HEC MONTRÉAL

Particle Filter Performance with High Frequency Option Prices

Par

Marie-Ève Malette

Sciences de la gestion
(Ingénierie financière)

Mémoire présenté en vue de l'obtention du grade de maîtrise ès sciences (M.Sc.)

Décembre 2017

Abstract

An important development in the financial econometrics literature has been the use of high frequency observations to assess the daily variability of an asset's return. This is known as the realized variance. In this article, we apply the concept of realized variance to high frequency implied volatilities and define a new observable measure: the realized implied volatility variance. We demonstrate that this new measure can be used as a model-free approximation of the implied volatility quadratic variation which can be decomposed into a diffusive and discontinuous part. This separation of risk factors is convenient because it enables us to use the realized implied volatility variance as a latent state detection tool in a particle filtering procedure. We show that it offers incremental information about the return and variance dynamics, and that its value added, when used as a source of information, resembles a great deal to that of the realized variance of option prices.

Résumé

L'utilisation de données à haute fréquence pour évaluer la variabilitée quotidienne du rendement d'un actif financier constitue un développement important en économetrie financière, connue sous le nom de variance réalisée. Dans ce travail, nous appliquons le concept de variance réalisée aux volatilités implicites à haute fréquence et définissons une nouvelle mesure observable, à savoir, la variance réalisée de la volatilité implicite. Nous démontrons que cette nouvelle mesure peut être utilisée comme une approximation non-paramétrique de la variation quadratique de la volatilité implicite décomposable en une partie diffusive et une partie discontinue. Cette décomposition en facteur de risque est avantageuse, car elle nous permet l'utilisation de la variance réalisée de la volatilité implicite comme un outil de détection d'état latent via un filtre à particule. Nous montrons que la mesure que nous proposons offre de l'information supplémentaire sur les dynamiques de rendements et de la variance. Nous montrons, de plus, que la valeur ajoutée par notre mesure, lorsqu'utilisée comme source d'information, est comparable à celle ajoutée par la variance realisée des prix d'options.

Acknowledgements

I give special thanks to my supervisors, Geneviève Gauthier and Jean-François Bégin for their help, time and support. The completion of this project would never have been possible without them. I also appreciate the help of my parents who supported me during the entire project. Lastly, I want to thank my boyfriend Ishmael for his encouragements, help and comments on the final drafts.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Understanding and managing risk is crucial for investors. They make decisions based on their prediction of how markets are going to evolve in the future, but they face various forms of uncertainty. This is evidenced by the emergence over the last two decades of derivatives that give investors exposure to specific risk factors. From a research perspective, this implies that studying derivatives data can give us valuable insight about the different types of risks inherent to underlying assets. In this article, we are primarily interested in volatility risk, and more precisely in the detection of abrupt shifts in its dynamic. The ability to identify volatility jumps as they occur is important because it in turn affects the reward investors are given for bearing such a risk. A study made by Santa-Clara and Yan (2010), among others, reports that jump risk embedded in option data has a significant impact on the overall equity risk premium.

There is however a challenge with the estimation of variables like volatility jumps: How can we estimate something we cannot observe? Solutions to this type of problem are not new to the literature and are referred to as the filtering problem; a term used in the theory of stochastic processes to designate a mathematical model for state estimation. The main idea is to estimate a latent variable that can be modeled by using a potentially noisy set of observations. The primary objective of this paper is to document the efficacy of various observed measures in the filtration of modeled latent variables.

In this context, three main themes emerge: the model specification, the observations and the filtering procedure. What follows outlays the literature that was used to build the foundation of this article with respect to those three themes.

First, regarding model specification, we use a special case of the general affine jump-diffusion class of Duffie et al. (2000) to model log-equity price and variance. This sub-case was used by Amaya et al. (2017) in a similar study. Key features include the stochastic volatility of Heston (1993) with log-equity jumps and variance jumps, as well as the stochastic jump intensity specification of Eraker (2004). Early work on the mispecification of models include Bakshi et al. (1997), Bates (2000), Pan (2002), Duffie et al. (2000) and Eraker et al. (2003). They all report the inability of continuous-time models with stochastic volatility and jumps in return only to fit simultaneously return and option data. Conversely, the inclusion of variance jumps reduces the estimated volatility of the variance process to an empirically valid level. Additionally, the stochastic jump intensity is rationalized by the tendency of large movements to cluster over time.

As pointed out, single-factor models are not sufficient to characterize the variance dynamic. To this end, various additional model features have been studied. Although we did not incorporate them in ours, it is worth mentioning a few interesting findings regarding those features. For instance, Bardgett et al. (2014) study the inclusion of a stochastic level of reversion for the variance and report improvements in the estimation of its risk neutral distributions, the term structure of volatility smiles and of variance risk premia. Along similar lines, Andersen et al. (2015b) find that left tail risk cannot be captured by volatility and its components alone and propose a new tail factor. Another noteworthy property of return and variance processes, reported by Bandi and Renò (2016), is the existence of co-jumps between the two dynamics. The authors find empirical evidence of a strong negative dependence between the size of return and variance jumps. This feature is somewhat captured in our model by the stochastic log-equity jump intensity which depends on the variance level. Hence, if the variance jumps, the log-equity price is more likely

to jump as well. Bandi and Renò (2016) also emphasize that the literature that does not employ option data is less conclusive about the existence of such dependencies.

Second, with respect to observations, we use two sources of information: the S&P 500 Index prices as well as corresponding option prices. The key is to take advantage of high frequency data which have been revolutionary in econometrics. For example, the realized variance and bipower variation measures can be computed via high frequency index prices and are powerful model-free tools for detecting return jumps (Andersen et al., 2001 and Barndorff-Nielsen and Shephard, 2004). Christoffersen et al. (2012) and Christoffersen et al. (2015) show that the use of these measures significantly reduces option pricing errors.

Similarly, Amaya et al. (2017) incorporate the realized variance and bipower variation in their filtering procedure, along with other measures derived from the second source of information. They propose a new observable measure, the realized option variance (ROV) computed via the high frequency of option prices and demonstrate that its inclusion in the filtering process improves the detection of variance jumps. Our work is a continuation of theirs. As an alternative to the realized option variance, we define a similar observable measure, the realized implied volatility variance ($RIVV$), and compare its efficacy in the estimation of latent variables.

Another use of the realized variance of option prices has been documented by Audrino and Fengler (2015). Instead of using it as a jump detection tool (as in Amaya et al., 2017), the authors demonstrate its use as a means to validate option pricing models. They also provide evidence that the realized variance of underlying asset returns has a high predictive power when high frequency option data is not available. This conclusion corroborates the findings reported by Amaya et al. (2017), as well as our own: both the ROV and $RIVV$ are redundant with the RV when markets are calm (in the absence of jumps). However, in times of financial turmoil (in the presence of jumps), the options-related observable measures bring additional information about latent variables. Similar conclusions are reported by Bardgett et al. (2014), study-

ing the information content of the S&P 500 index and the VIX markets on volatility of the S&P 500 returns.

Additionnally, we want to highlight the work of Andersen et al. (2015a) and Todorov and Tauchen (2011) who both employ VIX Index data. The former construct a bundle of implied volatility measures across different moneynesses and maturities and report that it reacts differently to various latent variables. The latter conclude that the VIX index preserves important information about the behavior of latent instantaneous variance over small time scales. More precisely, the authors conclude that the VIX index, and hence the instantaneous variance process, can be described as a pure-jump process with infinite variation jumps. We find these conclusions interesting because in our empirical study, we observe that for long dated options, the $RIVV$ measure is often near zero, but in times of financial stress, it spikes intensely.

Third, with regards to the filtering process, we use Gordon et al. (1993)'s SIR-type filter, the bootstrap filter with multinomial resampling. Our approach is similar to that of Amaya et al. (2017). In fact, our main objective is to compare the use of their proposed ROV measure to our proposed $RIVV$ measure. To this end, we reproduce part of their work and test both measures' impact on the filter's performance. There is some disadvantage to the particle filter we use. For instance, it can suffer from sample impoverishment and weight degeneracy. Many variations of filters that mitigate this type of problem exist and have been documented. For instance, Bardgett et al. (2014) used the auxiliary particle filter of Pitt and Shephard (1999). Johannes et al. (2009) studied both methods and reported that the former is better at jump detection. Our aim however is to compare two observable measures related to option prices, the ROV of Amaya et al. (2017) and the $RIVV$. Hence, whether we use an auxiliary particle filter or standard particle filter does not matter. Creal (2012) provides a complete review of particle filters with applications in finance.

To summarize, the objective of this article is to compare the particle filter's performance in the estimation of latent variables using various sources of high frequency data and a SIR-type filtering procedure. More precisely, we want to

know if the information content of our proposed $RIVV$ measure differs from that of Amaya et al. (2017)'s ROV measure. Keeping in mind, we make the following contributions:

- We define a new observable variable to use in the filtering process, the realized implied volatility variance ($RIVV$), and its theoretical equivalent, the implied volatility quadratic variation. We also provide some intuition of the $RIVV$'s behavior based on its theoretical dynamic (Chapter 2).

- We test the $RIVV$'s ability to improve the accuracy of the estimation of latent variables when it is included in the filtering process along with other observables (Chapter 3).

- We demonstrate the $RIVV$'s redundancy with the realized variance of Barndorff-Nielsen and Shephard (2004) in the absence of jumps. Conversely, we also demonstrate the $RIVV$'s value added relative to the realized variance in the presence of jumps (Chapter 3).

- We provide proof that the realized implied volatility variance is a good approximation of the implied volatility quadratic variation, especially in the presence of jumps (Chapter 3).

- We show that the information content of the realized implied volatility variance resembles to that of the realized option variance from a theoretical perspective (Chapter 3 and 4).

- We provide a global picture of the empirical properties of the realized implied volatility variance and make some comparisons with the realized option variance (Chapter 4).

Chapter 2

Setup

In this chapter, we present and explain the model specification we used in our study. We then move on to explain how latent states can be approximated using empirically observable measures with regards to two sources of information: Index and options data. We formally define our proposed observable variable to use in the filtering process, the realized implied volatility variance, and its theoretical equivalent, the implied volatility quadratic variation. We also provide some intuition of its behavior based on its theoretical dynamic.

2.1 Model

Our model is part of the general affine jump-diffusion class of Duffie et al. (2000) which has the benefit of providing a semi-closed form solution for option pricing. Since our goal is to estimate latent variables, namely the instantaneous variance, log-equity jumps and variance jumps, our model incorporates those features. More precisely, it includes the square root mean reverting stochastic volatility specification of Heston (1993). Additionally, a compound Poisson process with Gaussian jumps is included in the log-price to account for fat tails of return distributions. Similarly, exponentially distributed jumps are incorporated in the variance dynamic. The omission of this feature tends to result in misspecified models as noted by Bakshi et al. (1997), Bates (2000)

and Pan (2002).

The asset model under the physical probability measure $\mathbb{P}$ is given by the following stochastic differential equations:

$$
\begin{aligned}
dY_t &= \alpha^{\mathbb{P}}_{t^-} dt + \sqrt{V_{t^-}} dW_{Y,t} + dJ_{Y,t} && (2.1)\\
dV_t &= \kappa(\theta - V_{t^-})dt + \sigma\sqrt{V_{t^-}} dW_{V,t} + dJ_{V,t} && (2.2)\\
W_{Y,t} &= \rho W_{V,t} + \sqrt{1-\rho^2} W_{\perp,t} \\
J_{Y,t} &= \sum_{n=1}^{N_{Y,t}} Z_{Y,n}, \qquad Z_{Y,n} \sim \mathcal{N}(\mu_Y,\, \sigma_Y^2) \\
J_{V,t} &= \sum_{n=1}^{N_{V,t}} Z_{V,n}, \qquad Z_{V,n} \sim Exp(\mu_V)
\end{aligned}
$$

where $W_{V,t}$ and $W_{\perp,t}$ are two independent standard Brownian motion under $\mathbb{P}$, and $\alpha^{\mathbb{P}}_{t^-}$ is the drift of the log-price Y[1]:

$$
\alpha^{\mathbb{P}}_{t^-} = r - q + \Big(\eta_Y - \frac{1}{2}\Big) V_{t^-} + (\gamma_Y - (\varphi^{\mathbb{P}}_{Z_Y}(1) - 1))\lambda_{Y,t^-}. \tag{2.3}
$$

In Equation (2.3), r and q correspond to the risk-free rate and dividend yield, respectively. In addition, η_Y and γ_Y represent parameters constituting the diffusive and jump risk premium of Y, and finally, $\varphi^{\mathbb{P}}_{Z_Y}(1)$ is the cumulant generating function of Z_Y evaluated at 1.

Log-equity jumps are generated with Cox process, $\{N_{Y,t}\}_{t\geqslant 0}$, with a stochastic intensity that depends on the instantaneous variance: $\lambda_{Y,t^-} = \lambda_{Y,0} + \lambda_{Y,1} V_{t^-}$. On the other hand, variance jumps are generated with a Poisson process, $\{N_{V,t}\}_{t\geqslant 0}$, with constant intensity $\lambda_{V,t^-} = \lambda_{V,0}$. The size $J_{Y,t}$ of log-price jumps follows a Gaussian distribution with mean μ_Y and standard deviation σ_Y, whereas the size $J_{V,t}$ of variance jumps is exponentially distributed with mean μ_V.

[1] We refer the reader to Appendix A of Amaya et al. (2017) for details on the choice of Radon-Nykodym derivative, the model under the risk-neutral measure and its impact on option pricing.

2.2 Separating Latent Variables: Theory and Empirics

We now analyze the theoretical measures derived from our model framework, along with the observable variables used in the filtering of latent states. What follows reproduces the work of Amaya et al. (2017) and extends it with an additional observable, named the realized implied volatility variance ($RIVV$).

2.2.1 Index-related Information

This subsection presents and analyzes three index-related sources of information about latent processes: log-price values, realized variances and bipower variation of the underlying asset log-price process. The first and most basic observable variable is the log-price value, Y_t:

Observable and Theoretical Measure #1: Y_t. (2.4)

Merton (1980) pointed out that by increasing the sample frequency, log-prices reveal information about the volatility process. Additionally, Barndorff-Nielsen and Shephard (2004) mention that log-prices are informative not only about the volatility process, but also about jumps in returns. Similar to our approach, Johannes et al. (2009) implement a particle filter using returns as an observable to infer latent states.

The second observable we use is the realized variance of log-prices obtained from intraday prices:

$$\textbf{Observable \#2: } RV_{t-\Delta,t} = \sum_{j=1}^{N} \left(Y_{t-\Delta+j\frac{\Delta}{N}} - Y_{(t-\Delta)(j-1)\frac{\Delta}{N}} \right)^2 \tag{2.5}$$

where Δ is the time interval and N is the number of steps in that interval. It was initially proposed by Andersen et al. (2001) as a good approximation of

the log-price quadratic variation of log-prices:

$$\textbf{Theoretical Measure \#2: } \Delta QV_{t-\Delta,t} = \int_{t-\Delta}^{t} V_{s^-} ds + \sum_{n=N_{Y,t-\Delta}+1}^{N_{Y,t}} \left(Z_{Y,n}\right)^2. \quad (2.6)$$

Protter (2004) argued that as the sample frequency N increases, the realized variance converges to the quadratic variation if we forget about market microstructure noise. The quadratic variation, the theoritical counterpart of the realized variance, is said to be a global measure of price volatility. More precisely, the price variation it captures is twofold: the volatility due to the diffusive component and the volatility due to log-equity price jumps.

The third observable used in this study is the realized bipower variation of the log-price:

$$\begin{aligned}\textbf{Observable \#3: } BV_{t-\Delta,t} = \frac{\pi}{2} \sum_{j=2}^{N} & \left|Y_{t-\Delta+\frac{j\Delta}{N}} - Y_{t-\Delta+\frac{(j-1)\Delta}{N}}\right| \\ \times & \left|Y_{t-\Delta+\frac{(j-1)\Delta}{N}} - Y_{t-\Delta+\frac{(j-2)\Delta}{N}}\right|, \end{aligned} \quad (2.7)$$

initially introduced by Barndorff-Nielsen and Shephard (2004) as a good approximation of the integrated variance:

$$\textbf{Theoretical Measure \#3: } \Delta I_{t-\Delta,t} = \int_{t-\Delta}^{t} V_{s^-} ds. \quad (2.8)$$

Contrary to the realized variance, which captures both types of price variation, the bipower variation only captures the price volatility due the diffusive term. This means that when compared, one can isolate the jumps in the return dynamics. More precisely, if the difference between the realized variance and the bipower variation for a given period is significantly different than zero, then at least one log-equity jump occurred over that time interval.

