Option Pricing Based on Implied Volatility Forecasting and No-Arbitrage Projection Correction: Evidence from SSE 50ETF Options

Manyi Liu[1]

[1] Beijing-Dublin International College, Beijing University of Technology, Beijing, 100124, China
liumanyi@emails.bjut.edu.cn

Abstract. As option markets expand, pricing and risk assessment grow increasingly critical for investor decisions and price discovery. In practice, predictive accuracy and no-arbitrage constraints are often hard to satisfy simultaneously. Using SSE 50ETF options as samples, this paper proposes a two-stage prediction-projection framework. First, XGBoost is applied in the implied volatility space to learn the volatility surface shape, followed by rigorous out-of-sample prediction. Predicted implied volatilities are deterministically mapped to option prices via the Black-Scholes model to generate a price surface. In the price space, projection corrections are imposed on three static no-arbitrage conditions via minimizing squared deviations from predicted prices. Empirical results show implied volatility prediction errors are broadly stable, with no noticeable increase after price mapping. However, point predictions alone lead to clear no-arbitrage violations, which are mostly eliminated after correction, accompanied by reduced pricing errors. Further one-day Delta hedging tests reveal that for highly liquid samples, corrected hedging performance outperforms prediction-only results, indicating the correction ensures constraint compliance and more stable risk measurement. Robustness checks confirm the conclusion's reliability. This paper develops an intuitive framework integrating statistical prediction with financial constraints and proposes an embeddable quality-control scheme for pricing and risk-management systems.

Keywords: Option Pricing, Implied Volatility Forecasting, XGBoost, No-Arbitrage Projection.

1 Introduction

In recent years, the options market has continued to expand, and option prices not only affect investors' hedging and asset allocation strategies but also largely reflect how the market prices uncertainty, thereby being closely related to the efficiency of price discovery. Therefore, obtaining option pricing and forecasting results that are stable, interpretable, and applicable to risk management remains a key issue in both derivatives research and practical applications. Against this background, the Black–

© The Author(s) 2026
L. C. Yoong (ed.), *Proceedings of the 2026 2nd International Conference on Financial Innovation and Marketing Management (FIMM 2026)*, Advances in Economics, Business and Management Research 408,
https://doi.org/10.2991/978-94-6239-801-6_37

Scholes model, as a classical pricing framework, establishes a clear mapping between prices and volatility and has long been used in practice as a tool for computing implied volatility and risk sensitivities. However, when dealing with complex surface shapes and significant data noise, relying only on point predictions or local fitting often causes price structures to violate basic financial constraints. This motivates the study of a modeling framework that maintains prediction accuracy, follows static no-arbitrage rules, and improves the stability of surface outputs for risk measurement.

Existing studies have examined implied volatility surface forecasting from the perspective of machine learning and high-dimensional features, pointing out that time-series feature engineering may lead to potential information leakage and thus bias evaluation results [1]. Therefore, testing must be conducted strictly under out-of-sample conditions. A related body of work has also focused on the construction and correction of no-arbitrage structures, including the generation and fitting of arbitrage-free volatility surfaces [2, 3], as well as arbitrage filtering strategies for real-time market quotes [4]. However, for Chinese ETF options, a research gap remains in empirical studies employing an integrated framework of statistical prediction, price mapping, and no-arbitrage projection in the price space, while conducting closed-loop evaluation using forecast errors, no-arbitrage violation measures, and hedging performance. Based on this, taking SSE 50ETF options as an example, this paper uses XGBoost to forecast the implied volatility surface, maps the forecasts into the price space through the Black–Scholes model, and then applies least-squares projection correction to the predicted price curve under boundary, monotonicity, and convexity constraints. It then compares, under a strict out-of-sample setting, the improvements of this model over a prediction-only model in terms of error levels, no-arbitrage consistency, and the stability of one-day Delta hedging. This paper not only enriches the theory of option surface modeling under a framework that separates statistical prediction from economic constraints, but also improves in practice the usability of surface forecasts and the consistency of risk control.

