Intraday VIX Hedging via Deep Filtering: A State-Space Approach to Volatility Risk Management

Ningning Liu *

School of Mathematics and Statistics, Henan University, Henan, China

Abstract. This paper develops a deep filtering framework for intraday VIX derivatives hedging that generalizes traditional stochastic volatility models. Volatility is mod-eled as a latent process in a partially observed state-space system, where neural networks serve as non-linear Bayesian filters to infer hidden states from high-frequency market data. The methodology directly optimizes hedging portfolio variance rather than price prediction accuracy. Using tick-by-tick data from 2010-2024, the approach demonstrates 42% reduction in hedging error variance, 68% higher Sharpe ratios, and superior performance across all market regimes compared to traditional Delta hedging methods. The framework maintains robust per-formance after accounting for transaction costs and market microstructure effects.

Keywords: Deep Filtering; Stochastic Volatility Models; Latent Volatility Process; Neural Network Bayesian Filters; High-Frequency Data.

1. Introduction

Volatility risk management through VIX derivatives represents a fundamental chal-lenge in modern quantitative finance. The Chicago Board Options Exchange Volatility Index (VIX) and its associated derivatives have become essential tools for managing equity portfolio risk, with daily trading volume exceeding 800,000 contracts. How-ever, traditional hedging approaches based on parametric stochastic volatility models face significant limitations in high-frequency settings where market microstructure ef-fects, non-stationary dynamics, and complex dependency structures dominate price formation [1].

The core methodological challenge stems from the latent nature of volatility processes. Unlike equity prices that are directly observable, volatility must be inferred from op-tion prices or high-frequency returns, introducing substantial estimation risk and model misspecification concerns [2]. This problem intensifies at intraday frequencies, where tra-ditional filtering methods like Kalman filters and particle methods struggle with the high-dimensional, noisy data characteristic of modern electronic markets. Existing approaches based on Heston [3] and other parametric models often fail during volatility regime shifts, precisely when effective hedging is most critical for risk man-agement.

This research addresses these limitations through a novel deep filtering framework that combines rigorous state-space modeling with modern neural network architec-tures. The approach formalizes volatility hedging as a partially observed stochastic control problem, where the hedging agent must simultaneously infer latent volatility states from noisy market observations while optimizing portfolio positions. The deep filter operates as a non-linear Bayesian estimator, learning the mapping from observed market variables to optimal hedging strategies without relying on restrictive paramet-ric assumptions about the underlying data-generating process.

The methodological contributions proceed along three dimensions. First, we develop a general deep filtering architecture for volatility state estimation that explicitly han-dles the complex temporal dependencies, market microstructure noise, and regime-switching behavior present in high-frequency VIX data. Second, we introduce a direct portfolio optimization objective that minimizes hedging error variance rather than fo-cusing on intermediate price prediction targets, ensuring alignment between model training and practical risk management objectives. Third, we provide comprehen-sive empirical validation demonstrating the framework's robustness across different market regimes, including detailed analysis of transaction costs, liquidity effects, and implementation considerations.

From a practical perspective, this research provides volatility-dependent portfolios with more effective risk management tools during periods of market stress. The ability to maintain hedging performance across different volatility regimes addresses a crit-ical need in modern portfolio management, particularly as VIX products become in-creasingly integrated into systematic investment strategies and risk parity approaches.

* Corresponding author: liuningning990722@163.com

2. Methodology

2.1 State-Space Formulation for Volatility Dynamics

The deep filtering framework begins with a general state-space representation of volatil-ity dynamics. Let vt ∈ Rd denote the latent volatility state vector and yt ∈ Rm repre-sent observed market variables at time t. The system evolves according to the follow-ing non-linear state-space model:

$$v_t = f_\theta(v_{t-1}, \epsilon_t), \epsilon_t \sim \mathcal{N}(0, Q_t) \quad (1)$$

$$y_t = g_\phi(v_t, \eta_t), \eta_t \sim \mathcal{N}(0, R_t) \quad (2)$$

where $f_\theta(\cdot)$ and $g_\phi(\cdot)$ are non-linear functions parameterized by neural networks with parameters θ and ϕ respectively. The state transition equation (1) models the evo-lution of latent volatility states, while the observation equation (2) relates these states to observable market quantities.

The state vector vt captures multiple dimensions of volatility dynamics: short-term fluctuations (1-5 day horizon), medium-term trends (1-3 week horizon), long-term regimes (1-6 month horizon), and volatility-of-volatility components. This multi-scale representation enables the model to capture the term structure effects that are crucial for VIX derivatives pricing and hedging across different maturities.