2.2.2 Option-related Information

We now present and analyze three option-based sources of information: the implied volatiliy, the realized option variance and the realized implied volatility variance.

The first observable related to option prices is the time series (of the cross-section) of Black and Scholes implied volatilities derived from end of day information:

$$\textbf{Observable \#4: } \sigma_{t,i}^{BS}, i \in 1, 2, ...n_t, \tag{2.9}$$

where n_t is the number of options considered at the end of day t. The implied volatility helps capture how option prices react to latent variables with different moneyness levels. Another notable advantage of using the implied volatility is that it is invariant to the price level. For more details on the benefits of the implied volatility, we refer the reader to Renault (1997). The theoretical counterpart of the implied volatility is straightforward:

$$\textbf{Theoretical Measure \#4: } IV^{BS}(O_{t,i}), i \in 1, 2, ...n_t, \tag{2.10}$$

where IV^{BS} is a function that computes the Black-Scholes implied volatility of an option from its price O_t, and n_t is the number of options considered at the end of day t. The option price O_t in Equation (2.10) is itself a function of the log-equity price Y_t and the instantaneous variance V_t, and comes from the option pricing framework provided in Appendix A.

Next, we also include the realized option variance (ROV) proposed by Amaya et al. (2017). It is computed using the high frequency time series of intraday option prices:

$$\textbf{Observable \#5: } ROV_{t-\Delta,t,i} = \sum_{j=1}^{N} \left(O_{t-\Delta+j\frac{\Delta}{N},i} - O_{t-\Delta+(j-1)\frac{\Delta}{N},i} \right)^2, \tag{2.11}$$

where $i = \{1, 2, ...n_t\}$ and where n_t is the number of options considered at the end of day t. The option price O_t in Equation (2.11) represents the observed

option price at time t, in contrast with the theoretical option price of Equation (2.10). The authors show that the *ROV* can approximate the quadratic variation of option prices which is given by the following equation for a European option:

$$\begin{aligned}\langle O, O\rangle_t = \int_0^t \Bigg(\left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\right)^2 + \sigma^2 \left(\frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\right)^2 \\ + 2\sigma\rho \left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\right)\left(\frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\right)\Bigg) V_{u^-} du \\ + \sum_{0<u<=t} \Big\{ O_u(Y_u, V_u) - O_u(Y_{u^-}, V_{u^-}) \Big\}^2. \end{aligned} \tag{2.12}$$

It is worthy to mention that the last term of Equation (2.12) captures the variation due to jumps and can be broken down in two parts: the variation due to jumps in the log-price dynamics and jumps in the variance process:

$$\begin{aligned} &\sum_{0<u\leqslant t} \big\{ O_u(Y_u, Y_u) - O_u(Y_{u^-}, V_{u^-}) \big\}^2 \\ &= \sum_{0<u\leqslant t} \big\{ O_u(Y_u, V_u) - O_{u^-}(Y_{u^-}, V_u) \big\}^2 + \sum_{0<u\leqslant t} \big\{ O_u(Y_u, V_u) - O_{u^-}(Y_u, V_{u^-}) \big\}^2. \end{aligned} \tag{2.13}$$

To simplify matters, we rewrite the previous equation with the following terminology:

$$\sum_{0<u\leqslant t} (\Delta O_u)^2 = \sum_{0<u\leqslant t} (\Delta_Y O_u)^2 + \sum_{0<u\leqslant t} (\Delta_V O_u)^2. \tag{2.14}$$

The realized option variance given by Equation (2.11) is a good approximation for the incremental option quadratic variation over the same time interval:

$$\textbf{Theoretical Measure \#5: } \Delta OQV_{t-\Delta,t,i} = \langle O, O\rangle_{t,i} - \langle O, O\rangle_{t-\Delta,i}, \tag{2.15}$$

where $i = \{1, 2, ...n_t\}$ and n_t is the number of options considered at the end of day t.

Before introducing the final observable used in this paper, we summarize

a few properties inherent to the realized option variance, which are covered in more detail in Amaya et al. (2017). Based on Equation (2.12), the realized option variance is a function of the option delta and variance-vega, and hence, will vary according to the moneyness of the option at stake.[2] Keeping that in mind, Amaya et al. (2017) explain that the realized option variance is redundant with the realized variance of the log-price dynamic in the case of both deep in-the-money (ITM) and at-the-money (ATM) options. However, for a short-maturity deep out-of-the-money (OTM) option, the realized option variance behaves like a pure-jump process, a property also mentioned by Andersen et al. (2015a) about option prices. Therefore, the realized option variance of a short-term out-the-money option is approximated by Equation (2.14) such that:

$$ROV_t^{\text{OTM}} \simeq \langle O, O \rangle_t^{\text{OTM}} \simeq \sum_{0<u\leqslant t} (\Delta O_u)^2.$$

which helps us identify jumps in the volatility process because it can be written as:

$$ROV_t^{\text{OTM}} \simeq \langle O, O \rangle_t^{\text{OTM}} \simeq \sum_{0<u\leqslant t} (\Delta_Y O_u)^2 + \sum_{0<u\leqslant t} (\Delta_V O_u)^2.$$

More precisely, for a given time interval, if the difference between the realized variance and bipower variation is statistically null, then there is no log-equity jump. Furthermore, if the realized option variance of OTM option for the same time interval is significantly different than 0, it means a variance jump occurred.

Finally, we introduce the last observable used in this study, the realized implied volatility variance. It is computed using the time series (of the cross

[2]The option delta is given by the partial derivative of the option price relative to the underlying asset price S. Hence, $\frac{\partial O_t}{\partial S} = \frac{\partial O_t}{\partial y}\frac{\partial Y_t}{\partial S} = \frac{\partial O_t}{\partial S}\frac{1}{S_t}$, since $Y_t = \log(S_t)$.

section) of intraday Black and Scholes implied volatilities:

$$\textbf{Observable \#6: } RIVV_{t-\Delta,t,i} = \sum_{j=1}^{N} \left(\sigma^{BS}_{t-\Delta+j\Delta/N} - \sigma^{BS}_{t-\Delta+(j-1)\Delta/N}\right)^2, \tag{2.16}$$

where $i = \{1, 2, ...n_t\}$ and n_t is the number of options considered on day t. The implied volatility realized variance serves as an approximation of the implied volatility quadratic variation, which is defined by the following expression for a European option in our model[3]:

$$\begin{aligned}\langle IV, IV\rangle_t = \int_0^t \Bigg(\left(\frac{\partial IV_u}{\partial y}(O_{u^-})\right)^2 + 2\sigma\rho\left(\frac{\partial IV_u}{\partial y}(O_{u^-})\right)\left(\frac{\partial IV_u}{\partial v}(O_{u^-})\right) \\ + \sigma^2\left(\frac{\partial IV_u}{\partial v}(O_{u^-})\right)^2\Bigg) V_{u^-}du + \sum_{0<u\leqslant t} \left\{IV_u(O_u) - IV_u(O_{u^-})\right\}^2. \end{aligned} \tag{2.17}$$

The last term of Equation (2.17) captures the variation due to jumps and can be broken down into two parts: the variation due to log-price jumps and variance jumps. We can also rewrite the implied volatility as a function of Y_t and V_t instead of the option price O_t. The last term of Equation (2.17) becomes:

$$\begin{aligned} &\sum_{0<u\leqslant t} \left\{IV_u(O_u) - IV_u(O_{u^-})\right\}^2 \\ &= \sum_{0<u\leqslant t} \left\{IV_u(Y_u, V_u) - IV_u(Y_{u^-}, V_u)\right\}^2 + \sum_{0<u\leqslant t} \left\{IV_u(Y_u, V_u) - IV_u(Y_u, V_{u^-})\right\}^2 \\ &= \sum_{0<u\leqslant t} (\Delta_Y IV_u)^2 + \sum_{0<u\leqslant t} (\Delta_V IV_u)^2 \end{aligned}$$

The difference between the implied volatility quadratic variation at time t and $t - \Delta$ is the theoretical equivalent of the realized implied volatility variance over the same time interval:

$$\textbf{Theoretical Measure \#6: } \Delta IVQV_{t-\Delta,t,i} = \langle IV, IV\rangle_{t,i} - \langle IV, IV\rangle_{t-\Delta,i},$$

[3]Details on how to derive Equation (2.17) using Itō's lemma along with explanations on how to compute the derivatives are available in Appendix B.1.

where $i = \{1, 2, ...n_t\}$ and n_t is the number of options considered on day t.

Using Equation (2.17), we want to give some intuition on the behaviour of the realized implied volatility variance with respect to the moneyness of the option at stake. The key is to rewrite the partial derivatives using the chain rule such that

$$\frac{\partial IV_t}{\partial y} = \frac{\partial IV_t}{\partial O}\frac{\partial O_t}{\partial S}\frac{\partial S_t}{\partial y} \quad \text{and} \quad \frac{\partial IV_t}{\partial v} = \frac{\partial IV_t}{\partial O}\frac{\partial O_t}{\partial v},$$

where $\frac{\partial O_t}{\partial S}$ and $\frac{\partial O_t}{\partial v}$ are the Black Scholes delta and the variance-vega respectively. Moreover, note that $\frac{\partial S_t}{\partial y} = e^{Y_t}$. Then, the missing link is the partial derivative $\frac{\partial IV_t}{\partial O}$ which is in fact the inverse of the Black and Scholes vega and is defined as[4]:

$$\frac{\partial IV_t}{\partial O} = \frac{1}{\exp(Y_t)\phi(d_t)\sqrt{T-t}}.$$

As a European call option becomes deeper in-the-money, its delta tends towards 1, whereas both the variance-vega and Black and Scholes vega, $\frac{\partial O_t}{\partial IV}$, should be close to zero. Indeed, an ITM option is insensitive to changes in volatility or variance. Consequently, the inverse of the Black and Scholes vega should tend towards infinity. Therefore, the realized implied volatility variance of an ITM option is approximately:

$$RIVV_t^{\text{ITM}} \simeq \int_0^t \Big(\frac{\partial IV_u}{\partial O}\Big)^2 e^{2Y_{u^-}} V_{u^-} du + \sum_{0<u\leqslant t} (\Delta IV_u)^2,$$

where $\lim_{\exp(Y_{u^-})/K\to\infty} \frac{\partial IV_u}{\partial O} = \infty$, and K is the strike price.

For an ATM option, the inverse of the Black and Scholes vega is given by:

$$\frac{\partial IV_t}{\partial O} = \frac{C}{\exp(Y_{t^-})\sqrt{T-t}},$$

where C is a known constant, i.e. $C = \frac{1}{\phi(0)} \simeq 2.506$, where $\phi(\cdot)$ is the standard

[4]We provide more details in Appendix B.2.1 for further details.

normal density function. The delta is close to $\frac{1}{2}$ and the variance-vega is at its highest. Hence, the $RIVV$ behaves according to:

$$RIVV_t^{\text{ATM}} = \int_0^t \frac{C^2}{4\exp(2Y_{u^-})(T-u)} \left\{ 1 + \rho\sigma \frac{\partial O_u}{\partial v} + \sigma^2 \left(\frac{\partial O_u}{\partial v} \right)^2 \right\} V_{u^-} du + \sum_{0<u\leqslant t} (\Delta IV_u)^2$$

Notice that the integral also depends on the time-to-maturity $T-t$. It means that when maturity is imminent, the term in the integral will increase exponentially. On the other hand, when the maturity is very long, it will tend towards zero. Our empirical study validates this behaviour (see Chapter 5).

Lastly, as an option gets deeper out-of-the-money, it becomes again insensitive to changes in volatility and variance. Thus, the Black and Scholes vega as well as the variance-vega should be close to 0. Consequently, the inverse of the Black and Scholes vega should tend towards infinity. However, because the delta is close to 0, the integral of Equation (2.17) will tend towards 0, with only the jump component remaining:

$$RIVV_t^{\text{OTM}} \simeq \sum_{0<u\leqslant t} (\Delta IV_u)^2.$$

Our empirical study also confirms this behaviour for options that are not close to maturity.

To summarize, we have shown that our proposed $RIVV$ observable can be broken down into two parts: a diffusive component that is also impacted by the time to maturity of the option at stake, and a jump component, that is further divisible into a return and variance part. These components can be isolated by using options with different moneynesses. More precisely, we showed that the $RIVV$ of an OTM option behaves like a pure-jump process. Recall that the ROV of Amaya et al. (2017) exhibits a similar behaviour.

Chapter 3

The Filtering Process

In dynamic models, the use of Bayesian state estimation is a common approach to infer latent states, such as the instantaneous variance and jumps. Bayesian filtering involves an unobserved vector of variables, $L_{1:t} = \{L_1, ...L_t\}$ that evolves according to a dynamic model and a set of observations $d_{1:t} = \{d_1, ..., d_t\}$ where the subscript t denotes the time period. The goal of this method is to estimate the posterior distribution of latent variables given the observed data $f(L_{1:t}|d_{1:t})$.

Depending on the model, exact inference methods such as Kalman filters and hidden Markov model filters can be used. However, for complex financial models that contain non-linearities in their dynamics or non-Gaussianities in their innovations, approximation techniques are often required. This chapter examines one of the most commonly used approximation methods, the sequential Monte Carlo method, also called particle filtering. Simply put, the filtering process works in two recursive steps: the prediction, followed by the update. The first step consists of predicting a sample of the latent variables, $\{L_t\}_{i=1}^N$, called particles, according to the chosen dynamic model and assuming the latent variables up to time $t-1$ are known. Then, the update step consists of updating the importance weights $\{\check{w}_t\}_{i=1}^N$ of each particle according to the observed data d_t at time t. The resulting samples from each iteration represent an approximation for the posterior distribution of latent state variables given

the observed data, $f(L_{1:t}|d_{1:t})$.

The rest of this chapter is divided into three sections: first, we review the theory behind filtering. Second, we describe a space-type of particle filter, called the bootstrap filter, and finally, we describe how the filtering algorithm is implemented given our model.

3.1 The Theory Behind Filtering

It is convenient to cast a model in a state space form when dealing with filtering applications. A state space model is defined by two equations, the observation equation and the transition equation, respectively, given by

$$d_t = g_t(L_t, \omega_t), \tag{3.1}$$

$$L_t = m_t(L_{t-1}, \zeta_t), \tag{3.2}$$

where ω_t and ζ_t denote the noise innovations. The densities corresponding to both Equations (3.1) and (3.2) are defined as the observation density $f(d_t|L_t)$ and the transition density $f(L_t|L_{t-1})$, respectively.

There are two main concepts involved in the filtering process: Bayesian inference and sequential importance sampling. The first allows us to rewrite the posterior distribution of latent state variables given the observed data as a recursion using Bayes' theorem:

$$\begin{aligned} f(L_{1:t}|d_{1:t}) &= \frac{f(d_t|L_{1:t}, d_{1:t-1})f(L_t|L_{1:t-1}, d_{1:t-1})}{f(d_t|d_{1:t-1})} f(L_{1:t-1}|d_{1:t-1}) \\ &= \frac{f(d_t|L_t)f(L_t|L_{t-1})}{f(d_t|d_{1:t-1})} f(L_{1:t-1}|d_{1:t-1}). \end{aligned} \tag{3.3}$$

The second concept consists of defining a Monte Carlo estimator of the posterior latent states' distribution given the observed data. This is done by means of the sequential importance sampling (SIS). To that end, let's first suppose that we want to estimate a function g that depends on the latent

variables. This would be equivalent to approximating the following integral

$$\mathbb{E}[g(L_{1:t})] = \int g(L_{1:t}) f(L_{1:t}|d_{1:t}) dL_{1:t}, \tag{3.4}$$

using a Monte Carlo estimator. This could be achieved sampling paths $L_{1:t}^{(i)}$ directly from the posterior distribution $f(L_{1:t}|d_{1:t})$. However, because it is often impossible to sample directly from the posterior distribution, importance sampling allows us to estimate g by sampling instead from a proposal distribution $q(L_{1:t}|d_{1:t})$. Hence, Equation (3.4) is equivalent to

$$\mathbb{E}[g(L_{1:t})] = \int g(L_{1:t}) \frac{f(L_{1:t}|d_{1:t})}{q(L_{1:t}|d_{1:t})} q(L_{1:t}|d_{1:t}) dL_{1:t}. \tag{3.5}$$

Then, the Monte Carlo estimator

$$\mathbb{E}[g(L_{1:t})] \approx \sum_{i=1}^{N} g(L_{1:t}^{(i)}) \hat{w}_t^{(i)}, \;\; \hat{w}_t^{(i)} = \frac{\check{w}_t^{(i)}}{\sum_{j=1}^{N} \check{w}_t^{(j)}}, \;\; \check{w}_t^{(i)} \propto w_t^{(i)} = \frac{f(L_{1:t}^{(i)}|d_{1:t})}{q(L_{1:t}^{(i)}|d_{1:t})} \tag{3.6}$$

can be used to approximate the function g by sampling N particles, $L_t^{(i)}$ from the proposal distribution $q(L_{1:t}|d_{1:t})$. Since the integration constant is unknown, the importance weights $\{\check{w}_t^{(i)}\}_{i=1}^{N}$ are normalized.

