

Uncovering Conditional Alpha: A Divide-and-Conquer Approach to Volatility Trading in Crypto-Proxy Markets

Jiayi Zhou

School of Economics, Shanghai University of Finance and Economics, Shanghai, China
ethan.z@stu.sufe.edu.cn

Abstract. This study develops a comprehensive framework for evaluating volatility trading strategies in crypto-proxy option markets under realistic market conditions. A high-fidelity backtesting environment is constructed that explicitly incorporates commissions, bid–ask spreads, slippage, and liquidity limits, ensuring that reported performance metrics reflect true implementable outcomes rather than theoretical gains. The methodological progression begins with a simple slope-based trading rule and extends to generalized machine learning classifiers and reinforcement learning agents. Empirical results demonstrate that the baseline heuristic and universal machine learning models quickly lose profitability once realistic frictions are introduced, highlighting the fragility of universal alpha in frictional markets. By contrast, sector-specific approaches uncover conditional profitability that remains robust under practical constraints. In particular, XGBoost yields positive returns in Bitcoin spot exchange-traded funds, while reinforcement learning agents generate consistent gains in mining and semiconductor sectors. These results collectively support a conditional efficiency paradigm, whereby market inefficiencies are not broad-based but localized, context-dependent, and exploitable only through specialized, adaptive modeling techniques. The framework thus contributes to both academic research and practical trading strategy design, offering insights into how algorithmic trading in emerging digital asset markets must evolve to remain viable in the presence of persistent market frictions.

Keywords: Implied Volatility Surface, Crypto-Proxy Markets, Machine Learning, Reinforcement Learning

1. Introduction

The rise of cryptocurrency-linked exchange-traded products has created a complex ecosystem of “crypto-proxies,” including Bitcoin spot and futures ETFs, mining corporations, and technology firms with digital asset exposure. These instruments exhibit high volatility and sentiment sensitivity but also face wide bid–ask spreads and liquidity constraints, raising doubts about the feasibility of volatility-based strategies.

Extensive research has examined the predictive content of the implied volatility surface, focusing on term structure slope and skewness [1,2]. However, many anomalies collapse once trading frictions are included [3]. This tension is particularly relevant in emerging crypto-proxy markets, where frictions are more severe than in established asset classes.

This study develops a high-fidelity backtesting framework that embeds commissions, spreads, slippage, and liquidity limits into performance evaluation. Three paradigms are tested: a slope-based rule, machine learning classifiers, and reinforcement learning agents. Results show that the baseline strategy fails under costs, generalized machine learning models cannot overcome heterogeneity, and conditional alpha appears only in sector-specific contexts. XGBoost achieves profitability in Bitcoin spot ETFs, while reinforcement learning succeeds in mining and semiconductor sectors.

The contribution is both methodological and empirical: embedding frictions ensures credible evaluation, and the findings support a conditional efficiency paradigm, where inefficiencies exist but only in localized, structurally specific markets [4].

2. Data and methodology

2.1. Data and asset universe

The analysis uses 14 U.S.-listed optionable assets with material cryptocurrency exposure, including Bitcoin spot ETFs, futures ETFs, mining corporations, exchanges, and semiconductors as a control group. This classification reflects the hypothesis that predictability is conditional on sectoral drivers. End-of-day option chain data provide implied volatility and trading activity measures, consistent with established practices in option pricing research [5].

2.2. Feature engineering

Four features summarize the volatility surface: the at-the-money implied volatility of the nearest maturity, the slope of the term structure, the skew between out-of-the-money calls and puts, and volume-based metrics. Skewness, often linked to net buying pressure, has been shown to influence option demand [6]. Together, these features capture structural and behavioral aspects of option markets. Figures 1 and 2 illustrate the temporal dynamics and correlations among features.




