Testing for Jumps and Modeling Volatility in Asset Prices

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Johan Bjursell
Master of Science
University of Southern Mississippi, 2002
Bachelor of Science
Göteborg Unversity, 2000

Director: Dr. James E. Gentle, Professor
Department of Computational and Data Sciences

Spring Semester 2009
George Mason University
Fairfax, VA

Dedication

My mom, dad and sister

Dagmar and Knut

Britten and Carl

Kumiko

Acknowledgments

I want to acknowledge the following people for their contributions and continuing support throughout my graduate studies at George Mason University.

Dr. James E. Gentle, my advisor, who always keeps his door open. He has in our weekly meetings patiently listened to my ideas, offered his opinion on results and provided suggestions for extensions and modifications. As a result, I conclude this dissertation with an extensive list of future research directions. Besides his academic support, I also appreciate several great culinary experiences owed to his cooking.

Dr. George. H. K. Wang, committee member, who as the Deputy Chief Economist at the Commodity Futures Trading Commission, Washington DC, offered me an internship position. I truly appreciate the opportunity to learn about futures markets, work with trading data, and develop code libraries for analyzing such data. I have also benefited from having worked on a number of projects with Dr. Wang over the past few years.

Dr. Daniel B. Carr, Dr. Michael G. Ferri, and Dr. Igor Griva, committee members, with whom I have had several interesting discussions which I have benefited greatly from. They have generously offered suggestions for research directions and corrections of earlier drafts.

Finally, friends and family for their constant encouragement.

Johan Bjursell
May 01, 2009

Table of Contents

List of Tables

List of Figures

Abstract

TESTING FOR JUMPS AND MODELING VOLATILITY IN ASSET PRICES

Johan Bjursell, PhD

George Mason University, 2009

Dissertation Director: Dr. James E. Gentle

Observers of financial markets have long noted that asset prices are very volatile and commonly exhibit jumps (price spikes). Thus, the assumption of a continuous process for asset price behavior is often violated in practice. Although empirical studies have found that the impact of such jumps is transitory, the shortterm effect in the volatility may nonetheless be considerable with important financial implications for the valuation of derivatives, asset allocation and risk management.

This dissertation contributes to the literature in two areas. First, I evaluate the small sample properties of a nonparametric method for identifying jumps. I focus on the implication of adding noise to the prices and recent methods developed to contend with such market frictions. Initially, I examine the properties and convergence results of the power variations that constitute the jump statistics. Then I document the asymptotic results of these jump statistics. Finally, I estimate their size and power. I examine these properties using a stochastic volatility model incorporating alternative noise and jump processes. I find that the properties of the statistics remain close to the asymptotics when methods for managing the effects of noise are applied judiciously. Improper use leads to invalid tests or tests with low power. Empirical evidence

demonstrates that the nonparametric method performs well for alternative models, noise processes, and jump distributions.

In the second essay, I present a study on market data from U.S. energy futures markets. I apply a nonparametric method to identify jumps in futures prices of crude oil, heating oil and natural gas contracts traded on the New York Mercantile Exchange. The sample period of the intraday data covers January 1990 to January 2008. Alternative methods such as staggered returns and optimal sampling frequency methods are used to remove the effects of microstructure noise which biases the tests against detecting jumps.

I obtain several important empirical results: (*i*) The realized volatility of natural gas futures exceeds that of heating oil and crude oil. (*ii*) In these commodities, large volatility days are often associated with large jump components and large jump components are often associated with weekly announcements of inventory levels. (*iii*) The realized volatility and smooth volatility components in natural gas and heating oil futures are higher in winter months than in summer months. Moreover, cold weather and inventory surprises cause the volatility in natural gas and heating oil to increase during the winter season. (*iv*) The jump component produces a transitory surge in total volatility, and there is a strong reversal in volatility on days following a significant jump day. (*v*) I find that including jump and seasonal components as explanatory variables significantly improves the modeling and forecasting of the realized volatility.

Chapter 1: Introduction

Observers of financial markets have long noted that asset prices are very volatile and often exhibit jumps (price spikes). Thus, the assumption of a continuous diffusion process for asset price behavior is often violated in practice. Although empirical studies often note that the impact of such jumps generally is transitory, the short-term effect in the volatility may nonetheless be considerable with important financial implications for valuation of derivatives (Merton (1976)), asset allocation (Jarrow and Rosenfeld (1984)) and risk management (Duffie and Pan (2001)).

A number of studies has shown that models including both a discontinuous jump component and a continuous component fits the data better than only a continuous process. For example, Cox and Rubinstein (1985) compare the Black-Scholes formula with Merton's option pricing formula (Merton (1976)) and show that for large and frequent jumps in the price process of the underlying assets, Black-Scholes significantly undervalues out-of-the-money and at-the-money options. In a more recent study, Bakshi et al. (1997) compare several parametric models with and without a jump component based on model fitting, pricing, and hedging. They report that it is essential to include a jump component for pricing and internal consistency. Eraker et al. (2003) observe that jumps in the returns occur less frequently than what is reported in most literature but are nevertheless still significant. Maheu and McCurdy (2004) assume that jumps in stock market returns are generated by a nonhomogenous Poisson process and find that the addition of the jump component improves forecasts of volatility.

An increase in the availability of high-frequency or transaction data has produced a growing literature on nonparametric methods to identify jumps such as Barndorff-Nielsen and Shephard (2004, 2006), Fan and Wang (2007), Jiang and Oomen (2008) and Sen (2008). Literature using nonparametric methods include Huang and Tauchen (2005), who provide evidence that jumps account for seven percent of the S&P 500 index's realized variance. Andersen et al. (2007) provide empirical evidence that the volatility jump component is both highly significant and less persistent than the continuous component in foreign exchange rate spot (DM/$) market, S&P 500 index futures and thirty-year US Treasury bond futures. Jiang et al. (2008) study treasury bond futures and find that about seventy percent of jumps can be associated with macroeconomic news releases.

Parametric models are generally applied to daily observations while nonparametric approaches are based on intraday data. Clearly, intraday data is richer in information and thus presumably may produce more efficient estimates. However, the utilization of intraday data is hampered by the presence of market microstructure noise. Such frictions come from trade mechanisms and rules that govern the markets. On a daily or longer time horizon, such noise is small compared to the volatility due to information, but may dominate estimates at high intraday sampling frequencies. Consequently, methods that are based on transaction data need to contend with the effects of such noise. Huang and Tauchen (2005) examine the impact of noise in a small sample study on the nonparametric method proposed by Barndorff-Nielsen and Shephard (2004, 2006); however, recent methods for filtering the effects of noise have not been applied and evaluated in this context. I seek to fill this gap in Chapter 2, *"Detecting Jumps in Asset Prices Using Bipower Variation"*.

In Chapter 3, *"Volatility and Jump Dynamics in U.S. Energy Futures Markets"*, I evaluate a more recent nonparametric method proposed by Jiang et al. (2008) and

apply the method to a dataset from U.S. energy futures markets. I document jump processes, study their seasonal and intraday trends, and examine their contribution to the total volatility.

Detecting Jumps in Asset Prices Using Bipower Variation

In this chapter, I evaluate nonparametric statistics by Barndorff-Nielsen and Shephard (2004, 2006) that can be applied to identify days with jumps in a price process. I evaluate the finite sample properties of the test statistics for noisy prices and particularly examine whether recently proposed methods for reducing the impact of noise improve the tests.

First, I examine the properties and convergence results of the statistics that constitute the jump statistics. Second, I evaluate whether methods that Bandi and Russell (2006) and Zhang et al. (2005) propose to reduce the impact of noise in estimates of the daily integrated variance apply to other intraday variations, specifically, to the bipower and tripower variations. Third, I use the methods by Bandi and Russell (2006) and Zhang et al. (2005) to test for jumps in price processes with noise. Moreover, I combine these methods with using staggered returns, which previously have been applied in the literature (see Andersen et al. (2007) and Huang and Tauchen (2005)). The methods by Bandi and Russell (2006) and Zhang et al. (2005) have not, to the best of by knowledge, previously been applied to the jump statistics.[1] I also propose and evaluate a modified version of the method by Bandi and Russell (2006) to make it more robust to jumps. Fourth, I consider alternative noise processes recently proposed by Aït-Sahalia et al. (2006) and Li and Mykland (2007), and study the finite sample properties of the jump test statistics under these processes. I also evaluate the statistics for alternative jump distributions. Specifically, I generate

[1]See Andersen et al. (2007), for example, who call attention to the lack of such a study.

jumps from normal, skewed-normal and double exponential distributions. Finally, while the methods by Bandi and Russell (2006) and Zhang et al. (2005) are based on sampling the price process in an optimal manner to lessen the bias due to noise, they nevertheless discard large fractions of the data. I empirically evaluate another method developed by Zhang et al. (2005) that uses all data to estimate the daily integrated variations.

I obtain several interesting results:

1. The statistics converge to the limiting normal distribution with zero mean and unit variance as the sampling interval approaches zero for efficient (noise-free) prices. The statistics have converged at a one-minute sampling interval. The convergence results are highly influenced by noise, however, in which case the limiting distribution remains normal but the moments become strongly biased.

2. Noise biases the statistics against identifying jumps, which is consistent with the findings by Huang and Tauchen (2005). The optimal sampling methods by Bandi and Russell (2006) and Zhang et al. (2005) address the bias against finding jumps and increases the power of the test statistics. These methods perform similarly to applying staggered returns, which Huang and Tauchen (2005) evaluate. Combining the optimal sampling methods with staggered returns generally leads to invalid tests.

3. Bandi and Russell (2006) give two equations for computing the optimal sampling rate; one that is exact and one approximation. The former requires an optimization routine while the second has a simple closed-form solution. I find that the two methods perform equivalently, thus there is no significant loss to use the approximation which is faster to compute.

4. I find that a modified version of the method by Bandi and Russell (2006) corrects the original tests from being slightly anti-conservative, and produces more powerful jump statistics.

5. The size and power of the test statistics are similar for the three noise processes, that is, adding serial correlation to the error process and introducing rounding errors do not have a significant impact beyond the effects of a normal iid noise process.

6. The finite sample properties under the alternative hypothesis are consistent for alternative jump distributions.

Volatility and Jump Dynamics in U.S. Energy Futures Markets

Barndorff-Nielsen and Shephard (2004, 2006) and Jiang and Oomen (2008) propose nonparametric procedures for identifying jumps in high-frequency intraday financial time series. Jiang et al. (2008) show that the methods can be combined to produce a test that remains powerful but is more robust to noise in the price series. Previously, these methods have been applied to markets such as U.S. treasury, foreign exchange and equity, but there is no empirical work using the newly developed procedures to investigate the presence of jumps over time and the relative contribution of jumps to the volatility of energy futures prices. The second essay (Chapter 3) seeks to fill this gap. I apply nonparametric methods to identify jumps in futures prices of crude oil, heating oil and natural gas contracts traded on the New York Mercantile Exchange. I document the jump components in these markets and investigate the impact on the total volatility.

Previous literature on investigating volatility behavior of energy futures prices include Pindyck (2004), Linn and Zhu (2004), Ates and Wang (2007), Mu (2007),

Wang et al. (2008) and others. Pindyck (2004) documents a significant positive trend in natural gas futures during the sample period from May 2, 1990 to February 2, 2003. Linn and Zhu (2004) report an increase in volatility before and after the release of inventory reports by the Energy Information Administration. Ates and Wang (2007) document that extreme cold weather surprises and inventory surprises are the short-run demand and supply factors that affect the spot and futures price change volatility in natural gas and heating oil markets. Mu (2007) finds that extreme weather conditions and low inventories are important factors affecting natural gas futures volatility. Wang et al. (2008) examine the realized volatility and correlation of crude oil and natural gas futures. They provide evidence that realized crude oil futures volatility increases in the weeks immediately before OPEC recommends price increases. However, none of these papers dealing with energy price volatility have separated the volatility jump component from the smooth volatility component and examined the relative importance of jump versus smooth components in the total price volatility.

This study makes several contributions to the literature on detecting jump components and in analyzing the time series properties of jumps in energy futures prices. I examine the realized volatility behavior of natural gas, heating oil and crude oil futures contracts traded on the New York Mercantile Exchange (NYMEX) using high-frequency intraday data from January 1990 to January 2008. I apply a nonparametric test statistic proposed by Jiang et al. (2008), and identify significant jump components in energy futures prices and estimate the relative contribution of jumps to the realized variance in the three futures contracts. I investigate whether significant jumps are typically associated with Energy Information Administration's inventory news announcement dates and extreme cold weather periods. I test whether including jump and seasonal components, and weather and inventory surprises as explanatory

variables improve the modeling and forecasting of energy futures volatility.

I obtain several interesting empirical results:

1. For the whole sample period, I find that the means of annualized volatility for natural gas futures, crude oil futures and heating oil futures are 39.4, 26.0 and 26.5 percent, respectively. Thus, natural gas is the most volatile among these price series. There are upward trends in volatility of the three series during the sample period; for natural gas the increase is primarily due to the jump component while the smooth component dominates the increase in the crude oil and heating oil markets. There are significant jumps (price spikes) in all three price series and the number of days with significant jumps per year ranges from 5 to 34 for natural gas, 5 to 28 for heating oil and 4 to 20 days for crude oil.

2. I document that the total realized volatility and smooth sample component for natural gas and heating oil are higher in the winter months than during the summer months. These results are consistent with the general hypothesis that when short run demand for natural gas and heating oil is suddenly shifted higher due to extreme cold weather during the winter, the short run supply is inelastic due to low inventories at this time of the year. These two factors are the ones largely responsible for generating volatility in the winter months.

3. I document in an intraday analysis that the volatility is higher during inventory news announcement periods and that many jumps are associated with these announcement dates. Furthermore, it is interesting to observe that for all markets, the volatility returns to preannouncement levels faster when there is a jump in the futures price changes than when there is no jump. The volatility remains elevated for about thirty minutes or shorter on days with a jump at

the announcement and longer otherwise.

4. I find that including the jump component as an explanatory variable improves the performance of a realized volatility forecasting model. The coefficient of the jump component attains the largest value at the daily lag and decreases for corresponding weekly and monthly regression estimates. Furthermore, all of the coefficients of jumps are negative and most are significant. The above two results indicate that the jump component in the price process produces transitory surges in volatility and that there is a strong reversal in the volatility on the subsequent days of a jump.

5. Cold weather and inventory surprises lead to an increase in volatility in natural gas and heating oil markets. Furthermore, the lagged interest-rate adjusted spread may be a suitable proxy for the negative inventory periods since the significance of the weather and inventory surprise variables drops while the spread remains highly significant when including all three variables. The spread also reduces the significance of the jump component.

Organization

The remainder of the dissertation is organized as follows. Chapter 2 presents the essay *"Detecting Jumps in Asset Prices Using Bipower Variation"*. Thereafter, Chapter 3 reports the study *"Volatility and Jump Dynamics in U.S. Energy Futures Markets"*. Chapter 4 concludes and offers directions for future work.

Bibliography

Bibliography

Aït-Sahalia, Y., Mykland, P. A., and Zhang, L. (2006). Ultra high frequency volatility estimation with dependent microstructure noise. NBER Working Paper No. W11380.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility. *Review of Economics and Statistics*, 89(4):701–720.

Ates, A. and Wang, G. H. K. (2007). Price dynamics in energy spot and futures markets: The role of inventory and weather. Presented at Financial Management Association Annual Meeting.

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

Bandi, F. M. and Russell, J. R. (2006). Separating microstructure noise from volatility. *Journal of Financial Economics*, 79:655–692.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–48.

Barndorff-Nielsen, O. E. and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1):1–30.

Cox, J. C. and Rubinstein, M. (1985). *Options Markets*. Prentice-Hall, Inc.

Duffie, D. and Pan, J. (2001). Analytical value-at-risk with jumps and credit risk. *Finance and Stochastics*, 5:155–180.

Eraker, B., Johannes, M., and Polson, N. (2003). The impact of jumps in volatility and returns. *Journal of Finance*, 53(3):1269–1300.

Fan, J. and Wang, Y. (2007). Multi-scale jump and volatility analysis for high-frequency financial data. *Journal of the American Statistical Association*, 102(480):1349–1362.

Huang, X. and Tauchen, G. (2005). The relative contribution of jumps to total price variation. *Journal of Financial Econometrics*, 3:456–499.

Jarrow, R. A. and Rosenfeld, E. R. (1984). Jump risks and the intertemporal capital asset pricing model. *Journal of Business*, 57(3):337–351.

Jiang, G. J., Lo, I., and Verdelhan, A. (2008). Information shocks and bond price jumps: Evidence from the U.S. treasury market. Working Paper.

Jiang, G. J. and Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise - A "Swap Variance" approach. *Journal of Econometrics*, 144(2):352–370.

Li, Y. and Mykland, P. A. (2007). Are volatility estimators robust with respect to modeling assumptions? *Bernoulli*, 13(3):601–622.

Linn, S. C. and Zhu, Z. (2004). Natural gas prices and the gas storage report: Public news and volatility in energy futures markets. *Journal of Futures Markets*, 24(3):283–313.

Maheu, J. M. and McCurdy, T. H. (2004). News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance*, 59:755–793.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125.

Mu, X. (2007). Weather, storage, and natural gas price dynamics: Fundamentals and volatility. *Energy Economics*, 29:46–63.

Pindyck, R. S. (2004). Volatility in natural gas and oil markets. *Journal of Energy and Development*, 30(1):1–19.

Sen, R. (2008). Jumps and microstructure noise in stock price volatility. In Gregoriou, G. N., editor, *Stock Market Volatility*. Chapman Hall-CRC/Taylor and Francis.

Wang, T., Wu, J., and Yang, J. (2008). Realized volatility and correlation in energy futures markets. *Journal of Futures Markets*, 28(10):993–1011.

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*, 100:1394–1411.

Chapter 2: Detecting Jumps in Asset Prices Using Bipower Variation

2.1 Introduction

It is essential to understand the dynamics of the volatility process for decision making in many financial applications such as derivative pricing, hedging, and portfolio rebalancing. While assuming that a continuous diffusion process that is based on the Brownian motion drives the price process simplifies the theoretical analysis as well as estimation of volatility, in practice, the observed price process for many financial assets and derivatives exhibits events that cause discontinuities or jumps in the returns. Although empirical studies often note that the impact of such jumps generally is transitory, the short-term effect in the volatility may nonetheless be considerable with important financial implications for valuation of derivatives (Merton (1976)), asset allocation (Jarrow and Rosenfeld (1984)) and risk management (Duffie and Pan (2001)).

A number of studies has shown that models including a discontinuous jump component in the return process separately from the diffusion process fits the data better than only a continuous process. Cox and Rubinstein (1985), for example, compare the Black-Scholes formula with Merton's option pricing formula (Merton (1976)) and show that with large and frequent jumps in the price process of the underlying assets, Black-Scholes significantly undervalues out-of-the-money and at-the-money options. In a more recent study, Bakshi et al. (1997) compare several parametric models with and without a jump component based on model fitting, pricing, and hedging. They

report that it is essential to include a jump component for pricing and internal consistency. Eraker et al. (2003) observe that jumps in the returns occur less frequently than what is reported in most literature but are nevertheless still significant. Maheu and McCurdy (2004) assume that jumps in stock market returns are generated by a nonhomogenous Poisson process and find that the model improves forecasts of volatility.

These studies are all based on parametric models with smooth and jump components. Recently, motivated by an increase in the availability of high-frequency or transaction data, a growing literature on nonparametric methods has emerged. Barndorff-Nielsen and Shephard (2004b, 2006) propose a number of statistics based on realized power variations to test for jumps and to estimate the contribution of jumps to the total variation. Fan and Wang (2007) develop a method based on wavelets. Another nonparametric statistic is proposed by Jiang and Oomen (2008), which relies on the asymptotic differences between logarithmic returns and percentage returns. Sen (2008) bases a jump test on principal component analysis. These nonparametric approaches are all based on intraday data while most parametric models are applied to daily observations. Clearly, intraday data is richer in information and thus presumably may produce more efficient estimates. The utilization of intraday data, however, is hampered by the presence of market microstructure noise. Such frictions come from trade mechanisms and rules that govern the markets. Measurement errors arise, for example, due to price rounding and stale prices while the bid-ask spread and minimum tick size discretize the prices that the models typically assume are continuous. Furthermore, the bid-ask spread leads to negative serial correlation as the traded prices fluctuate around the fair price, while the practice of large traders to split their orders into smaller trades to hide information generates positive serial correlation. On a daily or longer time horizon, such noise is small compared to the

volatility due to information, but may dominate at high intraday sampling frequencies. Consequently, statistics based on transaction data need to contend with the effects of such noise.

In this study, I evaluate nonparametric statistics by Barndorff-Nielsen and Shephard (2004b, 2006) that are applied to identify days with jumps in the price process. The main contribution is to evaluate the finite sample properties of the tests on noisy prices and particularly whether recently proposed methods for reducing the impact of noise improve the tests. Huang and Tauchen (2005) carry out a simulation study on these statistics, which I extend in a number of directions. Initially I consider the same cases as Huang and Tauchen (2005), but I provide a more thorough examination as to how the market microstructure noise impacts the jump statistics. Second, Bandi and Russell (2006) and Zhang et al. (2005) recently proposed methods for reducing the impact of noise for estimating the daily integrated variance using high-frequency data. I carry out an extensive empirical simulation study to determine whether the optimal sampling rates for computing the realized variation also apply for other intraday variations, specifically, for the bipower and tripower variations, which both are variables in the nonparametric jump statistics. Third, I apply the methods by Bandi and Russell (2006) and Zhang et al. (2005) to test for jumps in price processes with noise. Moreover, I combine these methods with using staggered returns, which previously have been applied in the literature (see Andersen et al. (2007) and Huang and Tauchen (2005)). The methods by Bandi and Russell (2006) and Zhang et al. (2005) have not, to the best of by knowledge, previously been applied to the jump statistics.[1] I also propose and evaluate a modified version of the method by Bandi and Russell (2006) to make it more robust to jumps. Fourth, I consider two alternative noise processes recently proposed by Aït-Sahalia et al. (2006) and Li and Mykland

[1]See Andersen et al. (2007), for example, who call attention to the lack of such a study.

(2007), and study the size and power of the jump test statistics under these processes. These have not, to the best of my knowledge, previously been evaluated in this context. I also evaluate the statistics for alternative jump distributions. Specifically, I generate jumps from normal, skewed-normal and double exponential distributions. Finally, while the methods by Bandi and Russell (2006) and Zhang et al. (2005) are based on sampling the price process in an optimal manner to lessen the bias due to noise, they nevertheless discard a large fraction of the data. I empirically evaluate another method developed by Zhang et al. (2005) that uses all data to estimate the daily integrated variations.

I obtain several interesting results. Noise biases the statistics against identifying jumps, which is consistent with the findings by Huang and Tauchen (2005). The optimal sampling methods by Bandi and Russell (2006) and Zhang et al. (2005) address the bias against finding jumps and increases the power of the test statistics. These methods perform similarly to applying staggered returns, which Huang and Tauchen (2005) evaluate. Combining the optimal sampling methods with staggered returns generally leads to invalid tests. Second, Bandi and Russell (2006) give two equations for computing the optimal sampling rate; one that is exact and one approximation. The former requires an optimization routine while the second has a simple closed-form solution. I find that the two methods perform equivalently, thus there is no significant loss to use the approximation which is faster to compute. Third, I find that a modified version of the method by Bandi and Russell (2006) produces more powerful jump statistics. Fourth, the size and power of the test statistics are similar for the three noise processes that I consider, that is, adding serial correlation to the error process and introducing rounding errors do not have a significant impact beyond the affects of a normal iid noise process.

In Section 2.2, I describe the underlying theoretical framework and review the

nonparametric jump statistics by Barndorff-Nielsen and Shephard (2004b, 2006). I also describe methods to reduce bias due to market microstructure noise. Section 2.3, thereafter, sets up the experimental design of the simulation study followed by the empirical results in Section 2.4. Finally, Section 2.5 concludes the work.

2.2 Review of Nonparametric Test Statistics

This section provides a background of the nonparametric procedure to test for jumps in asset prices by Barndorff-Nielsen and Shephard (2004b, 2006). Thereafter, I discuss the sources and implications of market microstructure noise. In particular, I review recent advances of methods to address the bias in estimating realized variations in contaminated prices.

2.2.1 Power Variations and Jump Test Statistics

Let $X_t = \log S_t$ denote the logarithmic price where S_t is the observed price at time t. Assume that the logarithmic price process, X_t, follows a continuous-time diffusion process coupled with a discrete process defined as,

$$\mathrm{d}X_t = \mu_t \mathrm{d}t + \sigma_t \mathrm{d}W_t + \kappa_t \mathrm{d}q_t, \tag{2.1}$$

where μ_t is the instantaneous drift process and σ_t is the diffusion process; W_t is the standard Wiener process; q_t is a counting process with intensity λ_t, that is, $P(\mathrm{d}q_t = 1) = \lambda_t \mathrm{d}t$; and κ_t is the size of the price jump at time t if a jump occurred. If X_{t-} denotes the price immediately prior to the jump at time t, then $\kappa_t = X_t - X_{t-}$.

Define the *intraday return*, r_{t_j}, as the difference between two logarithmic prices,

$$r_{t_j} = X_{t_j} - X_{t_{j-1}}, \tag{2.2}$$

where t_j denotes the jth intraday observation on the tth day. Importantly, X_{t_j} and $X_{t_{j-1}}$ are not necessarily two subsequently observed logarithmic prices. Let Δ denote the discrete intraday sample period of length, $t_j - t_{j-1}$. Then, X_{t_j} is the observed price at time $t_{j\Delta}$ where Δ is assumed to be constant.

The nonparametric jump statistics are based on the difference between two estimators of the daily integrated variation. The *realized variance* is defined as the sum of squared intraday returns,

$$\mathrm{RV}_t = \sum_{j=1}^{m_t} r_{t_j}^2, \tag{2.3}$$

where m_t is the number of Δ-returns during the tth time horizon (such as a trading day) and is assumed to be an integer. Jacod and Shiryaev (1987) show that the realized (quadratic) variation converges to the integrated variation assuming that the underlying process follows equation (2.1) without jumps ($\lambda = 0$). Furthermore, in the presence of jumps ($\lambda > 0$), the realized volatility converges in probability to the total variation as $\Delta \to 0$,

$$\mathrm{RV}_t \xrightarrow{p} \int_{t-1}^{t} \sigma_s^2 \mathrm{d}s + \sum_{t<s<t+1} \kappa^2(s). \tag{2.4}$$

Hence, the realized variation captures the effects of both the continuous and the discrete processes where the first term in equation (2.4) is the return variation from the diffusion process and the second term is due to the jump component.

The second estimator of the integrated variance is the *realized bipower variation*,

which is defined as,

$$\mathrm{BV}_t = \mu_1^{-1} \frac{m_t}{m_t - 1} \sum_{j=2}^{m_t} |r_{t_j}||r_{t_{j-1}}|, \tag{2.5}$$

where μ_1 is a constant given by,

$$\mu_k = \frac{2^{k/2}}{\sqrt{\pi}} \Gamma\left(\frac{k+1}{2}\right), \tag{2.6}$$

where Γ is the Gamma function. Barndorff-Nielsen and Shephard (2004b) show that as $\Delta \to 0$,

$$\mathrm{BV}_t \xrightarrow{p} \int_{t-1}^{t} \sigma_s^2 \mathrm{d}s, \tag{2.7}$$

where the underlying price process is defined by the jump-diffusion process in equation (2.1). The result follows from that only a finite number of terms in the sum in equation (2.5) are affected by jumps while the remaining returns go to zero in probability. Since the probability of jumps goes to zero as $\Delta \to 0$, those terms do not impact the limiting probability. Hence, the asymptotic convergence of the bipower variation captures only the effects of the continuous process even in the presence of jumps. Importantly, this result is robust in that it does not make any additional assumptions regarding the counting process, the jump size distribution, and the relationship between the jump process and the volatility component, σ_t. By combining the results from equations (2.4) and (2.7), the contribution of the jump process in the total quadratic variation can be estimated by the difference between these two variations where,

$$\mathrm{RV}_t - \mathrm{BV}_t \xrightarrow{p} \sum_{t<s<t+1} \kappa^2(s), \tag{2.8}$$

as $\Delta \to 0$. Hence, equation (2.8) estimates the integrated variation due to the jump

component and, as such, provides the basis for a nonparametric statistic for identifying jumps.

Barndorff-Nielsen and Shephard (2004b, 2006) and Barndorff-Nielsen et al. (2006) show that in the absence of jumps in the price process,

$$\Delta^{-1/2}\frac{\mathrm{RV}_t - \mathrm{BV}_t}{\left(\left(\nu_{bb} - \nu_{qq}\right)\int_{t-1}^{t}\sigma^4(s)\mathrm{d}s\right)^{1/2}} \xrightarrow{p} \mathrm{N}(0,1), \tag{2.9}$$

as $\Delta \to 0$ where RV_t and BV_t are defined in equations (2.3) and (2.5) and $\nu_{bb} = \pi^2/2 + \pi - 3$ and $\nu_{qq} = 2$. The integral in the denominator, called the *integrated quarticity*, is unobservable. From the work by Barndorff-Nielsen and Shephard (2004b) on multipower variations, Andersen et al. (2007) propose to estimate the integrated quarticity using the *realized tripower quarticity*, TP_t, which is defined as,

$$\mathrm{TP}_t = m_t\mu_{4/3}^{-3}\frac{m_t}{m_t - 2}\sum_{j=3}^{m_t}\prod_{i=0}^{2}|r_{t_{j-i}}|^{4/3}, \tag{2.10}$$

where $\mu_{4/3}$ is defined in equation (2.6). Asymptotically, as $\Delta \to 0$,

$$\mathrm{TP}_t \xrightarrow{p} \int_{t-1}^{t}\sigma_s^4\mathrm{d}s. \tag{2.11}$$

Another estimator of the integrated quarticity from Barndorff-Nielsen and Shephard (2004b) is the *realized quadpower quarticity*, QP_t,

$$\mathrm{QP}_t = m_t\mu_1^{-4}\frac{m_t}{m_t - 3}\sum_{j=4}^{m_t}\prod_{i=0}^{3}|r_{t_{j-i}}|, \tag{2.12}$$

where μ_1 is given by equation (2.6). Hence, a test statistic based on equation (2.9) is

given by,

$$\Delta^{-1/2}\frac{\mathrm{RV}_t - \mathrm{BV}_t}{\left(\left(\nu_{bb} - \nu_{qq}\right)\mathrm{TP}_t\right)^{1/2}}, \tag{2.13}$$

where QP_t provides an alternative to TP_t.

Statistics

Barndorff-Nielsen and Shephard (2004b, 2006) propose a number of variations of the statistic in equation (2.13), all of which asymptotically have a standard normal distribution. A logarithmic form of the statistic is given by,

$$Z_{\mathrm{TPL},t} = \frac{\log(\mathrm{RV}_t) - \log(\mathrm{BV}_t)}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\frac{\mathrm{TP}_t}{\mathrm{BV}_t^2}}}, \tag{2.14}$$

and a similar form with an added maximum adjustment due to a Jensen's inequality argument (Barndorff-Nielsen and Shephard (2004a)),

$$Z_{\mathrm{TPLM},t} = \frac{\log(\mathrm{RV}_t) - \log(\mathrm{BV}_t)}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\max\left\{1, \frac{\mathrm{TP}_t}{\mathrm{BV}_t^2}\right\}}}. \tag{2.15}$$

Analogous statistics are given based on the quadpower variation, QP_t (equation (2.12)),

$$Z_{\mathrm{QP},t} = \frac{\mathrm{RV}_t - \mathrm{BV}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\mathrm{QP}_t}}, \tag{2.16}$$

$$Z_{\mathrm{QPL},t} = \frac{\log(\mathrm{RV}_t) - \log(\mathrm{BV}_t)}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\frac{\mathrm{QP}_t}{\mathrm{BV}_t^2}}}, \tag{2.17}$$

and,

$$Z_{\text{QPLM},t} = \frac{\log(\text{RV}_t) - \log(\text{BV}_t)}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\max\left\{1, \frac{\text{QP}_t}{\text{BV}_t^2}\right\}}}. \tag{2.18}$$

Andersen et al. (2007) and Huang and Tauchen (2005) favor replacing the logarithmic difference between RV_t and BV_t in the statistics above with the ratio,

$$\text{RJ}_t = \frac{\text{RV}_t - \text{BV}_t}{\text{RV}_t}. \tag{2.19}$$

Notice that the ratio, RJ_t, is an estimator of the relative contribution of the jump component to the total variance since the difference between RV_t and BV_t estimates the jump component and RV_t estimates the total variance. The following four statistics are based on the ratio, RJ_t,

$$Z_{\text{TPR},t} = \frac{\text{RJ}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\frac{\text{TP}_t}{\text{BV}_t^2}}}, \tag{2.20}$$

$$Z_{\text{TPRM},t} = \frac{\text{RJ}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\max\left\{1, \frac{\text{TP}_t}{\text{BV}_t^2}\right\}}}, \tag{2.21}$$

$$Z_{\text{QPR},t} = \frac{\text{RJ}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\frac{\text{QP}_t}{\text{BV}_t^2}}}, \tag{2.22}$$

and,

$$Z_{\text{QPRM},t} = \frac{\text{RJ}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\max\left\{1, \frac{\text{QP}_t}{\text{BV}_t^2}\right\}}}. \tag{2.23}$$

Hypothesis

I apply these statistics to test the null hypothesis that there is no jump in the return process during an interval t, where the hypothesis is rejected for large values of the statistics relative to the standard normal distribution. The test is one-sided since the statistics are based on the difference between two variances where the difference is zero under the null hypothesis and greater than zero otherwise. Importantly, the alternative hypothesis is the finding of *detectable* jumps. Small jumps relative to the diffusion or noise processes are unlikely to be discernible.

