CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Economic Studies

Volatility Risk Premium Across Multiple Asset Classes

Masters's thesis

Author: Bc. Štěpán Havel

Study program: Economy and finance

Supervisor: doc. PhDr. Jozef Baruník, Ph.D.

Year of defense: 2024

Declaration of Authorship

The author hereby declares that he compiled this thesis independently, using only the listed resources and literature, and the thesis has not been used to obtain any other academic title.

The author grants to Charles University permission to reproduce and to distribute copies of this thesis in whole or in part and agrees with the thesis being used for study and scientific purposes.

In Prague
April 30, 2024

Bc. Štěpán Havel

Abstract

This paper analyzes the utilization of volatility risk premium for volatility forecasting across 18 distinct markets in two time frame periods. We show that VRP can be used as a viable tool for volatility forecasting. Our research suggests that RIV models, which adjust current implied volatility for volatility risk premium, can yield significantly more accurate volatility forecasts compared to other forecasting models. We examine various methodologies for deriving RIV models, assessing their effectiveness and limitations. Our findings hold true for the majority of assets analyzed and are supported by various loss functions and other statistical tests, underscoring the robustness of our conclusions.

JEL Classification	C22, C53, G13
Keywords	Volatility Risk Premium, Volatility Forecasting

Title	Volatility Risk Premium Across Multiple Asset Classes

Author's e-mail	havel.stepan@gmail.com
Supervisor's e-mail	barunik@fsv.cuni.cz

Abstrakt

Tato práce analyzuje využití rizikové prémie za volatilitu pro předpovídání volatility na 18 odlišných trzích, ve dvou časových obdobích. Ukazuje, že VRP lze použít jako vhodný prostředek k předpovídání volatility. Výzkum naznačuje, že RIV modely, které upravují současnou implikovanou volatilitu o rizikovou prémii, dosahují výrazně přesnějších predikcí budoucí volatility ve srovnání s jinými modely. Práce zkoumá různé metody výpočtu RIV modelů, hodnotí jejich efektivitu a hranice. Závěry výzkumu platí pro většinu zkoumaných aktiv a jsou podloženy různými statistickými testy.

Klasifikace JEL	C22, C53, G13
Klíčová slova	Riziková Prémie Volatility, Predikce Volatility

Název práce	Riziková Prémie Volatility Napříč Ruznými Druhy Aktivy

E-mail autora	havel.stepan@gmail.com
E-mail vedoucího práce	barunik@fsv.cuni.cz

Acknowledgments

Here I would like to express my sincere gratitude to doc. PhDr. Jozef Baruník, Ph.D., for his expertise, patience, useful comments and willingness to help. I also thank my family for their unwavering encouragement, and my friends for their support. Special thanks to the Institute of Economic Studies for fostering a conducive academic environment without which this paper could not have been created.

Typeset in FSV LaTeX template with great thanks to prof. Zuzana Havrankova and prof. Tomas Havranek of Institute of Economic Studies, Faculty of Social Sciences, Charles University.

Bibliographic Record

Bc. Štěpán Havel: *Volatility Risk Premium Across Multiple Asset Classes*. Masters's thesis. Charles University, Faculty of Social Sciences, Institute of Economic Studies, Prague. 2024, pages 81. Advisor: doc. PhDr. Jozef Baruník, Ph.D.

Contents

List of Tables

List of Figures

Acronyms

ATM	At the Money
ARVRP	Average Relative Volatility Risk Premium
CBOE	Chicago Board of Option Exchange
DM	Diebold-Mariano
EMA	Exponential Moving Average
IV	Implied Volatility
ITM	In the Money
LTCM	Long-Term Capital Management
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MSE	Mean Squared Error
MSPE	Mean Squared Percentage Error
OTM	Out of the Money
RIV	Risk Premium Adjusted Implied Volatility
RV	Realized Volatility
RVRP	Relative Volatility Risk Premium
SPX	S&P 500 Index
VRP	Volatility Risk Premium

Masters's Thesis Proposal

Author	Bc. Štěpán Havel
Supervisor	doc. PhDr. Jozef Baruník, Ph.D.
Proposed topic	Volatility Risk Premium Across Multiple Asset Classes

Motivation and research question The efficient market hypothesis suggests that there is no way to consistently achieve superior trading results. Risk and reward should be perfectly symmetrical, giving all investments an expected value of zero. This intuitively suggests volatility risk premium (VRP) might be zero. However, for several reasons, this is not the case and the VRP is persistently positive.

We most commonly measure the volatility risk premium as the difference between the implied volatility (IV) and subsequent realized volatility (RV) of the underlying security. The volatility risk premium refers to the phenomenon where option implied volatility tends to exceed realized volatility of the same asset over time. This possibly creates profit opportunity for volatility sellers as buyers keep overpaying for options.

The aim of this paper is to conduct analysis of VRP of multiple assets from different classes – mainly commodities. There are several papers shoving that currency VRP has substantial predictive power for the cross-section of currency returns. However, literature exploring similar dependencies in commodity markets seems limited. Therefore, the thesis will try to find connection between volatility risk premium and commodity returns. Since positive dependencies have been found in the currency market, we expect to find similar connections in the commodities market. The magnitude is however unknown in advance. The paper will further study the magnitude of premiums paid through time and occurrences where VRP turned out to be negative.

Contribution Having a proper grasp of volatility risk premium gives us a better understanding of options markets. Why sellers of volatility keep demanding premiums for their issued insurance through options and why buyers keep paying it. The size and persistence of this insurance risk premium reflects behavioral factors (loss aversion), structural factors (imbalance of natural longs and shorts), and economic factors (added value of convex assets with negative correlations).

The literature studying VRP usually deals with only one asset class. In the paper, we will conduct analysis of volatility risk premium of multiple different assets (equities, commodities, currencies etc.). The paper's goal is to examine which asset class tends to exhibit the largest VRP. Opposite to stock market, where the volatility is connected mainly to equity risk, we see different types of risks causing increased volatility. For example, volatility in natural gas prices may be caused by weather or volatility in agricultural prices can be caused by anticipated hurricane. All of which is mostly uncorrelated to equity risk.

Methodolody Main data set consists of calculated VRP multiple asset classes. The paper will work with representatives from the currency, equity and commodity markets. Volatility indexes of oil, gold, silver, corn, wheat, soybean, S&P 500 index and others. Close to close realized volatility and intraday realized volatility will be calculated and compared to the implied volatility by the volatility index. Through this study we will be able to analyze where we find some markets where investors keep overpaying for insurance. Another part of the paper will be conducting regression analysis whether the VPR of multiple assets have any predicting power with subsequent return of commodities.

Preliminary Outline

1. Introduction
2. Theoretical framework
    - a Literature review
    - b Options
        - i Basics (+ Greeks)
        - ii Option Pricing
    - c Volatility
        - i Volatility Risk Premium
        - ii Implied Volatility
    - d Efficient market hypothesis
3. Empirical research
    - a Data description
    - b Predictive power of VRP
    - c VRP Analysis
    - d Results
4. Conclusion

Core Bibliography

1. Bennett, C. (2014): Trading Volatility: Trading volatility, Correlation, Term Structure and Skew

2. Eraker, B. (2009): The Volatility Premium. Wisconsin School of Business, University of Wisconsin

3. Goyal, A., Saretto, A. (2009) Cross-section of option returns and volatility. Journal of Financial Economics 97:310–326

4. McMillan, L.G. (1980): Options as a strategic investment : Lawrence G. McMillan, (New York Institute of Finance, New York)

5. Han, B., Cao, J. (2013) Cross section of option returns and idiosyncratic stock volatility. Journal of Financial Economics 108(1):231–249

6. Lorondo, J., Zhou, H. (2012) Variance Risk Premiums and the Forward Premium Puzzle International Finance Discussion Papers No.:1068

Chapter 1

Introduction

Volatility risk premium refers to the phenomenon where option implied volatility tends to be priced higher than what subsequent realized volatility would justify. The existence of positive VRP has been largely documented, and imply, that option buyers tend to overpay for protection against volatility fluctuations (Trolle & Schwartz 2009; Han 2011; Lu 2020). There could be several reasons why market participants are not able to price options correctly such as supply and demand dynamics or liquidity conditions. Nevertheless, recognizing potential discrepancies in option prices is crucial in the realm of risk management, portfolio allocation or asset pricing. While numerous studies have validated the existence of the volatility risk premium, its practical utilization remains a subject of ongoing exploration. Volatility forecasting has significant attention within the domain of financial markets. Yet, comparatively less focus is directed towards leveraging the knowledge of volatility risk premium for such forecasting purposes. Prior research has mainly focused on a limited number of markets and the information content of volatility risk premium across a wider array of assets is still unexplored.

Implied volatility is the current market expectation of future volatility, and we are able to extract it at any time from the option prices. However, this estimate can be seen as biased by the size of VRP. This paper is the first to analyze the information content of volatility risk premium for volatility forecasting on 18 distinct markets. Previous studies showed that the knowledge of VRP can be beneficial for this task, although this behavior was yet to be analyzed across a broader spectrum of assets (Prokopczuk & Simen 2014). Our findings suggest that the knowledge of VRP is a statistically meaningful metric for volatility forecasting and deliver new insight into its possible employment.

In our paper, we adjust the current implied volatility for the past volatility risk premium, aiming to improve its forecasting accuracy of realized volatility. In this manner, we obtain the risk adjusted implied volatility, introducing a novel forecasting technique. We refer to the forecasting methodologies based on this approach as RIV models. In extension, we propose a new procedure for extracting historical volatility premiums, which utilizes the fact that volatility is heteroskedastic. We perform in sample and out of sample forecasts and assess the forecasting accuracy by employing various loss functions and Diebold-Mariano tests. Our findings indicate that RIV models exhibit greater effectiveness in volatility forecasting compared to both implied volatility and GARCH models. These conclusions hold true for the majority of assets analyzed.

Furthermore, we illustrate that utilizing the ratio of implied and realized volatility produces volatility risk premiums which can be compared between assets with varying volatility levels. Additionally, the distribution of such volatility risk premium more closely resembles a normal distribution. This observation could lead to improvements of financial models reliant on the normality assumption. Lastly, we demonstrate that volume appears to have a direct effect on the volatility risk premium, and that illiquid markets can be the cause for its unusual behavior.

The rest of the paper is structured as follows. First, we introduce the literature review related to our work. Then, we proceed with the development of the theoretical framework. We start with a brief introduction of options and volatility. We further familiarize the reader with the concept of volatility risk premium and present the statistical procedures employed in the paper. We follow with the empirical research, where we firstly present volatility risk premiums for each asset and display different methodologies for calculation of the VRP. We also analyze the effect of traded volume on volatility risk premium and describe the process of adjusting implied volatility for the past premiums. Last, we present the test results for both forecasting windows.

Chapter 2

Literature Review

2.1 Volatility Risk Premium

The fact that volatility of equity returns is stochastic has a firm ground in the realm of in financial economics. However, less work has been given to study whether volatility risk is compensated for. There are several evidences that market volatility risk premium might me nonzero. Jackwerth & Rubinstein (1996) showed that at-the-money implied volatilities of S&P 500 are consistently and systematically higher than realized volatilities. An economic interpretation of these findings is that market participants buying volatility are willing to pay a premium for downside protection. Further, using a sample of S&P 500 options, Bakshi & Kapadia (2003) examined the statistical properties of the delta-hedged option portfolios to show existence of VRP.

Carr & Wu (2009) measured the VRP as a difference between the realized volatility and the variance swap rate, which is identical to implied volatility. They found that the VRP was strongly positive for the S&P 500 and Dow Jones indices. In other words, the counterparty of the swap got compensated for bearing the risk of uncertain variance. Bollerslev et al. (2011) further proposed a method for constructing a volatility risk premium index utilizing the sample moments of model-free realized and option-implied volatility measures. Further they showed that the extracted volatility risk premium helps to predict future stock market returns.

There are also several investment strategies focusing on harvesting the volatility risk premium. It has been shown that the VRP of S&P 500 ranges from 5 to 10 percentage points during regular market conditions (Feri 2017). Zhou (2018) further showed that VRP can help to predict equity, bond or currency returns over short horizons and the analysis of Lu (2020) documents the existence of volatility risk premium under the Black Scholes model.

2.2 Volatility Forecasting

Early research on the predicting power of implied volatility (IV) on realized volatility (RV) found that IV explains variation in future volatilities better than RV (Latane & Rendleman 1976; Beckers 1981). However, further study conducted on S&P100 index options found, that despite the fact that the IVs and future RVs are related, the implied volatility has little to no correlation with future volatility, and it does not incorporate the information contained in recent observed volatility. (Canina & Figlewski 1993). Christensen & Prabhala (1988) also found that implied volatility outperforms past volatility in forecasting future volatility. Furthermore, it also subsumes the information content of past volatility in some of the model specifications. This change in results might have been caused by the shift in the option pricing methods after the 1987 stock market crash (Hull 2003). Jorion (1995) further focuses on the information content and predictive power of implied standard deviations of currency option.

Szakmary et al. (2003) followed with using data from 35 futures option markets with a conclusion that the implied volatility outperforms the historical volatilities as a predictor of RV for a large majority of the commodities studied. In addition, they also showed that GARCH forecasts were not superior to those of IV. However, these findings are not supported by Agnolucci (2009) who studied the predictive power of IV and GARCH models for crude oil futures. Christensen & Hansen (2010) further conclude that implied volatility is an efficient estimator of realized return volatility on S&P100 index.

Prokopczuk & Simen (2014) study the role of the volatility risk premium for the forecasting performance of implied volatility on energy markets. They try to improve prediction power of implied volatility on future volatility by accounting for the volatility risk premium. They find that adjusting current implied volatility for the past average volatility risk premium results in better forecasts of realized volatility compared to implied volatility.

Diavatopoulos et al. (2020) examines the characteristics of real estate investment trust equity options and finds that implied volatility is significantly related to future realized volatility. Chen & Li (2023) estimate that approximately one third of the predictive power of implied volatility on future realized volatility can be attributed to its ability to predict the arrival of both scheduled and unscheduled news.

Chapter 3

Theoretical Framework

This chapter presents the key theory behind the empirical research. We start with the introduction of options, their history, fundamental theory and pricing models. We follow with establishing the concept of volatility, where we show the differences between historical and implied volatility. We also introduce different methodologies used to obtain them. In the following section, we introduce the concept of volatility risk premium, defined as the difference between implied volatility and realized volatility. This also includes the introduction of RIV models, which utilize the past volatility risk premium for volatility forecasting and are of main attention during the empirical research. Further, we provide an introduction to futures as a financial instrument and we conclude this chapter with an overview of the statistical procedures utilized in the paper.

3.1 Options

An option is a security giving the buyer the right to buy or sell an asset, subject to certain conditions, within a specified period of time. A call option grants the holder the right to buy at the strike price, while a put option gives the buyer the right to sell at the strike price. The price, that the option sells for, is referred to as the premium. We can differentiate between European and American options based on whether the right can be exercised solely at the end of the option period or throughout the entire option life, respectively. The determinants of an option contract include the expiration date, the strike price, the underlying asset, and the quantity of the asset subject to purchase or sale (Black & Scholes 1973).

The option is considered exercised when the holder decides to utilize their right to convert the option into the purchase or sale of the underlying asset. In-the-money (ITM) options are those where the strike price is favorable compared to the current market price. On the contrary, out-of-the-money (OTM) options have strike prices less favorable than the current market price. Options are said to be at-the-money (ATM), if the strike price is equal to the current market price.

Although American type option, such as a stock option, can be exercised at any time during the life of the contract, it will be ultimately exercised at the end of its expiracy, when it is ITM and possesses some intrinsic value. For the call option to be ITM, the price of the underlying would need to exceed the strike price. If that is the case, the holder is able to buy at cheaper price than the market price. His counterparty, the seller of the option, would occur loss as he would be obliged to honor the trade and sell the asset at the strike price. If the opposite happens, and the options expires OTM, the underlining right to exercise the option has no value and the option expires worthless. In this case, the option seller keeps the premium that was paid by the buyer for the option. Put options work similarly to call options, where they instead provide the right to sell at the strike price. Therefore, the put options work in the opposite direction than call options. A holder of call option would profit from a price increase of the underlying while a holder of put option benefits from a price decrease of the underlying.

3.1.1 History

One of the first attempts to describe traded options was written by Castelli (1877). His work also included a description of some commonly used trading and hedging strategies such as call-of-more and the call-and-put.[1] The foundations of option pricing can be traced back to the 19th century, when French mathematician Louis Bachelier introduced the concept of Brownian motion in his doctoral thesis, which later became the basis for modeling price movements in financial markets. In 1900, Bachelier first employed higher mathematics to price options (Davis & Etheridge 2006). Although the paper was an interesting academic study, it had little to none practical applications as no organized option market existed at the time. In the same year the French mathematician Paul Lévy introduced the concept of a geometric Brownian motion, which

[1] These strategies are today known as a covered write and a straddle.

became a fundamental building block in the development of option pricing models.

In 1973, the Chicago Board Option Exchange (CBOE) opened. It was the first exchange to list standardized, exchange traded stock options. During that time, academics tried to build on the work of Bachelier to develop a first theoretical model for option pricing. In the same year the famous Black-Scholes model was introduced by economists Fischer Black and Myron Scholes (Black & Scholes 1973).[2] This model revolutionized option pricing by providing a closed-form solution to valuing European options. Although the Black-Scholes model still remains widely used, more sophisticated models have emerged since such as the binomial model or the stochastic volatility option pricing model (Cox 1979; Shephard 2010).

3.1.2 Fundamentals

Certain terms describe the relationship between the option strike price K and the price of the underlying S. The price of an option can be divided in two components. The intrinsic value and the time value. An option has intrinsic value when it is trading ITM and its exercising can provide the holder an access to better prices than are currently offered at the market.

$$\text{Option Price} = \text{Intrinsic Value} + \text{Time Value} \tag{3.1}$$

The time value of an option is the amount by which the option price exceeds its intrinsic value. Normally, an option has the highest time value when it is trading ATM. In this scenario, a small move of the underlying results in the option being either profitable of worthless. The final outcome is in this case the most uncertain, which is expressed by the large time value. As an option becomes deeply ITM or OTM, its time value decreases substantially. That is because it is already known, with a decent probability, whether the option will expire worthless or not. The term "time value" refers to the fact that part of an option's premium is completely due to the time remaining until its expiration. At the option's expiracy, there will be no more time value left and the whole price of the option will be credited to intrinsic value.

