Erasmus School of Economics

Master Thesis Quantitative Finance - MSc Econometrics and Management Science

The Volatility Knows Where It Is, Because It Knows Where It Isn't

Author:

Burak Mercan (508907)

Second Assessor:

Prof. Dr. Chen Zhou

Supervisors:

Dr. Rutger-Jan Lange

Corjan Overbeeke

Anne Bakx

July 18th 2023

Abstract

This research compares different filtering methods applied to an adjusted Heston model to estimate commodity volatilities, for which the realized volatility (RV) is utilized as a proxy. The study focuses on the Extended Kalman Filter (EKF), Unscented Kalman Filter (UKF), and Particle Filter (PF), making use of the CME Group Volatility Index (CVOL), as the observed variable. Benchmark models include the HAR-RV and AR(1) models. Evaluation involves comparing the in-sample forecasts in graphs and using robust loss metrics. Results indicate that the EKF and UKF methods perform well, closely aligning with the RV. The PF method outperforms EKF and UKF, capturing volatility dynamics more accurately. However, both UKF and PF occasionally underestimate the RV. The PF consistently performs well, surpassing benchmark models, while the HAR-RV model shows good accuracy for specific commodities. The research enhances volatility estimation techniques in commodity markets.

Keywords: Stochastic Volatility - Heston Model - Filtering Methods

Contents

1 Introduction

"The missile knows where it is at all times. It knows this because it knows where it isn't." This famous segment of an Air Force training video from 1997 summarizes the idea of using filtering methods. It also shows the wide field of application in which the filtering methods can be utilized. Especially, in the field of finance it is applicable, namely to handle methods and models in which latent variables are an important factor. Examples include, but are not limited to, probability of default in credit risk models, interest rate models and stochastic volatility models (SVM). In this paper we will focus on the latter, and implement filtering methods on an SVM.

SVM are utilized for various financial applications, for example in risk management, portfolio optimization and derivative pricing. The main assumption made by SVM is that the volatility of an asset follows a dynamic process, instead of being fixed, which was the assumption made by early literature models like Black-Scholes (BS), in 1973. The need for these models arose from the fact that the assumption of constant volatility violates behavior of assets in reality. The Heston model (HM) is an example of an SVM with appropriate features that captures characteristics of assets in practice. Furthermore, because true underlying volatilities are unknown in SVMs, filtering methods can be implemented to estimate these latent volatilities.

We attempt to utilize different filtering methods on an adjusted HM to estimate the underlying volatilities. In all filtering methods we iteratively predict the state of the dynamic volatility system based on the HM, and consequently filter this predicted state based on noisy volatility observations from the available information set. Specifically, three different filtering methods are researched. Because of the non-linearity of the HM, filtering methods that can handle the non-linearity are required. First, we have the extended Kalman filter (EKF). The EKF is a variant of the traditional Kalman filter (KF) (Kalman, 1960). It extends the KF by applying a first-order Taylor series expansion on the model to approximate the non-linearity of the system. Second, we examine the unscented Kalman filter (UKF). The UKF is another variant of the KF to handle non-linear models. Instead of linearization, which the EKF utilizes, the UKF approximates the probability distribution of the system's state with a Gaussian probability distribution using a set of weighted sampling points. Third, we study the particle filter (PF). The PF makes use of the same idea as the UKF, the difference is that the PF allows for

non-Gaussian approximation. Because the true volatilities are unknown, exact maximum likelihood estimation is computationally expensive to estimate the parameters of the HM. To solve this problem we first employ the quasi-maximum likelihood estimation (QMLE) to calibrate the parameters, proposed by Wedderburn (1974), and more recently utilized in Alizadeh et al. (2002) and Xiu (2010) for this subject. Thereafter, the filtering methods are used to update the volatilities.

The Realized Volatility (RV) of the commodities is utilized as a proxy to evaluate the CME Group Volatility Index (CVOL) and its, by filtering methods, improved version. The CVOL is a measure of the implied volatility (IV) of commodities to quantify the market's expectation of the 30-day forward risk. The performance of the CVOL and the three different filtering methods are assessed by robust loss values between the RV and the in-sample volatility estimates. Subsequently, they are compared to each other to determine the best performing filtering method and utilized to robustify the CVOL. Moreover, benchmark models, such as the HAR-RV model, proposed by Corsi (2009) and AR(1) model, are used to evaluate the effectiveness of the filtering methods on the HM.

The aim of this research is to investigate which filtering method performs best. We expect the CVOL to overestimate the RV, because of its inherent risk premium for the seller. So, we make use of the filtering methods to find a more robust IV measure, which has a better predictive accuracy of the RV than the CVOL, by using the CVOL as a noisy IV measure. Therefore, the main question of this research is formulated as follows: *"Which filtering method in combination with quasi-maximum likelihood estimation performs best, and can it improve the CVOL of different commodities in practice?"*.

We are provided data by the company Transtrend. Transtrend is a commodity trading advisor firm with expertise in trend following investments. The data contains the CVOL observations of various commodities for the period October 2013 up until and including December 2022. On top of that, futures price data of the corresponding commodities are provided for the same time frame.

Our analysis indicates that the EKF, UKF, and PF filtered estimates outperform the CVOL measure in estimating the RV for various commodities. The EKF and UKF models demonstrate a closer alignment with the RV and provide more accurate estimates compared to the CVOL measure. However, the UKF consistently outperforms the EKF, capturing non-linearity better. The PF method exhibits the highest accuracy, closely

aligning with the RV estimates for corn, crude oil, gold, natural gas, and soybean, and achieves the lowest loss values for cases in which under- and over-prediction are penalized more or equally. Nonetheless, both the UKF and PF methods occasionally underestimate the RV. Overall, these findings highlight the superiority of the UKF and PF techniques in capturing complex volatility dynamics. The PF filtered CVOL consistently outperformed the benchmark models, displaying a close alignment with RV values and strong predictive accuracy. The HAR-RV model also performed well, particularly excelling in in-sample forecasting corn and soybean. Although the AR(1) benchmark model showed reasonable performance and better in-sample accuracy than the CVOL, it fell slightly behind the other models.

The empirical results from this research can be used in practice. For example, a filtering method combined with an SVM can function as a method to improve IV measures. Moreover, the formulated methods and models can be used in academic research as benchmark models in other empirical studies where filtering methods are used on SVMs to develop volatility estimates.

The remainder of the paper proceeds as follows. In Section 2, we provide a brief literature review. The relevant data is described in Section 3. Next, our methods are introduced in Sections 4 and 5. Afterwards, the results are presented in Section 6. Finally, we conclude and discuss our results in Section 7.

2 Literature

The research of modelling volatility has been a key subject in financial literature and applications. There are different types of volatility that are researched in the literature. The most common, and also the ones that will by analyzed in this paper, are historical volatility and IV. The RV is widely used as a benchmark to assess how well volatility estimates are doing when it comes to historical volatility. This is because it is an unbiased proxy of the actual quadratic variation. The RV was initially explored and used in works such as Andersen et al. (2005) and Andersen and Teräsvirta (2009) after being first introduced in Andersen and Bollerslev (1998). The RV is a more accurate volatility estimator than other historical volatility measures that are derived using daily returns rather than intra-day returns (evidence in Appendix A). Furthermore, Liu et al. (2015)

asserts that alternative volatility estimates do not, on average, significantly outperform the 5-minute RV.

On the other hand we have IV, which is a measure that captures the market's view on the volatility of a certain financial instrument in the future. The literature workhorse model to calculate the IV from market prices of options was first introduced in Black and Scholes (1973), called the BS model. The BS model assumes that the price of an equity follows a geometric Brownian motion with parameters that are constant. Although the BS model has nice features, for example it gives a closed-form solution for pricing European options, it does not capture characteristics of actual market data. In the literature we can read about these characteristics, that are not described by the BS model. First of all, Cont (2001) found that the returns on equity prices follow a distribution that has high peaks and is fat-tailed. In addition, Engle and Ng (1993) conclude that equity prices and volatilities are negatively correlated. Last of all, the volatilities of returns in equity prices have mean-reverting and clustering features, stated in Engle and Patton (2001) and Fama (1965), respectively.

As these properties are not captured by the BS model, the literature tried to improve on it and come up with developments that do incorporate these features. For example, Wiggins (1987) and Hull and White (1987) adjusted the BS model to allow for stochastic volatility, which are build on two stochastic diffusion processes that are linked, namely the equity and volatility processes. Despite the fact that these models are successful in describing the equity price movements, they do not have closed-form solutions. Another method was introduced in Heston (1993), which we refer to as the HM, that utilizes a stochastic volatility model, which follows a square root mean-reverting process. In this way the HM is able to capture the features of real market data, the volatilities are non-negative, by construction, and on top of that the HM can accommodate for the "smile" effect of IV, discussed in Weron and Wystup (2005) and Dupire (1997). The most significant advantage, for practical use of the HM, is that it has a closed-form solution for pricing. However, because the parameters of the model and the true underlying volatilities are not observed, the estimation procedure is not straightforward and needs special attention.

In the early literature the idea of moment matching was often used. For example, the method of moments proposed by Taylor (2008), the generalized method of moments dis-

cussed by Andersen and Sørensen (1996) and Melino and Turnbull (1990) and the efficient method of moments stated in Andersen et al. (1999). The disadvantage of these methods is that they are computationally expensive, due to optimization in high-dimensional parameter spaces, handling of non-analytical solutions, and accommodation for model complexity.

More state-of-the-art research focuses on filtering methods. For instance, in Ruiz (1994) the asymptotic and finite sample properties of a quasi-maximum likelihood estimator on linear SVMs based on the KF was analyzed. However, because the HM is non-linear, Javaheri et al. (2003) looked into estimating the model using the EKF. Furthermore, PF methods, that deal better with the non-linearity of the HM, are utilized on SVMs in for example Aihara et al. (2009). But, because of the computational limitations of the PF methods, an approach called the UKF could be utilized, which is considered better in performance but has the same computational properties as the EKF. Corresponding literature that makes use of and discusses the application of the UKF on SVMs are for example Zoeter et al. (2004) and Li (2013).

We distinguish from literature by implementing the filtering methods on an adjusted HM to analyze and robustify volatilities of commodities, while most papers only assess the performance of the various methods by a simulation study, and sometimes include applications on stock indices.

3 Data

3.1 Description

The data for this research is provided by Transtrend. The data consists of tick and daily prices of various commodity futures. The data also includes daily observations of the CVOL. Furthermore, the time period that is used to perform research is October 2013 up until and including December 2022. This sample period is chosen because of availability of CVOL data.

The proxy that is utilized to evaluate the CVOL is the RV. The methodology on how the RV is calculated in this research is described below. For the details of the calculation of the CVOL and its methodology we refer to the CME Group Volatility Index Benchmark

Methodology.[1]

3.2 Realized Volatility

For the calculation of the RV we exploit the 5-minute intra-day natural logarithm returns during market hours such that the return is

$$r_{i,j,t} = \log(p_{i,j,t}/p_{i,j-1,t}). \tag{1}$$

In this equation $p_{i,j,t}$ is the j^{th} price observation of commodity i on day t. Using this definition, the RV is defined as follows

$$RV_{i,t} = \sum_{j} r^2_{i,j,t}. \tag{2}$$

However, this measure for volatility neglects information contained in overnight returns. To incorporate this information to get a more robust RV, Hansen and Lunde (2005) introduced a scaled version of the RV, which is of the form

$$RV^{'}_{i,t} = \hat{c} RV_{i,t}, \tag{3}$$

where

$$\hat{c} = \frac{\sum_{t=1}^{T}(r_{i,t} - \hat{\mu}_i)^2}{\sum_{t=1}^{T} RV_{i,t}}. \tag{4}$$

Here we have that $r_{i,t}$ is the daily settle-to-settle natural logarithm return of commodity i, and $\hat{\mu}_i$ is the mean of $r_{i,t}$ over the sample period.

3.3 Future Selection

In this section we discuss the future selection procedure (FSP) to determine which future contracts to use to calculate the realized volatility and daily settle-to-settle returns. As the various commodities have characteristics that differ, the FSP for each of them will also be different.

First, we discuss the FSP of gold, the precious metal commodity we investigate in our research. In case of gold futures, most of the time, there is only one contract of main

[1] https://www.cmegroup.com/market-data/cme-group-benchmark-administration/files/cvol-methodology.pdf

interest and traded the most. So, the FSP of gold for a certain day is to sort out the contract that has the most trading volume on daily basis and utilize that one.

Second, we have the agricultural commodities. These are corn and soybean. The FSP for these products is the same as for gold, however there is a small difference. This follows from the fact that there are harvest seasons for the grain products. Throughout the year, contracts with different expiration dates also have different crops as an underlying. For corn futures the underlying changes to the new crop in the December contract and for soybean in the November contracts. So basically, these contracts have different underlying assets, as the products are the same but the crops are different, so these contracts could have different characteristics. To take this into account, the FSP for agricultural commodities does not switch to the December contract, with underlying the new crops, until at least 30 days before days to expire (DTE) of the current contract, with underlying the current crops, even though if the December contract has the most trading volume.

Last of all, the FSP of energies is examined. In case of energy commodities, which are crude oil and natural gas in our analysis, we often observe that two contracts have a high volume of trades at the same time. To incorporate information from both contracts in the daily returns and realized volatility, a future continuation scheme is proposed and weighted accordingly. To conclude, the FSP for these energy commodities is to consider the two most active future contracts, consequently weigh them on their importance based on their DTE. The further the expiration date of the contract is from 30 days in the future, the less important the contract becomes. The weights for the daily returns and realized volatilities are calculated the same way as is done in the CME CVOL methodology for IV, which is

$$w_1 = \frac{|DTE_2 - (t+30)|}{|DTE_2 - (t+30)| + |(t+30) - DTE_1|}, \tag{5}$$

$$w_2 = \frac{|(t+30) - DTE_1|}{|DTE_2 - (t+30)| + |(t+30) - DTE_1|}, \tag{6}$$

where t is the current date on which the calculation is being performed. DTE_1 and DTE_2 are the DTE for the most and second most active future contract, respectively. And finally, w_1 and w_2 are the weights for the most and second most active future contract, respectively.

3.4 Summary Statistics

In Table 1 and Table 2 the CVOL and RV summary statistics of different commodities are outlined. We observe that for all commodities the RV is on average lower than the CVOL. From this we can conclude that the markets view on the volatility of these commodities, in the upcoming 30 days, is on average higher than what is realized. This fact is in line with what generally is observed in practice and research, that IV is higher than RV. The reason behind these observed results are thought to be to compensate for the risk the seller is taking and hedging costs. We also notice that the standard deviation of the RV is higher than their IV counterparts, except for corn and natural gas. This means that the RV does differ more from period to period than the CVOL for crude oil, gold, and soybean and less for corn and natural gas. Another noticeable aspect is that the CVOL has both a higher skewness and a larger kurtosis than the RV for the energy commodities. This infers the fact that the RV is more positively skewed and more leptokurtic than the CVOL, which corresponds to the tails of the distribution of the RV being larger than that of the CVOL, for crude oil and natural gas. This could be due to the fact that the markets view on volatility for the upcoming 30 days, in peak volatility periods is exaggerated and highly overestimated. For corn, gold and soybean the opposite holds. This probably follows from the fact that the precious metal and agricultural commodities have less peaking volatility periods. For energy commodities the level of volatility peaked much more from time to time. Examples include but are not limited to the energy crisis and the price of a certain, close to expiration, oil future going negative. These occurrences caused the volatility to peak a lot for energy commodities. Last of all, we note that for all commodities the maximum and minimum values are lower for the RV.

