tactical strategies. First, to address the friction of continuous rebalancing, the Conditional Hedging Strategy (OHR_on/off) utilizes the predicted probability to conditionally activate the integrated VECM-Robust EWMA hedge ratio. By initiating hedging solely during periods of elevated behavioral risk, OHR_on/off suspends hedging during stable market conditions, thereby mitigating the friction of unnecessary trading and preserving upside portfolio returns.

Hypothesis 1 (H1). The Conditional Hedging Strategy (OHR_on/off), driven by out-of-sample herding-state predictions, delivers superior trading efficiency compared to conventional continuously rebalanced dynamic strategies by preserving upside portfolio returns during stable market conditions.

While OHR_on/off dictates the timing of hedge activation, it does not prescribe the optimal intensity required to manage asymmetric downside risk. Financial returns exhibit asymmetric dependence structures, with asset correlations increasing substantially during market downturns (Ang and Chen, 2002). Furthermore, Prospect Theory dictates that loss-averse investors penalize portfolio drawdowns more heavily than they value equivalent gains (Kahneman and Tversky, 1992). Conventional variance-minimization strategies fail to incorporate these asymmetric preferences. Therefore, the Enhanced Behavioral Hedging Strategy (OHR_alpha) leverages the same out-of-sample MRS-CSAD filter to adaptively scale hedge intensity. By integrating a behavioral premium directly proportional to the predicted probability of panic, OHR_alpha aggressively protects the portfolio during severe distress, aligning the hedging decision with the loss-averse objectives of investors.

Hypothesis 2 (H2). The Enhanced Behavioral Hedging Strategy (OHR_alpha), which incorporates an adaptive behavioral premium along with predicted panic probability, provides superior risk-adjusted portfolio performance during periods of elevated behavioral risk while better accommodating the loss-averse preferences implied by Prospect Theory.

3. Data and Methodology

3.1. Data Properties and Dynamic Portfolio Simulation

This study utilizes daily closing prices of the VN30 Index and its continuous one-month futures contract (VN30F1M) from 2 January 2019 to 31 December 2025, sourced from the FiinProX database. To construct the cross-sectional absolute deviation (CSAD) behavioral indicator without introducing survivorship bias, a dynamic portfolio simulation algorithm was implemented in the R environment. Leveraging the complete historical price series of all 46 stocks that appeared in the index during the sample period, the algorithm systematically processes official semiannual revision announcements to reconstruct the exact 30-stock constituent basket active on each specific trading day.

Table 1 summarizes the statistical properties of the daily spot returns, futures returns, and the dynamically computed CSAD series.

Table 1. Descriptive statistics for the Rate of Return and CSAD series

Variables	Mean	Std. Dev	Skewness	Kurtosis	Jarque-Bera (JB_Stat)	P- value
R_s	0.049	1.307	-0.892	7.935	2006.196	0
R_f	0.051	1.490	-0.93	10.23	4060.52	0
CSAD	1.13	0.44	1.32	5.33	888.24	0

Notes: R_s and R_f denote continuously compounded daily returns computed as $\ln(P_t/P_{t-1}) \times 100$. A reported p-value of 0.00 indicates statistical significance at the 1% level.

Both markets exhibit comparable mean returns, indicating a stable long-run co-movement, while the futures market displays higher volatility, reflecting its role in more rapidly incorporating new information. Crucially, both return series exhibit significant negative skewness, aligning with the asymmetric market dependence documented by [Ang and Chen \(2002\)](#), whereby downside market movements tend to be considerably more severe than upside fluctuations. Regarding the behavioral indicator, the CSAD series yields a mean of 1.13 with a positive skewness of 1.32 and an excess kurtosis of 5.33. These characteristics indicate an asymmetric, right-skewed distribution with heavy tails, demonstrating that cross-sectional dispersion frequently experiences extreme positive deviations. Such distributional properties suggest that standard linear frameworks may not fully capture the behavior of CSAD, thereby motivating the exploration of non-linear models such as the MRS-CSAD. Furthermore, the excess kurtosis across all variables, combined with the highly significant Jarque-Bera statistics, decisively rejects the normality assumption. This pronounced heavy-tailed and non-Gaussian behavior empirically justifies the subsequent implementation of the Robust EWMA framework ([Harris and Shen, 2003](#)), which is specifically engineered to provide stable volatility estimates under extreme market shocks.

3.2. Two-Stage Strictly Out-of-Sample Research Design

To preclude look-ahead bias, the full dataset is asymmetrically partitioned into an initial estimation window of 250 trading days and an out-of-sample evaluation period of 1476 days. The initial 250-day window is utilized to calibrate the baseline architectures prior to forecasting: it establishes the cointegrating vector and lag structure for the VECM, and initializes the MRS-CSAD framework by setting the starting transition matrix and identifying the baseline panic-herding regime (empirically defined as the state exhibiting the most negative γ_2 coefficient – more details at [3.4](#)).

Following this initialization, the study employs a recursive two-stage forecasting procedure. In the first stage, both models are re-estimated iteratively across the 1476 out-of-sample days. For any given day t , the parameter estimation utilizes an expanding information set spanning strictly from the first observation up to day $t - 1$. This generates true one-step-ahead outputs for day t : the VECM extracts the out-of-sample market innovations, while the MRS-CSAD yields the predicted probability of transitioning into the panic regime. In the second stage, these ex-ante outputs of VECM feed directly into the Robust EWMA model to dynamically update the conditional variance-covariance matrix for day t . Consequently, all tactical hedging positions are executed sequentially, relying exclusively on information available at the time of the decision.

3.3. Cointegrated VECM and Baseline Dynamic Hedge Ratio

To establish the econometric foundation for the dynamic hedging framework, the integration properties of the logarithmic spot ($\ln S_t$) and futures ($\ln F_t$) price series are first evaluated using the Augmented Dickey-Fuller (ADF) test. Upon confirming that the series are integrated of the same order, an

unrestricted Vector Autoregression (VAR) system is constructed to determine the optimal lag length via standard information criteria (AIC, HQ, FPE) and to verify system stability. Subsequently, the Johansen trace test is applied to determine the existence and rank of the long-run equilibrium relationship.