Importance sampling can also be performed sequentially to avoid the evaluation of the entire expression at each period. SIS consists of writing the proposal distribution as a recursion, i.e.

$$q(L_{1:t}|d_{1:t}) = q(L_t|L_{1:t-1}, d_{1:t}) q(L_{1:t-1}|d_{1:t-1}). \tag{3.7}$$

Replacing Equations (3.3) and (3.7) in the importance weights of Equation

(3.6) as follows

$$\begin{aligned} \check{w}_t^{(i)} \propto \quad w_t^{(i)} &= \frac{f(d_t|L_t^{(i)}) f(L_t^{(i)}|L_{t-1}^{(i)}) f(L_{1:t-1}^{(i)}|d_{1:t-1})}{f(d_t|d_{1:t-1}) q(L_t^{(i)}|L_{1:t-1}^{(i)}, d_{1:t}) q(L_{1:t-1}^{(i)}|d_{1:t-1})} \\ &= \frac{f(d_t|L_t^{(i)}) f(L_t^{(i)}|L_{t-1}^{(i)})}{f(d_t|d_{1:t-1}) q(L_t^{(i)}|L_{1:t-1}^{(i)}, d_{1:t})} w_{t-1}^{(i)}, \end{aligned} \tag{3.8}$$

yields to the following recursion:

$$\check{w}_t^{(i)} \propto \; w_t^{(i)} = \; \tilde{w}_t^{(i)} w_{t-1}^{(i)}, \text{ where } \; \tilde{w}_t^{(i)} = \frac{f(d_t|L_t^{(i)}) f(L_t^{(i)}|L_{t-1}^{(i)})}{q(L_t^{(i)}|L_{1:t-1}^{(i)}, d_{1:t})},$$

and $\tilde{w}_t^{(i)}$ is called the incremental importance weight. Notice that $\tilde{w}_t^{(i)}$ depends on known densities: the observation and transition distributions (i.e. $f(d_t|L_t)$ and $f(L_t|L_{t-1})$) as well as the proposal distribution $q(L_t|L_{1:t-1}, d_{1:t})$. Finally, we define the Monte Carlo estimator of the posterior distribution of the latent variables as

$$f(L_{0:t}|d_{1:t};\theta) \approx \sum_{i=1}^{N} \delta(L_{1:t}) \hat{w}_t^{(i)}, \;\; \hat{w}_t^{(i)} = \frac{\check{w}_t^{(i)}}{\sum_{j=1}^{N} \check{w}_t^{(j)}}, \;\; \check{w}_t^{(i)} \propto w_t^{(i)} = \tilde{w}_{t-1}^{(i)} w_{t-1}^{(i)}, \tag{3.9}$$

where $\delta(\cdot)$ is the Dirac delta function.

Writing the proposal distribution as a recursion makes the evaluation of the importance weight easier for the entire period up to time t and allows us to only compute the incremental importance weight. However, it creates a problem referred to as weight degeneracy, characterized by the increasing variance of importance weights at each iteration. Algorithm 1 summarizes the SIS algorithm.

Algorithm 1 SIS Algorithm.

Initialization:

1. For $i = 1, ..., N$, set $w_0^{(i)} = \frac{1}{N}$.

Sequential Importance Sampling:

2. For $i = 1, ..., N$, sample from the proposal distribution
$$L_t^{(i)} \sim q(\ \cdot\ | d_t, L_{t-1}^{(i)}).$$
3. For $i = 1, ..., N$, compute the importance weights
$$\check{w}_t^{(i)} = \frac{f(d_t|L_t^{(i)}) f(L_t^{(i)}|L_{t-1}^{(i)})}{q(L_t^{(i)}|L_{1:t-1}^{(i)}, d_{1:t})} w_{t-1}^{(i)}.$$
4. Normalize the importance weights
$$\hat{w}_t^{(i)} = \frac{\check{w}_t^{(i)}}{\sum_{j=1}^N \check{w}_t^{(j)}}.$$

3.2 The Bootstrap Filter

Gordon et al. (1993) suggest adding a resampling step to the SIS algorithm to mitigate the weight degeneracy problem. There approach is often referred to as the bootstrap filter in the literature. The resampling step consists of selecting a new population of particles from the existing population according to their normalized importance weights. Concretely, the resampling step of the bootstrap filter consists of drawing with replacement a sample of N random variables from a multinomial distribution with weights $\{\hat{w}_t^{(i)}\}_{i=1}^N$. This means that the particles with the largest importance weights should be resampled multiple times. After the resampling step, the weights of each particle are equal: we set $\check{w}_t^{(i)} = \frac{1}{N}$ for $i = 1, ..., N$.

Another property of the bootstrap filter is that the transition density $f(L_t|L_{t-1})$ is used as the proposal distribution. The incremental importance

weights thus become

$$\begin{aligned}\tilde{w}_{t-1}^{(i)} &= \frac{f(d_t|L_t^{(i)})f(L_t^{(i)}|L_{t-1}^{(i)})}{q(L_t^{(i)}|L_{1:t-1}^{(i)}, d_{1:t})} \\ &= \frac{f(d_t|L_t^{(i)})f(L_t^{(i)}|L_{t-1}^{(i)})}{f(L_t^{(i)}|L_{1:t-1}^{(i)})} = f(d_t|L_t^{(i)}),\end{aligned}$$

which is the weighted contribution to the likelihood. Algorithm 2 describes the bootstrap filter algorithm.

Algorithm 2 Bootstrap Filter Algorithm.

Initialization:

1. For $i = 1, ..., N$, set $w_0^{(i)} = \frac{1}{N}$.

Sequential Importance Sampling:

2. For $i = 1, ..., N$, sample from the transition density
$$L_t^{(i)} \sim f(L_t|L_{t-1}^{(i)}).$$
3. For $i = 1, ..., N$, compute the importance weights
$$\check{w}_t^{(i)} = f(d_t|L_t^{(i)})w_{t-1}^{(i)} \ .$$
4. Normalize the importance weights
$$\hat{w}_t^{(i)} = \frac{\check{w}_t^{(i)}}{\sum_{j=1}^{N} \check{w}_t^{(j)}}.$$

Resampling:

5. Resample with replacement N particles with probabilities $\left\{\hat{w}_t^{(i)}\right\}_{i=1}^{N}$ and for $i = 1, ..., N$, set the new iteration weights $w_t^{(i)} = \frac{1}{N}$.

3.3 The Filtering Algorithm

In this section, we explain how to implement the bootstrap filter when applied to our model framework. We start the filtering process with a set of daily observed samples $d_{1:t} = \{d_1, ..., d_t\}$, where d_t is defined as the list of observables

available on day t, as explained in Chapter 2:

$$d_t = \big[Y_t, RV_{t-\Delta,t}, BV_{t-\Delta,t}, \boldsymbol{\sigma_t^{BS}}, \boldsymbol{ROV}_{t-\Delta,t}, \boldsymbol{RIVV}_{t-\Delta,t}\big],$$

where $\boldsymbol{\sigma_t^{BS}} = \{\sigma_{t,1}^{BS}, ..., \sigma_{t,n_t}^{BS}\}$ is a vector of n_t implied volatilities on day t. Similarly, $\boldsymbol{ROV}_{t-\Delta,t}$, $\boldsymbol{RIVV}_{t-\Delta,t}$ are vectors of realized option variances and realized implied volatility variances for the same n_t options on day t. In this study, we use daily aggregated measures, and we fix $\Delta = 1/252$.

Using Amaya et al. (2017)'s discretization and the aggregation scheme of Equation (2.1) and (2.2), N particles are simulated with M intraday steps. At the end of each day, we obtain a set of simulated particles $L_t = \{L_t^{(1)}, L_t^{(2)}, ...L_t^{(N)}\}$, where $L_t^{(i)}$ is defined as the list of theoretical measures detailed in Chapter 2:

$$L_t^{(i)} = \big[Y_t^{(i)}, V_t^{(i)}, \Delta QV_{t-\Delta,t}^{(i)}, \Delta I_{t-\Delta,t}^{(i)}, \boldsymbol{IV}_t^{(i)}(Y_t, V_t), \boldsymbol{\Delta OQV}_{t-\Delta,t}^{(i)}, \boldsymbol{\Delta IVQV}_{t-\Delta,t}^{(i)}\big].$$

The latent states $\Delta QV_{t-\Delta,t}^{(i)}$, $\Delta I_{t-\Delta,t}^{(i)}$, $\boldsymbol{\Delta OQV}_{t-\Delta,t}^{(i)}$ and $\boldsymbol{\Delta IVQV}_{t-\Delta,t}^{(i)}$ are aggregated at the end of the day and $\boldsymbol{IV}_t^{(i)}(Y_t, V_t)$ is computed using end of day information, i.e. $Y_t^{(i)}$ and $V_t^{(i)}$.

The time t contribution to the likelihood $f(d_t|L_t^{(i)})$ is computed as in Amaya et al. (2017). More precisely, we assume that the relative error on each observable is normally distributed with a mean of zero and fixed standard deviation.

The relative error of the realized variance ($RVRE$) is given by :

$$RVRE_t^{(i)} = \frac{\Delta QV_{t-h,t}^{(i)} - RV_{t-h,t}}{RV_{t-h,t}} \quad , \quad RVRE_t^{(i)} \sim \mathcal{N}(0, \varepsilon_1^2).$$

The relative error of the bipower variation ($BVRE$) is given by :

$$BVRE_t^{(i)} = \frac{\Delta I_{t-h,t}^{(i)} - BV_{t-h,t}}{BV_{t-h,t}} \quad , \quad BVRE_t^{(i)} \sim \mathcal{N}(0, \varepsilon_2^2).$$

The relative error of implied volatility ($IVRE$) for the k^{th} option is given by :

$$IVRE_{t,k}^{(i)} = \frac{IV_{t,k}^{(i)}(Y_t, V_t) - \sigma_{t,k}^{BS}}{\sigma_{t,k}^{BS}} \quad , \quad IVRE_{t,k}^{(i)} \sim \mathcal{N}(0, \varepsilon_3^2).$$

The relative error of realized option variance ($ROVRE$) for the k^{th} option is given by :

$$ROVRE_{t,k}^{(i)} = \frac{\Delta OQV_{t-h,t,k}^{(i)} - ROV_{t-h,t,k}}{\Delta ROV_{t-h,t,k}} \quad , \quad ROVRE_{t,k}^{(i)} \sim \mathcal{N}(0, \varepsilon_4^2).$$

The relative error of implied volatility realized variance ($RIVVRE$) for the k^{th} option is given by :

$$RIVVRE_{t,k}^{(i)} = \frac{\Delta IVQV_{t-h,t,k}^{(i)} - RIVV_{t-h,t,k}}{RIVV_{t-h,t,k}} \quad , \quad RIVVRE_{t,k}^{(i)} \sim \mathcal{N}(0, \varepsilon_5^2).$$

Using the assumptions described above, the time t weighted contribution to the likelihood function is defined by

$$\begin{aligned} f(d_t|L_t^{(i)}) = \; & \phi(Y_t; \mu_t, \sigma_t^2) \;\; \phi(RVRE_t^{(i)}; 0, \varepsilon_1^2) \;\; \phi(BVRE_t^{(i)}; 0, \varepsilon_2^2) \\ & \Big(\prod_{k=1}^{n_t} \phi(IVRE_{t,k}^{(i)}; 0, \varepsilon_3^2) \Big)^{1/n_t} \Big(\prod_{k=1}^{n_t} \phi(ROVRE_{t,k}^{(i)}; 0, \varepsilon_4^2) \Big)^{1/n_t} \\ & \Big(\prod_{k=1}^{n_t} \phi(RIVVRE_{t,k}^{(i)}; 0, \varepsilon_5^2) \Big)^{1/n_t}, \end{aligned} \tag{3.10}$$

where $\phi(x; \mu, \sigma^2)$ stands for the normal density function with mean μ_t and variance σ_t^2 (see Appendix C). Also note that, as a result of using a panel of n_t options, the weights of option-related observations are powered by one over n_t to avoid these sources of information to dominate the total time t weighted contribution to the likelihood function.

This concludes Chapter 3. We have now given and explained all the details necessary to implement a bootstrap filter applied to our model.

Chapter 4

Simulation-based Results

The goal of this chapter is to present the results of a series of simulation-based experiments. More precisely, we perform five experiments designed to answer the following questions:

1. How does the use of the realized implied volatility variance affect the particle filter's performance? (Experiment 1)

2. Is the information conveyed by the realized implied volatility variance redundant with that conveyed by the realized variance in the absence of jumps? (Experiment 2)

3. Is the information conveyed by the realized implied volatility variance redundant with that conveyed by the realized variance in the presence of jumps? (Experiment 3)

4. How robust is the realized implied volatility variance in approximating the implied volatility quadratic variation? (Experiment 4)

5. Is the information contained in the realized implied volatility variance and the realized option variance the same? (Experiment 5)

Throughout this chapter, we use the model parameters shown in Table 4.1. They correspond to the parameter estimates reported by Amaya et al. (2017)

when they include the realized option variance as a source of information in the filtering procedure. We set the standard deviations of error terms to the same arbitrary value (i.e. 0.05) in order to obtain comparable results from one filter to the next. The values attributed to the jump intensities $\lambda_{Y,0}$, $\lambda_{Y,1}$ and $\lambda_{V,0}$ consequently affects the number of intraday steps needed to effectively detect the occurrence of a jump. A higher jump intensity means that more jumps are occurring for a given period of time. Hence, taking smaller steps in the simulation (increasing the number of intraday steps M) will improve the detection of jumps.

Table 4.1: **Model parameters used in the simulation study.**

Log-price process		Variance process		SD of error terms	
η_Y	0.764	η_V	-0.230	ε_1	0.05
γ_Y	0.006	Γ_V	0.581	ε_2	0.05
r	0.01	κ	5.781	ε_3	0.05
q	0.01	θ	0.009	ε_4	0.05
ρ	-0.434	σ	0.894	ε_5	0.05
$\lambda_{Y,0}$	2.150	$\lambda_{V,0}$	7.650		
$\lambda_{Y,1}$	29.070	μ_V	0.021		
μ_Y	-0.005	V_0	0.012		
σ_Y	0.0160				
Y_0	log(1000)				

The first and second column display the parameters relating to Equation (2.1) and (2.2), respectively. They correspond to the parameter estimates obtained by Amaya et al. (2017) when they include the *ROV* as a source of information in the filtering procedure. The last column refers to the standard deviation (SD) of error terms defined in the assumptions embedded in Equation (3.10).

4.1 Experiment 1: Filter's Performance

The goal of the first experiment is to determine how the filter performs with different sources of information. To this end, we simulate 100 one-year paths of log-prices and variances from our model setup using discretization of 1,560 intraday steps and using the Broadie and Kaya (2006) drift-interpolation scheme of Van Haastrecht and Pelsser (2010).[5] Note that those one-year simulated

[5]When time steps are very small (i.e. $M = 1560$), the method for computing the integrated variance in the data generating process described in Appendix C.1 is not robust.

paths are used as observations in this experiment. For each simulated path, we compute the realized variance and the bipower variation according to Equations (2.5) and (2.7). We compute intraday option prices for a panel of ten call options with call-equivalent deltas of 0.2, 0.35, 0.5, 0.65, 0.8 and for maturities of 30 days and 90 days. From those option prices, we also compute the intraday Black and Scholes implied volatilities. We compute the realized option variance and the realized implied volatility variance by aggregating the intraday information according to Equations (2.11) and (2.16) respectively. To account for measurement errors, we add error terms to each observable as defined in Chapter 3.3. We then run the filter using 50,000 particles with different sets of observables and different data aggregation periods, i.e. $M = \{1, 2, 3, 5, 10\}$.

Let us first show a graphical example of the densities of the filtered instantaneous variance in Figure 4.1. Out of the 100 one-year simulated paths, we randomly pick a day for which a variance jump occurred and use Gaussian kernels to obtain a smooth density from the discrete sample of resampled variances. The aim of this exercise is to illustrate the filter's ability to detect variance jumps under different information sets. Ideally, we want a filter that produces a narrow density centered around the true value without bias, as this indicates that the mean of resampled variances is close to the true value and that their standard deviation is small. We observe that the density obtained from running the filter with the log-price (Y) alone is centered far away from the true value of the instantaneous variance represented by the vertical line, so far in fact that we only see a flat line at the bottom of the graph. Adding the realized variance (RV) to the information set drastically improves the estimation. Once the bipower variation (BV) is added to the set, the density peaks slightly higher and moves somewhat closer to the true variance. We observe similar improvements each time a subsequent observable is added to the set.