Table 1: Summary statistics CVOL of the various commodities

Commodity	Mean	St. Dev.	Skew	Kurt	Max	Min
Corn	24.77	7.82	1.05	1.70	66.99	10.37
Crude Oil	39.47	23.63	6.40	72.32	435.29	13.80
Gold	15.51	4.32	1.40	5.65	51.03	8.18
Natural Gas	51.74	24.25	1.31	1.48	152.23	19.90
Soybean	20.33	5.00	0.71	0.70	43.19	10.55

Table 2: Summary statistics realized volatilities of the various commodities

Commodity	Mean	St. Dev.	Skew	Kurt	Max	Min
Corn	21.56	6.85	1.31	2.35	49.82	9.52
Crude Oil	36.35	28.99	6.00	45.98	292.36	9.51
Gold	14.02	4.42	2.48	11.81	43.20	7.78
Natural Gas	45.19	20.54	1.03	0.28	117.00	17.25
Soybean	18.63	5.18	0.83	0.87	39.79	9.20

3.5 CVOL/RV Comparison

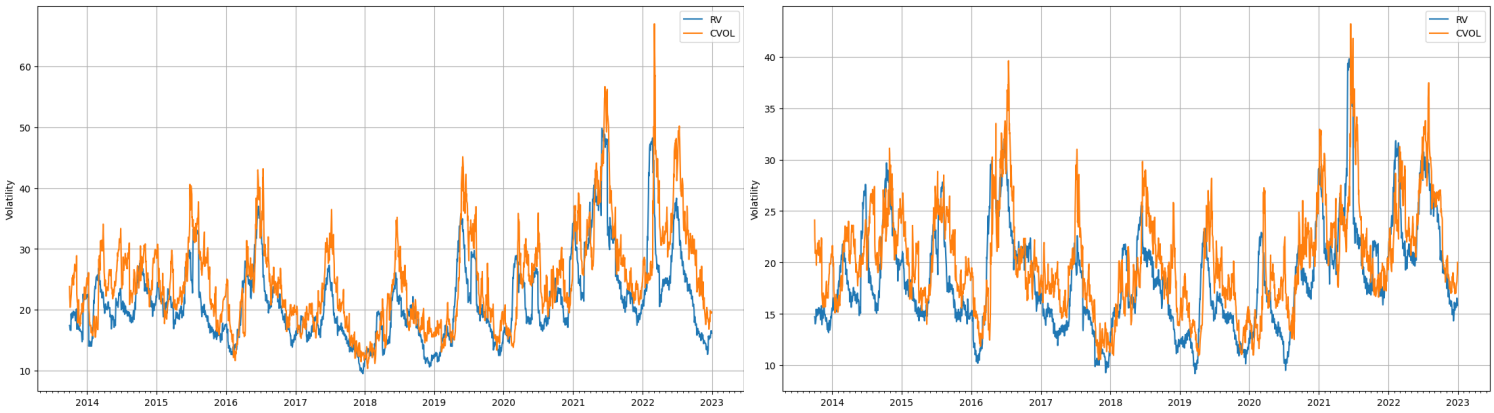

(a) Corn **(b)** Soybean

Figure 1: CVOL and RV of agriculture products

In Figure 1 we observe the CVOL and RV of agricultural commodities. The first thing that we notice is that the RV for both corn and soybean are in general lower than the CVOL. This corresponds with their means. It seems though that in normal volatility periods the CVOL approaches the RV, however in peak volatility periods it overshoots the RV a lot, like in early 2022 at the start of the Russo-Ukrainian War, and in low volatility periods it overestimates the volatility. Both volatility measures also seem to be mean-reverting.

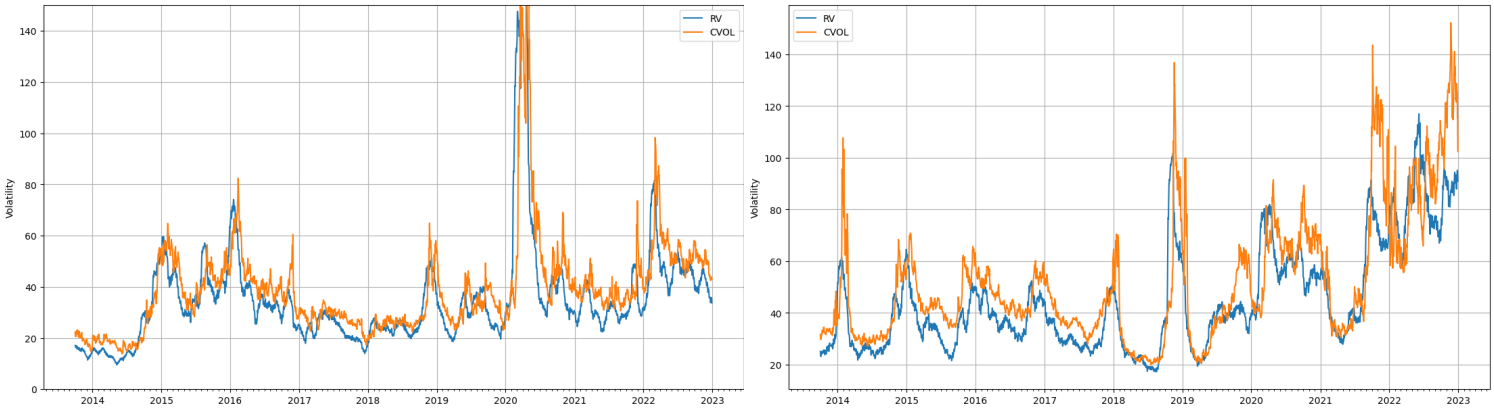

(a) Crude Oil **(b)** Natural Gas

Figure 2: CVOL and RV of energy products

From Figure 2 we can conclude that for energy commodities the same results hold as for the agricultural commodities. For crude oil we observe a very high peak early 2020, and as a consequence the very large kurtosis of crude oil compared with the other commodities. This is probably because of the corona-crisis and the fact that oil prices went negative. Although crude oil seems mean-reverting, the mean of natural gas seems to have changed at the end of the time period, after 2020, as it seems to revert back to another level then before. The Russo-Ukrainian War is likely to be the cause of this.

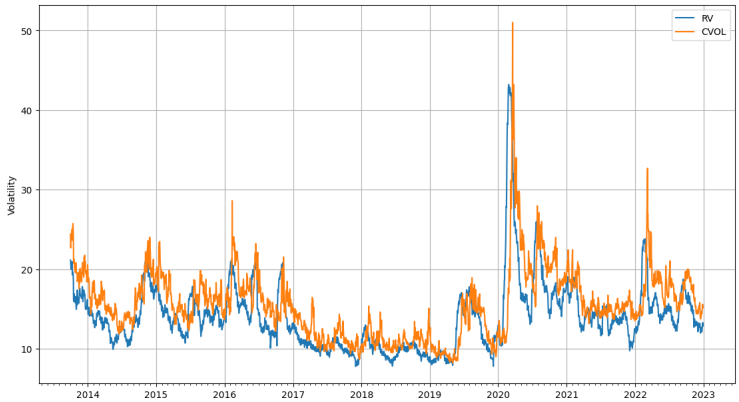

Figure 3: CVOL and RV of Gold

Lastly, Figure 3 shows the RV and CVOL of gold. Also gold incorporates characteristics like the other commodities, namely, mean-reverting, IV is generally higher than RV and peaks early 2020 due to the corona crisis. However, the gold RV and CVOL do seem less volatile than the other commodities, next to having its peaks from time to time.

All in all, reviewing the graphs we again observe that the CVOL is on average higher than the RV for the various commodities. As we know, the CVOL is in itself the market's

prediction of the volatility for the upcoming 30 days. We notice that the CVOL does move according to the RV and captures its various dynamics, although making some errors from time to time. By utilizing the RV as a proxy for the volatility that is predicted by the CVOL, we can assess its performance. In further results we will utilize the CVOL as a noisy prediction variable of the volatility to estimate the latent volatility, namely the RV, by making use of the various filtering methods in combination with the HM.

4 Filtering Methods

Relations between the observed and latent variables can be represented in a state-space form, which is a system of equations describing how the observed variables are related to the latent variables and potentially to exogenous variables. Filtering methods are procedures that utilize various techniques to extract information from these observed variables to estimate the latent variables. The methods focus on two central steps. Namely, the prediction and update step. In the prediction step, the filter predicts the latent variable for the current period based on the previous estimate. In the update step, the filter combines the predicted estimate with the current observed variable to get an estimate for the current latent variable. This procedure is done iteratively for each period.

In this section we present the filtering methods that we make use of. We first discuss the KF, since some of the methods we utilize are based on this method. We cannot use the KF in our analysis, as it is not applicable to non-linear systems. The extensions that we research are the EKF and UKF, which are appropriate filtering methods that can handle non-linear systems with Gaussian noise. On top of the the Kalman-like filters, we also review the PF, another filtering method that is applicable to non-linear systems but considers non-Gaussian noise. Furthermore, the system of equations of the latent and observed variables consist of unknown parameters. To estimate these parameters we will explain and apply the quasi-maximum likelihood estimation (QMLE) method.

4.1 Kalman Filter

The state-space representation consists of two sort of equations. Specifically, these are the observations equation(s) and the state equation(s). In case of the KF, these equations are linear of form and their corresponding noises are Gaussian. This system of equations

can be presented as follows

$$y_t = H\xi_t + B^o X_t^o + w_t, \tag{7}$$

$$\xi_t = F\xi_{t-1} + B^s X_t^s + v_t. \tag{8}$$

Here, equation 7 is the observation equation, with y_t the observed variable, ξ_t the latent variable, also called state variable, X_t^o the optional exogenous variables, and $w_t \sim N(0, R)$ the Gaussian observation noise. Equation 8 is the state equation, with its corresponding optional exogenous variables X_t^s, and $v_t \sim N(0, Q)$ the state Gaussian noise. H describes the link between the state and observed variable. F describes how the state variable evolves over time. B^o and B^s are the parameters of their corresponding exogenous variables. The subscripts on each variable correspond to their time periods.

The goal is to calculate the estimate of the latent variable ξ_t, the RV in our case, given the observations y_t, the CVOL in our case. We move forward using the two crucial steps in Kalman filtering: prediction and update. Prediction of the state variable is estimated from the previously determined states $\hat{\xi}_{t-1}$ in the prediction stage. We get

$$\hat{\xi}_{t|t-1} = E_{t-1}[F\xi_{t-1} + v_t] = F\hat{\xi}_{t-1|t-1} + B^s X_t^s, \tag{9}$$

where $t|t-1$ implies that the previous state estimate is utilized to calculate the current state estimate. So, in our case $\hat{\xi}_{t-1|t-1}$ is the filtered CVOL, with which we try to predict and estimate the RV.

Moreover, we compute the estimation error by

$$e_{t|t-1} = \xi_t - \hat{\xi}_{t|t-1}, \tag{10}$$

with its corresponding covariance being equal to

$$P_{t|t-1} = E_{t-1}[e_{t|t-1}e'_{t|t-1}] = FP_{t-1|t-1}F' + Q. \tag{11}$$

Furthermore, we calculate the observation prediction, for our research the prediction of the CVOL, which is provided by

$$\hat{y}_{t|t-1} = E_{t-1}[H\xi_t + w_t] = H\hat{\xi}_{t|t-1} + B^o X_t^o. \tag{12}$$

Next we follow the KF in the update step. In this step information from the actual observation y_t, predicted observation $\hat{y}_{t|t-1}$, and predicted current state estimate $\hat{\xi}_{t|t-1}$ are

utilized to update the current state estimate and get $\hat{\xi}_{t|t}$, the filtered CVOL of the new time period. This is depicted in the following equation

$$\hat{\xi}_{t|t} = \hat{\xi}_{t|t-1} + K_t(y_t - \hat{y}_{t|t-1}), \tag{13}$$

where K_t is the so-called Kalman gain, a key component in the Kalman-like filters. The estimation error of the updated current state estimate is

$$e_{t|t} = \xi_t - \hat{\xi}_{t|t}. \tag{14}$$

The covariance of $e_{t|t}$ is noted as

$$P_{t|t} = E_t[e_{t|t}e'_{t|t}] = P_{t|t-1} - K_t H' P_{t|t-1}. \tag{15}$$

All the variables are known, except for K_t. K_t can be seen as a weighting factor. The optimal Kalman gain equals

$$K_t = P_{t|t-1}H(H'P_{t|t-1}H + R)^{-1}. \tag{16}$$

For the exact derivations of the equations and further explanation on the KF, refer to Durbin and Koopman (2012). As this is an iterative procedure, the values of the latent variable ξ_0 and its corresponding error covariance matrix P_0 need to be initialized.

4.2 Extended Kalman Filter

The EKF is a filtering technique that is based on the KF, which is applicable to non-linear state-space models. We utilize the methodology of Wan and van der Merwe (2001) to discuss the EKF. Instead of having a linear system of equations, like in 7 and 8, we assume a general function form for the state and observation equations:

$$y_t = h(\xi_t, w_t), \tag{17}$$

$$\xi_t = f(\xi_{t-1}, v_t), \tag{18}$$

where h_t and f_t are non-linear functions of the latent variable ξ_t and Gaussian noise variables w_t and v_t, respectively. We again assume that $w_t \sim N(0, R)$ and that $v_t \sim N(0, Q)$.

The EKF is also split into the prediction and update step. The prediction step is the same as in KF. Namely, the latent variable is predicted as

$$\hat{\xi}_{t|t-1} = E_{t-1}[f(\xi_{t-1}, w_t)] \approx f(\hat{\xi}_{t-1|t-1}, 0). \tag{19}$$

And again let the estimation error be

$$e_{t|t-1} = \xi_t - \hat{\xi}_{t|t-1}. \tag{20}$$

However, the calculation of the covariance is slightly different. The EKF makes use of linearization of the non-linear observation and state transition equations. The Jacobian matrices are utilized to perform this linearization. The Jacobians are:

$$H_t = \frac{\partial h}{\partial \xi_t},\ F_t = \frac{\partial f}{\partial \xi_t},\ W_t = \frac{\partial h}{\partial w_t},\ V_t = \frac{\partial f}{\partial v_t}. \tag{21}$$

In this way we get for the covariance of the predicted state the following

$$P_{t|t-1} = E_{t-1}[e_{t|t-1}e'_{t|t-1}] = F_t P_{t-1|t-1} F'_t + W_t Q W'_t. \tag{22}$$

The observation prediction is provided by

$$\hat{y}_{t|t-1} = E_{t-1}[h(\xi_t, w_t)] \approx h(\hat{\xi}_{t|t-1}, 0), \tag{23}$$

with its covariance being equal to

$$P^{\hat{y}}_{t|t-1} = H'_t P_{t|t-1} H_t + V_t R V'_t. \tag{24}$$

The update step of the EKF looks like the update step of the KF, however the calculations of the covariance and Kalman gain matrices are different due to the linearization of the non-linear state-space model. The updated current state estimate is again noted as

$$\hat{\xi}_{t|t} = \hat{\xi}_{t|t-1} + K_t(y_t - \hat{y}_{t|t-1}), \tag{25}$$

but the optimal Kalman gain is now computed as

$$K_t = P_{t|t-1} H_t (H'_t P_{t|t-1} H_t + V_t R V'_t)^{-1}. \tag{26}$$

Estimation error of the updated current state estimate is again

$$e_{t|t} = \xi_t - \hat{\xi}_{t|t}, \tag{27}$$

with its covariance equal to

$$P_{t|t} = E_t[e_{t|t}e'_{t|t}] = P_{t|t-1} - K_t H'_t P_{t|t-1}. \tag{28}$$

And again, just like in the KF, initialization of ξ_0 and P_0 is required, as it is a recursive process.

4.3 Unscented Kalman Filter

A filtering technique called the UKF was proposed by Julier and Uhlmann (1997). They claimed that the UKF outperforms the extended Kalman filter for estimating highly non-linear systems with Gaussian distributions. This is due to the EKFs inferior state covariance approximation compared to the UKFs.

In contrast to the EKF, the UKF employs actual non-linear models and approximates the distribution of the latent variable rather than approximating the non-linear process and observation models. A collection of well selected deterministic sample points, called sigma points, can be used to estimate the distribution of the states. There are two weights assigned to each sigma point. The real mean and covariance of the states are entirely represented by the sigma points. We demonstrate how to create the sigma points and their related weights below. We make use of the methodology of Wan and van der Merwe (2001) to explain the steps taken in the UKF.

First we introduce the unscented transformation (UT). A technique for figuring out the statistics of a random variable x (with dimension L) that has undergone a non-linear transformation is known as the UT. Assume a random variable x is propagated through a non-linear function $y = g(x)$. Suppose x has a mean of $\bar{x}$ and a covariance of P_x. We create a matrix χ from $2L + 1$ sigma vectors χ_i (with corresponding weights W_i) to determine the statistics of y as follows:

$$\chi_0 = \bar{x}, \tag{29}$$

$$\chi_i = \bar{x} + (\sqrt{(L+\lambda)P_x})_i \qquad i = 1, ..., L, \tag{30}$$

$$\chi_i = \bar{x} - (\sqrt{(L+\lambda)P_x})_i \qquad i = L+1, ..., 2L, \tag{31}$$

$$W_0^{(m)} = \lambda/(L+\lambda), \tag{32}$$

$$W_0^{(c)} = \lambda/(L+\lambda) + (1 - \alpha^2 + \beta), \tag{33}$$

$$W_i^{(m)} = W_i^{(c)} = 1/(2L + 2\lambda) \qquad i = 1, ..., 2L. \tag{34}$$

Here we have that $\lambda = \alpha^2(L + \kappa) - L$ is a scaling parameter. The value of α, which is often set to a small positive number, defines the dispersion of the sigma points around $\bar{x}$. (e.g., 1e-3). β is used to take into account past knowledge of the distribution of x (for Gaussian distributions $\beta = 2$ is optimal), while κ is a supplementary scaling parameter

that is often set to 0. $(\sqrt{(L+\lambda)P_x})_i$ is the i^{th} row of the matrix square root. These sigma vectors are propagated through the non-linear function,

$$\Upsilon_i = g(\chi_i) \qquad i = 0, ..., 2L. \tag{35}$$

Furthermore, utilizing a weighted sample mean and covariance of the posterior sigma points, the mean and covariance approximations for y are

$$\bar{y} \approx \sum_{i=0}^{2L} W_i^{(m)} \Upsilon_i, \tag{36}$$

$$P_y \approx \sum_{i=0}^{2L} W_i^{(c)} (\Upsilon_i - \bar{y})(\Upsilon_i - \bar{y})'. \tag{37}$$

The UKF is an iterative process of the UT following equation 13. Instead of ξ the concatenation of the state and noise variables, $\xi_t^a = [\xi_t \ v_t \ w_t]'$, are examined. The sigma point selection process and computation of the corresponding sigma matrix, χ_t^a, along with the prediction and update step of the UKF are depicted below.