Conditional on the evidence of cointegration, a Vector Error Correction Model (VECM) is employed to capture both the long-run equilibrium and the short-run adjustment dynamics. Let $\Delta Y_t = [R_{s,t}, R_{f,t}]'$ denote the vector of observed spot and futures returns; the VECM is specified as:

$$\Delta Y_t = \Pi Y_{t-1} + \Gamma_1 \Delta Y_{t-1} + \varepsilon_t \quad (1)$$

where Π is the long-run impact matrix and Γ_1 is the coefficient matrix associated with the short-run dynamics. Through the recursive expanding-window procedure (Stage 1), the VECM extracts genuine one-step-ahead forecast innovations for the spot ($\varepsilon_{spot,t}$) and futures ($\varepsilon_{fut,t}$) markets.

To appropriately specify the robust volatility estimator empirically rather than imposing it a priori, these out-of-sample innovations are pooled to assess their underlying distribution. The pooled sample is evaluated using descriptive statistics (skewness and kurtosis) and the Jarque-Bera normality test. Maximum-likelihood estimation is then performed under both Gaussian and Laplace distributional assumptions, with the associated Akaike Information Criterion (AIC) values compared to identify the best empirical fit and determine the specific geometric scaling constant ($g(k)$).

Following this distributional identification, the time-varying volatility is estimated using the Robust Exponentially Weighted Moving Average (Robust EWMA) framework of [Harris and Shen \(2003\)](#):

$$\sigma_t^k = \lambda \sigma_{t-1}^k + (1 - \lambda) g(k) |\varepsilon_{t-1}|^k \quad (2)$$

where ε_{t-1} denotes the one-step-ahead forecast innovation extracted from the VECM, and $\lambda = 0.94$ represents the standard RiskMetrics decay factor. Crucially, setting $k = 1$ specifies first-order conditional absolute deviations (σ_t) rather than traditional squared variances ($k = 2$), structurally dampening the influence of extreme outliers inherent in heavy-tailed financial returns. The geometric scaling constant $g(k)$ adjusts to the empirical error distribution: $g(1) = \sqrt{\pi/2} \approx 1.2533$ under Gaussianity, and $g(1) = \sqrt{2} \approx 1.4142$ under the heavy-tailed Laplace specification. The recursive estimation is initialized using the sample standard deviation of the first 30 out-of-sample innovations.

The corresponding dynamic covariance between spot (s) and futures (f) innovations, $\sigma_{sf,t}$, is subsequently recovered via the polarization identity applied to the smoothed volatility estimates of the sum and difference innovation processes:

$$\sigma_{sf,t} = \frac{1}{4} (\sigma_{(s+f),t}^2 - \sigma_{(s-f),t}^2) \quad (3)$$

where $\sigma_{(s+f),t}^2$ and $\sigma_{(s-f),t}^2$ represent the dynamic variances derived from Equation (2) for the summed innovation series ($\varepsilon_{s,t} + \varepsilon_{f,t}$) and the differenced innovation series ($\varepsilon_{s,t} - \varepsilon_{f,t}$), respectively.

Finally, the baseline optimal hedge ratio is derived as the ratio between the estimated conditional covariance and the conditional variance of the futures market:

$$OHR_{base,t} = \frac{\sigma_{sf,t}}{\sigma_{f,t}^2} \quad (4)$$

3.4. Market State Prediction and Active Behavioral Hedging Strategies

To extract the real-time behavioral signal driving the tactical hedging decisions, the Cross-Sectional Absolute Deviation (CSAD) is computed daily using the dynamic 30-stock VN30 basket:

$$CSAD_t = \frac{1}{30} \sum_{i=1}^{30} |R_{i,t} - R_{m,t}| \quad (5)$$

where $R_{i,t}$ is the return of constituent stock i , and $R_{m,t}$ is the equally weighted market return. Rather than employing a static regression, this study captures the time-varying nature of herding behavior by embedding the non-linear CSAD specification of [Chang et al. \(2000\)](#) within a two-state Markov Regime-Switching (MRS) framework:

$$CSAD_t = \gamma_{0,S_t} + \gamma_{1,S_t} |R_{m,t}| + \gamma_{2,S_t} R_{m,t}^2 + \varepsilon_t \quad (6)$$

where $S_t \in \{1,2\}$ denotes the unobservable market regime. To resolve the label-switching problem during recursive estimation, an objective regime-identification rule is applied at every iteration: the latent state exhibiting the more negative quadratic coefficient (γ_2) is consistently classified as the panic-herding regime, reflecting stronger return compression during large market movements. To ensure numerical stability, each estimation iteration is executed from five independent random initializations, retaining the converged solution with the highest log-likelihood. If convergence fails, the most recently converged parameter vector is temporarily carried forward. This expanding-window filter ultimately outputs the one-step-ahead predicted probability of the panic-herding regime, $p_{panic,t}$, serving as the out-of-sample behavioral signal.

Building upon the baseline hedge ratio ($OHR_{base,t}$), this behavioral signal explicitly dictates two independent active hedging strategies.

The Conditional State-Switching Strategy (OHR_on/off) addresses transaction-cost inefficiency by establishing an intentional no-trading region. Hedging is exclusively activated when the predicted panic probability exceeds the Maximum A Posteriori (MAP) classification threshold of 0.5:

$$OHR_{on/off,t} = OHR_{base,t} \times D_t \quad (7)$$

where $D_t = 1$ if $p_{panic,t} > 0.5$, and 0 otherwise. This binary mechanism suspends hedging during normal conditions, mitigating unnecessary trading activities.

The Enhanced Behavioral Hedging Strategy (OHR_alpha) mitigates downside tail risk by adaptively scaling the hedge intensity in continuous proportion to the predicted market panic:

$$OHR_{\alpha,t} = OHR_{base,t} (1 + \alpha_t \times p_{panic,t}) \quad (8)$$

To ensure the adjustment solely strengthens (never weakens) the baseline position, the adaptive behavioral premium, α_t , is constrained to be non-negative:

$$\alpha_t = \max(0, x_t)$$

where x_t quantifies the historical relative difference in hedging intensity required during past panic versus normal periods:

$$x_t = \frac{Mean(OHR_{base,t} | D_t = 1) - Mean(OHR_{base,t} | D_t = 0)}{Mean(OHR_{base,t} | D_t = 0)} \quad t \in [1, \dots, t-1] \quad (9)$$

Both regime-specific averages are computed exclusively from the expanding historical information set available up to day $t-1$. As a conservative numerical safeguard, this adaptive learning mechanism

remains strictly inactive ($\alpha_t = 0$) until a minimum threshold of 30 historical panic-state observations is accumulated, preventing unstable regime-specific estimates during the early stages of the out-of-sample evaluation. It also remains inactive if the normal-regime benchmark is numerically close to zero, avoiding ill-conditioned updates.