2.2.2 Methods to Contend with Market Microstructure Noise

This subsection discusses the implications of market microstructure noise. I briefly discuss the sources of such market frictions and thereafter focus on methods to limit the effects.

Optimal Sampling Rate: Bandi and Russell

The test statistics in Section 2.2.1 rely on estimates of integrated variations, which are obtained with model-free methods on high-frequency intraday data. The asymptotic results hinge on efficient (noise-free) price processes. Observed prices, however, are noisy due to market microstructure. Thus, the variation in intraday returns can be attributed to two components: the efficient price returns and the microstructure frictions. The variance generated by market frictions is the result of price formation under specific trade mechanisms and rules, such as discrete price grids and bid-ask bounce effects. Such noise introduces bias in the variance estimates, which becomes particularly severe at high sampling rates. The variance due to noise rather than the integrated variance will dominate the estimate as the sampling interval goes to zero.

One approach that is used in the applied literature to alleviate the bias is simply to sample the price process at lower frequencies than what the data permits. The sampling intervals are typically arbitrarily chosen and commonly in the range of five to thirty minutes. Bandi and Russell (2006), however, propose a method that finds an optimal sampling rate for estimating the realized volatility.

Let X_{t_j} denote the (unobservable) efficient logarithmic price, and define the noisy logarithmic price process, Y_{t_i}, which is observed in the market by,

$$Y_{t_j} = X_{t_j} + \epsilon_{t_j}, \tag{2.24}$$

where ϵ_{t_j} denotes the microstructure noise process. The observed returns, $\widetilde{r}_{t_j}$, are then given by,

$$\widetilde{r}_{t_j} = Y_{t_j} - Y_{t_{j-1}} = r_{t_j} + \eta_{t_j}, \tag{2.25}$$

where as before r_{t_j} denotes the efficient returns,

$$r_{t_j} = X_{t_j} - X_{t_{j-1}}. \tag{2.26}$$

The microstructure noise in the observed return process is given by,

$$\eta_{t_j} = \epsilon_{t_j} - \epsilon_{t_{j-1}}. \tag{2.27}$$

The random shocks, ϵ_{t_j}, are assumed to be iid with mean zero and variance σ_ϵ^2. Furthermore, the true price return process, r_{t_j}, and the noise process, ϵ_{t_j}, are assumed to be independent. The noise component in the return process, η_{t_j}, has a moving average structure of order one as defined in equation (2.27). Hence, higher-order

serial correlations are under this model assumption restricted to zero. Bandi and Russell (2006) argue that these assumptions are valid in decentralized markets where the random arrival of trade requests generate noisy prices that are approximately independent. In a single specialist market structure, however, the appropriateness may be questionable since autocovariances of orders higher than one may be non-zero. They claim that even under such circumstances the impact by the improper model is marginal.

Under the assumptions imposed on the efficient price process and the market structure, they show that efficient returns are of order $\mathrm{O}(\sqrt{\Delta})$. The result follows from the definition of the true price returns in equation (2.26) and the properties of the standard Brownian motion. Meanwhile, the microstructure noise, η_{t_j}, is of order O(1). The independence from the time duration in the microstructure noise component is motivated by that adjustments of observed prices (such as the bid-ask spread) are fixed in size regardless of how short the time interval is.

Hence, the variance in the noise component dominates the realized variance estimate when the returns are sampled at high frequencies. For lower frequencies, however, the noise component is small compared to the variance in the efficient return process. As a result, high frequencies can be used to estimate the noise component, σ_ϵ^2, while the integrated variance of the underlying efficient price, $\int_{i-1}^{i} \sigma_s^2 \mathrm{d}s$, can be estimated at lower frequencies.

By equation (2.25), summing the squared observed returns over the daily subperiods gives,

$$\sum_{j=1}^{m_t} \widetilde{r}_{t_i}^2 = \sum_{j=1}^{m_t} r_{t_i}^2 + 2\sum_{j=1}^{m_t} r_{t_i}\eta_{t_i}. + \sum_{j=1}^{m_t} \eta_{t_i}^2. \tag{2.28}$$

Hence, for short sampling intervals, Δ, the true return process component vanishes

while the microstructure noise component converges to the second moment by equation (2.29), which Bandi and Russell (2006) formalize in Proposition 1a on page 661. They conclude that the second moment of the noisy price returns, $E(\eta_t^2)$, can be consistently estimated by,

$$\frac{\sum_{j=1}^{m_t} \widetilde{r}_{t_j}^2}{m_t} \underset{m_t \to \infty}{\xrightarrow{p}} E(\eta_t^2), \tag{2.29}$$

where the price process should be sampled as frequently as possible. By the assumptions of iid random shocks, η_{t_j}, it follows that,

$$E\left(\eta_t^2\right) = E\left((\epsilon_t - \epsilon_{t-1})^2\right) = E\left(\epsilon_t^2 - 2\epsilon_t\epsilon_{t-1} + \epsilon_{t-1}^2\right) = 2E\left(\epsilon^2\right), \tag{2.30}$$

since $E(\epsilon_t\epsilon_{t-1}) = 0$. Hence,

$$\frac{\sum_{j=1}^{m_t} \widetilde{r}_{j,i}^2}{2m_t} \underset{m_t \to \infty}{\xrightarrow{p}} E(\epsilon^2). \tag{2.31}$$

As noted above, the accumulated noise dominates the realized variance at high sampling rates, whereas at lower sample rates the variance of the efficient price process is proportionally larger compared to the component due to noise. An optimal sampling rate is obtained by minimizing the conditional mean-square error (MSE), which Bandi and Russell (2006) show can be written as,

$$E\left(\sum_{j=1}^{m_t} \widetilde{r}_{t_i}^2 - \int_{t-1}^{t} \sigma_s^2 \mathrm{d}s\right)^2 = 2\frac{1}{m_t}(Q_t + \mathrm{o}(1)) + m_t\beta + {m_t}^2\alpha + \gamma, \tag{2.32}$$

where Q_t denotes the quarticity, $\int_{t-1}^{t} \sigma^4 \mathrm{d}s$. The three other parameters are defined

as,

$$\begin{aligned}
\alpha &= \left(E\left(\eta_t^2\right)\right)^2, \\
\beta &= 2E\left(\eta_t^4\right) - 3\left(E\left(\eta_t^2\right)\right)^2, \\
\gamma &= 4E\left(\eta_t^2\right)\int_{t-1}^{t}\sigma_s^2 \mathrm{d}s - E\left(\eta_t^4\right) - 2\left(E\left(\eta_t^2\right)\right)^2.
\end{aligned}$$

The optimal number of samples, m_t, is obtained by minimizing the MSE in equation (2.32), which requires estimates of the second and fourth moments of the noise returns, η_t, and the quarticity. An estimator of $E(\eta_t^2)$ is provided by equation (2.29). Bandi and Russell (2006) show in Proposition 2a on page 662 that a daily estimator of the fourth moment is obtained by its sample moment, $\sum_{j=1}^{m_t}\widetilde{r}_{t_j}^4/m_t$. The estimator of the quarticity suggested by Barndorff-Nielsen and Shephard (2002), $m_t/3\sum_{j=1}^{m_t}\widetilde{r}_{t_j}^4$, is not consistent under market structure noise due to the accumulation of noise at high-frequencies. Therefore, Bandi and Russell (2006) recommend to sample at a conservative sampling interval (such as every fifteen minutes) to reduce the bias. While the authors recognize that such a low frequency may be conservative, they argue that the improvement of using the optimal frequency for estimating the quarticity is negligible. They summarize the results in Proposition 2a on page 662, which states that the optimal sampling interval, $\Delta_0 = 1/m_0$, is obtained by solving,

$$m_0 = m_t : 2{m_t}^3\widehat{\alpha} + {m_t}^2\widehat{\beta} - 2\widehat{Q} = 0. \tag{2.33}$$

Furthermore, Bandi and Russell (2006) show that m_0 can be approximated by,

$$m_0 \sim \left(\frac{Q_t}{(E(\eta^2))^2}\right)^{1/3}, \tag{2.34}$$

when the optimal sampling frequency is high. Notice that the approximation does

not depend on the fourth moment and has a closed-form solution. Intuitively, the approximation seems reasonable since for large estimates of the second moment of the microstructure noise component, η_{t_j} (that is, the more contaminated the series is), the lower the sampling frequency should be.

A Robust Optimal Sampling Rate

The optimal sampling frequency by Bandi and Russell (2006) is positively related to the integrated quarticity, $Q_t = \int_{t-1}^{t} \sigma^4 \mathrm{d}s$, which is particularly evident from the approximate sampling frequency given by equation (2.34) where Q_t appears in the numerator. The authors propose to estimate the quarticity by $m_t/3 \sum_{j=1}^{m_t} \widetilde{r}_{t_j}^4$, which is not robust to jumps. As a result, the method may produce a sampling frequency that is suboptimal under the alternative hypothesis. This may in turn lead to biased estimates of the power variations and consequently impact the jump statistics.

Instead, I propose to estimate the integrated quarticity by the QP_t estimator (equation (2.12)) proposed by Barndorff-Nielsen and Shephard (2004b), which is robust to jumps.

Optimal Sampling Rate: Zhang, Mykland and Aït-Sahalia

Zhang et al. (2005) independently from Bandi and Russell (2006) under slightly different assumptions suggest a similar approach to estimate an optimal sampling rate by minimizing the appropriate mean square error. Their motivation and derivation resemble closely that of Bandi and Russell (2006), hence I only report the mean square error that they minimize to find an optimal sampling rate. Let

$$[Y, Y]_t \equiv \sum_{t_i} \left(Y_{t_{i+1}} - Y_{t_i} \right)^2, \tag{2.35}$$

where Y_{t_j} is the observed noisy price process defined in equation (2.24). The MSE is then given by,

$$\begin{aligned}\text{MSE} = &\left(2m_t E\left(\epsilon^2\right)\right)^2 \\ &+ \left[4m_t E\left(\epsilon^4\right) + \left(8[Y,Y]_t E\left(\epsilon^2\right) - 2\sigma_{\epsilon^2}^2\right) + \frac{T}{m_t}\int_{t-1}^{t} 2\sigma_s^4 \mathrm{d}s\right],\end{aligned} \tag{2.36}$$

where estimators for the unobservable parameters are discussed in the previous section.

Optimal Sampling Using All Data

Zhang et al. (2005) point out that while sparse sampling helps to alleviate the bias generated by the market microstructure noise, the estimators become inefficient by discarding large fractions of data. Instead, the authors propose an estimator that uses more data without being dominated by the market microstructure noise. They obtain multiple estimates of the integrated variance per day by computing estimates from intraday return series that starts at different transactions, k. That is, create the first intraday return series by starting with the first transaction ($k = 1$) and sample the price process every Δ time unit; obtain a second intraday series by starting at the second transaction ($k = 2$) and so on for $k = 1, \ldots, K$. The mean of these K daily estimates, $[Y,Y]_t^{(\text{avg})}$, utilizes more (potentially all) data and thus reduces the variance by averaging.

The authors show that,

$$[Y,Y]_t^{(\text{avg})} \xrightarrow{L} \langle X,X\rangle_t + 2\overline{m}_t E\left(\epsilon^2\right) + \left[4\frac{\overline{m}_t}{K}E\left(\epsilon^4\right) + \frac{4}{3\overline{m}_t}\int_{t-1}^{t}\sigma_s^4 \mathrm{d}s\right]^{1/2} Z_{(\text{total})}, \tag{2.37}$$

where $\langle X,X\rangle_t$ denotes the estimand and $\overline{m}_t$ denotes the average number of samples

of the K grids where,

$$\overline{m}_t = \frac{1}{K}\sum_{k=1}^{K} m_{t,k} = \frac{m_t - K + 1}{K}. \tag{2.38}$$

The authors show that the optimal number of intraday samples, $\overline{m}_t^*$ is,

$$\overline{m}_t^* = \left(\frac{1}{6(E(\epsilon^2))^2}\int_{t-1}^{t}\sigma_s^4 \mathrm{d}s\right)^{1/3} = \left(\frac{\xi_t}{8(E(\epsilon^2))^2}\right)^{1/3}, \tag{2.39}$$

where,

$$\xi_t^2 = \frac{4}{3}\int_{t-1}^{t}\sigma_s^4 \mathrm{d}s. \tag{2.40}$$

Staggered Returns

Andersen et al. (2007), Barndorff-Nielsen and Shephard (2006) and Huang and Tauchen (2005) evaluate a different approach to reduce the impact of microstructure noise. Specifically, the method addresses the bias generated by spurious correlations in the returns due to noise, such as the bid-ask bounce, which generates negative correlations as trades are executed at the spread slightly above and below the fair value. Moreover, the practice to split large trades into several smaller trades executed over a relatively short horizon may induce positive serial correlation. Any correlation structure in the returns may bias the bipower, tripower and quadpower estimators since these are functions of adjacent returns. The method, referred to as *staggered returns*, attempts to break up or at least reduce the correlation structure by skipping one or more returns when computing the estimators. The bipower variation using staggered returns becomes,

$$\mathrm{BV}_{t+i} = \frac{\pi}{2}\frac{m_t}{m_t - 1 - i}\sum_{j=2+i}^{m_t}|r_{t_j}||r_{t_{j-1-i}}|. \tag{2.41}$$

The offset, i, is chosen based on the order of the autocorrelation in the return process;

if the autocorrelation is only significant at one lag, an offset $i = 1$ is sufficient. Similarly, the definitions of the quarticity estimators are modified to allow for staggered returns. The staggered version of the tripower quarticity is defined by,

$$\mathrm{TP}_t = m_t \mu_{4/3}^{-3} \frac{m_t}{m_t - 2(1+i)} \sum_{j=1+2(1+i)}^{m_t} \prod_{k=0}^{2} |r_{t_{j-k(1+i)}}|^{4/3}, \tag{2.42}$$

and the quadpower quarticity is given by,

$$\mathrm{QP}_t = m_t \mu_1^{-4} \frac{m_t}{m_t - 3(1+i)} \sum_{j=1+3(1+i)}^{m_t} \prod_{k=0}^{3} |r_{t_{j-k(1+i)}}|. \tag{2.43}$$

2.3 Design of Simulation Study

This section describes the experimental design of the simulation study. Firstly, I define the data-generation price process. Thereafter, I introduce three market microstructure noise models.

2.3.1 Data-Generating Price Processes

The setup follows Huang and Tauchen (2005), who consider a one-factor stochastic volatility jump-diffusion model written as,

$$\begin{aligned} \mathrm{d}X_t &= \mu \mathrm{d}t + \mathrm{e}^{\beta_0 + \beta_1 v_t} \mathrm{d}w_{p,t} + \kappa_t \mathrm{d}q_t, \\ \mathrm{d}v_t &= \alpha_v v_t \mathrm{d}t + \mathrm{d}w_{v,t}, \end{aligned} \tag{2.44}$$

where v_t is a stochastic volatility factor; α_v is the mean reversion parameter; and $\mathrm{d}w_p$ and $\mathrm{d}w_v$ are standard Brownian motions with correlation, ρ. q_t is a discontinuous jump process where jumps occur at a rate denoted by λ. κ_t is the size of the jumps. In the following, I refer to the model defined in equation (2.44) as SV1F for $\lambda_t = 0$,

that is, when no jumps are simulated, and SV1FJ otherwise.

Table 2.1 presents values of the parameters in the data-generating processes that I consider. The values are obtained from Huang and Tauchen (2005), who select the values based on empirical studies reported in literature.

Table 2.1: The experimental design for SV1F and SV1FJ (equation (2.44)) where the jump rate, λ, is set to zero for SV1F.

Parameter	Value
μ	0.030
β_0	0.000
β_1	0.125
α_v	-0.100
ρ	-0.620
λ	{0.000, 0.014, 0.118, 0.500, 1.000, 1.500, 2.000}
σ_{jmp}	{0.000, 0.500, . . . , 2.500}

I simulate observed prices per second from the stochastic differential equation following the Euler scheme. The number of simulated prices per interval t is equivalent to six hours and a half of trading, that is, t corresponds to a typical trading day. I compute intraday price returns for time intervals ranging from one second to thirty minutes. I assume that the number of jumps in the SV1FJ model has a Poisson distribution; hence, the interarrival times have an exponential distribution with parameter λ. Initially, I assume that the size of the jumps, κ, has a normal distribution with zero mean and variance, σ^2_{jmp}. This jump model produces the asymmetric leptokurtic features of the return distribution that is typical for market data. I also evaluate alternative distributions of the jump size; specifically, I consider double exponential and skewed normal distributions. Notice that I generate jumps in the log form of the process, X_t, equation (2.44), and not in the price series form, S_t.

Figure 2.1 graphs realizations of 10000 simulated days from SV1F (Panel A) and SV1FJ (Panel B). The values of the model parameters are listed in Table 2.1. For SV1FJ, the parameters λ and σ_{jmp} are 0.014 and 1.50, respectively. The top panel plots daily closing prices; the second panel plots daily price returns, the third panel

plots the volatility factor, v_t, and the bottom panel plots the jump component, $\kappa_t \mathrm{d}q_t$.

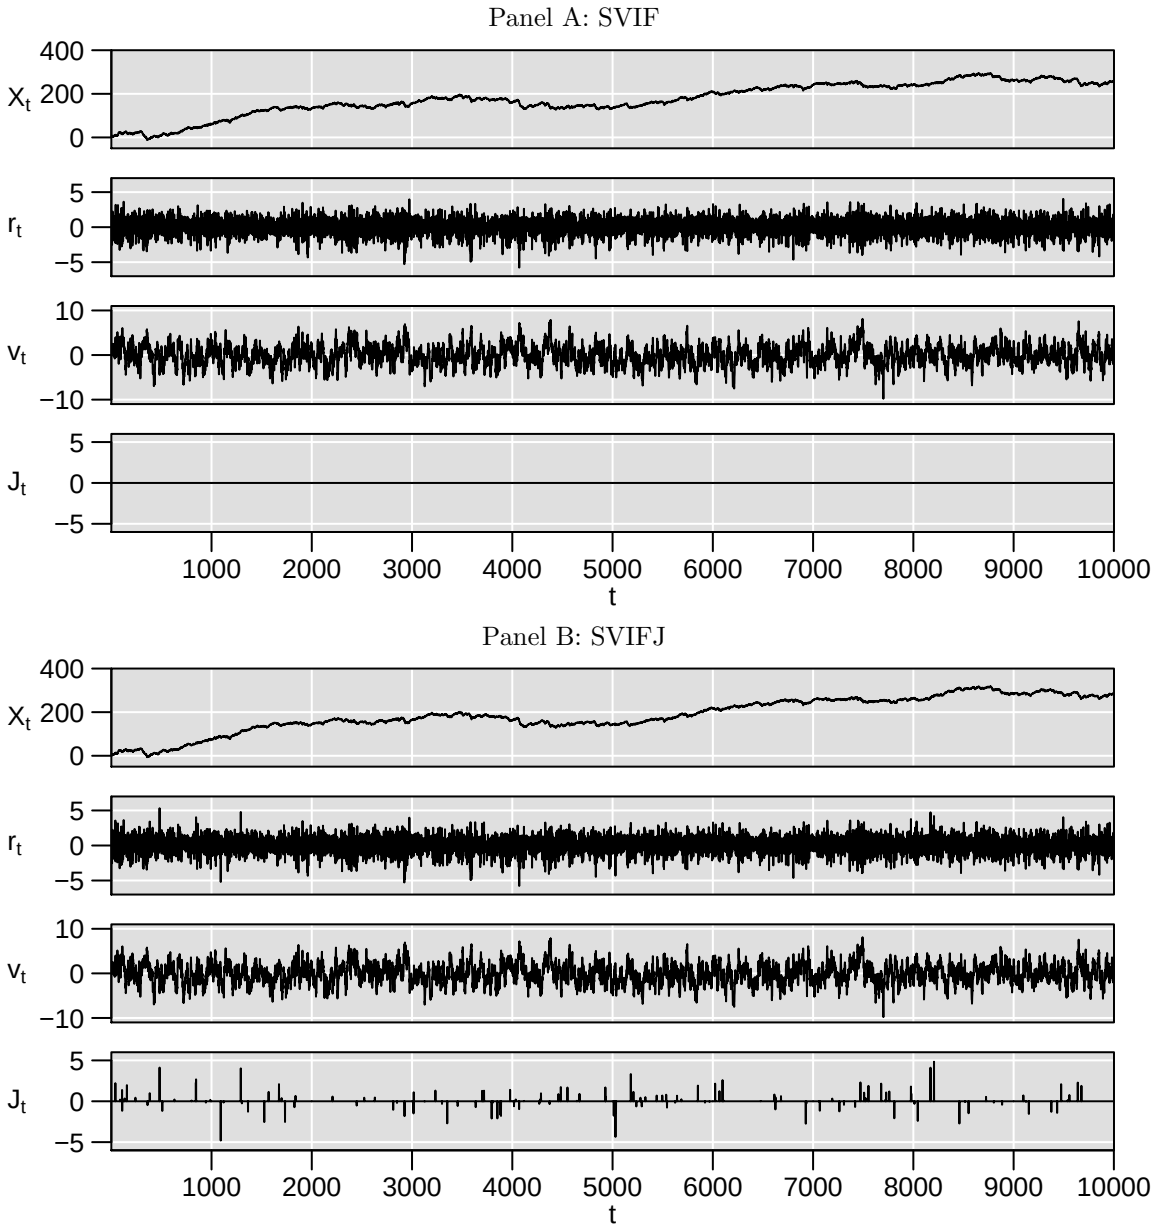

Figure 2.1: The figure plots results based on realizations of 10000 simulated days of the SV1F (Panel A) and SV1FJ (Panel B) models, equation (2.44) (page 30). The experimental design is described in Table 2.1 with $\lambda = 0.014$ and $\sigma_{\text{jmp}} = 1.50$. The top panel is the daily closing price; the second panel is the daily price returns given by the logarithmic difference between the last and first price, the third panel plots the volatility factor, v_t, and the bottom panel plots the jump process.

2.3.2 Data-Generating Process with Market Microstructure Noise

An iid normal process is a common model for market microstructure noise in the applied literature. This model also underlies the discussion in Section 2.2.2. Let Y_{t_i} denote the observed noisy price process and let ϵ_{t_i} denote the microstructure noise. Y_{t_i} is then given by,

$$Y_{t_i} = X_{t_i} + \epsilon_{t_i}. \tag{2.45}$$

where ϵ_{t_i} have a zero mean normal distribution and are independent.

In addition, I evaluate two noise processes which arguably provide more realistic settings. The models are recently proposed by Aït-Sahalia et al. (2006) and Li and Mykland (2007), respectively, and have not, to the best of my knowledge, previously been evaluated in this context. Aït-Sahalia et al. (2006) assume that the observed price, Y_{t_i}, is given by equation (2.45) but define a dependent microstructure noise model, ϵ_{t_i}, as,

$$\epsilon_{t_i} = U_{t_i} + V_{t_i}, \tag{2.46}$$

where U_{t_i} is independently identically distributed and V_{t_i} is an autoregressive process of order one, AR(1), defined by $V_{t_i} = \phi V_{t_{i-1}} + Z_{t_i}$. They assume that U and V are independent and that $Z \sim N(0, \sigma_Z^2)$ and $U \sim N(0, \sigma_U^2)$.

Li and Mykland (2007) consider another form of contamination by adding iid errors to the latent price and, thereafter, round these values to generate observed prices. Jacod et al. (2007) consider a variation of this process as well. Specifically, the price process in equation (2.45) is rounded to a multiple of γ, which is motivated by the price grid and minimum tick size rule that reside in many markets. The

observed logarithmic price is modeled by,

$$Y_{t_i} = \log\left(\gamma\left[\frac{\mathrm{e}^{X_{t_i}+\epsilon_{t_i}}}{\gamma}\right]\right), \tag{2.47}$$

where the value in the squared brackets $[\cdot]$ is rounded to the nearest integer. Hence, the expression in parentheses gives the price rounded to the nearest multiple of γ. Notice that in practice, the price cannot be allowed to drop below 0.5γ in order to impose positive values of the argument to the natural logarithm although the model does not prevent that from occurring. To account for this, I generate values by an acceptance/rejection approach where I reject price changes that yield non-positive prices.

2.4 Empirical Finite Sample Results

The empirical results are organized into four sections. First I examine the power variations that constitute the jump statistics; that is, the realized variance (RV_t), bipower variation (BV_t), tripower variation (TP_t) and ratio (RJ_t). Thereafter, I shift focus to the test statistics for jumps and examine their distribution under the null hypothesis and the convergence to asymptotic results. Finally, I estimate their size and power.

2.4.1 Power and Bipower Variations

Distributions

Before examining the properties of the jump statistics, I examine the power variations RV_t, BV_t, TP_t and RJ_t. I document their distributions and convergence to the asymptotics for efficient and noisy prices.

The results are obtained for 10000 realizations of price return series from the SV1J model as defined in equation (2.44) on page 30 without initially adding any noise to the prices. Each return series is equivalent to six and a half hours of trading. The values of the model parameters are listed in Table 2.1 on page 31.

The two first rows in Figure 2.2 plot normal QQ plots (left panel) and kernel density estimates (right panel)[2] of RV_t and BV_t. Since the data-generating process does not include discontinuous jumps, the distributions are asymptotically equal. Correspondingly, the plots show little difference. The distributions are heavy tailed and positively skewed.

The third row in Figure 2.2 plots results for the tripower estimator, TP_t, which estimates the integrated quarticity, $\int_{t-1}^{t} \sigma^4 \mathrm{d}s$. Similarly to RV_t and BV_t, the distribution is positively skewed. The density estimate for QP_t which also is an estimator of the integrated quarticity is nearly identical to TP_t and, therefore, excluded.

The ratio statistic, RJ_t, is asymptotically normal. The graphs in the fourth panel corroborates the normality although the QQ plot (left panel) reveals slightly thin tails. The continuous line in the density plot (right panel) is the density estimate for RJ_t; the dashed line is the density of a normal distribution with zero mean and variance 0.81. The densities are closely matched. Furthermore, the normality hypothesis is not rejected by any of the six tests: Anderson-Darling, Cramer-von Mises, Jarque-Bera, Lilliefors (Kolmogorov-Smirnov), Shapiro-Francia, and the Energy (Székely and Rizzo (2005)); the p-values range from 0.28 to 0.38, see Table 2.2.

To evaluate the behavior for noisy prices, I add an iid normal noise process to the SVIF model. From the discussion in Section 2.2.2, I expect the estimators to be dominated by noise, particularly at high sampling frequencies. QQ normal plots (not graphed here) do not reveal a significant change in the distributional form but

[2] I compute the density estimates in `R` (R Development Core Team (2008)) using the `density` function with a Gaussian kernel.

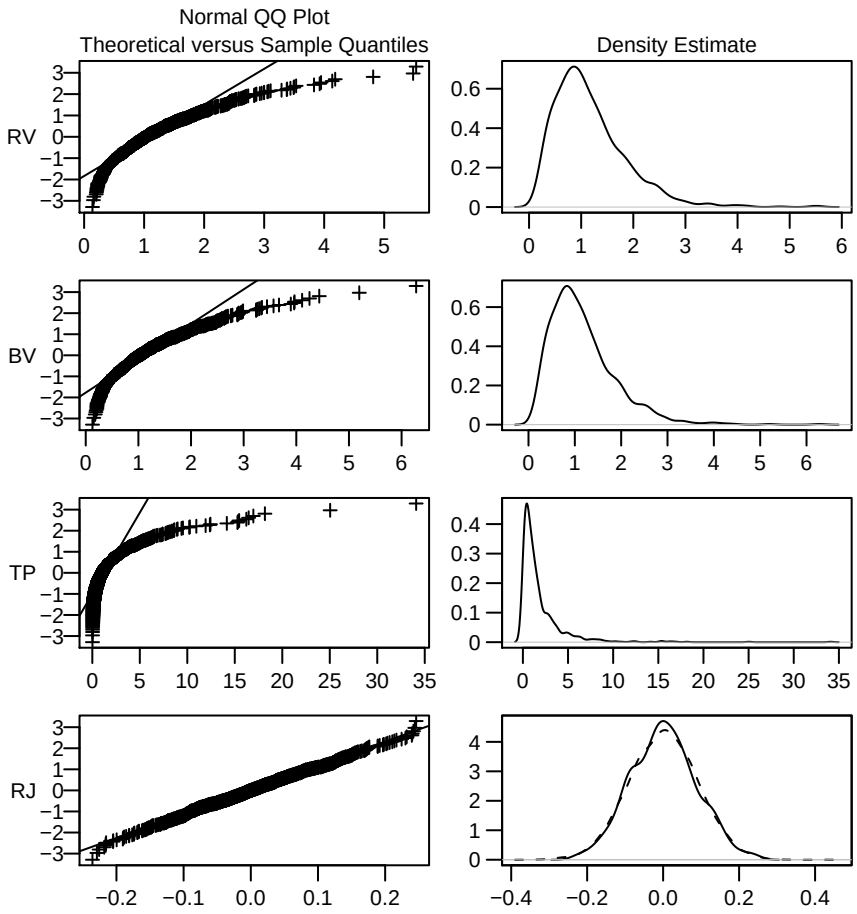

Figure 2.2: The left panel plots normal QQ plots for the realized variance, RV_t, (equation (2.3)), bipower variation, BV_t, (equation (2.5)), tripower variation, TP_t, (equation (2.10)), and ratio, RJ_t, (equation (2.19)) based on 1000 simulated days of the SV1F model, equation (2.44) (page 30). The values are in percentage form. The experimental design is described in Table 2.1 (page 31). The right panel plots kernel density estimates, which are estimated in `R` using the `density` function with a Gaussian kernel. The right panel for the RJ_t statistic also plots the density of a normal distribution with standard deviation 0.09 (dashed line). The statistic is computed based on five-minute intraday returns.

Table 2.2: The table presents p-values for normality tests of the RJ_t statistic. The column labels denote the following tests: the Anderson-Darling (AD), Cramer-von Mises (CM), Jarque-Bera (JB), Lilliefors (Kolmogorov-Smirnov) (KS), Shapiro-Francia (SF), and Energy (Energy) tests.

	AD	CM	KS	SF	JB	Energy
RJ	0.284	0.349	0.384	0.345	0.364	0.311

suggest that the location and scale may have changed. Therefore, I compute means and standard deviations of the intraday variations RV_t, BV_t and $\mathrm{RV}_t - \mathrm{BV}_t$. I include

the difference since it constitutes the basis of the jump statistics. The estimates are based on 10000 realizations.

The results in Table 2.3 indicate that the estimators are severely biased at high sampling frequencies. The first column in Panels A and B shows that RV_t and BV_t converge to about 1.17 for efficient prices. The second column presents estimates for prices with a small noise component. In spite of the small noise variance, the estimators are highly positively biased at high sampling frequencies. Moreover, the difference between the estimators (Panel C), which is nearly zero without noise, becomes negatively biased for noisy prices. The bias intensifies as the noise variance increases.

Staggered returns is proposed above as a method to address the bias. Panel D shows the difference when the bipower estimator is computed by offsetting returns by one lag. While BV_t remains biased, the difference with RV_t is nearly zero, which is the asymptotic property that the jump statistic is based on. Hence, this support that the staggered method adequately corrects for the bias due to noise.

Notice that the differences without offsetting the returns (Panel C) indicate that while they are biased at high sampling rates, they are close to zero for longer sampling intervals. The appropriate sampling frequency varies with the noise variance. The optimal sampling interval methods introduced in Section 2.2.2 attempt to find these rates. I evaluate the performance of these methods next.

Optimal Sampling Rates

Bandi and Russell (2006) and Zhang et al. (2005) derive optimal sampling frequencies for estimating the integrated variance, $\int_{t-1}^{t} \sigma_s^2 \mathrm{d}s$, using the realized variation estimator, RV_t. The jump statistics, however, also require the bipower, BV_t, and tripower, TP_t (or quadpower, QP_t) estimators, all of which are based on intraday sampling. I

Table 2.3: The table presents means and standard deviations of the realized variance (RV_t), bipower variation (BV_t), and the difference $\mathrm{RV}_t - \mathrm{BV}_t$ for 10000 price paths. The prices are generated from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. The panel label "Offset 0" denotes that no staggering is applied and "Offset 1" denotes that the intraday returns are offset by one lag for the bipower estimator.