Hence, we use the following approximation: $\Delta I_{t,t+h} = \frac{1}{2}(V_t + V_{(t+h)^-})$.

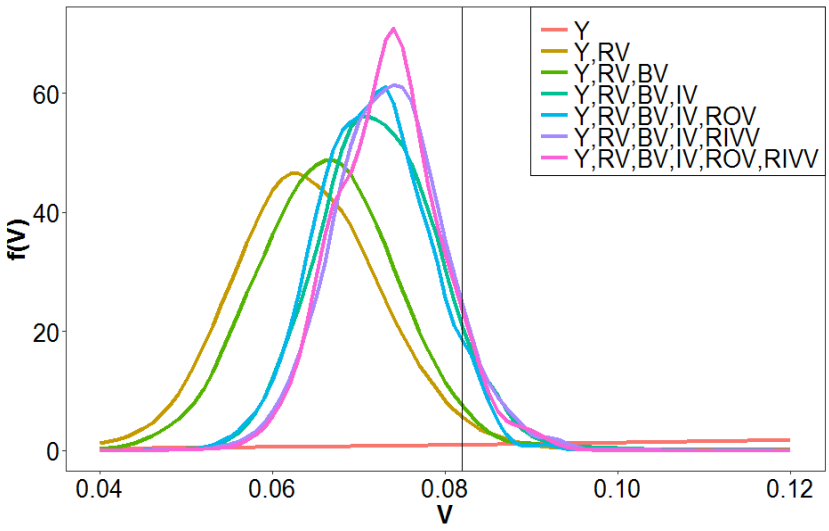

Figure 4.1: **Example of instantaneous variance using different information data sets.** On a random day where a variance jump occurred, each density is derived from the particles obtained after resampling. Gaussian kernels are used to smooth the densities. The true instantaneous variance is represented by the vertical line.

Continuing with a more detailed analysis, we now want to evaluate the filter's performance in the estimation of the five following latent states: the instantaneous variance, the integrated variance, the quadratric variation, the log-equity jumps and the variance jumps. Working with simulated observations is practical, because we know the true value of each latent state. Hence, we employ the root mean square error (RMSE) to evaluate and compare the performance of the filter with different sets of observables. At the end of each day and for each latent state aforementioned, we compute the RMSE of resampled particles. We end up with 25,200 RMSEs (252 days for 100 one-year paths) which are then averaged. We report our results in Table 4.2.

First, we note that the number of intraday steps M appears to have diminishing impact on the results as more observables are included in the information set. More specifically, across all panels (A to E), we observe that whether 5 or 10 intraday steps are used make no significant difference on the estimation results. Hence, increasing the number of intraday steps to $M = 10$ is unnecessary, which is good news because increasing the number a intraday steps requires a higher computational budget.

In Panel A, we observe that adding RV to the information set provides drastic improvement in the instantaneous variance estimation. Indeed, RMSEs are reduced by a factor as large as 26. This reduction is far greater than the one reported by Amaya et al. (2017). It could be rationalized by the use of a different parameter set, but it appears that the filter does not identify the latent states at all. Adding BV and then IV to the information set further reduces the RMSEs by about 10% and 60% respectively. Then, whether ROV, $RIVV$ or both measures are included in the filtering process, the RMSE is again marginally improved. However, the information set containing only ROV yields overall better results. Panel B and Panel C show similar patterns for the estimation of the quadratic variation and the integrated variance respectively: adding RV and BV improves the estimation result, but the inclusion of subsequent observables have no significant impact.

Panel D shows the RMSE for the size of log-equity jumps. When RV is combined with Y, the RMSE shrinks by a factor of 12. Similarly, when BV is added to the mix, the RMSE is reduced by a factor of almost three. However, when other observables are added to the information set, the improvements are less important. For instance, adding IV appears to make no difference at all. Then, incorporating ROV, $RIVV$ or both combined show additional reductions of 20%, 35% and 100% respectively.

Lastly, Panel E shows the RMSE for the size of variance jumps. The inclusion of IV, ROV and $RIVV$ appears to have more impact than what we report in Panel D. Indeed, adding the IV to the set of information reduces the RMSE by more than half, whereas adding the ROV reduces it by a factor as large as two. Additionally, the use of $RIVV$ instead of ROV yields even better results. This is also true when $RIVV$ is used in combination with ROV.

From this experiment, we conclude that ROV and $RIVV$ yield overall similar results. We observe instances where the inclusion ROV as a source of information results in the lowest RMSEs such as for the instantaneous variance. Other instances, namely for log-equity and variance jumps, $RIVV$ leads to better results, although marginally.

Table 4.2: **RMSE for different sets of observables and different data aggregation periods.**

Panel A: Instantaneous Variance					
	RMSE				
M	1	2	3	5	10
Y	165.8550	160.6421	159.0516	157.4972	154.5703
Y, RV	6.9519	6.7240	6.6574	6.0354	6.0704
Y, RV, BV	6.3110	6.0925	6.0322	5.5142	5.5066
Y, RV, BV, IV	4.2191	4.1633	4.1221	3.7783	3.8270
Y, RV, BV, IV, ROV	3.9290	3.8319	3.7939	3.4878	3.5454
$Y, RV, BV, IV, RIVV$	4.0167	4.0065	3.9669	3.5718	3.6197
$Y, RV, BV, IV, ROV, RIVV$	4.0298	3.9812	3.9418	3.5739	3.6097
Panel B: Quadratic Variation					
	RMSE				
M	1	2	3	5	10
Y	0.6679	0.6464	0.6400	0.6348	0.6229
Y, RV	0.0078	0.0084	0.0083	0.0072	0.0074
Y, RV, BV	0.0060	0.0063	0.0063	0.0054	0.0056
Y, RV, BV, IV	0.0058	0.0061	0.0061	0.0053	0.0054
Y, RV, BV, IV, ROV	0.0058	0.0059	0.0059	0.0051	0.0052
$Y, RV, BV, IV, RIVV$	0.0064	0.0065	0.0064	0.0056	0.0056
$Y, RV, BV, IV, ROV, RIVV$	0.0066	0.0063	0.0062	0.0054	0.0055
Panel C: Integrated Variance Increments					
	RMSE				
M	1	2	3	5	10
Y	0.6603	0.6399	0.6336	0.6275	0.6157
Y, RV	0.0099	0.0098	0.0097	0.0083	0.0090
Y, RV, BV	0.0059	0.0058	0.0057	0.0049	0.0051
Y, RV, BV, IV	0.0059	0.0055	0.0055	0.0047	0.0049
Y, RV, BV, IV, ROV	0.0056	0.0054	0.0053	0.0046	0.0048
$Y, RV, BV, IV, RIVV$	0.0061	0.0059	0.0058	0.0050	0.0052
$Y, RV, BV, IV, ROV, RIVV$	0.0061	0.0057	0.0057	0.0049	0.0051

Table continues.

Table 4.2 (Continued): **RMSE for different sets of observables and different data aggregation periods.**

Panel D: Log-Equity Jumps					
	RMSE				
M	1	2	3	5	10
Y	0.0636	0.0637	0.0631	0.0631	0.0624
Y, RV	0.0051	0.0046	0.0046	0.0039	0.0038
Y, RV, BV	0.0019	0.0016	0.0016	0.0014	0.0014
Y, RV, BV, IV	0.0020	0.0016	0.0015	0.0014	0.0013
Y, RV, BV, IV, ROV	0.0016	0.0012	0.0012	0.0011	0.0012
$Y, RV, BV, IV, RIVV$	0.0013	0.0014	0.0014	0.0011	0.0010
$Y, RV, BV, IV, ROV, RIVV$	0.0010	0.0011	0.0011	0.0009	0.0008
Panel E: Instantaneous Variance Jumps					
	RMSE				
M	1	2	3	5	10
Y	0.2444	0.2547	0.2522	0.2624	0.2601
Y, RV	0.1085	0.0996	0.0986	0.0843	0.0803
Y, RV, BV	0.1064	0.0983	0.0974	0.0776	0.0754
Y, RV, BV, IV	0.0366	0.0435	0.0431	0.0345	0.0306
Y, RV, BV, IV, ROV	0.0168	0.0129	0.0127	0.0134	0.0126
$Y, RV, BV, IV, RIVV$	0.0124	0.0116	0.0115	0.0099	0.0121
$Y, RV, BV, IV, ROV, RIVV$	0.0116	0.0114	0.0113	0.0117	0.0120

[1] Y stands for daily log-price, RV for realized variance, BV for bipower variation, IV for implied volatility,ROV for realized option variance, $RIVV$ for realized implied volatility variance, and M is number of intraday steps. Results are multiplied by 1000.

4.2 Experiment 2: Information in $RIVV$ in the Absence of Jumps

We now look at the information contained in the realized implied volatility variance in the absence of jumps. More precisely, we want to know if the realized implied volatility variance brings information not already provided by the realized variance. To this end, we simulate 1000 one-day path with 1,560 intraday steps for which there are no jumps. For each of these paths, we generate option parameters such that the call-equivalent delta of the option is uniformly distributed between values 0.1 and 0.9, and such that the initial variance of each path is also uniformly distributed between the values 0.01 and 0.1. The moneyness ratios are computed by numerical inversion from the randomly generated deltas and variances, and then, strike prices K are derived from the results. Using the same strike price for the entire path, we compute the option prices and corresponding implied volatilities throughout the day. Finally, from the intraday implied volatilities, the $RIVV$ for each path are computed. We then run the following regression:

$$RIVV_i = \gamma_1 \Big(\frac{\partial IV_i}{\partial y}\Big)^2 RV_i + \gamma_2 \Big(2\frac{\partial IV_i}{\partial y}\frac{\partial IV_i}{\partial v}\Big) RV_i + \gamma_3 \Big(\frac{\partial IV_i}{\partial v}\Big)^2 RV_i + \epsilon_i, \quad (4.1)$$

which is designed to resemble Equation (2.17) without the jump component. Indeed, if we replace RV and $RIVV$ in Equation (4.1) with their theoretical equivalents, the integrated variance and the implied volatility quadratic variation, we find an equation quite similar to Equation (2.17). Therefore, we should expect the regression coefficients γ_1 to be equal to 1, γ_2 to be equal to $\rho\sigma = -0.388$, and γ_3 to be equal to $\sigma^2 = 0.799$. As shown in Table 4.3, we find that none of these hypothesis can be rejected with a significance level of 5%. As a result, we conclude that $RIVV$ behaves as prescribed by Equation (2.17), assuming the implied volatility derivatives are kept constant. Moreover, the R-squared of 0.9818 is evidence of the quality of the fit between $RIVV$ and RV. From this experiment, we conclude that $RIVV$ does not convey information not already conveyed by RV in the absence of jumps.

Table 4.3: **Information content of $RIVV$ in the absence of jumps.**

	Coefficients	Standard Error
γ_1	1.5215	0.2401
γ_2	-0.3651	0.0166
γ_3	0.7117	0.0057
R^2	0.9818	

The regression of Equation (4.1) is applied using beginning-of-the-day information to compute the implied volatility derivatives. Results are obtained with a sample of 1000 observations. None of the values are statistically different (at a significance level of 5%) from their theoretical values (i.e. 1, -0.388, and 0.799 respectively).

4.3 Experiment 3: Information in $RIVV$ in the presence of jumps

We now look at the information conveyed by $RIVV$ in the presence of jumps. To this end, we simulate 1000 one-day path with a frequency of 1,560 steps in a day. For each path, we compute intraday option prices and corresponding implied volatilities for a panel of five call options with call-equivalent deltas of 0.20, 0.35, 0.50, 0.65, 0.80. We then compute $RIVV$ and ROV, and we run the following regressions:

$$\log(RIVV_i) = \beta_0 + \beta_1 \log(RV_i) + \beta_2 \Delta N_{Y,i} + \beta_3 \Delta N_{V,i} + \epsilon_i \tag{4.2}$$

$$\log(ROV_i) = \beta_0 + \beta_1 \log(RV_i) + \beta_2 \Delta N_{Y,i} + \beta_3 \Delta N_{V,i} + \epsilon_i, , \tag{4.3}$$

where $\Delta N_{Y,i}$ and $\Delta N_{V,i}$ represent the number of log-equity and variance jumps respectively. Again, these models are inspired by Equations (2.17) and (2.12). In section 2.2.2, we show that both the implied volatility quadratic variation and option quadratic variation can be separated in three components: the integral which is driven by the integrated variance, the log-equity jump component and similarly the variance jump component. Hence, in the regressions above, the integral part is represented by RV and the jump components are represented by $\Delta N_{Y,i}$ and $\Delta N_{V,i}$.

We run two versions of the each regression. First, we force β_3 to 0 (Regression (1)), and second, we allowed β_3 to differ from 0 (Regression (2)). This will confirm whether variance jumps is a driving factor of $RIVV$. Assuming $RIVV$ brings valuable information about variance jumps, we would expect the R-squareds to be higher in Regression (2). Table 4.4 reports the results of both regressions. As expected, when variance jumps are added to the regression, the R-squared are almost doubled. The results of the regressions with $RIVV$ contrast those obtained with ROV in two ways: first, the R-squareds are smaller in Regression (1). This is indicative that ROV responds more to log-equity jumps than does $RIVV$. Second, the R-squareds do not vary much from one moneyness to the other, whereas with ROV, the quality of the fit improves significantly for in-the-money options.

Additionally, we want to point out that the results of the regression with ROV are consistent with those reported by Amaya et al. (2017). However, we only show results with regards to call options, as using put options yields identical results for the regression with $RIVV$. This is due to the fact that the Black-Scholes implied volatility is the same for a call and put option with otherwise identical features and therefore results in identical $RIVV$, at least in a simulation context.

From this experiment, we conclude that $RIVV$ conveys information not provided by RV in the presence of jumps. This is an indication that $RIVV$ could be used as an alternative to ROV as a jump detection tool.

Table 4.4: **Information content of $RIVV$ in the presence of jumps.**

		Regression with $RIVV$				Regression with ROV			
		Regression (1)		Regression (2)		Regression (1)		Regression (2)	
Δ^e		Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
0.20	β_0	1.665	(0.345)	-0.640	(0.218)	14.185	(0.070)	12.149	(0.057)
	β_1	1.016	(0.034)	0.792	(0.022)	1.342	(0.007)	1.142	(0.006)
	β_2	-1.333	(0.121)	-1.055	(0.074)	-0.891	(0.027)	-0.578	(0.021)
	β_3			1.598	(0.039)			1.510	(0.012)
	R^2	0.467		0.783		0.594		0.750	
0.35	β_0	0.497	(0.350)	-1.879	(0.215)	12.621	(0.059)	10.936	(0.049)
	β_1	0.876	(0.035)	0.644	(0.021)	1.077	(0.006)	0.912	(0.005)
	β_2	-1.249	(0.123)	-0.964	(0.073)	-0.376	(0.022)	-0.118	(0.018)
	β_3			1.648	(0.039)			1.250	(0.010)
	R^2	0.386		0.783		0.581		0.736	
0.50	β_0	0.459	(0.356)	-1.944	(0.222)	11.891	(0.044)	10.600	(0.035)
	β_1	0.871	(0.036)	0.637	(0.022)	0.936	(0.004)	0.809	(0.004)
	β_2	-1.271	(0.125)	-0.982	(0.076)	-0.020	(0.017)	0.178	(0.013)
	β_3			1.666	(0.040)			0.957	(0.007)
	R^2	0.376		0.775		0.673		0.803	
0.65	β_0	0.930	(0.352)	-1.470	(0.215)	12.141	(0.031)	11.228	(0.025)
	β_1	0.928	(0.350)	0.695	(0.021)	0.920	(0.003)	0.831	(0.003)
	β_2	-1.163	(0.124)	-0.875	(0.073)	0.118	(0.012)	0.258	(0.009)
	β_3			1.664	(0.039)			0.677	(0.005)
	R^2	0.413		0.796		0.803		0.880	
0.80	β_0	0.775	(0.349)	-1.628	(0.209)	12.464	(0.021)	11.887	(0.018)
	β_1	0.921	(0.035)	0.687	(0.021)	0.923	(0.002)	0.867	(0.002)
	β_2	-1.005	(0.123)	-0.715	(0.071)	0.166	(0.008)	0.254	(0.006)
	β_3			1.667	(0.037)			0.428	(0.004)
	R^2	0.417		0.805		0.901		0.940	

The regressions from Equations (4.2) and (4.3) are applied. Two versions of both equations are reported: In Regression (1), $\beta_3 = 0$, and in Regression (2), β_3 is allowed to differ from 0. Standard errors (SE) are in parentheses, and all coefficients in this table are statistically significant with a significance level of 5%.