Intrinsic value is dependent only on the relationship between stock price S and strike price K. The intrinsic value of an ITM call is the amount by which

[2] Robert Merton is also often credited as being the first one to publish an academic paper on the topic.

the stock price S exceeds the striking price K. Generally we can say that the intrinsic value of ITM options is defined as the absolute difference between the stock price and the strike price. The intrinsic value of OTM options is zero.

$$\text{Intrinsic Value}_{ITM} = |S - K| \tag{3.2}$$

$$\text{Intrinsic Value}_{OTM} = 0 \tag{3.3}$$

Option price is dependent on several other factors. Although each of them contributes to the overall price with different magnitudes, they all play an important role. Besides the strike price, the expected price fluctuations of the underlying have the largest effect on an option premium. More volatile assets have more room to appreciate or depreciate until the expiration, thus creating larger chance of price moves which could make OTM options be ITM and the opposite. For that reason, sellers of options will require larger prices for the potential risk. Expected volatility of the underlying effects puts and calls similarly.

The payment of dividends has an indirect impact on the option through the price of the underlying. We know that on the ex-dividend day, the stock price will decrease by the dividend amount. With this anticipation, option prices adjust in order to accommodate for this expected price decline. The effect of dividend on option premium can be profound in stocks with high dividend yield, otherwise it is negligible. Lastly, the effect of interest rate on option prices is minuscule, however cannot be omitted. Holding all else equal, call option premiums generally rise when interest rates increase, and put option premiums generally decrease when interest rates increase.[3]

3.1.3 Pricing Models

Since options are widely used on both stocks and futures contracts, a proper valuation technique was desired by both academics and practitioners. The option pricing is most commonly performed by the application of the Black-

[3] One stock option controls 100 shares of the underlying. Thus, holding stock option can serve as a decent substitute to holding the stock. Buying 100 shares would be much more costly than buying a single option. Therefore, one could decide to buy cheaper call option and invest the rest at the risk free rate. In rising rates environment the opportunity cost of buying the stock outright increases. That results in either decrease of the number of option sellers, increase of the number of option buyers or both. Ultimately the price of calls increases. The opposite analogy holds true for puts.

Scholes formula in the Equation 3.4.

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0 \tag{3.4}$$

Variable V represents the option price as a function of stock price S, r is the continuously compounded risk-free rate, t is time and σ is the standard deviation of the underlying returns. The key insight of the formula is that a participant on the financial markets is able to perfectly hedge an option by buying or selling the underlying asset and holding cash in such a way that completely eliminates risk. The optimal combination of the underlying and cash holdings changes as the underlying fluctuates in price and must be therefore adjusted continuously in order to eliminate risk. This method of continuously adjusting the underlying position is called continuously revised delta hedging. This hedge implies that there is only one solution for the price of the option. By solving the equation for the corresponding terminal and boundary condition, the Black-Scholes formula calculates the price of European put and call option (Black & Scholes 1973). The price of a call C and put P option for a non-dividend paying underlying is defined by the Black-Scholes model as:

$$C(S,t) = SN(d_1) - Ke^{-rt}N(d_2) \tag{3.5}$$

$$P(S,t) = -SN(-d_1) + Ke^{-rt}N(-d_2) \tag{3.6}$$

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}} \tag{3.7}$$

$$d_2 = d_1 - \sigma\sqrt{t} \tag{3.8}$$

Variable S is the price of the underlying, K is the strike price, t is the time remaining until expiration, expressed as a percentage of a year, r is current risk-free rate, σ is the future volatility measured by annual standard deviation and $N(x)$ is the cumulative normal density function. Even though the model is widely used due to its simplicity, it assumes ideal market conditions, which are not usually met in reality. Most importantly, it assumes the returns to be normally distributed. However, the returns of financial time series are often observed to exhibit certain statistical properties that deviate from a normal distribution such as fat tails, volatility clustering or non-stationarity. Further, it assumes no transaction costs in buying or selling the stock or the option.

The formula is also only defined for European options and does not penalize short selling. Nevertheless, it is still a great guide to explain how changes of the underlying variables effect the option prices (Mandelbrot 1963; Fama 1965).

3.2 Volatility

Volatility is the term used to describe how fast a stock, futures or index changes in price. It is denoted by the Greek letter σ and is usually expressed on yearly basis. Normally, higher volatility of the underlying is associated with greater risk and a greater potential for returns. When one speaks of volatility in connection to options, there are two types of volatility that are important. The first is the historical volatility, which is a measure of how fast the underlying had been changing in price in the past. The second is the implied volatility, which is the option market's prediction of the volatility of the underlying over the life of the option contract (McMillan 1980).

3.2.1 Realized Volatility

Realized volatility is defined as the standard deviation of continuously compounded logarithmic returns. The logarithmic transformation helps to mitigate the effect of exponential growth in the markets and linearise the time series. In the case of assets with large volatilities, the differences between arithmetic and geometric returns differ the most (Hudson & Gregoriou 2010). The daily rate of continuously compounded log returns r_t is calculated using the daily closing prices[4] c_t and ordinary dividend d_t by:

$$r_t = \log\left(\frac{c_t + d_t}{c_{t-1}}\right) \tag{3.9}$$

We now utilize the daily returns r_t to calculate the historical volatility. The size of volatility in different periods depends on the events that were occuring in the markets in the corresponding periods. Historical volatility utilizing time series of returns for n observations is defined as follows:

$$\sigma = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(r_i - \bar{r})^2} \tag{3.10}$$

[4]Closing prices are commonly used when dealing with daily data. However, any other time series of discrete observations can be used.

Variable $\bar{r}$ represents the expected value of r_t. In order to have the ability to compare between volatilities of different time lengths, we often express volatility on yearly basis. This would be done by dividing the volatility σ_t by the square root of total days in the time frame.[5]

$$\frac{\sigma_t}{\sqrt{t}} = \sigma_{annually} \tag{3.11}$$

Another feature of volatility is that it is additive. The volatility of a longer time frame can be approximated by summing the volatilities of its constituent shorter time frames as expressed in Equation 3.12. This feature of volatility enables us to utilize several daily volatilities to estimate the volatility of a longer time frame. Most notably, we will be interested in 30 day historical volatility as it corresponds to the 30 day implied volatility calculated by volatility indexes such as VIX.

$$\sigma = \sqrt{\frac{252 \sum_{i=1}^{n} \sigma_i^2}{n}} \tag{3.12}$$

With the emergence of high-frequency price data, many new estimators of asset return volatility were constructed. With the usage of intraday data and Equation 3.10 we are able to calculate the daily realized volatility. Liu et al. (2012) studied the accuracy of a wide variety of estimators of asset price volatility constructed from high-frequency data, and compare them with a simple realized volatility estimator. They conclude that daily realized volatility calculated from 5-min data performs decently and little evidence is found that 5-min RV estimate is significantly outperformed by any of the other measures.

3.2.2 Implied Volatility

Implied volatility is a term specific to the realm of options. As opposed to historical volatility, where the number of historical days in the calculation changes the resulting output, there is only one fixed value of the implied volatility for a certain strike and expiration date. Implied volatility is the market expectation of future volatility derived from the price of an option contract.

In the Black-Scholes formula from Equations 3.5 and 3.6 all other variables besides the future volatility are known. One can thus input the current market price of an option in the formula and estimate the volatility which would be required to obtain the known price. This volatility is refereed to as the implied

[5] Expressed in years. It is common practice to assume 252 trading days in one year.

volatility. Thus, the size of the implied volatility will be dependent on the expectations of volatility of the underlying during the life of an option contract. More volatile assets will, on average, have higher implied volatilities.

We can look at the current implied volatility in a relation to its past values. If implied volatility is relatively high, the option will be deemed overpriced. That is, they will be relatively expensive. If, on the contrary, the implied volatility is too low, the option will appear as underpriced. However, in the literature we usually reference to these concepts as that the options are trading with a high or low implied volatility. Since we have some sense where implied volatility has been in the past we can compare the current implied volatility to the historical volatility. If market participants believe that the volatility of the underlying will be higher than currently implied by option prices, they view the option as cheaper than it should be. Therefore, they will bid up the option prices, making them more costly. This simultaneously results in an increase of option implied volatility as option price and implied volatility move in the same direction. The same argument would apply for expectations of lower volatility than implied by options.

It is important to note that market participants have no actual way of knowing, how volatile is the underlying going to be during the life of an option. Thus, every option trader is forced to make a guess about the future volatility when he buys or sells an option. The market participants are often interested just in the hedging capabilities of options. In that case, they just pay the price, which seem reasonable based on historical volatility or any other metric.

The Black Scholes option pricing formula utilizes several inputs to calculate the option price P and the majority of them are already known. These are the stock price S, strike price K, time to expiration t, interest rate r and dividends d. The only variable in the formula that we do not know ahead is the future volatility σ. Mathematically, this would be defined as:

$$P = f(S, K, t, r, d, \sigma) \tag{3.13}$$

We can use the current option price P and the Equation 3.13 to calculate the implied volatility σ of the option contract. From the option pricing formula, the values of implied volatilities should be independent on the strike price of the option. In theory, implied volatility should produce a straight line on a chart with moneyness on x axis and implied volatility on y axis. However, this is not what we observe on financial markets.

Figure 3.1: Volatility Surface

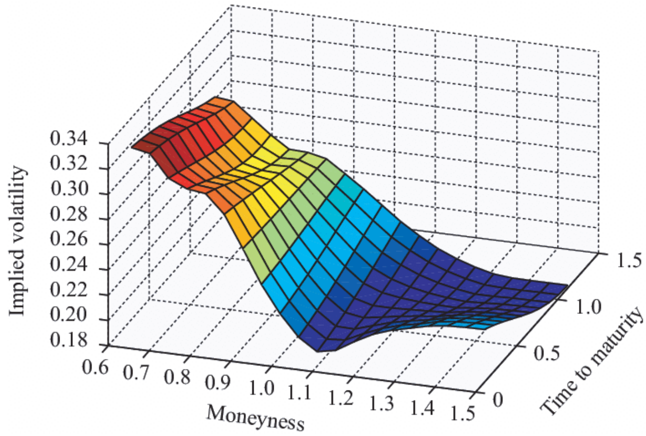

Source: Cont & Fonseca (2002)

The Figure 3.1 presents an implied volatility surface. On the z axis is the option moneynes. Moneyness of 1 means that the option is trading at the money. The x axis represents the time to maturity and y axis is the corresponding implied volatility. If we consider only one specific time to expiration, the resulting curve of implied volatility is very close to a shape of a convex curve.[6] Similarly, option implied volatilities will also differ depending on the time to maturity. If there is an anticipated event in a couple of months, which could have a profound effect on the price moves of the underlying, the options will be priced accordingly.

As the implied volatilites differ throughout the volatility surface it would be convenient to comprise all of the available information from the option chain. Since 1993, CBOE calculates the CBOE Volatility Index, commonly known as VIX, which is derived from options on S&P 500 index futures. VIX represents the 30 day expected volatility of the underlying, which is calculated as a weighted average of OTM call and put options on the S&P 500 futures.[7] The calculation process involves taking two strips of nearest monthly futures option expiration dates and obtaining the 30 day implied volatility by weighted averaging the two strips.

The exact VIX calculation is not crucial for this paper. Is is however important, that the same approach that is used to calculate the VIX index can be applied on any other market. The only condition is that continuous bid and ask offers are being made in the two strips of options in the two front months.

[6] As this shape resembles a smile, this feature is commonly referred to as the volatility smile.

[7] More precisely the expected square root of realized variance.

We can thereafter compare implied volatilities across different markets. Calmer markets will produce smaller implied volatilities while turbulent markets higher ones. As a result, a VIX-like calculation of volatility can be made for nearly every listed stock, index or futures contract. In our thesis we are mostly interested in the analysis of 30 day implied volatility of futures contracts as it is a common standardized measure.

3.3 Volatility Risk Premium

Volatility risk premium refers to the phenomenon where, on average, option implied volatility tends to exceed realized volatility of the same underlying asset. In theory, the volatility risk premium creates an opportunity to capitalize on the gap between implied and realized volatility and capture the premium through option trades. When comparing option implied volatilities with subsequent realized volatilities, it is important to notice that implied volatility is obtained under risk-neutral measure Q and realized volatility is is seen under the physical measure P. Therefore, it unfeasible to directly compare the two time series together with the assumption that volatility price is zero (Carr & Wu 2009). Furthermore, volatility risk premium is not constant but exhibits variability over time as demonstrated by Driessen et al. (2009) or Trolle & Schwartz (2009). According to Lu (2020), the difference between implied and realized volatility is most apparent in broad equity indexes like the S&P 500.

Bollerslev et al. (2009) approximate the market volatility risk premium $VRP_{t,T}$ between t and T as the difference between risk-neutral expectations and physical expectations of volatility as:

$$VRP_{t,T} = \mathbb{E}_t^Q(\sigma_{t,T}) - \mathbb{E}_t^P(\sigma_{t,T}) \tag{3.14}$$

$\mathbb{E}_t^Q(\sigma_{t,T})$ is the ex ante forecast of volatility under the risk-neutral measure Q, that is the equivalent of IV. $\mathbb{E}_t^P(\sigma_{t,T})$ is the ex ante forecast of volatility calculated under the physical measure P. This can be estimated as ex post RV.

If the volatility risk premium was positive, it would mean that traders on the short side of the trade made a theoretical profit. According to Bennett (2014), shorting volatility on equity markets should have a similar returns profile as going long equity. Since volatility is negatively correlated to equity returns, a short volatility position will result in long equity risk. As equity returns are expected to yield an equity risk premium over a risk free asset, shorting volatility should, by consequence, be also abnormally profitable. Thus, risk neutral implied volatility should be expected to be above the true realized volatility.

Due to upward sloping term structure, long term options tend to be most overpriced.[8] This can be explained by the demand for long term protection or the utilization of options in structured products. Long positions of investors or providers of annuity products need to have counterparts. The liquidity on the market is normally provided by market makers, who are able to remain market neutral through various types of hedging.[9] Since market makers need to be compensated for the costs of hedging, they price in an extra charge.

3.3.1 RIV Models

Instead of calculating the difference as outlined in Equation 3.14, an alternative approach involves computing the ratio between the expected volatility under the risk-neutral measure and the expectation of volatility under the physical measure (Trolle & Schwartz 2009). This ratio is defined as the relative volatility risk premium $RVRP_{t,T}$. The benefit of using the ratio instead on the difference is that it is independent of the level of volatility. Consequently, it enables a more accurate comparison of volatility premiums between time series that exhibit varying volatility levels. The $RVRP_{t,T}$ is defined as follows:

$$RVRP_{t,T} = \frac{\mathbb{E}_t^Q(\sigma_{t,T})}{\mathbb{E}_t^P(\sigma_{t,T})} \tag{3.15}$$

Volatility indexes strive to standardize implied volatility for a specific asset as the 30-day expected volatility. Consequently, we can compare the IV on day t with the subsequent RV over the next 30 calendar days. To accomplish this, we can employ Equations 3.14 and 3.15 for volatility risk premiums and define the 30-day volatility risk premium and 30-day relative volatility risk premium.[10]

[8] However, as more short term options can be sold in the same time as one long term option, greater returns can be captured through shorting near-dated products. This naturally comes with greater risk.

[9] Such as the previously mentioned delta hedging.

[10] It is important to note that 30-day volatility incorporates weekend days. In practice, we assume 21 trading days in one month.

$$VRP_t = IV_t - \sqrt{\frac{252\sum_{i=t}^{t+20} RV_i^2}{21}} \tag{3.16}$$

$$RVRP_t = \frac{IV_t}{\sqrt{\frac{252\sum_{i=t}^{t+20} RV_i^2}{21}}} \tag{3.17}$$

The study conducted by Prokopczuk & Simen (2014) has demonstrated that utilizing the past average of RVRP can enhance the predictive capability of implied volatility for realized volatility in energy markets. One of the research objectives of this paper is to examine whether similar outcomes hold true across a broader spectrum of securities. By averaging RVRP through a rolling window of length n, we define average relative volatility risk premium A-RVRP as:

$$A\text{-}RVRP_t = \frac{1}{n}\sum_{i=t-\tau-n}^{t-\tau} RVRP_i \tag{3.18}$$

An important characteristic of A-RVRP is its non-constant nature over time. During certain periods, there is a pattern of elevated average volatility premiums, which is subsequently succeeded by periods characterized by smaller average volatility premiums. The choice of the number of days, denoted as n, used to compute the average, is crucial. It should be long enough to ensure sufficient data points for an accurate estimation of the volatility risk premium. Conversely, if n is too large, recent market conditions may not be adequately captured. Experiments conducted to determine the appropriate selection of n concluded that a period of just under one year provides a reasonable tradeoff. In this thesis, we will adopt the assumption that 231 trading days correspond to this definition.

It is crucial to consider that the actual realized volatility, which is used for comparison with implied volatility, becomes known only after 21 trading days, the equivalent of 30 calendar days. Therefore, at the current date, we do not know the past 21 volatility risk premium observations for which we do not have the realized volatility yet. Consequently, the equation needs to be adjusted to account for this delay through the introduction of a parameter τ. When adding another 21 trading days for which we currently lack realized volatilities to the 231 days used for determining the past VRP average, we arrive at a total of 252 trading days. This corresponds to one calendar year. The averaging period for today thus starts one year ago and ends 30 days ago.

It can be shown that the magnitude of RVRP exhibits long term patterns similar to volatility clusters. This insight allows us to use the recent size of volatility risk premium to adjust the implied volatility to account for the embedded volatility premium. Consequently, we can assess whether the current implied volatility is appropriately priced given recent market developments. The implied volatility, adjusted with the recent sizes of volatility risk premium, is defined as the average risk premium adjusted implied volatility $A\text{-}RIV_t$.

$$A\text{-}RIV_t = \frac{IV_t}{A\text{-}RVRP_t} \tag{3.19}$$

It was shown that adjusting the implied volatility with A-RVRP tends to yield forecasts that significantly outperform commonly used forecasting methods in the analysis of energy markets. Given the autocorrelation tendencies observed in volatility, there is an interest in exploring whether a more refined selection of the averaging function can result in improved estimates (Engle & Bollerslev 1986). Instead of relying on the mean of volatility risk premium in the past year, we can give greater emphasis on recent volatility risk premiums compared to those in the more distant past. To address this, we propose employing an exponential moving average function to obtain the estimate of past volatility risk premium. As far as the authors' knowledge extends, this adjustment has not been implemented in any empirical paper and we are thus pioneers in utilizing this approach to enhance the predictive power of implied volatility on realized volatility.