3.5. Multidimensional Performance Evaluation Framework

To provide a comprehensive assessment of the proposed framework, the out-of-sample performance of the unhedged portfolio and the three dynamic hedging strategies (OHR_base, OHR_on/off, and OHR_alpha) is evaluated using six complementary metrics that jointly capture risk reduction, investment performance, trading efficiency, and investor utility.

Portfolio variance, $Var(R_p)$, serves as the primary indicator of total residual risk. The proportional risk reduction relative to the unhedged benchmark is quantified by Hedging Effectiveness (HE):

$$HE = 1 - \frac{Var(R_{hedged})}{Var(R_{unhedged})} \quad (10)$$

Investment performance is captured by the average out-of-sample daily portfolio return (R_p), while risk-adjusted performance is evaluated using the daily Sharpe Ratio, assuming a zero risk-free rate.

To explicitly evaluate the economic benefit of reducing unnecessary hedge activation, a Trading Efficiency metric is introduced:

$$Efficiency = \frac{HE}{D} \quad (11)$$

where D represents the proportion of the $T_n = 1476$ out-of-sample trading days on which an active hedge position ($OHR_t > 0$) is maintained. For continuously hedged portfolios (OHR_base and OHR_alpha), $D = 1$, rendering this metric equivalent to standard HE. Conversely, for the conditional strategy (OHR_on/off), $D < 1$, thereby quantifying the variance reduction achieved per unit of active hedge participation.

The economic value of each strategy is further evaluated using the Certainty Equivalent Return (CER) under a mean-variance utility framework:

$$CER = \bar{R}_p - \frac{\gamma}{2} Var(R_p) = \bar{R}_p - 2Var(R_p) \quad (12)$$

where the coefficient of relative risk aversion is fixed at $\gamma = 4$, representing a moderately risk-averse institutional investor.

Finally, cumulative portfolio returns are plotted to provide an intuitive visual comparison of long-run wealth evolution across the competing strategies throughout the evaluation period.

4. Empirical Results and Hedging Performance

This section evaluates the out-of-sample performance of the proposed behavioral hedging strategies (OHR_on/off and OHR_alpha) against the continuous baseline model (OHR_base). To maintain focus on the economic implications of the hedging frameworks, all preliminary econometric diagnostics – including unit root tests, Johansen cointegration analysis, VECM specification, and residual distributional checks – are fully reported in Appendices A and B. These diagnostics confirm that the price series share a single long-run equilibrium and that the out-of-sample innovations exhibit heavy-tailed Laplace distributions, thereby justifying the employment of the VECM-Robust EWMA framework.

4.1. Herding Regime Dynamics and Market States

To implement the state-dependent hedging rules, the recursive out-of-sample MRS-CSAD model is utilized to classify trading days into either a normal state or a high-volatility herding regime. Following standard regime-switching literature, a trading day t is classified into the herding regime if the predicted probability exceeds the 0.5 threshold.

Figure 1 illustrates the dynamic evolution of this extracted probability signal. The series demonstrates highly responsive behavior, characterized by abrupt structural shifts that rapidly cross the 0.5 threshold during market turbulence. Once crossed, the probabilities do not merely oscillate erratically but form persistent clusters near the upper boundary (approaching 1.0). This structural clustering confirms that the model captures distinct, persistent behavioral regimes rather than transient daily noise. As supplementary validation, these intense probability clusters temporally align with well-documented periods of macroeconomic distress, initiating with the COVID-19 shock in early 2020 and subsequently spiking during the domestic liquidity crunch in the fourth quarter of 2022.




Figure 1. Recursive One-Step-Ahead Herding Probabilities

The statistical robustness of this classification is corroborated by the regime transition matrix reported in Table 2. The transition probabilities reveal substantial persistence in both the normal state ($P_{11}=0.9298$) and the herding state ($P_{22}=0.8925$). Applying Hamilton's expected duration framework, these probabilities indicate that normal market conditions last for approximately 14.2 consecutive days, while herding regimes persist for an average of 9.3 days.

Table 2. Out-of-Sample Probability Evaluation Metrics

	P_{11}	P_{22}
Probability	0.9298	0.8925

Aggregating these daily classifications over the entire out-of-sample evaluation window, the market operates within the herding regime for 39.05% of the total trading days, while stable conditions prevail for the remaining 60.95% of the period. This empirical segmentation formally establishes the macroeconomic environment, dictating exactly when the subsequent tactical hedging models will activate their behavioral risk premiums. Further comprehensive diagnostics confirming the strict bimodal quality and dynamic

stability of these probabilistic classifications (including the Regime Classification Measure and Probability Stability Index) are detailed in Appendix A.

4.2. Comparative Behavior of Optimal Hedge Ratios

Driven by the behavioral state classifications established above, the tactical framework computes a time-varying adaptive risk premium (α_t) to dynamically adjust the hedging positions. As reported in Table 3 and illustrated in Figure 2, this adaptive premium is characterized by a median of zero and a mean of 0.017. Methodologically, α_t remains strictly at zero during normal market conditions and exhibits pronounced sensitivity during volatile phases, reaching a maximum magnitude of 0.092 when panic clusters are identified.

Table 3. Descriptive Statistics of the Adaptive Learning Parameter and Out-of-Sample Hedge Ratios

	OHR_BASE	ALPHA	OHR_ALPHA	OHR_ON_OFF
Mean	0.892	0.017	0.90	0.345
Median	0.914	0	0.917	0
Maximum	1.242	0.092	1.265	1.232
Minimum	0.437	0	0.437	0
Std. Dev.	0.122	0.023	0.125	0.438
Skewness	-0.626	0.952	-0.546	0.548
Kurtosis	3.821	2.421	3.871	1.415
Jarque-Bera	135.26	238.98	117.76	223.84
Probability	0	0	0	0
Sum	1290.417	24.727	1302.324	498.629
Sum Sq. Dev.	21.48	0.74	22.73	277.89
Observations	1447	1447	1447	1447




Figure 2. Adaptive Behavioural Premium (α_t)

The integration of this adaptive premium with the regime signals produces two structurally distinct series of dynamic hedge ratios. Figure 3 provides a visual comparison of their operational trajectories.

The baseline strategy (OHR_base) and the Enhanced Behavioral Hedging strategy (OHR_alpha) both exhibit continuous, relatively smooth operational profiles, fluctuating around a mean hedging weight of 0.892 and 0.90, respectively. However, by dynamically incorporating α_t , OHR_alpha structurally deviates upward from the baseline during high-volatility states, achieving a higher maximum hedging weight of 1.265. This divergence objectively demonstrates the execution of the behavioral risk premium, aggressively increasing the short futures position to absorb anticipated downside shocks during panic regimes.