4.4 Experiment 4: Robustness of $RIVV$ as an Approximation of $\Delta IVQV$

The next experiment tests the robustness of $RIVV$ as an approximation of the implied volatility quadratic variation with different values of intraday steps M. Using the parameters given in Table 4.1, we simulate 500 one-day paths with 1,560 intraday steps. For each of these paths, we compute intraday option prices and corresponding implied volatilities for a panel of five call options with call-equivalent deltas of 0.20, 0.35, 0.50, 0.65, 0.80, and with maturities of 30 and 90 days. We then compute $RIVV$ along each path. We also compute simultaneously its theoretical equivalent, $\Delta IVQV$, using the time discretization of Equation (2.17) with different values of intraday steps $M = \{1, 2, 3, 5, 10, 1560\}$.[6] The goal is to compare $RIVV$ obtained via high frequency implied volatilities with the lower frequency approximations of $\Delta IVQV$ by using RMSE and relative RMSE. We also run a regression of $\Delta IVQV$ on $RIVV$ and report the R-squareds as a measure of the quality of the fit.

Panel A of Table 4.5 reports the RMSE on days without jumps and on days with jumps. The results are similar to those reported by Amaya et al. (2017) in three ways: first, the RMSEs are smaller on days without jumps; second, deep OTM options yield better approximations, and third, as the frequency increases, the approximation improves. The relative RMSEs reported in Panel B paint a similar picture. However, at low frequencies and for OTM options, the relative RMSE are not systematically higher on days with jumps.

The R-squared values resulting from the regressions reported in Panel C are also similar to those obtained by Amaya et al. (2017) with ROV. One notable difference however is the quality of the fit at the lowest frequency (i.e. $M = 1$). Indeed, it is twice as good with $RIVV$ as with ROV. To visualize the result of this experiment, the $RIVV$ and corresponding $\Delta IVQV$ obtained with different frequencies are plotted in Figure 4.2. The better the approximation, the closer the dots should be to the diagonal line. Notice also that the scale of the graphs in the top panel is much smaller than in the bottom panel due to

[6] The time discretization of Equation (2.17) is provided is Appendix C.2.

the jump component.

The goal of Experiment 4 was to justify the use of $RIVV$ as an approximation for $\Delta IVQV$. From this experiment, we observe that the approximation of $\Delta IVQV$ with $RIVV$ improves as the number of intraday steps used to compute $\Delta IVQV$ increases, as well as when jumps occur. This means that as the number of intraday steps M increases to 1560, the two measures become almost equivalent. Therefore, we conclude that the approximation is robust. However, as we have shown in Experiment 1, this does not mean that increasing M beyond ten steps a day would improve the estimation of latent states. This is also a consequence of the chosen value for the jump intensity (i.e. 7.650 for the variance). Indeed, perhaps if we chose a much higher jump intensity, a higher number of intraday steps would be needed to effectively detect jumps.

Table 4.5: **RMSE, relative RMSE and R-squareds for $\Delta IVQV$ approximation across 500 days.**

Panel A. RMSE

	Days without jumps					Days with jumps				
Δ^e	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	0.027	0.041	0.040	0.035	0.038	0.574	1.022	1.527	1.256	1.774
2	0.015	0.021	0.020	0.018	0.019	0.393	0.723	1.007	0.815	2.850
3	0.012	0.016	0.015	0.014	0.014	0.193	0.429	0.780	0.619	3.099
5	0.009	0.012	0.012	0.010	0.010	0.169	0.333	0.445	0.407	1.889
10	0.008	0.010	0.010	0.009	0.008	0.117	0.242	0.341	0.301	0.845
1560	0.007	0.009	0.009	0.008	0.007	0.025	0.039	0.05	0.045	0.059

Panel B. Relative RMSE

	Days without jumps					Days with jumps				
Δ^e	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	40.910	47.789	53.023	57.125	70.087	19.340	32.529	82.811	56.208	101.545
2	18.779	22.590	25.318	27.133	32.667	10.626	29.359	48.444	33.198	98.953
3	14.808	16.929	18.321	19.611	23.541	2.266	15.797	44.691	38.830	97.984
5	8.090	9.209	9.908	10.833	13.059	0.722	8.548	19.507	19.995	57.382
10	3.754	4.479	4.828	5.334	6.346	0.281	6.848	10.361	9.718	30.312
1560	0.024	0.062	0.197	0.170	0.133	0.185	0.406	0.815	0.875	1.186

Panel C. Regression R-squareds for $\Delta IVQV$

	Days without jumps					Days with jumps				
Δ^e	0.2	0.35	0.5	0.65	0.8	0.2	0.35	0.5	0.65	0.8
1	0.893	0.921	0.919	0.913	0.888	0.961	0.950	0.940	0.943	0.909
2	0.917	0.929	0.930	0.925	0.914	0.983	0.976	0.971	0.974	0.786
3	0.928	0,939	0.940	0.935	0.928	0.995	0.991	0.985	0.988	0.812
5	0.938	0.946	0.947	0.942	0.937	0.996	0.994	0.994	0.994	0.889
10	0.946	0.953	0.953	0.949	0.946	0.998	0.997	0.996	0.997	0.974
1560	0.951	0.957	0.957	0.953	0.951	1.000	1.000	1.000	1.000	1.000

The real $RIVV$ value is computed using implied volatilities with 1,560 intraday steps. The difference between $\Delta IVQV$ and $RIVV$ correspond to the error. $\Delta^e = \{0.2, 0.35, 0.5, 0.65, 0.8\}$ are the call-equivalent deltas considered. All options are call options with a time-to-maturity of 30 days. $\Delta IVQV$ are computed with six different frequencies: $M = \{1, 2, 3, 5, 10, 1560\}$. Results in Panel A. and B. were multiplied by 1000. The R-squareds for $\Delta IVQV$ are computed by running the following regression: $\Delta IVQV = \beta_0 + \beta_1 RIVV_i + \epsilon_i$.

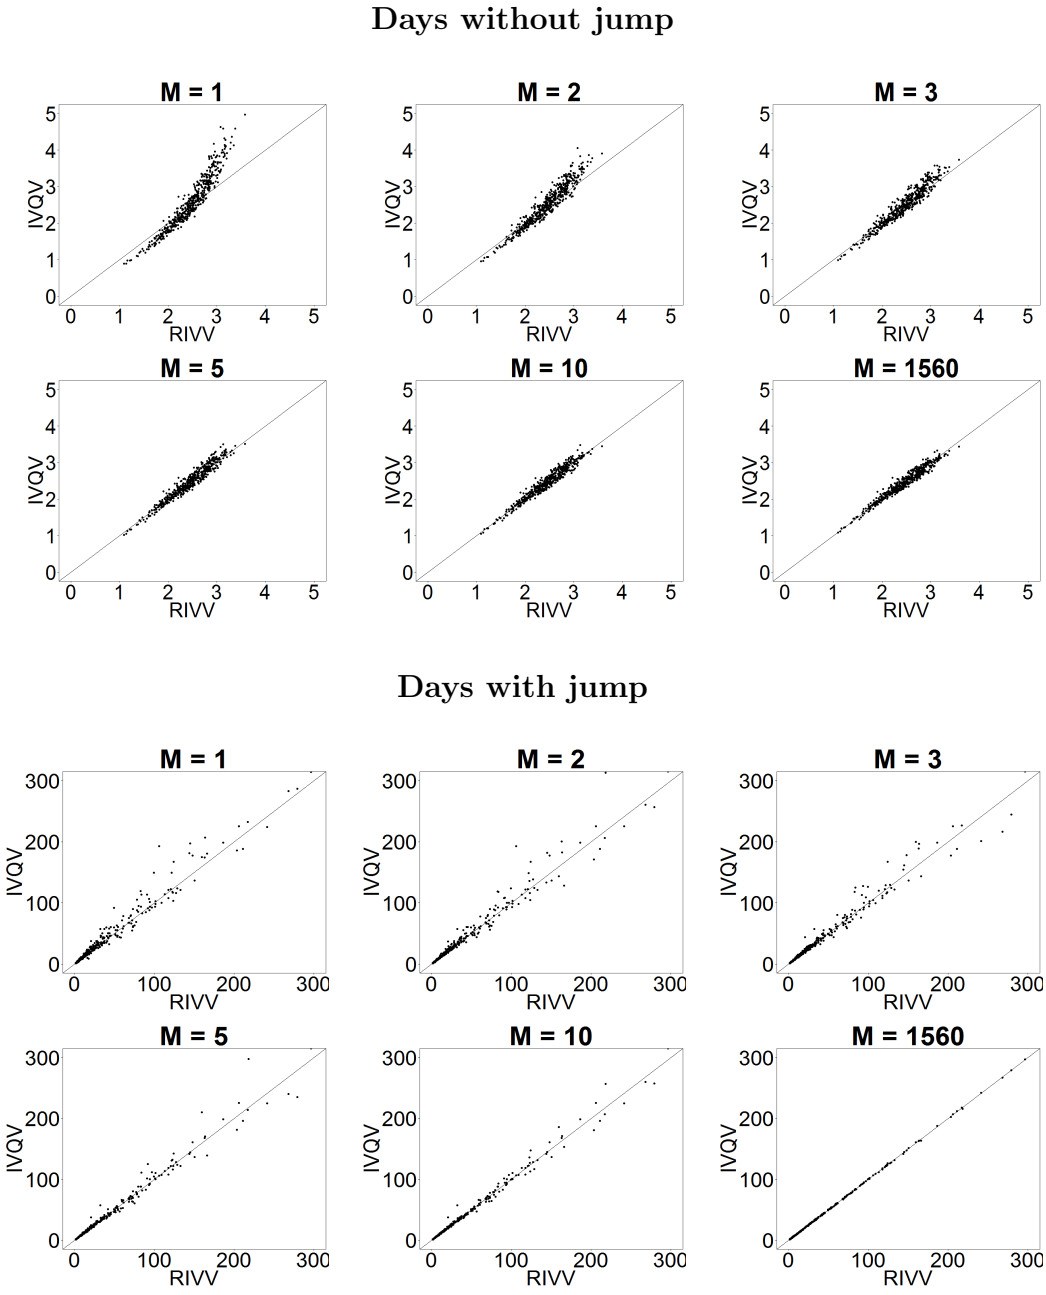

Figure 4.2: **Realized Implied Volatility Variance against implied volatility quadratic variation increments for a call-equivalent delta of 0.50 and a time-to-maturity of 30 days.** Top and bottom six figures correspond to the the $RIVV$ and $\Delta IVQV$ on days without jumps and days with jumps, respectively. $RIVV$ are computed using 1560 intraday steps. $\Delta IVQV$ are computed with intraday frequencies of $M = \{1, 2, 3, 5, 10, 1560\}$. Call options are used with parameters given in Table 4.1.

4.5 Experiment 5: Information Content of $RIVV$ vs ROV

The goal of Experiment 5 is to determine whether the realized implied volatility variance and the realized option variance convey the same information. It is important to take into account that maturity and moneyness are two driving factors of $RIVV$ and ROV. Hence, we perform a series of regressions to look at the information shared between the two measures in silos. In other words, we look at the information shared between the two measures separately for two differents maturities, 30 and 90 days, and for three different groups of moneyness OTM, ATM, ITM. Lastly, we look at the information shared under two different scenarios: on days with jumps and on days without jumps. In total, we run 12 separate regressions.

For each of the two scenarios, we simulate 1000 one-day paths with 1,560 intraday steps. That is, for the scenario "days without jumps", we have a 1000 one-day path for which no jump has occurred, and for the "days with jump" scenario, we have 1000 one-day paths for which either a jump in return or in variance occurred. For each of these paths, we generate option parameters such that the call-equivalent delta is uniformly distributed between 0.1 and 0.4 for the regression with out-of-the-money options, 0.4 and 0.6 for the regression at-the-money options and 0.6 and 0.9 for the regression with in-the-money options. The initial variance of each path is uniformly distributed between values 0.01 and 0.1. The moneyness ratios are computed by numerical inversion from the randomly generated deltas and variances, and then, strike prices K are derived from the results. Using the same strike price for the entire path, we compute options prices and corresponding implied volatilities throughout the day. Finally, we compute $RIVV$ and ROV for each path and run the following regression:

$$RIVV_i = \beta_0 + \beta_1 ROV_i + \epsilon_i \tag{4.4}$$

The results are reported in Table 4.6. Panel A displays the results for

out-of-the-money options, Panel B for at-the-money options and Panel C for in-the-money options. Although the model described in Equation (4.4) only accounts for the linear relationship between $RIVV$ and ROV, it still gives us some insight about the information content they share. Assuming they do share the same information content, ROV should be a good predictor of $RIVV$.

Overall, the R-squared values indicate that the quality of the fit for OTM options is generally better than for ATM and ITM options. This is not surprising. As we explain in section 2.2.2, both $RIVV$ and ROV of OTM options are solely driven by the jump components. If there is no jump, then the measure is close to zero.

Across all three Panels, the quality of the fit is generally better for 90-day options than 30-day options. This also does not surprise us. It is due to the fact that $RIVV$, contrary to ROV is influenced by the maturity of the option. The $RIVV$ surface increases exponentially as the maturity of the option approaches. This is true empirically as we demonstrate in Chapter 5, as well as theoretically as we explained in section 2.2.2.

From this experiment, we conclude that the information contained in $RIVV$ and ROV is the same when the maturity effect is tuned out.

This completes our five simulation-based experiments. To summarize our findings about the realized implied volatility variance, we found that its inclusion to the filtering procedure affects the particle filter's performance the same way than the realized option variance. We found that the information it conveys is redundant with the realized variance in the absence of jumps, but contains incremental information in the presence of jumps. We also showed that it is a good model-free approximation of the implied volatility quadratic variation. Finally, we found that its information content is the same than that of the realized option variance.

Table 4.6: **Information content of *RIVV* vs *ROV***

Panel A. Out-of-the-money options								
	Days without jumps				Days with jumps			
Maturity	30-day		90-day		30-day		90-day	
	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
β_0	1.507	(0.017)	0.800	(0.016)	1.741	(0.072)	1.851	(0.034)
β_1	0.156	(0.004)	0.150	(0.007)	0.494	(0.008)	0.183	(0.001)
R^2	0.635		0.889		0.810		0.945	
Panel B. At-the-money options								
	Days without jumps				Days with jumps			
Maturity	30-day		90-day		30-day		90-day	
	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
β_0	1.666	(0.041)	1.332	(0.039)	-0.748	(1.254)	-3.746	(0.048)
β_1	0.616	(0.003)	0.591	(0.002)	0.347	(0.008)	0.168	(0.001)
R^2	0.278		0.449		0.648		0.942	
Panel C. In-the-money options								
	Days without jumps				Days with jumps			
Maturity	30-day		90-day		30-day		90-day	
	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
β_0	1.355	(0.004)	1.290	(0.017)	6.529	(1.868)	-2.758	(1.071)
β_1	0.334	(0.001)	0.380	(0.006)	0.202	(0.009)	0.161	(0.003)
R^2	0.398		0.032		0.303		0.701	

Coeff. stands for Coefficient. Standard deviation of error terms (SE) are in parentheses. Results were multiplied by 10,000. All coefficients reported are significant with a significance level of 5%. The call-equivalent deltas ranges are $\Delta^e = [0.1, 0.4]$ for OTM options, $\Delta^e = [0.4, 0.6]$ for ATM options and $\Delta^e = [0.1, 0.4]$ for ITM option. The R-squareds are computed by running the regression of Equation (4.4)

Chapter 5

Empirical Results

In this chapter, we look at the proposed realized implied volatility variance empirically. We start with a description of our dataset. We then look at the properties of the realized implied volatility variance, and conclude with a principal component analysis of the $RIVV$ surface.

5.1 Data

The data used in this study cover the period starting July 2004 up to December 2012. During those years, there have been moments of low uncertainty as well as periods of high financial stress. Hence, it makes it a good period to study the behavior of $RIVV$ during various market environments.

We construct time series of $RIVV$ using tick-by-tick Level I quote data provided by Tick Data of European options on the S&P 500 Index. Essentially, we start with a basket of options each day. For each of these options, we have a total of 390 one minute bid and ask prices covering the period during which markets are opened (9AM to 4PM). To mitigate bid-ask spread variation noise, we compute midquotes (Amaya et al., 2017). For each option midquote we compute the Black and Scholes implied volatility. The risk-free rate is bootstrapped using Zero Coupon yield curve data from Option Metrics Database.[7]

[7]Option Metrics database is available through Wharton Research Data services.