The exponential moving average (EMA) places more weight on recent observations, offering a faster response to market changes. The EMA adapts swiftly to trends while maintaining a smoothed trend overview. In this framework we define the exponentially weighted average volatility risk premium E-RVRP as:

$$E\text{-}RVRP_t = \frac{\alpha}{\sum_{i=0}^{n-1} \alpha^i} \cdot \sum_{i=0}^{n-1} \alpha^i \cdot RVRP_{t-i-1-\tau} \tag{3.20}$$

Variable n defines the number of periods in the EMA window and α is the smoothing factor, usually calculated as $\alpha = \frac{2}{n+1}$. Similarly to Equation 3.18, we also have to shift the calculation by 21 days utilizing the parameter τ since we do not have realized volatilities for these dates. In the same manner as in Equation 3.19 we define the exponentially averaged risk premium adjusted implied volatility $E\text{-}RIV$ in Equation 3.21. In our paper we plan to examine whether the usage of exponentially weighted past volatility premium is able to produce better forecasts of future realized volatility than GARCH, IV and $A\text{-}RIV$ models.

$$E\text{-}RIV_t = \frac{IV_t}{E\text{-}RVRP_t} \tag{3.21}$$

Throughout the rest of the paper, we will collectively denote A-RIV and E-RIV as RIV models. Both of these estimation methods utilize the past volatility risk premium in adjusting current implied volatility. The initial letters, A or E, denote the underlying method for computing the long term average of the volatility risk premium. Specifically, 'A' denotes the simple average method, while 'E' represents the exponential averaging.

3.3.2 Negative VRP

As shown by many, the implied volatility often overestimates the future realized volatility (Carr & Wu 2009). There could be several reasons why option sellers demand a premium for taking volatility risk. However, situations where the future RV is larger than was predicted by implied volatility are common. They are just on average outweighed by periods where IV was higher. Nevertheless, if we concentrate on analyzing a long time frame period with enough observations, we might have a more accurate estimation whether IV and RV match. If they do not match, the difference would be either positive or negative volatility premium.

There are several reasons why the volatility risk premium could be negative in the short term. In calm and stable market conditions, investors may become complacent and expect low levels of future volatility. As a result, they might be willing to sell options at relatively low prices, which might underestimate the future price movements. Similarly, if investors believe that the risk of adverse events is small, they may not demand high enough premium for taking on volatility risk.

Also, in low interest rate environment, investors might start looking for additional income by selling options on assets they already own. This high supply of options would drive their price lower, underestimating the potential for market turbulence. In some cases, central banks or government policies aimed at stabilizing financial markets may lead to a perception of reduced risk, causing the volatility risk premium to turn negative. When investors believe that central banks will step in to prevent excessive market volatility, they may be less inclined to pay high premiums for options (Hull 2003; Gatheral 2006).

3.3.3 IV as Insurance

Since options are frequently used as a hedging instrument, there are many similarities between implied volatility and the insurance payments.[11] Both insurance and implied volatility involve probabilities and expected outcomes. Insurance companies calculate probabilities to determine appropriate insurance, while option pricing models utilize implied volatility to estimate the likelihood of various price movements. Thus, size of the payment for the protections is determined based on the perceived risk of the insured event occurring (Pasquale et al. 2016).

Normally, an insurance price could be determined as monthly or yearly payments. Typically of the same payment each period. The prices on the insurance market are similar between different companies, however there is no marketplace in which could the current demand and supply meet to produce the market price similarly to the option market. The prices thus not fully correspond to the perceived future risk as they do not update on current conditions. However, in the option market, the prices for protection, as expressed by implied volatility, change each day, constantly updating to reflect new information. This process thus results in always having the current market forecast for the future price moves. Therefore, IV oscillates around its theoretical mean values much more than insurance contracts do.

Implied volatility represents the best market expectation of future price fluctuations in the underlying asset. It reflects the collective wisdom of market participants regarding the uncertainty and potential volatility of the price movement of the underlying. Similar to how insurance premiums increase for riskier assets or events, implied volatility tends to rise in times of market uncertainty or when there are significant potential catalysts, such as earnings announcements or geopolitical events.

3.4 Futures

Futures contract is a standardized contract calling for the delivery of a specified quantity of a certain commodity at some future time. The primary economic

[11] Where price is determined based on the risk taken by the underwriter.

function of futures market is hedging. One can establish a futures position to offset some of the risk of actually owning the physical commodity. Futures contracts enable to lock in future prices at today's levels.

Futures contracts are listed on a wide variety of commodities and financial instruments. We can think of futures as a stock, which has an expiration date. Although they are both used for hedging purposes, we should not confuse them with options. Futures contracts can rise up substantially in price and can fall all the way to bottom, just as a stock can. On the contrary, with the ownership of option, the risk is limited. Therefore, the only real similarity is that they both have an expiration date. Nevertheless, both of the contracts provide a means to speculate on price movements, hedge against risk, and manage exposure in various asset classes.

Futures on physical commodities can be assigned, much like stock option can be assigned. However, they cannot be assigned at any time during their life, like American options can. Rather, there is a short period of time before the expiracy in which one take delivery. This period, sometimes referred to as the first notice, is usually around five weeks. Speculators, who do not wish to deal with the physical commodity, exit their position before this window. The rest of the trading is thereafter left to the actual hedgers dealing with the physical commodity.

In the futures market, one must be familiar with the specific terms of the contract. Two standardized future contracts on two markets could corresponds to completely different quantities of the underlying commodity. The Table A.2 summarizes the contract specifications for futures used in this paper. Further, one must know the methods of delivery if they expect to deal with the physical commodity. A significant aspect of futures is the concept of leverage. Market participants can initiate sizable positions with only fraction used as a collateral. This allows them to control a larger exposure to the underlying asset than their initial investment would typically allow (McMillan 1980).

3.4.1 Options on Futures

The commercial use of the option on futures is to lock in the worst case price, similarly like on the stock option market. A futures option is an option on the future contract, not on the underlying commodity. The options are always for one contract of the underlying commodity.

The actual dollar cost of the futures contract is not necessary for the futures option holder. He is more interested in knowing the breaking points where his option becomes profitable. If the futures in not settled by delivery of the physical commodity, it is settled in cash. This would be the case of futures on indices, where options and futures generally expire simultaneously at the end of the last trading day. However, options on physical commodities will expire before the first notice day of the futures contract. This has the effect that futures options often expire in the month preceding the month used to describe it (McMillan 1980).

In this paper we utilize the implied volatility of futures contracts to calculate the size of volatility risk premium in different markets. It is the implied volatility of futures options that enables us to compare the expected market move of the underlying futures with the actual realized move of the futures.

3.5 Statistical Procedures

3.5.1 Volatility Forecasting

Volatility forecasting models play a crucial role in financial markets as they help traders and investors anticipate and manage risk. The models assist market participants in estimating and predicting future volatility, allowing them to make informed decisions regarding risk management, portfolio allocation, and the pricing of derivative instruments. These models leverage historical data, market information, and statistical techniques to provide valuable insights into the dynamic nature of volatility in financial markets. There are several popular volatility forecasting models employed in practice. In this section we will talk about historical volatility models, implied volatility models and ARCH family models. Though there are many more volatility forecasting models, we will only discuss these three as they are the most important for this paper.

3.5.1.1 Historical Volatility

Historical volatility models are a class of models used to forecast future volatility based on past price movements of an asset. These models rely on historical data to calculate and analyze the realized volatility, which represents the actual level of price fluctuations over a given period. Historical volatility models typically use statistical techniques such as standard deviation or variance to quantify the dispersion of historical price returns. By examining past volatility patterns, these models aim to identify trends, cycles, or recurring patterns that can be used to make projections about future volatility. While historical volatility models provide a straightforward and intuitive approach to volatility forecasting, they have limitations, including the assumption that historical patterns will persist in the future and the inability to capture sudden shifts or structural changes in volatility.

3.5.1.2 Implied Volatility

The implied volatility models focus on using option prices to infer market expectations of future volatility. Since option price reflects the current anticipation of future volatility, the model takes the current market option implied volatility and uses it as a forecast of future realized volatility. The IV is said to be a better estimator of future volatility than historical volatility, as it also includes the current market conditions. Traders often compare the current implied volatility with the subsequently realized volatility to identify potential mispricings.

In our paper we further utilize the RIV models, which are simple model free implied volatilities adjusted for the past volatility risk premium. These are the average risk premium adjusted implied volatility A-RIV defined in Equation 3.19 and exponentially averaged risk premium adjusted implied volatility E-RIV from Equation 3.21.

3.5.1.3 GARCH

Engle (1982) introduced a new class of stochastic processes called autoregressive conditional heteroskedastic (ARCH) processes. These are zero mean, serially uncorrelated processes with nonconstant variances conditional on the past, but constant unconditional variances. Bollerslev (1986) continued and extended the ARCH model into generalized autoregressive conditional heteroskedasticity (GARCH) model. It permits the conditional variance σ_t^2 to rely not only on the preceding error terms ε_{t-j}^2, but also on the preceding values of its own lag σ_{t-j}^2. The GARCH(p,q) model, where p is the order of the GARCH terms σ^2 and q is the order of the ARCH terms ε^2, is given by:

$$r_t = \mu + \varepsilon_t \tag{3.22}$$

$$\varepsilon_t | \mathcal{F}_{t-1} \sim N(0, \sigma_t^2) \tag{3.23}$$

$$\sigma_t^2 = \alpha_0 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 \tag{3.24}$$

$\mathcal{F}_{t-1}$ denotes the information set up to time $t-1$ and also $\{p, q\} \geq 0$. Since variance is positive it must also hold $\{\alpha_i, \beta_j\} \geq 0$, $\forall i, j \in \{0, \ldots, q\}$. Parameter $\alpha = \sum_{i=1}^{q} \alpha_i$ measures the extent to which a volatility shock today feeds into the volatility of the next period. It usually ranges from 0.05 (stable market) to 0.1 (volatile market). Parameter $\beta = \sum_{j=1}^{p} \beta_j$ usually ranges between 0.85 and 0.98, opposite to α. Sum $\alpha+\beta$ measures the rate at which persistence of volatility dies over time. As shown by Chan (2010), persistence of volatility which does not die over time occurs when this summation is equal to one. Under this scenario, unconditional variance becomes infinite. For this reason the parameters should also be further constrained such that $\sum_{i=1}^{q} \alpha_i + \sum_{j=1}^{p} \beta_j < 1$ in order to imply weak stationarity (Engle 1995; Alexander 2008).

While the GARCH model effectively captures volatility clusters, it lacks the ability to distinguish between positive and negative shocks, as it solely considers squared residuals. In reality, volatility tends to escalate more rapidly in response to negative returns than positive ones. Several non-linear models, including EGARCH and TGARCH, have been introduced to address this limitation. However, empirical evidence suggests that these alternative models often fail to yield significantly improved results in estimating conditional volatility compared to the basic GARCH(1,1) model (Hansen & Lunde 2005).

3.5.2 Model Evaluation

In this paper we evaluate the forecasting errors with four commonly used loss functions. We employ the mean absolute error (MAE), mean squared error (MSE), mean absolute percentage error (MAPE) and mean squared percentage error (MSPE). These loss functions are defined as follows:

$$MAE = \frac{1}{n} \sum_{i=1}^{n} |RV_i - f_i| \tag{3.25}$$

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (RV_i - f_i)^2 \tag{3.26}$$

$$MAPE = \frac{1}{n} \sum_{i=1}^{n} \left| \frac{RV_i - f_i}{RV_i} \right| \tag{3.27}$$

$$MSPE = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{RV_i - f_i}{RV_i} \right)^2 \tag{3.28}$$

Variable n denotes the length of the forecasting window, RV_i is the realized volatility and f_i is the volatility forecast obtained from volatility forecasting model. A model which better reflects reality should minimize the loss function. In general, all loss functions could deliver similar results with small differences based on the underlying data. It is therefore important to consider all of them to draw a more refined conclusion.

3.5.3 Diebold-Mariano Test

While MAE and similar loss functions are frequently employed for evaluating forecast accuracy, there is also interest in a more rigorous statistical evaluation of differences in forecast. Diebold & Mariano (1995) introduced a test which compares the forecast accuracy of two forecasting methods. The Diebold-Mariano (DM) test considers the forecast errors of both models across multiple periods to determine if there exists a notable difference in their forecast accuracy. The DM statistic for testing the null hypothesis of equal forecast accuracy is defined as:

$$DM = \frac{\bar{d}}{\sqrt{\frac{2\pi \hat{f}_d(0)}{T}}} \tag{3.29}$$

where $\bar{d}$ is the sample mean loss differential and $\hat{f}_d(0)$ is the consistent estimate of the spectral density of the loss differential at frequency zero. For further details we refer the reader to the original paper. Under the null hypothesis, the DM statistic is normally distributed. The test can be performed as either one sided or two sided. In this paper, we will utilize the two tailed variant, as we are mainly interested whether the two forecasts differ. The sign of the test statistics tells us which of the two forecasts performed better. The test statistic is negative in the first model's supremacy, while positive in the second model's superiority.

Chapter 4

Empirical Research

Following the review of existing literature and the presentation of the underlying theory in the previous chapters, we will now shift our focus to the primary objective of this paper, the empirical research. Our initial examination involves analyzing the volatility risk premium across 18 distinct assets traded in the US market. We explore various methodologies for calculating the volatility risk premium, assess the impact of trading volume on this premium, and establish diverse approaches for deriving past averages of the volatility risk premium.

Further, we explore, whether leveraging knowledge of historical volatility risk premiums enhances our ability to make more accurate forecasts of future volatility, compared to commonly used models. It was shown that calculating averages of past volatility premiums and using them to adjust market volatility forecasts yields better estimates of future realized volatility than simple IV on energy markets (Prokopczuk & Simen 2014). In this paper we plan to answer the question, whether these findings hold true for a broader spectrum of assets. In addition we propose a new methodology of adjusting the current implied volatility for past volatility risk premium. Instead of computing the average of the past premiums, we obtain the exponentially weighted average. This approach places more emphasis on the recent volatility risk premiums compared to the later ones.

Implied volatility, as the current market's forecast of future volatility, serves as a natural point of comparison for RIV models, despite its acknowledged biases (Bennett 2014). Given the general acceptance of GARCH family models as robust tools for volatility forecasting, we also construct a GARCH model to forecast future realized volatility (Andersen & Bollerslev 1998; Brooks 2014). To asses the predictive power of all models, we utilize the loss functions specified above and the Diebold-Mariano tests.

The remainder of this chapter is structured as follows. The first section introduces our data set followed with the analysis of the volatility risk premium. We present both RVP and RVRP descriptive statistics and their distributions. Further, we briefly investigate whether volume of contracts shows to have any effect on VRP. Finally, we present the reader with the outcomes of the RV forecasting along with the associated forecasting errors. The first forecasting period is around three years and includes the GARCH model. Since this short timed period was marked by the covid-19 market crash, we will see how the models perform during turbulent market conditions, as this period witnessed a substantial surge in realized volatility across all assets. The second forecasting period, analyzing only IV and RIV models, spans around twelve years. In this time frame we will investigate how the models perform during mostly normal market conditions.

4.1 Data Description

In our study, we examine 18 assets distributed across 6 distinct asset classes: grains, energy, meats, metals, bonds, and equities. The data set comprises of time series data encompassing daily prices, implied volatilities, and realized volatilities. The observation period extends from 01/Jan/2010 to 01/Nov/2022, resulting in 3158 observations for each asset. Throughout the paper, volatility is presented on a yearly basis unless explicitly stated otherwise. All volatilities and returns are expressed as percentages. Daily realized volatilities are computed using intraday 5-minute data, following the methodology outlined by Liu et al. (2012). Daily implied volatilities are derived from OTM call and put option prices employing a model-free approach as detailed in Bakshi et al. (1997; 2003). For the utilization in the GARCH model, the daily prices are subsequently transformed into continuously compounded logarithmic returns using Equation 3.9.

Given the significance of implied volatility in our paper, Table 4.1 provides the reader with descriptive statistics of the implied volatility time series. Notably, the implied volatility of bonds stands out as the lowest among all assets, reflecting their status as relatively low-volatility instruments in the market. Additionally, the mean implied volatility for S&P 500, gold, and several meats falls within the mid-range. Conversely, the remaining assets exhibit relatively high mean implied volatilities, with natural gas reaching a maximum of 49.9.

Table 4.1: Implied Volatility - Descriptive statistic

	Statistic						
Asset	Mean	Median	St. dev.	Min	Max	Skew	Ex. kurt.
Grains							
Corn	27.738	26.42	9.30	9.64	62.06	0.58	2.93
Oats	32.301	30.59	7.33	17.9	84.67	1.41	6.09
Rough Rice	20.389	19.23	7.01	7.07	63.78	1.21	5.87
Soybean	22.599	21.19	7.76	9.59	61.42	1.33	5.03
Soybean oil	23.720	21.92	6.87	12.02	64.37	1.22	4.81
Wheat	31.111	29.03	9.53	15.2	177.2	2.04	19.6
Energy							
Crude Oil	36.180	32.13	18.58	12.81	485.8	6.75	109
Natural Gas	49.722	45.41	18.83	18.57	130.6	0.93	3.58
Meats							
Live Cattle	16.510	16.03	5.27	6.68	73.42	2.63	20.1
Feeder Cattle	14.100	13.25	4.68	3.81	60.97	2.07	14.1
Lean Hogs	23.413	21.84	8.18	6.05	108.99	2.87	20.6
Metals							
Gold	17.063	15.77	6.66	7.38	63.26	2.01	9.74
Copper	27.209	24.94	11.45	10.70	97.88	1.74	6.80
Bonds							
2Y UST	1.2554	1.09	0.79	0.20	5.58	1.49	5.49
5Y UST	3.5428	3.10	1.53	0.95	10.22	1.22	4.36
10Y UST	5.6366	4.96	2.06	2.53	14.37	1.32	4.65
30Y UST	10.199	9.56	3.34	4.98	33.87	1.43	6.22
Equities							
E-MINI SP500	17.411	14.96	8.69	5.69	103.82	2.65	14.8

Source: Author's computations

The descriptive statistics for log-returns are outlined in Table A.1. Figure A.1 provides visual representations of price movements for the previously mentioned assets, with a note that the y-axis may vary across charts. Additionally, Table A.2 details the specifications of the futures contracts utilized in this paper. Bonds and SP500 futures are also mentioned in the table, although the bonds are quoted as the yield on the underlying treasury security, typically expressed as a percentage, and futures price on S&P 500 are expressed in the nominal value of the index.