2 Theory and Model

2.1 Black-Scholes Mapping and the Implied Volatility Measure

This paper takes the cross-sectional quotes of SSE 50ETF options within each trading day as the object of analysis. To obtain a unified surface representation suitable for modeling, this paper does not treat the Black–Scholes analytical model as the true data-generating mechanism of the market, but rather uses it as a deterministic mapping tool linking option prices and volatility. This makes it possible to characterize the quote structure of contracts with different strike prices and times to maturity on a common scale, while also providing a comparable basis for subsequent out-of-sample evaluation.

Implied volatility is defined as the parameter establishing a one-to-one correspondence between price and volatility: given (S_t, K, τ, r_t, q_t), it is defined as the value of σ that makes the Black-Scholes price equal to the observed option price, where S_t denotes the underlying price, r_t the risk-free interest rate, q_t the dividend

yield or equivalent dividend adjustment, and σ the volatility parameter. To maintain consistency with the subsequent out-of-sample error evaluation and hedging tests, this paper uses the closing price as the observed price in the benchmark specification, and then compares it with settlement-price-based results in the robustness checks. The implied volatility is then defined as model (1):

$$\widehat{\sigma_t}(K,\tau) = \arg\min_{\sigma>0} \left| C_t^{BS}(K,\tau;\sigma) - C_t^{obs}(K,\tau) \right| \tag{1}$$

This representation transforms the complex price surface into a smoother, more learnable implied volatility surface and provides a unified interface for mapping forecasts back to the price space.

2.2 Static No-Arbitrage Constraints

If a study relies solely on machine learning to generate point forecasts of implied volatility and then maps them back into the price space, the resulting prices are highly likely to violate static no-arbitrage conditions [5]. This may manifest as violations of price bounds, non-monotonic relationships with strike prices, or non-convexity in strike prices, creating opportunities for deterministic-profit strategies such as vertical spread arbitrage or butterfly arbitrage [6]. Thus, under fixed maturity, three key strike-dimension constraints are prioritized: price bounds, strike monotonicity, and strike convexity. Expressed in difference form on a discrete strike grid, these conditions are characterized by linear inequalities, turning no-arbitrage correction into an implementable convex optimization or quadratic programming problem. Cross-maturity static no-arbitrage constraints are reserved for extended analysis due to complexity and interpretability concerns [7].

This paper uses a two-stage framework of prediction and projection. In the first stage, the study learns the surface structure in the implied volatility space and performs strict out-of-sample prediction [8]. In the second stage, the study maps the predicted results back to the price space and projects them into a no-arbitrage feasible set to remove static arbitrage violations and obtain surface outputs for pricing and risk control [9]. Unlike generative approaches that directly construct arbitrage-free implied volatility surfaces, this framework separates statistical prediction from economic constraints. This makes it possible to clearly see how much each component contributes, namely, prediction error and no-arbitrage correction, rather than having the two confounded with each other.

3 Data Processing and Model Construction

3.1 Data Sources and Research Sample

This paper focuses on SSE 50 ETF options traded on the Shanghai Stock Exchange, with a sample period spanning from February 9, 2015, to February 26, 2026. These options, first listed on February 9, 2015, are written on the SSE 50 ETF with ticker 510050, available in both call and put forms, each contract covering 10,000 units of

the underlying. Option trading records and pricing parameters are drawn from the CSMAR database, merged with daily ETF quotes, the China Bond Treasury yield curve, and dividend data, from which a daily panel at the contract level is constructed. To align model inputs and evaluation metrics, rolling volatility features are constructed from daily ETF returns to reflect underlying price behavior and market conditions [10]. Variable definitions are summarized in Table 1.

Table 1. Table captions should be placed above the tables.