2.2 Deep Filter Architecture and Implementation

The filtering network implements a recursive Bayesian estimation procedure. At each time step t, the filter maintains a belief state $p(v_t \mid y_{1:t})$ and updates this distribution as new observations arrive according to the Bayesian filtering equations:

$$p(v_t \mid y_{1:t-1}) = \int p(v_t \mid v_{t-1}) p(v_{t-1} \mid y_{1:t-1}) dv_{t-1} \quad (3)$$

$$p(v_t \mid y_{1:t}) = \frac{p(y_t|v_t)p(v_t|y_{1:t-1})}{\int p(y_t|v_t)p(v_t|y_{1:t-1})dv_t} \quad (4)$$

The network architecture consists of four main components implemented through spe-cialized neural network modules:

Table 1. Deep Filter Architecture Components

Component	Network Type	Function
Encoder Network	Temporal Convolutional Network	Extracts multi-scale features from high-frequency market data
Transition Network	Gated Recurrent Unit (256 units)	Models non-linear state evolution with uncertainty
Observation Network	4-layer MLP [512,256,128,64]	Maps latent states to observable market quantities
Attention Mechanism	Multi-head self-attention (8 heads)	Weights relevant historical information dynamically

The training objective directly minimizes hedging portfolio variance over the re-balancing horizon:

$$\mathcal{L}(\theta, \phi) = \mathbb{E}\left[\left(\Delta\Pi_t - \sum_{i=1}^{N} \delta_{i,t}\Delta S_{i,t}\right)^2\right] + \lambda \cdot \mathrm{TC}(\delta_t) \quad (5)$$

where ΔΠt is the change in portfolio value, δi,t are the hedge ratios output by the deep filter, ΔSi,t are price changes of hedging instruments, and TC(δt) represents trans-action costs with regularization parameter λ.

2.3 Data Processing and Feature Engineering

The analysis uses comprehensive intraday data from January 2010 through December 2024, comprising multiple data sources processed through a unified feature engineer-ing pipeline:

Table 2. Data Sources and Feature Specifications

Data Type	Frequency	Key Features
VIX Futures	Tick-by-tick	All maturities (1-9 months), bid-ask spreads, volume, order flow
SPX Options	1-minute	Implied volatility surface, skew, term structure slope
S&P 500 Index	1-second	Realized volatility, jump components, overnight returns
ETF Flows	5-minute	Volatility ETF flows, market maker positioning
Macro Releases	Event time	Economic surprises, central bank communications

Market microstructure noise is addressed through signature methods (Lyons et al., 2014) and volatility signature plots. Missing data in the high-frequency series is han-dled through a combination of forward-filling for short gaps and neural process im-putation for longer intervals. The final feature set comprises 247 dimensions across the different data sources.

2.4 Implementation and Computational Considerations

The deep filtering framework is implemented in PyTorch with custom CUDA ker-nels for high-performance inference. Training uses a multi-stage procedure: first pre-training on a large-scale synthetic dataset generated from various stochastic volatility models, then fine-tuning on actual market data with the hedging objective. The model processes approximately 2.5 million observations during training, with a typical train-ing time of 18 hours on a single NVIDIA A100 GPU.

For real-time deployment, the framework maintains a sliding window of the most recent 1,000 observations (approximately one trading day) and generates hedging sig-nals with a latency of under 50 milliseconds. This performance enables practical im-plementation in live trading environments while maintaining the model's ability to capture both short-term dynamics and longer-term regime patterns.

3. Empirical Results

3.1 Filtering Performance and State Estimation Accuracy

The deep filter demonstrates superior volatility state estimation compared to tradi-tional methods across multiple evaluation metrics. Table 3 summarizes the out-of-sample performance from January 2023 through December 2024, using VIX futures prices as the ground truth benchmark for volatility levels:

Table 3. Volatility State Estimation Performance (2023-2024)

Method	RMSE (%)	MAE (%)	Correlation	Regime Accuracy
Deep Filter	2.34	1.87	0.94	89.2%
Heston Kalman Filter	4.12	3.25	0.82	67.8%
Particle Filter	3.67	2.89	0.85	73.4%
GARCH(1,1)	5.23	4.17	0.74	58.3%
Realized Volatility	3.45	2.71	0.87	61.5%
RNN	3.12	2.45	0.89	76.8%

The improvement is particularly pronounced during high-volatility periods (VIX ¿ 30), where the deep filter achieves RMSE of 3.12% compared to 6.45% for the Heston filter and 5.23% for the particle filter. The model's multi-scale representation successfully captures both rapid mean reversion in short-term volatility and persistent trends in longer-term volatility regimes.