Although there was no problem with missing data in the case of implied volatilities and prices, there were several occurrences of missing data in the realized volatility time series. In the paper, we are working with realized volatilizes throughout the next 30 calendar days, in order to match it with implied volatilities from option prices. The presence of missing data points in the realized volatility time series did not significantly impact our computations. In instances where data was missing, we mitigated the impact by simply utilizing the average for the corresponding 30 calendar day period derived from the available data. This was done by employing the Equation 3.12. This approach ensured that the calculation of 30 day realized volatility remained robust despite occasional missing data points.

When dealing with financial time series data, a longer time period is generally preferable. Despite our data spanning only 12 years, we maintain that this duration is adequate for deriving meaningful conclusions. All calculations were performed using the R programming language, and we are prepared to provide readers with access to the code upon request.

4.2 VRP Analysis

In this section, we present the reader with an analysis of the volatility risk premium associated with the studied assets. The VRP values presented in Table 4.2 were derived through the application of Formula 3.16. From theory, positive volatility risk premium is expected. However, in several instances within our sample, this expectation was not met. Oats, rough rice, and feeder cattle futures exhibited a negative median of volatility risk premium. Additionally, bonds across all maturities displayed a negative volatility premium as well.

Table 4.2: Volatility Risk Premium - Descriptive statistic

	Statistic								
Asset	Mean	St.d.	Skew	Kurt	Min	1st Q	Median	3rd Q	Max
Grains									
Corn	1.50	4.76	0.44	1.80	-18.3	-1.34	1.06	4.02	27.0
Oats	-7.01	8.11	0.02	8.55	-49.3	-11.27	-6.68	-2.36	53.8
Rough Rice	-4.80	5.54	-0.25	4.09	-30.4	-8.15	-4.24	-1.25	17.4
Soybean	1.16	3.37	0.05	4.74	-12.7	-0.77	1.10	3.13	16.5
Soybean oil	-0.03	3.73	0.12	7.80	-18.4	-2.00	0.04	1.90	30.5
Wheat	0.32	5.29	1.93	54.20	-36.3	-2.26	0.28	2.85	102.4
Energy									
Crude Oil	1.37	13.96	-1.41	84.81	-147.7	-0.96	2.69	5.96	264.6
Natural Gas	2.43	8.75	-0.15	6.75	-38.2	-1.25	2.30	6.52	56.1
Meats									
Live Cattle	2.15	3.23	-0.27	10.69	-19.0	0.33	2.12	4.01	26.7
Feeder Cattle	-1.72	3.72	-1.44	15.66	-33.4	-3.42	-1.52	0.49	29.1
Lean Hogs	0.89	6.21	-0.20	14.07	-42.1	-1.72	1.17	3.79	50.2
Metals									
Gold	0.49	3.06	-2.21	15.74	-23.3	-0.55	0.65	2.07	14.4
Copper	0.75	4.54	-1.18	8.20	-27.4	-1.44	1.11	3.20	22.1
Bonds									
2Y UST	-0.29	0.26	0.24	5.62	-1.7	-0.45	-0.30	-0.15	0.8
5Y UST	-0.09	0.61	-1.35	10.05	-4.4	-0.34	-0.03	0.26	2.1
10Y UST	-0.16	0.96	-1.82	14.73	-7.7	-0.55	-0.08	0.35	3.9
30Y UST	-0.20	2.11	-3.73	36.71	-21.4	-0.87	-0.03	0.86	13.0
Equities									
E-MINI SP500	5.17	5.06	-1.22	21.83	-39.4	3.10	5.11	7.44	54.8

Source: Author's computations

From the grains asset class, corn and soybeans exhibit positive volatility risk premiums with median values of 1.06 and 1.10, respectively. Employing the median helps us mitigate the impact of extreme values in the tails of RVP distribution. Notably, corn displays a higher standard deviation in VRP, as evidenced in Figure A.2, aligning with its greater volatility based on log-returns from Table A.1. The fact that market makers price a markup for volatility risk is in line with the theory. Soybean oil and wheat have volatility risk premiums near zero, though the latter being slightly positive. On these markets the participants are, on average, able to correctly value the size of future price movements. If we look at the VRP of wheat we can spot a period of increased variance of VRP from 2011 to 2013 that was connected to large price swings of the underlying commodity. After 2014 the VRP was mostly oscillating around zero or negative. Oats and rough rice futures have both considerably negative median of VRP. This negative premium trend, apart from several small periods, persists throughout the sample length. In the Figure 4.1 we present the implied and realized volatility time series of oats futures. On average, there has been a negative VRP of −7.01 volatility points.

Figure 4.1: IV, RV & VRP of Oats Futures

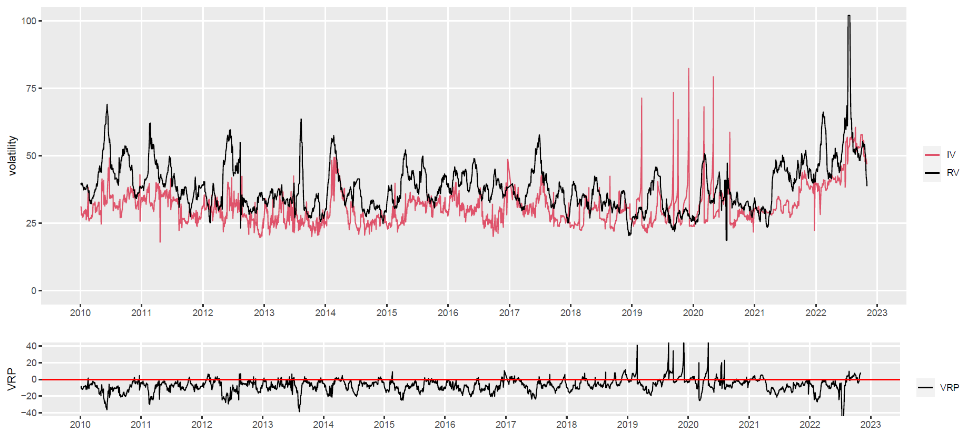

Source: Author's computations

Within the energy category, crude oil and natural gas exhibit positive volatility risk premiums of 1.37 and 2.43, respectively. They also have the highest standard deviation of VRP across the entire dataset, as depicted in Figure A.2. Both VRP charts for these assets show notable fluctuations throughout the observed period. An intriguing observation in natural gas VRP is the occurrence of sharp drops during the first couple of years, coinciding with periods of dramatic price moves, as evident in the price time series in Figure A.1. This volatility in prices during certain periods introduces uncertainty, potentially contributing to the disparities between implied and realized volatility. These unexpected moves likely account for the average positive VRP on these assets during calmer trading days.

An interesting comparison arises when examining the VRP of natural gas and oats, which share similar standard deviations. Apart from the turbulent times around the year 2020, VRP variance of oats appears constant over time. On the other hand, the VRP of natural gas exhibits VRP clustering,a pattern also observed in crude oil.[12] The identification of this clustering phenomenon in some time series and its absence in others presents an intriguing finding, warranting further in-depth exploration for potential insights into underlying market dynamics.

In the meats category, we find that both live cattle and lean hogs have positive median volatility risk premiums. Both time series show increased volatility of VRP during the second half of the dataset. Interestingly, the VRP variance of feeder cattle appears to remain constant over time. A notable and unexpected discovery is the opposite average VRP for live cattle and feeder cattle, with the former being positive and the latter negative. This contrast is surprising given the similarity of these assets. One potential explanation for this discrepancy could be related to futures volumes, and a detailed analysis of this issue is presented later in the paper.

Figure 4.2: IV, RV & VRP of Live Cattle Futures

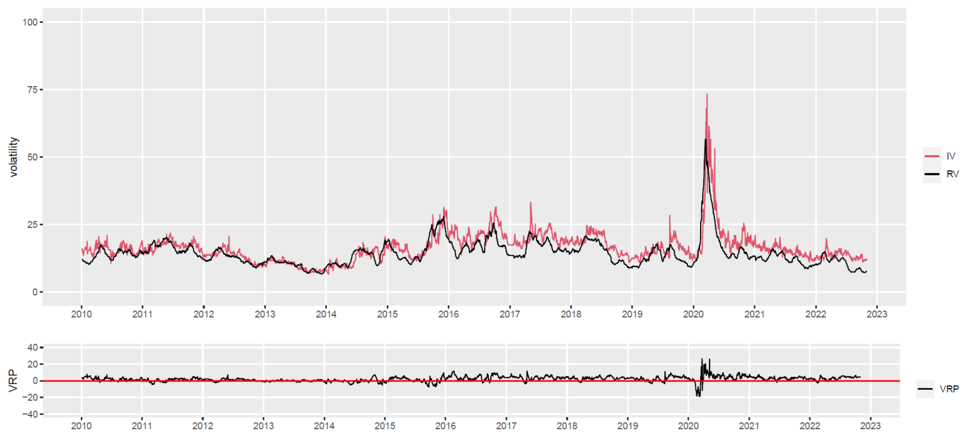

Source: Author's computations

[12]We refer to the similarities with volatility clustering of returns.

Figure 4.2 illustrates the implied and realized volatility of live cattle futures. On average, the implied volatility is 2.15 volatility points higher than the realized volatility. The consistent elevation of implied volatility is visually apparent in the graph. It is important to note that the realized volatility time series appears smoother due to our calculation method. We compute the 30-day realized volatility by averaging daily volatilities over 30 days. In this way we are able to align it with implied volatility from futures option prices. Consequently, this approach impacts a trajectory of the realized volatility time series, as a sustained increase in daily volatility over several days would be necessary to significantly impact the 30-day volatility average.

Both gold and copper exhibit slightly positive volatility risk premiums, with low VRP variances. Excluding bonds, gold has the least volatile VRP among all assets. The VRP chart of gold appears notably smooth compared to other assets, with some signs of VRP clustering. The prices of copper experienced a decline in the first half of the observed time period, reaching its lows around 2016. Subsequently, the price moved higher with several intermittent pullbacks. Yet, no significant change in VRP was observed during this period. With the exception of the years 2010 and 2011, the variance of VRP for copper seems to remain constant throughout the entire sample.

Bond prices have an inverse relationship with the size of the interest rate, which is left to the discretion of the Federal Reserve System. While various variables impact bond prices, current and expected future interest rates are paramount. For bond traders, accurately predicting future bond movements thus involves forecasting interest rate decisions. Notably, the volatility risk premium of bonds varies by maturity, with shorter maturities experiencing smaller absolute value VRPs, while longer maturities exhibit larger ones. Examining the 2 year treasury futures, a period of zero interest rates from 2012 to 2016 corresponds to consistently negative VRP. Although this trend continued post-2016, there were intermittent instances of VRP breaking upward of zero. In contrast, medium to long-term bonds displayed a distinct chart pattern, with premiums oscillating around mean, close to zero. Despite the high liquidity of the US bond market there is no apparent indication that market participants consistently overpay for protection through options.

The sole representative of equities in this study is the SP 500 futures. The mean volatility risk premium for the SP 500 is 5.17, a relatively high value compared to other assets in our dataset. However, this aligns well with findings from similar studies examining the VRP of equities (Feri 2017). Figure

4.3 illustrates that the VRP of the SP 500 remains consistently above zero throughout our time frame. There are only a few instances of dips below zero, most notably during the 2020 stock market crash triggered by the covid-19 pandemic.

Figure 4.3: IV, RV & VRP of SP 500 Futures

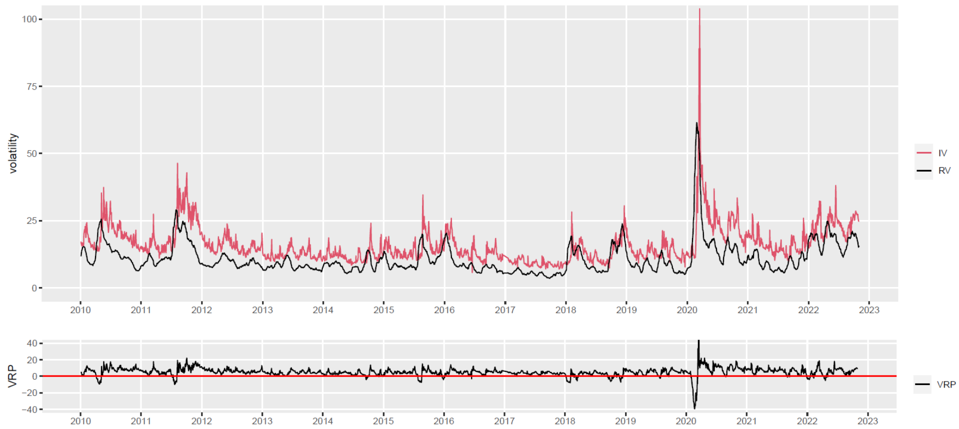

Source: Author's computations

4.2.1 Relative VRP

In the preceding section we observed that there, on average, exist differences between the IV and RV. However, it is difficult to compare the size of the premium among assets with different volatility levels. As indicated in Table 4.1, the average IV expressed on an annual basis ranged from 1.25 to 49.7 for different assets. The bounder assets were 2-year notes and natural gas, respectively. To address this issue and enable a more meaningful comparison, we employ the relative volatility risk premium (RVRP) using Equation 3.15. By taking the ratio of implied and realized volatilities, rather than the difference, we aim to standardize the measure and facilitate a more comparable assessment of volatility pricing across different markets. From the definition, we expect both VRP and RVRP to be either positive or negative.

Table 4.3 presents descriptive statistics for the relative volatility risk premium. Excluding equity futures, the average RVRP ranged from 0.82 to 1.17.[13] For mid to long-term bonds, the RVRP hovers near zero with a slightly negative premium, while the RVRP for short-term 2-year notes is 0.74. Typically,

[13]Rough rice and live cattle respectively.

Table 4.3: RVRP - Descriptive statistic

	Statistic								
Asset	Mean	St.d.	Skew	Kurt	Min	1st Q	Median	3rd Q	Max
Grains									
Corn	1.06	0.19	0.38	3.71	0.50	0.94	1.05	1.17	1.88
Oats	0.84	0.20	2.50	20.52	0.39	0.71	0.82	0.93	3.18
Rough Rice	0.82	0.21	0.73	4.49	0.28	0.68	0.80	0.94	1.82
Soybean	1.07	0.17	0.22	3.50	0.47	0.96	1.06	1.17	1.68
Soybean Oil	1.01	0.16	0.88	6.76	0.54	0.91	1.00	1.09	2.34
Wheat	1.02	0.15	0.47	5.45	0.47	0.92	1.01	1.11	2.37
Energy									
Crude Oil	1.09	0.20	-0.16	4.07	0.32	0.97	1.10	1.22	2.20
Natural Gas	1.07	0.17	0.01	3.70	0.51	0.97	1.07	1.17	1.84
Meats									
Live Cattle	1.17	0.21	0.28	3.10	0.44	1.02	1.15	1.31	1.98
Feeder Cattle	0.91	0.18	0.18	3.09	0.31	0.78	0.89	1.03	1.91
Lean Hogs	1.06	0.21	0.17	3.53	0.42	0.92	1.06	1.19	1.94
Metals									
Gold	1.05	0.17	-0.38	4.66	0.35	0.96	1.05	1.15	1.92
Copper	1.05	0.18	-0.11	3.37	0.44	0.93	1.06	1.17	1.65
Bonds									
2Y UST	0.74	0.20	0.29	3.09	0.21	0.60	0.74	0.87	1.56
5Y UST	0.99	0.16	-0.05	3.79	0.39	0.88	0.99	1.09	1.76
10Y UST	0.98	0.15	-0.08	4.08	0.37	0.89	0.98	1.07	1.59
30Y UST	0.99	0.15	-0.30	4.56	0.29	0.91	1.00	1.09	1.62
Equities									
S&P 500	1.59	0.41	-0.08	3.13	0.23	1.33	1.59	1.85	3.24

Source: Author's computations

a positive volatility premium is interpreted as a payment for downside protection. If we flip our narrative, one could view a negative volatility premium as a payment for upside protection. The negative RVRP for 2-year notes may be seen as compensation for the possibility of interest rates decreasing in the foreseeable future. However, during most of our studied time period, the federal funds rate was already near zero. On the other hand, in a low-interest environment, investors might leverage their bond holdings, selling options to generate additional income while potentially underestimating the possibilities of market turbulence.

We have seen in Table 4.2 the mean VRP of oats and rough rice to be -7.01 and -4.80 respectively. The corresponding standard deviations are 8.11 and 5.54. Without looking at the underlying volatility levels, we might get the impression that the volatility premiums differ. However, in terms of RVRP, both time series show very similar both means and standard deviations. Another pairs exhibiting similar tendencies are gold and copper or crude oil and natural gas. Excluding oats and rough rice, the rest of the grains market shows a similar mean relative volatility risk premium, ranging from 1.01 to 1.07. This contrasts with the findings in Table 4.2, where all four of these assets exhibited different volatility risk premiums. Thus, the use of RVRP, which normalizes for underlying volatility levels, provides a more consistent measure for comparing volatility premiums across different assets.

Figure 4.4: VRP & RVRP of Oats Futures

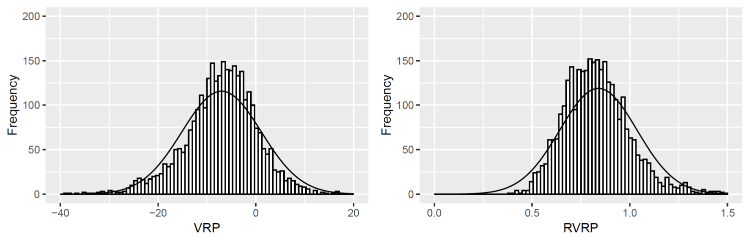

Source: Author's computations

Figure 4.4 displays the frequency distribution of both standard and relative volatility risk premiums of oats futures. For better comparation, normal distributions with the same first and second moment as the underlying data have been fitted on the histograms. Oats futures are notable for demonstrating a fairly significant negative volatility risk premium, with average values of -7.01 for VRP and 0.84 for RVRP. In both cases, the distributions exhibit considerable skewness or kurtosis relative to a normal distribution. Similar statistical properties are often observed in financial market time series, such as logarithmic returns. However, in this particular case, the use of ratios instead of differences in the calculation of volatility premium did not lead to a significant alteration in the underlying distribution of volatility risk premium. The fat-tailed and skewed characteristics persist in both the VRP and RVRP distributions for oats futures.