Data module	Frequency	Summary of core variables and their uses
Daily option trading data	Daily	Prices, trading volume, open interest, and trading status are used for pricing inputs, liquidity filtering, and robustness specifications
Option contract and pricing parameters	Daily	Strike price, maturity date, risk-free interest rate, implied volatility, and Greeks; used to construct (K, T, τ, r), supervised target IV, and hedging or sensitivity measures
Daily market data of the underlying ETF (510050)	Daily	Closing price and returns, used to construct log moneyness $m = \ln(S/K)$, nd rolling volatility features
Government bond yield curve	Daily	Curve type, maturity nodes, and yields; used to construct an alternative risk-free rate $r(t, \tau)$ through maturity matching or interpolation
ETF dividends	Event-based	Ex-dividend date and dividend per unit; used for robustness of dividend specifications

This paper first selects the daily records of option contracts written on the SSE 50ETF as the underlying asset and then further applies filters on trading validity and contract life: only formal trading records are retained, and observations with remaining time to maturity $\tau \leq 0$ are excluded. The final sample is merged by trading date, contract, and call or put type to obtain the core panel used for subsequent forecasting and no-arbitrage projection, containing 340,214 observations and 4,858 contracts.

3.2 Variable Definitions, Feature Construction, and the Supervised Target

First, the core input variables are log moneyness and time to maturity. Let S_t denote the closing price of the underlying ETF on trading day t, and let K_i denote the strike price of option contract i; log moneyness is then defined as the natural logarithm of the ratio of the closing price to the strike price. The maturity dimension is measured by the annualized remaining time to maturity τ. The pair (m, τ) forms the cross-sectional coordinates of the implied volatility surface and serves as the core input variable set in the forecasting stage. As shown in Fig. 1, this paper plots a heatmap of sample coverage in the (m, τ) feature space, and the results indicate that the sample is mainly concentrated around at-the-money regions and short to medium-term maturities, thereby providing effective data support for subsequent implied volatility surface learning and out-of-sample evaluation.

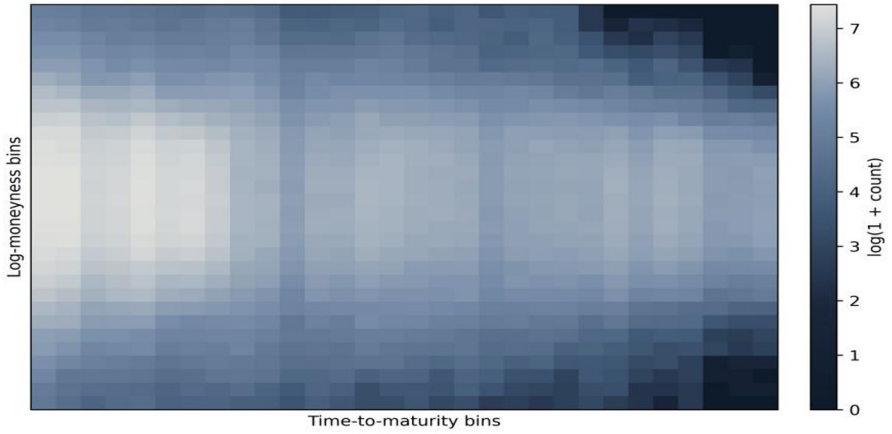


Fig. 1. Sample coverage distribution on log moneyness and expiry time (Picture credit: Original).

Second, the definition of the risk-free rate is matched with the term structure. The benchmark rate is the daily risk-free rate used in the option pricing parameters. In the robustness check, the study uses the government bond yield curve to construct an alternative interest rate measure. The published curve system includes yield-to-maturity, spot, and forward curves generated by interpolation models using data from market makers, brokers, and settlement prices, while accounting for liquidity constraints in sample selection. On each trading day, the remaining time to expiry τ is matched to the corresponding tenor node on the yield curve, and $r(t, \tau)$ is obtained via nearest-neighbor matching or linear interpolation for robustness checks.

Third, implementation details and robustness settings. As detailed in Section 2, predicted implied volatilities are mapped to option prices via the Black–Scholes model, after which projection correction is applied to enforce no-arbitrage consistency. Contract size conversion follows the standard units set by the Shanghai Stock Exchange for SSE 50 ETF options. Given that cash dividends can cause discrete price jumps in the underlying ETF, dividend events for the 510050 are collected and used in adjusted-versus-unadjusted comparisons as part of the robustness analysis.