	0.000	0.020	0.040	0.080	0.160	0.320
Panel A: RV (Offset 0)						
1	1.17	19.90	76.02	300.57	1198.93	4792.16
	(0.72)	(0.81)	(1.13)	(3.48)	(13.72)	(54.77)
2	1.17	10.55	38.59	150.87	600.12	2396.89
	(0.72)	(0.79)	(0.95)	(2.51)	(9.53)	(37.95)
4	1.17	5.87	19.88	76.03	300.72	1199.30
	(0.72)	(0.79)	(0.85)	(1.86)	(6.83)	(27.09)
8	1.17	3.53	10.52	38.59	150.94	600.19
	(0.72)	(0.78)	(0.79)	(1.42)	(4.86)	(19.14)
16	1.16	2.36	5.84	19.88	76.04	300.57
	(0.72)	(0.78)	(0.76)	(1.15)	(3.57)	(13.85)
32	1.16	1.78	3.50	10.52	38.60	150.82
	(0.72)	(0.79)	(0.75)	(0.97)	(2.59)	(9.71)
64	1.16	1.48	2.33	5.84	19.89	75.99
	(0.72)	(0.79)	(0.74)	(0.88)	(1.94)	(6.91)
128	1.16	1.33	1.74	3.49	10.51	38.46
	(0.72)	(0.80)	(0.74)	(0.82)	(1.53)	(4.95)
256	1.16	1.26	1.45	2.32	5.85	19.82
	(0.74)	(0.81)	(0.75)	(0.81)	(1.27)	(3.60)
512	1.14	1.22	1.29	1.72	3.49	10.40
	(0.76)	(0.84)	(0.76)	(0.80)	(1.14)	(2.70)
1024	1.12	1.17	1.19	1.40	2.29	5.68
	(0.79)	(0.89)	(0.80)	(0.83)	(1.09)	(2.11)
2048	1.10	1.15	1.13	1.24	1.70	3.40
	(0.87)	(0.98)	(0.88)	(0.90)	(1.12)	(1.78)
Panel B: BV (Offset 0)						
1	1.17	22.15	85.43	338.69	1351.90	5404.53
	(0.72)	(0.75)	(1.26)	(4.35)	(17.30)	(69.11)
2	1.17	11.61	43.23	169.85	676.45	2702.70
	(0.72)	(0.73)	(1.00)	(3.08)	(11.97)	(47.79)
4	1.17	6.35	22.13	85.44	338.87	1352.36
	(0.72)	(0.74)	(0.85)	(2.26)	(8.61)	(34.28)
8	1.17	3.73	11.58	43.22	169.92	676.56
	(0.72)	(0.75)	(0.77)	(1.67)	(6.08)	(24.13)
16	1.17	2.44	6.32	22.13	85.44	338.64
	(0.72)	(0.77)	(0.74)	(1.29)	(4.43)	(17.42)
32	1.16	1.80	3.71	11.58	43.21	169.72
	(0.72)	(0.78)	(0.73)	(1.05)	(3.17)	(12.18)
64	1.16	1.49	2.41	6.32	22.15	85.44
	(0.72)	(0.79)	(0.73)	(0.91)	(2.33)	(8.70)
128	1.16	1.34	1.77	3.69	11.57	43.09
	(0.73)	(0.80)	(0.74)	(0.84)	(1.79)	(6.21)
256	1.16	1.26	1.46	2.41	6.32	22.04
	(0.75)	(0.82)	(0.76)	(0.83)	(1.45)	(4.50)
512	1.15	1.22	1.29	1.75	3.69	11.44
	(0.78)	(0.86)	(0.78)	(0.83)	(1.25)	(3.33)
1024	1.12	1.17	1.19	1.41	2.37	6.15
	(0.83)	(0.91)	(0.84)	(0.88)	(1.19)	(2.57)
2048	1.10	1.15	1.13	1.24	1.74	3.60
	(0.95)	(1.05)	(0.96)	(0.99)	(1.25)	(2.15)

Table 2.3 continue

	0.000	0.020	0.040	0.080	0.160	0.320
			Panel C: RV-BV (Offset 0)			
1	0.00	−2.25	−9.42	−38.13	−152.97	−612.37
	(0.01)	(0.12)	(0.36)	(1.39)	(5.56)	(22.22)
2	0.00	−1.06	−4.63	−18.98	−76.34	−305.81
	(0.01)	(0.10)	(0.26)	(0.98)	(3.86)	(15.44)
4	−0.00	−0.48	−2.25	−9.42	−38.15	−153.06
	(0.01)	(0.08)	(0.20)	(0.71)	(2.78)	(11.09)
8	−0.00	−0.20	−1.06	−4.63	−18.98	−76.38
	(0.02)	(0.06)	(0.15)	(0.51)	(1.97)	(7.82)
16	−0.00	−0.08	−0.48	−2.25	−9.40	−38.07
	(0.03)	(0.05)	(0.12)	(0.38)	(1.40)	(5.53)
32	0.00	−0.03	−0.20	−1.07	−4.61	−18.90
	(0.04)	(0.06)	(0.10)	(0.29)	(1.00)	(3.89)
64	0.00	−0.01	−0.08	−0.48	−2.25	−9.45
	(0.06)	(0.07)	(0.10)	(0.23)	(0.75)	(2.82)
128	0.00	−0.00	−0.03	−0.20	−1.06	−4.63
	(0.08)	(0.09)	(0.11)	(0.20)	(0.57)	(2.04)
256	−0.00	−0.00	−0.01	−0.08	−0.47	−2.22
	(0.11)	(0.12)	(0.14)	(0.20)	(0.46)	(1.51)
512	−0.00	−0.00	−0.00	−0.03	−0.20	−1.05
	(0.16)	(0.18)	(0.17)	(0.22)	(0.41)	(1.15)
1024	0.00	0.00	−0.00	−0.01	−0.08	−0.47
	(0.23)	(0.24)	(0.25)	(0.28)	(0.41)	(0.95)
2048	−0.00	0.00	−0.00	−0.00	−0.03	−0.21
	(0.34)	(0.36)	(0.34)	(0.37)	(0.48)	(0.87)
			Panel D: RV-BV (Offset 1)			
1	−0.00	−0.00	−0.01	−0.02	−0.13	−0.50
	(0.01)	(0.10)	(0.39)	(1.52)	(6.09)	(24.35)
2	0.00	0.00	0.00	0.00	0.05	0.20
	(0.01)	(0.07)	(0.28)	(1.08)	(4.24)	(16.95)
4	−0.00	0.00	−0.00	−0.00	0.00	0.02
	(0.01)	(0.06)	(0.20)	(0.77)	(3.06)	(12.21)
8	0.00	−0.00	−0.00	−0.00	0.00	0.02
	(0.02)	(0.05)	(0.15)	(0.55)	(2.17)	(8.64)
16	−0.00	−0.00	0.00	0.00	−0.00	0.00
	(0.03)	(0.05)	(0.12)	(0.41)	(1.54)	(6.09)
32	−0.00	0.00	−0.00	−0.01	0.02	0.05
	(0.04)	(0.06)	(0.10)	(0.31)	(1.11)	(4.32)
64	0.00	−0.00	0.00	−0.00	−0.01	−0.04
	(0.06)	(0.07)	(0.10)	(0.24)	(0.82)	(3.12)
128	0.00	0.00	0.00	0.00	0.00	0.01
	(0.08)	(0.09)	(0.11)	(0.21)	(0.62)	(2.28)
256	−0.00	−0.00	−0.00	−0.00	0.00	−0.00
	(0.11)	(0.12)	(0.13)	(0.20)	(0.49)	(1.64)
512	−0.00	0.00	0.00	0.00	−0.00	−0.00
	(0.16)	(0.18)	(0.18)	(0.23)	(0.44)	(1.27)
1024	0.00	0.00	−0.00	−0.00	0.00	0.01
	(0.24)	(0.26)	(0.25)	(0.28)	(0.44)	(1.06)
2048	0.00	0.00	0.00	0.00	0.01	0.02
	(0.36)	(0.39)	(0.37)	(0.39)	(0.52)	(1.00)

evaluate how well the optimal sampling rates apply to these power variations.

Instead of using the stochastic model defined in Section 2.3.1, I assume that the return process, X_t, follows the geometric Brownian motion with constant drift and volatility so that the bias and mean-square error can be computed without any error.

Thus, let the data-generating process be,

$$\mathrm{d}X_t = \mu \mathrm{d}t + \sigma \mathrm{d}W_t, \tag{2.48}$$

where W_t is a standard Wiener process and the drift, μ, and volatility, σ, parameters are constant. Let Y_{t_i} denote the observed noisy price process given by,

$$Y_{t_i} = X_{t_i} + \epsilon_{t_i}, \tag{2.49}$$

where ϵ_{t_i} is normally distributed. The estimates are mean values of 1000 realized trading days per data point, where each trading day is equivalent to six and a half hours. The drift rate, μ, is zero and the volatility, σ, is one. The sampling intervals range from one to sixty minutes in increments of one minute. The results for the quadpower estimator, QP_t, are similar to those of the tripower estimator, TP_t, and are therefore excluded.

Panel A in Figure 2.3 plots the bias (first column), variance (second column) and mean square error (third column) for the realized variance (RV_t), bipower variation (BV_t) and tripower variation (TP_t) for a price process without noise. Under these conditions, the asymptotic theory states that the price process should be sampled as frequently as possible. Consistent with this, the MSE obtains its minimum at the highest frequency, that is, one minute and increases nearly linearly with the sampling interval. This is expected since the variance is negatively related to the sampling frequency.

Panel B graphs the equivalent results with the standard deviation of the noise component equal to 0.040. The pattern is consistent across all three estimators. The bias is large at high sampling frequencies but drops as the sampling interval increases by a few minutes and thereafter flattens out after about ten minutes. Similarly, as in

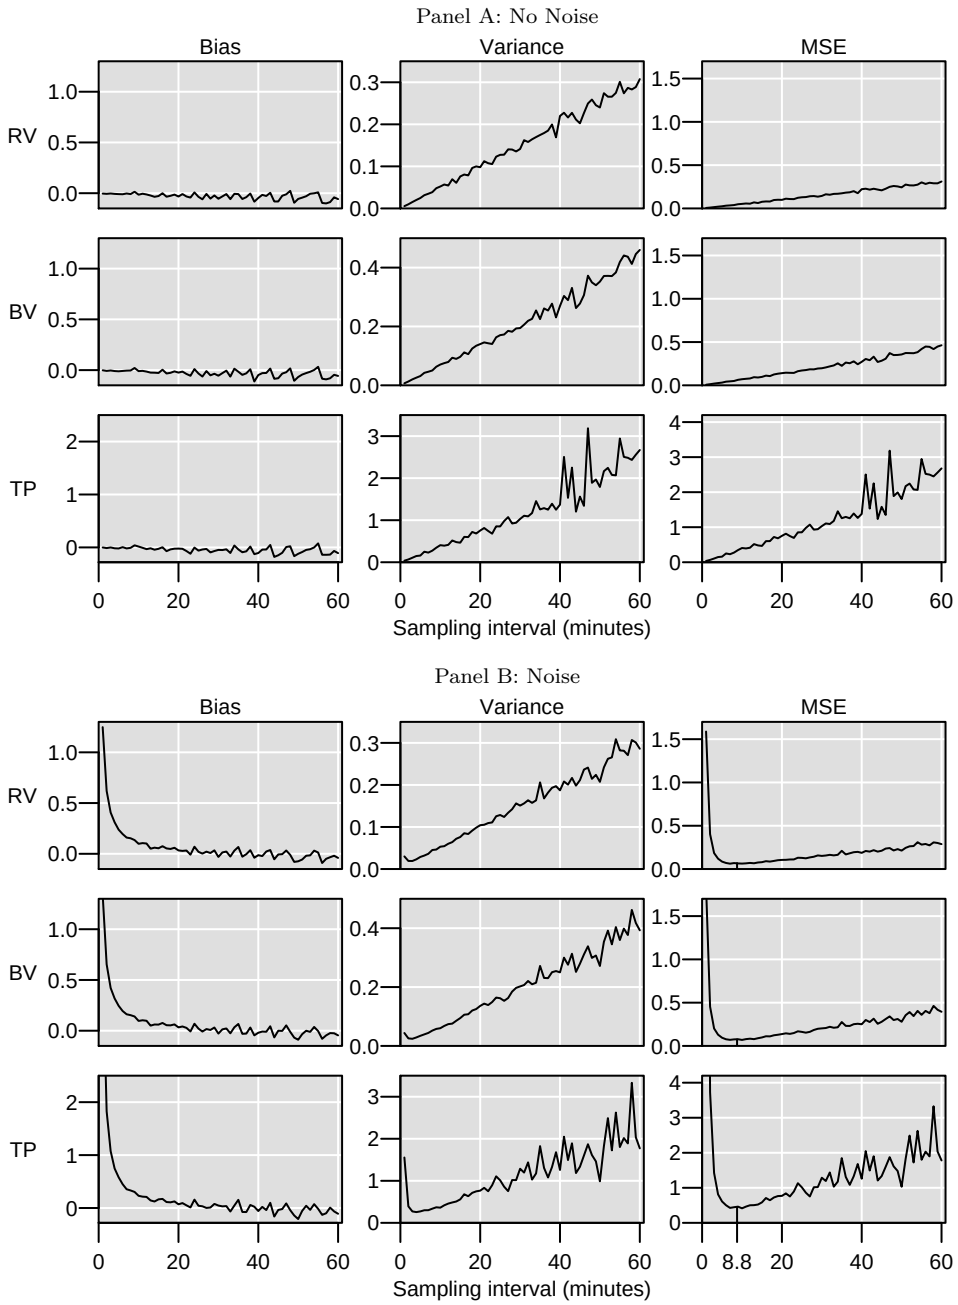

Figure 2.3: The figure plots bias, variance and mean-square error for three estimators: realized variance (RV_t), bipower (BV_t), and tripower (TP_t) variations. Prices are generated from the geometric Brownian motion model, equation (2.48), with $\mu = 0$ and $\sigma = 1$. Panel A plots results for efficient price series; Panel B presents results where an iid $\mathrm{N}(0, 0.040)$ noise component is added to the price process.

the previous case with no noise, the variance is low at the highest sampling frequencies

but increases nearly linearly as the sampling frequency drops. As a result, the mean-square error peaks at the shortest sampling interval but drops rapidly and reaches its minimum around seven to ten minutes when it turns upwards as the bias approaches zero but the variance begins to dominate. These results are consistent with the theoretical results in Section 2.2.2.

To evaluate whether the optimal sampling rates by Bandi and Russell (2006) (BR) and Zhang et al. (2005) (ZMA) coincide with these results, I estimate the optimal sampling rates and compare with the minimum point of the MSE. For the first set of results without noise, the optimal sampling intervals are about twenty seconds for BR and ZMA. Since I generate prices per second, the true optimal rate is to sample at that interval since the prices are free from noise. One explanation for this discrepancy is that by following the recommendations by Bandi and Russell (2006), I sample the price process every fifteen minutes for the tripower estimator to avoid bias due to noise, which clearly is suboptimal in this case since the price process does not contain any noise. Certainly, another explanation is that these are finite sample results while theory hinges on asymptotics.

Once I add noise to the observed prices, the MSE in Panel B in Figure 2.3 suggests that the optimal sampling interval is in the range of seven to ten minutes for all three estimators. The mean (standard deviation) of the sampling interval based on 1000 simulations using BR is 8.8 (1.8) minutes; I add a vertical line to the MSE plots to represent the BR estimate. Notice that the sampling rate given by BR coincides with the minimum of the MSE for all three estimators. ZMA produces similar rates.

Additional simulations for other values of the model parameters yield analogous results. In sum, these results provide new empirical support for that the optimal sampling rates derived for the realized variance, also are appropriate for the bipower and tripower estimators.

An important difference between the two methods to deal with noise is that staggered returns make RV_t and BV_t equally biased such that the difference is zero while the optimal sampling approaches attempt to find the sampling frequency when both are unbiased. Hence, staggered returns apply to these jump statistics, which are based on the difference between RV_t and BV_t; however, it does not correct for the bias in RV_t as an estimator of realized volatility. The optimal sampling rate methods, on the other hand, reduce the bias of the two estimators individually.

Overlapping Streams

The methods discussed thus far for dealing with noise are successful in reducing the bias; however, they discard large fractions of the data. I follow Zhang et al. (2005) (see Section 2.2.2 starting on page 28) and evaluate whether taking the mean of several estimates computed from overlapping streams reduces the mean square error.

Figure 2.4 graphs the mean-square errors for realized variance (RV_t), bipower (BV_t), and tripower (TP_t) variations for different sampling intervals. The price process is generated without noise. Importantly, the averaging method reduces the MSE for all cases and by a large percentage for some scenarios. This finding is not only influential for the jumps statistics, but for any application where these power variations are used.

2.4.2 Convergence to Asymptotics

This section examines the asymptotic properties of the jump statistics.

The left column in Figure 2.5 plots normal QQ plots of the Z_{TP}, Z_{TPLM}, and Z_{TPRM} statistics based on 1000 simulations from the SVIF model. The statistics are defined in Section 2.2 beginning on page 16. Plots for the equivalent statistics based on the quadpower (QP_t) estimator are graphed in the right column. The plot for

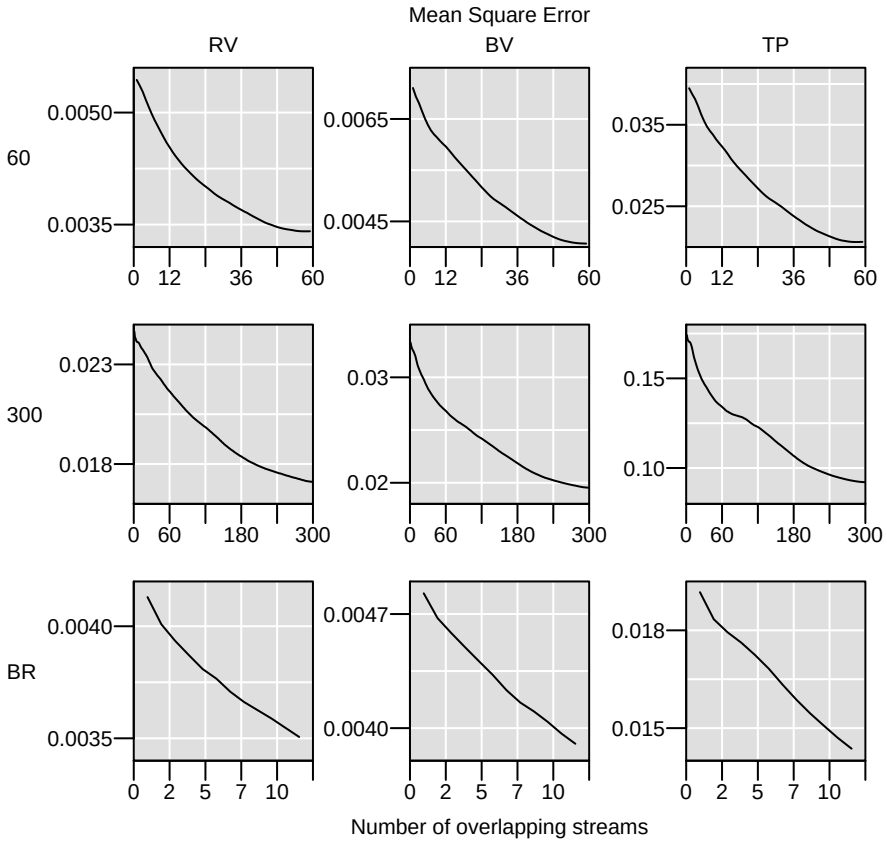

Figure 2.4: The figure graphs the mean-square error for different number of overlapping samples for three estimators: realized variance (RV_t), bipower (BV_t), and tripower (TP_t) variations. The prices are generated from the geometric Brownian motion model, equation (2.48), with $\mu = 0$ and $\sigma = 1$. The estimates in the first two rows are based on one and five-minute sampling intervals, respectively. The results in the third row are based on sampling at the interval given by Bandi and Russell (2006).

Z_{TP} indicates that its distribution has heavy tails, particularly the right tail, which is noteworthy as it is used to test for jumps. The two other statistics are both close to normal although Z_{TPLM} (Z_{QPLM}) may have a slightly heavy right tail. The QQ plots for Z_{TPRM}, however, are nearly linear. These conclusions are supported by six tests of normality. Table 2.4 shows that all tests strongly reject the normality hypothesis for Z_{TP} and Z_{TPLM} but not for Z_{TPRM}; the p-values range from 0.22 to 0.38 for Z_{TPRM}. These results confirm the findings reported by Huang and Tauchen (2005) who document that the logarithmic and ratio statistics with the maximum correction

Table 2.4: The table presents p-values for normality tests of the Z_{TP}, Z_{TPLM}, and Z_{TPRM} statistics. The column labels denote the following tests: the Anderson-Darling (AD), Cramer-von Mises (CM), Jarque-Bera (JB), Lilliefors (Kolmogorov-Smirnov) (KS), Shapiro-Francia (SF), and Energy (Energy) tests.

	AD	CM	KS	SF	JB	Energy
Z_{TP}	0.000	0.000	0.000	0.000	0.000	0.000
Z_{TPLM}	0.000	0.001	0.005	0.000	0.000	0.000
Z_{TPRM}	0.224	0.297	0.375	0.238	0.285	0.226
Z_{QP}	0.000	0.000	0.000	0.000	0.000	0.000
Z_{QPL}	0.000	0.000	0.007	0.000	0.000	0.000
Z_{QPRM}	0.154	0.236	0.382	0.146	0.162	0.183

deviate the least from the asymptotic normal distribution. The results for the two sets of statistics based on TP_t and QP_t are equivalent, which also is the case for the forthcoming results in this subsection; therefore, I limit the discussion to the statistics computed based on TP_t for the remainder of this section.

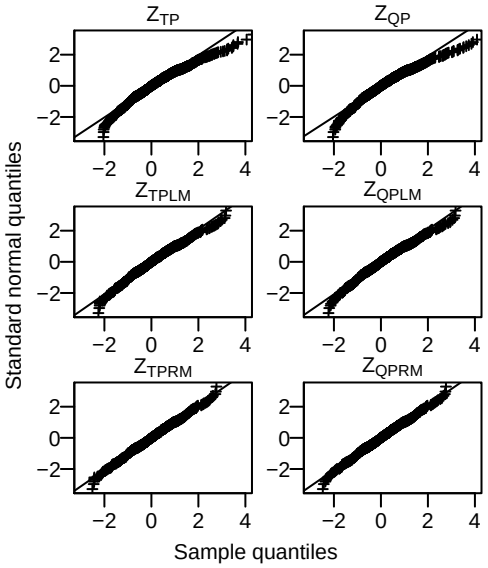

Figure 2.5: The figure graphs normal QQ plots for jump statistics based on 1000 simulated days of the SV1F model, equation (2.44) (page 30). The experimental design is described in Table 2.1 (page 31). The left column plots results for the three statistics Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The right column plots results for the equivalent statistics based on the quadpower (QP_t). The statistics are computed based on five-minute intraday returns.

I produce equivalent results for noisy price series. I particularly focus on the Z_{TPRM} statistic. I limit the discussion to the iid normal noise process and expand the study

to the other noise processes that I describe in Section 2.3.2 in later sections.

Figure 2.6 graphs normal QQ plots for three different values of the standard deviation of the noise component. I report results for σ_{mn} equal to 0, 0.052 and 0.080, respectively, in the first, second, and third row. The results in the left column are based on constant five-minute sampling intervals, which is near the optimal rates for small values of the noise variance but is likely too frequent for noisier prices. The second and third columns are results based on the optimal daily sampling rates by Bandi and Russell (2006) and Zhang et al. (2005), respectively, which I refer to as BR and ZMA. The plots are based on 1000 realized trading days.

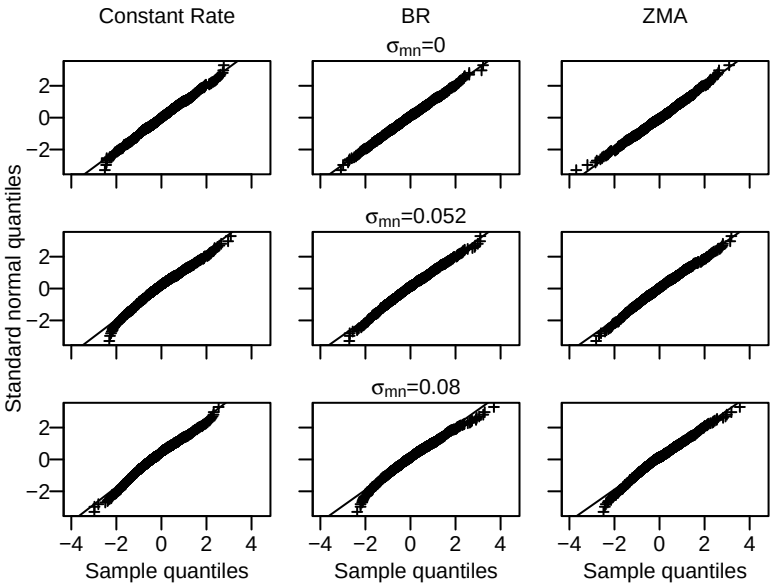

Figure 2.6: The figure plots normal QQ plots of realizations of the Z_{TPRM} statistic based on 1000 days which are simulated from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. An iid $\text{N}(0, \sigma_{mn})$ noise component is added to the price process; σ_{mn} takes the values 0.000 (first row), 0.052 (second row), and 0.080 (third row). The results in the left column are based on five-minute sampling intervals; the second column is based on sampling at the rates given by Bandi and Russell (2006) by solving equation (2.33); and the third column is based sampling at the rates given by Zhang et al. (2005).

The distributions do not appear to deviate largely from normality. The six normality test statistics do not reject the normal hypothesis for zero noise for any of the three sampling methods (Table 2.5). However, for $\sigma_{\text{mn}} = 0.052$ all statistics strongly

reject normality for constant sampling (p-value < 0.0005) but only weakly so for BR and ZMA with p-values in between 0.01 and 0.07. The tests strongly reject the normality hypothesis in all cases for the largest noise variance, $\sigma_{\text{mn}} = 0.080$.

Table 2.5: The table presents p-values for normality tests of the Z_{TPRM} statistic. The column labels denote the following tests: the Anderson-Darling (AD), Cramer-von Mises (CM), Jarque-Bera (JB), Lilliefors (Kolmogorov-Smirnov) (KS), Shapiro-Francia (SF), and Energy (Energy) tests.

	AD	CM	KS	SF	JB	Energy
			$\sigma_{\text{mn}} = 0.000$			
Constant	0.224	0.297	0.375	0.238	0.285	0.233
BR	0.921	0.857	0.855	0.975	0.662	0.912
ZMA	0.541	0.446	0.262	0.521	0.323	0.533
			$\sigma_{\text{mn}} = 0.052$			
Constant	0.000	0.000	0.001	0.000	0.000	0.000
BR	0.019	0.023	0.037	0.010	0.007	0.017
ZMA	0.032	0.061	0.073	0.019	0.012	0.045
			$\sigma_{\text{mn}} = 0.080$			
Constant	0.000	0.000	0.000	0.000	0.000	0.000
BR	0.000	0.000	0.006	0.000	0.000	0.000
ZMA	0.000	0.000	0.000	0.000	0.000	0.001

Figure 2.7 plots normal QQ plots when applying staggered returns with one lag offset. The normality tests now reject the null hypothesis for the constant sampling rate for zero noise while the p-values for BR and ZMA remain large (Table 2.6). For $\sigma_{\text{mn}} = 0.052$, the statistical inferences remain unchanged; however, the p-values for constant sampling increases slightly ranging from 0.005 (Anderson-Darling) to 0.053 (Jarque-Bera) while the p-values decrease for the two other sampling methods. The trend continues for the largest noise variance where the p-values for BR and ZMA are less than 0.02 while the normal hypothesis is no longer rejected for the constant sampling rate. The following subsection examines in more detail how the noise impacts the statistics.

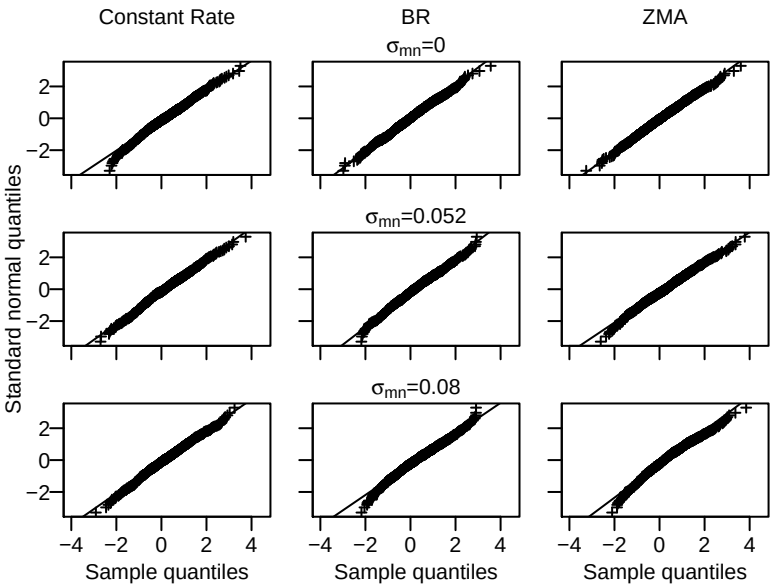

Figure 2.7: The figure graphs normal QQ plots of realizations of the Z_{TPRM} statistic based on 1000 trading days which are simulated from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. An iid $\text{N}(0, \sigma_{mn})$ noise component is added to the price process; σ_{mn} takes the values 0.000 (first row), 0.052 (second row), 0.080 (third row). The results in the left column are based on five-minute sampling intervals; the second column is based on sampling at the rates given by Bandi and Russell (2006) by solving equation (2.33); and the third column is based on sampling at the rates given by Zhang et al. (2005). The bipower and tripower estimates are based on staggered returns with one lag offset ($i = 1$).

Table 2.6: The table shows p-values for normality tests of the Z_{TPRM} statistics using staggered returns with one lag offset. The column labels denote the following tests: the Anderson-Darling (AD), Cramer-von Mises (CM), Jarque-Bera (JB), Lilliefors (Kolmogorov-Smirnov) (KS), Shapiro-Francia (SF), and Energy (Energy) tests.

	AD	CM	KS	SF	JB	Energy
			$\sigma_{\text{mn}} = 0.000$			
Constant	0.009	0.012	0.033	0.009	0.015	0.010
BR	0.128	0.183	0.383	0.227	0.335	0.155
ZMA	0.562	0.597	0.639	0.434	0.188	0.624
			$\sigma_{\text{mn}} = 0.052$			
Constant	0.007	0.009	0.005	0.033	0.053	0.005
BR	0.031	0.043	0.081	0.028	0.030	0.036
ZMA	0.274	0.379	0.510	0.089	0.098	0.309
			$\sigma_{\text{mn}} = 0.080$			
Constant	0.079	0.142	0.300	0.088	0.077	0.106
BR	0.000	0.001	0.022	0.000	0.001	0.001
ZMA	0.000	0.000	0.000	0.000	0.000	0.000

Asymptotics

This section documents the convergence to the asymptotics of the jump statistics. I assume a fixed trading time, T, and investigate the distribution of the Z_{TPRM} statistic as the sampling interval, Δ, approaches zero or, equivalently, as the number of

intraday return samples, m_t, goes to infinity since $m_t = T/\Delta$.

I examine the distribution of the Z_{TPRM} statistic as $\Delta \to 0$ by generating 100 trading days from the SV1F model and compute the jump statistics per day. Thereafter, I compute the p-value from the Kolmogorov-Smirnov test of normality.[3] I repeat these steps 100 times and examine the distribution of the p-values, which is uniform under the asymptotic theory. I produce these results for sampling intervals, Δ, ranging from one second to about thirty minutes; specifically, $\Delta = 1, 2, 4, \ldots, 2048$ seconds.

Panel A in Figure 2.8 plots histograms for prices without noise. The labels specifies the sampling interval, Δ, in seconds. The distributions appear to be uniform with the exception for the longest sampling intervals (1024 and 2048). Panel B graphs histograms for a price process with a relatively severe noise component, NIID$(0, 0.160)$. The pattern is consistent in that there are proportionally too many small p-values at the lower sampling frequencies. Interestingly, in spite of the severe noise process, the distribution converges to normal distribution for the high frequency estimates; even for the one-second sampling interval.

Since the asymptotic results show that the statistic in most cases appear to be normal, I examine the mean and standard deviations of the jump statistic; recall that their asymptotic values are zero and one, respectively. Panel A in Table 2.7 presents the mean and standard deviation of 10000 realizations of the Z_{TPRM} statistic. The row labels denote the sampling interval in seconds, and the column labels denote the standard deviation of the noise process. The means remain close to zero for efficient prices (first column); however, even for small levels of noise, the estimates become negatively biased at high sampling frequencies. This is consistent with the convergence results for the power variations, which show that noise causes the difference between RV_t and BV_t to become negatively biased. For the lowest levels of noise

[3]The results based on the Anderson-Darling and Shapiro-Francia tests are equivalent.