The S&P 500 Index dividend yield is also provided by Option Metrics. From one minute Black and Scholes implied volatilities, we compute daily annualized $RIVV$ for each option in our daily basket according to Equation (2.16) using Zhang et al. (2005) methodology. The authors explain that blindly computing realized variance from high frequency data can result in large observation errors. A common fix has been to sample less frequently, throwing away most of the data which can contain valuable information. Zhang et al. (2005) propose a new way of mitigating observation error by constructing five grids with five minute prices, computing the realized variance over each grid, and then taking the average.

We filtered the database with a few common exclusion rules: options with bid prices lower than zero, with condition codes or eligible for automatic execution were excluded. Additionally, we ruled out options that did not meet the Bakshi et al. (1997) arbitrage free rule. Lastly, we only included options with positive volume. Our final dataset contains 686,305 options.

For our empirical study, we focus on ATM and OTM options and excluded those for which $RIVV$ fell in the top 0.1%. Figure 5.1 illustrates a few properties of our dataset. The RV and BV are computed using Zhang et al. (2005) methodology with high frequency E-mini futures contracts. For the IV, ROV and $RIVV$, we construct the time series by selecting a 30-day ATM option on each date (i.e. $\Delta^e = 0.50$).

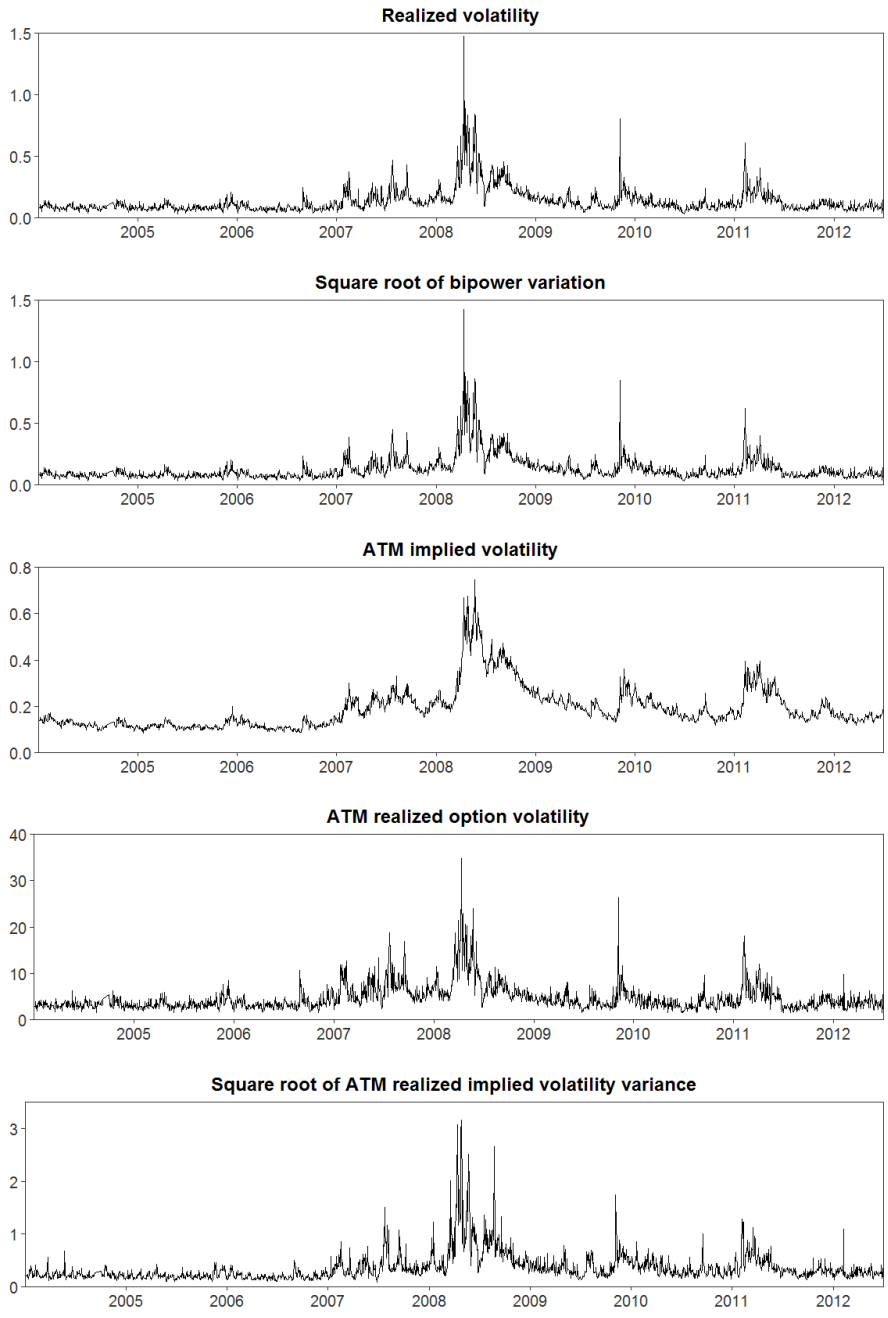

Figure 5.1: **Realized volaltility, Square root bipower variation, Realized option volatility and Realized implied volatility variance** Realized volatility and square root bipower variation are computed from intraday S&P 500 futures prices. Realized option volatility and realized implied volatility variance are derived from tick-by-tick Level I quote data from Tick Data. Midquotes and corresponding Black Scholes implied volatilities are computed from one minute option bid and ask prices. Zhang et al. (2005) methodology is used to compute microstructure-noise robust daily estimates. Each time series covers the period starting July 2004 up to December 2012.

5.2 Empirical Properties of $RIVV$

To get a clear picture of the empirical features of $RIVV$, we construct time series of its characteristics and compute a correlation matrix with other properties of our data, namely RV, $\max(RV - BV, 0)$ and ROV. Results are shown in Table 5.1. The level corresponds to the realized implied volatility variance of 30-day ATM options ($\Delta^e = 0.50$). The term structure (TS) is obtained by taking the difference between the realized implied volatility variance of 90-day ATM options and 30-day ATM options. The skew corresponds to the difference between the realized implied volatility variance of 30-day OTM put options ($\Delta^e = 0.90$) and 30-day OTM call options ($\Delta^e = 0.10$). Finally, the skew term structure is obtained by taking the difference between the 90-day skew and the 30-day skew. Figure 5.2 illustrates the time series of the properties we just defined along with RV and $\max(RV - BV, 0)$.

Table 5.1: **Correlation matrix for realized implied volatility variance, realized variance and realized jump variation.**

	$RIVV$	TS	Skew	Skew TS	RV	$\max(RV-BV,0)$	ROV
$RIVV$	1.000						
TS	-0.943	1.000					
Skew	0.053	-0.041	1.000				
Skew TS	0.021	-0.024	-0.915	1.000			
RV	0.844	-0.678	0.063	0.043	1.000		
$\max(RV-BV,0)$	0.596	-0.474	0.063	-0.040	0.703	1.000	
ROV	0.763	-0.638	0.038	0.094	0.877	0.608	1.000

$RIVV$ represents the realized implied volatility variance for ATM 30-day options. Term structure corresponds to the difference between 90-day and 30-day options. Skew is the difference between OTM 30-day put and OTM 30-day call. Skew TS is the difference between 90-day skew and 30-day skew. Realized variances and realized jump variations are multiplied by 1000. Zhang et al. (2005) methodology is used to compute microstructure-noise robust daily estimates.

We want to highlight three important results. First, the correlation between the level of ATM 30-day options and the realized variance, the realized option variance and the realized jump variance ($\max(RV - BV, 0)$) are high (i.e. 0.844, 0.763 and 0.596 respectively) indicating that $RIVV$ of 30-day

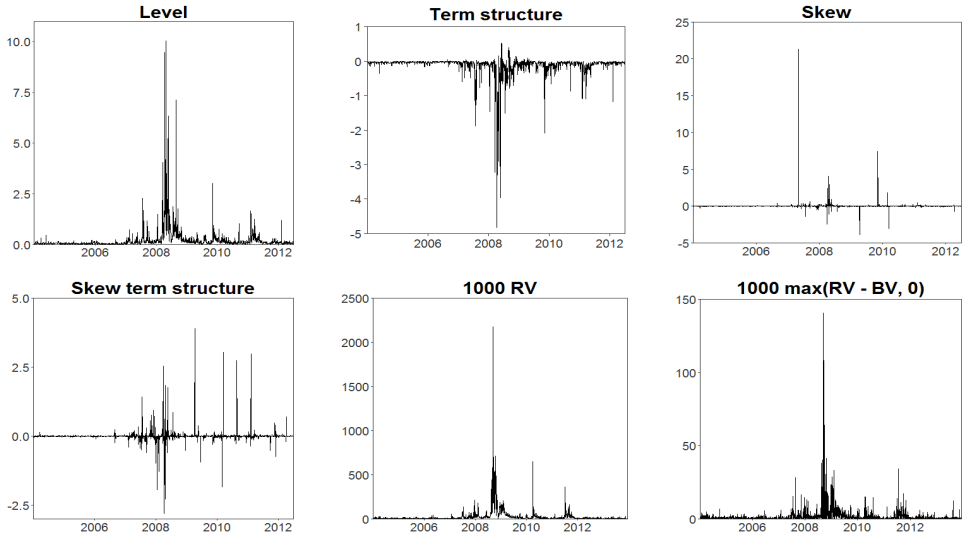

Figure 5.2: **Realized implied volatility variance characteristics, realized variance and realized jump variation.** Level represents the realized implied volatility variance for ATM 30-day options. Term structure corresponds to the difference between 90-day and 30-day options. Skew is the difference between OTM 30-day put and OTM 30-day call. Skew TS is the difference between 90-day skew and 30-day skew. Realized variances and realized jump variations are multiplied by 1000. Zhang et al. (2005) methodology is used to compute microstructure-noise robust daily estimates.

ATM options alone does not provide much new information. Second, there is a strong negative correlation between the level and the term structure, as well as between the skew and skew term structure. This highlights the fact that $RIVV$ increases exponentially as the expiration of the option approaches. Third, the correlation between the skew and the realized variance, the realized option variance and the realized jump variance are all near zero. The results are similar with the skew term structure. This is an indication that the cross-section of $RIVV$ contains valuable information not already provided by RV, ROV and $\max(RV - BV, 0)$.

5.2.1 The $RIVV$ Surface

Keeping in mind that the skew and skew term structure seem to reveal additional information, we now construct a sequence of surfaces as a function of the call-equivalent delta and business days-to-maturity. Each surface is ob-

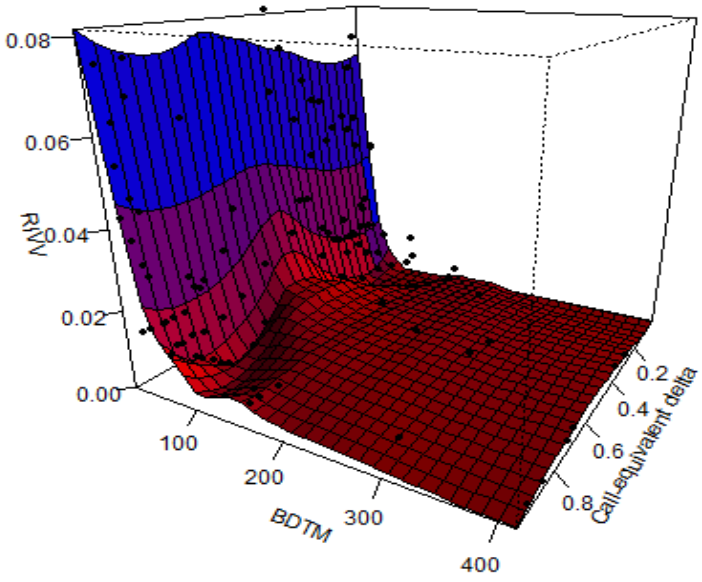

Figure 5.3: **Example of realized implied volatility variance and fitted surface as a function of business days-to-maturity and Black-Scholes call-equivalent delta.** The surface is obtained with a local smoothing regression method.*RIVV* is computed according to Equation (2.16). The data from July 6 2004 are used in this example. Zhang et al. (2005) methodology is used to compute microstructure-noise robust daily estimates.

tained using all *RIVV* values available on a given day and performing a locally weighted scatter plot smoothing across the two dimensions.[8] An example of such surface is illustrated in Figure 5.3; it is obtained with *RIVV* values on October 25th 2004. On this date specifically, we observe that the surface is flat for long dated options across all moneynesses. Moving along the business days-to-maturity (BDTM) axis towards zero, *RIVV* reaches a low for OTM options (i.e. $\Delta = 0.10$ for call options and $\Delta = 0.90$ for put options) before increasing exponentially as maturity approaches. This behavior is also validated from a theoretical perspective, as we saw in Chapter 2.

[8]We used the LOESS regression procedure in R. The algorithm performs a local regression using weighted linear least squares with second degree polynomial.

5.2.2 Principal Component Analysis of $RIVV$

Next, we perform a principal component analysis (PCA). The goal is to find out whether the $RIVV$ surface can be summarized by one principal component (PC) or if subsequent PCs are necessary to explain the surface. To this end, we interpolate 18 equally spaced points on each surface with nine different call-equivalent deltas ranging from 0.1 to 0.9, and for two different business days-to-maturity, 30 and 90 BDTM.

Figure 5.4 shows the six first PCs of the realized implied volatility variance surface. Notice that the first PC graph (top-left) resembles the level characteristic, RV and $\max(RV - BV, 0)$ graphs of Figure 5.2 (top-left, bottom-middle and bottom-right). This is indicative of the commonality between the variables. The proportion of the variance explained by the first PC is 71.55%, and as much as 13.85% by the second PC, followed by 7.55%, 3.52%, 1.10% and 0.95% for subsequent PCs. The $RIVV$ has a strong level effect, but is also influenced by other factors to a great extent.

Although similar, these results contrast with those reported by Amaya et al. (2017) with regards to ROV. They report that close to 95% of the ROV surface can be summarized by the first PC. One possible explanation could be the fact that the number of business days-to-maturity is a major source of variance among $RIVV$ values, contrary to the ROV surface.

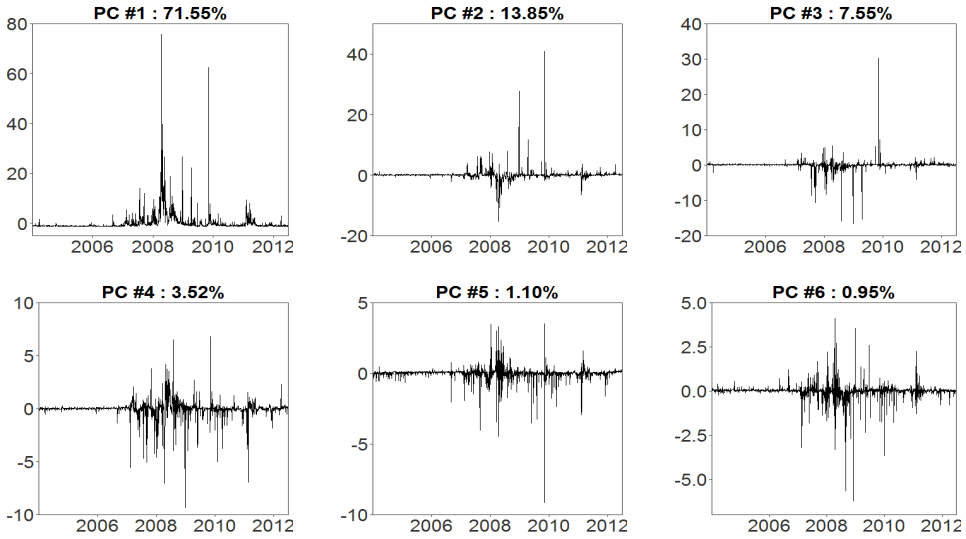

Figure 5.4: **Principal components of the realized implied volatility variance surface.** This figure shows the first six principal components of the S&P 500 interpolated from the realized implied volatility variance surface for the period starting in July 2004 and ending in December 2012. On each day of this period, 18 values of $RIVV$ are interpolated with nine different call-equivalent deltas ($\Delta^e = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$) and two different maturities (30 and 90 days) from the locally smoothed quadratic regression surface obtained with the loess procedure.