4 Empirical Analysis of Model Performance

4.1 Experimental Setup and Evaluation Criteria

This section takes the 510050 daily option panel from Section 3 as the foundation. Following consistent data cleaning criteria, only contracts with standard trade types and positive time to expiry are retained, and the sample is split chronologically by trading date. The first 80% of trading days serve as the training period, with the remaining 20% as the strict out-of-sample test period starting on December 5, 2023.

For market implied volatility forecasting, contract-day observations are treated as discrete points on the volatility surface, with log moneyness and remaining maturity as core features. Basic contract information, risk-free interest rate, and other related variables are incorporated to enrich the prediction feature set. In implementation, XGBoost is adopted to capture nonlinear relationships in implied volatility forecasting. Implied volatility forecasting is known to be sensitive to data-splitting schemes and feature construction choices, both of which can introduce information leakage if not carefully controlled. Strict chronological splitting is therefore maintained throughout, and multiple specifications are tested to verify the robustness of the results.

Predicted implied volatilities are then mapped to option prices via the Black–Scholes formula and subjected to projection correction under the three constraints outlined in Section 2. The correction is applied separately for each option type and remaining-maturity bucket, with least-squares projection used to minimize deviations from the predicted prices. Three categories of evaluation metrics are used: IV errors, price errors, and no-arbitrage violation degree.

4.2 Out-of-Sample Forecasting Performance

The out-of-sample period yields approximately 68,000 contract-day observations in total. The empirical results show that the model delivers overall stable out-of-sample implied volatility forecasts, with an implied volatility MAE of approximately 0.049 and an RMSE of approximately 0.068. After mapping the predicted implied volatilities into the option price space, the pricing errors also remain within a reasonable and interpretable range, with a price MAE of about 0.296 and a price RMSE of about 0.452.

To clearly present the cross-sectional heterogeneity in errors, this paper further regroups the errors along three dimensions: remaining-maturity intervals, log-moneyness intervals, and option type. The grouped results reveal a fairly clear overall pattern. Implied volatility errors are more concentrated in regions with short remaining maturities and deep out-of-the-money contracts far from at-the-money. This pattern largely reflects the market characteristics inherent to that region, rather than being an artifact of model fitting. The region tends to exhibit heavier quote noise and sharper volatility surface curvature, and training observations are relatively sparse there, all of which contribute to the weaker fit.

4.3 Static No-Arbitrage Violations and the Effect of Projection Correction

Looking at the three constraints in price space, namely monotonicity, convexity, and boundary conditions, the prediction-only surface shows clear violations across all three, with convexity and boundary conditions being the most problematic. This is not surprising. Point predictions do not embed any structural constraints, and given that market quotes are inherently noisy and discrete, such violations are difficult to avoid. Once the projection step is introduced, both the rate and magnitude of violations drop sharply, approaching zero. This shows that projection correction can bring the surface

back in line with fundamental financial constraints without touching the prediction framework.

The analysis then turns to how projection correction affects price fitting errors. Looking at the projected sample, price errors do not worsen relative to the prediction-only model; they actually come down slightly. The mean absolute error drops from around 0.296 to 0.264, and the root mean squared error from around 0.452 to 0.440. This suggests that the projection step substantially improves no-arbitrage consistency without any meaningful loss in price-space fitting accuracy.

4.4 Economic Significance Test

Beyond statistical error and no-arbitrage structure, this paper also conducts a more practice-oriented test using single-day Delta hedging, placing the model in a realistic trading setting to assess whether it can meaningfully reduce risk for investors. For each contract on each trading day, a hedge portfolio is formed using the Delta estimated on that day, and the hedging error is then computed from the next-day price changes of the underlying and the option.

This paper compares the hedging performance of three Delta measures. The first is market Delta, derived directly from market quotes, which serves as the benchmark. The second is predicted Delta, computed from model-predicted implied volatilities. The third is projected Delta, obtained by first backing out implied volatilities from the projection-corrected prices and then computing Delta accordingly.