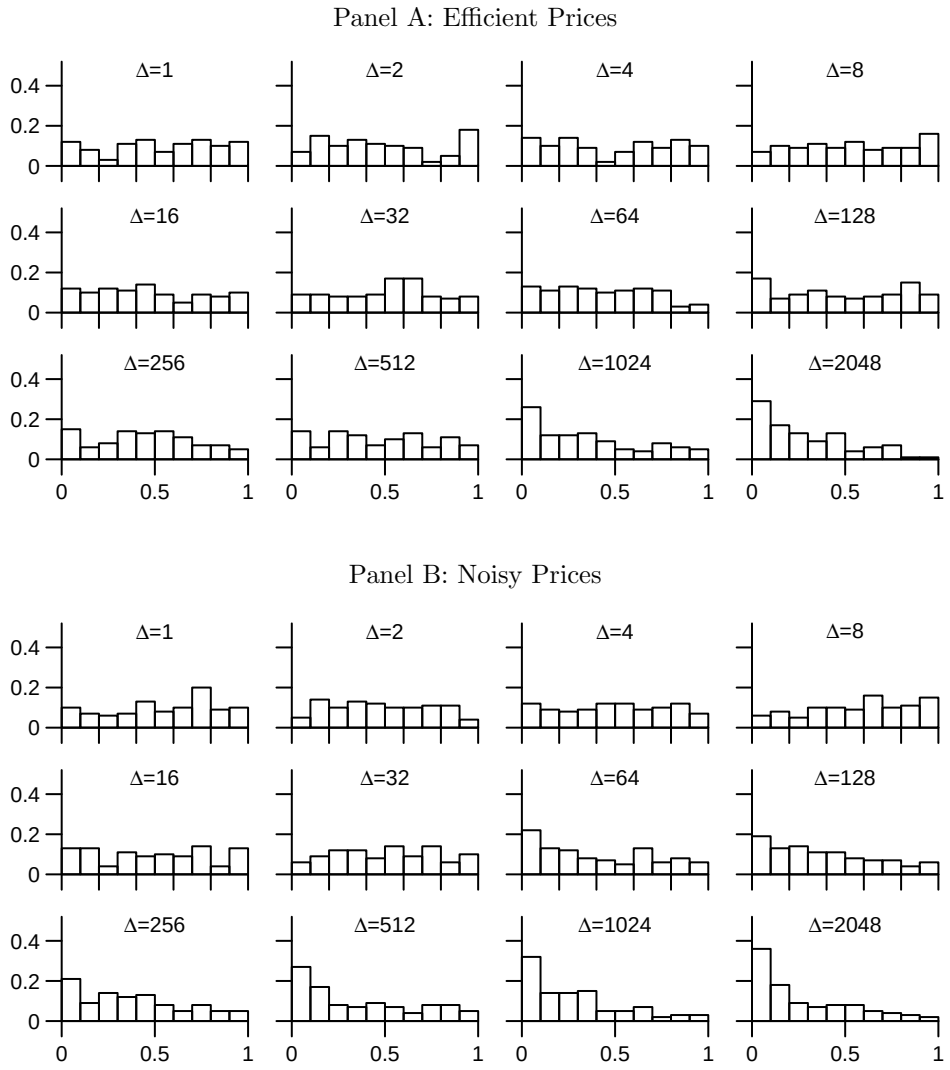

Figure 2.8: The figure graphs histograms of 100 p-values from the K-S test on 100 realizations of the Z_{TPRM} statistic. Prices are simulated from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. The plot labels denote the sampling interval, Δ, in seconds. Panel A plots results for efficient price series; Panel B presents results where an iid $\text{N}(0, 0.160)$ noise component is added to the price process.

($\sigma_{\text{mn}} = 0.027$), a sampling interval around three minutes or longer seems appropriate. The optimal frequency drops as the noise intensifies. Similarly, the standard deviation is biased. Panel B presents estimates for staggered intraday returns which are offset by one lag. Interestingly, the means are close to zero at any noise level and sampling frequency. Similarly, the standard deviations are close to one. Hence, offsetting the

intraday returns appears to adequately address the impact of noise.

The results without offsetting the intraday returns presented in Panel A in Table 2.7 suggest that for each noise variance, there exists a range of sampling intervals that produces estimates of the moments that are consistent with the asymptotic properties. The objective of the optimal sampling rate methods introduced in the previous section is to find these ranges. I compute the rates given by Bandi and Russell (2006) and find that these correspond well with the sampling intervals in Panel A that produces values close to the asymptotics. The estimated optimal sampling intervals for the first five noise processes are 31, 242, 569, 1227, and 2125 seconds, which all seem appropriate. The interval for the most severe noise component is 2827, and thus goes beyond the range covered in the table.

Overlapping Streams

Section 2.2.2 reviews a method by Zhang et al. (2005) where all data can be used to compute the daily realized variation without inflating the bias due to noise. I show in Section 2.4.1 that this approach reduces the mean-square error for the power variations. I evaluate whether the same method can be applied to the jump statistics. Following the same experimental design, I find that the jump statistics do not converge towards the normal distribution with mean zero and standard deviation one. A likely explanation is that while the staggered and optimal sampling methods reduce the bias, there is still some bias remaining which accumulates when estimates from overlapping streams are combined. Hence, I do not proceed with this method in the following.

2.4.3 Size

The following section evaluates the size of the statistics for different sampling rates, noise processes, and jump distributions.

Table 2.7: The table presents means and standard deviations (in parentheses) of the Z_{TPRM} statistic for 10000 price paths. The prices are generated from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. Panel A and Panel B report estimates without and with staggered returns by one lag.

	0.000	0.020	0.040	0.080	0.160	0.320
Panel A: No Staggering						
1	0.01	−21.04	−23.00	−23.53	−23.66	−23.69
	(1.00)	(1.61)	(0.84)	(0.73)	(0.72)	(0.72)
2	0.01	−13.37	−15.77	−16.50	−16.68	−16.73
	(1.01)	(1.77)	(0.91)	(0.73)	(0.72)	(0.71)
4	−0.01	−7.85	−10.54	−11.49	−11.76	−11.83
	(0.99)	(1.77)	(1.01)	(0.76)	(0.72)	(0.72)
8	0.01	−4.10	−6.71	−7.89	−8.25	−8.34
	(0.98)	(1.52)	(1.09)	(0.79)	(0.73)	(0.72)
16	−0.00	−1.81	−3.94	−5.27	−5.73	−5.87
	(1.00)	(1.24)	(1.12)	(0.83)	(0.73)	(0.72)
32	0.00	−0.65	−2.06	−3.36	−3.92	−4.11
	(1.00)	(1.04)	(1.06)	(0.87)	(0.75)	(0.72)
64	0.01	−0.19	−0.91	−1.97	−2.63	−2.87
	(0.99)	(1.00)	(1.00)	(0.90)	(0.78)	(0.74)
128	0.02	−0.03	−0.33	−1.01	−1.65	−1.95
	(0.98)	(0.98)	(0.98)	(0.93)	(0.83)	(0.76)
256	−0.01	0.00	−0.11	−0.45	−0.93	−1.27
	(0.98)	(0.97)	(0.98)	(0.95)	(0.86)	(0.79)
512	−0.02	−0.01	−0.02	−0.15	−0.47	−0.79
	(0.98)	(0.98)	(0.98)	(0.96)	(0.91)	(0.83)
1024	0.00	0.01	−0.01	−0.06	−0.19	−0.43
	(1.00)	(0.99)	(1.00)	(0.98)	(0.94)	(0.89)
2048	0.01	0.00	0.01	−0.01	−0.07	−0.21
	(0.99)	(1.00)	(0.99)	(0.99)	(0.98)	(0.95)
Panel B: Staggering						
1	−0.01	−0.02	−0.01	−0.01	−0.02	−0.02
	(1.00)	(1.00)	(0.99)	(0.99)	(0.99)	(0.99)
2	0.00	0.01	0.00	0.00	0.01	0.01
	(1.00)	(0.98)	(0.99)	(0.99)	(0.98)	(0.98)
4	−0.01	−0.00	0.00	0.00	0.00	0.00
	(1.00)	(1.00)	(0.99)	(0.98)	(0.99)	(0.99)
8	0.02	0.00	−0.01	−0.01	0.00	0.00
	(1.00)	(1.00)	(0.99)	(0.99)	(0.99)	(0.99)
16	−0.00	−0.00	0.01	0.01	0.00	0.00
	(1.01)	(0.99)	(1.00)	(0.99)	(0.98)	(0.98)
32	−0.00	0.02	−0.01	−0.02	0.02	0.01
	(1.00)	(0.98)	(0.99)	(1.00)	(0.98)	(0.98)
64	0.02	0.01	0.01	−0.00	−0.01	−0.01
	(0.99)	(1.00)	(1.00)	(0.99)	(0.99)	(0.99)
128	0.02	0.02	0.01	0.01	0.01	0.01
	(0.99)	(0.99)	(1.01)	(0.99)	(1.00)	(1.00)
256	−0.02	−0.00	−0.01	−0.01	0.01	0.01
	(0.99)	(0.99)	(0.99)	(0.99)	(0.99)	(0.98)
512	−0.00	−0.01	0.01	0.01	−0.00	0.01
	(1.01)	(1.01)	(1.01)	(1.01)	(1.00)	(1.00)
1024	0.01	0.01	0.00	0.00	0.01	0.01
	(1.04)	(1.04)	(1.04)	(1.03)	(1.04)	(1.03)
2048	0.01	0.01	0.00	0.00	0.01	0.02
	(1.07)	(1.07)	(1.07)	(1.07)	(1.08)	(1.08)

Figure 2.9 graphs time series of five statistics. The statistics are Z_{TP}, Z_{TPL}, Z_{TPLM},

Z_{TPR} and Z_{TPRM}, which are defined in Section 2.2 beginning on page 16. The continuous and dashed horizontal lines indicate the 0.99 and 0.999 critical values, respectively. Since there are no jumps in the process, 1 and 0.1 percent of the samples are expected to be greater than the critical values. Thus, on average the number of rejections should be 100 and 10 (out of a 10000 realizations). The actual rejection rates for the critical value 0.01 for Z_{TP}, Z_{TPL}, Z_{TPLM}, Z_{TPR} and Z_{TPRM} are 0.031, 0.024, 0.019, 0.015, and 0.011, respectively. All statistics except Z_{TPRM} reject the null too frequently. I present a more rigorous investigation of the rejection rates henceforth.

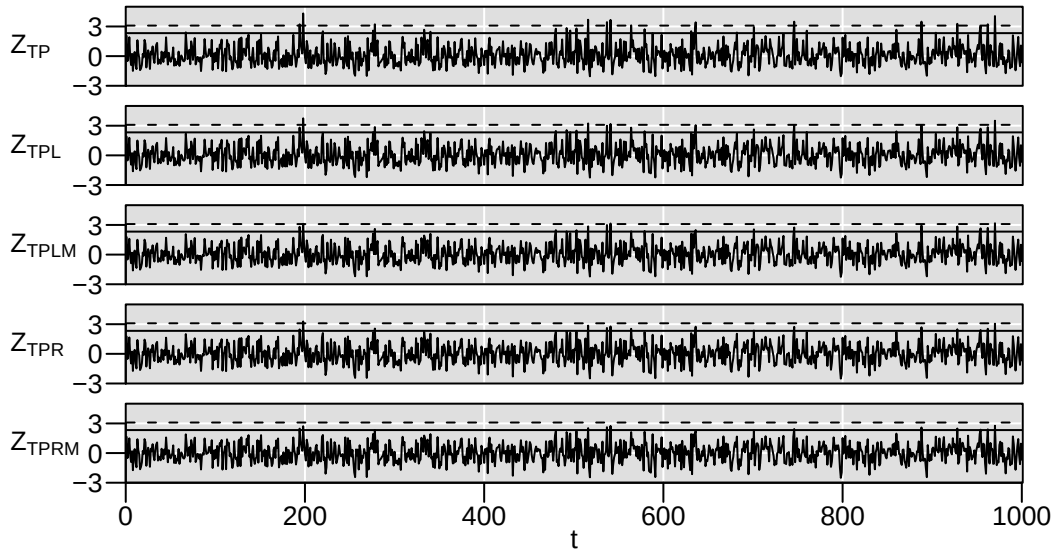

Figure 2.9: The figure plots realizations of five jump statistics based on 1000 simulated days from the SV1F model, equation (2.44). The experimental design is described in Table 2.1. The statistics are Z_{TP}, Z_{TPL}, Z_{TPLM}, Z_{TPR} and Z_{TPRM}. The statistics are computed based on five-minute intraday returns. The horizontal continuous and dashed lines are the 0.99 and 0.999 critical values, respectively.

Panel A in Figure 2.10 plots the size of five test statistics against the sampling intervals: one, three, five, ten, fifteen and thirty minutes. The nominal size is 0.05 for the upper-left panel, 0.01 for the upper-right, 0.005 for the lower-left, and 0.001 for the lower-right. The estimates are for 10000 simulated trading days from the SV1F

model without any noise process.[4]

The test size and sampling interval are positively related with the exception of Z_{TPRM}; as a result, the tests become increasingly anti-conservative as the sampling interval increases. The estimated size for the Z_{TP} statistic is in fact greater than the nominal value even for the highest sampling frequency. The rejection rate based on the Z_{TPRM} statistic, however, remains near the nominal size for all sizes and sampling intervals. The two statistics based on the ratio, RJ_t, deviate the least from the nominal size although the logarithmic adjusted statistic with a maximum correction, Z_{TPLM}, remains close to Z_{TPR}. Moreover, the maximum correction reduces the bias in both cases where it is applied.

Panel B plots the size for prices with noise. I add an iid $\text{N}(0, \sigma_{mn})$ process to the simulated prices with $\sigma_{mn} = 0.080$. The trends are similar to the behavior without noise in Panel A; however, the statistics are conservative at higher sampling frequencies before eventually turning invalid. The ranking among the statistics in terms of having a size near the nominal size remains the same as before where the estimated size for the Z_{TPRM} statistic is closest to the nominal size for sampling intervals at about ten minutes or longer; its size remains impressively close to the nominal size in spite of the relatively severe noise component. Furthermore, a sampling interval in the range from ten to thirty minutes appears appropriate with and without noise while a higher rate produces a very conservative statistic at high levels of noise. That is, the noise biases the statistics against identifying jumps if the price process is sampled too often. These findings agree with the applied literature on high-frequency data where the sampling interval typically is chosen in the five to thirty-minute range. From a

[4]Standard errors of the mean values are readily available since the estimates have a binomial distribution. The standard error is $(\widehat{p}(1-\widehat{p})/N)^{1/2}$, where $N = 10000$ is the number observations and $\widehat{p}$ is the estimated rejection rate. The standard error for $\widehat{p} = 0.1$, for example, is approximately 0.001.

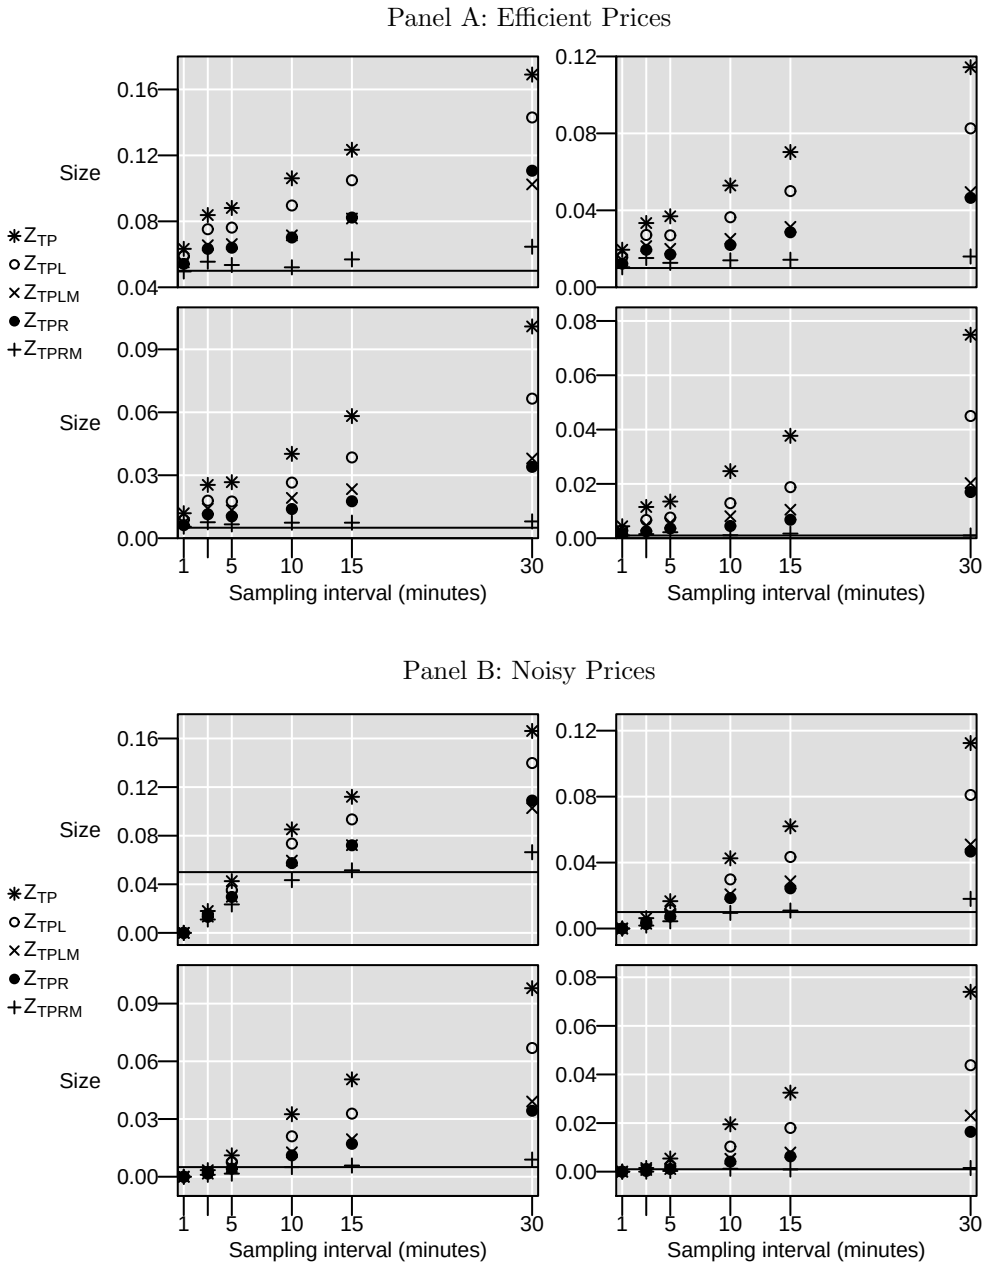

Figure 2.10: The figure plots the size of five jump statistics based on 10000 simulated days of the SV1F model, equation (2.44) (page 30). The experimental design is defined in Table 2.1 (page 31). The estimates are based on efficient prices in Panel A and noisy prices in Panel B where an iid $\mathrm{N}(0, \sigma_{mn})$ process is added to the simulated prices with $\sigma_{mn} = 0.080$. The sizes are plotted against sampling intervals which range from one to thirty minutes The statistics are Z_{TP}, Z_{TPL}, Z_{TPLM}, Z_{TPR} and Z_{TPRM}. The horizontal lines denote the nominal sizes, which are 0.05 (upper-left panel), 0.01 (upper-right panel), 0.005 (lower-left panel), and 0.001 (lower-right panel).

statistical point of view, however, it does not seem optimal to throw away such large amount of data unless absolutely needed. That is, if the noise level is sufficiently low so that a three-minute sampling interval does not produce biased estimates, then sampling at thirty minutes only uses a tenth of the data, thus potentially produces less efficient estimates. Henceforth, my objective is to evaluate whether the methods from Section 2.2.2 allow to use more data while preserving the asymptotic properties of the statistics.

Table 2.8 reports the rejection frequencies under the null hypothesis; the significance level is set to $\alpha = 0.99$. The columns report rejection rates for different values of the standard deviation of the noise process, σ_{mn}. The sampling rates are kept constant at one, three, five and thirty minutes. I include results for three statistics, Z_{TP}, Z_{TPLM} and Z_{TPRM}. The results based on the quadpower estimator, QP_t, are analogous to TP_t and therefore not reported. The three panels tabulate the rejection rates for the statistics where the bipower (equation (2.41)) and tripower (equation (2.42)) are computed using staggered returns with offset zero (panel $i = 0$), one (panel $i = 1$) and two (panel $i = 2$).

Several interesting results appear.

1. The first panel clearly shows that the noise has a considerable impact on the test sizes, particularly at high sampling frequencies. For one-minute sampling intervals, the statistics become biased against identifying jumps, which is consistent with the convergence results in Section 2.4.2. In fact, the rejection rates are less than 0.0005 for all three statistics for the three largest values of the noise variations although the nominal test size is 0.01. As the sampling interval increases to three minutes, the test size approaches the nominal size yet remains conservative for the larger values of σ_{mn}. Notice, however, that the statistics are becoming increasingly anti-conservative for no or minor noise at this sampling

rate. The same patterns hold for five-minute sampling. For the thirty-minute interval, however, the Z_{TP} and Z_{TPLM} statistics reject the null hypothesis too often while Z_{TPRM} is reasonably close to the nominal size.

Thus, I confirm that the optimal constant sampling rate is highly dependent on the noise variance. A high sampling rate yields test sizes that are closer to the true size without noise while the appropriate sampling frequency drops as the noise variance increases.

2. A cross-panel comparison indicates that applying staggered returns reduces the impact of noise considerably. Importantly, the estimated sizes are nearly constant across all values of the noise variations (see second panel: $i = 1$), and thus alleviate the user from having to gauge the level of noise in order to select an appropriate sampling rate.

   Offsetting the returns by two lags (panel: $i = 2$) does not enhance the results but rather worsen the tests by making them further anti-conservative. I find that the mean of the statistic becomes positively biased and thus causes the test to be invalid.

   A noteworthy finding is that the rejection rates for Z_{TPRM} at the highest sampling frequency when offsetting the returns by one lag is analogous to the thirty-minute sampling interval without staggering. That is, the former uses thirty times more data. I investigate below whether this translates into a more powerful test.

3. The Z_{TPRM} statistic consistently produces test sizes that are closest to the nominal size narrowly followed by Z_{TPLM}. The Z_{TP} is considerably worse.

Table 2.9 presents results based on the method by Bandi and Russell (2006). I compute sampling rates per day using their exact and approximate equations, see

Table 2.8: The sizes of three statistics are tabulated based on 10000 days simulated from the SV1F model, equation (2.44). An iid $\mathrm{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.040, 0.052, 0.065, 0.080. The panel labels $i = 0, 1, 2$ denote the staggered offset. The statistics are Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The return horizons are one, three, five and thirty minutes. The test size is 0.01.

		$\{\sigma_{mn}\}$					
Interval		0.000	0.027	0.040	0.052	0.065	0.080
				$(i = 0)$			
1 minutes	Z_{TP}	0.019	0.006	0.002	0.000	0.000	0.000
	Z_{TPLM}	0.014	0.004	0.001	0.000	0.000	0.000
	Z_{TPRM}	0.010	0.003	0.001	0.000	0.000	0.000
3 minutes	Z_{TP}	0.033	0.029	0.022	0.014	0.009	0.006
	Z_{TPLM}	0.021	0.018	0.013	0.009	0.005	0.003
	Z_{TPRM}	0.015	0.011	0.008	0.006	0.004	0.002
5 minutes	Z_{TP}	0.037	0.038	0.034	0.027	0.023	0.017
	Z_{TPLM}	0.020	0.019	0.019	0.016	0.013	0.009
	Z_{TPRM}	0.013	0.013	0.012	0.010	0.007	0.004
30 minutes	Z_{TP}	0.114	0.115	0.116	0.115	0.115	0.113
	Z_{TPLM}	0.049	0.050	0.050	0.049	0.052	0.051
	Z_{TPRM}	0.016	0.015	0.016	0.017	0.017	0.018
				$(i = 1)$			
1 minutes	Z_{TP}	0.025	0.025	0.024	0.024	0.024	0.023
	Z_{TPLM}	0.018	0.017	0.017	0.017	0.017	0.015
	Z_{TPRM}	0.014	0.014	0.013	0.013	0.013	0.011
3 minutes	Z_{TP}	0.042	0.041	0.040	0.038	0.036	0.034
	Z_{TPLM}	0.026	0.023	0.022	0.023	0.022	0.022
	Z_{TPRM}	0.017	0.015	0.015	0.015	0.016	0.015
5 minutes	Z_{TP}	0.048	0.049	0.048	0.050	0.048	0.049
	Z_{TPLM}	0.026	0.027	0.028	0.028	0.028	0.028
	Z_{TPRM}	0.016	0.016	0.017	0.017	0.018	0.017
30 minutes	Z_{TP}	0.192	0.192	0.192	0.192	0.190	0.190
	Z_{TPLM}	0.090	0.092	0.090	0.089	0.090	0.090
	Z_{TPRM}	0.033	0.033	0.034	0.033	0.034	0.034
				$(i = 2)$			
1 minutes	Z_{TP}	0.025	0.027	0.028	0.028	0.029	0.031
	Z_{TPLM}	0.018	0.020	0.020	0.021	0.022	0.022
	Z_{TPRM}	0.014	0.017	0.017	0.016	0.016	0.017
3 minutes	Z_{TP}	0.047	0.047	0.047	0.046	0.046	0.046
	Z_{TPLM}	0.027	0.029	0.030	0.029	0.028	0.028
	Z_{TPRM}	0.019	0.020	0.021	0.021	0.019	0.020
5 minutes	Z_{TP}	0.070	0.065	0.067	0.068	0.070	0.070
	Z_{TPLM}	0.040	0.040	0.039	0.040	0.041	0.043
	Z_{TPRM}	0.026	0.025	0.026	0.027	0.026	0.026
30 minutes	Z_{TP}	0.308	0.318	0.320	0.319	0.320	0.317
	Z_{TPLM}	0.161	0.159	0.161	0.163	0.165	0.164
	Z_{TPRM}	0.067	0.066	0.065	0.066	0.065	0.067

equations (2.33) and (2.34), which I refer to as BR1 and BR0, respectively. Notice that the optimal sampling rates are computed per day; that is, the sampling rate is adjusted per day. The benefit is that if the price process is noisier during certain periods, the sampling rate is appropriately adjusted.

In contrast to the results for constant sampling without staggered returns, the sizes

stay effectively constant near the nominal size across all standard deviations for the Z_{TPRM} statistic. That is, the noise does not bias the Z_{TPRM} statistic against rejecting the null hypothesis, which is remarkable considering the large bias resulting from sampling at constant sampling rates, see Table 2.8. Rather than being conservative, which is the case with constant sampling, the Z_{TP} and Z_{TPLM} statistics reject the null hypothesis too often, in particular as the noise increases. The size for the Z_{TPRM} statistic, however, is very close to the nominal size at all instances. The application of staggered returns combined with BR makes the test statistics anti-conservative and thus invalidates the test. Further analysis shows that the mean values of the statistics becomes positively biased, which results in too many rejection, when both methods are applied.

Bandi and Russell (2006) evaluate two methods for estimating the optimal sampling rate, one exact and one approximate. The authors find that the exact (BR1) and approximate (BR0) rates are close when it is optimal to sample the price process at a high frequency. As the noise component increases, however, the approximation tends to underestimate the optimal interval. I do not find any such disparity between the two equations in this context. This finding has an important practical implication since the approximating equation has a closed-form solution while the exact equation requires an optimization routine and consequently is computationally slower to obtain.

I document the sampling intervals estimated by BR1 to further explore what causes the difference between applying constant and optimal sampling rates with no staggering. For prices without noise, the optimal sampling interval predicted by BR1 is around 30 seconds. The interval gradually increases with the noise and reaches about 30 minutes for the largest noise variance. Interestingly, even though BR1 on average gives a 30-minute sampling rate for the largest noise variance, holding the

sampling interval constant at that rate across the whole sample period yields worse results (compare with the first panel in Table 2.8 for results with constant 30-minute sampling.) This supports the notion that estimating the sampling rate per trading day rather than across the full sample is beneficial since some intervals are more (or less) noisy and thus require longer (shorter) sampling intervals.

Table 2.9: The size is tabulated for three statistics based on 10000 days simulated from the SV1F model, equation (2.44). An iid N$(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.040, 0.052, 0.065, 0.080. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The panel labels $i = 0, 1, 2$, denote the staggered offset. The statistics are Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The test size is 0.01.

		$\{\sigma_{mn}\}$					
Interval		0.000	0.027	0.040	0.052	0.065	0.080
				$(i = 0)$			
BR0	Z_{TP}	0.016	0.040	0.048	0.062	0.068	0.074
	Z_{TPLM}	0.013	0.021	0.024	0.030	0.032	0.034
	Z_{TPRM}	0.011	0.013	0.013	0.014	0.013	0.014
BR1	Z_{TP}	0.016	0.040	0.050	0.057	0.066	0.079
	Z_{TPLM}	0.013	0.023	0.024	0.027	0.033	0.035
	Z_{TPRM}	0.011	0.013	0.013	0.012	0.013	0.012
				$(i = 1)$			
BR0	Z_{TP}	0.019	0.054	0.071	0.093	0.108	0.133
	Z_{TPLM}	0.014	0.031	0.038	0.045	0.051	0.062
	Z_{TPRM}	0.011	0.017	0.021	0.021	0.025	0.024
BR1	Z_{TP}	0.019	0.052	0.075	0.092	0.117	0.135
	Z_{TPLM}	0.014	0.029	0.039	0.046	0.057	0.066
	Z_{TPRM}	0.011	0.018	0.021	0.023	0.025	0.026
				$(i = 2)$			
BR0	Z_{TP}	0.021	0.071	0.104	0.130	0.167	0.198
	Z_{TPLM}	0.016	0.041	0.057	0.068	0.085	0.107
	Z_{TPRM}	0.014	0.025	0.031	0.034	0.042	0.047
BR1	Z_{TP}	0.021	0.070	0.103	0.139	0.169	0.197
	Z_{TPLM}	0.016	0.041	0.054	0.072	0.087	0.103
	Z_{TPRM}	0.014	0.024	0.030	0.034	0.045	0.046

Table 2.10 documents the test sizes based on the optimal sampling rate by Zhang et al. (2005). The size for the Z_{TPRM} statistic remains close to the nominal test size and is generally unaffected as the noise increases without staggering. Applying staggered returns does not improve the performance but rather worsen the test as the noise increases. Hence, the patterns are analogous to those of BR.

Table 2.10: The size is tabulated for three statistics based on 10000 days simulated from the SV1F model, equation (2.44). An iid $N(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.040, 0.052, 0.065, 0.080. The sampling rates are determined per day by the method by Zhang et al. (2005). The panel labels $i = 0, 1, 2$, denote the staggered offset. The statistics are Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The test size is 0.01.

		$\{\sigma_{mn}\}$					
Interval		0.000	0.027	0.040	0.052	0.065	0.080
				$(i = 0)$			
ZMA	Z_{TP}	0.016	0.043	0.054	0.064	0.071	0.080
	Z_{TPLM}	0.013	0.022	0.028	0.029	0.034	0.039
	Z_{TPRM}	0.011	0.012	0.014	0.014	0.014	0.015
				$(i = 1)$			
ZMA	Z_{TP}	0.019	0.061	0.082	0.097	0.115	0.129
	Z_{TPLM}	0.015	0.034	0.041	0.053	0.057	0.066
	Z_{TPRM}	0.013	0.019	0.022	0.027	0.025	0.027
				$(i = 2)$			
ZMA	Z_{TP}	0.022	0.081	0.112	0.148	0.172	0.201
	Z_{TPLM}	0.018	0.048	0.060	0.077	0.094	0.106
	Z_{TPRM}	0.015	0.030	0.033	0.040	0.045	0.047

Robust Optimal Sampling Rate

In Section 2.2.2, I propose a modification of the BR method by estimating the optimal sampling rate based on the QP_t estimator, which is robust to jumps. I refer to the robust version as RobustBR. In the following, I evaluate whether the validity of the test is impacted compared to the original BR version.

I constrain the results to the Z_{TPRM} statistic. Table 2.11 presents the estimated sizes for different standard deviations of the noise process. The table includes results for the approximate and exact solutions for BR and RobustBR, respectively. Interestingly, RobustBR is valid in most cases except for small noise variances while BR is slightly anti-conservative for all models. The difference between BR0 and RobustBR0 is significant for $\sigma_{\text{mn}} \geq 0.040$.[5] I find that even though the two estimators are asymptotically equal under the null hypothesis, the RobustBR produces a sampling frequency that on average is higher compared to BR irrespective of the noise component, which may explain the difference in their sizes. Applying staggered returns to RobustBR causes the test to be invalid, which is consistent with the findings

[5]P-values based on the large sample hypothesis test `prop.test` in `R` are less than 0.001.

for BR.

Table 2.11: The size is tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1F model, equation (2.44). An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.040, 0.052, 0.065, 0.080. BR1 and BR0 denote sampling rates that are obtained by solving equations (2.33) and (2.34), respectively. The results denoted by RobustBR are based on the sampling rates computed as described in Section 2.2.2. The test size is 0.01.

	σ_{mn}					
	0.000	0.027	0.040	0.052	0.065	0.080
BR0	0.011	0.013	0.013	0.014	0.013	0.014
Robust BR0	0.012	0.012	0.009	0.009	0.010	0.010
BR1	0.011	0.013	0.013	0.012	0.013	0.012
Robust BR1	0.012	0.011	0.011	0.012	0.009	0.010

Serially Correlated Noise and Rounding Errors

The previous studies of the jump statistics and daily variations that I extend evaluate the finite sample properties under a normal iid noise process. Thus, an important contribution is to evaluate the impact of other models of market frictions. In Section 2.3.2 on page 33, I introduce processes by Aït-Sahalia et al. (2006) with autocorrelated errors and by Li and Mykland (2007) with rounding errors. Henceforth, I refer to these by AMZ and LMTS.