Figure 3. Dynamic Hedge Ratios under Alternative Hedging Strategies

In fundamental contrast, the tactical on/off strategy (OHR_on/off) demonstrates a discrete, highly discontinuous trajectory. Exhibiting a median of zero, this strategy remains strictly inactive during normal states and activates solely during identified panic regimes. This structural design explicitly trades off continuous downside protection in exchange for a drastic reduction in portfolio turnover, an operational mechanism that directly impacts transaction cost efficiency and net performance.

4.3. Out-of-Sample Hedging Effectiveness and Economic Evaluation

Table 4 evaluates the out-of-sample performance of the alternative hedging models. To explicitly address the structural trade-off between downside protection and operational costs, the analysis examines absolute risk reduction (Panel A), risk-adjusted economic utility (Panel B), and transaction efficiency (Panel C).

Table 4. Summary of Performance Metrics for Out-of-Sample Hedging Strategies				
Indicator Group	Unhedged	OHR_base	OHR_alpha	OHR_on/off
Panel A: Absolute risk reduction				
Portfolio Variance ($\times 10^{-4}$)	1.89	0.184	0.186	0.768
Hedging Effectiveness (HE, %)	-	90.24	90.13	59.30
Panel B: Risk-adjusted Economic utility				
Mean Daily Portfolio Return ($\times 10^{-4}$)	3.07	4.63	4.85	7.65
Daily Sharpe Ratio	0.022	0.108	0.112	0.087
Certainty Equivalent Return (CER $\times 10^{-4}$)	-0.71	4.26	4.47	6.12
Panel C: Transaction efficiency				
Mean active hedging proportion (<i>MeanD</i>)	-	1	1	0.39
Trading Efficiency Index (Efficiency)	-	90.24	90.13	151.89

Note: All are reported on a daily basis. CER is calculated under a coefficient of relative risk aversion of $\gamma = 4$. MeanD denotes the proportion of trading days on which an active hedge position is maintained.

4.3.1. Absolute Risk Reduction and Portfolio Variance

The foundational objective of any hedging mechanism is the structural suppression of portfolio variance. Panel A demonstrates that all tactical strategies provide robust protective capabilities. The unhedged portfolio exhibits a substantial baseline variance of 1.89×10^{-4} . The implementation of the continuous baseline strategy (OHR_base) forcefully compresses this variance to 0.184×10^{-4} , translating to a Hedging Effectiveness (HE) of 90.24%.

Interestingly, incorporating the behavioral risk premium does not prioritize pure variance reduction. OHR_alpha yields a variance of 0.186×10^{-4} and an HE of 90.13%, marginally lower than the baseline. However, this is a theoretical feature, not a flaw: the objective of the time-varying α_t is not to over-hedge daily noise, but to strategically reallocate capital during specific panic shocks. Furthermore, the discrete OHR_on/off strategy successfully captures an HE of 59.30%. Despite its inherently intermittent nature,

retaining nearly 60% of the maximum hedging effectiveness confirms that a significant proportion of total market variance is heavily clustered within the identified behavioral panic regimes.

4.3.2. Risk-Adjusted Performance and Economic Utility

While variance reduction measures statistical mitigation, it fails to account for the return penalty – specifically, the transaction cost and the opportunity cost – incurred by maintaining short futures positions. Panel B addresses this fundamental trade-off by evaluating risk-adjusted returns and investor utility, assuming an institutional risk aversion degree of $\gamma = 4$.

It is in these metrics that the economic value of the behavioral signals becomes explicit. The unhedged portfolio yields a fundamentally weak daily Sharpe ratio of 0.022. The standard OHR_base improves this to 0.108. However, the behaviorally augmented OHR_alpha strategy delivers a Sharpe ratio of 0.112. This represents a 4.14% relative improvement over the continuous baseline and is more than five times higher than that of the unhedged position.

More strikingly, the tactical OHR_on/off strategy reveals a profound economic paradox. By accepting a higher variance (0.768×10^{-4}), it generates the highest mean return (7.65×10^{-4}) among all configurations – approximately 2.5 times the unhedged return. Consequently, it achieves the highest Certainty Equivalent Return (CER = 6.12×10^{-4}). Compared to the OHR_base (CER = 4.26×10^{-4}), the OHR_on/off framework improves institutional utility by an impressive 43.6%. This indicates that selectively mitigating the hedging drag during normal, upward-drifting markets can mathematically dominate continuous protection in terms of absolute wealth accumulation.

4.3.3. Transaction Costs and Trading Efficiency

The theoretical superiority of any financial model is frequently negated by excessive transaction costs derived from over-rebalancing. Panel C introduces this critical empirical constraint by evaluating the active hedging proportion (*MeanD*) and Trading Efficiency.

Here, the OHR_on/off strategy establishes its definitive structural dominance. While OHR_base and OHR_alpha continuously maintain and rebalance short futures positions (*MeanD* = 1), OHR_on/off exhibits a *MeanD* of 0.39. This signifies that the strategy remains completely inactive and entirely immune to trading frictions for approximately 61% of the evaluation period. Because OHR_on/off secures 59.30% of the hedging effectiveness while only incurring costs on 39.05% of the days, it achieves a staggering Trading Efficiency score of 151.89. This magnitude is 1.68 times (or 68.2%) greater than the continuous models, empirically proving that continuous hedging is economically suboptimal when operational frictions are explicitly modeled.

4.3.4. Cumulative Trajectories and Drawdown Mitigation

Figure 4 contextualizes the performance metrics reported in Table 4 by plotting the cumulative return trajectories of the four portfolios. This temporal visualization highlights how the core trade-off between hedging drag and downside protection unfolds across alternating market regimes. By tracking wealth accumulation over time, the cumulative trajectories empirically illustrate the distinct economic mechanics embedded in each model's design.

During the expansionary macroeconomic regime (2020-2022), the graph explicitly visualizes the concept of performance drag discussed in Panel B. The continuous models (OHR_base and OHR_alpha) heavily penalize the portfolio's upward mobility due to the friction of perpetual short positions. In stark contrast, the structural advantage of the tactical OHR_on/off strategy is realized here: by maintaining a

zero-hedge position ($MeanD = 0$), it successfully bypasses this drag, tightly mirroring the unhedged portfolio's capital growth.