Figure 4.5: VRP & RVRP of Live Cattle Futures

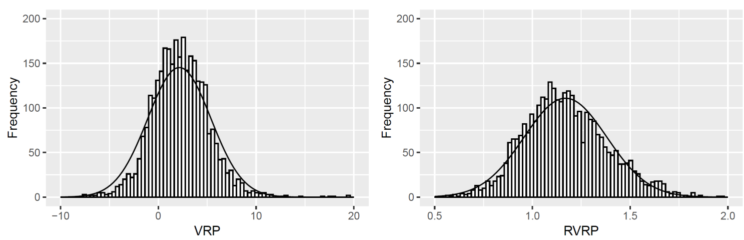

Source: Author's computations

Figure 4.5 illustrates the VRP and RVRP of live cattle futures. An intriguing observation is that the distribution of RVRP more closely resembles the density function of a Gaussian distribution. It is important to note that this feature was not observed in the previous case of oats futures, so making a general statement that the ratio of volatility risk premium is normally distributed might not be warranted. However, the unexpected negative volatility premium observed in oats futures might contribute to the non-normal distribution of RVRP. Examining all markets in Table 4.3 with positive mean RVRP, the third and fourth moments for most of them are very close to 0 and 3, respectively. This stands in a strong contrast to the findings in Table 4.2, where these statistical values were mostly ambiguous. The difference in the distribution characteristics between VRP and RVRP underscores the impact of the chosen metric on the observed statistical properties.

Figure 4.6: VRP & RVRP of SP 500 Futures

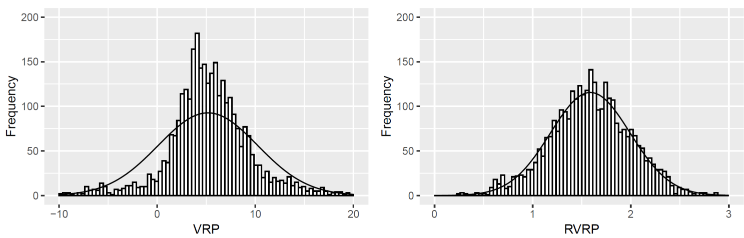

Source: Author's computations

Figure 4.6 displays the VRP and RVRP histograms for the SP 500. At a first glance, the difference between both distributions is apparent. In this case, the use of relative volatility risk premium RVRP results in a density distribution very close to Gaussian. The Shapiro–Wilk test with the null hypothesis of a sample of the population being normally distributed results in a p-value of 0.007. Consequently, it is challenging to assert whether the RVRP is actually normally distributed. Despite this, the evident change in distribution between VRP and RVRP is in this case quite obvious.

If we follow the analysis above, the SP 500 exhibited the largest average volatility risk premium of 5.17. Similar values have been previously reported by Feri (2017). Furthermore, it shows the largest RVRP of 1.59, which is significantly larger than the rest of the analyzed markets and aligns with the findings of Lu (2020). This might be influenced by the fact that the SP 500 typically trades at very low volatility levels. Thus, even a small volatility premium on an absolute basis would then translate into a relatively large premium. Another explanation could be the usage of the options in structured products.

The discovery that the relative volatility risk premium in markets with positive volatility risk premium might exhibit a normal distribution can have significant implications in the field of financial markets. Understanding the underlying distribution could prove valuable in risk management, where the true distribution of many studied time series is generally unknown. The observation that the market tends to overshoot and undershoot expected volatility with similar frequencies suggests two possibilities. First, it could indicate that market participants, on average, are adept at predicting future volatility, and the normal distribution stems from the averaging process of all trades. Alternatively, it might suggest that market participants are unable to predict future volatility at all, and this inherent unpredictability averages out to a normal distribution.

In conclusion, we have demonstrated that using the ratio instead of the difference between implied and historical volatility produces comparable estimates of volatility risk premium. Out of 18 markets analyzed, 14 exhibited the properties of a non-negative volatility risk premium.[14] Further, we plan to explore the connection between volatility volatility premium and the average volume of traded contracts for the underlying assets.

4.2.2 Volume Effect

The price determination process in the market works efficiently only when enough liquidity is provided. Markets with low volumes can exhibit several undesirable features such as reduced liquidity, wider bid-ask spreads or increased trading costs. Low-volume markets are also more susceptible to price manipulation and exaggerated price movements. The illiquidity can further lead to delays in executing trades and may result in investors being unable to exit positions at desired levels (Amihud 2002).

The average volumes of futures contracts analyzed in this paper are presented in the Table A.2. The daily volume is the average of the continuous futures volume of the first 2024 expiration. The time span of these contracts is from the beginning of 2022 to the end of 2023. This time frame is shorter and does not correspond to the time frame analyzed in the paper. However, we are only interested in a rough relative estimate of volume levels between different markets. We thus approximate the volume of options by the volume of futures, which should, on ordinal basis, not differ substantially. For these reasons, we will further use the volumes of futures and volumes of options interchangeably.

[14] From the data, we can easily imagine the RVRP of mid to long term bonds to be at or above zero.

It is also worth mentioning that volatility risk premium is obtained from the option prices and the volumes are obtained from futures prices. In this case, the futures volumes serve only as a proxy with the assumption that the higher the futures volume is, the higher the demand for hedging these positions through options should be.

Figure 4.7: Effects of Volume on RVRP

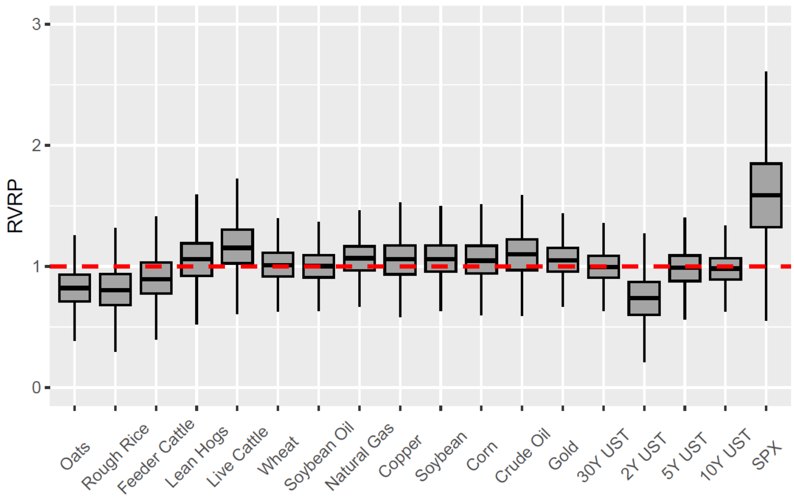

Source: Author's computations

The Figure 4.7 presents a box plot showing the connection between volume and relative volatility risk premium. The y axis depicts RVRP and the x axis shows the assets, which have been sorted by volume from the smallest on the left to the highest on the right. Three least liquid commodities turned out to be the only three commodities with negative RVRP. Feeder cattle and lean hogs on the other hand show slightly increased values compared to the rest of the data set. The middle part of the plot constitutes of assets with relatively similar RVRP between 1.01 and 1.09. There has been no connection between volatility premium and volume on the bond market.

We can also consider the overall notional value of contracts traded. Prices on the futures markets represent different amounts of the underlying commodity. If we were to compare the real commodity price of two assets, we would not be able to do so without knowing the contract specifications. Specifically, the contract sizes and price quotations. Note how almost all futures trade in different contract sizes, units or price quotations. With the knowledge of notional value of one contract we calculate the notional volume e.g. the total dollar value which is on average traded per day.[15] Box plot of the notional value of contracts traded is depicted in Figure A.3. However, we assume volume of contracts traded to be superior to the total notional value traded, as low volume itself is the cause for low market depth. It is interesting though, that both methods yield similar findings.

Our analysis suggests that volume has an effect on the size of volatility premium. We expect that many problems caused by illiquid markets would have a direct impact on the price of options, therefore on the size of volatility premium. Out of the subset of commodities with more liquid markets, all of them showed positive relative volatility risk premium - as one might expect from the theory. We could assume that the low trading volume creates an environment where supply and demand pressures on option prices are not strong enough to eliminate the discrepancies between implied and realized volatilities. However, further research that would be needed to scrutinize the problem more closely is beyond the scope of this paper. We can only conclude that volume has a direct effect on the volatility risk premium, or more precisely, that illiquid markets cause the volatility premium to behave unusually.

4.2.3 Average Relative VRP

The Figure 4.8 shows the heat map of RVRP one day correlations. As shown above, RVPR oscillates around its mean similarly to returns. It is interesting to see that nearly all correlations in the figure are positive. That means that the volatility risk premium of many analyzed assets move in the same direction on a day-to-day basis. Even though several assets exhibited close zero correlations, almost all of the correlations are positive and range from 0.3 to 0.8 on multiple occasions. These short term correlations might be caused by the market sentiment on each day which makes the implied volatility increase or decrease together among markets simultaneously.

Typically, realized volatility rises more rapidly during a decline in asset prices, leading to a negative volatility risk premium. Conversely, a positive VRP is more anticipated when prices remain stable. The nearly universal non-negative correlation in premiums across assets may imply either a collective movement of all asset prices in the same direction or a market misjudgment

[15]The asset price used in the calculation was taken on 11/Jan/2024. The usage of only one specific price could run into the same objection as above. However, similarly as in the previous case, the price action should not have a significant effect on the ordinal order of volumes.

Figure 4.8: Correlations of RVRP

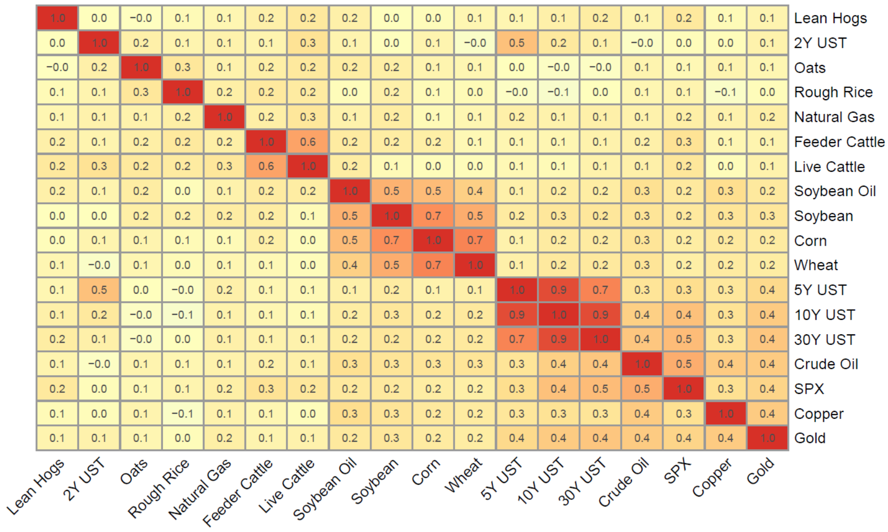

Source: Author's computations

regarding future volatility on all markets in the same direction. Both of these scenarios seem improbable. Alternatively, a third explanation could involve a nontrivial feature causing short-term correlations in volatility premiums. We may notice a parallel with the story of LTCM, where investments appeared uncorrelated over the long term but demonstrated short-term correlations.[16]

Several correlation clusters can be spotted near the diagonal. It is interesting to see that these clusters are usually made out of assets with similar characteristics such as feeder cattle and live cattle or medium and long term bonds. However, if we were to use the RVRP as a predictor for future realized volatility it would need to be independent on any other factors effecting all markets collectively. In this case, we would expect to see a similar number of negative and positive correlations. Since this is not true, the use of daily historical RVRP might not result in the best possible forecasts.[17] The observed imbalance between negative and positive correlations suggests that there might be a need for adjustment of RVRP in order for it to yield the most accurate forecasts.

In this paper we work with two methods of obtaining the medium term

[16] Long-Term Capital Management was a highly profitable hedge fund which employed a complex investment strategy centered on mathematical models and fixed-income arbitrage. However, their downfall occurred when unforeseen correlations among their leveraged assets during the 1998 financial crisis resulted in massive losses, prompting a major financial bailout to avert a broader market collapse (Lowenstein 2001).

[17] This could further imply that the usage of simple HV models to predict future IV might not be the most plausible one.

average. Prokopczuk & Simen (2014) already showed that using the average volatility risk premium A-RVRP improves the predictions of realized volatility. However, due to volatility clustering, the volatility observed in recent period should have a larger impact on future realized volatility than volatility further in the past. For these reasons we are interested in finding a procedure of enhancing this averaging process in order to produce a more accurate forecasting method of realized volatility. In our paper we propose the utilization of exponentially weighted moving average instead of simple average.

We thus first calculate the weighted averages of the past observations of RVRP to obtain average relative volatility risk premium A-RVRP as defined in Equation 3.18. The parameters used in this case were $= 231$ and $\tau = 21$. Further, we calculate the weighted exponential average relative volatility premium E-RVRP over the past 231 trading days for each day in our time frame utilizing the Equation 3.20. Also adjusting for the 21 days for which we do not have RV.

Figure 4.9: RVRP, A-RVRP & E-RVRP of Corn Futures

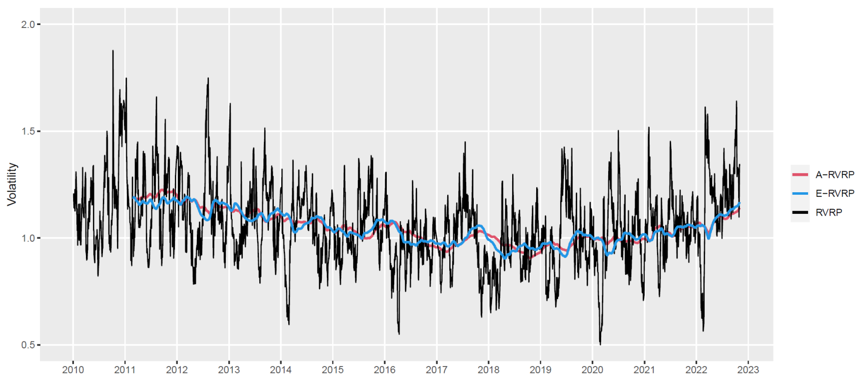

Source: Author's computations

For illustration purposes, we also present Figure 4.9 showing the connection between RVRP, A-RVRP and E-RVRP of corn futures. We see how both averaging methods slowly adjust to the changes in RVRP and how exponential averaging methodology (E-RVRP) is able to react faster to market turbulence. This could prove paramount in forecasting the future realized volatility. Nevertheless, the faster adjustment does not always correspond to the future levels of RVRP. For example, in the middle of 2012 the volatility premium of corn sharply dropped. The E-RVRP adjusted faster to reflect this drop in premium. However, following this decline the volatility premium increased substantially and thus the simple average (A-RVRP) was closer to the following surge in volatility premium. On the contrary, in 2018 the volatility premium declined for a substantial time period. In this case the E-RVRP was able to react faster to this decline than A-RVRP. Also note that the usage or rolling average shortened our data frame by 252 trading days.[18] This can be see in the first year of Figure 4.9 for which the A-RVRP and E-RVRP are missing.

Figure 4.10: Correlations of A-RVRP

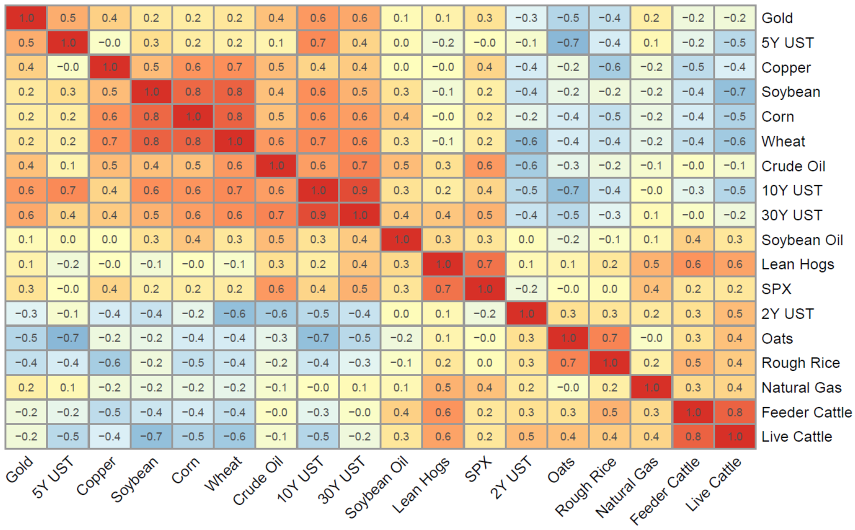

Source: Author's computations

In Figure 4.10 we have the correlation heat map of A-RVRP which has been taken for parameters $n = 231$ and $\tau = 21$. As opposed to the previous correlation matrix, we now have a similar amount of positive and negative correlations among assets. Intuitively, many of the assets exhibiting pairwise correlations in daily RVRP also show strong correlation in the A-RVRP framework. It is interesting to see that SP 500 is among the few assets which show positive correlation to almost all of the other assets. An interpretation could be that SP 500 is so essential in the market that its volatility premium can feed through the market on other assets, and these rises and falls of volatility premium of SP 500 slightly influence the rest of the market. However, since also soybean oil and lean hogs futures achieved similar correlations to the rest of the studied assets, this hypothesis would grant further investigation.

It would be challenging to make any meaningful conclusions based solely on the A-RVRP correlation matrix. Although some assets seem to have a reasonable correlation between each other, it would be hard to asses why, for example,

[18]The sum of 231 days used for the weighted moving average calculation and 21 past days, for which we do not already know the 30 day realized volatility.

the A-RVRP of oats moves in the opposite direction of 10 year treasury note. We could arguably come up with explanations about perceived market risk or future short term expectations. Since each futures market is made up of immense number of moving parts which determine the market prices, all of our efforts could prove unfruitful. Nevertheless, the correlation matrix shows that the computation method of A-RVRP is able to eliminate short term non negative correlations among volatility premium as seen in the Figure 4.8. Since the A-RVRP and E-RVRP time series proved to be very similar we only present the correlation matrix of A-RVRP, expecting the outcomes of E-RVRP to yield almost identical conclusions.