I simulate from AMZ for different values of $E[U^2]$, $E[V^2]$, and ϕ. Aït-Sahalia et al. (2006) estimate the parameters of their model for a few stocks and find that ϕ generally is negative. I report results for ϕ equal to -0.3 and -0.7. Specifically, I consider the following sets of parameter values: $(\sigma_U, \sigma_V, \phi) = \{(0.020, 0.020, -0.3), (0.040, 0.040, -0.3), (0.020, 0.020, -0.7), (0.040, 0.040, -0.7)\}$. These values give a total variation in the same range as those I consider above but with an added correlation structure introduced by ϕ.

Li and Mykland (2007) assume that the noise term, ϵ_{t_i}, is iid normal with zero mean and variance σ_ϵ^2. I report simulation results for $\sigma_\epsilon = \{0.027, 0.052, 0.080\}$ and $\gamma = \{0.01, 0.10\}$ where γ is applied to introduce rounding errors in the price process,

see equation (2.47) on page 34.

I find that the results based on the optimal sampling frequency proposed by Zhang et al. (2005) are similar to those by Bandi and Russell (2006); therefore, I excluded the former. For the same reason, I only report estimates for BR0 and exclude BR1. All estimates are based on 10000 simulations.

Table 2.12 tabulates the size for the Z_{TP}, Z_{TPLM} and Z_{TPRM} statistics for contaminated prices using the AMZ model. The test size does not appear to be sensitive to ϕ as the values are equivalent for the two values of ϕ as well as for other values not reported here. The variance, however, biases the statistic against identifying jumps, which is consistent with the previous results for iid errors. The tests becomes conservative without staggering at longer sampling intervals. The size for staggered returns at one lag offset is near the nominal level at the highest frequency for the Z_{TPRM} statistic. Combining staggering and sampling at longer intervals produces invalid tests, which is consistent with the results for iid normal noise. Similarly for the optimal rates by Bandi and Russell (2006), the autocorrelation does not seem to have a significant impact on the size.

The results based on prices that are contaminated following the model by Li and Mykland (2007) are consistent with the results for iid normal errors. Hence, rounding prices does not have a significant impact. Tables 2.14 and 2.15 show the test sizes for constant and optimal sampling rates and reveal modest differences from the iid case. Additional simulation runs with larger rounding errors have a similar effect to increasing the variance of the iid normal process.

2.4.4 Power

In this section, I add the jump component to the data-generating process and evaluate the power of the test statistics for different values of the jump intensity, λ, and the

Table 2.12: The size is tabulated for the Z_{TP}, Z_{TPLM}, and Z_{TPRM} statistics based on 10000 days simulated from the SV1F model, equation (2.44). The observed prices are generated following AMZ where $\{\sigma_U, \sigma_Z, \phi\} \in \{(\sigma, \sigma, \phi)\}$ for σ equal to 0.020 and 0.040 and ϕ equal to -0.3 and -0.7. The panel label i denotes the staggered offset. The return horizons are one, three, five and thirty minutes. The test size is 0.01.

		$\{\sigma_U, \sigma_Z, \phi\}$			
Interval		$\{20, 20, -0.3\}$	$\{40, 40, -0.3\}$	$\{20, 20, -0.7\}$	$\{40, 40, -0.7\}$
			$(i = 0)$		
1 minutes	Z_{TP}	0.005	0.001	0.006	0.001
	Z_{TPLM}	0.004	0.000	0.004	0.001
	Z_{TPRM}	0.003	0.000	0.003	0.000
3 minutes	Z_{TP}	0.026	0.011	0.026	0.009
	Z_{TPLM}	0.016	0.006	0.014	0.006
	Z_{TPRM}	0.011	0.004	0.009	0.004
5 minutes	Z_{TP}	0.036	0.023	0.036	0.023
	Z_{TPLM}	0.019	0.012	0.019	0.013
	Z_{TPRM}	0.012	0.007	0.012	0.007
30 minutes	Z_{TP}	0.111	0.109	0.111	0.107
	Z_{TPLM}	0.050	0.047	0.050	0.048
	Z_{TPRM}	0.014	0.014	0.014	0.015
			$(i = 1)$		
1 minutes	Z_{TP}	0.026	0.023	0.027	0.023
	Z_{TPLM}	0.019	0.017	0.019	0.017
	Z_{TPRM}	0.015	0.013	0.016	0.015
3 minutes	Z_{TP}	0.039	0.037	0.039	0.037
	Z_{TPLM}	0.022	0.022	0.023	0.023
	Z_{TPRM}	0.015	0.015	0.016	0.015
5 minutes	Z_{TP}	0.048	0.048	0.048	0.047
	Z_{TPLM}	0.027	0.028	0.028	0.027
	Z_{TPRM}	0.016	0.017	0.016	0.015
30 minutes	Z_{TP}	0.189	0.192	0.188	0.189
	Z_{TPLM}	0.087	0.089	0.088	0.088
	Z_{TPRM}	0.032	0.033	0.033	0.033

Table 2.13: The size is tabulated for the Z_{TP}, Z_{TPLM}, and Z_{TPRM} statistics based on 10000 days simulated from the SV1F, equation (2.44). The observed prices are generated following AMZ where $\{\sigma_U, \sigma_Z, \phi\} \in \{(\sigma, \sigma, \phi)\}$ for $\sigma = 0.020, 0.040$, and $\phi = -0.3, -0.7$. The sampling rates are obtained by solving equation (2.33). The panel labels i denotes the staggered offset. The test size is 0.01.

		$\{\sigma_U, \sigma_Z, \phi\}$			
Interval		$\{20, 20, -0.3\}$	$\{40, 40, -0.3\}$	$\{20, 20, -0.7\}$	$\{40, 40, -0.7\}$
			$(i = 0)$		
BR0	Z_{TP}	0.041	0.069	0.047	0.066
	Z_{TPLM}	0.023	0.033	0.025	0.031
	Z_{TPRM}	0.013	0.013	0.014	0.015
BR1	Z_{TP}	0.043	0.063	0.045	0.067
	Z_{TPLM}	0.023	0.029	0.025	0.030
	Z_{TPRM}	0.014	0.013	0.015	0.011
			$(i = 1)$		
BR0	Z_{TP}	0.055	0.105	0.065	0.111
	Z_{TPLM}	0.030	0.053	0.036	0.055
	Z_{TPRM}	0.017	0.025	0.021	0.026
BR1	Z_{TP}	0.060	0.109	0.066	0.115
	Z_{TPLM}	0.032	0.054	0.037	0.056
	Z_{TPRM}	0.019	0.023	0.022	0.026

Table 2.14: The size is tabulated for three statistics based on 10000 days simulated from the SV1FJ model, equation (2.44) with the jump rate, $\lambda = 0.014$, and standard deviation of the jump size is, $\sigma_{\text{jmp}} = 1.50$. The observed prices are generated following LMTS for $\sigma = 0.027, 0.052, 0.080$ and $\gamma = 0.01, 0.10$. Results are presented for the return horizons: one, three, five and thirty minutes. The statistics are Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The panel label i denotes the staggered offset. The test size is $\alpha = 0.01$.

		$\{\sigma_\epsilon, \gamma\}$					
Interval		$\{27, 0.01\}$	$\{52, 0.01\}$	$\{80, 0.01\}$	$\{27, 0.1\}$	$\{52, 0.1\}$	$\{80, 0.1\}$
				$(i = 0)$			
1 minutes	Z_{TP}	0.007	0.001	0.001	0.007	0.001	0.001
	Z_{TPLM}	0.005	0.001	0.001	0.005	0.001	0.001
	Z_{TPRM}	0.004	0.001	0.000	0.004	0.001	0.000
3 minutes	Z_{TP}	0.030	0.015	0.007	0.030	0.015	0.007
	Z_{TPLM}	0.018	0.010	0.003	0.018	0.010	0.003
	Z_{TPRM}	0.012	0.007	0.002	0.012	0.007	0.003
5 minutes	Z_{TP}	0.039	0.028	0.017	0.039	0.028	0.017
	Z_{TPLM}	0.020	0.016	0.009	0.020	0.016	0.009
	Z_{TPRM}	0.013	0.011	0.005	0.013	0.011	0.005
30 minutes	Z_{TP}	0.115	0.115	0.113	0.115	0.115	0.113
	Z_{TPLM}	0.050	0.049	0.051	0.050	0.050	0.051
	Z_{TPRM}	0.016	0.017	0.018	0.016	0.017	0.018
				$(i = 1)$			
1 minutes	Z_{TP}	0.027	0.025	0.024	0.027	0.026	0.024
	Z_{TPLM}	0.019	0.018	0.016	0.019	0.018	0.017
	Z_{TPRM}	0.015	0.014	0.013	0.015	0.014	0.013
3 minutes	Z_{TP}	0.042	0.039	0.035	0.042	0.039	0.035
	Z_{TPLM}	0.024	0.024	0.023	0.024	0.024	0.023
	Z_{TPRM}	0.016	0.016	0.016	0.016	0.016	0.016
5 minutes	Z_{TP}	0.050	0.051	0.050	0.050	0.051	0.050
	Z_{TPLM}	0.028	0.029	0.028	0.029	0.029	0.029
	Z_{TPRM}	0.016	0.018	0.017	0.017	0.018	0.017
30 minutes	Z_{TP}	0.192	0.192	0.190	0.192	0.192	0.191
	Z_{TPLM}	0.092	0.089	0.090	0.093	0.089	0.090
	Z_{TPRM}	0.033	0.033	0.034	0.034	0.033	0.034

standard deviation of the jump size, σ_{jmp}.

I generate prices from the jump-diffusion model, SVIFJ, as specified in equation (2.44) (page 30). The experimental design is described in Table 2.1 (page 31) with $\lambda = 0.014$ and $\sigma_{\text{jmp}} = 1.50$. I initially consider a price process without noise. The first five panels in Figure 2.11 graph time series of statistics and the bottom panel plots the jump process, $\kappa_t \mathrm{d}q$. The statistics appear to catch most of the larger jumps. There are, however, a number of smaller jumps that are not identified by the statistics at these significance levels, which is not surprising since these are difficult to distinguished from the diffusion process.

Table 2.16 presents confusion matrices for the Z_{TPRM} statistic for different values

Table 2.15: The size is tabulated for three statistics based on 10000 days simulated from the SV1FJ model, equation (2.44) with the jump rate, $\lambda = 0.014$, and standard deviation of the jump size is, $\sigma_{\text{jmp}} = 1.50$. The observed prices are generated following LMTS for $\sigma = 0.027, 0.052, 0.080$ and $\gamma = 0.01, 0.10$. The sampling rates are obtained by solving equation (2.33). The statistics are Z_{TP}, Z_{TPLM}, and Z_{TPRM}. The panel label i denotes the staggered offset. The test size is $\alpha = 0.01$.

		$\{\sigma_\epsilon, \gamma\}$					
Interval		$\{27, 0.01\}$	$\{52, 0.01\}$	$\{80, 0.01\}$	$\{27, 0.1\}$	$\{52, 0.1\}$	$\{80, 0.1\}$
				$(i = 0)$			
BR0	Z_{TP}	0.040	0.063	0.074	0.040	0.063	0.074
	Z_{TPLM}	0.022	0.031	0.034	0.022	0.031	0.034
	Z_{TPRM}	0.013	0.014	0.014	0.014	0.014	0.014
BR1	Z_{TP}	0.041	0.058	0.079	0.041	0.058	0.079
	Z_{TPLM}	0.023	0.028	0.035	0.023	0.028	0.036
	Z_{TPRM}	0.013	0.013	0.013	0.013	0.013	0.013
				$(i = 1)$			
BR0	Z_{TP}	0.055	0.094	0.133	0.055	0.094	0.134
	Z_{TPLM}	0.031	0.045	0.062	0.032	0.045	0.063
	Z_{TPRM}	0.018	0.022	0.024	0.018	0.022	0.024
BR1	Z_{TP}	0.053	0.093	0.136	0.053	0.093	0.135
	Z_{TPLM}	0.030	0.047	0.066	0.030	0.047	0.066
	Z_{TPRM}	0.018	0.023	0.027	0.018	0.024	0.027

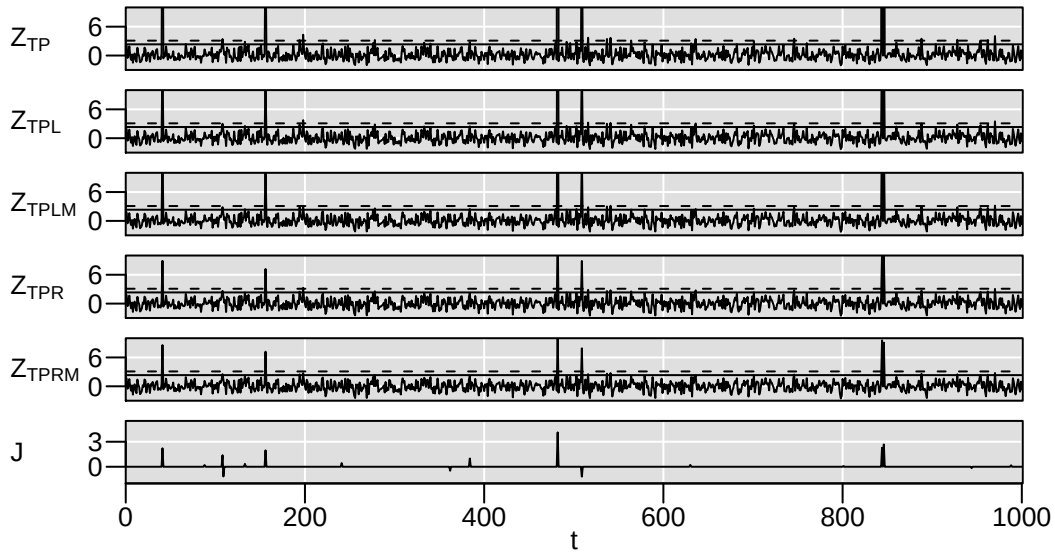

Figure 2.11: The five first panels plot realizations of five jump statistics for 1000 simulated days from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1 with $\lambda = 0.014$ and $\sigma_{\text{jmp}} = 1.50$. The bottom panel plots the jump process. The statistics are Z_{TP}, Z_{TPL}, Z_{TPLM}, Z_{TPR} and Z_{TPRM}. The statistics are computed based on five-minute intraday returns. The horizontal continuous and dashed lines are the 0.99 and 0.999 critical values, respectively.

of the jump intensity, λ. I exclude the Z_{TP} and Z_{TPLM} statistics since they are invalid, see Section 2.4.3. The labels, NJ and J, denote days without and with a jump, respectively. The rows represent the true events while the columns denote the statistical

inference. Hence, the rows for the 2×2 matrices add up to one where the 1×1 element is the fraction of correct non-rejections of the null (no-jump) hypothesis and the 1×2 element is the false rejection rate. Meanwhile, the 2×1 element is the false non-rejection of the null hypothesis and the 2×2 element is the correct rejection. The jump intensity, λ, is set to 0.014, 0.118, 1.000 and 2.000, respectively, while the standard deviation of the jump size, σ_{jmp}, is held constant at 1.50. The significance level, α, is 0.99.

Since the underlying prices are efficient, theory states that the price series should be sampled as frequently as possible. Consistently, the type I error is smallest and near the nominal test size for the highest frequency, that is, for the one-minute sampling interval. Furthermore, the test correctly rejects the null hypothesis at relatively high rates. As the sampling interval increases, the statistic is computed on fewer data points. Consequently, the test properties deteriorate as the variance increases. The type I error holds up reasonable well for the Z_{TPRM} statistic as the sampling rate decreases; the type II error, however, increases significantly. Hence, for efficient prices there is a considerable loss in power at low sampling rates. In fact, there is an evident drop in the power already at the five-minute sampling rate compared to the highest frequency.

Moreover, the observed patterns are nearly constant across the different jump intensities, which is anticipated since the nonparametric statistic is applied to each day individually. If the jump arrival rate were large enough to generate multiple jumps per day, the power should increase as the statistics would accrue the effects of several jumps. I expect the variance of the jump size, σ_{jmp}, however, to be positively related to the power of the test since larger jumps are easier to identify. To further explore the relationship between the power and the magnitude of the jump size, I simulate price processes for different values of the jump parameters. The rejection rates are

Table 2.16: Confusion matrices are tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1 $\sigma_{\text{jmp}} = 1.50$. The jump rates, λ, are $0.014, 0.118, 1.000$, and 2.000. Results are presented for four return horizons: one, three, five and thirty minutes. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\lambda = 0.014$		$\lambda = 0.118$		$\lambda = 1.000$		$\lambda = 2.000$	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
1 minutes	(NJ)	0.990	0.010	0.989	0.011	0.990	0.010	0.984	0.016
	(J)	0.239	0.761	0.208	0.792	0.211	0.789	0.211	0.789
3 minutes	(NJ)	0.985	0.015	0.985	0.015	0.986	0.014	0.990	0.010
	(J)	0.319	0.681	0.323	0.677	0.307	0.693	0.305	0.695
5 minutes	(NJ)	0.987	0.013	0.987	0.013	0.988	0.012	0.989	0.011
	(J)	0.377	0.623	0.404	0.596	0.375	0.625	0.373	0.627
30 minutes	(NJ)	0.984	0.016	0.984	0.016	0.980	0.020	0.985	0.015
	(J)	0.754	0.246	0.761	0.239	0.734	0.266	0.734	0.266

presented in Table 2.17. I hold the return horizon fixed at five minutes. There is a positive relationship between the power and variance for all three statistics.

Table 2.17: The power is tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1. The jump rate, λ, is 0.5, 1.0, 1.5 and 2.0. The standard deviation of the jumps, σ_{jmp}, ranges from 0.5 to 2.5 by increments of 0.5. The return horizon is five minutes. The test size is 0.01.

		σ_{jmp}				
	λ	0.5	1.0	1.5	2.0	2.5
Z_{TPRM}	0.5	0.197	0.465	0.615	0.700	0.749
	1.0	0.211	0.477	0.625	0.712	0.765
	1.5	0.206	0.484	0.634	0.715	0.768
	2.0	0.205	0.485	0.627	0.714	0.769

In order to determine the extent to which the power characteristics depend on the sampling frequency, Table 2.18 presents the equivalent results for the Z_{TPRM} statistic at different sampling intervals. It is remarkable how low the power drops for the 30-minute sampling intervals. Even going from one to five-minute sampling leads to a considerable reduction in power, which is significant since five-minute sampling intervals are commonplace in the applied empirical literature. Importantly, however, these results are for efficient prices.

To examine the impact of noise on the power of the test, I tabulate confusion matrices for different sampling intervals and noise variances based on 10000 simulated

Table 2.18: The power is tabulated for Z_{TPRM} per jump intensity and standard deviation of the jump size based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1. The jump rates, λ, are 0.5, 1.0, 1.5 and 2.0. The standard deviation of the jumps, σ_{jmp}, ranges from 0.5 to 2.5 by increments of 0.5. The return horizons are one, three, five and thirty minutes. The test size is 0.01.

		σ_{jmp}				
	λ	0.5	1.0	1.5	2.0	2.5
1 minutes	0.5	0.432	0.678	0.780	0.837	0.868
	1.0	0.439	0.691	0.789	0.842	0.871
	1.5	0.443	0.690	0.789	0.839	0.872
	2.0	0.438	0.692	0.789	0.841	0.871
3 minutes	0.5	0.278	0.552	0.679	0.756	0.806
	1.0	0.288	0.559	0.693	0.765	0.812
	1.5	0.289	0.569	0.698	0.768	0.812
	2.0	0.290	0.562	0.695	0.768	0.811
5 minutes	0.5	0.197	0.465	0.615	0.700	0.749
	1.0	0.211	0.477	0.625	0.712	0.765
	1.5	0.206	0.484	0.634	0.715	0.768
	2.0	0.205	0.485	0.627	0.714	0.769
30 minutes	0.5	0.038	0.138	0.252	0.349	0.426
	1.0	0.037	0.139	0.266	0.368	0.448
	1.5	0.041	0.147	0.273	0.377	0.462
	2.0	0.041	0.149	0.266	0.368	0.455

days from the SV1FJ model, equation (2.44). The jump intensity, λ, is 0.014 and the standard deviation of the jump size is 1.50.

Table 2.19 presents matrices for constant one, three, five and thirty-minute sampling intervals. For σ_{mn} equal to 0.052 and 0.080, the type I errors are less than 0.0005 at the highest sampling frequency. For the 30-minute sampling interval, the type I errors are near 0.01 for all values of σ_{mn}. The power decreases with the sampling frequency. Staggering the returns, however, increases the power. The type I errors remain nearly constant only narrowly exceeding 0.01. Without noise, the test rejects the false null about 75 percent of the time while the percentage drops to 50 percent for the largest noise variance for one-minute sampling. Offsetting the returns by two lags increases power slightly but causes the test to become invalid.

Table 2.20 presents the confusion matrices for the method by Bandi and Russell (2006). Analogous with the previous findings, the best results are obtained without staggering. Staggering the returns with one or two lags offset does not alter the power

Table 2.19: Confusion matrices are tabulated for Z_{TPRM} based on 10000 days simulated from the SV1FJ model, equation (2.44), with $\lambda = 0.014$, and $\sigma_{\text{jmp}} = 1.50$. An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.052, 0.080. Results are presented for four return horizons: one, three, five and thirty minutes. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
		$(i = 0)$							
1 minutes	(NJ)	0.990	0.010	0.997	0.003	1.000	0.000	1.000	0.000
	(J)	0.239	0.761	0.297	0.703	0.493	0.507	0.623	0.377
3 minutes	(NJ)	0.985	0.015	0.988	0.012	0.994	0.006	0.998	0.002
	(J)	0.319	0.681	0.399	0.601	0.471	0.529	0.572	0.428
5 minutes	(NJ)	0.987	0.013	0.987	0.013	0.990	0.010	0.996	0.004
	(J)	0.377	0.623	0.377	0.623	0.464	0.536	0.536	0.464
30 minutes	(NJ)	0.984	0.016	0.984	0.016	0.983	0.017	0.982	0.018
	(J)	0.754	0.246	0.775	0.225	0.761	0.239	0.768	0.232
		$(i = 1)$							
1 minutes	(NJ)	0.986	0.014	0.986	0.014	0.987	0.013	0.989	0.011
	(J)	0.246	0.754	0.297	0.703	0.413	0.587	0.500	0.500
3 minutes	(NJ)	0.983	0.017	0.985	0.015	0.985	0.015	0.985	0.015
	(J)	0.348	0.652	0.355	0.645	0.435	0.565	0.536	0.464
5 minutes	(NJ)	0.984	0.016	0.984	0.016	0.983	0.017	0.983	0.017
	(J)	0.362	0.638	0.391	0.609	0.449	0.551	0.493	0.507
30 minutes	(NJ)	0.968	0.032	0.967	0.033	0.967	0.033	0.966	0.034
	(J)	0.674	0.326	0.703	0.297	0.725	0.275	0.754	0.246
		$(i = 2)$							
1 minutes	(NJ)	0.985	0.015	0.983	0.017	0.984	0.016	0.983	0.017
	(J)	0.246	0.754	0.297	0.703	0.399	0.601	0.507	0.493
3 minutes	(NJ)	0.981	0.019	0.980	0.020	0.979	0.021	0.980	0.020
	(J)	0.341	0.659	0.362	0.638	0.435	0.565	0.493	0.507
5 minutes	(NJ)	0.974	0.026	0.974	0.026	0.973	0.027	0.974	0.026
	(J)	0.348	0.652	0.413	0.587	0.464	0.536	0.529	0.471
30 minutes	(NJ)	0.933	0.067	0.933	0.067	0.934	0.066	0.933	0.067
	(J)	0.681	0.319	0.681	0.319	0.696	0.304	0.703	0.297

significantly; however, the type I errors increase further beyond the nominal size. The rates for BR without offsetting the returns are equivalent to the values for constant sampling at the highest frequency with staggering the returns at one lag. Moreover, the results for BR0 and BR1 are equivalent. ZMA yields comparable results, see Table 2.21.

Robust Optimal Sampling Rate

The results in Section 2.4.3 showed that the RobustBR produces valid tests for most cases while the standard BR approach tends to be slightly anti-conservative. Table

Table 2.20: Confusion matrices are tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1 with $\lambda = 0.014$ and $\sigma_{jmp} = 1.50$. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
BR0	(NJ)	0.989	0.011	0.987	0.013	0.986	0.014	0.985	0.015
	(J)	0.203	0.797	0.370	0.630	0.493	0.507	0.572	0.428
BR1	(NJ)	0.989	0.011	0.987	0.013	0.988	0.012	0.988	0.012
	(J)	0.210	0.790	0.399	0.601	0.493	0.507	0.609	0.391
					$(i = 1)$				
BR0	(NJ)	0.989	0.011	0.983	0.017	0.979	0.021	0.976	0.024
	(J)	0.210	0.790	0.348	0.652	0.464	0.536	0.565	0.435
BR1	(NJ)	0.989	0.011	0.982	0.018	0.977	0.023	0.974	0.026
	(J)	0.232	0.768	0.391	0.609	0.457	0.543	0.609	0.391
					$(i = 2)$				
BR0	(NJ)	0.985	0.015	0.975	0.025	0.966	0.034	0.953	0.047
	(J)	0.217	0.783	0.326	0.674	0.486	0.514	0.565	0.435
BR1	(NJ)	0.985	0.015	0.976	0.024	0.965	0.035	0.953	0.047
	(J)	0.196	0.804	0.362	0.638	0.478	0.522	0.565	0.435

Table 2.21: Confusion matrices are tabulated for the Z_{TPRM} statistics based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1 with $\lambda = 0.014$ and $\sigma_{jmp} = 1.50$. The return sampling rates are obtained per day using the method by Zhang et al. (2005). Results are presented for four return horizons: one, three, five and thirty minutes. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
ZMA	(NJ)	0.989	0.011	0.987	0.013	0.985	0.015	0.985	0.015
	(J)	0.225	0.775	0.442	0.558	0.551	0.449	0.638	0.362
					$(i = 1)$				
ZMA	(NJ)	0.987	0.013	0.980	0.020	0.973	0.027	0.973	0.027
	(J)	0.203	0.797	0.428	0.572	0.572	0.428	0.572	0.428
					$(i = 2)$				
ZMA	(NJ)	0.985	0.015	0.970	0.030	0.960	0.040	0.953	0.047
	(J)	0.203	0.797	0.384	0.616	0.478	0.522	0.609	0.391

2.22 presents the power for RobustBR. The behavior is similar to the patterns for BR discussed surrounding Table 2.20. Notice that RobustBR has in most instances more power without being invalid while BR often is slightly anti-conservative.

Table 2.22: Confusion matrices are tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1 with $\lambda = 0.014$ and $\sigma_{jmp} = 1.50$. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The results denoted by RobustBR are computed by sampling rates as described in Section 2.2.2. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
RobustBR0	(NJ)	0.988	0.012	0.988	0.012	0.990	0.010	0.990	0.010
	(J)	0.200	0.800	0.393	0.607	0.481	0.519	0.570	0.430
RobustBR1	(NJ)	0.988	0.012	0.989	0.011	0.988	0.012	0.990	0.010
	(J)	0.178	0.822	0.333	0.667	0.504	0.496	0.585	0.415
					$(i = 1)$				
RobustBR0	(NJ)	0.987	0.013	0.987	0.013	0.984	0.016	0.982	0.018
	(J)	0.170	0.830	0.363	0.637	0.444	0.556	0.570	0.430
RobustBR1	(NJ)	0.986	0.014	0.987	0.013	0.984	0.016	0.983	0.017
	(J)	0.170	0.830	0.348	0.652	0.496	0.504	0.556	0.444

Since RobustBR is proposed as a more robust alternative to jumps compared to BR, I compare the power of the Z_{TPRM} statistic using these two methods for more severe jump processes. I compute the power for different jump rates and jump sizes for prices with and without noise, see Table 2.23. Panel A clearly shows that the RobustBR has higher power in all instances. The difference is smaller for noisy prices, see Panel B.

Hence, these empirical results suggest that the robust version of BR should be used in place of BR since (1) it remains valid for most cases in which BR is slight anti-conservative; and (2) it has higher power.

Alternative Jump Distributions

The jump distribution that I have considered thus far follows the jump-diffusion model proposed by Merton (1976), where the jump size, κ, is normally distributed. Alternative jump distributions that produce similar characteristics of the return distribution include the double exponential, which Kou (2002) and Glasserman (2004)

Table 2.23: The power is tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1. The jump rate, λ, is 0.5, 1.0, 1.5 and 2.0. The standard deviations of the jumps, σ_{jmp}, are 0.5, 1.5 and 2.5. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The results denoted by RobustBR are computed by sampling rates as described in Section 2.2.2. The test size is 0.01.

		σ_{jmp}		
		0.5	1.5	2.5
		Panel A: $\sigma_{mn} = 0.000$		
BR	0.5	0.525	0.814	0.885
	1.0	0.522	0.823	0.892
	1.5	0.523	0.829	0.896
	2.0	0.525	0.831	0.896
RobustBR	0.5	0.598	0.846	0.906
	1.0	0.593	0.855	0.910
	1.5	0.603	0.861	0.918
	2.0	0.600	0.857	0.917
		Panel B: $\sigma_{mn} = 0.080$		
BR	0.5	0.056	0.436	0.617
	1.0	0.064	0.437	0.627
	1.5	0.062	0.452	0.636
	2.0	0.062	0.443	0.638
RobustBR	0.5	0.064	0.440	0.621
	1.0	0.066	0.451	0.632
	1.5	0.064	0.456	0.647
	2.0	0.065	0.456	0.644

apply. Kou (2002) argues that the double exponential jump diffusion model (DExp) produces properties similar to those under the normal assumption but also enables closed-form solutions for a wider range of option-pricing problems. Ramezani and Zeng (1999) find that the same double exponential jump-diffusion model improves the empirical fit of the normal jump diffusion model to stock price data (Merton (1976)).

Furthermore, I also consider a skewed-normal jump distribution (SN) (Azzalini (1985)) to add a negative skewness to the jump returns. This distribution includes a shape parameter along with location and scale parameters, where the distribution reduces to the standard normal when the shape parameter is zero.[6] The skewed-normal distribution has previously been considered as a candidate for the return

[6]The probability density function is given by $f(z;\mu,\sigma^2,\alpha) = 2\phi(z;\mu,\sigma^2)\Phi(\alpha z;\mu,\sigma^2)$, where μ, σ^2, and α are the location, scale, and shape parameters, and $\phi(x)$ and $\Phi(x)$ are the normal probability and cumulative density functions.

distribution. Keel (2006), for example, finds that the skewed-normal provides a poor fit since it does not produce heavy tails. The distribution may, however, be more appropriate for the jump size.

Figure 2.12 compares estimates of density functions based on 1000 deviates from the normal, $\mathrm{N}(0, 1.5)$, double exponential, $\mathrm{DExp}(\eta_1 = 1.1, \eta_2 = 0.8)$, and skewed normal $\mathrm{SN}(\mu = 0, \sigma = 2.5, \alpha = -1)$, distributions.

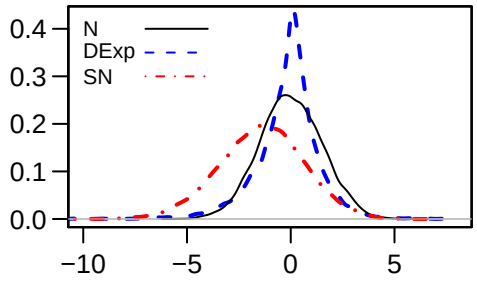

Figure 2.12: The figure graphs estimated density functions based on 1000 deviates from the $\mathrm{N}(0, 1.5)$, $\mathrm{DExp}(\eta_1 = 1.1, \eta_2 = 0.8)$, and $\mathrm{SN}(0, 2.5, -1)$.

Recall that the asymptotic properties of the nonparametric jump statistics do not rely on the normality assumption of the jump size, thus remain valid for these alternative distribution. Nevertheless, the small sample properties may vary. Therefore, I examine the power of the tests under these distributions. To the best of my knowledge, literature has only considered the normal jump distribution in this context hitherto.

Tables 2.24 - 2.27 present confusion matrices based on constant and optimal sampling (BR) rates for the two alternative jump size distributions. The power is higher under the skewed-normal distribution, which is consistent with that the jumps are generally larger in magnitude compared to under the double exponential distribution. The power is in fact low for the latter distribution which can be explained by that a large fraction of the jumps are small, thus indistinguishable from the diffusion and noise processes.