However, the empirical value of the dynamic frameworks is permanently crystallized during the protracted liquidity and panic cycles (late 2021 through 2023). As the unhedged portfolio suffers a catastrophic drawdown – eventually becoming the lowest-yielding asset – the distinct protective mechanics of the hedged models activate.

Crucially, the post-2023 divergence completely corroborates the Trading Efficiency metrics from Panel C. OHR_on/off aggressively dominates the performance hierarchy because it enters the recovery phase without the anchor of continuous short positions, mathematically explaining its superior Certainty Equivalent Return (CER). Meanwhile, although OHR_alpha and OHR_base exhibit tightly correlated trajectories, a persistent geometric gap emerges between them specifically during the severe drawdowns and is maintained thereafter. This visible premium explicitly confirms that the behavioral parameter (α_t) successfully translates temporary, state-dependent risk adjustments into permanent long-term capital preservation, fundamentally justifying its higher Sharpe ratio.




Figure 4. Out-of-sample cumulative returns.

4.3.5. Strategic Implications

Collectively, these empirical findings map a clear operational frontier. For institutional mandates strictly requiring absolute volatility suppression and downside protection, OHR_alpha is the optimal configuration, maximizing the Sharpe ratio without sacrificing HE. Conversely, for cost-sensitive practitioners aiming to maximize long-term utility, OHR_on/off provides a highly innovative solution, leveraging sentiment signals to eliminate 61% of active hedging days while securing peak Certainty Equivalent Returns.

5. Discussion

5.1. Beyond Variance Minimisation: Hedging as an Adaptive Decision Framework

Since the seminal framework of Ederington (1979), the dominant objective in optimal hedge ratio research has been the minimisation of portfolio variance. Subsequent methodological advancements have primarily focused on refining conditional covariance estimates to maximise hedging effectiveness. However, the empirical results presented in Section 4 indicate that variance reduction represents only one dimension of comprehensive hedging performance. While continuous estimation provides robust theoretical baselines, it inherently imposes persistent constraints on portfolio returns and accumulates trading frictions.

By contrasting the continuous strategies with the tactical OHR_{on/off} framework, this study contributes a shift in perspective to the existing literature: dynamic hedging should be evaluated as a state-dependent decision problem rather than strictly a parameter estimation problem. The empirical evidence demonstrates that maintaining a zero-hedge position during normal market states (approximately 61% of the out-of-sample period) allows the portfolio to capture the positive upward drift of the equity market. Conversely, activating the hedge exclusively during the 39.05% of days classified as high-risk panic regimes effectively truncates severe tail risk.

Theoretically, this finding directly addresses the structural limitations identified by Park and Switzer (1995) and Hou and Li (2013), who documented that the economic utility of dynamic hedging is frequently negated by excessive portfolio rebalancing. Instead of attempting to eliminate this friction through complex statistical smoothing penalties, the proposed framework explicitly integrates the cost-benefit trade-off into a behavioural state mechanism.

Practically, these results offer a quantifiable solution for institutional fund managers. The OHR_{on/off} strategy functions as a cost-effective tail-risk management framework. It provides a quantitative justification to minimise portfolio turnover during stable periods – thereby structurally mitigating the accumulation of transaction costs – while ensuring robust defensive positioning when systemic behavioural risks escalate.

Overall, these findings provide robust empirical support for Hypothesis 1. The evidence confirms that a behavioural timing mechanism can substantially reduce trading activity while preserving, and in terms of certainty-equivalent returns enhancing, favourable economic performance.

5.2. Operationalising Behavioural Risk: Asymmetric Downside Protection

The second major contribution of this study lies in bridging descriptive behavioural finance with prescriptive quantitative hedging. Traditional dynamic hedging models adjust hedge ratios in response to changes in conditional variance and covariance, implicitly assuming that market risk evolves solely through the information reflected in second moments.

The empirical performance of the Enhanced Behavioral Hedging strategy (OHR_{alpha}) demonstrates that behavioural conditions constitute a complementary, quantifiable dimension of market risk. Theoretically, this aligns with Prospect Theory (Kahneman and Tversky, 1992), which posits that investors exhibit loss aversion, and corroborates Ang and Chen (2002), who document asymmetric, intensified market co-movements during downturns. Because standard volatility models may structurally under-hedge during panic regimes – when behavioural distortions amplify fundamental declines – supplementing baseline volatility with a behavioural premium addresses a critical vulnerability in quantitative risk management.

Crucially, the methodology avoids abrupt, binary triggers. By integrating the behavioral premium proportionally with the predicted probability of herding, the additional hedging weight scales smoothly. This design explains why OHR_{alpha} yields relatively modest gains over the full sample: it inherently

functions as conditional tail-risk insurance, operating identically to the baseline during calm periods and concentrating its economic value strictly during high-stress regimes.

From an economic and practical perspective, this proportional integration transforms behavioural anomalies from abstract theoretical observations into tradable parameters. For wealth management platforms and algorithmic robo-advisors, OHR_alpha provides a systematic mechanism to mitigate the portfolio impact of behavioural shocks. By increasing derivative hedge ratios proportionally as panic probabilities rise, portfolio managers can secure downside protection without liquidating underlying spot assets. At a broader market level, the deployment of such state-contingent hedging frameworks functions as an endogenous stabilising mechanism; by shifting risk-mitigation activities to the derivatives market, it indirectly alleviates panic-induced selling pressures in the spot market, thereby contributing to the stabilisation of systemic liquidity during crisis events.

Taken together, these findings support [Hypothesis 2](#). The evidence confirms that incorporating an adaptive behavioural premium alongside predicted panic probability provides superior risk-adjusted portfolio performance specifically during periods of elevated behavioural risk, successfully accommodating the asymmetric, loss-averse preferences of market participants.

5.3. Methodological Limitations and Future Research

Despite the recursive out-of-sample framework implemented to prevent look-ahead bias, two key methodological limitations should be acknowledged.

First, the expanding-window estimation scheme continuously accumulates historical observations over time. As the sample size expands, parameter estimates gain statistical stability but inherently develop mathematical inertia. Consequently, the adaptive parameter (α_t) exhibits reduced responsiveness to newly emerging structural shocks during the later stages of the evaluation period. Future research could refine this mechanism by incorporating adaptive state-space models or decaying-memory algorithms – such as forgetting-factor filters or Bayesian updating – to balance parameter stability with sensitivity to evolving market conditions.