First we calculate the sigma points, and construct the sigma matrix

$$\chi^a_{0,t-1|t-1} = \hat{\xi}^a_{t-1|t-1}. \tag{38}$$

$$\chi^a_{i,t-1|t-1} = \hat{\xi}^a_{t-1|t-1} + \left(\sqrt{(L+\lambda)P^a_{t-1|t-1}}\right)_i \qquad i = 1, ..., L. \tag{39}$$

$$\chi^a_{i,t-1|t-1} = \hat{\xi}^a_{t-1|t-1} - \left(\sqrt{(L+\lambda)P^a_{t-1|t-1}}\right)_{i-L} \qquad i = L+1, ..., 2L. \tag{40}$$

where

$$P^a_{t-1|t-1} = \begin{bmatrix} P^{\chi}_{t-1|t-1} & 0 & 0 \\ 0 & P^{v}_{t-1|t-1} & 0 \\ 0 & 0 & P^{w}_{t-1|t-1} \end{bmatrix}, \tag{41}$$

with $P^{\chi}_{t-1|t-1}$, $P^{v}_{t-1|t-1}$ and $P^{w}_{t-1|t-1}$ covariance matrices of the corresponding state and noise variable.

Again, assume a general function form for the state and observation equations, like in equations 17 and 18. Following the construction of the sigma matrix, we have that $\chi^a_{t-1|t-1} = [\chi^{\xi}_{t-1|t-1} \ \chi^{v}_{t-1|t-1} \ \chi^{w}_{t-1|t-1}]'$. Then the equations for the prediction step of the latent variable $\chi^{\xi}_{t|t-1}$, in our case the filtered CVOL, and the prediction of the observation $\hat{y}_{t|t-1}$, in our case the CVOL, are

$$\chi^{\xi}_{t|t-1} = f(\chi^{\xi}_{t-1|t-1}, \chi^{v}_{t-1|t-1}), \tag{42}$$

$$\hat{\xi}_{t|t-1} = \sum_{i=0}^{2L} W_i^{(m)} \chi_{i,t|t-1}^{\xi}, \tag{43}$$

$$P_{t|t-1}^{\xi} = \sum_{i=0}^{2L} W_i^{(c)} [\chi_{i,t|t-1}^{\xi} - \hat{\xi}_{t|t-1}][\chi_{i,t|t-1}^{\xi} - \hat{\xi}_{t|t-1}]', \tag{44}$$

$$\Upsilon_{t|t-1} = h(\chi_{t|t-1}^{\xi}, \chi_{t-1|t-1}^{w}), \tag{45}$$

$$\hat{y}_{t|t-1} = \sum_{i=0}^{2L} W_i^{(m)} \Upsilon_{i,t|t-1}, \tag{46}$$

$$P_{t|t-1}^{\hat{y}} = \sum_{i=0}^{2L} W_i^{(c)} [\Upsilon_{i,t|t-1} - \hat{y}_{t|t-1}][\Upsilon_{i,t|t-1} - \hat{y}_{t|t-1}]', \tag{47}$$

$$P_{t|t-1}^{\xi\hat{y}} = \sum_{i=0}^{2L} W_i^{(c)} [\chi_{i,t|t-1}^{\xi} - \hat{\xi}_{t|t-1}][\Upsilon_{i,t|t-1} - \hat{y}_{t|t-1}]'. \tag{48}$$

And additionally, the update step equations for the filtered CVOL of the new time period $\hat{\xi}_{t|t}$, which is the prediction of the RV in our case, and its corresponding covariance $P_{t|t}$ are

$$K = P_{t|t-1}^{\xi\hat{y}} (P_{t|t-1}^{\hat{y}})^{-1}, \tag{49}$$

$$\hat{\xi}_{t|t} = \hat{\xi}_{t|t-1} + K(y_t - \hat{y}_{t|t-1}), \tag{50}$$

$$P_{t|t} = P_{t|t-1} - K P_{t|t-1}^{\hat{y}} K'. \tag{51}$$

Again, because it is an iterative procedure, initialization of ξ_0 and P_0 is necessary, just like before. In addition for the UKF, initialization of ξ_0^a and P_0^a is required.

4.4 Particle Filter

Two non-linear filtering methods that depend on Gaussian approximation have been demonstrated so far. The PF, also referred to as the Sequential Monte Carlo approach, is another broadly used filtering method, which in particular is useful for models that have non-Gaussian characteristics. The objective of the PF is to recursively approximate the posterior state distribution in order to implement Bayesian estimation. To review the PF, we make use of the methodology of Wan and van der Merwe (2001). The main idea behind the PF is to use Monte Carlo simulation to calculate the posterior probability distribution (PPD). The PPD is approximated using a collection of weighted samples, called particles, taken from a proposed distribution $q(x_{0:t}|y_{1:t})$. Suppose for each time period t we have N

particles $x_t^{(1)}, ..., x_t^{(N)}$ and they are independent and identically distributed (i.i.d). By the law of large numbers it holds that

$$1/N \sum_{i=1}^{N} f_t(x_{0:t}^{(i)}) \rightarrow \int f_t(x_{0:t}) P(dx_{0:t}|y_{1:t}) \quad as\ N \rightarrow \infty. \tag{52}$$

This shows that there is convergence approximating the PPD, utilizing the PF, as long as N is large enough. So, the PF consists of generating particles from a proposed distribution, with corresponding weights, recursively, and approximates the PPD to eventually have an estimate for the updated state variable. This procedure can be split up into three steps, namely, sampling, weight computation, and resampling. The latter is a step to tackle the problem of particle degeneration, which we will discuss in the upcoming paragraphs.

One of the most important considerations for the sampling algorithms is the selection of the proposal function. The proposal function is the probability distribution that generates new particles based on the predicted states and observed variable information. Doucet (1997) recommends choosing proposal functions that reduce the variance of the weights. In Doucet et al. (2001) it has been proved that "The proposal distribution $q(x_t|x_{0:t-1}, y_{1:t}) = p(x_t|x_{0:t-1}, y_{1:t})$ minimizes the variance of the importance weights conditional on $x_{0:t-1}$ and $y_{1:t}$." Other academics have also backed this idea, for example in Liu and Chen (1995), Zaritskii et al. (1976), and Kong et al. (1994). However, the most common option for a proposal function is $q(x_t|x_{0:t-1}, y_{1:t}) \approx p(x_t|x_{t-1})$, called the transition prior, which we will utilize. Because it does not take into account the most recent data, it results in greater variation than the optimal distribution $p(x_t|x_{0:t-1}, y_{1:t})$, but is typically simpler to execute. The model describes the development of the states, and the characteristics of the noise are used to determine the transition prior. After generating the particles with the transition prior, each particle is assigned a weight accordingly.

In the academic field it is found that the variance of the weights increases over time, which causes the simulations to become inaccurate. This is called particle degeneration, and proof for this can be found in Doucet et al. (2001) and Kong et al. (1994). In practice, this degeneration is demonstrated through the weights of some particles becoming close to zero, whilst others becoming really high (close to 1). This causes a set of particles to be irrelevant. To address this issue, resampling methods are used. Resampling involves removing all the particles with insignificant weights and selecting a new group of N equally weighted particles from the particles that are still present. A typical method for assessing

degeneracy is to use the following expression

$$\lambda_t^{deg} = \frac{1}{\sum_{i=1}^{N}(w_t^{(i)})^2}, \tag{53}$$

where $w_t^{(i)}$ is the weight of particle i at time step t. If $\lambda_t^{deg} < N/2$ the particles get resampled and the new set of particles all get the same weight, namely $1/N$.

We will now discuss the equations of the particle filter algorithm, as defined in Wan and van der Merwe (2001). For initialization, so $t = 0$, draw N particles $x_0^{(i)}$ from the prior $p(x_0)$. Because there is limited prior information about the initial state, in our research we initialize the latent variable x_0, namely the filtered CVOL, and utilize a uniform distribution to select the particles $x_0^{(i)}$ around this initialization, to represent the uncertainty across the state space. Furthermore, set the weights equal to $1/N$. Afterwards, in the sampling step, sample N particles $x_t^{(i)}$ from the proposal distribution $q(x_t|x_{0:t-1}, y_{1:t}) \approx p(x_t|x_{t-1})$. As the HM assumes Gaussian process noise, the transition prior for our case is $p(x_t|x_{t-1}) = N(f(\hat{\xi}_{t-1|t-1}, 0), R)$. Following the sampling step, compute the weights corresponding to the particles as

$$w_t^{(i)} = w_{t-1}^{(i)} \frac{p(y_t|x_t^{(i)})p(x_t^{(i)}|x_{t-1}^{(i)})}{q(x_t^{(i)}|x_{0:t-1}^{(i)}, y_{1:t})} \approx w_{t-1}^{(i)} p(y_t|x_t^{(i)}), \tag{54}$$

and normalize them according to

$$\tilde{w}_t^{(i)} = \frac{w_t^{(i)}}{\sum_{j=1}^{N} w_t^{(j)}}. \tag{55}$$

Furthermore, in the resampling step, if $\lambda_t^{deg} < N/2$, we sample N particles from the existing particles according to their normalized weights, which can be utilized as probabilities for choosing a particle. The N particles are equally weighted. And last of all, estimate the state variable, namely the updated filtered CVOL, which is the prediction of the RV, using

$$\hat{\xi}_{t|t} = \sum_{i=1}^{N} \tilde{w}_t^{(i)} x_t^{(i)}. \tag{56}$$

5 Heston Model

In this section we discuss the HM, with its implications. Furthermore, the state-space representation, with according discretization, is presented, to make the filtering methods applicable. Finally, we show how to adjust the resulting model to make it suitable to our research.

5.1 Original Heston Model

We examine the HM in a risk-neutral probability world with measure $\mathbb{Q}$, which corresponds to the real world probability martingale measure $\mathbb{P}$. The original HM is given by the following dynamic system

$$dS(t) = rS(t)dt + \sqrt{V(t)}S(t)dW_1(t), \tag{57}$$

$$dV(t) = \kappa(\theta - V(t))dt + \sigma\sqrt{V(t)}dW_2(t), \tag{58}$$

where we have that $S(t)$ is the equity price, and $V(t)$ is the variance of the equity at time t. The average yield of the equity is denoted by r, which in the risk-neutral space is the risk-free rate. The $\kappa(\theta - V(t))$ term is called the mean-reverting force, in which $\theta > 0$ is the long-run mean of $V(t)$, and $\kappa > 0$ is the long-run mean reversion rate of the variance. Also, we have $\sigma \geq 0$, which is the volatility of the variance. These parameters are assumed to be positive and constant. If it holds that $2\kappa\theta > \sigma^2$, the variance is always positive. Last of all, we have the two Brownian motions $W_1(t)$ and $W_2(t)$, the increments of which are correlated by a constant correlation coefficient $\rho \in [-1, 1]$.

In practice the equity returns are more of interest than the equity prices. So, instead of using $dS(t)$, $d\log(S(t))$ is regarded. By Itô equation it is stated that

$$d\log(S(t)) = \frac{1}{S(t)}dS(t) - \frac{1}{2}\frac{1}{S(t)^2}dS(t)dS(t). \tag{59}$$

By filling in 57 into 59, $d\log(S(t))$ is given by

$$d\log(S(t)) = \left(r - \frac{1}{2}V(t)\right)dt + \sqrt{V(t)}dW_1(t). \tag{60}$$

5.2 State-Space Representation and Discretization

Before we present the state-space representation of the HM model, we first transform the noise variables, which are the Brownian motions. This is necessary since the Brownian motions need to be independent. According to Javaheri (2011), the optimal method to do this is to subtract from the variance process (58) a multiple of the equity process (59) minus the log-equity $d\log(S(t))$, that equals zero. This is done by setting $W_1(t) = \sqrt{1-\rho^2}W_s(t) + \rho W_v(t)$ and $W_2(t) = W_v(t)$, where $W_s(t)$ and $W_v(t)$ are independent. Consequently, we can write the HM as a state-space representation in the following way

$$dV(t) = \kappa(\theta - V(t))dt + \sigma\sqrt{V(t)}dW_v(t), \tag{61}$$

$$d\log(S(t)) = \left(r - \frac{1}{2}V(t)\right)dt + \sqrt{(1-\rho^2)V(t)}dW_s(t) + \rho\sqrt{V(t)}dW_v(t). \tag{62}$$

As a result, the state-space representation of the HM can be discretized as follows

$$V_t = V_{t-1} + \kappa(\theta - V_{t-1})dt + \sigma\sqrt{V_{t-1}}\Delta W_{V_t}, \tag{63}$$

$$\log(S_t) = \log(S_{t-1}) + \left(r - \frac{1}{2}V_t\right)\Delta t + \sqrt{(1-\rho^2)V_t}\Delta W_{s_t} + \rho\sqrt{V_t}\Delta W_{V_t}. \tag{64}$$

5.3 Adjusted Heston Model

To apply the HM for our research, we need to adjust it. As a result of utilizing the CVOL, as the observation variable with noise, to estimate the RV variable, the system of equations changes. As we are now only interested in the volatility equation of the HM model, namely 63, and its characteristics to describe the volatility dynamics. The system of equations we utilize becomes

$$\xi_t = \xi_{t-1} + \kappa(\theta - \xi_{t-1})dt + \delta RV_{t-1} + \gamma\sqrt{\xi_{t-1}}\eta_t, \tag{65}$$

$$y_t = \alpha + \beta\xi_t + \phi\sqrt{\xi_{t-1}}\varepsilon_t, \tag{66}$$

where we have that y_t is the observed CVOL estimate, containing noise, and ξ_t is the state variable, namely the latent "true" volatility measure, which we also refer to as the filtered CVOL, so a forecast of the RV. On top of that, RV_{t-1} is the observed RV from the previous time period, which might contain information to better estimate the filtered CVOL. Again, we assume that $\eta_t \sim N(0, R)$ and that $\varepsilon_t \sim N(0, Q)$. We have that $\kappa, \theta, \delta, \gamma, \alpha, \beta, \phi > 0$ are the parameters of the model, and that $\kappa \leq 1$. We included α and β in equation 66, because when regressing the CVOL on the RV we observe values that are significantly different from zero and one, on a 5% level (see Appendix B), for these parameters. From this we can conclude that we cannot simple set α equal to zero, and β equal to one, as in that way the model loses the flexibility it requires, and as a consequence a loss of information.

5.4 Parameter Estimation

In this section we review the estimation procedure we utilize for the parameters of our adjusted HM. For our research we make use of the so-called QMLE procedure. The QMLE technique provides a robust and efficient framework for parameter estimation by

maximizing a pseudo-likelihood function that closely approximates the true likelihood. The pseudo-likelihood function is derived from the observed variables, which in our case is the CVOL, and the system of equation we use to describe its characteristics.

Let Ω be the set of all unknown parameters, which in case of the adjusted HM are $\kappa, \theta, \delta, \gamma, \alpha, \beta, \phi, Q, R$. Moreover, let y_t be the observed CVOL estimate. Then, the likelihood function in case of the EKF and UKF is

$$p(y_1, ..., y_T; \Omega) = \prod_{t=1}^{T} \frac{1}{\left(2\pi P_{t|t-1}^{\hat{y}}\right)^{1/2}} \exp\left(-\frac{1}{2}\frac{y_t - \hat{y}_{t|t-1}}{P_{t|t-1}^{\hat{y}}}\right), \tag{67}$$

where, for the EKF, $\hat{y}_{t|t-1}$ and $P_{t|t-1}^{\hat{y}}$ are calculated by equations 23 and 24, respectively, and for the UKF by equations 46 and 47, respectively. T is equal to the amount of observations of the CVOL. For numerical reasons, instead of maximizing the likelihood we minimize the negative log likelihood, which equals

$$-LogL = \sum_{t=1}^{T} \frac{1}{2}\log(2\pi) + \frac{1}{2}\log\left(P_{t|t-1}^{\hat{y}}\right) + \left(\frac{1}{2}\frac{y_t - \hat{y}_{t|t-1}}{P_{t|t-1}^{\hat{y}}}\right). \tag{68}$$

As for the EKF and UKF, we can use QMLE in order to estimate the parameter-set Ω. Due to the fact that the PF does not make any assumptions about the noise distribution, the likelihood function that needs to be maximized takes on a more general form compared to the ones used in the UKF and EKF. The likelihood at a time point t equals

$$L_t = p(y_t|y_{1:t-1}; \Omega) = \int p(z_t|x_t)p(x_t|z_{1:t-1})dx_t, \tag{69}$$

then the negative log likelihood to be minimized becomes

$$-LogL = -\sum_{t=1}^{T} \log(L_t). \tag{70}$$

The likelihood can be written as

$$L_t = \int p(z_t|x_t)\frac{p(x_t|z_{1:t-1})}{q(x_t|x_{t-1}, z_{1:t})}q(x_t|x_{t-1}, z_{1:t})dx_t, \tag{71}$$

and using that by construction the distribution of the $x_t^{(i)}$'s are conform to $q()$, the likelihood can be estimated by

$$\hat{L}_t = \sum_{i=1}^{N} w_t^{(i)}, \tag{72}$$

where N is the number of particles. For further explanation on the parameter estimation in case of the PF, we refer to Javaheri et al. (2003).