Continuing our investigation, we want to quantify the degree of commonality between the $RIVV$ surface characteristics, as well as other properties of the S&P 500 Index, and the principal components. To this end, we run six regressions; one on each of the characteristics (i.e. Level, TS, Skew, Skew TS, RV, $\max(RV - BV, 0)$) described earlier. The regressions are performed according to the following equation:

$$\text{Char}_t = \beta_0 + \beta_1 PC_{1,t} + \beta_2 PC_{2,t} + \beta_3 PC_{3,t} + \beta_4 PC_{4,t} + \beta_5 PC_{5,t} + \beta_6 PC_{6,t} + \varepsilon_t,$$

where Char_t is the characteristic at time t and $PC_{n,t}$ for $n = \{1, 2, 3, 4, 5, 6\}$ correspond to the principal components at time t. Residuals are adjusted for autocorrelation and heteroskedasticity (Newey and West, 1987).

Table 5.2 displays regression coefficients, standard errors, R-squares, and autocorrelations in the residuals. From these results, we observe that the level, the term structure, the realized variance and $\max(RV - BV, 0)$ present

a high degree of commonality with the the first and second principal components as shown by the statistically significant regression coefficients. Amaya et al. (2017) yield similar conclusion with ROV. However, after adjustment for autocorrelation and heterosckedasticity, the skew and skew term structure cannot be explained by the principal components as evidenced by both the low R-squared values of 0.027 and 0.240, and the fact that almost none of the coefficients appear to be statistically significant. This result contrasts those obtained by Amaya et al. (2017). They report that these characteristics can be explained to some extent by the first principal component of the ROV surface.

To summarize our empirical study, the correlation between the $RIVV$ skew and RV, $\max(RV - BV, 0)$ and ROV is closed to zero, showing some sign that non-redundant information may be embedded in the $RIVV$ skew. We reached similar conclusion with regards to the $RIVV$ skew term structure. Moreover, the principal component analysis reveals that the $RIVV$ surface cannot be summarized by one PC. It also appears that the factors driving the $RIVV$ level and term structure differ from those driving the skew and skew term structure. These findings suggest that multiple options with different maturities and moneynesses should be used in a parametric study to capture the intricacies of the $RIVV$ surface and the information it has to offer. Similar intuition with respect to the ROV surface is given by Amaya et al. (2017).

Table 5.2: **Principal component analysis and regressions of realized implied volatility variance, realized variance and realized jump variation.**

	Level	TS	Skew	Skew TS	RV	$\max(RV-BV,0)$
β_0	**0.201**	**-0.121**	0.002	0.019	**27.976**	**1.953**
	(0.007)	(0.007)	(0.025)	(0.026)	(0.675)	(0.077)
β_1	**0.125**	**-0.076**	0.012	0.054	**19.123**	**0.968**
	(0.013)	(0.013)	(0.025)	(0.033)	(1.241)	(0.119)
β_2	**-0.107**	**0.064**	-0.038	**0.216**	**13.632**	**-1.054**
	(0.015)	(-0.012)	(0.050)	(0.073)	(1.208)	(0.160)
β_3	-0.019	0.013	0.210	0.044	1.741	0.264
	(0.022)	(0.021)	(0.117)	(0.139)	(2.349)	(0.206)
β_4	0.037	-0.013	0.427	-0.427	3.277	0.357
	(0.036)	(0.041)	(0.351)	(0.356)	(2.900)	0.424
β_5	**0.085**	-0.050	-0.666	**0.857**	6.785	-0.474
	(0.043)	(0.042)	(0.389)	(0.389)	(5.474)	(0.851)
β_6	-0.059	-0.002	0.371	-0.330	**-20.883**	**-1.891**
	(0.096)	(0.093)	(0.234)	(0.241)	(9.000)	(0.914)
R^2	0.764	0.520	0.027	0.240	0.901	0.512
$AC(1)$	0.234	0.253	0.002	0.015	0.048	-0.038
$AC(2:10)$	-0.051	-0.038	-0.000	-0.016	0.063	0.003
$AC(11:20)$	0.017	0.022	0.001	-0.006	0.037	-0.021

Each day, between July 2004 and December 2012, 18 $RIVV$ values are interpolated on the $RIVV$ surface with nine different call-equivalent deltas ($\Delta^e = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$) and two different maturities (30 and 90 days) from the locally smoothed quadratic regression surface obtained with the loess procedure. A principal component analysis is performed on those 18 time series. A regression of the six first principal components is run on the $RIVV$ characteristics (i.e. Level, TS, Skew, Skew TS) as well as RV and $\max(RV - BV, 0)$ according to:

$$\text{Char}_t = \beta_0 + \beta_1 PC_{1,t} + \beta_2 PC_{2,t} + \beta_3 PC_{3,t} + \beta_4 PC_{4,t} + \beta_5 PC_{5,t} + \beta_6 PC_{6,t} + \varepsilon_t,$$

Standard errors are in parenthesis and are adjusted for autocorrelation and heteroscedasticity with Newey West (1987) method. $AC(1)$ corresponds to the first lag period autocorrelation coefficient, $AC(2:10)$ is the average of the autocorrelation coefficients for lag periods 2 to 10, and $AC(11:20)$ is the average of the autocorrelation coefficients for lag periods 11 to 20.

Chapter 6

Conclusion

In this paper, we carried out a simulation-based study designed to evaluate and compare the ability of various sources of information to estimate latent states, namely the instantaneous variance, log-equity jumps and variance jumps. We first defined an alternative observable measure to the realized option variance of Amaya et al. (2017), the realized implied volatility variance. We demonstrated that this new model-free observable is a good approximation of the implied volatility quadratic variation, enabling us to use it as a jump detection tool via particle filtering. Second, we analyzed nonparametrically the properties of the realized implied volatility variance using S&P 500 Index and underlying options data.

In a simulation environment, we determined that the filter's performance using the realized option variance, the realized implied volatility variance or both combined yield similar results, leading us to conclude that they can be used interchangeably in a filtering procedure. Moreover, we demonstrated that the realized implied volatility variance is redundant with the realized variance in the absence of jumps, but brings additional information in the presence of jumps. Similarly, we showed that the information contained in the realized option variance is the same as the one conveyed by the realized implied volatility variance. For ITM and ATM options, $RIVV$s convey information about both the diffusive and discontinuous part of the asset price and variance

dynamics. Conversely, OTM options are driven by the discontinuous part.

From an empirical perspective, we studied the properties of the realized implied volatility variance for options on the S&P 500 and how they contrast with the index variations. This is done by constructing a series of $RIVV$ surfaces and studying the features of this surface. We found that additional information is embedded in the cross-section of $RIVV$ that complements the one conveyed by RV. We also performed a principal component analysis of the realized implied volatility variance and found a high degree of commonality between index variations. However, the principal component analysis also revealed that $RIVV$ exhibits information not provided by RV, suggesting again its added value as a jump detection tool of underlying generating processes. Considering the similar results reported in Amaya et al. (2017) with regards to the realized option variance, our overall conclusion is that ROV and $RIVV$ convey similar information.

Appendix A

Option Pricing

The price of a European call option with strike price K and maturity T is

$$C_t(Y_t, V_t) = \exp(Y_t)\exp(-q(T-t))P_1(Y_t, V_t) - K\exp(-r(T-t)P_2(Y_t, V_t), \quad \text{(A.1)}$$

where

$$P_1(y, v) = \frac{1}{2} + \frac{1}{\pi}\int_0^t Re\left(\frac{\exp(-iuk - y)\varphi^{\mathbb{Q}}_{Y_T|Y_t,V_t}(ui+1, y, v)}{ui}\right)du,$$
$$P_2(y, v) = \frac{1}{2} + \frac{1}{\pi}\int_0^t Re\left(\frac{\exp(-iuk)\varphi^{\mathbb{Q}}_{Y_T|Y_t,V_t}(ui, y, v)}{ui}\right)du,$$

and $\varphi^{\mathbb{Q}}_{Y_T|Y_t,V_t}(u, y, v)$ is the moment generating function of Y_T conditional on information at time t. Also, note that $k = \log(K)$. The moment generating function is defined as:

$$\varphi^{\mathbb{Q}}_{Y_T|Y_t,V_t}(u, y, v) = exp(\mathcal{A}(u, t, T) + uY_t + \mathcal{C}(u, t, T)V_t),$$

where

$$\mathcal{C}(u,t,T) = \frac{2C_0(\exp(-C_2(T-t))-1)}{C_2(\exp(-C_2(T-t))+1) - C_1(\exp(-C_2(T-t))-1)}, \tag{A.2}$$

$$C_0 = \lambda^{\mathbb{Q}}_{Y,1}\Big(\varphi^{\mathbb{Q}}_{Z_Y}(1)-1\Big)u - \lambda^{\mathbb{Q}}_{Y,1}\Big(\varphi^{\mathbb{Q}}_{Z_Y}(u)-1\Big) + \frac{u-u^2}{2},$$

$$C_1 = \kappa^{\mathbb{Q}} - \rho\sigma u,$$

$$C_2 = \sqrt{C_1^2 + 2\sigma^2 C_0},$$

and

$$\mathcal{A}(u;t,T) = D_0(T-t) + \theta^{\mathbb{Q}}\kappa^{\mathbb{Q}} g_1(t,T) + \lambda^{\mathbb{Q}}_{V,0} g_2(t,T), \tag{A.3}$$

$$D_0 = -r + (r-q)u + \lambda^{\mathbb{Q}}_{Y,0}\Big(\varphi^{\mathbb{Q}}_{Z_Y}(u)-1\Big) - \lambda^{\mathbb{Q}}_{Y,0}\Big(\varphi^{\mathbb{Q}}_{Z_Y}(1)-1\Big)u,$$

$$g_1(t,T) = \frac{2\log(2C_2) - 2\log(C_1(e^{C_2(T-t)}-1)}{-\sigma^2} + \frac{C_2(e^{C_2(T-t)}+1)) + (C_1+C_2)(T-t)}{-\sigma^2},$$

$$g_2(t,T) = \frac{\mu^{\mathbb{Q}}_V}{2}\frac{2\log(2C_2) - 2\log\Big(C_1(e^{C_2(T-t)}-1) + C_2(e^{C_2(T-t)}+1))\Big)}{C_0\Big(\mu^{\mathbb{Q}}_V\Big)^2 + C_1\mu^{\mathbb{Q}}_V - \frac{\sigma^2}{2}} + \frac{\mu^{\mathbb{Q}}_V}{2}\frac{2C_0\mu^{\mathbb{Q}}_V(e^{C_2(T-t)}-1) + 2C_0\mu^{\mathbb{Q}}_V + C_1 - C_2}{C_0\Big(\mu^{\mathbb{Q}}_V\Big)^2 + C_1\mu^{\mathbb{Q}}_V - \frac{\sigma^2}{2}}(T-t).$$

Appendix B

Quadratic variations

B.1 Quadratic Variation of Option Prices

The quadratic variation of options prices is defined by the following expression:

$$
\begin{aligned}
&O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
&= \int_0^t \frac{\partial O_u}{\partial u}(Y_{u^-}, V_{u^-})du + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})dY_u \\
&+ \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})dV_u + \frac{1}{2}\int_0^t \frac{\partial^2 O_u}{\partial y^2}(Y_{u^-}, V_{u^-})d\langle Y, Y\rangle_u^C \\
&+ \frac{1}{2}\int_0^t \frac{\partial^2 O_u}{\partial v^2}(Y_{u^-}, V_{u^-})d\langle V, V\rangle_u^C + \int_0^t \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})d\langle Y, V\rangle_u^C \\
&+ \sum_{0<u\leqslant t}\left(O_u(Y_u, V_u) - O_u(O_{u^-}, V_{u^-})\right) - \sum_{0<u\leqslant t}\left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})(Y_u - Y_{u^-})\right) \\
&\qquad - \sum_{0<u\leqslant t}\left(\frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})(V_u - V_{u^-})\right),
\end{aligned}
$$

where $\langle\cdot,\cdot\rangle_t^C$ is the time t continuous part of the quadratic variation.

Then, replacing model Equations (2.1) and (2.2), leads to:

$$
\begin{aligned}
&O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
&= \int_0^t \frac{\partial O_u}{\partial u}(Y_{u^-}, V_{u^-})du + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\alpha_{u^-}du \\
&+ \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\rho\sqrt{V_{u^-}}dW_{V,u} + \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\sqrt{1-\rho^2}\sqrt{V_{u^-}}dW_{\perp,u} \\
&+ \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})dJ_{Y,u} + \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\kappa(\theta - V_{u^-})du \\
&+ \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\sigma\sqrt{V_{u^-}}dW_{V,u} + \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})dJ_{V,u} \\
&+ \frac{1}{2}\int_0^t \frac{\partial^2 O_u}{\partial y^2}(Y_{u^-}, V_{u^-})V_{u^-}du + \frac{1}{2}\int_0^t \frac{\partial^2 O_u}{\partial v^2}(Y_{u^-}, V_{u^-})\sigma^2 V_{u^-}du \\
&+ \int_0^t \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})\sigma\rho V_{u^-}du + \sum_{0<u\leqslant t}\left(O_u(Y_u, V_u) - O_u(O_{u^-}, V_{u^-})\right)^2 \\
&\qquad - \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})dJ_{Y,u} - \int_0^t \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})dJ_{V,u}.
\end{aligned}
$$

By reorganizing the previous equation, the jump integrals cancel each other out and we obtain:

$$
\begin{aligned}
&O_t(Y_t, V_t) - O_0(Y_0, V_0) \\
&= \int_0^t \Bigg\{ \frac{\partial O_u}{\partial u}(Y_{u^-}, V_{u^-}) + \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\alpha_{u^-} + \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\kappa(\theta - V_{u^-}) \\
&+ \left(\frac{1}{2}\frac{\partial^2 O_u}{\partial y^2}(Y_{u^-}, V_{u^-}) + \frac{1}{2}\frac{\partial^2 O_u}{\partial v^2}(Y_{u^-}, V_{u^-})\sigma^2 + \frac{\partial^2 O_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})\sigma\rho\right)V_{u^-}\Bigg\}du \\
&+ \int_0^t \left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\rho + \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\sigma\right)\sqrt{V_{u^-}}dW_{V,u} \\
&+ \int_0^t \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \\
&\qquad \sqrt{1-\rho^2}\sqrt{V_{u^-}}dW_{\perp,u} + \sum_{0<u\leqslant t}\left(O_u(Y_u, V_u) - O_u(O_{u^-}, V_{u^-})\right)^2.
\end{aligned}
$$

Finally, the option prices quadratic variation is

$$
\begin{aligned}
&\langle O, O\rangle_t \\
&= \int_0^t \left(\rho \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) + \sigma \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\right) V_{u^-} du \\
&+ \int_0^t \left(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\right)^2 (1-\rho^2) V_{u^-} du \\
&+ \sum_{0<u\leqslant t} \left(O_u(Y_u, V_u) - O_u(O_{u^-}, V_{u^-})\right)^2. \\
&= \int_0^t \Bigg(\Big(\frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-})\Big)^2 \\
&+ 2\sigma\rho \frac{\partial O_u}{\partial y}(Y_{u^-}, V_{u^-}) \frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-}) + \sigma^2 \Big(\frac{\partial O_u}{\partial v}(Y_{u^-}, V_{u^-})\Big)^2 \Bigg) V_{u^-} du \\
&\qquad\qquad + \sum_{0<u\leqslant t} \left(O_u(Y_u, V_u) - O_u(V_{u^-}, V_{u^-})\right)^2.
\end{aligned}
$$

B.2 Quadratic Variation of Implied Volatilities

The quadratic variation of implied volatility is defined by the following expression:

$$
\begin{aligned}
&IV_t(Y_t, V_t) - IV_0(Y_0, V_0) \\
&= \int_0^t \frac{\partial IV_u}{\partial u}(Y_{u^-}, V_{u^-})du + \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})dY_u \\
&+ \int_0^t \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})dV_u + \frac{1}{2}\int_0^t \frac{\partial^2 IV_u}{\partial y^2}(Y_{u^-}, V_{u^-})d\langle Y, Y\rangle_u^C \\
&+ \frac{1}{2}\int_0^t \frac{\partial^2 IV_u}{\partial v^2}(Y_{u^-}, V_{u^-})d\langle V, V\rangle_u^C + \int_0^t \frac{\partial^2 IV_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})d\langle Y, V\rangle_u^C \\
&+ \sum_{0<u\leqslant t}\left(IV_u(Y_u, V_u) - O_u(IV_{u^-}, V_{u^-})\right) - \sum_{0<u\leqslant t}\left(\frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})(Y_u - Y_{u^-})\right) \\
&\qquad - \sum_{0<u\leqslant t}\left(\frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})(V_u - V_{u^-})\right),
\end{aligned}
$$

where $\langle \cdot, \cdot \rangle_t^C$ is the time t continuous part of the quadratic variation.