3.2 Hedging Performance Analysis

Hedging performance is evaluated through a comprehensive set of metrics across dif-ferent VIX derivative types and market conditions. Table 4 presents results for hedging VIX call options with 30 days to maturity:

Table 4. VIX Call Options Hedging Performance

Metric	Deep Filter	Heston	BS	Improvement
Error Variance	0.142	0.328	0.415	56.7%
Sharpe Ratio	1.68	0.92	0.74	82.6%
Max Drawdown (%)	8.2	18.7	24.3	56.1%
Win Rate (%)	72.4	58.3	51.6	24.2%
P/L Skewness	0.45	-0.38	-0.62	-
Cost (bps/day)	3.2	5.8	7.1	44.8%
VaR (95%)	2.1%	4.8%	6.3%	56.3%

The 56.7% reduction in hedging error variance represents substantial risk reduction for volatility-dependent portfolios. The improvement is economically significant and persists across different rebalancing frequencies, from 5-minute to 1-hour intervals. The deep filter's superior performance stems from its ability to adapt hedge ratios dynamically in response to changing market conditions and volatility regimes.

3.3 Performance Across Option Characteristics

Table 5 examines how hedging performance varies across different option types and moneyness levels, providing insights into where the deep filtering approach provides the greatest advantages:

Table 5. Hedging Performance by Option Characteristics

Option Type	Moneyness	Deep Filter Error	Heston Error	Improvement
Call	0.8 (OTM)	0.156	0.412	62.1%
Call	0.9 (OTM)	0.143	0.356	59.8%
Call	1.0 (ATM)	0.138	0.298	53.7%
Call	1.1 (ITM)	0.135	0.267	49.4%
Call	1.2 (ITM)	0.132	0.254	48.0%
Put	0.8 (ITM)	0.141	0.267	47.2%
Put	0.9 (ITM)	0.144	0.289	50.2%
Put	1.0 (ATM)	0.145	0.312	53.5%
Put	1.1 (OTM)	0.152	0.378	59.8%
Put	1.2 (OTM)	0.163	0.435	62.5%

The results demonstrate that the deep filtering approach provides the greatest im-provement for out-of-money options, where traditional Delta hedging performs poor-est due to high gamma and vega sensitivity. This pattern holds for both calls and puts, suggesting that the framework's advantages are particularly valuable for options with non-linear payoffs.

3.4 Robustness Across Market Regimes

A critical test of any hedging methodology is its performance across different market environments. Table 6 examines how the deep filtering approach performs during various volatility regimes:

The deep filtering framework maintains strong performance across all regimes, with particularly significant improvements during high-volatility periods and regime tran-sitions. This robustness is crucial for practical risk management, as these are precisely the conditions where effective hedging is most valuable and traditional methods typ-ically fail.

Table 6. Performance Across Volatility Regimes

Volatility Regime	Deep Filter	Heston	Improvement	Regime Frequency
Low Volatility (VIX ¡ 15)	0.128	0.245	47.8%	42.3%
Normal Volatility (VIX 15-25)	0.146	0.334	56.3%	38.7%
High Volatility (VIX 25-40)	0.168	0.452	62.8%	14.2%
Extreme Volatility (VIX ¿ 40)	0.215	0.678	68.3%	4.8%
Rising Volatility	0.152	0.387	60.7%	21.5%
Falling Volatility	0.139	0.298	53.4%	23.8%
Volatility Regime Transition	0.161	0.423	61.9%	11.2%

3.5 Transaction Cost Analysis

Practical implementation requires consideration of transaction costs and market im-pact. Table 7 analyzes performance under different cost assumptions:

Table 7. Performance After Transaction Costs

Cost Scenario	Deep Filter	Heston	Improvement	Frequency
Low Costs (1-2 bps)	1.62	0.86	88.4%	5-min
Medium Costs (3-5 bps)	1.52	0.76	100.0%	15-min
High Costs (6-10 bps)	1.38	0.61	126.2%	30-min
Realistic Costs (2-8 bps)	1.42	0.68	108.8%	Mixed
With Market Impact	1.28	0.52	146.2%	Adaptive
Institutional Scale	1.15	0.41	180.5%	1-hour

The deep filtering approach maintains strong net performance even after accounting for realistic transaction costs. The framework's ability to generate smoother hedge ra-tio paths reduces unnecessary trading while still capturing essential risk adjustments, leading to better cost-adjusted performance compared to traditional methods.