Figure 1. Temporal dynamics of volatility surface features across sectors




Figure 2. Correlation structure of volatility surface features

2.3. Backtesting framework

The backtesting engine simulates one-day straddle trades with liquidation at the next close. Net return is defined as:

$$r_t = \frac{PnL_{t,net}}{P_{t-1}} \quad (1)$$

Transaction costs include commissions, spreads, and slippage:

$$C(N, P) = (4 \cdot N \cdot C_{fee}) + (2 \cdot N \cdot P \cdot M) \cdot (\beta_{spread} + \delta_{slippage}) \quad (2)$$

Liquidity constraints cap trade sizes at a fraction of daily volume, ensuring institutional feasibility.

2.4. Modeling approaches

Rule-based strategy. The baseline model trades on the slope signal, formally introduced in the results section.

Machine learning. XGBoost and multilayer perceptrons are trained to predict the sign of next-day changes in implied volatility. Following best practices in empirical asset pricing via machine learning [7], the binary cross-entropy loss is minimized:

$$\mathcal{L}(y, \hat{p}) = -\frac{1}{N} \sum_{i=1}^N \left[y_i \log(\hat{p}_i) + (1 - y_i) \log(1 - \hat{p}_i) \right] \quad (3)$$

Feature importance results are reported in Figure 3.




Figure 3. Feature importance rankings from sector-specific XGBoost models

Reinforcement learning. Sector-specific deep Q-networks are deployed, reflecting the suitability of RL for sequential policies in volatile markets [8]. The reward function is detailed in the results section.

3. Empirical results and analysis

3.1. Baseline strategy

The baseline strategy tests whether the slope of the implied volatility term structure provides tradable alpha. The rule is defined as:

$$Signal_t = sign(Slope_t) \quad (4)$$

Applied across all sectors, the strategy produces negative Sharpe ratios and persistent drawdowns once costs are considered. Figure 4 shows cumulative equity curves, where the baseline consistently underperforms. This confirms that slope-based rules, often profitable in frictionless tests, collapse under realistic frictions [1, 3].




Figure 4. Cumulative equity curve comparison across strategies

3.2. Machine learning models

Generalized XGBoost and multilayer perceptron models fail to deliver profitability across the heterogeneous asset universe. This supports the view that predictive relationships in option surfaces are not universal [7].

Sector-specific training, however, reveals conditional alpha. For Bitcoin spot ETFs, an XGBoost classifier generates positive risk-adjusted returns, robust to transaction costs. Feature importance analysis (Figure 3) highlights slope and skew as dominant drivers in ETFs, while equities rely more on volume metrics. In mining and exchange sectors, results remain unprofitable, reinforcing the conditional nature of alpha.

3.3. Reinforcement learning models

Reinforcement learning agents, implemented as deep Q-networks, show superior performance in sectors with path-dependent volatility dynamics. The reward function guiding training is:

$$r_{t+1} = (a_t - 1) \cdot \Delta IV_{t+1} \quad (5)$$

where $a_t \in \{0,1,2\}$ represents short, neutral, or long exposure.

Results indicate Sharpe ratios above unity for mining and semiconductor sectors, even after costs, as summarized in Figure 5. Performance in Bitcoin spot ETFs is weaker, reflecting that sequential adaptation adds limited value in markets dominated by direct underlying volatility. These outcomes align with advances in deep reinforcement learning, which highlight its ability to uncover sequential policies [8].




Figure 5. Performance metrics summary for baseline, ML, and RL models

3.4. Interpretation

The results yield three key insights. First, the slope-based baseline fails comprehensively, rejecting its value as a universal alpha factor. Second, machine learning confirms that volatility surface features only have predictive power in select sectors, supporting the conditional efficiency view [9]. Third, reinforcement learning agents succeed where static models do not, particularly in markets shaped by complex, path-dependent dynamics.

Overall, the evidence demonstrates that alpha in crypto-proxy markets is both fragile and conditional. Inefficiencies persist, but only in specific institutional and structural contexts, and can be exploited only through specialized, context-aware modeling.

4. Discussion

The findings highlight the fragility of volatility-based strategies under realistic frictions. The slope of the implied volatility term structure, though long studied, fails to yield tradable alpha once

transaction costs and liquidity limits are considered. This outcome mirrors earlier evidence that theoretical option models often diverge from empirical performance [5].