With the usage of A-RVRP and E-RVRP we are able to compute the corresponding RIV models defined in Equations 3.19 and 3.21. The next section will examine whether the adjustment of IV with A-RVRP or E-RVRP can yield significantly better forecasts of future volatility as compared to the IV or GARCH models.

4.3 Short Term Forecasting

We now shift our focus to the examination of the predictive capabilities of models employed for forecasting realized volatility over a 30 calendar day period. In this paper we compare four models. The GARCH, IV, A-RIV and E-RIV. In this section, we will outline the modeling processes for each methodology. Given that implied volatility is readily available from the market, there is no need to model it or describe the modeling method.

To ensure a direct comparison between the RIV models and the GARCH model, we must specify the same testing time frame. In this paper, we utilize a testing time frame spanning from 06/Jan/2020 to 01/Nov/2022, encompassing 710 data points from our original dataset length of 3138. As such, 77% of the dataset was allocated for training the GARCH model, while the remaining 23% was reserved for testing purposes. We implemented a one-day-ahead rolling GARCH approach, which updates the dataset after each prediction.

4.3.1 GARCH

Typically, when employing the GARCH model for forecasting, we divide our dataset into two parts. In the first part, we train the model to estimate specific parameters of the GARCH model from Equation 3.24. We are particularly interested in determining the α and β parameters, although all of the variables are required for the forecasting. Once the model parameters are obtained, we can use them to generate one day ahead forecasts for the future.

Utilizing model parameters solely from the training period for forecasting the entire testing period would overlook valuable information from the future that could aid in model updating. Therefore, on the first day of forecasting with the GARCH, we estimate the model parameters and utilize them to forecast realized volatility over the next 30 calendar days. On the subsequent day, as we obtain one additional realized volatility data point, we estimate new GARCH model parameters. While the parameters should not vary significantly on a day to day basis, sharp market moves may lead to substantial differences in model parameters, and thus in forecasts.

We repeat this process for all data points in our testing time frame, updating the GARCH model each day, to reflect the most recent information available. In total, 710 GARCH model parameter sets were estimated for each asset to generate the forecasts. We employed a GARCH(1,1) volatility model with ARMA(0,0) mean model. Considering the number of assets, we opted for the simplest and most robust GARCH family model. GARCH(1,1) can be regarded as a reasonable choice for our purposes (Hansen & Lunde 2005).

4.3.2 RIV Models

In the case of A-RIV and E-RIV forecasts, there is no necessity to obtain any specific parameters as the model is already defined as a mean of past observations. This stands as one of the key advantages of this approach because all we require are past data points of volatility risk premium and the current market implied volatility. The calculation of these forecasts was outlined in Equations 3.19 and 3.21, where we only substitute index t with $t+1$. Consequently, forecasting can be conducted solely based on past observations without the need to learn any model specific parameters, as is required in the case of GARCH.

The values of A-RVRP and E-RVRP are of paramount importance, as they are utilized to adjust implied volatility to generate more accurate forecasts. In Figure 4.9, we observed that E-RVRP responds more quickly to recent market turmoil. However, readers might be interested in comparing the values of A-RVRP and E-RVRP more closely. Since both values are employed to adjust current implied volatility in the same fashion, if the absolute difference between them is negligible, both methods would result in similar IV adjustments. Figure

4.11 presents these values for our forecasting period of 710 days for soybean futures.

Figure 4.11: A-RVRP & E-RVRP of Soybean Futures

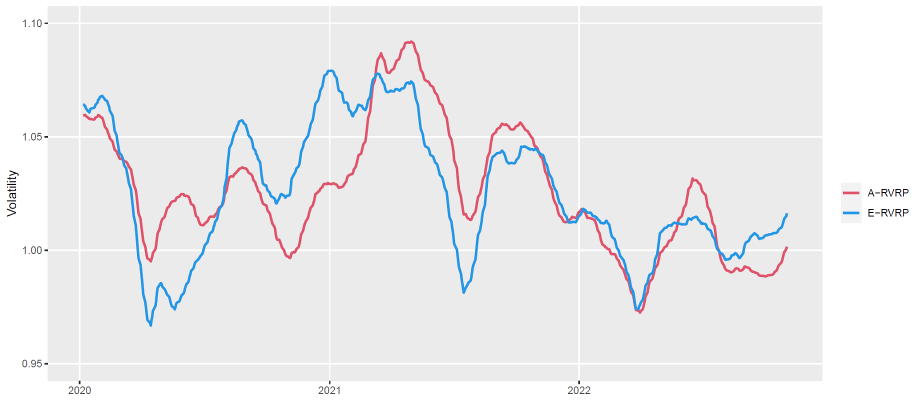

Source: Author's computations

Upon initial examination, it is evident that during the first half of our forecasting period, there were differences between the values of A-RVRP and E-RVRP. The rapid response of E-RVRP to the volatility premium drop during the March 2020 crash is clearly observable. However, it is crucial to consider the absolute difference between both averaging methods. Even during dates with the largest discrepancies between A-RVRP and E-RVRP, the difference was only around 5 percentage points. This implies that the adjusted implied volatility would also vary by only around 5 percentage points.[19] Furthermore, during the second half of our testing period, the difference was lower than 1 percentage point. Hence, the disparity between both time series is in this case negligible. These findings may vary for each studied market. Nonetheless, the magnitude of the maximum difference is never so substantial that it would result in entirely different IV predictions.

Based on these findings, it is reasonable to expect that the predictive power of both A-RIV and E-RIV should be very similar. During the analysis of the forecasting errors, even a slight outperformance of one model might translate into a significant improvement in the overall forecasting power, given that the adjustment was also considerably small. Therefore, even minor differences in forecasting accuracy between the two models could have notable implications for the effectiveness of the forecasting approach as a whole.

[19]For illustration, the difference between IV of 28.6% and 30%.

4.3.3 Forecasting Errors

In the Table 4.4 we present the mean average percentage errors (MAPE) of the forecasts of realized volatility over a horizon of 30 days. If a particular methodology yields the smallest MAPE, we denote this as a *hit*. All hits in the table have been highlighted. Additionally, Table A.3 in the appendix presents the remaining loss function errors. Due to the structure of our data, MAPE could potentially be the most practical loss function to analyze. The utilization of absolute errors mitigates the impact of outliers on forecasting evaluation, while the use of percentages helps us in comparing among errors between assets. Nonetheless, analyzing all loss functions provides additional validation of our findings. Remarkably, all loss functions yield highly similar results, leading to nearly identical conclusions.

Table 4.4: MAPE - Forecast #1

Asset	GARCH	IV	A-RIV	E-RIV
Corn	0.2642	0.1600	**0.1381**	0.1399
Oats	0.2052	0.1809	0.1730	**0.1672**
Rough Rice	0.4274	0.2996	0.2788	**0.2782**
Soybeans	0.2175	**0.1213**	0.1255	0.1265
Soybean Oil	0.1518	**0.1408**	0.1445	0.1470
Wheat	0.1497	**0.1221**	0.1303	0.1295
Crude Oil	0.3243	0.1986	**0.1807**	0.1836
Natural Gas	0.2518	0.1780	**0.1620**	0.1694
Live Cattle	0.4870	0.3156	0.1482	**0.1477**
Feeder Cattle	0.2780	**0.1617**	0.1784	0.1810
Lean Hogs	0.9260	0.2201	**0.1790**	0.1824
Gold	0.1940	**0.1395**	0.1435	0.1426
Copper	0.1460	0.1527	**0.1344**	0.1356
2Y UST	0.2140	0.2578	0.2206	**0.2106**
5Y UST	0.2170	**0.1437**	0.1511	0.1550
10Y UST	0.1980	**0.1266**	0.1390	0.1411
30Y UST	0.1910	**0.1244**	0.1436	0.1420
E-MINI SP500	0.5350	0.6393	0.2270	**0.2261**
Σ hits	-	8	5	5

Source: Author's computations

At the first glance, it is evident that the GARCH forecasting method did not achieve a single hit across all assets. This aligns with the findings of Szakmary et al. (2003). The most effective forecasting models were those of implied volatility or RIV models. Implied volatility emerges as the most reliable predictor with 8 hits, while A-RIV and E-RIV tie for second place with 5 hits each. Consequently, in 10 instances, RIV models managed to enhance the forecasting accuracy of implied volatility on realized volatility. This suggests that volatility premium was embedded in option prices, which were consequently not priced accurately.

Upon observation, we can identify groups of assets with similar characteristics that are best predicted by the same model. These groups include crude oil and natural gas, medium to long term bonds or soybean and soybean oil. The reasoning in why each model performs optimally for each asset could stem from differing volatility regimes or the magnitude of volatility premium. However, further analysis would be necessary to comprehensively examine this issue and could serve as a foundation for future research on analyzing the predicting power of IV on RV during turbulent market conditions.

Implied volatility was generally effective in bond markets, with the exception of the 2-year note where E-RIV performed better. In this scenario, the knowledge of past volatility premium did not contribute to enhancing IV for better forecasting power. The same holds for all assets where IV performed better. During turbulent times, IV emerged as a reliable predictor of future volatility, suggesting that volatility was accurately priced. An interesting observation is the effectiveness of implied volatility as the best predictor for gold, which shares similar characteristics with longer term bonds.

Both A-RIV and E-RIV have 5 hits each. This suggests that during periods of market turmoil, it is challenging to determine whether IV or RIV models offer superior forecasting of realized volatility. However, there seems to be merit in incorporating the size of past volatility risk premium in volatility forecasting. The relatively consistent number of hits for RIV models indicates that there is value in using past volatility risk premium in forecasting, as we would expect to see fewer hits if the methodology was entirely ineffective.

The errors of both A-RIV and E-RIV forecasts were generally comparable, aligning with our expectations from previous sections, given the small differences between A-RVRP and E-RVRP. The most substantial improvement over a basic implied volatility approach was observed in the SP 500 market, where the utilization of either adjusting method resulted in significant improvement. Moreover, the disparity between GARCH forecasts and the forecasts from the other IV methods was considerable. Forecasting realized volatility using the GARCH methodology does not yield superior results compared to obtaining simple implied volatility from the market. There were only two instances where GARCH forecasts scored a hit, which are soybean oil and 2-year treasuries forecasts when utilizing the MPSE loss function.

4.4 Long Term Forecasting

Given the less precise nature of GARCH forecasts, we might be interested in conducting a deeper analysis of RIV models. In the previous scenario, we dividend our entire data set of 3138 observations into two segments, reserving the larger portion for training the GARCH model. Since the computation of RIV models does not need any data for training, we can utilize the full size of our original data set. That could yield more efficient estimates of the effectiveness of various IV methodologies. Only the initial 252 data points will be excluded to facilitate the first average calculation. These omitted points can be seen on the left side of the Figure 4.9 showing the time series of corn RVRP.

Our original forecasting window spanned from 2020 to 2022 and included the 2020 covid market crash and its aftermath. This period could offer insights into the performance of RIV models in turbulent market conditions. However, we are also interested in observing how the RIV models perform in longer time frames, including mostly normal market conditions. The extended testing time period ranges from 24/Feb/2011 to 01/Nov/2022 and contains 2887 observations for each asset.

4.4.1 Forecasting Errors

The MAPE of long-term RIV models is presented in Table 4.5. Similarly as before, the results for the remaining three loss functions are included in Table A.4 in the appendix. During this longer term period, IV achieved 3 hits, A-RIV achieved 4 hits, and E-RIV achieved 11 hits. From the table, we see that RIV models collectively outperform model free implied volatility. Adjusting implied volatility with past sizes of volatility risk premium resulted in smaller prediction errors in 15 out of 18 instances. Also, the utilization of the exponential weighting approach E-RIV produced superior predictions, surpassing both IV and A-RIV combined. Similar conclusions could be made from the MAE loss function. In the case of squared loss functions MSE and MSPE, A-RIV and E-RIV achieved a comparable number of hits. This suggests that E-RIV may be more effected by bigger mismatches between IV and RV, as these would exert a greater influence on the results of squared loss functions. Nevertheless, the IV forecasts were inferior independently on the loss function applied.

Although A-RIV and E-RIV yielded error values that were close to each

Table 4.5: MAPE - Forecast #2

Asset	IV	A-RIV	E-RIV
Corn	0.1434	**0.1307**	0.1318
Oats	0.2055	0.1695	**0.1693**
Rough Rice	0.2294	0.1887	**0.1848**
Soybeans	0.1348	0.1255	**0.1236**
Soybean Oil	**0.1150**	0.1152	0.1154
Wheat	0.1115	0.1099	**0.1088**
Crude Oil	0.1768	0.1504	**0.1480**
Natural Gas	0.1393	0.1264	**0.1253**
Live Cattle	0.2172	0.1312	**0.1308**
Feeder Cattle	0.1715	**0.1661**	0.1679
Lean Hogs	0.1794	**0.1636**	0.1640
Gold	0.1320	0.1260	**0.1246**
Copper	0.1456	**0.1295**	0.1310
2Y UST	0.2906	0.1857	**0.1802**
5Y UST	**0.1252**	0.1287	0.1309
10Y UST	**0.1143**	0.1149	0.1158
30Y UST	0.1155	0.1141	**0.1136**
E-MINI SP500	0.6148	0.2075	**0.2054**
Σ hits	3	4	11

Source: Author's computations

other, it is important to notice the shift of hits in favor of E-RIV during the analysis of the longer time frame. This would suggest that during normal market conditions, usage of E-RIV further enhances the forecasting of realized volatility solely utilizing implied volatility and past observations of volatility risk premium. It is interesting to observe that during this time frame, E-RIV achieved hits in both crude oil and natural gas as well. In this scenario the E-RIV outperformed both IV and A-RIV, which suggest that the change of the averaging function can have a meaningful impact on the forecasting accuracy. Upon closer examination of the E-RIV columns in Table 4.5 and Table 4.4, smaller error values emerge for the longer time frame. This holds true for the both IV and A-RIV forecasts as well. A larger data size increased the accuracy of all models and showed that, on average, E-RIV serves as the better predictor of future realized volatility.

When examining the grains market, E-RIV consistently outperformed IV in nearly all instances, with the exception of soybean oil where the difference was marginal. Conversely, in the case of rough rice the difference in errors between IV and E-RIV was impressive 4.5 percentage points. One might think that this percentage corresponds to the volatility risk premium on these markets, but such an assumption would be inaccurate. Whether IV or E-RIV forecasts yield smaller errors does not tell us anything about the size of volatility risk premiums on these markets. It would be wrong to believe that the ability of E-RIV to achieve smaller predicting errors implies incorrect implied volatility and hence incorrectly priced options. The volatility risk premium for each asset was already provided in Table 4.3. What the Table 4.5 illustrates is the fact that the utilization of volatility premium, whether positive or negative, can be employed to adjust current IV, bringing it closer to where it should be, in order to better reflect future realized volatility. Therefore, the 4.5 percentage point difference in MAPE should be interpreted solely as the discrepancy between two forecasts, unrelated to the average size of volatility risk premium.

In the energy market, both crude oil and natural gas saw E-RIV emerge as the most accurate forecasting model. Meats and metals markets showed A-RIV and E-RIV achieving very similar results, with negligible differences between the two methods. However, both models significantly outperformed simple IV forecasts. The bond market stands out as the only market where simple IV forecasts noticeably outperformed both alternative models, particularly on medium term maturities. This suggests that in these markets, there is generally a small usage of past volatility risk premium in making more accurate predictions. As anticipated, in the case of the SP 500, where initially observed volatility risk premium was notably high, any method adjusting to this premium would probably yield favorable results. Notably, E-RIV outperformed A-RIV in this scenario, although the difference was not substantial.

A potential room for further research could involve analyzing the optimal length of historical data required for forecasting. Specifically, exploring the optimal length of the averaging window used to calculate the A-RVRP and E-RVRP could shed light on the effectiveness of different models. In our study, we utilized a 252 day period. However, it would be valuable to investigate whether the newly proposed E-RIV model might benefit from alternative averaging lengths. Given that the averaging process of E-RIV is exponentially decreasing, excluding several observations at the end of averaging period should potentially not have a huge impact in the overall model efficiency. Additionally, exploring alternative averaging functions such as weighted moving averages or smoothed moving averages could result in insightful discoveries.

Another option for further research is to conduct a comprehensive analysis to determine the best averaging methodology which is most suitable for each market. We believe it is unlikely that there exists a universal averaging method that outperforms others across all markets. The unique characteristics of each market may favor different averaging functions, making certain methodologies more efficient in specific contexts. With these findings, a market participant would be able to choose the proper methodology for the prediction of realized volatility based on the market of his choice.

In summary, RIV models consistently achieved superior forecasts compared to IV, indicating that past volatility premium is indeed a valuable metric for adjusting current IV in the markets. The exponential weighting forecasts E-RIV outperformed both IV and A-RIV, suggesting its viability as a tool for generating more accurate forecasts of realized volatility. Moreover, it appears that recent volatility risk premium has a more pronounced impact on forecasts than historical VRP. Put differently, the influence of VRP on present conditions diminishes as we look further back in time. Therefore, understanding the current VRP level in the markets could be crucial for making informed predictions about the future.

4.4.2 Diebold-Mariano Test

To gather additional information about the relationship between different forecasts we employ the Diebold-Mariano test. This test acts as a supplementary measure to assess the robustness of the conclusions drawn from the previous loss functions. Under the null hypothesis, the forecasts of both models do not differ. We reject the null hypothesis at a significance level of $p \leq 0.05$.

Initially, we examine whether IV and RIV models forecasts exhibit differences. We conduct two DM tests: IV vs A-RIV (black) and IV vs E-RIV (red). Figure 4.12 depicts the DM test statistics, two for each asset. If the test statistic falls outside of the range from -1.96 to 1.96, we infer that the two forecasts differ significantly. Dashed lines have been included to denote the significance threshold. For specific test statistics and p-values we refer the reader to Table A.5.