4. Discussion and Conclusion

The deep filtering framework developed in this research represents a substantive method-ological advance in volatility risk management. By combining rigorous state-space modeling with modern neural network architectures, the approach addresses funda-mental limitations of parametric stochastic volatility models while maintaining the-oretical coherence through the Bayesian filtering formulation. The empirical results demonstrate clear and economically significant improvements over traditional hedg-ing methods across all evaluation metrics and market conditions.

The 56.7% reduction in hedging error variance and 82.6% improvement in Sharpe ratios provide compelling evidence for the framework's effectiveness in practical ap-plications. These performance advantages are particularly pronounced during high-volatility periods and for out-of-money options, precisely the situations where tradi-tional Delta hedging performs poorest and effective risk management is most critical. The robustness across different market regimes, transaction cost environments, and implementation scales further supports the framework's practical utility.

Methodologically, this research contributes to the growing literature at the intersec-tion of deep learning and financial economics. Unlike approaches that treat neural networks as black-box predictors, our framework embeds financial theory through the state-space formulation and maintains interpretability through the explicit model-ing of latent volatility states. The direct optimization of hedging performance rather than intermediate prediction targets ensures alignment between model training and practical risk management objectives.

Several limitations and directions for future research warrant consideration. The com-putational requirements, while manageable for institutional applications, may present challenges for smaller market participants. The interpretability of the learned state representations, while improved through attention mechanisms and state visualiza-tion techniques, remains more complex than parametric alternatives. Future work could explore hybrid approaches that combine the flexibility of deep learning with the interpretability of parametric models.

From a practical perspective, this research provides volatility-dependent portfolios with more effective risk management tools, particularly during periods of market stress. The ability to maintain hedging performance across different volatility regimes addresses a critical need in modern portfolio management as VIX products become increasingly integrated into systematic investment strategies. The framework's flexi-bility suggests potential applications beyond VIX derivatives to other markets where latent factors drive pricing and risk.

In conclusion, the deep filtering approach developed in this paper offers a power-ful and theoretically grounded framework for intraday VIX derivatives hedging. By successfully combining the flexibility of neural networks with the rigor of state-space modeling, the methodology provides substantial improvements in hedging perfor-mance while maintaining robustness across different market conditions. This research represents a meaningful step forward in the application of modern machine learning techniques to fundamental problems in financial risk management.

References

1. Luo R, Zhang W, Xu X, et al. A neural stochastic volatility model[C]//proceedings of the AAAI conference on artificial intelligence. 2018, 32(1). https://doi.org/10.1609/aaai.v32i1.12124
2. Gao K, Weston S, Vytelingum P, et al. Deeper hedging: A new agent-based model for effective deep hedging[C]//Proceedings of the Fourth ACM International Conference on AI in Finance. 2023: 270-278. https://doi.org/10.1145/3604237.3626913
3. Heston S L. A closed-form solution for options with stochastic volatility with applications to bond and currency options[J]. The review of financial studies, 1993, 6(2): 327-343.
   https://doi.org/10.1093/rfs/6.2.327

Appendix: Additional Results and Specifications

Table 8. Hyperparameter Specifications and Sensitivity Analysis

Parameter	Description	Value	Sensitivity
State Dimension	Latent volatility factors	64	Medium
GRU Hidden Units	Transition network size	256	Low
TCN Receptive Field	Maximum historical context	1 day	High
Attention Heads	Multi-head attention	8	Low
Learning Rate	Adam optimizer	0.001	Medium
Batch Size	Training mini-batches	32	Low
Sequence Length	Training sequences	1000	High
Dropout Rate	Regularization	0.1	Medium
Lambda	Transaction cost weight	0.01	High

Table 9. Statistical Significance of Performance Differences

Performance Metric	Deep Filter	Heston	t-statistic	p-value
Hedging Error Mean	0.008	0.015	4.23	¡0.001
Hedging Error Variance	0.142	0.328	5.67	¡0.001
Sharpe Ratio	1.68	0.92	3.89	0.002
Maximum Drawdown	0.082	0.187	4.12	¡0.001
VaR (95%)	0.021	0.048	4.78	¡0.001
Expected Shortfall	0.035	0.072	4.45	¡0.001

Table 10. Ablation Study: Component Importance

Model Variant	Hedging Error	Performance Drop	Component Importance
Full Model	0.142	-	-
Without Attention	0.187	31.7%	High
Without Multi-scale	0.168	18.3%	Medium
Without Transaction Costs	0.149	4.9%	Low
State Dimension 32	0.156	9.9%	Low
State Dimension 128	0.141	-0.7%	Low
MLP Observer Only	0.234	64.8%	High
Kalman Filter Baseline	0.328	131.0%	-