Across the full sample, market Delta produces the smallest hedging errors, which is not surprising given that it is derived directly from market quotes and thus naturally reflects prevailing market conditions. In the high-liquidity subsample, projected Delta yields smaller hedging errors than predicted Delta. This indicates that the projection step not only brings the surface closer to no-arbitrage compliance but also leads to more stable risk management outcomes. Structural correction goes beyond satisfying mathematical consistency; it also carries real economic value in practice. As shown in Fig. 2, the distribution of single-day hedging errors for projected Delta is more concentrated than that of predicted Delta in the high-liquidity subsample.

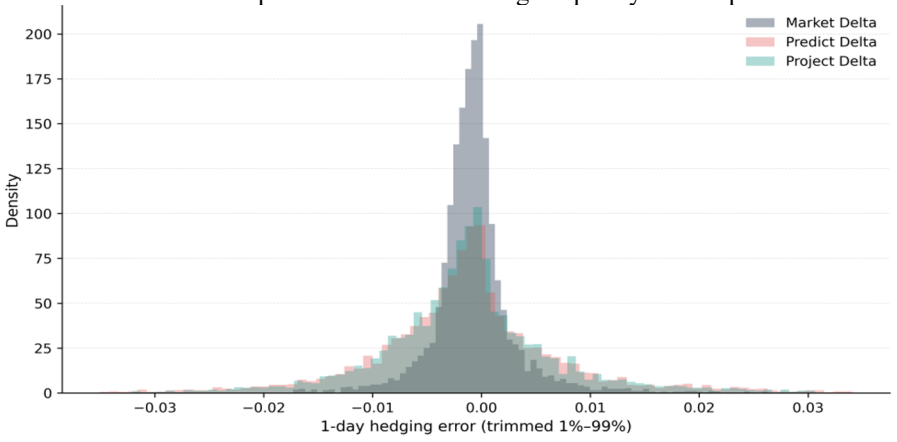


Fig. 2. Distribution of one-day Delta hedging errors (Picture credit: Original).

5 Robustness Checks

Since implied volatility forecasts on time-series data are sensitive to sample splitting and feature construction, robustness checks focus on three core aspects—temporal splitting, price specification, and sample selection—with additional comparisons for dividend events that may trigger discrete underlying price jumps. First, the time partitioning criterion is modified: the baseline 80/20 training-test split is compared with a 70/30 split to test model stability over a longer out-of-sample period. Second, the price specification is replaced, using settlement prices instead of closing prices to rule out the impact of price measurement differences. Third, liquidity filtering is tightened by requiring positive trading volume and open interest to verify whether results are driven by low-liquidity contracts. For dividend events, two controls are designed: excluding five trading days around ex-dividend dates and adjusting log moneyness for continuity, testing the stability of price-space conclusions.

Table 2 reports the out-of-sample errors under different specifications. Under the benchmark specification, both IV errors and price errors remain at stable levels. After replacing the price specification, the errors remain almost unchanged, and the results are likewise stable after imposing liquidity filters. Under the 70/30 split, price errors increase noticeably, indicating that when a longer test period covers a broader range of market-state changes, price-space errors become more sensitive to intertemporal drift; this pattern is consistent with recent discussions on how time splitting and distribution shift affect IV forecasting evaluation. Excluding the five trading days before and after each ex-dividend date leaves price errors almost unchanged, indicating that the price-space conclusions in Section 4 are independent of ex-dividend samples. Furthermore, after applying an ex-dividend continuity adjustment to the log-moneyness input, price errors also remain stable. The increase in IV errors is consistent with the label specification: because market IV is backed out from the contemporaneous underlying-price specification, changing the definition of log moneyness at the input stage weakens the fit to the original labels, but does not alter the stability of price-space errors. Overall, dividend events do not change the conclusions of the predict and project framework.

Table 2. Robustness checks.