6 Results

6.1 Filtering Models

In this section, we present the results of our analysis on the performance of the EKF, UKF and the PF filtered CVOL measures to the CVOL measure of the various commodities. Our evaluation includes the examination of graphs, formal checks to examine significance in difference between model forecasts, and the calculation of robust loss metrics, providing a comprehensive understanding of the accuracy of these volatility estimation techniques. We make use of the Diebold-Mariano (DM) test, from Diebold and Mariano (2002), to check whether the forecasts of the various models significantly differ. To correct for possible heteroskedasticity and autocorrelation, we incorporate heteroskedasticity-consistent (HAC) standard errors, also called Newey-West, in our analysis. We utilize the robust loss functions of Patton (2011) as evaluation metrics, as these take on functional forms in which requirements that must be met, for the loss function, are included. For further details we refer to Patton (2011). The optimized parameters can be found in Appendix C.

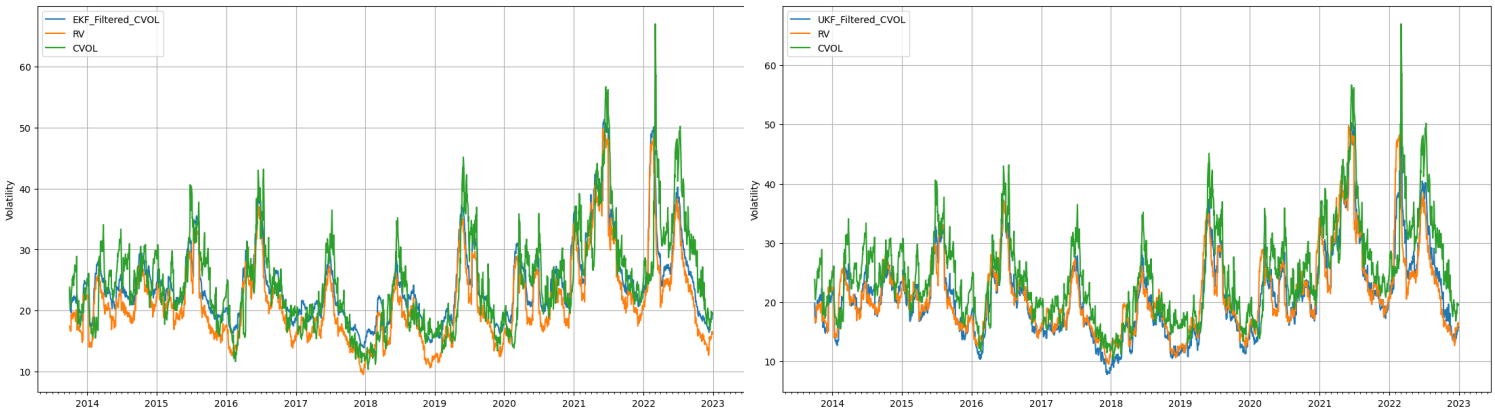

(a) EKF Filtered CVOL Corn **(b)** UKF Filtered CVOL Corn

Figure 4: Kalman Filtered CVOL Corn

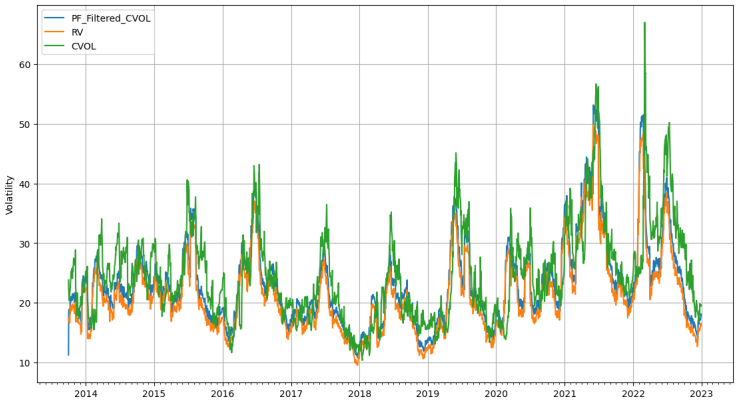

Figure 5: PF Filtered CVOL Corn

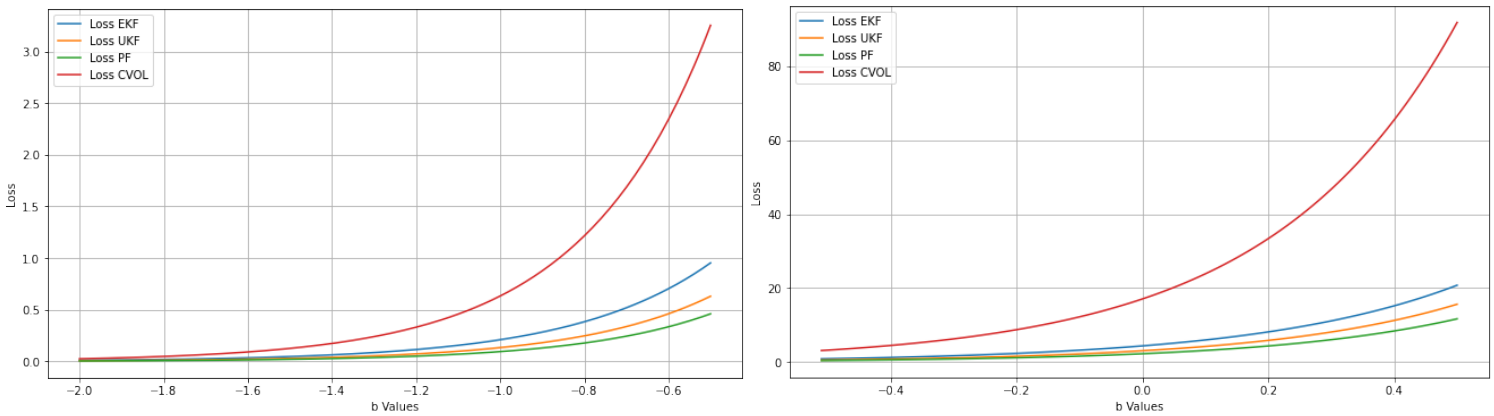

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 6: Robust Loss of Filtering Methods Corn

Table 3: Robust loss of (filtered) CVOL for Corn

	CVOL	**EKF**	**UKF**	**PF**
MSE	34.18*	8.82*	6.18*	4.60
QLike	0.013*	0.006*	0.004*	0.002
b-Value(0.5)	45.89*	10.39*	7.83*	5.87

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

For the KF methods we observe in Figure 4 a consistent pattern where both the EKF and UKF measures exhibit a significantly closer alignment with the RV compared to the CVOL measures for corn. This finding indicates that the EKF and UKF models are capable of capturing the underlying volatility dynamics more accurately, leading to improved estimation results.

Moreover, the robust loss calculations, in Figure 6, further confirm the superior performance of the EKF and UKF measures compared to the CVOL measures. The loss for both EKF and UKF measures are consistently lower, indicating a smaller deviation from the RV. In Figure 6 the loss for various b-Values ranging from -2.0 to 0.5 are depicted. Lower b-Values are corresponding to the loss function penalizing under-prediction more, and higher b-Values to penalizing over-prediction more. In Table 3 special cases are highlighted, namely the MSE ($b = 0$), Qlike ($b = -2.0$), and a case in which $b = 0.5$. So, in this way we cover three cases, when the penalization is equal for over- and under-prediction, the MSE, when under-prediction is penalized, the QLike, and when over-prediction is penalized, b-Value(0.5). This signifies that the EKF and UKF models provide more precise estimates of corn volatility, enhancing their reliability as volatility estimation tools. On top of that, we observe that both the EKF and UKF still overestimate the RV from time to time. This can be due to the fact that the filtered CVOL utilizes the CVOL, which overestimates the RV, to gain information about the RV. As a consequence, the optimized α, the level parameter between the CVOL and filtered CVOL, is not high enough, and the optimized κ, the mean-reverting parameter of the filtered CVOL, is not high enough to revert as fast as the RV. These values are stated in Appendix C.

When comparing the performance of the EKF and UKF measures, we observe that the UKF consistently outperforms the EKF. The UKF demonstrates a more accurate estimation of corn volatility, evident from both the figures and the lower MSE, Qlike and b-Value(0.5) values. This is probably due to the fact that the UKF captures the nonlinearity better than the EKF. This suggests that the UKF is better equipped to capture the complex volatility dynamics inherent in the corn market, surpassing the performance of the EKF.

When inspecting Figure 5, it is evident that the PF measures exhibit an exceptionally close alignment with the RV, surpassing both the EKF and UKF measures. This indicates that the PF model effectively captures the intricate volatility dynamics of the underlying asset, resulting in highly accurate volatility estimation for the given dataset of corn. As in the case of the EKF and the UKF, we note that the PF also overestimates the RV occasionally. Furthermore, in case of the PF the consequence of the initialisation is apparent, which was less observable for the UKF and EKF. This sometimes leads to the first filtered CVOL value to deviate a great amount from either the RV or CVOL.

Additionally, the calculation of MSE, Qlike and b-Value(0.5) values further confirm the stronger performance of the PF compared to the EKF and UKF measures. The PF consistently displays lower loss values, signifying a smaller deviation from the RV. This probably follows from the fact that the PF also allows for non-Gaussianity, and because the volatility process follows a noncentral chi-square distribution (Cox et al., 1985), the PF better represents the distribution of the underlying volatility compared to the EKF and UKF. For the corn market, all forecasts of the filtering methods and the CVOL, based on the three special cases of robust losses, differ significantly from each other. The t-statistics can be found in Appendix D.

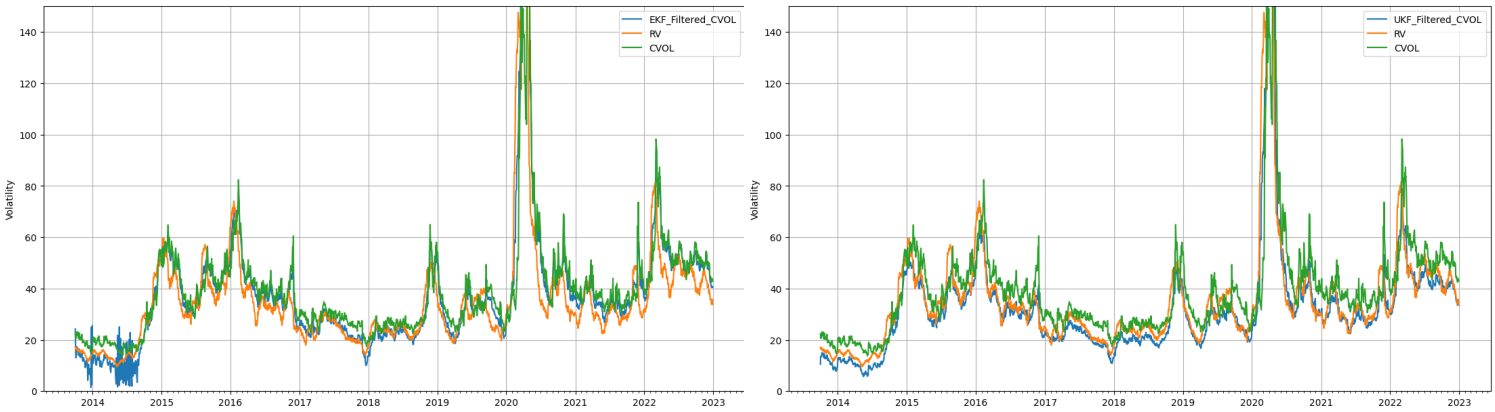

(a) EKF Filtered CVOL Crude Oil **(b)** UKF Filtered CVOL Crude Oil

Figure 7: Kalman Filtered CVOL Crude Oil

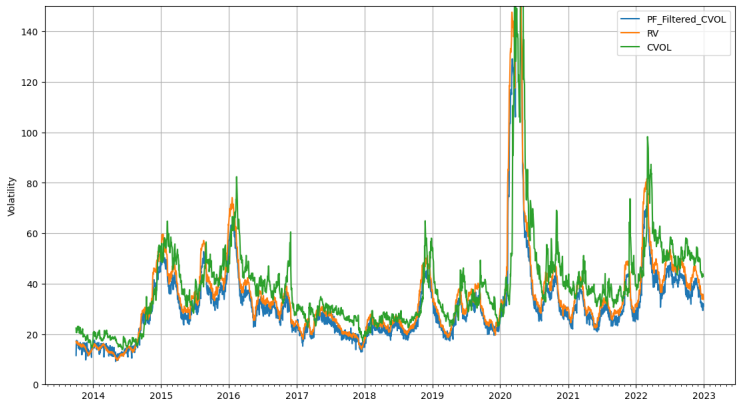

Figure 8: PF Filtered CVOL Crude Oil

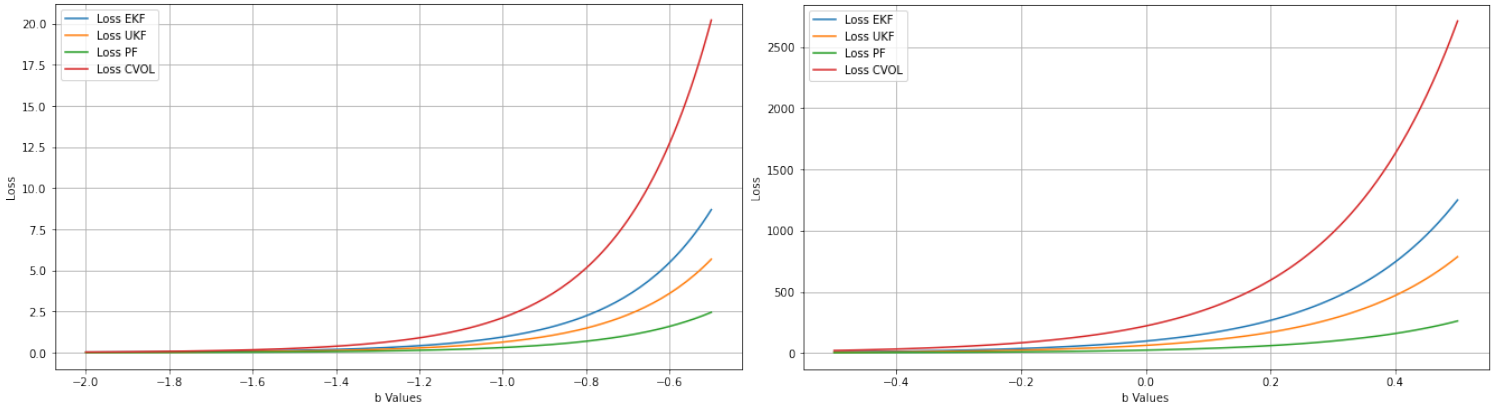

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 9: Robust Loss of Filtering Methods Crude Oil

Table 4: Robust loss of (filtered) CVOL for Crude Oil

	CVOL	**EKF**	**UKF**	**PF**
MSE	442.16*	194.62*	124.52*	46.52
QLike	0.037*	0.027*	0.018*	0.008
b-Value(0.5)	2713.14*	1250.12	786.44*	262.35

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

In analyzing crude oil volatility, both the EKF and the UKF measures consistently outperform the CVOL measures, for both under- and over-prediction, similar to our findings for corn, looking at the graphs in Figure 7 and 9, and the MSE, QLike and b-Value(0.5) in Table 4. Although, when under-prediction is penalized (QLike case), the loss between the EKF and CVOL, following the forecasts, seem to not significantly differ. However, when under- and over-prediction are equally penalized (MSE case), and over-prediction is penalized (b-Value(0.5) case), the EKF forecasts significantly outperform the CVOL. The UKF outperforms the CVOL significantly in all three cases, following the values in Appendix D. So, the EKF and UKF models demonstrate closer alignment with the RV, indicating improved estimation accuracy. We do observe in case of the EKF that in the early measures the filtered CVOL is volatile. This is probably due to the fact that the variance of the CVOL and RV (Table 1 and 2) are much higher than the other commodities, which leads to the estimation error covariance of the updated current state estimate to be high in early observations, and decreasing as more observations come in. Moreover,

in contrast to the UKF filtered CVOL of corn, the UKF filtered CVOL of crude oil under-estimates the RV most of the time. We can trace this back to a combination of the lagged RV parameter and long-run mean parameter, δ and θ respectively, being low, compared to those of other commodities. Thus, the level of the filtered CVOL is lower than when the δ and θ are higher. The values are depicted in Appendix C.