Upon initial observation, we see that both A-RIV and E-RIV DM statistics tend to be closely aligned. The DM test statistics skew more towards positivity. There are numerous instances of very large DM statistics exceeding 5, while none fall below -5. Additionally, there are a few points between the dashed lines where definitive conclusions cannot be drawn. In total, A-RIV outperformed IV in 11 cases, while E-RIV outperformed IV in 9 instances. These findings indicate that the adoption of RIV models either enhances or does not make worse the forecasting accuracy of IV on RV. There were only two cases, where

Figure 4.12: DM Test: A-RIV vs IV & E-RIV vs IV

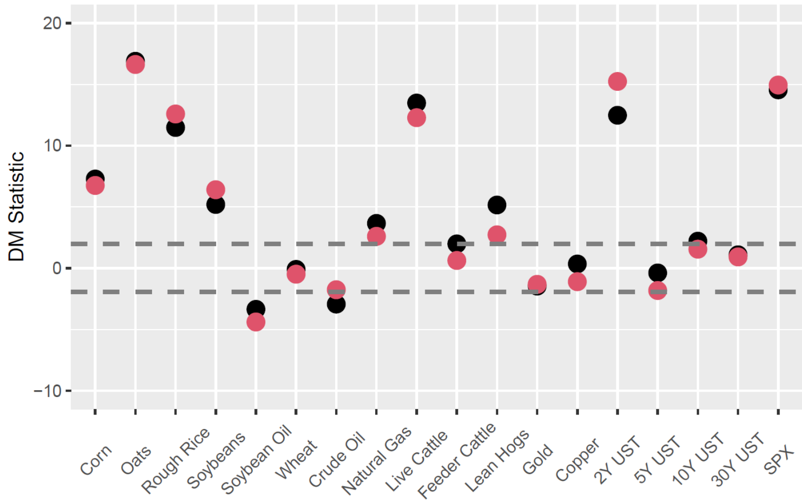

Source: Author's computations

the usage of IV was actually preferable. Soybean oil, which aligns from the findings of MAPE where IV also performed the best. Conversely, in the analysis of crude oil within the MAPE framework, both RIV models outperformed IV. However, the DM test statistics were significantly negative for both, indicating that IV forecasts were superior. There could be several reasons for this discrepancy, nonetheless, the DM test results generally align with the findings from Table 4.5. Adjusting for the past volatility premium can notably enhance the forecasting accuracy of IV.

Since we have observed the potential benefits of using RIV models, we are now curious about whether there exists a significant difference between the A-RIV and E-RIV forecasts. While the analysis of loss functions favored E-RIV as the superior model, the test statistics were often close to each other. Thus, the analysis with the Dienbold-Mariano test could potentially offer additional insights. The Figure 4.13 presents the DM test statistics of testing if A-RIV and E-RIV forecasts are significantly different. Interestingly, while the loss function statistics were often close to each other there were only four occurences in which we can not reject the null hypothesis. Overall we could say that E-RIV does not have a statistically better forecasts than A-RIV. However, there were several instances where it did show a better performance. Although the results are not as conclusive as those from the previous analysis, we see that in certain cases, the E-RIV model could be potentially advantageous. Notably, one of these cases is the SP 500 futures.

Figure 4.13: DM Test: A-RIV vs E-RIV

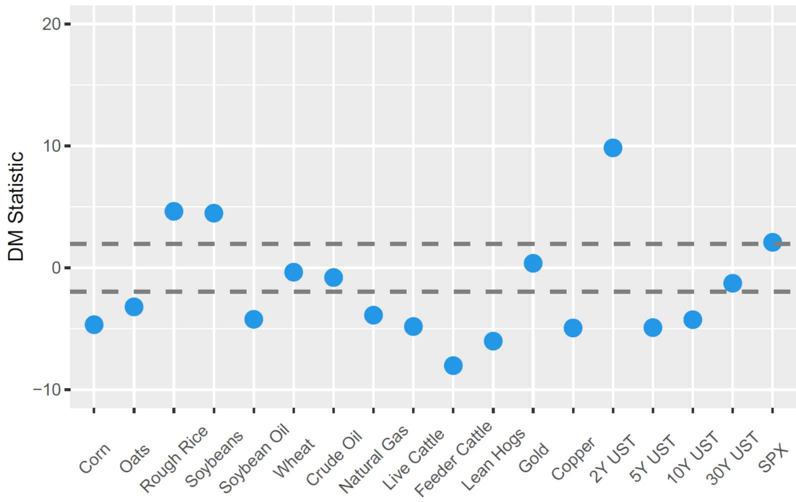

Source: Author's computations

The Diebold-Mariano tests have provided further confirmation that adjusting the current implied volatility for the past volatility risk premium can lead to improved forecasts in multiple instances. In these markets, it appears that the volatility risk premium is embedded in option prices and the knowledge of the past volatility level can help us adjust the current option price for the volatility premium. In the comparison between A-RIV and E-RIV, both models demonstrated effectiveness across different assets. This further confirms the notion that various markets may exhibit different patterns in the pricing of volatility risk. Further sensitivity analysis would be needed to examine the differences between both RIV methodologies on different markets. Exploring alternative averaging time frames or applying different functions could also potentially yield a deeper understanding of the effect of past volatility risk premium.

Chapter 5

Conclusion

This paper provides a focused review of the empirical finding that volatility risk premium can be used as a viable tool for volatility forecasting. Previous research has primarily focused on a limited set of assets, leaving gaps in our understanding. We are thus pioneers in analyzing this problem on a large number of markets, 18 in total. Our newly introduced RIV models, which adjust the current implied volatility for the volatility risk premium, exhibit the greatest forecasting effectiveness compared to IV and GARCH estimates. Our observation holds true across both short term and long term time frames and has been verified through the use of various loss functions and Diebold-Mariano tests. We introduce two distinct RIV models employing different methodologies to capture historical volatility risk premium. In most cases, either A-RIV or E-RIV emerged as the most effective predictive models. This also shows that each market might require a special treatment for the process of extracting past volatility risk premium for volatility forecasting. Our results prove the effectiveness of this approach on a variety of assets, further strengthening our conclusions.

As anticipated, the volatility risk premium was mostly positive across the majority of assets. Additionally, we demonstrate a connection between the magnitude of VRP and the trading volume of the underlying futures. Markets characterized by robust trading activity tended to generate positive VRP. Instances of negative VRP, observed in a few markets, are primarily associated with lower volumes in the underlying futures. This suggests that markets with lower trading volumes may lack the necessary depth to establish optimal option prices. Furthermore, our analysis illustrates that by computing the VRP as a ratio, we can effectively compare premiums across markets regardless of their average volatility levels.

Our paper contributes novel empirical insights and advances the knowledge of practical application of the volatility risk premium in financial markets. However, avenues for further research remain open as new questions emerge. While VRP has proven to be a viable metric, additional methods of its utilization warrant exploration. Within the framework of our study, further investigation into different averaging functions or averaging lengths could offer insights into the behavior of past VRP. Given the significant interest in volatility prediction in financial markets, it is likely that new forecasting methods will continue to emerge. These challenging issues are left for further research.

Bibliography

Agnolucci, P. (2003): "Volatility in crude oil futures: A comparison of the predictive ability of GARCH and implied volatility models." *Energy Economics* 531(2): pp. 316-321

Amihud, Y. (2002): "Illiquidity and stock returns: cross-section and time-series effects." *Journal of Financial Markets* 5(1): pp. 31-56

Alexander, C, (2008): "Market Risk Analysis, Pricing, Hedging and Trading Financial Instruments." *John Wiley & Sons* : pp. 283

Andersen, T.G. & Bollerslev, T. (1998): "Answering the skeptics: Yes, standard volatility models do provide accurate forecasts." *International Economic Review* 39(4): pp. 885–905

Bakshi, G. & Cao, C. & Chen, Z.(1997): "Empirical performance of alternative option pricing models." *Journal of Finance* 52: pp. 2003–2049

Bakshi, G. & Kapadia, N. & Medan, D.(2003): "Stock return characteristics, skew laws, and the differential pricing of individual equity options." *Review of Financial Studies* 16: pp. 101–143

Bakshi, G. & Kapadia, N. (2003): "Delta-hedged gains and the negative market volatility risk premium." *Review of Financial Studies* 16(2): pp. 527–566

Beckers, S. (1981): "Standard deviations implied in option prices as predictors of future stock price variability." *Journal of Banking & Finance* 5(3): pp. 363–381

Bennett, C. (2014): " Trading Volatility: Trading volatility, Correlation, Term Structure and Skew." *ISBN-13: 978-149920607*

Black, F. & Scholes, M. (1973): "The Pricing of Options and Corporate Liabilities." *The Journal of Political Economy* 81(3): pp. 637–654

Bollerslev T. (1986): "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31: pp. 307–327

Bollerslev T. & Tauchen, G. & Zhou, H. (2009): "Expected stock returns and variance risk premia." *Review of Financial Studies* 22(11): pp. 4463–4492

Bollerslev T. et al. (2011): "Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities." *Journal of Econometrics* 160: pp. 235–245

Brooks, C. (2014): "Introductory econometrics for finance." *Cambridge University press*, Vol.3

Carr, P. & Wu, L.(2009): "Variance risk premiums." *Review of Financial Studies* 22(3): pp. 1311–1341

Castelli, Ch. (1877): " The Theory of "options" in Stocks and Shares." *F.C. Mathieson, 1877; Wisconsin University*

Canina, L. & Figlewski, S. (1993): "The Informational Content of Implied Volatility." *Review of Financial Studies* 6(3): pp. 659–681

Chan, N.H. (2010): "Time Series: Applications to Finance with R and S-Plus." *Wiley, 2 edition*: pp. 210

Chen, S. & Li, G. (2023): "Why does Option Implied Volatility Forecast Realized Volatility? Evidence from News Events." *Journal of Banking & Finance* 156(1)

Cont, R. & Fonseca, J.d. (2002): "Dynamics of implied volatility surfaces" *Quantitative Finance* 2(1): pp. 45–60

Cont, R. (2005): "Volatility Clustering in Financial Markets: Empirical Facts and Agent–Based Models." Appeared in *Long memory in economics* by Kirman & Teyssiere

Cox J.C. et al. (1979): "Option pricing: A simplified approach." *Journal of Financial Economics* 7(3): pp. 229–263

Christensen, B.J. & Prabhala, N.R. (1988): "The relation between implied and realized volatility." *Journal of Financial Economics* 50(2): pp. 125–150

Christensen, B.J. & Hansen, C.S. (2010): "New evidence on the implied-realized volatility relation." *The European Journal of Finance* 8(2): pp. 187–205

Davis, M. & Etheridge, A. (2006): "Louis Bachelier's Theory of Speculation: The Origins of Modern Finance." *Princeton University Press*

Diavatopoulos D. et al. (2020): "The Predictive Power of REIT Implied Volatility and Implied Idiosyncratic Volatility." *Journal of Real Estate Portfolio Management* 16(1): pp. 29–38

Diebold, F.X. & Mariano, R.S. (1995): "Comparing predictive accuracy." *Journal of Business and Economic Statistics* 13: pp. 253–263

Driessen, J. & Maenhout, P. & Vilkov, G. (2009): "The price of correlation risk: Evidence from equity options." *Journal of Finance* 64(3): pp. 1377–1406

Engle, R.F. (1982): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50(4): pp. 987–1007

Engle, R.F. & Bollerslev, T. (1986): "Modelling the persistence of conditional variances." *Econometric Reviews* 5(1): pp. 1–50

Engle, R. (1995): "ARCH: Selected Readings." *Oxford University Press* ISBN 0-19-877431-1

Fama, E.F. (1965): "The Behavior of Stock-Market Prices." *The Journal of Business* 38(1): pp. 34–105

Fama, E.F. (1965): "Random Walks in Stock Market Prices." *Financial Analysts Journal* 21(5): pp. 55–59

Feri Trust GMBH(2017): "Harvesting the S&P 500 Volatility Risk Premium." *The Hedge Fund Journal* 121

Gatheral, J. (2006): "The Volatility Surface: A Practitioner's Guide." John Wiley & Sons, ISBN 0471792519

Ge, W. (2016): "A Survey of Three Derivative-Based Methods to Harvest the Volatility Premium in Equity Markets." *The Journal of Investing* 25(3): pp. 48–58

Granger, C.W.J. (1969): "Investigating Causal Relations by Econometric Models and Cross-spectral Methods." *Econometria* 37(3): pp. 424–438

Han, Y. (2011): "On the relation between the market risk premium and market volatility." *Applied Financial Economics* 21(22): pp. 1711–1723

Hansen, P.R. & Lunde, A. (2005): "A forecast comparison of volatility models: does anything beat a GARCH(1,1)?." *Journal of Applied Econometrics* 20(7): pp. 873–889

Hudson, R. & Gregoriou, A. (2010): "Calculating and Comparing Security Returns is Harder than you Think: A Comparison between Logarithmic and Simple Returns." *International Review of Financial Analysis* 38

Hull, J.C. (2003): "Options, Futures and Other Derivatives." *(5th ed.). Upper Saddle River, NJ: Prentice Hall. ISBN 0-13-046592-5*

Jackwerth, J. & Rubinstein, M. (1996): "Recovering Probability Distributions from Option Prices." *The Journal of Finance* 51: pp. 1611–1631

Jorion, P. (1995): "Predicting Volatility in the Foreign Exchange Market." *The Journal of Finance* 50(2): pp. 507–528

Kambouroudis D. S. et al. (2016): "Forecasting Stock Return Volatility: A Comparison of GARCH, Implied Volatility, and Realized Volatility Models." *The Journal of Future Markets* 36(12): pp. 1127–1163

Latane, H. A. & Rendleman, R. J. (1976): "Standard Deviations of Stock Price Ratios Implied in Option Prices." *The Journal of Finance* 31(2): pp. 369–381

Liu L. et al. (2012): "Does Anything Beat 5-Minute RV? A Comparison of Realized Measures Across Multiple Asset Classes ." *Duke University and University of Oxford*

Lowenstein, R. (2001): "When Genius Failed: The Rise and Fall of Long-Term Capital Management." *(Fourth Estate). ISBN 978-1-84115-504-3*

Lu, S. (2020): "Harversing Volatility Risk Premium." *Imperial College London - Department of Mathematics*

Mandelbrot, B. (1963): "The Variation of Some Other Speculative Prices." *The Journal of Business* 40(4): pp. 393–413

McMillan, L.G. (1980): "Options as a strategic investment." *New York Institute of Finance*

Pasquale, D. C. et al. (2016): "Volatility risk premia and exchange rate predictability." *Journal of Financial Economics* 120(1): pp. 21–40

Prokopczuk, M. & Simen, C. W. (2014): "The importance of the volatility risk premium for volatility forecasting." *Journal of Banking & Finance* 40: pp. 303–320

Shephard, N. (2010): "Stochastic volatility models." *Macroeconometrics and Time Series Analysis. The New Palgrave Economics Collection. Palgrave Macmillan, London*

Szakmary, A. et al. (2003): "The predictive power of implied volatility: Evidence from 35 futures markets." *Journal of Banking & Finance* 27(11): pp. 2151–2175

Toda, H.Y. & Yamamoto, T. (1995): "Statistical inference in vector autoregressions with possibly integrated processes." *Journal of Econometrics* 66(1): pp. 225–250

Trolle, A. & Schwartz, E. (2009): "Unspanned stochastic volatility and the pricing of commodity derivatives." *Review of Financial Studies* 22(11): pp. 4423–4461

Zhou, H. (2018): "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty." *Annual Review of Financial Economics* 10: pp. 481–497

Appendix A

Appendix

Table A.1: Log-returns - Descriptive statistics

	Statistic						
Asset	Mean	Median	St. d.	Min	Max	Skew	Ex. Kurt.
Grains							
Corn	0.0132	0.0423	1.82	−26.86	9.17	−1.56	24.77
Oats	0.0064	0.0000	2.36	−34.20	18.82	−1.23	24.87
Rough Rice	0.0050	−0.0316	1.59	−29.97	9.80	−2.15	47.15
Soybean	0.0103	0.0753	1.43	−11.09	7.34	−0.90	10.15
Soybean Oil	0.0150	−0.0305	1.47	−10.71	7.34	−0.38	6.51
Wheat	0.0088	−0.0484	2.05	−12.84	19.70	0.37	7.95
Energy							
Crude Oil	−0.0010	0.0948	2.90	−60.17	31.96	−2.73	79.62
Natural Gas	0.0000	0.0000	3.38	−30.05	38.17	0.27	12.76
Meats							
Live Cattle	0.0183	0.0445	1.19	−15.65	7.00	−1.29	19.70
Feeder Cattle	0.0206	0.0198	1.12	−8.61	10.47	0.39	13.47
Lean Hogs	0.0071	0.0668	2.50	−23.46	23.63	−1.08	25.73
Metals							
Gold	0.0151	0.0335	1.04	−9.82	6.24	−0.53	9.28
Copper	0.0039	0.0000	1.46	−7.55	8.57	−0.04	5.24
Bonds							
2Y UST	−0.0019	0.0000	0.07	−0.84	0.55	−1.10	17.09
5Y UST	−0.0020	0.0062	0.21	−1.48	1.39	−0.40	7.28
10Y UST	−0.0005	0.0119	0.33	−2.18	1.96	−0.21	5.86
30Y UST	0.0035	0.0205	0.67	−4.11	9.88	0.88	19.71
Equities							
E-MINI SP 500	0.0406	0.0694	1.15	−10.96	9.34	−0.82	15.11