Table 2.24: Confusion matrices are tabulated for Z_{TPRM} based on 10000 days simulated from the SV1FJ model, equation (2.44) with $\lambda = 0.014$ and $\kappa \sim \text{DExp}(\eta_1 = 1.1, \eta_2 = 0.8)$. An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.052, 0.080. Results are presented for four return horizons: one, three, five and thirty minutes. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
		$(i = 0)$							
1 minutes	(NJ)	0.989	0.011	0.996	0.004	0.999	0.001	1.000	0.000
	(J)	0.331	0.669	0.434	0.566	0.574	0.426	0.676	0.324
3 minutes	(NJ)	0.988	0.012	0.990	0.010	0.995	0.005	0.999	0.001
	(J)	0.456	0.544	0.471	0.529	0.551	0.449	0.610	0.390
5 minutes	(NJ)	0.986	0.014	0.988	0.012	0.990	0.010	0.994	0.006
	(J)	0.493	0.507	0.478	0.522	0.529	0.471	0.596	0.404
30 minutes	(NJ)	0.988	0.012	0.987	0.013	0.986	0.014	0.987	0.013
	(J)	0.779	0.221	0.779	0.221	0.757	0.243	0.772	0.228
		$(i = 1)$							
1 minutes	(NJ)	0.989	0.011	0.989	0.011	0.990	0.010	0.990	0.010
	(J)	0.353	0.647	0.375	0.625	0.507	0.493	0.610	0.390
3 minutes	(NJ)	0.987	0.013	0.986	0.014	0.986	0.014	0.988	0.012
	(J)	0.456	0.544	0.456	0.544	0.537	0.463	0.574	0.426
5 minutes	(NJ)	0.985	0.015	0.985	0.015	0.984	0.016	0.984	0.016
	(J)	0.493	0.507	0.507	0.493	0.566	0.434	0.618	0.382
30 minutes	(NJ)	0.980	0.020	0.983	0.017	0.980	0.020	0.982	0.018
	(J)	0.787	0.213	0.765	0.235	0.735	0.265	0.743	0.257

Table 2.25: Confusion matrices are tabulated for the Z_{TPRM} statistic based on 10000 days simulated from the SV1FJ model, equation (2.44). The experimental design is described in Table 2.1, with $\lambda = 0.014$ and $\kappa \sim \text{DExp}(\eta_1 = 1.1, \eta_2 = 0.8)$. An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.052, 0.080. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
		$(i = 0)$							
BR0	(NJ)	0.989	0.011	0.986	0.014	0.987	0.013	0.986	0.014
	(J)	0.309	0.691	0.515	0.485	0.588	0.412	0.662	0.338
BR1	(NJ)	0.989	0.011	0.986	0.014	0.988	0.012	0.987	0.013
	(J)	0.309	0.691	0.471	0.529	0.610	0.390	0.654	0.346
		$(i = 1)$							
BR0	(NJ)	0.989	0.011	0.987	0.013	0.983	0.017	0.984	0.016
	(J)	0.316	0.684	0.500	0.500	0.603	0.397	0.669	0.331
BR1	(NJ)	0.989	0.011	0.986	0.014	0.985	0.015	0.983	0.017
	(J)	0.287	0.713	0.493	0.507	0.610	0.390	0.684	0.316

Table 2.26: Confusion matrices are tabulated for Z_{TPRM} based on 10000 days simulated from the SV1FJ model, equation (2.44) with $\lambda = 0.014$ and $\kappa \sim \text{SN}(\mu = 0, \sigma = 2.5, \alpha = -1)$. An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.052, 0.080. Results are presented for four return horizons: one, three, five and thirty minutes. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
1 minutes	(NJ)	0.989	0.011	0.996	0.004	0.999	0.001	1.000	0.000
	(J)	0.140	0.860	0.199	0.801	0.294	0.706	0.419	0.581
3 minutes	(NJ)	0.988	0.012	0.990	0.010	0.995	0.005	0.999	0.001
	(J)	0.221	0.779	0.250	0.750	0.294	0.706	0.353	0.647
5 minutes	(NJ)	0.986	0.014	0.988	0.012	0.990	0.010	0.994	0.006
	(J)	0.235	0.765	0.257	0.743	0.316	0.684	0.397	0.603
30 minutes	(NJ)	0.988	0.012	0.987	0.013	0.986	0.014	0.987	0.013
	(J)	0.640	0.360	0.632	0.368	0.654	0.346	0.662	0.338
					$(i = 1)$				
1 minutes	(NJ)	0.989	0.011	0.989	0.011	0.990	0.010	0.990	0.010
	(J)	0.140	0.860	0.184	0.816	0.235	0.765	0.316	0.684
3 minutes	(NJ)	0.987	0.013	0.986	0.014	0.986	0.014	0.988	0.012
	(J)	0.221	0.779	0.235	0.765	0.279	0.721	0.316	0.684
5 minutes	(NJ)	0.985	0.015	0.985	0.015	0.984	0.016	0.984	0.016
	(J)	0.250	0.750	0.265	0.735	0.331	0.669	0.346	0.654
30 minutes	(NJ)	0.980	0.020	0.983	0.017	0.980	0.020	0.982	0.018
	(J)	0.574	0.426	0.581	0.419	0.574	0.426	0.603	0.397

Table 2.27: Confusion matrices are tabulated for Z_{TPRM} based on 10000 days simulated from the SV1FJ model, equation (2.44) with $\lambda = 0.014$ and $\kappa \sim \text{SN}(\mu = 0, \sigma = 2.5, \alpha = -1)$. An iid $\text{N}(0, \sigma_{mn})$ noise process is added to the simulated prices; σ_{mn} is set to 0.000, 0.027, 0.052, 0.080. BR1 and BR0 denote sampling rates that are obtained by solving equation (2.33) and by equation (2.34), respectively. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is 0.01.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
BR0	(NJ)	0.989	0.011	0.986	0.014	0.987	0.013	0.986	0.014
	(J)	0.132	0.868	0.250	0.750	0.360	0.640	0.412	0.588
BR1	(NJ)	0.989	0.011	0.986	0.014	0.988	0.012	0.987	0.013
	(J)	0.140	0.860	0.250	0.750	0.338	0.662	0.426	0.574
					$(i = 1)$				
BR0	(NJ)	0.989	0.011	0.987	0.013	0.983	0.017	0.984	0.016
	(J)	0.132	0.868	0.243	0.757	0.353	0.647	0.397	0.603
BR1	(NJ)	0.989	0.011	0.986	0.014	0.985	0.015	0.983	0.017
	(J)	0.125	0.875	0.235	0.765	0.338	0.662	0.397	0.603

2.5 Summary and Conclusions

In this study, I examine small samples properties of nonparametric statistics developed by Barndorff-Nielsen and Shephard (2004b, 2006) to test for jumps in asset

prices. I particularly study the impact of adding noise to the price process and recent methods to contend with such contamination.

I provide several contributions to the literature. Previous studies have established that market microstructure noise biases the statistics against identifying jumps. I provide empirical evidence showing that this can be attributed to that both the mean and variance of the test statistics become negatively biased as the noise increases. As a result, the rejection rates decreases since the test is based on the right tail.

Second, Bandi and Russell (2006) and Zhang et al. (2005) propose methods for reducing the impact of noise when estimating the daily integrated variance using high-frequency data. I show that these methods also perform well to determine the optimal sampling rates for the bipower and tripower variations.

Third, by applying the optimal sampling methods by Bandi and Russell (2006) and Zhang et al. (2005), the test size is closer to the asymptotic results under the null hypothesis and also increases the power of the test statistics. The two methods perform similarly to applying staggered returns, which Huang and Tauchen (2005) evaluate. I propose a modified version of the method by Bandi and Russell (2006) to improve its robustness to jumps. I find that the modified method produces valid jump statistics in most scenarios where the original procedure is slightly anti-conservative. Furthermore, the modified approach has greater power.

Fourth, Bandi and Russell (2006) give two equations for computing the optimal sampling rate; one that is exact and one approximation. The former requires an optimization routine while the second has a simple closed-form solution. I find that the two methods produces equivalently results, thus there is no significant loss to use the approximation, which is faster to compute.

Fifth, the size and power of the test statistics are similar for the three noise processes that I consider. Previous literature has evaluated the statistics for a normal

iid noise process. I find that the jump statistics perform equivalently when adding serial correlation to the iid noise process and introducing rounding errors of the prices.

A number of future research directions has emerged that I have not pursued in this work. The experimental design generates prices per second throughout the trading day. In practice, there are longer and irregular gaps in between trades depending on the liquidity of the markets. Certainly since the statistics rely on intraday data, this approach can only be applied to liquid markets; nonetheless, it is doubtful that prices can be observed at such high and regular frequencies on a daily basis. Hence, an interesting extension is to consider irregular durations in between price observations. Second, I determine the optimal sampling rate per day, which relies on that the market is sufficiently liquid, which is not true in practice in many markets. Future research may determine whether it is beneficial to estimate the rate over a longer time horizon. As a result, the method may yield more robust estimates of the optimal sampling rates and still update the sampling rate sufficiently often. Third, more recent nonparametric methods based on high-frequency data have been proposed, such as Fan and Wang (2007), Jiang and Oomen (2008), and Sen (2008). The finite sample properties of these and the statistics that I consider in this study can be compared.

Bibliography

Bibliography

Aït-Sahalia, Y., Mykland, P. A., and Zhang, L. (2006). Ultra high frequency volatility estimation with dependent microstructure noise. NBER Working Paper No. W11380.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility. *Review of Economics and Statistics*, 89(4):701–720.

Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 12:171–178.

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

Bandi, F. M. and Russell, J. R. (2006). Separating microstructure noise from volatility. *Journal of Financial Economics*, 79:655–692.

Barndorff-Nielsen, O. E., Graversen, S. E., Jacod, J., Podolskij, M., and Shephard, N. (2006). *From Stochastic Analysis to Mathematical Finance, Festschrift for Albert Shiryaev*. Springer Verlag.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society*, Series B 64:253–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2004a). How accurate is the asymptotic approximation to the distribution of realized volatility. In Andrews, E., Powell, J., Ruud, P., and Stock, J., editors, *Identification and Inference for Econometric Models. A Festschrift in Honour of T. J. Rothenberg*. Cambridge University Press.

Barndorff-Nielsen, O. E. and Shephard, N. (2004b). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–48.

Barndorff-Nielsen, O. E. and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1):1–30.

Cox, J. C. and Rubinstein, M. (1985). *Options Markets*. Prentice-Hall, Inc.

Duffie, D. and Pan, J. (2001). Analytical value-at-risk with jumps and credit risk. *Finance and Stochastics*, 5:155–180.

Eraker, B., Johannes, M., and Polson, N. (2003). The impact of jumps in volatility and returns. *Journal of Finance*, 53(3):1269–1300.

Fan, J. and Wang, Y. (2007). Multi-scale jump and volatility analysis for high-frequency financial data. *Journal of the American Statistical Association*, 102(480):1349–1362.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.

Huang, X. and Tauchen, G. (2005). The relative contribution of jumps to total price variation. *Journal of Financial Econometrics*, 3:456–499.

Jacod, J., Li, Y., Mykland, P. A., Podolskij, M., and Vetter, M. (2007). Microstructure noise in the continuous case: The pre-averaging approach. Manuscript University of Chicago.

Jacod, J. and Shiryaev, A. N. (1987). *Limit Theorems for Stochastic Processes*. Springer Verlag.

Jarrow, R. A. and Rosenfeld, E. R. (1984). Jump risks and the intertemporal capital asset pricing model. *Journal of Business*, 57(3):337–351.

Jiang, G. J. and Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise - A "Swap Variance" approach. *Journal of Econometrics*, 144(2):352–370.

Keel, S. T. (2006). *Optimal Portfolio Construction and Active Portfolio Management Including Alternative Investments*. PhD thesis, Swiss Federal Institute of Technology Zurich.

Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101.

Li, Y. and Mykland, P. A. (2007). Are volatility estimators robust with respect to modeling assumptions? *Bernoulli*, 13(3):601–622.

Maheu, J. M. and McCurdy, T. H. (2004). News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance*, 59:755–793.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125.

R Development Core Team (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Ramezani, C. A. and Zeng, Y. (1999). Maximum likelihood estimation of asymmetric jump-diffusion processes: Application to security prices. SSRN: Working Paper Series.

Sen, R. (2008). Jumps and microstructure noise in stock price volatility. In Gregoriou, G. N., editor, *Stock Market Volatility*. Chapman Hall-CRC/Taylor and Francis.

Székely, G. J. and Rizzo, M. L. (2005). A new test for multivariate normality. *Journal of Multivariate Analysis*, 93(1):58–80.

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*, 100:1394–1411.

Chapter 3: Volatility and Jump Dynamics in U.S. Energy Futures Markets

3.1 Introduction

Observers of energy futures markets have long noted that energy futures prices are very volatile and often exhibit jumps (price spikes) during news event periods. Thus, the assumption of a continuous diffusion process for asset price behavior is often violated in practice. Since volatility behavior is the central topic for option pricing, risk management and asset allocation strategies, market participants, regulators and academics have a strong interest in the identification of jumps over time and measuring the relative importance of the jump component versus the smooth sample path component as contributors to total volatility. Motivated by the increase in the availability of high-frequency data (tick by tick data), Barndorff-Nielsen and Shephard (2004, 2006) and Jiang and Oomen (2008) have developed nonparametric procedures for detecting the presence of jumps in high-frequency intraday financial time series. Jiang et al. (2008) show that these two nonparametric methods can be combined to produce a test that remains powerful but more robust to noise in the price series. Despite this, there has been no empirical work using this newly developed procedure to investigate the presence of jumps over time and the relative contribution of jumps to the volatility of energy futures prices. This paper seeks to fill this gap in the literature.

Recent literature based on parametric approaches to identify and model jumps in stock returns include Chan and Maheu (2002), Maheu and McCurdy (2004) and

others. Chan and Maheu (2002) propose an autoregressive conditional jump intensity within a GARCH model approach to detect and model jumps in seventy-two years of daily stock returns. They find a significant time variation in the conditional jump intensity and in the jump size in stock returns during the sample period. Maheu and McCurdy (2004) model conditional variance of returns as a combination of jumps and smoothly changing components.

Literature using nonparametric methods include Huang and Tauchen (2005), who perform a Monte Carlo study on small sample properties of the nonparametric procedure for detecting jumps by Barndorff-Nielsen and Shephard (2004, 2006). Their results indicate that microstructure noise biases the test against detecting jumps and suggest applying a simple lagging strategy to correct the bias. They also provide evidence that jumps account for seven percent of the S&P 500 index's realized variance. Using the same nonparametric approach, Andersen et al. (2007) provide empirical evidence that the volatility jump component is both highly significant and less persistent than the continuous sample path component in foreign exchange rate spot (DM/$) market, US S&P 500 index futures and thirty-year US Treasury bond futures.

Previous literature on investigating volatility behavior of energy futures prices include Pindyck (2004), Linn and Zhu (2004), Ates and Wang (2007), Mu (2007), Wang et al. (2008) and others. Pindyck (2004) documents a significant positive trend in natural gas futures during the sample period from May 2, 1990 to February 2, 2003. Linn and Zhu (2004) report an increase in volatility before and after the release of inventory reports by the Energy Information Administration. Ates and Wang (2007) use a nonlinear error-correction model with a multivariate GARCH errors process to document that extreme cold weather surprises and inventory surprises are the short-run demand and supply factors that affect the spot and futures price change volatility

in natural gas and heating oil markets. Mu (2007) finds that extreme weather conditions and low inventories are important factors affecting natural gas futures volatility within a single equation model with a GARCH error process. Wang et al. (2008) examine the realized volatility and correlation of crude oil and natural gas futures. They provide evidence that realized crude oil futures volatility increases in the weeks immediately before OPEC recommends price increases. However, none of these papers dealing with energy price volatility have separated the volatility jump component from the volatility smooth component and examined the relative importance of jump versus smooth components in the total price volatility.

This paper makes several contributions to the literature on detecting jump components and in analyzing the time series properties of jumps in energy futures prices. First, I examine the realized volatility behavior of natural gas, heating oil and crude oil futures contracts traded on the New York Mercantile Exchange (NYMEX) using high-frequency intraday data from January 1990 to January 2008. Second, using a nonparametric test statistic proposed by Jiang et al. (2008), I identify significant jump components in energy futures prices and estimate the relative contribution of jumps to the realized variance in the three futures contracts. Third, I investigate whether significant jumps are often associated with Energy Information Administration (EIA)'s inventory news announcement dates and extreme cold weather periods. Fourth, I test whether including jump and seasonal components as explanatory variables improve the modeling and forecasting of energy futures volatility. Finally, I evaluate the effects of weather and inventory surprises as short-run demand and supply factors on the realized volatility. In addition, I test whether the spread nearby and first deferred contracts can be used as a proxy for these surprises.

The remainder of the paper is organized as follows. Section 3.2 provides the background for the statistical methodology used in this paper. Section 3.3 describes

the data and contract specifications of natural gas, heating oil and crude oil futures. Empirical results are reported in Section 3.4, and Section 3.5 presents a summary and conclusions.

3.2 Background of Statistical Methodology

3.2.1 Asset Price Dynamics and Jump Statistics

Let $X_t = \log S_t$ denote the logarithmic price where S_t is the observed price at time t. Assume that the logarithmic price process, X_t, follows a continuous-time diffusion process coupled with a discrete process defined as,

$$\mathrm{d}X_t = \mu_t \mathrm{d}t + \sigma_t \mathrm{d}W_t + \kappa_t \mathrm{d}q_t, \tag{3.1}$$

where μ_t is the instantaneous drift process and σ_t is the diffusion process; W_t is the standard Wiener process; $\mathrm{d}q_t$ is a Poisson jump process with intensity λ_t, that is, $P(\mathrm{d}q_t = 1) = \lambda_t \mathrm{d}t$; and κ_t is the logarithmic size of the price jump at time t if a jump occurred. If X_{t-} denotes the price immediately prior to the jump at time t, then $\kappa_t = X_t - X_{t-}$. Define the intraday return, r_{t_j}, as the difference between logarithmic prices, $r_{t_j} = X_{t_j} - X_{t_{j-1}}$, where $t_j - t_{j-1}$ is the discrete intraday sample period, Δ.

Bipower Variations

Barndorff-Nielsen and Shephard (2004, 2006) propose a number of nonparametric statistics based on realized power variations to test for jumps and to estimate the contribution of jumps to the total variation. Specifically, the statistics are based on the difference between two estimators of the integrated daily cumulative variation.

The *realized (quadratic) variance* is defined as the sum of squared intraday returns,

$$\mathrm{RV}_t = \sum_{j=1}^{m_t} r_{t_j}^2, \tag{3.2}$$

where m_t is the number of intraday Δ-returns during the tth trading day and is assumed to be an integer. Jacod and Shiryaev (1987) show that the realized variation converges to the integrated daily variation under the assumption that the underlying process follows equation (3.1) without jumps. Furthermore, in the presence of jumps ($\lambda > 0$), the realized variance converges in probability to the total variation as $\Delta \to 0$,

$$\mathrm{RV}_t \xrightarrow{p} \int_{t-1}^{t} \sigma_s^2 \mathrm{d}s + \sum_{t<s<t+1} \kappa^2(s). \tag{3.3}$$

Hence, the realized variance captures the variance of both the continuous and the discrete processes. A second estimator of the integrated daily variance is the *realized bipower variation*, which is defined as,

$$\mathrm{BV}_t = \frac{\pi}{2} \frac{m_t}{m_t - 1} \sum_{j=2}^{m_t} |r_{t_j}||r_{t_{j-1}}|. \tag{3.4}$$

Barndorff-Nielsen and Shephard (2004) show that as $\Delta \to 0$,

$$\mathrm{BV}_t \xrightarrow{p} \int_{t-1}^{t} \sigma_s^2 \mathrm{d}s. \tag{3.5}$$

Hence, the asymptotic convergence of the bipower variation captures only the effects of the continuous process even in the presence of a jump process.[1] By combining the results from equations (3.3) and (3.5), the contribution of the jump process in the

[1]The result follows from that only a finite number of terms in the sum in equation (3.4) are affected by jumps while the remaining returns converges to zero in probability. Since the probability of jumps goes to zero as $\Delta \to 0$, those terms do not impact the limiting probability.

total quadratic variation can be estimated by the difference of these two variations where,

$$\mathrm{RV}_t - \mathrm{BV}_t \xrightarrow{p} \sum_{t<s<t+1} \kappa^2(s), \tag{3.6}$$

as $\Delta \to 0$. Hence, equation (3.6) estimates the integrated variation due to the jump component and, as such, provides the basis for a nonparametric test for identifying jumps.

Barndorff-Nielsen and Shephard (2004, 2006) and Barndorff-Nielsen et al. (2006) show that in the absence of jumps in the price process,

$$\Delta^{-1/2} \frac{\mathrm{RV}_t - \mathrm{BV}_t}{\left(\left(\nu_{bb} - \nu_{qq}\right) \int_{t-1}^{t} \sigma^4(s) \mathrm{d}s\right)^{1/2}} \xrightarrow{p} \mathrm{N}(0,1), \tag{3.7}$$

as $\Delta \to 0$ where RV_t and BV_t are defined in equations (3.2) and (3.4) and $\nu_{bb} = \pi^2/2 + \pi - 3$ and $\nu_{qq} = 2$. The integral in the denominator, the *integrated quarticity*, is unobservable. From the work by Barndorff-Nielsen and Shephard (2004) on multipower variations, Andersen et al. (2007) propose to estimate the integrated quarticity using the *realized tripower quarticity*, TP_t, which is defined as,

$$\mathrm{TP}_t = m_t \mu_{4/3}^{-3} \frac{m_t}{m_t - 2} \sum_{j=3}^{m_t} \prod_{i=0}^{2} |r_{t_{j-i}}|^{4/3}, \tag{3.8}$$

where $\mu_{4/3}$ is a constant given by,

$$\mu_k = 2^{k/2} \frac{\Gamma((k+1)/2)}{\Gamma(1/2)}. \tag{3.9}$$

Asymptotically, as $\Delta \to 0$,

$$\mathrm{TP}_t \xrightarrow{p} \int_{t-1}^{t} \sigma_s^4 \mathrm{d}s. \tag{3.10}$$

Hence, a test statistic based on equation (3.7) is given by,

$$\Delta^{-1/2}\frac{\mathrm{RV}_t - \mathrm{BV}_t}{\left(\left(\nu_{bb} - \nu_{qq}\right)\mathrm{TP}_t\right)^{1/2}}. \tag{3.11}$$

Barndorff-Nielsen and Shephard (2004, 2006) propose a number of variations of the statistic in equation (3.11), all of which asymptotically have a standard normal distribution. The small sample properties of these statistics are evaluated by Huang and Tauchen (2005) in a Monte Carlo study. They find that the following statistic has the best finite sample properties,

$$Z_{\mathrm{bns},t} = \frac{\mathrm{RJ}_t}{\sqrt{\left(\nu_{bb} - \nu_{qq}\right)\frac{1}{m_t}\max\left\{1, \frac{\mathrm{TP}_t}{\mathrm{BV}_t^2}\right\}}}, \tag{3.12}$$

where,

$$\mathrm{RJ}_t = \frac{\mathrm{RV}_t - \mathrm{BV}_t}{\mathrm{RV}_t}. \tag{3.13}$$

The $Z_{\mathrm{bns},t}$ statistic in equation (3.12) can be applied to test the null hypothesis that there is no jump in the return process during a trading day, t, where the hypothesis is rejected for large positive values of the statistic relative to the standard normal distribution. The test is one-sided since the statistic is based on the difference between variances where the difference is zero under the null hypothesis and greater than zero otherwise.

Swap Variance

Jiang and Oomen (2008) base a statistic to test for jumps in asset prices on the variance swap replication strategy (Neuberger (1994)). This strategy allows traders to hedge their exposure to volatility risk more effectively than by using traditional

put or call options. The hedge portfolio is based on that the accumulated difference between the simple return and the logarithmic return is one half of the integrated variance under the assumption that the asset price process is continuous. The relation between the two return measures breaks down, however, if the data-generating process has discontinuities in the price process, which Jiang and Oomen (2008) use to develop a test for jumps.

The price process in equation (3.1) (page 86) with $S_t = \exp(X_t)$ can be written as,

$$\frac{\mathrm{d}S_t}{S_t} = \left(\mu_t + \frac{1}{2}\sigma^2\right)\mathrm{d}t + \sigma_t \mathrm{d}W_t + (\mathrm{e}^{\kappa_t} - 1)\,\mathrm{d}q_t, \tag{3.14}$$

which can be shown to be,

$$2\int_0^1 \left(\frac{\mathrm{d}S_t}{S_t} - \mathrm{d}X_t\right) = \sigma^2_{(0,1)} + 2\int_0^1 (\mathrm{e}^{\kappa_T} - \kappa_t - 1)\,\mathrm{d}q_t. \tag{3.15}$$

In the discrete case, the left-hand side of equation (3.15) is the *swap variance*, which can be estimated by,

$$\mathrm{SwV}_t = 2\sum_{i=1}^{m_t}(R_{t_i} - r_{t_i}), \tag{3.16}$$

where $R_{t_i} = (S_{t_i} - S_{t_{i-1}})/S_{t_{i-1}}$ is the ith intraday *simple return*, r_{t_i} is the equivalent *logarithmic return*, and m_t is the number of intraday returns. Asymptotically, as $m_t \to \infty$,

$$\mathrm{SwV}_t - \mathrm{RV}_t \xrightarrow{p} \begin{cases} 0, & \text{if no jump;} \\ 2\int_{t-1}^{t} \left(e^{\kappa_t} - \frac{1}{2}\kappa_t^2 - \kappa_t - 1\right)\mathrm{d}q, & \text{if jump,} \end{cases} \tag{3.17}$$

where RV_t is the realized variation (equation (3.2)). The result in equation (3.17) follows from equation (3.15) and that $\mathrm{RV}_t \to \int_{t-1}^{t} \sigma_s \mathrm{d}s + \sum_{t<s<t+1} \kappa^2(s)$ (Jacod and Shiryaev (1987)).

Jiang and Oomen (2008) consider three statistics based on equation (3.17). They find in a simulation study that a ratio statistic defined as,

$$Z_{\text{swv},t} = \frac{\widehat{\sigma}_t^2 m_t}{\sqrt{\widehat{\Omega}_t}} \left(1 - \frac{\text{RV}_t}{\text{SwV}_t}\right), \tag{3.18}$$

has the best small sample properties. $\widehat{\Omega}_t$ is an estimator of,

$$\Omega_t = \frac{\mu_6}{9} \int_{t-1}^{t} \left(\sigma_u^2\right)^3 \mathrm{d}u, \tag{3.19}$$

where σ_t is the volatility term in the data-generating process defined in equation (3.1) (page 86) and μ_6 is a constant given by equation (3.9) (page 88). The estimator, $\widehat{\Omega}_t$, is defined by,

$$\widehat{\Omega}_t^{(p)} = \frac{\mu_6}{9} \frac{m_t^3 \mu_{6/p}^{-p}}{m_t - p + 1} \sum_{j=0}^{N-p} \prod_{k=1}^{p} |r_{t+k}|^{6/p}. \tag{3.20}$$

Jiang and Oomen (2008) conclude in simulations studies that four and six are appropriate choices for p.

Combined Statistic

While both the $Z_{\text{bns},t}$ and $Z_{\text{swv},t}$ statistics are based on intraday returns, they have important asymptotic differences. The BNS framework compares the realized variation, RV_t, with an estimator that is robust to jumps (BV_t); the swap-variance approach, on the other hand, compares RV_t with the SwV_t estimator where the latter is sensitive to jumps. Moreover, the swap-variance test is primarily driven by the third power of the returns while the BNS framework is based on the second moment, which has the implication that the former test is two-sided as opposed to BNS which is one-sided. According to Jiang and Oomen (2008), the $Z_{\text{swv},t}$ statistic generally has the same

sign as the jump although that may not be the case. A weakness with the swap-variance test is that multiple jumps with different signs may cancel out the effect of each other while the impact of multiple daily jumps accrues in the BNS framework. Furthermore, SwV_t requires estimating daily integrated variations at a higher order than BNS does. Generally, estimators of higher order terms are less robust to jumps and noise.

Simulation studies (for example Huang and Tauchen (2005) and Jiang et al. (2008)) on these two statistics have shown that both methods may become anti-conservative, which is particularly evident when applying methods to contend with noisy prices. Jiang et al. (2008) propose to address this by only rejecting the null hypothesis when both tests reject. They provide empirical evidence suggesting that the combined version is conservative and powerful. I report a small sample simulation study on the combined statistic in Appendix A.III (page 150). I find that although the tests based on the individual statistics are invalid, the combined method is highly conservative but nearly as powerful. I apply this combined method to the energy market data to increase the validity of the results. In addition, I produce all results based on the BNS framework for robustness.

3.2.2 Decomposing Total Variation

The daily variance due to the jump component is estimated by the difference between RV_t and BV_t, equation (3.6), where RV_t estimates the total variation including the contribution due to the jump component, whereas BV_t is robust to jumps and only captures the variation due to the continuous component. Hence, the difference is zero in absence of jumps and greater than zero otherwise. However, due to measurement errors, the difference can be negative. Barndorff-Nielsen and Shephard (2004) suggest imposing a lower bound at zero by letting the variance due to the jump component

be given by,

$$J_t = \max[\mathrm{RV}_t - \mathrm{BV}_t, 0]. \tag{3.21}$$

Furthermore, since small values of J_t may be due to noise rather than discontinuities in the price process, I identify the variance contributed by *significant* jumps as,

$$J_{t,\alpha} = (\mathrm{RV}_t - \mathrm{BV}_t)\, \mathrm{I}_{(p<1-\alpha)}, \tag{3.22}$$

where p is the p-value which is set to the maximum value of the p-values based on the $Z_{\mathrm{bns},t}$ and $Z_{\mathrm{swv},t}$ statistics; α is the significance level; and I is the indicator function, which is equal to one if the test rejects the null hypothesis and zero otherwise. The variation that is contributed by the continuous sample path component can then be estimated by,

$$C_{t,\alpha} = \mathrm{I}_{(p<1-\alpha)}\mathrm{RV}_t + \mathrm{I}_{(p\geq 1-\alpha)}\mathrm{BV}_t. \tag{3.23}$$

By this definition, the sum of $J_{t,\alpha}$ and $C_{t,\alpha}$ adds up to the total variation, RV_t.

3.2.3 Contending with Market Microstructure Noise

The $Z_{\mathrm{bns},t}$ and $Z_{\mathrm{swv},t}$ test statistics depend on estimates of daily integrated variations, which are obtained with model-free methods on high-frequency intraday data. The asymptotic properties of the realized variance, RV_t, and bipower variation, BV_t, assume an efficient price process. Observed prices, however, are noisy due to market microstructure. As a result, the variance in high-frequency returns can be attributed to two components: efficient price returns and microstructure frictions. The variance generated by market frictions is the result of the price formation under specific trade mechanisms and rules, such as the discrete price grid and bid-ask bounce effects. Such noise introduces bias in the variance estimates and becomes particularly severe

at high sampling rates. In fact, Bandi and Russell (2006) among others show that as the sampling interval goes to zero, the variance due to noise rather than the integrated variance will dominate the estimate.

The conventional approach relied on in the applied literature to alleviate the bias is simply to sample the price process at lower frequencies than what the data permit. The sampling intervals are typically arbitrarily chosen and commonly in the range of five to thirty minutes.

In addition to sparse sampling, Andersen et al. (2007), Barndorff-Nielsen and Shephard (2006), and Huang and Tauchen (2005) propose and evaluate a method, referred to as staggered returns to further reduce the impact of the microstructure noise.[2] In particular, the method addresses the bias generated by spurious correlations in the returns due to noise. The bid-ask bounce, for example, may induce negative autocorrelations in the intraday asset price returns as the trades are executed at the spread slightly above and below the fair value. Moreover, the practice to split large trades into several smaller trades that are executed during a relatively short time horizon may induce positive autocorrelation. Any such autocorrelation structure in the returns may bias the bipower and tripower estimators since these are functions of adjacent returns. The proposed method attempts to break or at least reduce the correlation structure by skipping one or more intraday returns when computing these estimates rather than taking adjacent returns. The bipower variation using staggered returns becomes,

$$\mathrm{BV}_{t+i} = \frac{\pi}{2} \frac{m_t}{m_t - 1 - i} \sum_{j=2+i}^{m} |r_{t_j}||r_{t_{j-1-i}}|. \tag{3.24}$$

[2]Bandi and Russell (2006) and Zhang et al. (2005) independently propose alternative methods for reducing the impact of noise when estimating the daily integrated variance using high-frequency data. They find an optimal sampling rate for estimating the realized variance by quantifying the trade-off between the bias due to noise at high sampling frequencies and the variance due to infrequent sampling.

The offset, i, is chosen based on the order of the autocorrelation in the return process; for example, if the autocorrelation is only significant at one lag, an offset $i = 1$ may be sufficient. Similarly, the definition of the tripower estimator is modified to allow for staggered returns as,

$$\mathrm{TP}_t = m_t \mu_{4/3}^{-3} \frac{m_t}{m_t - 2(1+i)} \sum_{j=1+2(1+i)}^{m_t} \prod_{k=0}^{2} |r_{t_{j-k(1+i)}}|^{4/3}. \tag{3.25}$$

3.3 Contract Specifications and Data

In this study, I examine price series for three contracts from the U.S. energy futures markets. The contracts are on natural gas, crude oil, and heating oil, all of which are traded on the New York Mercantile Exchange (NYMEX).

The natural gas futures contract is commonly cited as the benchmark for the spot market, which accounts for nearly 25 percent of the energy consumption in the U.S. The futures contract began trading on April 3, 1990 and is based on delivery at the Henry Hub in Louisiana. The contract trades in units of $10,000$ million British thermal units (mmBtu) and is quoted in dollars and cents per mmBtu. Contracts are traded for about thirteen years forward (the current calendar year plus the next twelve years). The minimum price fluctuation is \$0.001 per mmBtu (\$10.00 per contract) and a five-minute intraday trading halt is triggered by a price movement of \$3.00 per mmBtu (\$30,000 per contract).