Similar to our observations for corn, in the analysis of crude oil volatility estimation, the PF filtering method appears to be the top performer. The PF consistently exhibits the closest alignment with the RV, surpassing the performance of both the EKF and the UKF measures, which is apparent when comparing the robust loss values in Table 4, and when comparing the graphs in Figures 7 and 8. A noticeable fact is that the loss of EKF forecasts does not differ significantly from the loss of the PF and UKF forecasts, when there is a penalization on over-prediction. This can be due to the fact that the EKF under-estimates the RV heavily in the beginning, and after that period starts to moderately over-estimate the RV. Furthermore, similar to the case of the UKF, the PF measures also seem to underestimate the RV. The PF does seem more volatile over the given time frame compared to previous measures, for example the PF estimates of corn and UKF estimates of crude oil. This can be accounted to the γ parameter being relatively high for a PF method, because we observe lower γs for the other PF filtered commodities, except for soybean.

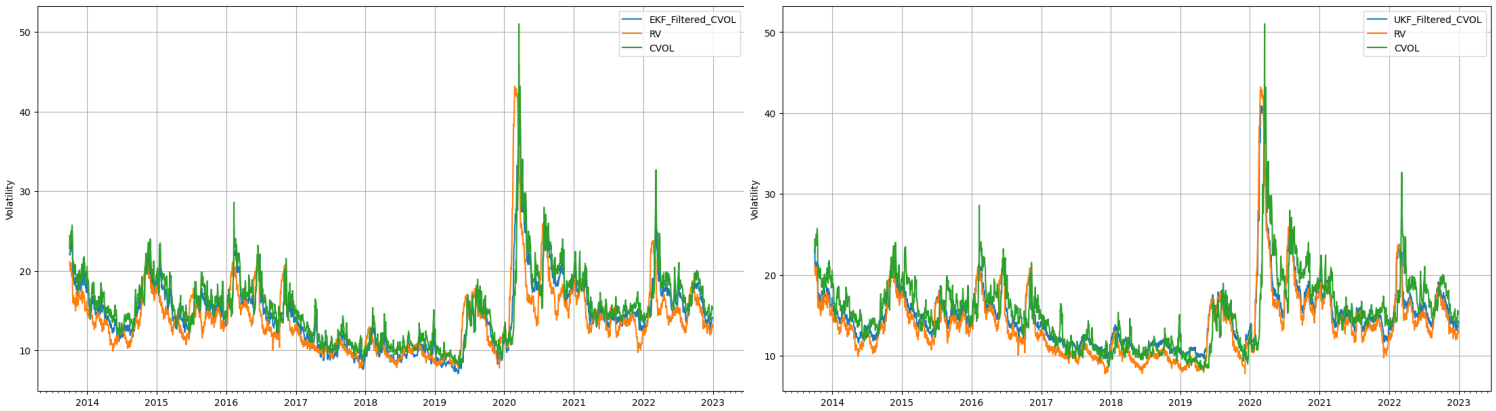

(a) EKF Filtered CVOL Gold **(b)** UKF Filtered CVOL Gold

Figure 10: Kalman Filtered CVOL Gold

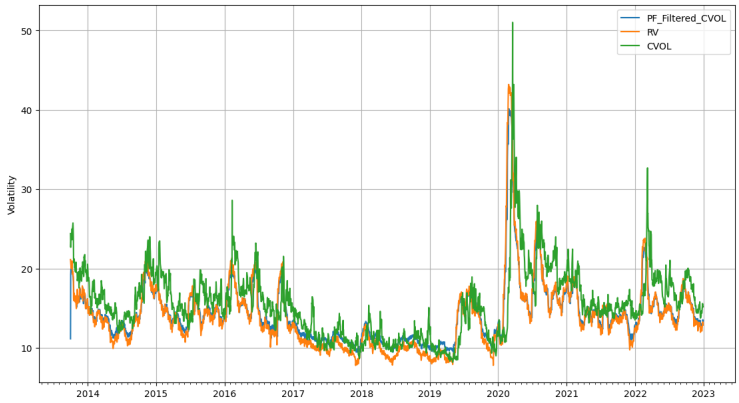

Figure 11: PF Filtered CVOL Gold

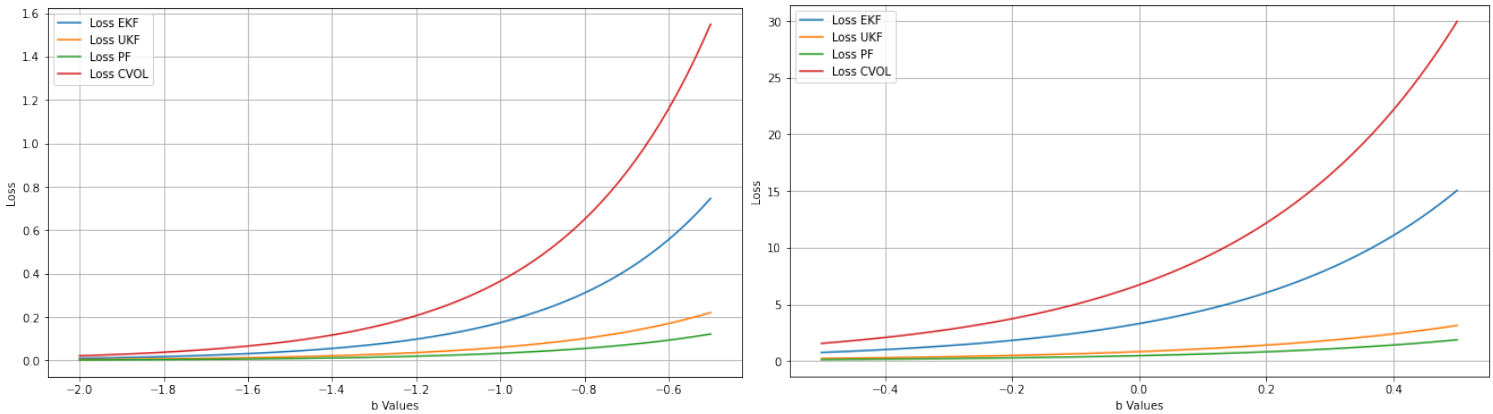

(a) b-Value ranging from -2.0 to -0.5

(b) b-Value ranging from -0.5 to 0.5

Figure 12: Robust Loss of Filtering Methods Gold

Table 5: Robust loss of (filtered) CVOL for Gold

	CVOL	**EKF**	**UKF**	**PF**
MSE	13.44*	6.60*	1.64*	0.93
QLike	0.022*	0.011*	0.005*	0.003
b-Value(0.5)	29.97*	15.05*	3.13*	1.87

We again observe that the PF filtered CVOL of gold performs best compared to the CVOL, EKF filtered, and UKF filtered CVOL according to the graphs from Figures 10, 11 and 12, and the loss values from Table 5. In case of gold, the filtered measures seem to again all overestimate the RV, just like for corn. However, in peak volatility time periods, like in early 2020 and 2022, the UKF and PF filtered CVOL of gold underestimates the RV, contrary to the UKF and PF filtered corn estimates.

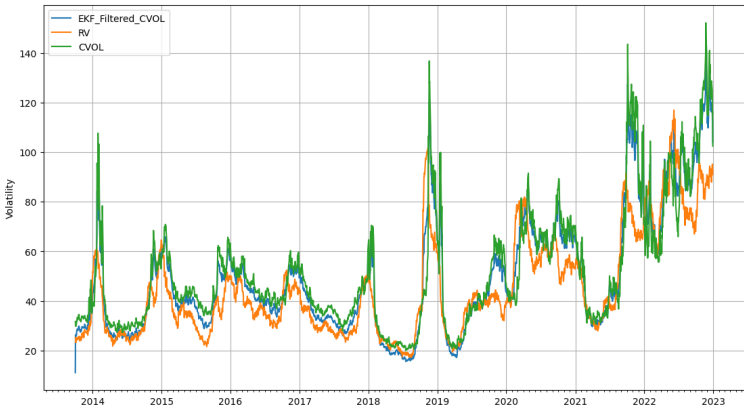

(a) EKF Filtered CVOL Natural Gas

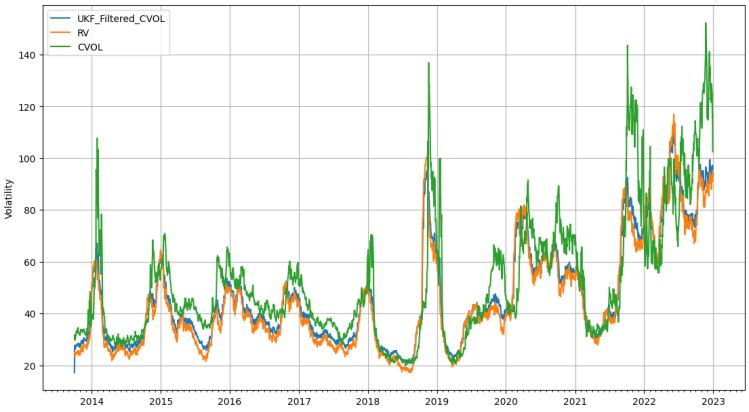

(b) UKF Filtered CVOL Natural Gas

Figure 13: Kalman Filtered CVOL Natural Gas

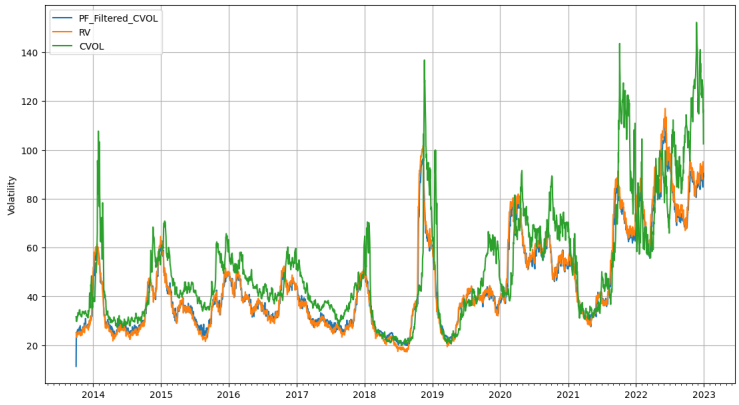

Figure 14: PF Filtered CVOL Natural Gas

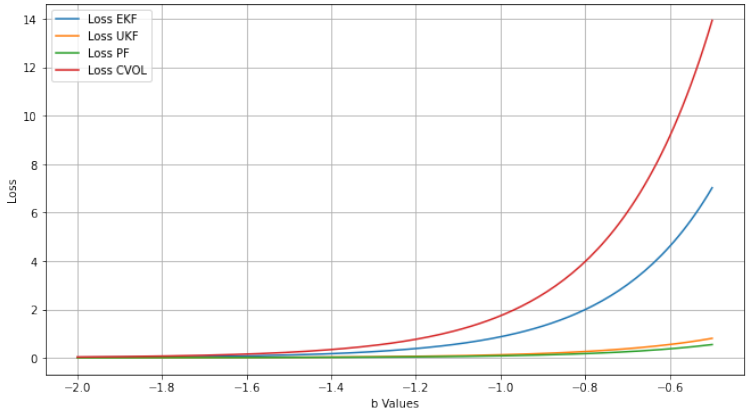

(a) b-Value ranging from -2.0 to -0.5

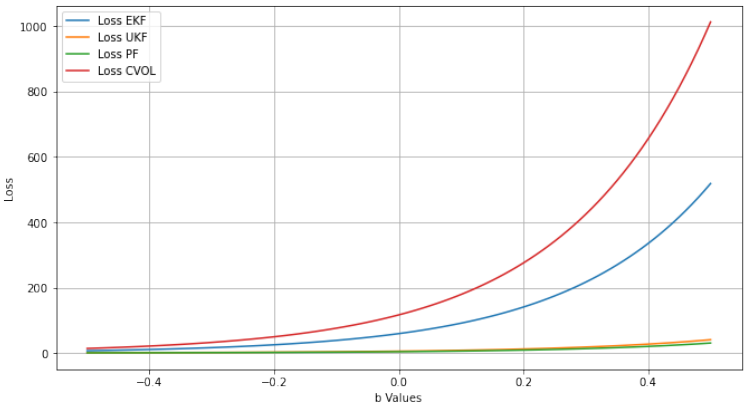

(b) b-Value ranging from -0.5 to 0.5

Figure 15: Robust Loss of Filtering Methods Natural Gas

Table 6: Robust loss of (filtered) CVOL for Natural Gas

	CVOL	**EKF**	**UKF**	**PF**
MSE	233.33*	118.73*	11.15*	7.87
QLike	0.031*	0.015*	0.003*	0.002
b-Value(0.5)	1012.68*	518.36*	40.71*	30.21

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Analysis reaffirms that also for natural gas the EKF, UKF, and PF measures outperform the CVOL estimate. Through both graphical analysis from Figures 13 and 14, and loss calculations in Figure 15 Table 6, these filtering methods exhibit a better alignment with the RV. Among them, the PF performs the best, followed by the UKF, while the EKF demonstrates relatively lower accuracy in RV estimation. All losses following the forecasts of the RV from the CVOL, EKF, UKF and PF, are significantly different from each other, according to the DM-test.

We observe that, although the EKF filtered CVOL of natural gas comes closer to the RV, it more closely follows the CVOL rather than the RV. This stems from the characteristic of the EKF, that becomes evident here, which is unable to accurately represent the non-linear behavior of the underlying latent variable, the RV in this context. This limitation becomes particularly pronounced when the RV of natural gas exhibits highly non-linear dynamics. Consequently, the EKF tends to align more closely with the observed variable, the CVOL in this case. This also is apparent from the MSE, QLike and b-Value(0.5) values of the EKF filtered CVOL, which are relatively higher than their UKF and PF counterparts and closer to the loss values of the CVOL. Thus, most of the time the EKF overestimates the RV, however in very low volatility periods it underestimates the RV. In contrast, the UKF filtered CVOL more closely follows the RV. The UKF also still has positive deviations from the RV. The PF filtered CVOL comes really close to the RV compared to the UKF and EKF, and captures its various characteristics, like the peaks and low volatility periods. This can be attributed to the PF's ability to closely approximate the distribution of the volatility process. On top of that, it also moves according to the RV and captures the dynamics.

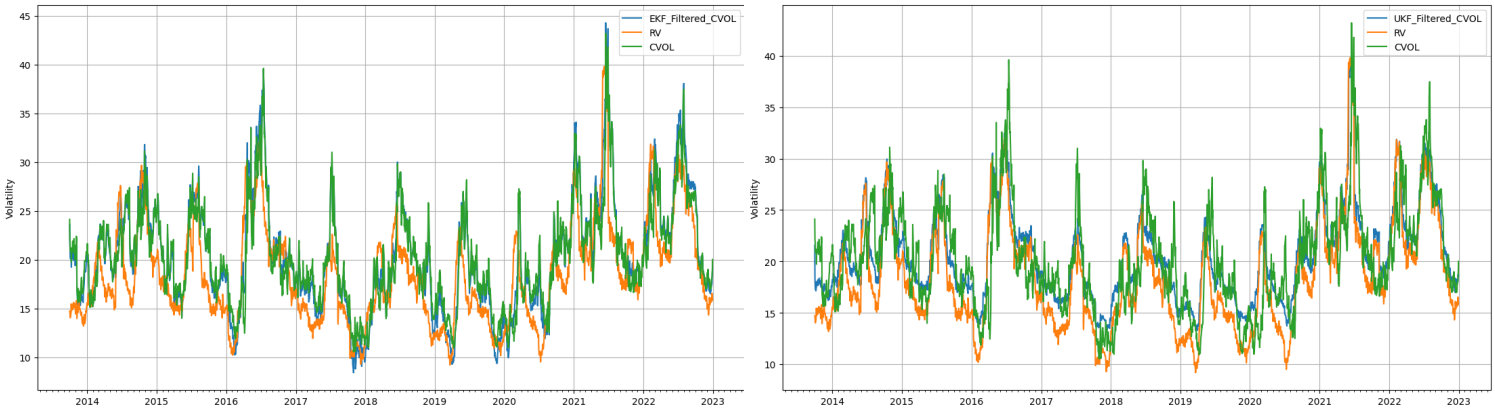

(a) EKF Filtered CVOL Soybean **(b)** UKF Filtered CVOL Soybean

Figure 16: Kalman Filtered CVOL Soybean

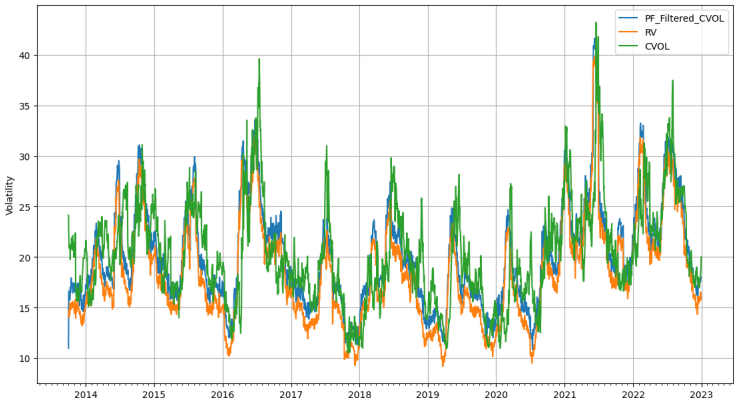

Figure 17: PF Filtered CVOL Soybean

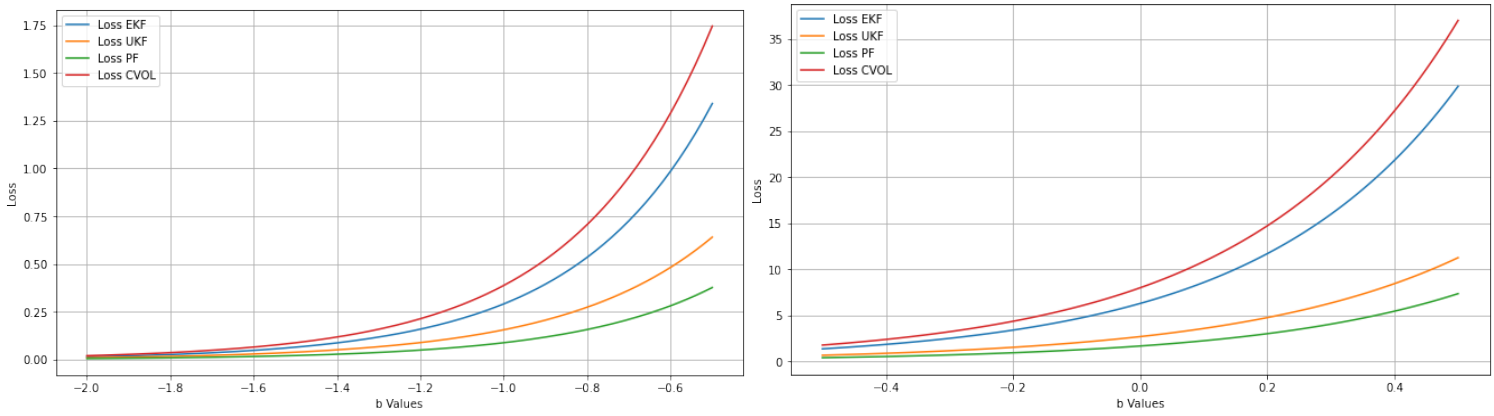

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 18: Robust Loss of Filtering Methods Soybean