Source: Author's computations

Figure A.1: Prices

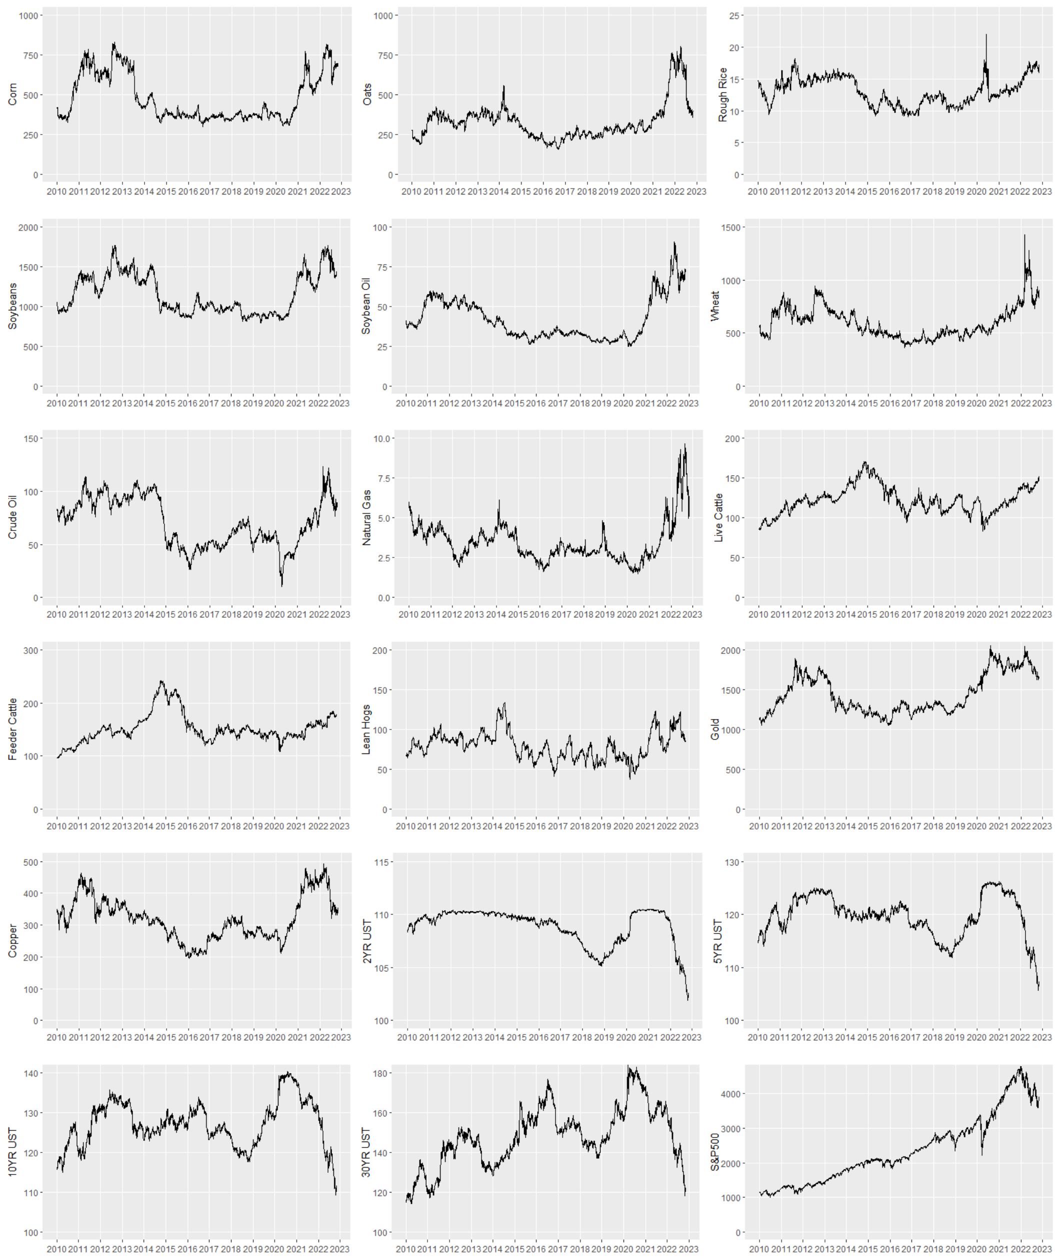

Source: Author's computations

Figure A.2: Volatility Risk Premium

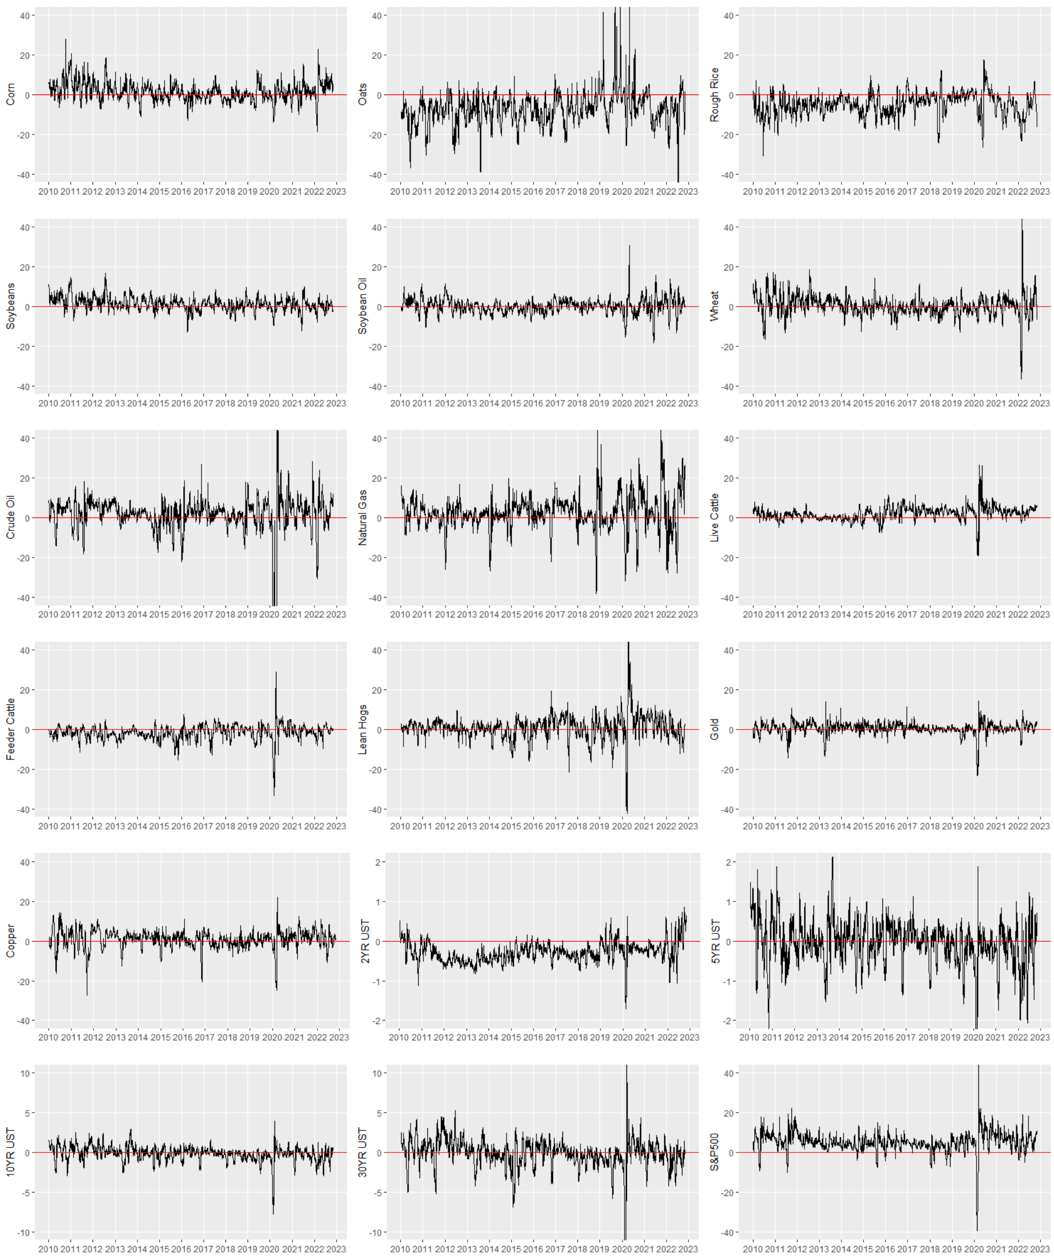

Source: Author's computations

Table A.2: Futures Contracts Specification and Volume

Asset	Contract Size	Unit	Price Quotation	Daily Volume	Contract Notional	Notional Volume[a]
Corn	5 000	bushels	cent/bushel	66 967	22 950	1 537
Oats	5 000	bushels	cent/bushel	178	19 150	3
Rough Rice	2 000	cwt	cent/cwt	276	35 000	10
Soybean	5 000	bushels	cent/bushel	51 000	61 800	3 185
Soybean Oil	60 000	pounds	cent/pound	25 000	28 800	728
Wheat	5 000	bushels	cent/bushel	24 467	30 500	746
Crude Oil	1 000	barrel	dollars/barrel	92 667	71 000	6 579
Natural Gas	10 000	MMBtu	dollars/MMbtu	36 167	30 000	1 085
Live Cattle	40 000	pounds	cents/pound	11 433	28 800	329
Feeder Cattle	50 000	pounds	cents/pound	11 867	68 000	807
Lean Hogs	40 000	pounds	cents/pound	3 167	113 000	358
Gold	100	troy ounces	dollars/ounce	41 067	1 000	41
Copper	25 000	pounds	dollars/pound	109 000	2 000	218
2Y UST	200	dollars	% of par	323 000	204 000	65 960
5Y UST	100	dollars	% of par	680 000	108 000	73 440
10Y UST	100	dollars	% of par	1 003 000	111 000	111 370
30Y UST	100	dollars	% of par	163 000	122 000	19 927
E-MINI SP500	50	dollars	SPX	1 030 000	240 000	248 000

Source: CME Group, Interactive Brokers, Author's computations

[a] Millions USD

Figure A.3: Effects of Notional Value on RVRP

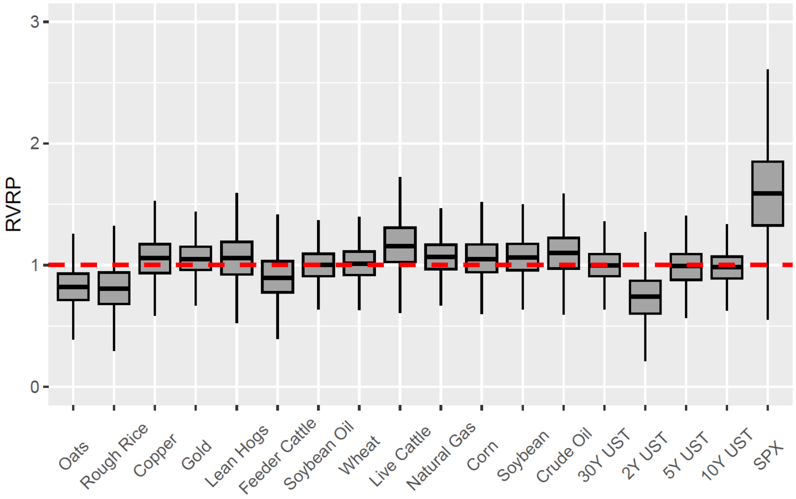

Source: Author's computations

Table A.3: Loss Functions - Forecast #1

	MAE				**MSE**				**MSPE**			
Asset	GARCH	IV	A-RIV	E-RIV	GARCH	IV	A-RIV	E-RIV	GARCH	IV	A-RIV	E-RIV
Corn	7.0669	4.2759	**3.7300**	3.7706	116.362	32.668	**27.550**	28.306	0.1432	0.0427	**0.0340**	0.0355
Oats	8.7976	8.1178	7.5400	**7.3306**	152.307	125.724	**100.070**	100.912	0.1288	0.0532	0.0490	**0.0488**
Rough Rice	8.8145	6.9958	6.2000	**6.1899**	322.002	78.377	65.270	**64.782**	0.8859	**0.1223**	0.1300	0.1271
Soybean	4.7369	**2.5072**	2.6100	2.6152	44.917	**11.258**	12.430	12.324	0.0828	**0.0243**	0.0260	0.0263
Soybean Oil	4.8490	**4.3068**	4.4100	4.4829	39.647	**33.712**	36.360	37.597	**0.0344**	0.0353	0.0390	0.0405
Wheat	5.4566	**4.3921**	4.6500	4.6396	88.314	**58.108**	61.270	62.160	0.0432	**0.0259**	0.0290	0.0295
Crude Oil	19.4908	12.6203	**12.4100**	12.6178	1540.901	**746.629**	778.940	795.038	0.2504	0.0667	**0.0600**	0.0624
Natural Gas	15.4023	10.8471	**10.2000**	10.5859	614.164	187.724	**169.690**	182.233	0.1431	0.0470	**0.0380**	0.0422
Live Cattle	6.4700	4.3020	**2.4000**	2.4286	77.550	30.008	**17.580**	18.676	0.3860	0.1340	**0.0350**	0.0358
Feeder Cattle	5.0300	**3.1871**	3.4100	3.4800	57.180	**30.199**	32.370	34.045	0.1320	**0.0433**	0.0540	0.0554
Lean Hogs	22.8100	6.4627	**5.5900**	5.6873	550.260	97.739	**88.040**	90.688	0.9700	0.0778	**0.0550**	0.0566
Gold	3.3800	2.4891	**2.6000**	2.6059	25.520	**17.883**	18.790	19.248	0.0600	**0.0368**	0.0390	0.0391
Copper	3.9400	3.8469	**3.5600**	3.5721	34.930	28.682	**27.230**	27.797	0.0380	0.0380	**0.0300**	0.0316
2Y UST	0.3100	0.2904	0.3000	**0.2849**	0.220	**0.146**	0.190	0.175	0.0680	0.0893	0.0800	**0.0736**
5Y UST	0.7200	**0.5344**	0.5600	0.5641	1.150	0.711	**0.670**	0.686	0.0930	**0.0346**	0.0360	0.0387
10Y UST	1.1900	**0.8170**	0.8600	0.8681	3.460	1.958	**1.870**	1.911	0.0790	**0.0295**	0.0350	0.0362
30Y UST	2.3600	**1.6343**	1.7900	1.7853	18.270	**11.415**	11.620	11.765	0.0820	**0.0314**	0.0420	0.0414
E-MINI SP500	7.9800	8.1992	4.0200	**3.9986**	162.330	104.586	56.760	**56.445**	0.4760	0.5659	0.0790	**0.0785**
Σ hits	0	6	8	4	0	8	7	2	1	8	6	3

Source: Author's computations

Table A.4: Loss Functions - Forecast #2

Asset	MAE			MSE			MSPE		
	IV	A-RIV	E-RIV	IV	A-RIV	E-RIV	IV	A-RIV	E-RIV
Corn	3.4687	**3.1888**	3.2110	22.323	**18.700**	19.243	0.0348	**0.0278**	0.0289
Oats	8.1267	6.4115	**6.3957**	109.370	**72.672**	73.853	0.0642	**0.0519**	0.0532
Rough Rice	5.5644	4.3941	**4.2936**	50.605	37.595	**36.282**	0.0755	0.0660	**0.0636**
Soybean	2.5129	2.3803	**2.3494**	11.063	9.872	**9.621**	0.0304	0.0255	**0.0246**
Soybean Oil	**2.6333**	2.6551	2.6703	**13.680**	14.405	14.694	**0.0232**	0.0239	0.0240
Wheat	3.2918	3.2752	**3.2531**	**25.682**	25.715	25.807	0.0206	0.0200	**0.0197**
Crude Oil	6.6039	6.1695	**6.1361**	**210.385**	214.900	218.037	0.0502	0.0376	**0.0372**
Natural Gas	6.3930	6.0411	**6.0308**	85.671	**80.046**	81.990	0.0323	**0.0266**	0.0267
Live Cattle	3.0169	**1.9588**	1.9678	15.673	**8.892**	9.232	0.0754	**0.0298**	0.0301
Feeder Cattle	2.9560	**2.8303**	2.8853	17.751	**16.595**	17.372	0.0437	**0.0434**	0.0442
Lean Hogs	4.3596	**4.0823**	4.1092	42.160	**39.679**	40.744	0.0521	**0.0431**	0.0435
Gold	2.0306	1.9950	**1.9751**	**9.817**	10.020	10.000	0.0304	0.0281	**0.0276**
Copper	3.1854	**2.9728**	3.0044	19.529	**19.434**	19.836	0.0333	**0.0278**	0.0286
2Y UST	0.3367	0.2297	**0.2230**	0.153	0.101	**0.095**	0.1123	0.0595	**0.0555**
5Y UST	**0.4015**	0.4124	0.4194	**0.347**	0.350	0.359	**0.0262**	0.0277	0.0286
10Y UST	0.6328	**0.6313**	0.6367	0.937	**0.886**	0.902	0.0230	**0.0229**	0.0233
30Y UST	1.2621	1.2402	**1.2388**	4.632	**4.486**	4.509	**0.0231**	0.0239	0.0236
E-MINI SP500	5.8444	2.5170	**2.4951**	51.219	20.857	**20.664**	0.5044	0.0686	**0.0675**
Σ hits	2	6	10	5	9	4	3	8	7

Source: Author's computations

Table A.5: Diebold-Mariano Test Results

Asset	**DM #1** IV / A-RIV	**DM #2** IV / E-RIV	**DM #3** A-RIV / E-RIV
Corn	**7.27** (0.00)	**6.74** (0.00)	-4.67 (1.00)
Oats	**16.88** (0.00)	**16.61** (0.00)	-3.21 (1.00)
Rough Rice	**11.47** (0.00)	**12.58** (0.00)	**4.63** (0.00)
Soybean	**5.19** (0.00)	**6.40** (0.00)	**4.48** (0.00)
Soybean Oil	-3.36 (1.00)	-4.40 (1.00)	-4.23 (1.00)
Wheat	-0.10 (0.54)	-0.50 (0.69)	-0.36 (0.64)
Crude Oil	-2.94 (1.00)	-1.76 (0.96)	-0.80 (0.79)
Natural Gas	**3.65** (0.00)	**2.59** (0.00)	-3.89 (1.00)
Live Cattle	**13.47** (0.00)	**12.27** (0.00)	-4.82 (1.00)
Feeder Cattle	**1.98** (0.02)	0.62 (0.27)	-8.03 (1.00)
Lean Hogs	**5.15** (0.00)	**2.72** (0.00)	-6.02 (1.00)
Gold	-1.47 (0.93)	-1.32 (0.91)	0.37 (0.36)
Copper	0.34 (0.37)	-1.11 (0.87)	-4.94 (1.00)
2Y UST	**12.47** (0.00)	**15.23** (0.00)	**9.84** (0.00)
5Y UST	-0.39 (0.65)	-1.83 (0.97)	-4.91 (1.00)
10Y UST	**2.20** (0.01)	1.54 (0.18)	-4.26 (0.90)
30Y UST	1.09 (0.14)	0.91 (0.18)	-1.27 (0.90)
E-MINI SP500	**14.54** (0.00)	**14.94** (0.00)	**2.09** (0.02)
Σ hits	11	9	4

Source: Author's computations