Specification	Training Share	Test Start Date	n_{test}	IV MAE	IV RMSE	Price MAE	Price RMSE
Comparison of out-of-sample errors (IV and price space)							
Benchmark	0.8	2023-12-05	68,008	0.049	0.067	0.296	0.452
Alternative split	0.7	2022-10-31	102,569	0.045	0.064	0.325	0.503
Alternative price specification	0.8	2023-12-05	68,008	0.049	0.067	0.296	0.452
Liquidity filter	0.8	2023-12-05	67,708	0.049	0.067	0.295	0.451
Effect of ex-dividend treatment on out-of-sample errors							
Benchmark	-	2023-12-05	68,008	0.049	0.067	0.296	0.452

Exclusion of the ex-dividend event window	-	2023-12-06	64,515	0.050	0.069	0.296	0.452
Ex-dividend adjustment	-	2023-12-05	68,008	0.071	0.097	0.296	0.452

6 Conclusion

Using the SSE 50ETF options as the research object, this paper develops an operational two-stage framework to address the challenge of jointly achieving predictability and no-arbitrage consistency in implied volatility surface forecasting. The framework first forecasts the surface in the implied volatility space, then maps it into the price space via the Black–Scholes model, and finally applies projection correction under three static no-arbitrage constraints. Out-of-sample results show forecast errors remain stable and reasonable after price mapping. Point predictions alone lead to clear no-arbitrage violations, which are mostly eliminated by projection correction, with pricing errors further reduced. Economic significance tests via one-day Delta hedging indicate that in high-liquidity samples, corrected hedging performance outperforms the prediction-only model, confirming that the correction enhances both no-arbitrage consistency and risk control stability. Multiple robustness checks verify the conclusion's reliability.

This paper's contributions lie in two aspects: first, it defines the Black–Scholes model as a deterministic mapping mechanism between implied volatility and prices, enabling unified comparison of surface forecasting results in the price space and avoiding over-reliance on implied volatility while neglecting practical pricing deviations. Second, it introduces a price-space projection-correction mechanism, providing a feasible solution to structural arbitrage violations in machine learning-based point forecasts. The study's limitations include unaddressed cross-maturity no-arbitrage constraints and simplified hedging tests that exclude market frictions. Future research may integrate cross-maturity constraints, extend the evaluation framework to friction-rich environments, and explore liquidity-weighted projection objectives to enhance model stability in extreme or low-liquidity markets.

References

1. Arratia, A., El Daou, M., Kagerhuber, J., et al.: Examining challenges in implied volatility forecasting: a critical review of data leakage and feature engineering combined with high-complexity models. Computational economics 1–23 (2025).
2. Campbell, T. C., Gallmeyer, M., Petkevich, A.: The only constant is change: nonconstant volatility and implied volatility spreads. Journal of financial and quantitative analysis 58(5), 2190–2227 (2023).
3. Cont, R., Vuletić, M.: Simulation of arbitrage-free implied volatility surfaces. Applied mathematical finance 30(2), 94–121 (2023).
4. Guterding, D.: Sparse modeling approach to the arbitrage-free interpolation of plain-vanilla option prices and implied volatilities. Risks 11(5), 83 (2023).

5. Medvedev, N., Wang, Z.: Multistep forecast of the implied volatility surface using deep learning. Journal of futures markets 42(4), 645–667 (2022).
6. Moussa, K.: Arbitrage filtering of option prices: a simple real-time approach. Quantitative finance 1–16 (2025).
7. Ning, B., Jaimungal, S., Zhang, X., et al.: Arbitrage-free implied volatility surface generation with variational autoencoders. SIAM journal on financial mathematics 14(4), 1004–1027 (2023).
8. Pallotta, A., Ciciretti, V.: Should you use GARCH models for forecasting volatility? A comparison to GRU neural networks. Studies in nonlinear dynamics & econometrics 28(5), 725–738 (2024).
9. Vuletić, M., Cont, R.: VolGAN: A generative model for arbitrage-free implied volatility surfaces. Applied mathematical finance 31(4), 203–238 (2025).
10. Zhang, W., Li, L., Zhang, G.: A two-step framework for arbitrage-free prediction of the implied volatility surface. Quantitative finance 23(1), 21–34 (2023).

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