Second, trading frictions are proxied via portfolio rebalancing frequency (*MeanD*) rather than an explicit transaction cost function. Although *MeanD* effectively captures relative operational intensity, it does not differentiate between fixed commission fees, proportional trading charges, bid-ask spreads, or market impact. Extending this framework to integrate multi-component transaction cost functions represents a natural direction for future research, enabling a comprehensive net-of-cost performance evaluation of behavioral timing strategies.

6. Conclusion

This study proposed a behavioural tactical hedging framework that integrates behavioural finance into a conventional quantitative risk management system to address two persistent challenges in dynamic hedging: trading frictions arising from frequent portfolio rebalancing and the inadequate treatment of tail risk during periods of market stress. By combining a Markov-switching CSAD model for identifying herd-behaviour regimes with a VECM-Robust EWMA volatility framework, the study developed and evaluated two alternative out-of-sample hedging strategies under a rigorous expanding-window recursive estimation scheme. This design ensured that all model estimation and hedge-ratio forecasts were generated without look-ahead bias, thereby strengthening the practical credibility of the empirical findings.

The empirical evidence contributes to the hedging literature in two important respects. First, the proposed framework demonstrates that behavioural information can be incorporated into quantitative hedge-ratio determination through an explicit decision mechanism rather than being treated solely as an

explanatory characteristic of financial markets. By conditioning hedging decisions on recursively estimated market regimes, the framework extends conventional volatility-based hedging into a state-dependent decision process that allocates hedging effort according to prevailing market conditions.

Second, the empirical results illustrate two distinct implementations of this behavioural framework. The conditional strategy, OHR_on/off, substantially reduced trading activity by restricting hedge positions to periods characterised by elevated herding risk. Although this design resulted in higher portfolio variance than continuously hedged strategies, the evidence suggests that this reflects an intentional trade-off rather than inferior risk management, as maintaining unhedged exposure during favourable market conditions allowed the portfolio to retain greater upside potential. Meanwhile, the enhanced strategy, OHR_alpha, incorporated behavioural information as a dynamic adjustment to the conventional hedge ratio. By increasing hedge intensity selectively during periods of elevated behavioural risk, it achieved modest but consistent improvements in risk-adjusted performance, supporting the view that behavioural risk constitutes an economically relevant component of hedge-ratio determination.

Overall, the findings suggest that no single hedging strategy universally dominates across all investment objectives. Rather than seeking a universally optimal hedge ratio, the proposed framework highlights the importance of aligning hedging decisions with institutional objectives and prevailing market conditions. Continuous adaptive hedging may be more appropriate for investors prioritising downside risk control, whereas tactical hedging may offer advantages for investors willing to accept greater return variability in exchange for lower trading intensity and higher economic utility.

Several avenues remain for future research. Extending the framework to incorporate more flexible recursive estimation techniques, explicit nonlinear transaction-cost models, or data-driven optimisation of behavioural activation thresholds may further improve both the adaptability and practical applicability of behavioural hedging strategies. More broadly, the present study suggests that future developments in quantitative risk management may benefit less from pursuing increasingly complex volatility models alone than from designing decision frameworks capable of integrating market behaviour into the hedging process in a systematic and economically interpretable manner.

References

- Ang, A. and Chen, J. (2002), "Asymmetric correlations of equity portfolios", *Journal of Financial Economics*, Vol. 63 No. 3, pp. 443-494, doi: [https://doi.org/10.1016/S0304-405X\(02\)00068-5](https://doi.org/10.1016/S0304-405X(02)00068-5).
- Bohl, M.T., Klein, A.C. and Siklos, P.L. (2016), "A Markov switching approach to herding", *Credit and Capital Markets-Kredit und Kapital*, Vol. 49 No. 2, pp. 193-220, doi: <https://doi.org/10.3790/ccm.49.2.193>.
- Brooks, C., Henry, O.T. and Persaud, G. (2002), "The effect of asymmetries on optimal hedge ratios", *The Journal of Business*, Vol. 75 No. 2, pp. 333-352, doi: <https://doi.org/10.1086/338484>.
- Buyukkara, G., Kucukozmen, C.C. and Uysal, E.T. (2022), "Optimal hedge ratios and hedging effectiveness: an analysis of the Turkish futures market", *Borsa Istanbul Review*, Vol. 22 No. 1, pp. 92-102, doi: <https://doi.org/10.1016/j.bir.2021.02.002>.
- Cao, L., Sun, R., Ma, T. and Liu, C. (2023), "On asymmetric correlations and their applications in financial markets", *Journal of Risk and Financial Management*, Vol. 16 No. 3, p. 187, doi: <https://doi.org/10.3390/jrfm16030187>.
- Chang, E.C., Cheng, J.W. and Khorana, A. (2000), "An examination of herd behavior in equity markets: an international perspective", *Journal of Banking and Finance*, Vol. 24 No. 10, pp. 1651-1679, doi: [https://doi.org/10.1016/S0378-4266\(99\)00096-5](https://doi.org/10.1016/S0378-4266(99)00096-5).

Doherty, O. (2018), “Informational cascades in financial markets: review and synthesis”, *Review of Behavioral Finance*, Vol. 10 No. 1, pp. 53-69, doi: <https://doi.org/10.1108/RBF-05-2016-0030>.

Ederington, L.H. (1979), “The hedging performance of the new futures markets”, *The Journal of Finance*, Vol. 34 No. 1, pp. 157-170, doi: <https://doi.org/10.1111/j.1540-6261.1979.tb02077.x>.

Fabozzi, F.J., Gupta, F. and Markowitz, H.M. (2002), “The legacy of modern portfolio theory”, *The Journal of Investing*, Vol. 11 No. 3, pp. 7-22, doi: <https://doi.org/10.3905/joi.2002.319510>.

Gong, X.L., Liu, X.H. and Xiong, X. (2019), “Measuring tail risk with GAS time varying copula, fat tailed GARCH model and hedging for crude oil futures”, *Pacific-Basin Finance Journal*, Vol. 55, pp. 95-109, doi: <https://doi.org/10.1016/j.pacfin.2019.03.010>.

Hao, N.X. and Nhung, N.T. (2023), “Impacts of the derivatives market on the underlying stock market: the case of VN30”, doi: <https://doi.org/10.57110/jebvn.v3i3.222>.

Harris, R.D. and Shen, J. (2003), “Robust estimation of the optimal hedge ratio”, *Journal of Futures Markets*, Vol. 23 No. 8, pp. 799-816, doi: <https://doi.org/10.1002/fut.10085>.