Volatility surface features such as skew further illustrate conditionality. While skew has been linked to order flow imbalances and net buying pressure, its predictive power is evident only in Bitcoin spot ETFs, where volatility directly reflects cryptocurrency demand [6]. In mining and exchange sectors, skew loses relevance amid stronger idiosyncratic influences.

The weak performance of generalized machine learning models supports the critique of universal predictors [10]. By contrast, sector-specific models succeed by tailoring weights to structural contexts, consistent with evidence that empirical asset pricing via machine learning requires careful specialization [7].

Reinforcement learning agents achieve the most robust results, uncovering sequential policies in mining and semiconductor assets. This aligns with advances in deep reinforcement learning, which emphasize adaptability to path-dependent dynamics [8].

Taken together, the results reinforce a conditional efficiency paradigm: inefficiencies exist, but only in local, sector-specific contexts, demanding rigorous backtesting and adaptive modeling for credible exploitation [4].

5. Conclusion

This study introduced a methodological framework for evaluating volatility trading strategies in crypto-proxy markets under realistic conditions. By embedding commissions, spreads, slippage, and liquidity constraints into a high-fidelity backtesting engine, the framework ensures that reported returns reflect implementable performance rather than artifacts of frictionless models.

The evidence delivers three key conclusions. First, the slope of the volatility term structure, though widely cited, fails to generate positive risk-adjusted returns once frictions are considered, echoing concerns over the fragility of simple strategies. Second, generalized machine learning models cannot overcome costs when trained across heterogeneous assets, consistent with warnings about the limitations of universal predictors. Third, conditional alpha emerges when models are specialized: XGBoost identifies profitability in Bitcoin spot ETFs, while reinforcement learning agents, aligned with advances in sequential decision-making, succeed in mining and semiconductor sectors.

These results support a conditional efficiency paradigm, where inefficiencies persist but only within particular institutional and structural contexts. Volatility features such as slope and skew provide value selectively, and their effectiveness is shaped by market frictions and structural dynamics.

Methodologically, the study demonstrates the necessity of friction-adjusted backtesting for credible evaluation. Empirically, it shows that sustainable alpha requires sectoral specialization and adaptive modeling rather than reliance on universal trading rules.

Future research should extend the feature set and explore advanced reinforcement learning architectures, but the main lesson is clear: alpha in crypto-proxy markets is conditional, localized, and accessible only through rigorous, context-aware approaches.

References

- [1] Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*, 50(2), 125–150. [https://doi.org/10.1016/S0304-405X\(98\)00034-8](https://doi.org/10.1016/S0304-405X(98)00034-8)
- [2] Jiang, G., & Tian, Y. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4), 1305–1342. <https://doi.org/10.1093/rfs/hhi033>

- [3] Harvey, C. R., & Liu, Y. (2019). Lucky factors. *Journal of Financial Economics*, 133(2), 353–378. <https://doi.org/10.1016/j.jfineco.2019.02.008>
- [4] Cochrane, J. H. (2011). Presidential address: Discount rates. *Journal of Finance*, 66(4), 1047–1108. <https://doi.org/10.1111/j.1540-6261.2011.01671.x>
- [5] Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, 52(5), 2003–2049. <https://doi.org/10.1111/j.1540-6261.1997.tb02724.x>
- [6] Bollen, N. P. B., & Whaley, R. E. (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance*, 59(2), 711–753. <https://doi.org/10.1111/j.1540-6261.2004.00645.x>
- [7] Gu, S., Kelly, B., & Xiu, D. (2016). Empirical asset pricing via machine learning. Chicago Booth Research Paper No. 16-19. <https://doi.org/10.2139/ssrn.2839283>
- [8] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- [9] Kelly, B., Kuznetsov, B., Malamud, S., & Xu, T. A. (2023). Deep Learning from Implied Volatility Surfaces. Working paper, Swiss Finance Institute.
- [10] International Review of Financial Analysis. (2024). Machine-learning stock market volatility: Predictability, drivers, and economic value. *International Review of Financial Analysis*, 94, Article 103286. <https://doi.org/10.1016/j.irfa.2024.103286>