Table 7: Robust loss of (filtered) CVOL for Soybean

	CVOL	**EKF**	**UKF**	**PF**
MSE	15.95*	12.54*	5.33*	3.24
QLike	0.020*	0.015*	0.010*	0.005
b-Value(0.5)	37.01*	29.86*	11.24*	7.21

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

For soybean volatility estimation, similar to all the other commodities examined before, the EKF, UKF, and PF methods once again surpass the IV estimate. Careful examination of the graphs in Figures 16 and 17, and the assessment of the loss values in Table 7 and Figure 18 reinforce their better performance. Notably, the PF is the foremost performer among the three methods, followed by the UKF, while the EKF displays comparatively inferior accuracy. And again, the DM-test shows that all the forecasts significantly differ from each other (t-statistics in Appendix D).

In case of the EKF filtered CVOL of soybean we observe that it, just like for natural gas, closely follows the CVOL, such that it overestimates the RV most of the time, and in low volatility periods it underestimates the RV, which is apparent from its graph in Figure 16. As a consequence, the MSE, QLike and b-Value(0.5) are also not greatly improved compared to that of the CVOL. For the UKF filtered CVOL of soybean, we do observe a more noteworthy improvement of the loss values. Although it still overestimates the RV most of the time, the UKF estimates align more with the RV and follows its dynamics better. The same holds for the PF estimates, and shows better alignment with the RV than the UKF estimates. Although again, the PF estimates seem to be more volatile, like the PF estimates of crude oil, than the other filtered CVOL estimates of soybean and the other commodities. This probably again follows from the fact that the PF method has a relatively high γ, for a PF method.

6.2 Benchmark Models

In this section, we present the results of our analysis on the performance between the (PF filtered) CVOL and the benchmark models, by reviewing its in-sample forecasts.

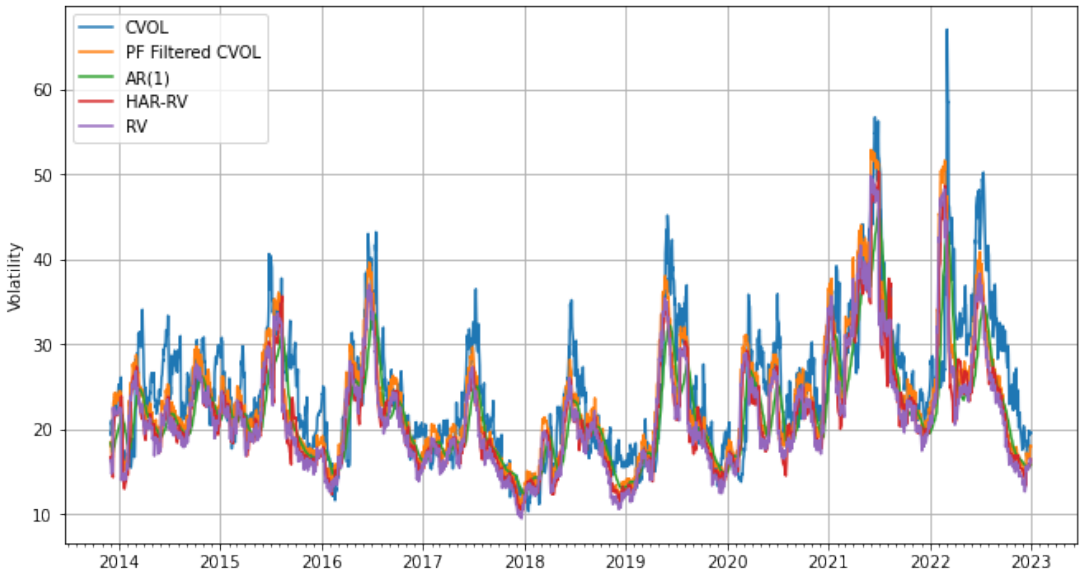

Figure 19: (PF Filtered) CVOL and Benchmark Models Corn

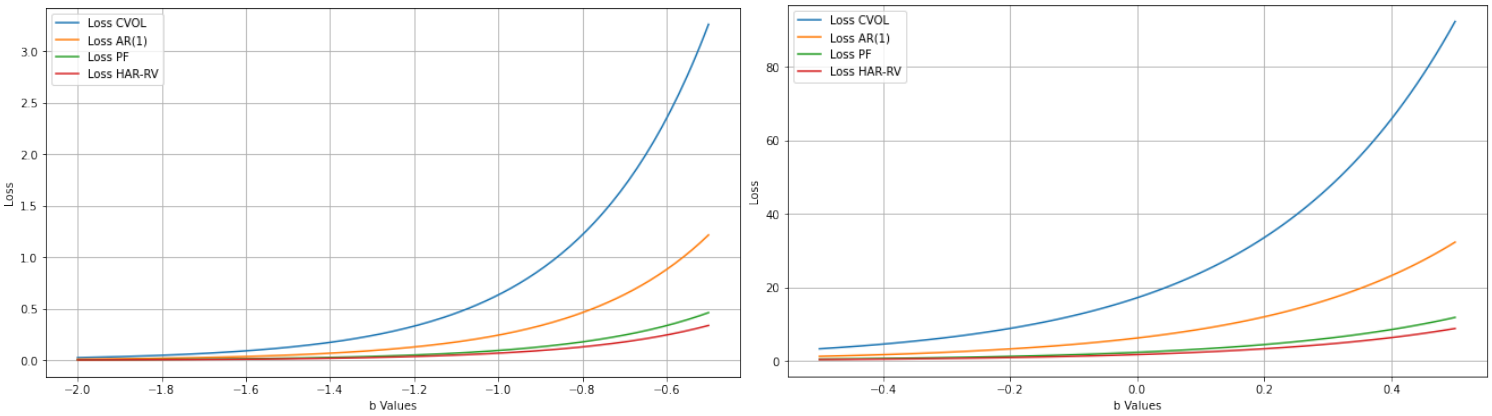

(a) b-Value ranging from -2.0 to -0.5

(b) b-Value ranging from -0.5 to 0.5

Figure 20: Robust Loss of Filtering Methods Corn

Table 8: Robust loss of (filtered) CVOL and Benchmark Models for Corn

	CVOL	**PF**	**HAR-RV**	**AR(1)**
MSE	34.30*	4.61	3.40*	12.40*
QLike	0.026*	0.005	0.003*	0.011*
b-Value(0.5)	92.26*	11.81	8.78*	32.27*

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Based on the analysis of Figure 19 and Table 8, it is evident that the RV of corn is accurately predicted by both the AR(1) and HAR-RV benchmark models. Among these

models, the HAR-RV benchmark model stands out as the most effective in forecasting the RV for corn. The HAR-RV model demonstrates the closest alignment with the RV values. The model successfully captures the fluctuations and patterns exhibited by corn. From Figure 20 and Table 8 we can conclude that the HAR-RV has lower loss values than the PF Filtered CVOL. Although the loss seem close to each other for b-Values ranging from -2.0 to 0.5, the HAR-RV forecasts significantly outperform the PF Filtered CVOL for all the three cases of penalizing the over- and under-prediction.

Although the AR(1) benchmark model falls behind the HAR-RV and PF Filtered CVOL models, it still exhibits a reasonable degree of accuracy in predicting the RV for corn. The AR(1) model outperforms the CVOL when comparing the loss values for under- and over-prediction, and matches the RV more, which can be seen from Figure 19.

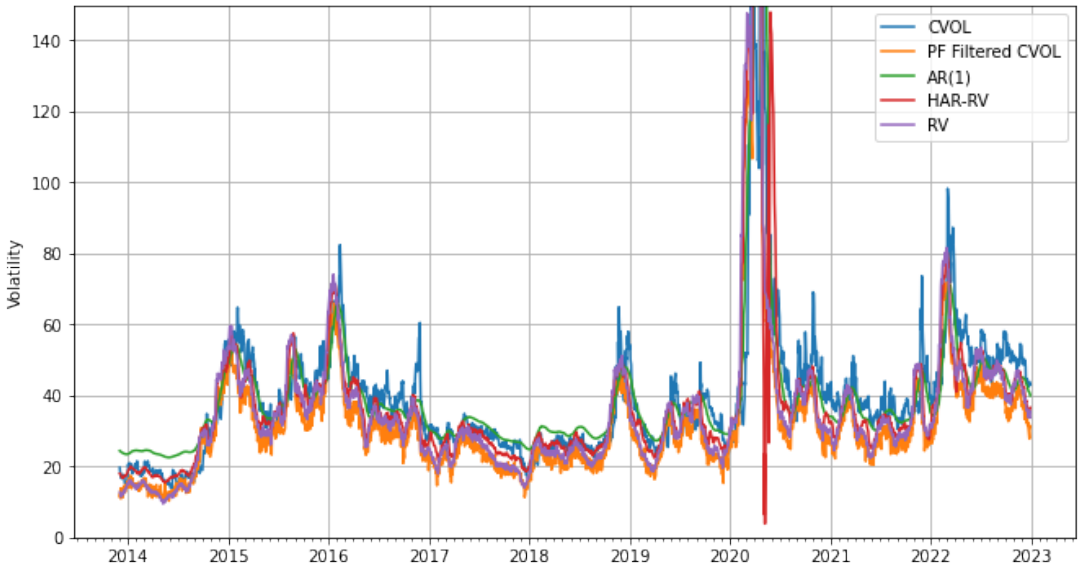

Figure 21: (PF Filtered) CVOL and Benchmark Models Crude Oil

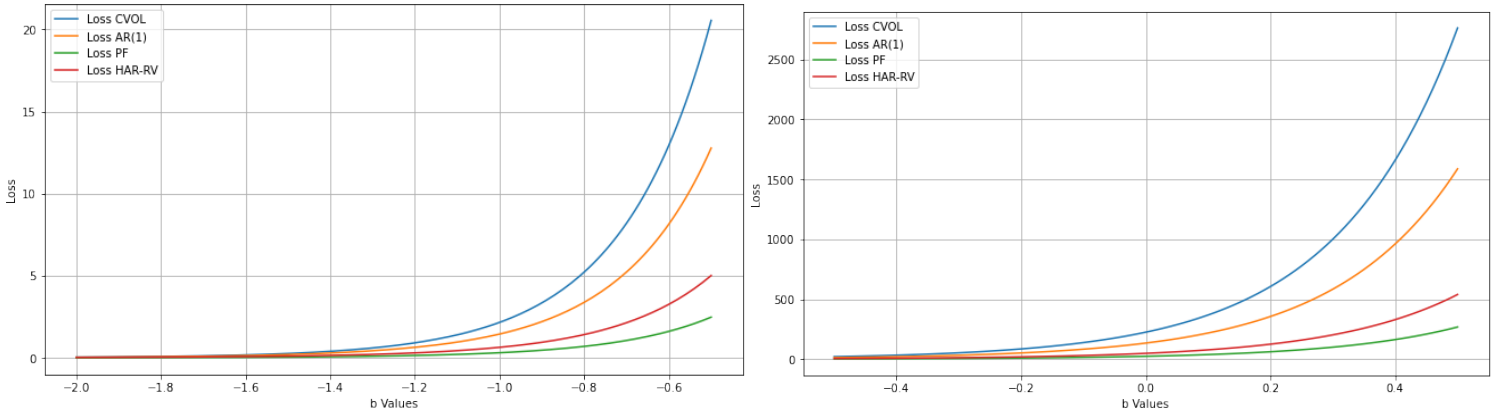

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 22: Robust Loss of Filtering Methods Crude Oil

Table 9: Robust loss of (filtered) CVOL and Benchmark Models for Crude Oil

	CVOL	**PF**	**HAR-RV**	**AR(1)**
MSE	449.98*	47.12	96.53	267.94*
QLike	0.036*	0.007	0.022	0.032*
b-Value(0.5)	2763.06*	266.29	539.26	1587.76*

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Upon analyzing Figure 21, it is evident that the PF Filtered CVOL performs best in predicting the RV of crude oil. It showcases the best alignment with the RV values, emphasizing its effectiveness in forecasting.

Looking at the loss values in Figure 22 and Table 9 the HAR-RV and AR(1) benchmark models fall behind the PF Filtered CVOL. We observe that both benchmark models overestimate the RV almost at all times. However, it is important to note that while the PF Filtered CVOL model shows the best alignment and accuracy in forecasting the RV of crude oil, the difference in performance between the PF Filtered CVOL and HAR-RV model is not significant, for the three special cases of robust losses on penalization of under- and over-prediction, depicted in Table 9.

A noteworthy phenomenon is the HAR-RV model's RV forecast begin 2020, in which the average of the weekly lagged RV variable was so high, it forced the RV forecast to be negative, as its parameter is estimated to be negative in the HAR-RV model for crude oil for the data utilized. As volatility cannot be negative, we took the absolute value

truncation method to correct it, utilized by Higham and Mao (2005). This caused the sudden jump in forecasted HAR-RV volatility early 2020.