The futures contract on crude oil began trading in 1983 and, according to NYMEX, is the world's most liquid futures contract on a physical commodity. The contract calls for delivery of both domestic as well as international crude oils of different grades in Cushing, Oklahoma. The contract, which is listed nine years forward, trades in units of $1,000$ U.S. barrels ($42,000$ gallons) and is quoted in U.S. dollars and cents

per barrel. The minimum price fluctuation for crude oil is $0.01 per barrel ($10.00 per contract) and five-minute intraday trading halts are triggered by price movements of $10.00 per barrel ($10,000 per contract) in either direction.

The heating oil futures contract began trading on November 14, 1978. Heating oil currently accounts for about a fourth of the yield of a barrel of crude oil, second only to gasoline. The futures contract trades in units of 42,000 gallons (1,000 barrels) and is based on delivery in the New York Harbor. Separate contracts are traded for 36 consecutive months. The price is quoted in dollars and cents per gallon. The minimum price fluctuation is $0.0001 per gallon ($4.20 per contract) and intraday trading halts are triggered by price movements of $0.25 per gallon ($10,500 per contract) in either direction. Detailed specifications for all three contracts are presented in Appendix A.I.

Part of the intraday data is from the Commodity Futures Trading Commission's database. The data series for crude oil and heating oil range from January 1, 1990 to May 31, 2002 and the series for the natural gas contract span from January 1, 1993 to March 31, 2004. This dataset is supplemented with more recent data obtained from the Institute for Financial Markets, Futures Industry Association, and extend the crude oil and heating oil series to December 31, 2007 and the natural gas series to January 31, 2008. Each transaction includes a date and time stamp and the transaction price. Prior to September 11, 2001, the open outcry trading hours were 9:45 AM to 3:10 PM. Following 9/11, the trading commenced on September 17, 2001 with varying trading times until October 1, 2001, when trading hours were set at 10:00 AM to 2:30 PM. Since January 31, 2007, the trading hours have been set at 9:00 AM to 2:30 PM.

The contracts began trading electronically via the Globex trading platform in the spring of 2007. While liquidity initially was low on the Globex platform, it increased

rapidly during the summer of 2007, and electronic trading became consistently higher than pit trading in September of 2007 for all three contracts, and has since remained the predominant trading platform. Consequently, I switch from the open outcry series to the electronic dataset from September 2007 forward. The electronic trading takes place from 6:00 PM to 5:45 PM the following day; however, for consistency I consider only the transactions for the same hours during which the pit trading takes place.

Furthermore, I use the data series from nearby contract months. During the maturity month, I shift to the first deferred contract month, using the daily trading frequency as the switch indicator. The data are filtered to limit any biased results due to illiquid trading.[3]

Data on stock levels are obtained from the Department of Energy's (DOE) Energy Information Administration (EIA). The inventories of crude oil and heating oil are published by EIA as part of their weekly petroleum status report.[4] EIA also provides estimations of weekly volumes of working gas in U.S. underground natural gas storage facilities in its historical weekly storage estimates database. Historical daily weather data is obtained from the National Oceanic and Atmospheric Administration (NOAA) database.[5,6] The database includes daily minimum and maximum temperatures in degrees Fahrenheit for Chicago and New York City. Chicago weather is used for the natural gas analysis since Chicago is one of the main consumption areas. Heating oil's main market is the Northeastern U.S. with New York State having the highest demand. Therefore, New York weather is chosen for the heating oil and crude oil analysis. I match the sample ranges of the stock levels and temperatures with the

[3]Since the daily variation estimators rely on frequent intraday trading, days with fewer than three hours of trading and days with a trading gap longer than two hours are excluded from the study.

[4]Distillate oil can be categorized into low and high sulfur classes where the latter class often is referred to as heating oil. Hence, I extract the historical stocks from this class from the petroleum status report as the heating oil inventory.

[5]`http://www.noaa.gov/`

[6]I thank Professor Sheryl Beach at George Mason University and Professor Cary J. Mock at the University of South Carolina for help with acquiring the weather data.

transaction data.

3.4 Empirical Results

This section first examines the persistency and dynamics in the total realized variation and the jump process. Thereafter, seasonal and intraday patterns of the variations and the impact of the jump process on the total realized variation are considered. Finally, I present results on modeling and forecasting total realized variation by incorporating the discontinuous jump component as an explanatory variable in the model. In addition, I add weather and inventory surprises and interest-rate adjusted spread to the model.

3.4.1 Realized Variations and Jump Dynamics

The time series behavior of daily closing prices (top panel) and log-returns (bottom panel) for crude oil, heating oil and natural gas are presented in Panels A-C in Figure 3.1. It clearly exhibits that the closing prices of the three energy markets have generally increased since around 1999.

The Augmented Dickey-Fuller (ADF) test is used to test for the presence of a unit root in realized variance, realized volatility (realized variance in standard deviation form), and log transformation of realized variance and the same forms of the jump component. The first row of Table 3.1 reports the ADF test statistics which indicate that the null hypothesis of unit root is rejected at the 1% level of significance for all series.

The top panel in Figure 2, A-C, shows daily volatilities (realized variance in standard deviation form) for the crude oil, heating oil and natural gas series. Each of the three series exhibits strong autocorrelation. This is confirmed by the Ljung-Box

statistic (LB_{10}), which is equal to $10,926$ for crude oil, $9,263$ for heating oil and $6,184$ for natural gas (see the bottom row of Panel A-C in Table 3.1). A cross-market comparison shows that the natural gas market is the most volatile market; the annualized realized volatilities are 39.4 percent for natural gas, 26.5 percent for heating oil and 26.0 percent for crude oil. The equivalent values for the S&P 500 and the thirty-year U.S. Treasury bond futures over the sample period $1990-2002$ are 14.7 and 8.0 percent, respectively (Andersen et al. (2007)). Hence, the energy futures markets, and particularly the natural gas market, are highly volatile during this sample period. The maximum of the daily volatility, however, is higher for the two oil markets with a daily maximum of 0.1950 for crude oil, which occurred during the period following the invasion of Iraq in 1990. Based on the skewness and excess kurtosis, the logarithmic form appears to be the most normally distributed, which is consistent with previous empirical findings in the equity and foreign exchange markets (Andersen et al. (2007)) although the Jarque-Bera test statistic rejects normality for all forms and markets at the 1% significance level.

The second panel in Figure 3.2, A-C, plots the separate measurement of the jump components in standard deviation form. The jump component is defined as the difference between the realized and bipower variations with a lower bound at zero (equation (3.21)). The mean of the daily volatility due to the jump component is equivalent for crude and heating oil at 0.0033 and 0.0038, respectively, while it is larger for natural gas at 0.0061; the corresponding annualized volatilities are 5.2, 6.0 and 9.7 percent, respectively. The jump component is highly positively skewed with a large kurtosis in all three markets. The Ljung-Box test statistics reported in the bottom row are significant although considerably smaller than for the total volatility. The Ljung-Box statistics for the standard deviation form of the jump components are between 190 and 290 for the three markets while the corresponding statistics

are greater than $6,000$ for the realized volatility of each of the three series. Hence, the smooth component appears to contribute more to the persistency in the total volatility. The second panel in Figure 3.2, A-C, clearly shows that there is some persistence in the jump component.

Since the jump component in Table 3.1 is computed by the difference defined in equation (3.21), the properties and in particular the prevalence of autocorrelation may partially be due to that the estimator captures some of the smooth process on days without jumps.[7] Hence, to alleviate such potential bias, I examine the properties for significant jumps as defined by equation (3.22). The significant jumps are determined by the combined statistic (see Section 3.2) where the bipower and tripower estimators are obtained using staggered returns with one lag offset (equations (3.24) and (3.25)) to reduce the impact of market microstructure noise.[8] The significant jump components based on the test level α set to 0.99 are plotted in the last panel in Figure 3.2, A-C, which clearly exhibits that large volatility often can be associated with a large jump component.

Table 3.2 reports yearly statistics of the significant jump components for α equal to 0.99. There are significant jumps in all three price series. The number of days with a jump ranges from 5 to 34 for natural gas, 5 to 28 for heating oil and 4 to 20 days for crude oil. The proportion of days with jumps in natural gas is higher during the second half of the sample period; the other markets do not reveal the same trend. The table also includes daily summary statistics per year for the relative contribution for days with a significant jump. The relative contribution of the jump component to the total variance ranges from 23 to 87 percent for natural gas futures, 23 to 64 percent for crude oil futures and 23 to 74 percent for heating oil futures for days with

[7]It is generally assumed that the jump component dominates the diffusion component on days with significant jumps.

[8]Equivalent results are obtained when sampling at the rate given by the method by Bandi and Russell (2006), see footnote 1.

jumps. Hence, jumps have a significant impact in all three markets.

Table 3.3 presents a regression analysis to test for linear trends in the daily time series of the realized volatility, the smooth and significant jump components and the relative contribution. The model is specified as,

$$V_t = \beta_0 + \beta_1 \text{Trend}_t + \epsilon_t, \tag{3.26}$$

where V_t denotes the realized variations and ratio, respectively, and Trend_t denotes the trend. The trends in the total volatility are positive and significant for all three markets. In natural gas, the jump component emerges as the driving factor of the trend in the total variation, whereas in the crude and heating oil markets the smooth component is positive and significant while the jump component contributes less. This result is also reflected in the trend in the relative contribution, which is positive and significant for natural gas while negative in the oil markets.

Moreover, I test for linear trends in the monthly estimates of the jump intensity and size of the significant jump component by estimating the model specified in equation (3.26). V_t denotes the intensity and volatility, respectively, and Trend_t denotes the (monthly) trend. The regression results are presented in Table 3.4. The arrival rates of jumps are computed per month by taking the ratio of days with a jump over the total number of trading days. The trends are positive and significant for both series in the natural gas market, while they are negative for crude oil. There are no significant trends in the heating oil market. The time series behavior of the exponentially smoothed series of the intensity and size for the three contracts are graphed in Figure 3.3.

To further examine the jump dynamics, I consider different levels of α ranging from 0.5 to 0.9999. The empirical results are reported in Table 3.5. The first row

tabulates the number of days with a significant jump. As a comparison, the total number of trading days for the complete sample period for natural gas is $3,676$, for crude oil is $4,510$, and for heating oil is $4,449$. As expected, the proportion of days with significant jumps declines from 0.49 to 0.02 for natural gas, 0.49 to 0.01 for heating oil, and from 0.44 to 0.01 for crude oil, as the level of α increases from 0.5 to 0.9999. Andersen et al. (2007) report that the equivalent values for S&P 500 futures and thirty-year U.S. Treasury bond futures are $0.737 - 0.051$ and $0.860 - 0.076$, respectively; thus, jumps are more frequent in the latter markets.[9] Based on the proportions of days with a jump for the energy futures markets, the test statistic consistently rejects the null hypothesis too frequently for the larger test sizes had the underlying data generating process been a continuous diffusion process. For natural gas, 13 percent of the days are identified as having a jump for $\alpha = 0.95$ and 7 percent for $\alpha = 0.99$. Similar percentages hold for the other markets. The sample mean and standard deviations are daily values of the volatility due to significant jumps where the estimates are computed only over days with significant jumps. Hence, the average jump size increases as the significance level increases. The annualized estimates range from 16.0 to 47.1 percent for natural gas, 9.68 to 22.6 percent for crude oil and 10.0 to 24.9 percent for heating oil. The Ljung-Box test statistics for significant jumps ($\text{LB}_{10}, J_\alpha^{1/2}$) are lower than the equivalent values for jumps defined by equation (3.21) reported in Table 3.1. Consistently, the Ljung-Box statistics decrease as the size of α increases. Yet, even as the number of jumps declines, the Ljung-Box statistics indicate that some persistency remains in the jump component. The p-values are less than 0.01 for $\alpha = 0.999$ for all markets and less than 0.01 for $\alpha = 0.9999$ for natural gas. The time series plot of the significant jump component is graphed in the fourth

[9]Andersen et al. (2007) identifies jumps by the BNS framework, which partially explain the differences. The rates using this statistic for the energy markets are 0.64 to 0.02 for natural gas and heating oil, and from 0.44 to 0.01 for crude oil.

panel of Figure 3.2, A-C.

Finally, Table 3.6 presents summary statistics for jump returns conditioned on the sign of the returns. Since the test statistic does not provide the direction of the price change, I define the largest (in magnitude) intraday price return as the jump for each day for which the test rejects the null hypothesis and thus obtain the size and sign of the jump return. From Table 3.6, I observe that there are more negative than positive jumps for all three energy futures markets. The mean and median values are equivalent, however.

In summary, using high-frequency data, I have applied a nonparametric statistical procedure proposed by Barndorff-Nielsen and Shephard (2004, 2006) to decompose total volatility into a smooth sample path component and a jump component for three markets. I find that jump components are less persistent than smooth components and large volatility is often associated with a large jump component. Across the three markets, natural gas futures is the most volatile, followed by heating oil and then by crude oil futures.

3.4.2 Seasonal Effects in Smooth and Jump Components

Previous literature (see Ates and Wang (2007) and others) has documented that extreme cold weather surprises and inventory surprises are the short run demand and supply factors affecting the spot and futures price change volatility in natural gas and heating oil markets. In this section, I investigate whether there is a seasonal (monthly) pattern in the total realized volatility, the smooth and jump components, and the relative contribution of jumps in these markets.

Table 3.7, Panel A-C, presents regression results from testing the equality of monthly means of the total realized volatility, the smooth and jump components,

and the relative contribution of jumps. I estimate the regression model,

$$V_t = \beta_0 + \sum_{i=1, i\neq 7}^{12} \beta_i D_i + \epsilon_t, \tag{3.27}$$

where V_t denotes the realized variations and ratio, respectively, and D_i denotes a dummy variable which is equal to one for the ith calendar month and zero otherwise. The intercept denotes the monthly value of July which is used as the basis of comparison. The monthly coefficients for the realized volatility and smooth sample path component are positive and higher during the winter months compared to those in the summer months for all markets. The top two panels in Figure 3.4, A-C, exhibit U-shaped curves for the realized volatility and smooth sample path components for all three markets during the whole year.

The jump component does not display the same consistent seasonal behavior as the smooth component. The jump components for natural gas and crude oil have peaks in both winter and summer. The monthly jump intensity is graphed in the fifth row in Figure 3.4, A-C. The intensity is highest in November through January for natural gas while the peak occurs during the summer for crude oil. I cannot identify a consistent pattern of the jump intensity for crude oil and heating oil futures.

The sixth panel plots the monthly stock levels as reported by EIA. The inventories for heating oil and natural gas exhibit inverted U-shaped curves, reaching their seasonal lows towards the end of the winter season, after which inventories begin to build up until the stock peaks around November. From the seventh plot in Figure 3.4, Panels A-C, I observe that the monthly mean of daily minimum temperatures clearly exhibits inverted U-shaped curves. Accordingly, during the winter season, when the stock levels of the underlying assets are low, sudden increases in demand due to extreme cold weather can often generate high volatility for natural gas and

heating oil. Section 3.4.5 provides a more detailed analysis on these effects.

In short, I have identified that the monthly volatility and the smooth path component exhibit seasonal variation during the year. They are relatively high in winter months and low in summer months for all three markets. I do not find an evident seasonal patterns in the jump component for any of the three markets during the sample period of this study.

3.4.3 EIA's Inventory Announcements, Intraday Realized Volatility and the Jump Component

Previous literature documents that jumps in asset prices are often associated with news events. Barndorff-Nielsen and Shephard (2006), for example, document that days with a jump in the DM/$ foreign exchange market can be linked with macroeconomic news. This is consistent with findings by Andersen et al. (2007), who report that macroeconomic announcements lead to large intraday price moves in a dataset from the foreign exchange markets. Johannes (2004) uses a parametric approach and shows that jumps in daily Treasury bill rates are associated with specific news announcements. Jiang et al. (2008) find that about 70 percent of the jumps in the U.S. Treasury markets occur at prescheduled macroeconomic announcements. This section examines the impact of the Energy Information Administration's (EIA) inventory announcements on intraday realized volatility and explores whether significant jumps can be linked to surprises in inventory announcements.

Before summarizing the full-sample results for the energy futures markets, I consider three specific days with large price movements. Figure 3.5, Panel C, plots cumulative intraday five-minute returns for natural gas on May 2, 2001. The $Z_{\text{bns},t}$ statistic is highly significant with a value of 6.61 for this day. The figure displays a

large jump around 2 PM, which coincides with the release time of the storage report published by EIA. On May 3, 2001, the Los Angeles Times printed the headline: *"Gas and crude oil prices dive as U.S. inventory levels surge"*, referring to industry reports from the previous day showing an unexpected increase in the stock-piles.[10] Figure 3.5, Panel A, graphs price movements for crude oil on Thursday, November 15, 2007. The $Z_{\text{bns},t}$ statistic identifies a jump for this day at the five percent significance level. In the afternoon, Bloomberg reported that the crude oil stock-piles rose unexpectedly due to the storage report released by EIA.[11,12] Bloomberg wrote that imports rose to their highest levels in the past few months allowing refineries to increase their runs and still increase their inventories. Figure 3.5, Panel B, plots the intraday behavior in heating oil futures prices on February 14, 2007. It shows that the price dropped considerably with a smooth sample path for about an hour following the announcement time at 10:30 AM. CNN reported later that day that while the crude oil supply declined more than expected, the heating oil reserve fell less than forecast. On the same day, analysts also reported that weather forecasts predict milder temperatures than expected.[13] Hence, the gradual price change may be due to the conflicting implications of the surprises in the stock-piles and the change in weather. The jump test statistic does not indicate a significant jump (the p-value is 0.9). In short, the three cases mentioned above illustrate the working of the jump statistic.

EIA releases weekly reports on the inventory status of heating oil, crude oil and natural gas. From May 2003 to December 2007, a smaller version of the inventory report for heating oil and crude oil with highlights and summarizing tables are released after 10:30 AM on Wednesdays; a full report is published after 1:00 PM on the same

[10] `http://articles.latimes.com/2001/may/03/business/fi-58758`

[11] `http://www.bloomberg.com`

[12] The petroleum report, which typically is released on Wednesdays, was postponed one day due to Veterans' day holiday.

[13] `http://money.cnn.com/2007/02/14/markets/oil_eia/`

day.[14] The EIA also compiles and releases a weekly natural gas storage report with estimates of natural gas in underground storage. From March 2000 to May 2002, the inventory report was released between 2:00 and 2:15 PM on Wednesdays. Since May 2002, EIA has released the report around 10:30 AM on Thursdays.[15]

I use one way analysis of variance models (that is, a regression model with dummy variables) to test whether the volatility increases at the announcement time on days with and without significant jump, respectively. I estimate the regression model,

$$\text{Vol}_{t,k} = \beta_0 + \sum_{i \neq l} \beta_i D_i + \epsilon_{t,k}. \tag{3.28}$$

The dependent variable, $\text{Vol}_{t,k}$, is the five-minute intraday volatility estimated by the absolute value of the difference between the logarithm of the closing and opening prices of the interval where $k = 1, \ldots, K$, denotes the intraday periods. The independent variables, D_i, are dummy variables denoting the five-minute intervals. The index l denotes a five-minute interval prior to the announcement time, which serves as the benchmark. It is implicitly assumed to be the equilibrium price absent any information about the content of the news release.

Table 3.8B, Panel B, presents the regression results of natural gas on Thursday mornings between 9:55 - 12:10 AM where the time period 10:05 - 10:10 serves as the benchmark.[16] The coefficients representing the release time of the storage report at 10:30 - 10:35 AM are positive and highly significant compared to the benchmark equilibrium price. The coefficients are 0.0227 (38.60) with significant jump days ($\alpha = 0.99$) compared to 0.0062 (24.07) for days without significant jumps, with t-statistics in parentheses. The volatility remains elevated for about thirty minutes

[14] The report was released off trading hours prior to May 2003.

[15] Before March 2, 2000, the report was released after the close of NYMEX on Wednesdays. Further discussion on the inventory reports is referred to EIA's website: `http://www.eia.doe`.

[16] Alternative benchmarks yield analogous results.

on days with jumps and sixty minutes on days without jumps. I also find that the coefficients prior to the announcement are insignificant and thus provide no evidence on leakage of inventory information related to the natural gas futures market. The results are consistent with the results of Linn and Zhu (2004), who examine the impact of the storage report on natural gas futures price volatility using transaction data from January 1999 to October 2002. They report a considerable increase in volatility at the time of the release and that the volatility remains elevated for up to thirty minutes following the announcement. They do not, however, consider jumps.

The regression results in Table 3.8B, Panel A, report the behavior in volatility when the release time of the natural gas storage report was at 2:00 PM on Wednesday afternoons from March 2000 to May 2002. The results are analogous to those for Thursdays. There is a large surge in volatility at the release time. Figure 3.6, Panel C, clearly demonstrates the time series behavior of intraday volatility of natural gas during the announcement periods.

Turning to the regression results in Table 3.8A, Panels A and B, I examine the intraday regression results for crude oil and heating oil markets on Wednesdays at the release time of the petroleum report at 10:30 AM. I observe that the crude oil coefficient is 0.0116 (23.62) for days with jumps and 0.0040 (24.73) for days without jumps. For heating oil, the coefficient is 0.0139 (17.13) for days with jumps and 0.0042 (15.06) for days without jumps; the parentheses report corresponding t-statistics. It is interesting to observe that volatility returns to prelevel volatilities approximately within ten minutes in both crude oil and heating oil markets on announcement days with jumps, while the volatility remains elevated beyond sixty minutes on announcement days without jumps. Panels A and B in Figure 3.6 plot the time series behavior of realized volatility during announcement days with jump and without jump components.

Furthermore, I examine whether the arrival rate of jumps is higher during the inventory event periods by determining the intraday timing of jumps. The timings are established by a procedure similar to one used by Jiang et al. (2008).[17] During sample period June 2002 to December 2007, I identify that there are 72 jumps on Thursdays in the natural gas market and 60 of these occur between 10:25-10:40 AM; there are a total of 167 jumps during this sample period. There are 45 jumps on Wednesdays in the heating oil market and 35 occur at the inventory announcement; there are 121 jumps for the whole sample period. Similarly, there are 36 jumps on Wednesdays in the crude oil market and 27 of these occur the announcement; there are a total of 91 during the whole sample period. The results strongly suggest that there are clusters in jumps during the inventory announcement dates in all three energy futures markets.

In summary, I provide three new interesting empirical findings. First, the intraday volatility increases during the period immediately following the inventory announcement; furthermore, the volatility is higher during the announcement period on days with a significant jump component than on days without a jump. Second, the volatility returns to previous levels faster on days with a jump at the announcement period compared to days without a jump. Third, jumps often occur during the inventory announcement days in energy futures markets.

[17] Let r_{tj} denote the jth intraday return on the tth day. (1) Find the largest (in magnitude) return, r^*_{tj}, and record that return as a jump. (2) To test whether there are additional jumps on the same day, set r^*_{tj} to the median value of the intraday returns for that day since the median is robust to jumps. (3) Recompute the test statistic for the tth day. If the test still rejects the null hypothesis, return to step (1) to determine the time of an additional jump. Otherwise, all jumps have been identified for the tth day.

3.4.4 Modeling Realized Variation with Jump Component

A number of studies in different markets have found that the volatility process is a member of the long-memory time series by employing ARFIMA models or stochastic volatility models. Corsi (2004) proposes the HAR-RV model to represent the time series behavior of the realized volatility process. The HAR-RV model is motivated by the heterogeneous market hypothesis suggested by Müller et al. (1997), which recognizes the heterogeneity of the traders due to differences in the endowment, degree of information, temporal horizons, risk profile, and so on. These differences lead trader groups to settle for different prices and decide to trade at different market situations; hence, they create volatility. Corsi (2004) incorporates realized volatility over different time horizons in his model such as daily, weekly and monthly volatility components to capture the volatility due to heterogeneous beliefs among trader groups. While the HAR-RV model does not explicitly incorporate the long-memory property, Corsi (2004) demonstrates that his model can reproduce the memory persistence observed in many markets. Andersen et al. (2007) extend the HAR-RV model by including a jump component as an explanatory variable, and demonstrate that the augmented model (HAR-RV-J model) significantly improves the performance of the HAR-RV model for modeling and forecasting volatility in DM/$, S&P 500 futures and T-bond futures.

Following the Andersen et al. (2007) approach, I specify the HAR-RV-J model for three energy markets as follows. First, I define the multi-period normalized realized variation as,

$$\mathrm{RV}_{t,t+h} = h^{-1} \sum_{i=1}^{h} \mathrm{RV}_{t+i}. \tag{3.29}$$

I consider $h = 1$, $h = 5$ and $h = 22$, which correspond to daily, weekly and monthly

lags. For these time horizons, the HAR-RV model is written as,

$$\begin{aligned}\mathrm{RV}_{t,t+h} =&\beta_0 + \beta_D \mathrm{RV}_t + \beta_W \mathrm{RV}_{t-5,t} + \beta_M \mathrm{RV}_{t-22,t} + \\ &\sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{252}\right) + \xi_q \cos\left(\frac{2\pi q T}{252}\right)\right) + \epsilon_{t,t+h}.\end{aligned} \tag{3.30}$$

I add sine and cosine terms to capture the seasonal variation. Andersen et al. (2007) extend the model by adding a variable for the jump component, J_t, that is,

$$\begin{aligned}\mathrm{RV}_{t,t+h} =&\beta_0 + \beta_D \mathrm{RV}_t + \beta_W \mathrm{RV}_{t-5,t} + \beta_M \mathrm{RV}_{t-22,t} + \beta_J J_t + \\ &\sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{252}\right) + \xi_q \cos\left(\frac{2\pi q T}{252}\right)\right) + \epsilon_{t,t+h}.\end{aligned} \tag{3.31}$$

Henceforth, this model is referred to as the HAR-RV-J model.

Besides estimating these models in the realized variance form, I also consider two nonlinear transformations. A standard deviation (volatility) form of the HAR-RV-J model is defined as,

$$\begin{aligned}\mathrm{RV}^{1/2}_{t,t+h} =&\beta_0 + \beta_D \mathrm{RV}^{1/2}_t + \beta_W \mathrm{RV}^{1/2}_{t-5,t} + \beta_M \mathrm{RV}^{1/2}_{t-22,t} + \beta_J J^{1/2}_t + \\ &\sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{252}\right) + \xi_q \cos\left(\frac{2\pi q T}{252}\right)\right) + \epsilon_{t,t+h},\end{aligned} \tag{3.32}$$

and a logarithmic transformation is written as,

$$\begin{aligned}\log(\mathrm{RV}_{t,t+h}) =&\beta_0 + \beta_D \log(\mathrm{RV}_t) + \beta_W \log(\mathrm{RV}_{t-5,t}) + \beta_M \log(\mathrm{RV}_{t-22,t}) + \\ &\beta_J \log(J_t + 1) + \sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{252}\right) + \xi_q \cos\left(\frac{2\pi q T}{252}\right)\right) + \epsilon_{t,t+h}.\end{aligned} \tag{3.33}$$

The nonlinear forms of the HAR-RV model are defined analogously.[18]

The ordinary least squares procedure is used to estimate the parameters of the models, and Newey and West (1987) is applied to estimate consistent standard errors of the parameters. The regression results are reported in Tables 3.9A and 3.9B. The major results are summarized as follows:

1. The coefficients of the volatility form of the HAR-RV model are consistently positive and in most cases highly significant, see Table 3.9A. The daily lag coefficient, β_D, decreases as h increases in all markets but remains significant. For natural gas futures, the values drop from 0.212 (5.17) to 0.167 (5.59) to 0.087 (7.08) for h equal to 1, 5 and 22. The values in parentheses are t-statistics. This relationship also holds for the weekly lag, β_W.

2. Table 3.9B tabulates the coefficients for the HAR-RV-J model where the jump component is included as an explanatory variable. All coefficients for the total variation remain positive and generally highly significant. The estimates of β_J are consistently negative and significant for all values of h for all models and markets. Thus, the impact of the lagged realized variation is reduced by the jump component. For natural gas, the coefficient β_J in the standard deviation form of the model is -0.097 (-2.55) for the daily horizon, and -0.073 (-2.56) for the weekly horizon, and -0.040 (-2.20) for the monthly horizon. These results suggest that the jump component in the price process produces transient surges in the volatility with a strong reversal in the volatility on the subsequent day of a jump. The coefficient remains significant at the weekly horizon but only weakly so over the monthly period, which is consistent with the low persistency in the jump component.

[18]The purposes of log transformation are: (1) the distribution of logarithmic realized volatility is nearly normal, and (2) estimated realized volatility transformed back from logarithmic realized volatility is always positive.

3. The adjusted R^2 values remain nearly the same or are slightly improved for HAR-RV-J compared to HAR-RV for all three markets. The extension to include the Fourier terms for seasonality improves the adjusted R^2, particularly for $h = 22$. The adjusted R^2 values increases by 0.05 for natural gas, 0.03 for heating oil, and 0.04 for crude oil.

4. The results for the variance and logarithmic forms are generally consistent with the standard deviation form. The logarithmic transformation has in most instances the highest adjusted R^2.

In short, I provide additional empirical evidence to support the improved performance of the HAR-RV-J model over the HAR-RV model in energy futures markets. The impact of the jump component on current realized volatility is transitory and there is a strong reversal in subsequent realized volatility.

3.4.5 HAR-RV Model and Short-Run Supply and Demand Factors

The financial press often refer to extreme weather conditions and declining inventory levels as factors driving short-run increases in prices and volatility in energy markets. Ates and Wang (2007) provide empirical evidence supporting these claims in the natural gas and heating oil markets; particularly, they find that extreme cold weather surprises and inventory surprises account for variation in the interest-rate adjusted basis, and spot and futures price changes. In the following, I extend the HAR-RV model with these variables to assess their explanatory power in this context.

Since natural gas and heating oil are primarily used for heating purposes, the demand in these commodities is highest during the colder winter months and lower in the summer. The production, however, remains fairly constant, which leads to increasing

inventories in the late summer and into the fall while the stock declines throughout the winter and early spring. Consequently, changes in the daily low temperature can serve as a proxy for changes in the short-run demand in these commodities, and a negative inventory surprise provides a proxy for quantifying the inelastic short-run supply during the low inventory season.

From the theory of storage (Working (1949), Brennan (1958), and Telser (1958)) the difference between futures and cash is equal to the cost of storage minus the marginal convenience yield of holding the commodity. The yield is negatively related to inventory levels and thus may cause the basis to become negative during periods of low inventories when there are sudden increases in demand due to cold weather. The basis remains positive, however, during warmer seasons with high inventory levels since sudden increases in demand of the commodities can be absorbed. As a result, the basis can be applied as a proxy to explain changes in prices due to changes in the short-run supply and demand. Ng and Pirrong (1994), for example, use the interest-rate adjusted basis as a proxy for the demand and supply factors in their empirical model. Since the spot markets in these commodities generally are less liquid than the futures markets, the difference between the nearby and first deferred contracts (that is, the spread) can be used as an alternative to the basis since the nearby contract trades very close to the spot price. Fama and French (1988), for example, use negative spread periods as a proxy for periods of low inventory levels. Hence, I compute the negative interest-rate adjusted spread and examine whether it captures the short-run supply and demand effects due to inventory and cold weather surprises.

I compute the interest-rate adjusted spread by,

$$\text{spread}_t = F_{2,t,T_2} - F_{1,t,T_1} \mathrm{e}^{r(T_2 - T_1)/365}, \tag{3.34}$$

where F_{1,t,T_1} is the nearby futures price for a contract that expires at time T_1, F_{2,t,T_1} is the first deferred contract that expires at time T_2, and r is the three-month Treasury-bill rate. I obtain historical daily interest rate series from the Federal Reserve database.[19] Since the negative spread is used as a proxy for low inventory periods, I create the variable spread_t^- by setting positive spreads to zero.

I construct the extreme low weather surprise variable (elws_t) to measure the nonlinear relationship between the commodities and the extreme cold temperatures by first estimating a model for daily low temperatures proposed by Campbell and Diebold (2005). They model the conditional mean by a trend and seasonal and cyclical components by approximating the seasonal volatility component using a Fourier series and the cyclical volatility component using a GARCH process, see the full model specification in Appendix A.II on page 149.[20] Second, I compute the empirical cumulative distribution of the residuals and find the tenth percentile, which I use as the threshold to define the extreme cold temperature surprises. Third, I set the values above the tenth percentile to zero.