Then, replacing model Equations (2.1) and (2.2), leads to:

$$
\begin{aligned}
&IV_t(Y_t, V_t) - IV_0(Y_0, V_0) \\
&= \int_0^t \frac{\partial IV_u}{\partial u}(Y_{u^-}, V_{u^-})du + \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\alpha_{u^-}du \\
&+ \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\rho\sqrt{V_{u^-}}dW_{V,u} + \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\sqrt{1-\rho^2}\sqrt{V_{u^-}}dW_{\perp,u} \\
&+ \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})dJ_{Y,u} + \int_0^t \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\kappa(\theta - V_{u^-})du \\
&+ \int_0^t \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\sigma\sqrt{V_{u^-}}dW_{V,u} + \int_0^t \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})dJ_{V,u} \\
&+ \frac{1}{2}\int_0^t \frac{\partial^2 IV_u}{\partial y^2}(Y_{u^-}, V_{u^-})V_{u^-}du + \frac{1}{2}\int_0^t \frac{\partial^2 IV_u}{\partial v^2}(Y_{u^-}, V_{u^-})\sigma^2 V_{u^-}du \\
&+ \int_0^t \frac{\partial^2 IV_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})\sigma\rho V_{u^-}du + \sum_{0<u\leqslant t}\left(IV_u(Y_u, V_u) - IV_u(IV_{u^-}, V_{u^-})\right)^2 \\
&\qquad - \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})dJ_{Y,u} - \int_0^t \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})dJ_{V,u}.
\end{aligned}
$$

By reorganizing the previous equation, the jump integrals cancel each other out and we obtain:

$$
\begin{aligned}
&IV_t(Y_t, V_t) - IV_0(Y_0, V_0) \\
&= \int_0^t \Bigg\{ \frac{\partial IV_u}{\partial u}(Y_{u^-}, V_{u^-}) + \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\alpha_{u^-} + \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\kappa(\theta - V_{u^-}) \\
&+ \left(\frac{1}{2}\frac{\partial^2 IV_u}{\partial y^2}(Y_{u^-}, V_{u^-}) + \frac{1}{2}\frac{\partial^2 IV_u}{\partial v^2}(Y_{u^-}, V_{u^-})\sigma^2 + \frac{\partial^2 IV_u}{\partial v \partial y}(Y_{u^-}, V_{u^-})\sigma\rho\right)V_{u^-} \Bigg\}du \\
&+ \int_0^t \left(\frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\rho + \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\sigma\right)\sqrt{V_{u^-}}dW_{V,u} \\
&+ \int_0^t \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-}) \\
&\qquad \sqrt{1-\rho^2}\sqrt{V_{u^-}}dW_{\perp,u} + \sum_{0<u\leqslant t}\left(IV_u(Y_u, V_u) - IV_u(IV_{u^-}, V_{u^-})\right)^2.
\end{aligned}
$$

Finally, the implied volatility quadratic variation is

$$
\begin{aligned}
&\langle IV, IV\rangle_t \\
&= \int_0^t \left(\rho \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-}) + \sigma \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\right) V_{u^-} du \\
&+ \int_0^t \left(\frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\right)^2 (1-\rho^2) V_{u^-} du \\
&+ \sum_{0<u\leqslant t} \left(IV_u(Y_u, V_u) - IV_u(IV_{u^-}, V_{u^-})\right)^2. \\
&= \int_0^t \Bigg(\Big(\frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-})\Big)^2 \\
&+ 2\sigma\rho \frac{\partial IV_u}{\partial y}(Y_{u^-}, V_{u^-}) \frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-}) + \sigma^2 \Big(\frac{\partial IV_u}{\partial v}(Y_{u^-}, V_{u^-})\Big)^2 \Bigg) V_{u^-} du \\
&\qquad + \sum_{0<u\leqslant t} \left(IV_u(Y_u, V_u) - IV_u(V_{u^-}, V_{u^-})\right)^2.
\end{aligned}
$$

B.2.1 Derivative Calculation for $\Delta IVQV$

The derivatives of $IV_t(Y_t, V_t)$ can be rewritten as follows

$$
\begin{aligned}
\frac{\partial IV_t}{\partial y}(Y_{t^-}, V_{t^-}) &= \left(\frac{\partial IV_t}{\partial O}(Y_{t^-}, V_{t^-})\right)\left(\frac{\partial O_t}{\partial y}(Y_{t^-}, V_{t^-})\right) \quad \text{and,} \\
\frac{\partial IV_t}{\partial v}(Y_{t^-}, V_{t^-}) &= \left(\frac{\partial IV_t}{\partial O}(Y_{t^-}, V_{t^-})\right)\left(\frac{\partial O_t}{\partial v}(Y_{t^-}, V_{t^-})\right).
\end{aligned}
$$

Internet Appendix B.2.1 of Amaya et al. (2017) shows how the derivatives $\frac{\partial O_t}{\partial y}$ and $\frac{\partial O_t}{\partial v}$ can be calculated. The derivative $\frac{\partial IV_t}{\partial O}$ can be rewritten as the inverse Black and Scholes vega:

$$
\frac{\partial IV_t}{\partial O}(Y_{t^-}, V_{t^-}) = \frac{1}{\exp(Y_t)\phi(d_t)\sqrt{T-t}}
$$

where $\phi(\cdot)$ is the normal density function and

$$d_t = \frac{\log(\frac{\exp(Y_t)}{K}) + (r - q + \frac{\sigma^2}{2})(T - t)}{\sigma\sqrt{T - t}}$$

is evaluated at $\sigma = IV_t(Y_t, V_t)$. Parameters K and T are the exercise price and the maturity of the option considered, respectively. Constants r and q are the risk-free rate and the dividend yield, respectively.

Appendix C

The filtering procedure

C.1 Intraday Simulation

The intraday step h is set to Δ/M, where Δ is one day (1/252) in our setting and M is the number of steps in a day. Starting from known values of Y_t, V_t and other model parameters, the next Y_{t+h} V_{t+h} and $\Delta I_{t,t+h}$ are generated as follows:

1. Generate a log-price jump indicator $\mathbb{1}_{\{N_{V,t+h}-N_{V,t}=1\}}$ and a variance jump indicator $\mathbb{1}_{\{N_{V,t+h}-N_{V,t}=1\}}$ from independent Bernouilli random variables with probability of success $\lambda_{Y,t}h$ and $\lambda_{V,t}h$, respectively.

2. Generate a log-price jump size $Z_{Y,t+h} \sim \mathcal{N}(\mu_Y,\, \sigma_Y^2)$ and a variance jump size variable $Z_{V,n} \sim \text{Exp}(\mu_V)$ if the jump indicators call for it. Otherwise, jump sizes are set to 0.

3. The new variance at $t+h$ is obtained with the Euler approximation of Equation (2.2):

$$V_{t+h} \cong V_t + \kappa(\theta - V_t)h + \sigma\sqrt{V_t}(W_{V,t+h} - W_{V,t}) + Z_{V,t+h}. \qquad \text{(C.1)}$$

4. The integrated variance $\Delta I_{t,t+h} = \int_t^{t+h} V_{u^-} du$ is simulated conditionally on V_t and $V_{(t+h)^-}$. We take advantage of the fact that the cumulant

generating function of $\Delta I_{t,t+h}$ can be approximated by Gaussian distribution. See Amaya et al. (2017) for more details on how the conditional density function of $\Delta I_{t,t+h}$ can be approximated. Tse and Wan (2013) provide a closed-form solution for the computation of $\Delta I_{t,t+h}$.

5. The new log-price at $t+h$ is simulated as follows:

$$\begin{aligned} Y_{t+h} =& Y_t + c_1 h + c_2 \Delta I_{t,t+h} + \frac{\rho}{\sigma}(V_{(t+h)^-} - V_t) \\ &+ \sqrt{1-\rho^2} \int_t^{t+h} \sqrt{V_{u^-}} dW_{\perp,u} \\ &- \frac{\rho}{\sigma} \left(\sum_{n=N_{V,t}+1}^{N_{V,t+h}} Z_{V,n} \right) + \sum_{n=N_{Y,t}+1}^{N_{Y,t+h}} Z_{Y,n}, \end{aligned} \tag{C.2}$$

where

$$\begin{aligned} c_1 &= r - q + \varphi^{\mathbb{P}}_{Z_Y}(\Gamma_Y) - \varphi^{\mathbb{P}}_{Z_Y}(1+\Gamma_Y)\lambda_{Y,0} - \frac{\rho\kappa\theta}{\sigma}, \\ c_2 &= \eta_Y - \frac{1}{2} + \varphi^{\mathbb{P}}_{Z_Y}(\Gamma_Y) - \varphi^{\mathbb{P}}_{Z_Y}(1+\Gamma_Y)\lambda_{Y,1} + \frac{\rho\kappa}{\sigma}. \end{aligned}$$

C.2 Time Aggregation

After completing the intraday simulation in Appendix C.1, the integrated variance and quadratic variation of the simulated particles are aggregated at the end of the day:

$$\begin{aligned} &\Delta I_{t,t+\Delta} \simeq \sum_{j=1}^{M} \Delta I_{t+(j-1)h,t+jh} \quad \text{and,} \\ &\Delta QV_{t,t+\Delta} \simeq \Delta I_{t,t+\Delta} + \sum_{j=1}^{M} Z_{Y,t+jh} \mathbb{1}_{\{N_{Y,t+h} - N_{Y,t} = 1\}}. \end{aligned}$$

It is important to note that the time aggregation of Equation (C.2) means the

daily returns can be approximated by:

$$Y_{t+\Delta} - Y_t = \sum_{j=1}^{M} (Y_{t+jh} - Y_{t+(j-1)h}) \simeq \mu_{t+\Delta} + \sigma^2_{t+\Delta}\epsilon_{t+h}$$

where ϵ_{t+h} is a standard normal random variable,

$$\mu_{t+\Delta} = c_1\Delta + c_2\Delta I_{t,t+\Delta} + \frac{\rho}{\sigma}(V_{t+\Delta} - V_t) \tag{C.3}$$

$$- \frac{\rho}{\sigma}\sum_{j=1}^{M} Z_{V,t+jh}\mathbb{1}_{\{N_{V,t+h}-N_{V,t}=1\}} + Z_{Y,t+jh}\mathbb{1}_{\{N_{Y,t+h}-N_{Y,t}=1\}}, \tag{C.4}$$

and $\sigma^2_{t+\Delta} = (1-\rho^2)\Delta I_{t,t+\Delta}$. $IV_{t+\Delta}(Y_{t+\Delta}, V_{t+\Delta})$ is the Black and Scholes implied volatility computed based on the model option price $O_{t+\Delta}(Y_{t+\Delta}, V_{t+\Delta})$ as given by Equation (A.1).

Both the option quadratic variation and implied volatility quadratic variation are approximated with a Euler discretization of Equations (2.12) and (2.17), which are provided in Appendices B.1, and B.2, respectively:

$$\begin{aligned}\Delta OQV_{t,t+\Delta} = h\sum_{j=1}^{M}\Bigg(&\Big(\frac{\partial O_t}{\partial y}(Y_t, V_t)\Big)^2 + \sigma^2\Big(\frac{\partial O_t}{\partial v}(Y_t, V_t)\Big)^2 \\ &+ 2\sigma\rho\Big(\frac{\partial O_t}{\partial y}(Y_t, V_t)\Big)\Big(\frac{\partial O_t}{\partial v}(Y_t, V_t)\Big)\Bigg)\Delta I_{t+(j-1)h,t+jh} \\ &+ \sum_{j=1}^{M}\Bigg(O_{t+jh}(Y_{t+jh}, V_{t+jh}) - O_{t+jh}(Y_{(t+jh)^-}, V_{(t+jh)^-})\Bigg)^2\end{aligned} \tag{C.5}$$

$$\begin{aligned}\Delta IVQV_{t,t+\Delta} = h\sum_{j=1}^{M}\Bigg(&\Big(\frac{\partial IV_t}{\partial y}(Y_t,V_t)\Big)^2+\sigma^2\Big(\frac{\partial IV_t}{\partial v}(Y_t,V_t)\Big)^2 \\ &+2\sigma\rho\Big(\frac{\partial IV_t}{\partial y}(Y_t,V_t)\Big)\Big(\frac{\partial IV_t}{\partial v}(Y_t,V_t)\Big)\Bigg)\Delta I_{t+(j-1)h,t+jh} \\ &+\sum_{j=1}^{M}\Bigg(IV_{t+jh}(Y_{t+jh},V_{t+jh})-IV_{t+jh}(Y_{(t+jh)^-},V_{(t+jh)^-})\Bigg)^2\end{aligned} \tag{C.6}$$

Appendix B.2.1 also provides details on how partial derivatives of the implied volatility are computed. We refer the reader to the Internet Appendix B.2.1 of Amaya et al. (2017) for the partial derivatives of the option price.

After the time aggregation is completed, we obtain N_L particles, $L^{(i)}_{t+\Delta}$, for $i = 1, ..., N_L$, at the end of the day:

$$\begin{aligned}L^{(i)}_{t+\Delta} = \Bigg\{&Y^{(i)}_{t+\Delta}, V^{(i)}_{t+\Delta}, \Delta I^{(i)}_{t,t+\Delta}, \Delta QV^{(i)}_{t,t+\Delta}, \\ &O^{(i)}_{t+\Delta,k}(Y_{t+\Delta},V_{t+\Delta}), \Delta OQV^{(i)}_{t,t+\Delta,k}, \Delta IVQV^{(i)}_{t,t+\Delta,k} : k \in 1,2,...,n_t\Bigg\}.\end{aligned}$$

Bibliography

Amaya, D., Bégin, J.-F., and Gauthier, G. (2017). Extracting latent states from high frequency option prices. Working Paper.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of financial economics*, 61(1):43–76.

Andersen, T. G., Bondarenko, O., Todorov, V., and Tauchen, G. (2015a). The fine structure of equity-index option dynamics. *Journal of Econometrics*, 187(2):532–546.

Andersen, T. G., Fusari, N., and Todorov, V. (2015b). The risk premia embedded in index options. *Journal of Financial Economics*, 117(3):558–584.

Audrino, F. and Fengler, M. R. (2015). Are classical option pricing models consistent with observed option second-order moments? evidence from high-frequency data. *Journal of Banking & Finance*, 61:46–63.

Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049.

Bandi, F. M. and Renò, R. (2016). Price and volatility co-jumps. *Journal of Financial Economics*, 119(1):107–146.

Bardgett, C., Gourier, E., and Leippold, M. (2014). Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of financial econometrics*, 2(1):1–37.

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1):181–238.

Broadie, M. and Kaya, Ö. (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations research*, 54(2):217–231.

Christoffersen, P., Feunou, B., and Jeon, Y. (2015). Option valuation with observable volatility and jump dynamics. *Journal of Banking & Finance*, 61:S101–S120.

Christoffersen, P., Jacobs, K., and Ornthanalai, C. (2012). Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options. *Journal of Financial Economics*, 106(3):447–472.

Creal, D. (2012). A survey of sequential monte carlo methods for economics and finance. *Econometric reviews*, 31(3):245–296.

Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.

Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *The Journal of Finance*, 59(3):1367–1403.

Eraker, B., Johannes, M., and Polson, N. (2003). The impact of jumps in volatility and returns. *The Journal of Finance*, 58(3):1269–1300.

Gordon, N. J., Salmond, D. J., and Smith, A. F. (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE Proceedings F (Radar and Signal Processing)*, volume 140, pages 107–113. IET.

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.

Johannes, M. S., Polson, N. G., and Stroud, J. R. (2009). Optimal filtering of jump diffusions: Extracting latent states from asset prices. *The Review of Financial Studies*, 22(7):2759–2799.

Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of financial economics*, 8(4):323–361.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of financial economics*, 63(1):3–50.

Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American statistical association*, 94(446):590–599.

Protter, P. E. (2004). Stochastic integration and differential equations, volume 21 of applications of mathematics: Stochastic modelling and applied probability.

Renault, E. (1997). Econometric models of option pricing errors. *Econometric Society Monographs*, 28:223–278.

Santa-Clara, P. and Yan, S. (2010). Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *The Review of Economics and Statistics*, 92(2):435–451.

Todorov, V. and Tauchen, G. (2011). Volatility jumps. *Journal of Business & Economic Statistics*, 29(3):356–371.

Tse, S. T. and Wan, J. W. (2013). Low-bias simulation scheme for the heston model by inverse gaussian approximation. *Quantitative finance*, 13(6):919–937.

Van Haastrecht, A. and Pelsser, A. (2010). Efficient, almost exact simulation of the heston stochastic volatility model. *International Journal of Theoretical and Applied Finance*, 13(01):1–43.

Zhang, L., Mykland, P. A., and Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 100(472):1394–1411.