Hirshleifer, D. and Teoh, S.H. (2003), “Herd behaviour and cascading in capital markets: a review and synthesis”, *European Financial Management*, Vol. 9 No. 1, pp. 25-66, doi: <https://doi.org/10.1111/1468-036X.00207>.

Hou, Y. and Li, S. (2013), “Hedging performance of Chinese stock index futures: an empirical analysis using wavelet analysis and flexible bivariate GARCH approaches”, *Pacific-Basin Finance Journal*, Vol. 24, pp. 109-131, doi: <https://doi.org/10.1016/j.pacfin.2013.04.001>.

Johnson, L.L. (1960), “The theory of hedging and speculation in commodity futures”, *The Review of Economic Studies*, Vol. 27 No. 3, pp. 139-151, doi: <https://doi.org/10.2307/2296076>.

Lien, D. (1996), “The effect of the cointegration relationship on futures hedging: a note”, *Journal of Futures Markets*, Vol. 16 No. 7, pp. 773-780, doi: [https://doi.org/10.1002/\(SICI\)1096-9934\(199610\)16:7<773::AID-FUT3>3.0.CO;2-L](https://doi.org/10.1002/(SICI)1096-9934(199610)16:7<773::AID-FUT3>3.0.CO;2-L).

Lien, D. and Luo, X. (1994), “Multiperiod hedging in the presence of conditional heteroskedasticity”, *Journal of Futures Markets*, Vol. 14 No. 8, pp. 927-955, doi: <https://doi.org/10.1002/fut.3990140806>.

Louhichi, W. and Rais, H. (2019), “Refinement of the hedging ratio using copula-GARCH models”, *Journal of Asset Management*, Vol. 20 No. 5, pp. 403-411, doi: <https://doi.org/10.1057/s41260-019-00133-5>.

Nguyen, Y.V.B. and Vo, A.H.K. (2024), “Herding behavior before and after COVID-19 pandemic: evidence from the Vietnam stock market”, *Journal of Economic Studies*, Vol. 51 No. 2, pp. 357-374, doi: <https://doi.org/10.1108/JES-01-2023-0031>.

Park, T.H. and Switzer, L.N. (1995), “Time-varying distributions and the optimal hedge ratios for stock index futures”, *Applied Financial Economics*, Vol. 5 No. 3, pp. 131-137, doi: <https://doi.org/10.1080/758523000a>.

Qu, H., Wang, T., Zhang, Y. and Sun, P. (2019), “Dynamic hedging using the realized minimum-variance hedge ratio approach—examination of the CSI 300 index futures”, *Pacific-Basin Finance Journal*, Vol. 57, p. 101048, doi: <https://doi.org/10.1016/j.pacfin.2018.08.002>.

Tversky, A. and Kahneman, D. (1992), “Advances in prospect theory: cumulative representation of uncertainty”, *Journal of Risk and Uncertainty*, Vol. 5 No. 4, pp. 297-323, doi: <https://doi.org/10.1007/BF00122574>.

Yang, W. and Allen, D.E. (2005), “Multivariate GARCH hedge ratios and hedging effectiveness in Australian futures markets”, *Accounting and Finance*, Vol. 45 No. 2, pp. 301-321, doi: <https://doi.org/10.1111/j.1467-629x.2004.00119.x>.

Appendix

Preliminary Econometric Diagnostics and VECM Specification

A.1. Unit Root and Stationarity Tests

Before estimating the dynamic hedging strategies, the integration properties of the logarithmic spot ($\ln S_t$) and futures ($\ln F_t$) price series are evaluated. As reported in Table A1, the Augmented Dickey-Fuller (ADF) tests indicate that both price series are non-stationary in levels but stationary in first differences, confirming they are integrated of order one, $I(1)$. This finding satisfies the fundamental prerequisite for cointegration analysis.

Table A1. Results of the Augmented Dickey-Fuller (ADF) Unit Root Test

Variable	Deterministic Terms	Test Stat	5% Critical Value	Inference
$\ln(S_t)$ (level)	Trend and intercept	-1.3737	-3.41	Non-stationary at level
$\ln(F_t)$ (level)	Trend and intercept	-1.4318	-3.41	Non-stationary at level
$\Delta \ln(S_t) = R_s$ (1 st Diff)	Drift	-28.8565	-2.86	Stationary (I(0))
$\Delta \ln(F_t) = R_f$ (1 st Diff)	Drift	-29.0936	-2.86	Stationary (I(0))

Note: The optimal lag length in each Augmented Dickey–Fuller (ADF) regression is selected automatically according to the Bayesian Information Criterion (BIC).

A.2. Johansen Cointegration Analysis

Given the $I(1)$ properties, an unrestricted Vector Autoregression (VAR) system is constructed to determine the optimal lag structure. As shown in Table A2, the Akaike Information Criterion (AIC), Hannan-Quinn (HQ), and Final Prediction Error (FPE) uniformly select an optimal lag length of 2. The estimated characteristic roots all lie strictly inside the unit circle, confirming that the unrestricted VAR(2) satisfies the stability condition and provides an appropriate basis for Johansen cointegration testing.

Table A2. Lag Selection and Stability Diagnostics for the Unrestricted VAR Model

AIC(n)	HQ(n)	SC(n)	FPE(n)
2	2	1	2

Root_Index	Modulus_Value	Is_Stable
1	0.9541	Yes

Conditional on the VAR(2) specification, the Johansen trace test (Table A3) decisively rejects the null hypothesis of no cointegration ($r = 0$) at the 1% significance level, while failing to reject the hypothesis of at most one cointegrating vector ($r \leq 1$) at the 5% level.

Table A3. Johansen Cointegration Test (Trace Statistic)

Hypothesis	Test_Statistic	Crit_10pct	Crit_5pct	Crit_1pct
$r \leq 1$	7.8161	7.52	9.24	12.97
$r = 0$	27.3223	17.85	19.96	24.6

These results provide empirical evidence of a single long-run equilibrium relationship between the VN30 spot index and the VN30F1M futures contract. This finding formally justifies the specification of a Vector Error Correction Model (VECM) with a cointegration rank of one, allowing the joint modeling of the long-run equilibrium relationship and the short-run adjustment dynamics between the two markets.

A.3. Residual Distributional Diagnostics

The recursive expanding-window VECM procedure generates strictly out-of-sample, one-step-ahead forecast errors for the spot (ε_{spot}) and futures (ε_{fut}) markets. Analyzing the statistical properties of these innovations is essential for correctly specifying the subsequent conditional volatility framework.