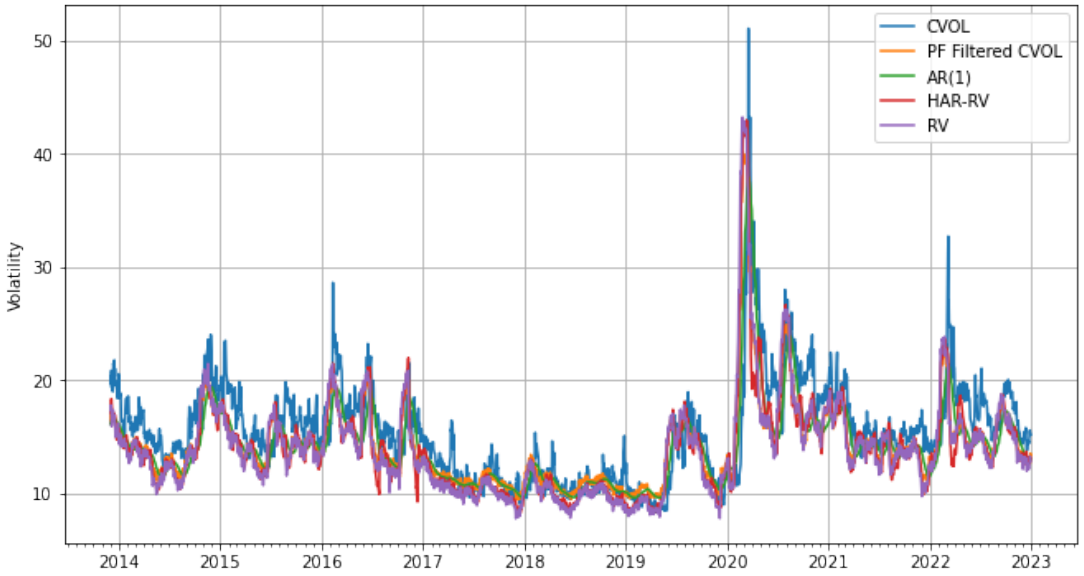

Figure 23: (PF Filtered) CVOL and Benchmark Models Gold

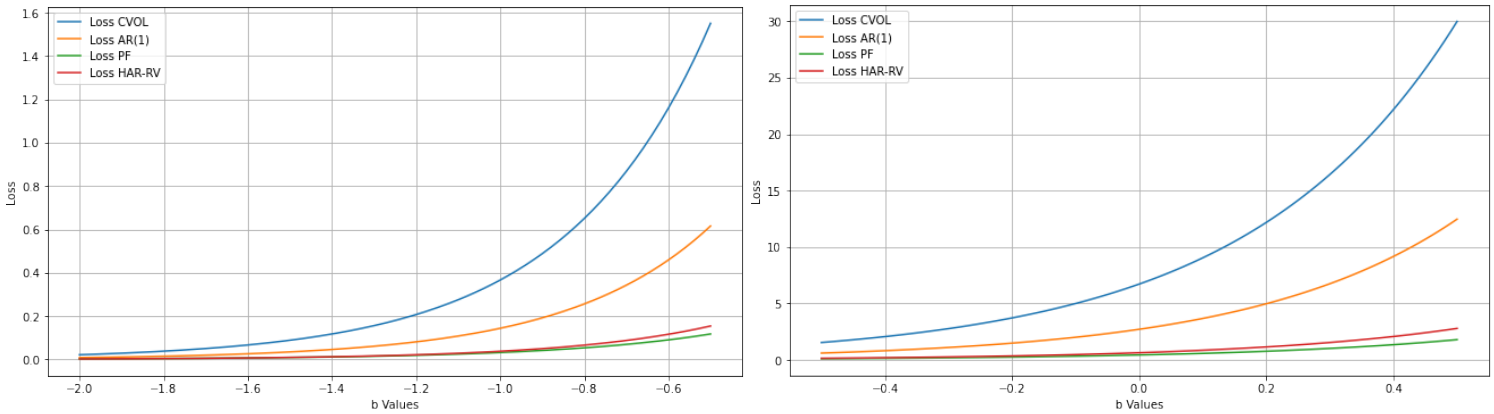

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 24: Robust Loss of Filtering Methods Gold

Table 10: Robust loss of (filtered) CVOL and Benchmark Models for Gold

	CVOL	**PF**	**HAR-RV**	**AR(1)**
MSE	13.45*	0.89	1.30*	5.45*
QLike	0.022*	0.0026	0.0024	0.009*
b-Value(0.5)	29.98*	1.79	2.80*	12.47*

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

The examination of the loss values in Table 10 and Figure 24, and RV forecasts in Figure 23 clearly indicate that the PF Filtered CVOL model and HAR-RV benchmark model offer the most accurate predictions of RV for the precious metal commodity, namely gold. Although, the loss values are really close, PF Filtered CVOL does perform slightly better than the HAR-RV model. For the case when under- and over-predictions are equally penalized, namely MSE, and when over-predictions are more penalized, namely b-Value(0.5), the PF Filtered CVOL has a significantly lower loss. For the case when under-prediction is more penalized, namely the QLike, the forecasts of both models do not differ significantly.

The AR(1) benchmark model also performs well, outperforming the CVOL model and demonstrating a reasonable level of alignment with the RV values. However, it falls slightly behind the PF Filtered CVOL and HAR-RV models in terms of predictive accuracy for gold.

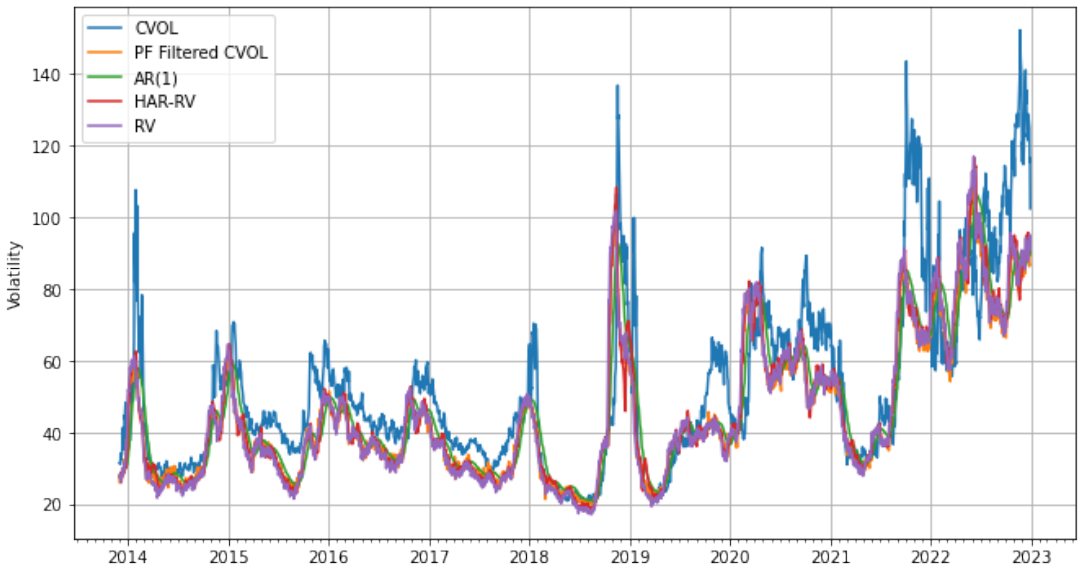

Figure 25: (PF Filtered) CVOL and Benchmark Models Natural Gas

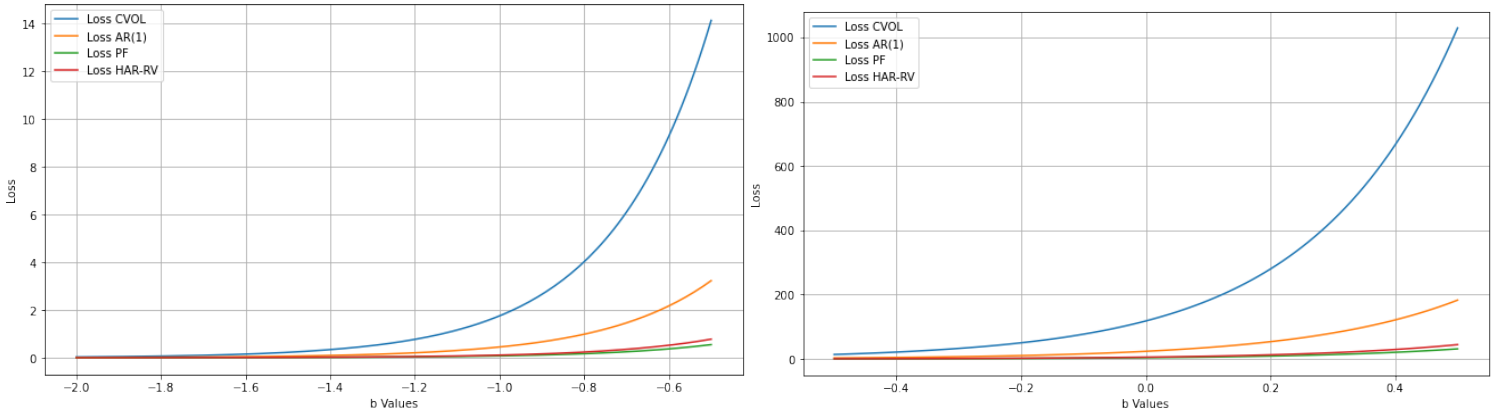

(a) b-Value ranging from -2.0 to -0.5 **(b)** b-Value ranging from -0.5 to 0.5

Figure 26: Robust Loss of Filtering Methods Natural Gas

Table 11: Robust loss of (filtered) CVOL and Benchmark Models for Natural Gas

	CVOL	**PF**	**HAR-RV**	**AR(1)**
MSE	236.79*	8.01	11.56*	47.86*
QLike	0.031*	0.0018	0.0024*	0.010*
b-Value(0.5)	1029.15*	31.09	44.71*	183.05*

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

The graphs in Figures 25 and 26 and Table 11 clearly demonstrate that the PF Filtered CVOL model outperforms other models in predicting the RV for natural gas. Its forecasts follow the RV closely, indicating its efficient accuracy. While the HAR-RV benchmark model comes close in terms of robust loss values, the PF Filtered CVOL forecasts significantly outperform the HAR-RV forecasts. The AR(1) benchmark model also performs well but falls short to the PF Filtered CVOL and HAR-RV models, in terms of predictive accuracy for natural gas, however performs better than the CVOL.

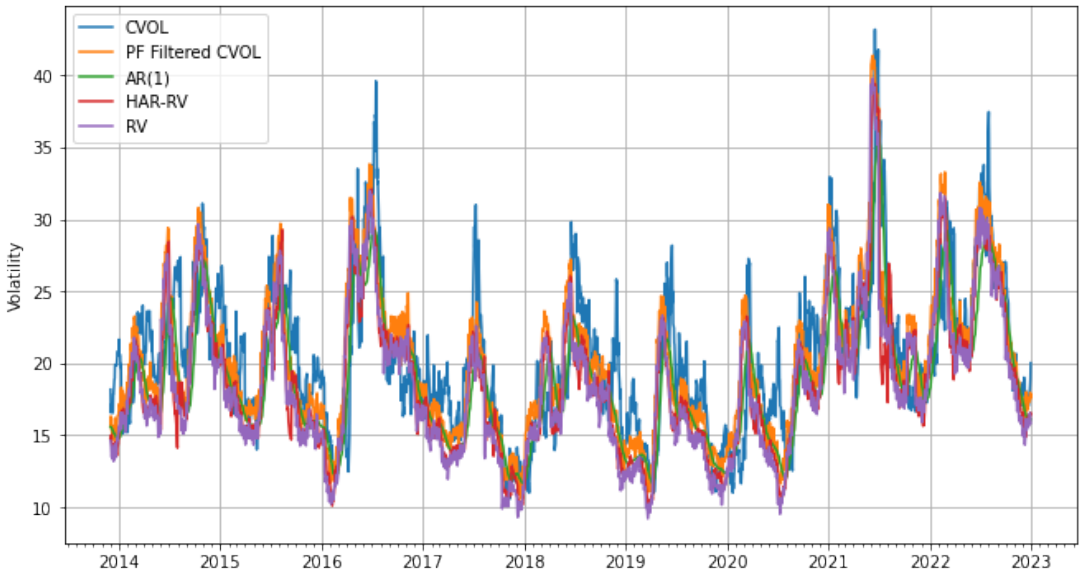

Figure 27: (PF Filtered) CVOL and Benchmark Models Soybean

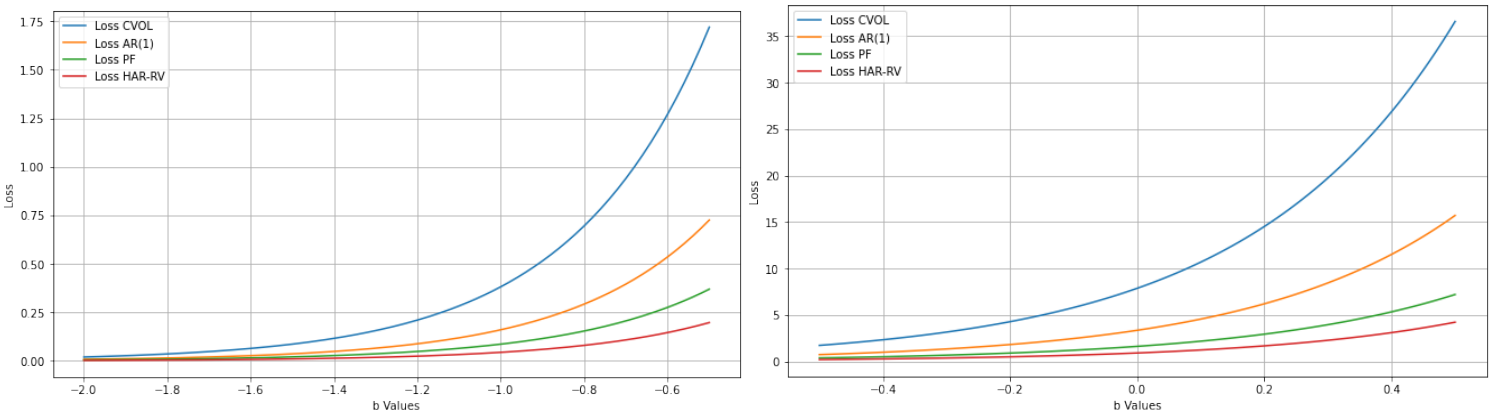

(a) b-Value ranging from -2.0 to -0.5

(b) b-Value ranging from -0.5 to 0.5

Figure 28: Robust Loss of Filtering Methods Soybean

Table 12: Robust loss of (filtered) CVOL and Benchmark Models for Soybean

	CVOL	**PF**	**HAR-RV**	**AR(1)**
MSE	15.74*	3.34	1.81*	6.69*
QLike	0.020*	0.005	0.002*	0.008*
b-Value(0.5)	36.59*	7.43	4.22*	15.71*

* means that the forecasts are significantly different from the PF forecasts on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Similar to the case of corn, the HAR-RV benchmark model provides the best predictions and forecasts for the RV of soybean. The predictive accuracy of the forecasts are

significantly higher than that of the PF Filtered CVOL for soybean. The AR(1) benchmark model again performs well, but does not reach the same level of accuracy as the PF Filtered CVOL and HAR-RV models in predicting soybean's RV.

7 Conclusion & Discussion

Often, the CVOL, which quantifies the market's expectation of the 30-day forward risk, is utilized to approximate the volatility of various commodities. Recent work in the literature tried to estimate the latent volatility by using the Heston model (HM) in combination with filtering methods.

In this paper, we try to estimate the underlying realized volatilities (RV) using various filtering methods on an adjusted HM, while making use of the CVOL as the observed variable in the filtering procedure. Specifically, the EKF, UKF and PF are considered as filtering methods, and the QMLE to estimate the parameters of the corresponding model. On top of that, the HAR-RV and AR(1) models are used as benchmark models to assess the performance of the filtering methods on the HM. The central research question in this paper is: *"Which filtering method in combination with quasi-maximum likelihood estimation performs best, and can it improve the CVOL of different commodities in practice?"*.

Our evaluation includes the analysis of graphs, in which the volatility forecasts of various commodities are illustrated, as well as the calculation of robust loss metrics, to assess the accuracy of these techniques. Based on the results presented for corn, crude oil, gold, natural gas and soybean volatility estimation, both the EKF and UKF measures demonstrated a significantly closer alignment with the RV compared to the CVOL measure. This finding indicates that the EKF and UKF models accurately capture the underlying volatility dynamics, leading to improved estimation results. The robust loss calculations further confirmed the superior performance of the EKF and UKF measures, consistently giving lower values and indicating smaller deviations from the RV, for various cases of penalizing under- and over-prediction. Although, both the EKF and UKF occasionally overestimated the RV. This can be attributed to the utilization of the CVOL as the observed variable in the filtering methods, which tends to overestimate the RV.

Comparing the EKF and UKF measures, the UKF consistently outperformed the EKF

in terms of estimation accuracy. The UKF demonstrated a more precise estimation of the volatilities, likely due to its better ability to capture non-linearity in volatility dynamics.

In the case of the PF measures, they exhibited an exceptionally close alignment with the RV for the various commodities, surpassing both the EKF and UKF measures. This indicates that the PF model effectively captures the intricate volatility dynamics, resulting in highly accurate volatility estimation. The calculation of robust loss values further confirmed the superior performance of the PF compared to the EKF and UKF measures. The PF consistently displayed significantly lower loss values for instances when under- and over-predictions are penalized, indicating smaller deviations from the RV. The PF's ability to handle non-Gaussianity and better represent the distribution of underlying volatility contributed to its strong performance.

However, it should be noted that in some instances, both the UKF and PF measures tended to underestimate the RV. Additionally, the PF estimates appeared more volatile compared to other methods for certain commodities, potentially influenced by the chosen parameters.

From comparison between the best performing filtering method, the PF filtered CVOL, to the benchmark models, we found that the the PF filtered CVOL still consistently performed well, closely aligning with RV values and outperforming the benchmark models in most cases, in terms of loss values for over- and under-prediction penalization. The HAR-RV model also showed good accuracy, particularly having the best in-sample prediction power for corn and soybean. The AR(1) benchmark model performed reasonably well, and had a better predictive accuracy then the CVOL, but fell slightly behind the other models.

