## Executive Summary

Since the introduction of the Black and Scholes (1973) and Merton (1973) BSM pricing model, the modelling of the implied volatility (IV) of an option has remained an important and challenging task in the financial sector. Understanding the dynamics underlying the movements in the implied volatility of an option provides insight into the factors influencing the market's expectations of future price movements of an option's underlying asset and is crucial in managing the price risk of an option.

Traditionally, most attempts at modelling IV have relied on a priori assumptions about its dynamics. In this paper, our aim is to contribute to the emerging field of applying machine learning methods to address classic financial problems. We build upon the work of Cao et al. (2018) by investigating the suitability of machine learning methods, which do not require a pre-determined functional form, for forecasting changes in IV. In addition, we evaluate the effectiveness of incorporating these methods into a minimum variance hedging strategy, providing insights into the economic application of these models.

The research paper is structured into three stages. In the first stage, we investigate the predictive capabilities of machine learning models in forecasting changes in IV at specific points on a standardised IV surface. We begin by parameterising an eSSVI IV surface for each date using market data of options written on S&P 500 futures. In response to challenges arising from significant discontinuities between call and put options near-the-money, we propose a novel midpoint method for anchoring the slicewise parameterisation of the surface. The new anchoring method appears to enhance the robustness of the parameterisation, providing satisfactory surfaces for our study. Our research reviews three models: a neural network, a random forest, and an XGBoost model. These models consider a mix of macroeconomic and option-specific predictors in their decision making process. The models' hyperparameters are chosen through a time-series cross-validation, and the best models of each type are used to predict daily changes in the IV at specific points on the IV surface. The performance of these models is compared with the performance of a naive predictor, and an adapted version of the Hull and White (2017) model. The machine learning models significantly outperform the naive predictor, with comparable performance to the Hull-White model.

In the second stage of our research, we explore the predictive capabilities of the machine learning

models in forecasting daily changes in the IV of specific options. The daily change in IV of an option is calculated as the change in IV of the same option on consecutive trading days. In this experiment, the machine learning models continue to outperform the naive predictor, although the neural network's performance declines significantly in comparison with the other models. The XGBoost model demonstrates the strongest performance of the machine learning models, maintaining similar results to the Hull-White model. Using the XGBoost model, we further identify the six month interest rate spread, the level of the IV relative to the preceding days, and the past returns of the underlying asset as important predictors. The resulting predictions suggest that IV exhibits a short term mean reversion effect. Furthermore, the results indicate that the state of economic growth is crucial in determining the magnitude of changes in IV.

In the final stage of our study, we aim to determine whether the incorporation of the machine learning predictions can enhance a minimum variance hedging strategy. We incorporate the predictions from the previous stage into an adaption of the minimum variance framework suggested by Hull and White (2017). The distributions of daily hedging errors for each model are compared against those produced by the BSM delta and Hull-White minimum variance delta hedging strategies. While the performance of the machine learning models is somewhat constrained by the ad-hoc incorporation of the predictions into the minimum variance strategy, the XGBoost model demonstrates a hedging strategy characterised by lower variance and downside variance of its hedging errors than the BSM strategy. This observation is supported by both an F-test and a Bartlett test. The results of the Levene test are less conclusive. Moreover, the XGBoost also results in smaller extreme gains and losses than all the other models. The neural network, by contrast, provides little improvement on the BSM model.

Overall, this study demonstrates the value of incorporating machine learning methods into the IV modelling problem. The XGBoost model, in particular, shows consistently strong performance in all of the tasks considered in the study. As a result of the study, we have not only reaffirmed some characteristics of implied volatility, such as the mean-reversion effect, but we have also presented new insights into the behaviour of IV. The findings of the paper show the potential and value of incorporating machine learning in solving financial problems.