As reported in Table A4, the sample means of both residual series are economically negligible (1.58×10^{-4} and -2.2×10^{-4} , respectively), which is consistent with the asymptotic unbiasedness of the recursive forecasting scheme. However, higher-order moments reveal pronounced departures from Gaussianity. Both series exhibit negative skewness (-0.7496 and -0.8027), indicating an asymmetric distribution with a higher incidence of large negative forecast errors. Furthermore, the substantial excess kurtosis (4.212 and 6.057) highlights the presence of heavy-tailed behavior, meaning extreme market shocks occur more frequently than a normal distribution would imply. Consequently, the Jarque-Bera test decisively rejects the null hypothesis of normality at the 1% significance level for both markets.

Table A4. Distributional Properties of the Out-of-Sample VECM Forecast Errors

Series	e_spot (VN30 Index)	e_fut (VN30F1M)
Mean ($\times 10^{-4}$)	1.58×10^{-4}	-2.2×10^{-4}
Std_Dev	0.0139	0.0159
Skewness	-0.7496	-0.8027
Kurtosis	7.212	9.057
Excess_Kurtosis	4.212	6.057
JB_Stat	1229.519	2415.053
JB_P_Value	0	0

Normal_AIC	-8423.63	-8030.84
Laplace_AIC	-8786.94	-8484.45

To formally identify an appropriate distributional assumption for the volatility estimator, the empirical fit of the Gaussian and Laplace distributions is evaluated and compared using the Akaike Information Criterion (AIC). For both residual series, the Laplace specification yields substantially lower AIC values than the Gaussian alternative. This provides quantitative evidence that the out-of-sample innovations are more accurately characterized by a double-exponential (Laplace) distribution. This distributional finding is visually corroborated by Figure A1. The empirical density displays a sharper central peak and heavier tails than the normal distribution, closely aligning with the fitted Laplace density curve. Additionally, the Q-Q plots reveal systematic deviations from the theoretical normal reference line primarily in the tails, confirming that the non-normality is driven by extreme outliers rather than observations around the center of the distribution.




Figure A1. Empirical Distribution and Q–Q Diagnostics of the Out-of-Sample VECM Forecast Errors

Beyond distributional shape, residual diagnostic tests are conducted to assess serial dependence and volatility clustering within the forecast errors (Table A5). The Ljung-Box test indicates no statistically significant serial correlation for the spot residuals at the 5% level, although a modest degree of autocorrelation remains in the futures residuals. More importantly, the ARCH-LM test strongly rejects the null hypothesis of homoskedastic disturbances for both series, with p-values approaching zero. This confirms the presence of significant conditional heteroskedasticity and volatility clustering in the out-of-sample innovations.

Table A5. Residual Diagnostic Tests for the Recursive VECM Forecast Errors

Series	Ljung_Box_Pval	ARCH_LM_Pval
VECM OOS e_spot	0.08435144	1.2873×10^{-33}

VECM OOS e _{fut}	0.01566897	1.5666×10^{-31}
---------------------------	------------	--------------------------

Collectively, these empirical properties provide a rigorous econometric justification for the dynamic hedging design employed in this study. First, the strong evidence of conditional heteroskedasticity necessitates the use of a time-varying volatility estimator rather than assuming constant second moments. Second, the established heavy-tailed, non-Gaussian behavior best captured by the Laplace distribution supports the adoption of the Robust EWMA framework of Harris and Shen (2003). By relying on absolute deviations rather than squared residuals, this framework structurally mitigates the disruptive impact of the extreme observations inherent in these financial return series. Consequently, based on the superior empirical fit of the Laplace distribution confirmed by the AIC and kurtosis diagnostics, the geometric scaling constant for the Robust EWMA volatility estimator is strictly set to $g(1) = \sqrt{2} \approx 1.4142$ for all subsequent dynamic hedging weight calculations.

Appendix B: Robustness of the Regime Classification Mechanism

B.1. Methodological Definitions of Signal Quality Metrics

To rigorously evaluate the quality and stability of the panic-herding probabilities extracted from the recursive MRS-CSAD framework, two primary diagnostics are employed.

The first is the Regime Classification Measure (RCM), which evaluates the sharpness of the out-of-sample predicted probabilities. Adapted from the continuous nature of the state probabilities, the RCM is defined as:

$$RCM = 400 \times \frac{1}{N} \sum_{t=1}^N p_t (1 - p_t)$$

where p_t is the predicted probability of the herding regime at time t , and N is the total number of out-of-sample observations. The RCM scale ranges from 0 to 100. A score of 0 indicates perfect binary classification (probabilities are strictly 0 or 1), while a score of 100 indicates total ambiguity (probabilities remain constant at 0.5).

The second is the Probability Stability Index (PSI), designed to measure the dynamic smoothness and lack of erratic jitter in the day-to-day signal. The daily probability step variation is defined as $\Delta p_t = p_t - p_{t-1}$. The PSI is evaluated through the variance of Δp_t , alongside the median and mean of the absolute deviations $|\Delta p_t|$. Values strictly converging toward zero indicate a highly stable signal that does not induce excessive spurious toggling between regimes.

B.2. Diagnostic Results and Empirical Implications

Table B1 reports the out-of-sample diagnostic metrics for the regime classification signal.

Table B1. Out-of-Sample Probability Evaluation Metrics

Indicator Group	Metric Specification	Value
Regime Classification Measure (RCM)	Out-of-Sample Classification Sharpness	84.86
	Variance of Δp_t	0.0068
Probability Stability Index (PSI)	Median of $ \Delta p_t $	0.0061

Mean of $ \Delta p_t $	0.0399
------------------------	--------

The RCM score of 84.86 reflects the continuous, probabilistic nature of out-of-sample behavioral forecasting. Rather than generating artificial binary (0 or 1) polarization – which often symptomatic of in-sample overfitting – this metric indicates that the recursive model correctly preserves the inherent uncertainty associated with real-time regime transitions in financial markets. The model successfully segregates broad behavioral states without forcing deterministic certainty on complex human behaviors.

Crucially, the signal exhibits exceptional dynamic stability. The PSI metrics reveal a near-zero variance (0.0068) and an extremely constrained median daily change (0.0061). This empirical smoothness confirms that the predicted probabilities do not suffer from erratic, day-to-day fluctuations. Economically, this lack of signal jitter is the direct catalyst for the operational efficiency observed in Section 4: it prevents the tactical OHR_on/off strategy from incurring the prohibitive transaction costs associated with continuous, false-positive regime toggling.