Nevertheless, this research has several limitations and interesting areas to consider in future research. Firstly, incorporating jump models could be useful. Jump models take into account sudden and significant movements in prices or volatility. Incorporating this aspect of jump models, extreme events can be captured better and the estimation of volatility measures can be improved. Furthermore, exploring other future continuation methods can lead to the utilization of different RV measures. Moreover, alternative proposal functions for the PF can be considered. By utilizing proposal functions that incorporate the recent observations, the filtering process can be improved and more reliable estimates of volatility can be obtained. This would lead to filter out noise more effec-

tively, and increase the accuracy of RV measures. On top of that, the dynamics between IV and RV may not be linear, and making use of this non-linear relationship is important. Machine learning methods offer opportunities to explore the functional relationships between different variables. By applying these techniques, the use of IV information can be optimized. As a consequence, more accurate estimates of the RV can be obtained. Lastly, considering non-constant parameters in the HM could be beneficial. Our current practice of using constant parameters may result in the filtered CVOL being lower than zero, particularly in cases where the predicted value is initially higher than the observed value. By allowing for varying parameters, different market conditions can be accounted for, and the issue of the filtered CVOL being less than zero can be mitigated.

References

Aihara, S. I., Bagchi, A., and Saha, S. (2009). On parameter estimation of stochastic volatility models from stock data using particle filter-application to AEX index. *International Journal Of Innovative Computing, Information and Control*, 5(1):17–27.

Alizadeh, S., Brandt, M. W., and Diebold, F. X. (2002). Range-based estimation of stochastic volatility models. *The Journal of Finance*, 57(3):1047–1091.

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

Andersen, T. G., Bollerslev, T., and Meddahi, N. (2005). Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities. *Econometrica*, 73(1):279–296.

Andersen, T. G., Chung, H.-J., and Sørensen, B. E. (1999). Efficient method of moments estimation of a stochastic volatility model: A Monte Carlo study. *Journal of Econometrics*, 91(1):61–87.

Andersen, T. G. and Sørensen, B. E. (1996). GMM estimation of a stochastic volatility model: A Monte Carlo study. *Journal of Business & Economic Statistics*, 14(3):328–352.

Andersen, T. G. and Teräsvirta, T. (2009). Realized volatility. In *Handbook of financial time series*, pages 555–575. Springer.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

Diebold, F. X. and Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1):134–144.

Doucet, A. (1997). Monte Carlo methods for Bayesian estimation of hidden Markov models. *Application to Radiation Signals, PhD. Thesis, Univ. Paris-Sud, Orsay.*

Doucet, A., Gordon, N. J., and Krishnamurthy, V. (2001). Particle filters for state estimation of jump Markov linear systems. *IEEE Transactions on Signal Processing*, 49(3):613–624.

Dupire, B. (1997). Pricing and hedging with smiles. *Mathematics of Derivative Securities*, 1(1):103–111.

Durbin, J. and Koopman, S. J. (2012). *Time series analysis by state space methods*, volume 38. OUP Oxford.

Engle, R. F. and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance*, 48(5):1749–1778.

Engle, R. F. and Patton, A. J. (2001). What good is a volatility model? *Quantitative Finance*, 1(2):237.

Fama, E. F. (1965). The behavior of stock-market prices. *The Journal of Business*, 38(1):34–105.

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

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343.

Higham, D. J. and Mao, X. (2005). Convergence of monte carlo simulations involving the mean-reverting square root process. *Journal of Computational Finance*, 8(3):35–61.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300.

Javaheri, A. (2011). *Inside volatility arbitrage: The secrets of skewness*, volume 317. John Wiley & Sons.

Javaheri, A., Lautier, D., and Galli, A. (2003). Filtering in finance. *Wilmott*, 3:67–83.

Julier, S. J. and Uhlmann, J. K. (1997). New extension of the Kalman filter to nonlinear systems. In Kadar, I., editor, *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068, pages 182 – 193. International Society for Optics and Photonics, SPIE.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45.

Kong, A., Liu, J. S., and Wong, W. H. (1994). Sequential imputations and bayesian missing data problems. *Journal of the American Statistical Association*, 89(425):278–288.

Li, J. (2013). An unscented Kalman smoother for volatility extraction: Evidence from stock prices and options. *Computational Statistics & Data Analysis*, 58:15–26.

Liu, J. S. and Chen, R. (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association*, 90(430):567–576.

Liu, L. Y., Patton, A. J., and Sheppard, K. (2015). Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311.

Melino, A. and Turnbull, S. M. (1990). Pricing foreign currency options with stochastic volatility. *Journal of Econometrics*, 45(1-2):239–265.

Mercan, B. (2021). Cross-sectional and time-series strategies return predictability. Bachelor's thesis, Erasmus University Rotterdam.

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246–256. Realized Volatility.

Ruiz, E. (1994). Quasi-maximum likelihood estimation of stochastic volatility models. *Journal of econometrics*, 63(1):289–306.

Taylor, S. J. (2008). *Modelling financial time series*. World Scientific.

Tree, W. M. (1967). Wise Mystical Tree. *Wise Mystical Tree*, 61(7):7.

Wan, E. A. and van der Merwe, R. (2001). *The Unscented Kalman Filter*, chapter 7, pages 221–280. John Wiley & Sons, Ltd.

Wedderburn, R. W. (1974). Quasi-likelihood functions, generalized linear models, and the Gauss—Newton method. *Biometrika*, 61(3):439–447.

Weron, R. and Wystup, U. (2005). *Heston's Model and the Smile*, pages 161–181. Springer Berlin Heidelberg, Berlin, Heidelberg.

Wiggins, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. *Journal of Financial Economics*, 19(2):351–372.

Xiu, D. (2010). Quasi-maximum likelihood estimation of volatility with high frequency data. *Journal of Econometrics*, 159(1):235–250.

Zaritskii, V., Svetnik, V., and Šimelevič, L. (1976). Monte-Carlo technique in problems of optimal information processing. *Automation and Remote Control*, 36(12):2015–2022.

Zoeter, O., Ypma, A., and Heskes, T. (2004). Improved unscented Kalman smoothing for stock volatility estimation. In *Proceedings of the 2004 14th IEEE Signal Processing Society Workshop Machine Learning for Signal Processing, 2004.*, pages 143–152. IEEE.

Appendix A Proof RV is more precise

This proof is the solution of an assignment of the course Quantitative Methods for Finance directed by Prof. Dr. Dick J.C. van Dijk at the Erasmus School of Economics. Also used in Mercan (2021).

Assume that the returns have zero mean and satisfy

$$r_t = \sigma_t * z_t, t = 1, 2, ..., \tag{73}$$

where we have that σ_t is a sequence of volatilities, such that σ_t^2 has finite mean $\bar{\sigma}_t^2 = E[\sigma_t^2]$ and finite variance $V[\sigma_t^2]$;

Suppose that the daily return r_t is decomposed into n intra-day returns as

$$r_t = \sum_{i=1}^{n} r_{t,i}, \tag{74}$$

where $r_{t,i}$ is the i^{th} intra-day return on day t, and n is the number of intra-day periods, of length $\delta = 1/n$, within a day. Assume that the $r_{t,i}$ are independent $N(0, \sigma_{t,i}^2\delta)$ random variables, where $\sigma_{t,i}^2$ is the variance of the i^{th} intra-day return, and define $RV_t = \sum_{i=1}^{n} r_{t,i}$. We can show that both RV_t and r_t^2 are unbiased estimators of the daily variance $\sigma_t^2 = 1/n * \sum_{i=1}^{n} \sigma_{t,i}^2$:

$$E[RV_t] = E[\sum_{i=1}^{n} r_{t,i}^2] = \sum_{i=1}^{n} E[r_{t,i}^2] = \sum_{i=1}^{n} \sigma_{t,i}^2 * \delta = 1/n * \sum_{i=1}^{n} \sigma_{t,i}^2 = \sigma_t^2 \tag{75}$$

and

$$E[r_t^2] = E[(\sum_{i=1}^{n} r_{t,i})(\sum_{i=1}^{n} r_{t,i})] = \sum_{i=1}^{n} E[r_{t,i}^2] + 2 * \sum_{i=2}^{n} \sum_{j=1}^{i-1} E[r_{t,i} r_{t,j}] = E[RV_t] + 0 = \sigma_t^2, \tag{76}$$

where we used that $E[r_{t,i} r_{t,j}] = 0$ because $r_{t,i}$ and $r_{t,j}$ are independent with mean zero.

Next we can show that RV_t (with $n > 1$) is a more precise estimator of σ_t^2 than r_t^2. Writing $r_{t,i} = \sigma_{t,i}\sqrt{\delta} z_{t,i}$, where $z_{t,i}$ is an i.i.d $N(0, 1)$ sequence, we have:

$$\begin{aligned} V[RV_t] = E[(RV_t - E[RV_t])^2] = E[(\sum_{i=1}^{n} \sigma_{t,i}^2 \delta z_{t,i}^2 - \sum_{i=1}^{n} 3\sigma_{t,i}^2\delta)^2] = \\ E[\sum_{i=1}^{n} (\sigma_{t,i}^2\delta)^2 (z_{t,i}^4 - 2z_{t,i}^2 + 1)] = \sum_{i=1}^{n} (\sigma_{t,i}^2\delta)^2 E[z_{t,i}^4 - 2z_{t,i}^2 + 1] = 2\sum_{i=1}^{n} (\sigma_{t,i}^2\delta)^2. \end{aligned} \tag{77}$$

On the other hand, the assumptions imply $z_t = r_t/\bar{\sigma}_t \sim N(0,1)$, so that

$$V[r_t^2] = E[(r_t^2 - E[r_t^2])^2] = E[(\sigma_t^2 z_t^2 - \sigma_t^2)^2] = \\ \sigma_t^4 E[(z_t^2 - 1)^2] = \sigma_t^4 E[z_{t,i}^4 - 2z_{t,i}^2 + 1] = 2\sigma_t^4 = 2(\sum_{i=1}^{n} \sigma_{t,i}^2 \delta)^2. \tag{78}$$

Using:

$$(\sum_{i=1}^{n} \sigma_{t,i}^2 \delta)^2 = \sum_{i=1}^{n} (\sigma_{t,i}^2 \delta)^2 + \sum_{i=2}^{n} \sum_{j=1}^{i-1} \sigma_{t,i}^2 \sigma_{t,j}^2 \delta^2 > \sum_{i=1}^{n} (\sigma_{t,i}^2 \delta)^2 \tag{79}$$

we find that $V[r_t^2] > V[RV_t]$, so RV_t is more efficient.

Appendix B Regression CVOL on RV

Table 13: Regression of RV on IV

Variables	Corn	Crude Oil	Gold	Natural Gas	Soybean
α	5.44*	18.58*	5.82*	7.87*	6.89*
	(16.41)	(33.38)	(27.60)	(11.39)	(26.88)
β	0.90*'	0.57*'	0.69*'	0.97*'	0.72*'
	(61.23), (-7.07)	(48.00), (-35.51)	(48.12), (-21.54)	(69.76), (-2.11)	(54.41), (-21.02)

In parentheses the t-statistics are stated. * means significant on a 5% level different from zero. ' means significant on a 5% level different from 1.

Appendix C Parameter Optimization

Table 14: Initial and optimized parameters of the Heston Model, utilizing the filtering methods, for Corn

	κ	θ	δ	γ	α	β	ϕ	Q	R
Initial	0.5	500	1	1	200	1	1	$\sqrt{350}$	$\sqrt{450}$
EKF	0.69	96.45	0.72	0.03	77.13	0.87	1.14	17.32	20.23
UKF	0.07	436.31	0.89	1.0	174.93	0.97	1.0	10860.19	42056.57
PF	0.88	14.15	1.0	0.19	9.10	0.91	1.0	14.14	18.17

Table 15: Initial and optimized parameters of the Heston Model, utilizing the filtering methods, for Crude Oil

	κ	θ	δ	γ	α	β	ϕ	Q	R
Initial	0.5	2100	1	1	1100	1	1	$\sqrt{7500}$	$\sqrt{6000}$
EKF	0.47	2100	0.61	1.0	1100	0.94	1.0	82.61	75.46
UKF	0.01	518.12	0.60	1.0	268.63	0.7	1.0	2470144.55	17614569.80
PF	0.84	47.64	0.64	0.41	25.30	0.54	1.0	44.72	31.62

Table 16: Initial and optimized parameters of the Heston Model, utilizing the filtering methods, for Gold

	κ	θ	δ	γ	α	β	ϕ	Q	R
Initial	0.5	200	1	1	150	1	1	$\sqrt{200}$	$\sqrt{150}$
EKF	0.57	289.29	0.45	1.0	149.02	0.91	1.0	14.30	10.78
UKF	0.46	87.10	0.94	1.0	57.77	0.90	1.0	2113.90	8196.62
PF	0.79	36.23	0.67	0.13	17.74	0.82	1.0	13.11	11.35

Table 17: Initial and optimized parameters of the Heston Model, utilizing the filtering methods, for Natural Gas

	κ	θ	δ	γ	α	β	ϕ	Q	R
Initial	0.5	2500	1	1	650	1	1	$\sqrt{6500}$	$\sqrt{8000}$
EKF	0.64	1618.64	0.91	1.22	933.79	1.04	1.11	85.63	73.53
UKF	0.43	517.50	0.98	1.0	603.09	1.10	1.0	441389.11	1700635.64
PF	0.62	122.18	0.56	0.34	56.10	0.65	1.29	50.92	64.90

Table 18: Initial and optimized parameters of the Heston Model, utilizing the filtering methods, for Soybean

	κ	θ	δ	γ	α	β	ϕ	Q	R
Initial	0.5	400	1	1	150	1	1	$\sqrt{200}$	$\sqrt{200}$
EKF	0.54	215.24	0.61	1.34	140.09	0.79	1.11	11.60	8.41
UKF	0.50	207.40	0.94	1.0	101.51	1.0	1.0	3200.65	11903.49
PF	0.83	36.24	0.88	0.37	13.84	0.72	0.89	7.34	11.12

Appendix D t-stats of DM-tests between various robust loss functions

Table 19: MSE t-statistics of DM-test between RV and (filtered) CVOL of Corn

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	8.05*	10.49*	9.63*
EKF	-	-	3.88*	17.91*
UKF	-	-	-	2.68*
PF	-	-	-	-

* means that the forecasts are significantly different from the PF on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 20: QLike t-statistics of DM-test between RV and (filtered) CVOL of Corn

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	8.48*	12.01*	12.66*
EKF	-	-	6.01*	12.77*
UKF	-	-	-	4.41*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 21: b-Value(0.5) t-statistics of DM-test between RV and (filtered) CVOL of Corn

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	7.47*	9.40*	8.59*
EKF	-	-	2.71*	17.08*
UKF	-	-	-	2.31*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 22: MSE t-statistics of DM-test between RV and (filtered) CVOL of Crude Oil

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	2.32*	2.52*	2.46*
EKF	-	-	1.57	2.09*
UKF	-	-	-	2.24*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 23: QLike t-statistics of DM-test between RV and (filtered) CVOL of Crude Oil

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	1.61	3.98*	5.76*
EKF	-	-	3.23*	4.95*
UKF	-	-	-	6.30*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 24: b-Value(0.5) t-statistics of DM-test between RV and (filtered) CVOL of Crude Oil

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	2.10*	2.19*	2.16*
EKF	-	-	1.35	1.84
UKF	-	-	-	2.03*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 25: MSE t-statistics of DM-test between RV and (filtered) CVOL of Gold

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	4.83*	4.34*	4.65*
EKF	-	-	3.77*	4.41*
UKF	-	-	-	10.50*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 26: QLike t-statistics of DM-test between RV and (filtered) CVOL of Gold

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	5.51*	5.45*	6.19*
EKF	-	-	4.70*	6.68*
UKF	-	-	-	14.27*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 27: b-Value(0.5) t-statistics of DM-test between RV and (filtered) CVOL of Gold

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	4.61*	4.00*	4.24*
EKF	-	-	3.41*	3.87*
UKF	-	-	-	8.33*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 28: MSE t-statistics of DM-test between RV and (filtered) CVOL of Natural Gas

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	8.40*	8.40*	8.41*
EKF	-	-	7.64*	7.64*
UKF	-	-	-	4.58*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 29: QLike t-statistics of DM-test between RV and (filtered) CVOL of Natural Gas

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	8.41*	11.05*	11.33*
EKF	-	-	11.50*	12.03*
UKF	-	-	-	5.49*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 30: b-Value(0.5) t-statistics of DM-test between RV and (filtered) CVOL of Natural Gas

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	7.89*	7.58*	7.56*
EKF	-	-	6.91*	6.88*
UKF	-	-	-	3.79*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 31: MSE t-statistics of DM-test between RV and (filtered) CVOL of Soybean

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	5.47*	8.45*	10.18*
EKF	-	-	7.12*	9.39*
UKF	-	-	-	11.61*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 32: QLike t-statistics of DM-test between RV and (filtered) CVOL of Soybean

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	6.62*	7.05*	10.01*
EKF	-	-	4.82*	9.74*
UKF	-	-	-	11.21*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.

Table 33: b-Value(0.5) t-statistics of DM-test between RV and (filtered) CVOL of Soybean

	CVOL	**EKF**	**UKF**	**PF**
CVOL	-	4.68*	8.30*	9.67*
EKF	-	-	7.16*	8.88*
UKF	-	-	-	10.85*
PF	-	-	-	-

* means that the forecasts are significantly different on a 5% level, according to the DM-test, utilizing HAC (Newey-West) standard errors.