I estimate a regression model to determine the inventory surprise variable (invs_t) which is given by the residuals. The regression model is defined as,

$$\text{Inv}_t = \beta_0 + \sum_{q=1}^{Q_I} \left(\delta_q \sin\left(\frac{2\pi q T}{365}\right) + \xi_q \cos\left(\frac{2\pi q T}{365}\right) \right) + \sum_{l=1}^{L_I} \beta_{t-l} \text{Inv}_{t-l} + \epsilon_{I,t}, \quad (3.36)$$

where Inv_t is the stock level reported by EIA on day t and Inv_{t-l} denotes lagged

[19] `http://www.federalreserve.gov/econresdata/releases/statisticsdata.htm`

[20] I obtain equivalent results based on a dynamic regression model,

$$\text{Tmin}_t = \beta_0 + \sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{365}\right) + \xi_q \cos\left(\frac{2\pi q T}{365}\right) \right) + \sum_{l=1}^{L} \beta_{t-l} \text{Tmin}_{t-l} + \epsilon_{W,t}, \quad (3.35)$$

where Tmin is the daily low temperature.

values.[21]

The extended version of the HAR-RV-J model is written as,

$$\begin{aligned}\mathrm{RV}_{t,t+h} =& \beta_0 + \beta_D \mathrm{RV}_t + \beta_W \mathrm{RV}_{t-5,t} + \beta_M \mathrm{RV}_{t-22,t} + \beta_J J_t + \\ & \sum_{r=0}^{R} \beta_{elws,r} |\mathrm{elws}_{t-r+1}| + \beta_{invs,r} |\mathrm{invs}_{t+1}| + \beta_{spread^-} |\mathrm{spread}_t^-| + \\ & \sum_{q=1}^{Q} \left(\delta_q \sin\left(\frac{2\pi q T}{252}\right) + \xi_q \cos\left(\frac{2\pi q T}{252}\right) \right) + \epsilon_{t,t+h}.\end{aligned} \tag{3.37}$$

where elws_{t+1} is the extreme low temperature surprise; invs_{t+1} is the inventory surprise; and, spread_t^- denotes the interest-rate adjusted negative spread.

In the following, I refer to model (1) through (5) where (1) denotes to the HAR-RV model; (2) is the HAR-RV-J model; (3) adds the negative interest-rate adjusted spread to HAR-RV-J; (4) adds the weather and inventory surprise variables to the HAR-RV-J model; and (5) is the full model defined by equation (3.37). Table 3.10 reports regression estimates of the five models for $Q = 3$ and $R = 2$. The estimates for the weather and inventory surprises are scaled by 10^5. The ordinary least squares procedure is used to estimate the parameters of the models, and Newey and West (1987) is used to estimate consistent standard errors of the parameters to account for autoregressive and heteroskedastic errors. The Fourier terms are included in all models but not presented to conserve space.

The regression results on the monthly variations reported in Table 3.7 show that the inventory increases throughout the summer and fall, and begins to decline in December in the natural gas and heating oil markets. Furthermore, the interest-rate adjusted spread remains negative over the same period as the inventory is declining

[21]Since the HAR-RV regression is estimated on a daily frequency and EIA releases its storage reports on a weekly basis, I obtain daily estimates of the inventory by spline interpolation, which is likely to smooth out the series.

in both markets. These patterns are also apparent from Figure 3.4.

The following results are obtained from the estimation of the five models.

1. I discuss models (1) and (2) in Section 3.4.4 but include them as a reference. The third model adds the negative interest-rate adjusted spread. The results for natural gas, Panel C, show that the absolute value of the negative spread is positive and highly significant. Together with the seasonal variations in the spread, these findings are consistent with that volatility is higher during the colder season when the inventories are low. These conclusions also hold for the heating oil market. The seasonality in the basis is not as apparent in the crude oil market, which is consistent with that this commodity is used more evenly throughout the year.

2. The fourth model adds the low temperature and inventory surprise variables to the HAR-RV-J model. The estimate of the absolute value of the current weather surprise, ewls_{t+1}, is positive for all markets and thus agrees with the hypothesis that low temperature surprises lead to an increase in volatility. It is weakly significant in the natural gas and heating oil markets but insignificant in the crude oil market. The inventory surprise is positive, which supports the hypothesis that inventory surprises generate volatility; the estimates are insignificant, however, in all markets.

3. Interestingly, the estimate of the fifth model suggests that the lagged negative spread may be a suitable proxy for the negative inventory periods since the significance of the weather and inventory surprise at time $t + 1$ drops while the spread estimate remains highly significant in all markets.

4. Lastly, the jump component is negative and significant in the HAR-RV-J model (2) and remains so in all models. Notice, however, that the estimate increases

somewhat and loses in significance as the negative interest-rate adjusted spread is added to the model. The addition of the weather and inventory surprise variables, however, leaves the jump estimate nearly unchanged. This suggest that the short-run supply and demand factors as proxied by the spread explain some of the jump component during the colder winter season.

3.5 Summary and Conclusions

This paper applies a nonparametric method based on realized and bipower variations calculated from intraday high-frequency data to identify jumps in daily futures prices in three energy futures contracts. The futures contracts are crude oil, heating oil and natural gas, which are traded on the NYMEX. The sample period of the high-frequency intraday data spans from January 1990 to January 2008. Alternative methods such as staggered returns and optimal sampling frequency methods are used to remove the effects of microstructure noise, which biases the tests against detecting jumps. I obtain several interesting empirical results:

First, for the whole sample period, I find that the means of annualized volatility for natural gas futures, crude oil futures and heating oil futures are 39.4, 26.0 and 26.5 percent, respectively. Thus, natural gas is the most volatile among these three price series. There are upward trends in volatility for all series during the sample period; for natural gas the increase is primarily due to the jump component while the smooth component dominates the increase in the crude oil and heating oil markets. There are significant jumps (price spikes) in the price series. The number of significant jump days per year ranges from 5 to 34 for natural gas, 5 to 28 for heating oil and 4 to 20 days for crude oil. I find that days with large volatility are often associated with large jumps. The relative contribution of the jump component to the total variance

on days with jumps ranges from 23% to 87% for natural gas futures, 23% to 74% for heating oil futures and 23% to 64% for crude oil futures. In addition, the results suggest that jump dynamics are much less persistent than continuous sample path dynamics.

Second, I document that the total realized volatility and smooth sample component for natural gas and heating oil are higher in the winter months than during the summer months. These results are consistent with the general hypothesis that when short run demand for natural gas and heating oil are suddenly shifted higher due to extreme cold weather during the winter, the short run supply is inelastic because of low inventories at this time of the year. These two factors are the ones largely responsible for generating volatility in the winter months.

Third, in an intraday analysis, I document that the volatility is higher during inventory news announcement periods and that many jumps are associated with these announcement dates. Furthermore, it is interesting to observe that for all markets, the volatility returns to pre-announcement levels faster when there is a jump in the futures price changes than when there is no jump. The volatility remains elevated for about thirty minutes or shorter on days with a jump at the announcement and longer otherwise.

Fourth, I find that including the jump component in the HAR-RV model improves the performance of the realized volatility forecasting model. The coefficient of the jump component attains the largest value at the daily lag and decreases for corresponding weekly and monthly regression estimates. Furthermore, most of the coefficients of jumps are negative and significant. The above two results indicate that the jump component in the price process produces transitory surges in volatility and that there is a strong reversal in the volatility on the subsequent days of a jump.

Fifth, cold weather and inventory surprises lead to an increase in volatility in natural gas and heating oil markets. Furthermore, the lagged interested-rate adjusted spread may be a suitable proxy for the negative inventory periods since the significance of the weather and inventory variables drops while the spread remains highly significant when including all three variables. The spread also reduces the significance of the jump component slightly.

The empirical results have several important implications for market participants in energy futures markets. For option traders, option pricing models with jumps are preferred over the Black option model in these three energy futures contracts during the winter months. Market participants may prefer to employ an optimal hedging ratio with jumps to hedge their exposure to energy price risk. Finally, market risk managers should be aware that the shapes of the return distributions of energy futures prices will change over time due to the presence of significant jumps in energy futures prices. Feng and Brooks (2002) and Sullivan et al. (2002) for example discuss risk management in the natural gas futures market but do not address jumps.

Bibliography

Bibliography

Aït-Sahalia, Y., Mykland, P. A., and Zhang, L. (2006). Ultra high frequency volatility estimation with dependent microstructure noise. NBER Working Paper No. W11380.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility. *Review of Economics and Statistics*, 89(4):701–720.

Ates, A. and Wang, G. H. K. (2007). Price dynamics in energy spot and futures markets: The role of inventory and weather. Presented at Financial Management Association Annual Meeting.

Bandi, F. M. and Russell, J. R. (2006). Separating microstructure noise from volatility. *Journal of Financial Economics*, 79:655–692.

Barndorff-Nielsen, O. E., Graversen, S. E., Jacod, J., Podolskij, M., and Shephard, N. (2006). *From Stochastic Analysis to Mathematical Finance, Festschrift for Albert Shiryaev*. Springer Verlag.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–48.

Barndorff-Nielsen, O. E. and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1):1–30.

Brennan, M. J. (1958). The supply of storage. *American Economic Review*, 48(1):50–72.

Campbell, S. D. and Diebold, F. X. (2005). Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, 100(469):6–16.

Chan, W. H. and Maheu, J. M. (2002). Conditional jump dynamics in stock market returns. *Journal of Business and Economic Statistics*, 20:377–389.

Corsi, F. (2004). A simple long memory model of realized volatility. SSRN Manuscript.

Fama, E. F. and French, K. R. (1988). Business cycles and the behavior of metals prices. *Journal of Finance*, 43(5):1075–1093.

Feng, W. and Brooks, R. (2002). Risk management of natural gas exposure. Working Paper No. 01-10-02.

Huang, X. and Tauchen, G. (2005). The relative contribution of jumps to total price variation. *Journal of Financial Econometrics*, 3:456–499.

Jacod, J. and Shiryaev, A. N. (1987). *Limit Theorems for Stochastic Processes*. Springer Verlag.

Jiang, G. J., Lo, I., and Verdelhan, A. (2008). Information shocks and bond price jumps: Evidence from the U.S. treasury market. Working Paper.

Jiang, G. J. and Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise - A "Swap Variance" approach. *Journal of Econometrics*, 144(2):352–370.

Johannes, M. (2004). The statistical and economic role of jumps in continuous-time interest rate models. *Journal of Finance*, 59:227–260.

Li, Y. and Mykland, P. A. (2007). Are volatility estimators robust with respect to modeling assumptions? *Bernoulli*, 13(3):601–622.

Linn, S. C. and Zhu, Z. (2004). Natural gas prices and the gas storage report: Public news and volatility in energy futures markets. *Journal of Futures Markets*, 24(3):283–313.

Maheu, J. M. and McCurdy, T. H. (2004). News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance*, 59:755–793.

Mu, X. (2007). Weather, storage, and natural gas price dynamics: Fundamentals and volatility. *Energy Economics*, 29:46–63.

Müller, U. A., Dacorogna, M. M., Davè, R. D., Olsen, R. B., Puctet, O. V., and von Weizsäcker, J. (1997). Volatilities of different time resolutions - Analyzing the dynamics of market components. *Journal of Empirical Finance*, 4:213–239.

Neuberger, A. (1994). The log contract - A new instrument to hedge volatility. *Journal of Portfolio Management*, 20(2):74–80.

Newey, W. and West, K. (1987). A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708.

Ng, V. K. and Pirrong, S. C. (1994). Fundamentals and volatility: Storage, spreads and dynamics of metal prices. *Journal of Business*, 67:203–230.

Pindyck, R. S. (2004). Volatility in natural gas and oil markets. *Journal of Energy and Development*, 30(1):1–19.

Sullivan, J. H., Brooks, R., and Stoumbos, Z. G. (2002). Structured assessment of risk systems and value at risk (VaR). Working Paper No. 02-06-01.

Telser, L. G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economics*, 66:233–255.

Wang, T., Wu, J., and Yang, J. (2008). Realized volatility and correlation in energy futures markets. *Journal of Futures Markets*, 28(10):993–1011.

Working, H. (1949). The theory of price storage. *American Economic Review*, 39:1254–1262.

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*, 100:1394–1411.

3.6 Tables

Table 3.1: Daily summary statistics for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C) for realized variance, RV_t (equation (3.2)) and jump component, J_t (equation (3.21)). Statistics are also computed in standard deviation form, $\mathrm{RV}_t^{1/2}$ ($J_t^{1/2}$), and logarithmic form, $\log(\mathrm{RV}_t)$ ($\log(J_t+1)$). ADF denotes the augmented Dickey-Fuller statistic. Min and Max are minimum and maximum daily values. JB is the Jarque-Bera test statistic for normality. LB_{10} denotes the Ljung-Box tenth-order serial correlation test statistic. Kurtosis denotes excess kurtosis. The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

	RV_t	$\mathrm{RV}_t^{1/2}$	$\log(\mathrm{RV}_t)$	J_t	$\mathrm{J}_t^{1/2}$	$\log(\mathrm{J}_t+1)$
Panel A: Crude Oil						
ADF[1]	−16.04	−6.75	−5.67	−33.99	−19.34	−33.96
Mean	0.0003	0.0164	−8.3774	0.0000	0.0033	0.0000
Std Dev	0.0007	0.0072	0.7718	0.0003	0.0045	0.0003
Skewness	44.30	4.87	0.04	59.82	7.34	59.71
Kurtosis	2534.14	88.62	1.06	3835.89	175.34	3825.63
Min	0.0000	0.0030	−11.6462	0.0000	0.0000	0.0000
Max	0.0381	0.1953	−3.2664	0.0188	0.1370	0.0186
JB	$1.21E+09$	$1.49E+06$	$2.13E+02$	$2.77E+09$	$5.82E+06$	$2.76E+09$
LB_{10}	968	10926	16947	91	283	93
Panel B: Heating Oil						
ADF[1]	−15.35	−6.80	−4.95	−27.02	−23.85	−27.00
Mean	0.0003	0.0167	−8.3128	0.0000	0.0038	0.0000
Std Dev	0.0004	0.0064	0.6897	0.0002	0.0044	0.0002
Skewness	27.81	3.24	0.17	50.19	3.83	50.07
Kurtosis	1286.59	39.97	0.83	2998.58	56.75	2988.21
Min	0.0000	0.0034	−11.3906	0.0000	0.0000	0.0000
Max	0.0207	0.1439	−3.8779	0.0103	0.1017	0.0103
JB	$3.08E+08$	$3.04E+05$	$1.50E+02$	$1.67E+09$	$6.08E+05$	$1.66E+09$
LB_{10}	1873	9263	13033	137	193	138
Panel C: Natural Gas						
ADF[1]	−10.91	−8.63	−7.53	−21.34	−13.62	−21.33
Mean	0.0007	0.0248	−7.5419	0.0001	0.0061	0.0001
Std Dev	0.0008	0.0105	0.7556	0.0003	0.0075	0.0003
Skewness	6.73	2.16	0.22	11.83	2.74	11.80
Kurtosis	81.85	10.01	0.66	207.69	14.62	206.60
Min	0.0000	0.0038	−11.1209	0.0000	0.0000	0.0000
Max	0.0165	0.1286	−4.1015	0.0073	0.0852	0.0072
JB	$1.06E+06$	$1.82E+04$	$9.60E+01$	$6.70E+06$	$3.74E+04$	$6.63E+06$
LB_{10}	2912	6184	8503	194	231	194

Note: 1. H_0: Unit root. The lag orders are determined by Schwartz criterion. Only intercepts are included in the level series. The critical value for the ADF test for the 1% (5%) significance level is −3.4393 (−2.8654).

Table 3.2: Yearly estimates for crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). No. Days denotes the number of trading days, No. Jumps denotes the number of days with jumps, and Prop denotes the proportion of days with jumps. Min, Mean, Median and Max are daily statistics of the relative contribution of the jump component to the total realized variance (equation (3.13)) computed for days with a significant jump component for $\alpha = 0.99$.

				RJ on Jump Days (%)			
	No. Days	No. Jumps	Prop	Min	Mean	Median	Max
Panel A: Crude Oil							
1993	250	13	0.052	22.61	30.01	28.92	40.14
1994	250	17	0.068	26.05	35.28	33.09	45.92
1995	250	10	0.040	22.58	37.07	34.74	62.32
1996	251	13	0.052	25.63	34.40	31.33	62.08
1997	251	11	0.044	23.75	33.78	31.55	47.48
1998	251	8	0.032	25.26	29.78	29.47	38.68
1999	250	9	0.036	22.92	29.94	28.61	39.73
2000	249	7	0.028	23.32	30.93	25.65	44.61
2001	246	4	0.016	26.84	31.49	31.18	36.76
2002	250	5	0.020	26.48	32.34	33.55	38.89
2003	250	15	0.060	30.52	41.21	37.17	58.77
2004	249	9	0.036	28.51	36.48	36.02	47.22
2005	251	12	0.048	24.97	32.97	30.17	44.20
2006	250	6	0.024	28.26	38.03	36.23	48.73
2007	258	20	0.078	22.79	32.47	29.40	64.18
Panel B: Heating Oil							
1993	236	14	0.059	25.39	34.83	30.48	57.86
1994	249	18	0.072	25.23	36.20	32.16	58.79
1995	237	12	0.051	24.37	32.25	31.96	43.12
1996	251	10	0.040	27.91	37.74	34.30	56.65
1997	248	16	0.065	24.16	34.77	30.25	52.06
1998	251	17	0.068	23.06	33.63	31.24	63.76
1999	248	7	0.028	26.41	35.31	37.65	43.31
2000	247	11	0.045	23.09	32.29	35.23	42.80
2001	246	14	0.057	23.06	36.40	36.89	53.63
2002	249	5	0.020	27.58	31.88	30.15	38.25
2003	250	18	0.072	27.02	39.07	37.00	61.52
2004	249	17	0.068	25.90	39.15	35.41	64.98
2005	251	20	0.080	28.09	42.37	41.00	73.56
2006	251	19	0.076	26.50	37.74	37.67	58.43
2007	258	28	0.109	22.83	33.89	31.30	65.06

Table 3.2 continue

	No. Days	No. Jumps	Prop	RJ on Jump Days (%) Min	Mean	Median	Max
Panel C: Natural Gas							
1993	250	5	0.020	31.72	46.17	46.58	60.52
1994	248	11	0.044	25.18	34.49	34.53	54.62
1995	250	8	0.032	26.42	39.34	33.76	75.23
1996	248	15	0.060	26.62	37.22	36.43	61.08
1997	213	8	0.038	28.84	38.65	33.20	73.60
1998	240	11	0.046	26.47	42.90	37.51	78.50
1999	232	12	0.052	25.32	33.53	32.12	55.07
2000	235	17	0.072	28.23	48.46	48.03	87.47
2001	236	34	0.144	23.56	45.76	44.06	85.92
2002	245	17	0.069	28.12	46.05	43.43	72.97
2003	249	25	0.100	25.89	38.51	34.75	77.15
2004	249	26	0.104	26.45	42.05	37.26	69.19
2005	251	19	0.076	26.50	42.05	40.37	68.96
2006	250	23	0.092	25.47	41.88	42.09	62.96
2007	258	14	0.054	23.39	33.81	32.18	52.13

Table 3.3: Regression analysis is used to test for trends in daily time series of realized volatility, smooth and jump components, and relative contribution in crude oil, heating oil and natural gas. OLS is used to estimate the regression model,

$$V_t = \beta_0 + \beta_1 \text{Trend}_t + \epsilon_t,$$

where V_t denotes the realized variations and ratio, respectively, and Trend_t denotes the trend. The significance level α to test for jumps is 0.99. The dependent variables are scaled by 10^6. t-statistics are reported in parentheses.

	$\text{RV}_t^{1/2}$	$\text{C}_{t,\alpha}^{1/2}$	$\text{J}_{t,\alpha}^{1/2}$	$\text{RJ}_t^{1/2}$
		Panel A: Crude Oil		
Intercept	16356.0	16182.0	521.3	38141.0
	(152.79)	(159.99)	(10.56)	(18.18)
Trend	0.383	0.395	−0.036	−6.207
	(4.61)	(5.04)	(−0.93)	(−3.82)
R^2 Adj	0.00	0.01	0.00	0.00
F Stat	21.26	25.40	0.87	14.56
		Panel B: Heating Oil		
Intercept	16646.0	16397.0	720.9	56079.0
	(173.84)	(179.68)	(14.14)	(25.13)
Trend	0.485	0.467	0.048	−3.057
	(6.44)	(6.52)	(1.20)	(−1.74)
R^2 Adj	0.01	0.01	0.00	0.00
F Stat	41.51	42.47	1.43	3.04
		Panel C: Natural Gas		
Intercept	24830.0	24288.0	1370.0	67217.0
	(143.45)	(148.02)	(13.82)	(24.98)
Trend	0.326	0.160	0.418	10.307
	(1.98)	(1.03)	(4.43)	(4.03)
R^2 Adj	0.00	0.00	0.01	0.00
F Stat	3.91	1.05	19.65	16.20

Table 3.4: Regression analysis is used to test for trends in the monthly intensity of jumps and volatility due to significant jumps, $J^{1/2}_{t,\alpha=0.99}$, in crude oil, heating oil and natural gas. OLS is used to estimate the regression model,

$$V_t = \beta_0 + \beta_1 \text{Trend}_t + \epsilon_t,$$

where V_t denotes the intensity and volatility, respectively, and Trend_t denotes the trend. The intensity is estimated by the number of jumps over the total number of trading days per month. The volatility is estimated by the mean of the daily jump volatility per month. t-statistics are reported in parentheses.

	Jump intensity	Jump volatility
	Panel A: Crude Oil	
Intercept	0.04104	0.00050
	(15.30)	(9.48)
Trend	−0.00006	−0.00000
	(−1.31)	(−1.27)
R^2 Adj	0.01	0.01
F Stat	1.72	1.62
	Panel B: Heating Oil	
Intercept	0.05668	0.00069
	(16.90)	(12.94)
Trend	0.00008	0.00000
	(1.54)	(1.09)
R^2 Adj	0.01	0.01
F Stat	2.37	1.18
	Panel C: Natural Gas	
Intercept	0.06181	0.00129
	(16.50)	(12.49)
Trend	0.00036	0.00001
	(5.05)	(4.41)
R^2 Adj	0.13	0.10
F Stat	25.55	19.42

Table 3.5: Summary statistics for significant jumps, $J_{t,\alpha}^{1/2}$ (equation (3.22)), for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). Jumps are identified by the Z_t statistic (equation (3.12)) for α equal to $0.500, 0.950, 0.990, 0.999, 0.9999$. No. Jumps denotes the number of jumps in the complete sample. Proportion denotes the ratio of days with a jump. The sample consists of $4,510$ trading days for crude oil, $4,449$ for heating oil, and $3,676$ for natural gas. Mean and Std Dev are the mean and standard deviation of the daily jump component, $\mathrm{J}_{t,\alpha}^{1/2}$. $\mathrm{LB}_{10}, \mathrm{J}_{t,\alpha}^{1/2}$ denotes the Ljung-Box tenth-order autocorrelation test statistic. The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

α	0.500	0.950	0.990	0.999	0.9999
			Panel A: Crude Oil		
No. Jumps	1993	440	197	80	37
Proportion	0.44	0.10	0.04	0.02	0.01
Mean ($J_{t,\alpha}^{1/2}$)	0.0061	0.0100	0.0121	0.0152	0.0144
Std Dev	0.0051	0.0082	0.0109	0.0159	0.0079
$\mathrm{LB}_{10}, J_{t,\alpha}^{1/2}$	75	71	59	58	0
			Panel B: Heating Oil		
No. Jumps	2161	502	272	115	66
Proportion	0.49	0.11	0.06	0.03	0.01
Mean ($J_{t,\alpha}^{1/2}$)	0.0063	0.0103	0.0121	0.0144	0.0157
Std Dev	0.0046	0.0064	0.0077	0.0096	0.0116
$\mathrm{LB}_{10}, J_{t,\alpha}^{1/2}$	124	101	105	41	0
			Panel C: Natural Gas		
No. Jumps	1816	470	246	121	75
Proportion	0.49	0.13	0.07	0.03	0.02
Mean ($J_{t,\alpha}^{1/2}$)	0.0101	0.0171	0.0207	0.0263	0.0297
Std Dev	0.0079	0.0103	0.0120	0.0137	0.0135
$\mathrm{LB}_{10}, J_{t,\alpha}^{1/2}$	179	241	216	222	38

Table 3.6: Summary statistics for jump returns for days with significant jumps ($\alpha = 0.99$) for crude oil, heating oil and natural gas. N denotes the number of jumps. The largest (in magnitude) five-minute intraday return per day with a significant jump is identified as the jump return. The statistics are computed for positive and negative returns, respectively.

Contract	N	Mean	Median	StdDev	Max	Min
			Positive Jumps			
Crude Oil	89	0.012	0.009	0.015	0.137	0.003
Heating Oil	101	0.012	0.010	0.010	0.102	0.005
Natural Gas	89	0.021	0.016	0.014	0.083	0.007
			Negative Jumps			
Crude Oil	107	0.012	0.011	0.005	0.031	0.003
Heating Oil	165	0.012	0.011	0.005	0.033	0.004
Natural Gas	153	0.020	0.018	0.011	0.067	0.006

Table 3.7: Regression analysis is used to analyze seasonal (monthly) variations. The dependent variables are realized volatility, $\mathrm{RV}_t^{1/2}$ (equation (3.2)), volatility due to the smooth component, $C_{t,\alpha=0.99}^{1/2}$ (equation (3.23)), volatility due to the significant jump component, $J_{t,\alpha=0.99}^{1/2}$ (equation (3.22)), relative contribution of the jump component, RJ_t (equation (3.13)), inventory level, temperature, and spread (equation (3.34)). OLS is used to estimate the regression model,

$$V_t = \beta_0 + \sum_{i=1, i\neq 7}^{12} \beta_i D_i + \epsilon_t,$$

where V_t denotes the dependent variable, and D_i is a dummy variable set to one for the ith month and zero otherwise. The intercept denotes the month of July which is used as the basis for comparison.

	$\mathrm{RV}_t^{1/2}$	$\mathrm{C}_t^{1/2}$	$\mathrm{J}_t^{1/2}$	RJ_t	Stock	Temp	Spread
			Panel A: Crude Oil				
Intercept	0.0146	0.0144	0.0005	0.0367	321164.0	68.3	−0.095
	(39.66)	(41.72)	(3.19)	(5.03)	(62.04)	(115.18)	(−2.94)
D_1	0.0047	0.0045	0.0003	−0.0028	−8894.9	−41.2	−0.178
	(9.05)	(9.22)	(1.19)	(−0.27)	(−1.22)	(−49.13)	(−3.83)
D_2	0.0021	0.0023	−0.0002	−0.0119	−6516.7	−40.7	−0.146
	(4.01)	(4.55)	(−0.99)	(−1.13)	(−0.89)	(−48.47)	(−3.10)
D_3	0.0022	0.0023	−0.0001	0.0018	4208.4	−34.6	−0.021
	(4.37)	(4.80)	(−0.60)	(0.17)	(0.57)	(−40.68)	(−0.46)
D_4	0.0019	0.0018	0.0002	0.0224	9889.4	−25.2	0.068
	(3.70)	(3.73)	(0.81)	(2.16)	(1.35)	(−29.62)	(1.47)
D_5	0.0008	0.0008	0.0000	0.0074	10031.0	−15.8	0.140
	(1.55)	(1.60)	(0.18)	(0.72)	(1.37)	(−18.87)	(3.05)
D_6	0.0010	0.0010	0.0001	0.0004	4222.6	−5.6	0.060
	(1.91)	(1.97)	(0.28)	(0.04)	(0.58)	(−6.69)	(1.32)
D_8	0.0007	0.0007	−0.0001	0.0063	−5302.2	−0.3	−0.076
	(1.40)	(1.53)	(−0.23)	(0.62)	(−0.71)	(−0.33)	(−1.68)
D_9	0.0011	0.0008	−0.0003	−0.0022	−12609.0	−7.2	−0.146
	(2.13)	(1.62)	(−1.22)	(−0.21)	(−1.70)	(−8.54)	(−3.14)
D_{10}	0.0014	0.0014	−0.0002	−0.0058	−4690.3	−18.6	−0.133
	(2.63)	(2.88)	(−0.83)	(−0.57)	(−0.63)	(−22.18)	(−2.37)
D_{11}	0.0022	0.0022	−0.0001	0.0065	−5650.0	−27.7	−0.094
	(4.10)	(4.46)	(−0.46)	(0.62)	(−0.76)	(−32.97)	(−1.56)
D_{12}	0.0031	0.0031	0.0002	−0.0043	−16178.0	−36.1	−0.021
	(5.92)	(6.29)	(0.63)	(−0.42)	(−2.18)	(−42.97)	(−0.45)
R^2 Adj	0.03	0.03	0.00	0.00	0.12	0.97	0.02
F Stat	11.68	12.85	1.12	1.37	2.58	668.92	8.89

Table 3.7 continue

	$RV_t^{1/2}$	$C_t^{1/2}$	$J_t^{1/2}$	RJ_t	Stock	Temp	Spread
			Panel B: Heating Oil				
Intercept	0.0145	0.0143	0.0006	0.0417	126390.0	68.3	0.010
	(45.11)	(46.72)	(3.47)	(5.35)	(51.64)	(115.18)	(10.46)
D_1	0.0057	0.0056	0.0007	0.0220	−1516.2	−41.2	−0.025
	(12.54)	(12.70)	(2.82)	(1.98)	(−0.44)	(−49.13)	(−18.84)
D_2	0.0040	0.0041	−0.0001	0.0088	−11433.0	−40.7	−0.033
	(8.65)	(9.31)	(−0.39)	(0.78)	(−3.30)	(−48.47)	(−24.08)
D_3	0.0040	0.0040	0.0003	0.0093	−19706.0	−34.6	−0.026
	(8.90)	(9.37)	(1.35)	(0.85)	(−5.69)	(−40.68)	(−19.58)
D_4	0.0023	0.0022	0.0004	0.0240	−19891.0	−25.2	−0.018
	(5.10)	(5.13)	(1.61)	(2.17)	(−5.75)	(−29.62)	(−13.66)
D_5	0.0010	0.0010	−0.0001	0.0137	−13826.0	−15.8	−0.009
	(2.13)	(2.37)	(−0.32)	(1.24)	(−3.99)	(−18.87)	(−7.00)
D_6	0.0006	0.0006	0.0002	0.0177	−8764.6	−5.6	−0.003
	(1.40)	(1.44)	(0.64)	(1.61)	(−2.53)	(−6.69)	(−2.57)
D_8	0.0011	0.0012	−0.0001	0.0265	4410.4	−0.3	−0.001
	(2.40)	(2.72)	(−0.30)	(2.44)	(1.26)	(−0.33)	(−0.49)
D_9	0.0014	0.0011	0.0000	0.0023	6823.5	−7.2	−0.001
	(3.01)	(2.63)	(0.09)	(0.21)	(1.94)	(−8.54)	(−1.09)
D_{10}	0.0010	0.0010	0.0001	0.0121	6414.7	−18.6	−0.000
	(2.29)	(2.45)	(0.30)	(1.12)	(1.83)	(−22.18)	(−0.19)
D_{11}	0.0016	0.0016	−0.0001	0.0257	10570.0	−27.7	−0.017
	(3.38)	(3.76)	(−0.32)	(2.29)	(3.01)	(−32.97)	(−7.11)
D_{12}	0.0029	0.0030	−0.0000	0.0119	9853.3	−36.1	−0.013
	(6.45)	(6.98)	(−0.20)	(1.08)	(2.81)	(−42.97)	(−10.07)
R^2 Adj	0.06	0.07	0.00	0.00	0.52	0.97	0.28
F Stat	28.01	30.01	1.99	1.27	20.56	668.92	140.40
			Panel C: Natural Gas				
Intercept	0.0224	0.0217	0.0016	0.0732	2218.7	64.2	0.032
	(38.83)	(39.88)	(4.74)	(7.78)	(56.85)	(74.77)	(2.64)
D_1	0.0071	0.0075	−0.0004	−0.0085	−37.2	−45.4	−0.118
	(8.91)	(9.82)	(−0.91)	(−0.65)	(−0.67)	(−37.44)	(−6.90)
D_2	0.0048	0.0048	0.0003	0.0004	−626.9	−42.8	−0.114
	(5.82)	(6.12)	(0.50)	(0.03)	(−11.18)	(−35.29)	(−6.46)
D_3	0.0007	0.0011	−0.0008	−0.0242	−981.5	−34.8	−0.020
	(0.90)	(1.49)	(−1.54)	(−1.83)	(−17.92)	(−28.27)	(−1.16)
D_4	−0.0016	−0.0013	−0.0007	−0.0025	−1005.9	−24.9	0.010
	(−1.98)	(−1.71)	(−1.38)	(−0.18)	(−18.15)	(−20.25)	(0.57)
D_5	−0.0021	−0.0019	−0.0006	−0.0084	−740.4	−15.6	0.022
	(−2.65)	(−2.55)	(−1.20)	(−0.64)	(−13.47)	(−12.82)	(1.31)
D_6	−0.0005	−0.0003	−0.0004	0.0008	−342.2	−5.4	0.020
	(−0.59)	(−0.33)	(−0.90)	(0.06)	(−6.20)	(−4.49)	(1.15)
D_8	0.0030	0.0025	0.0007	0.0038	259.5	−0.9	0.069
	(3.73)	(3.35)	(1.36)	(0.29)	(4.74)	(−0.78)	(4.09)
D_9	0.0039	0.0043	−0.0007	−0.0140	559.2	−9.6	0.342
	(4.75)	(5.54)	(−1.50)	(−1.05)	(10.09)	(−7.87)	(19.73)
D_{10}	0.0032	0.0034	−0.0004	−0.0103	822.8	−20.6	0.327
	(4.02)	(4.46)	(−0.80)	(−0.79)	(14.67)	(−16.94)	(18.17)
D_{11}	0.0024	0.0026	−0.0003	−0.0101	880.1	−31.2	−0.112
	(2.98)	(3.36)	(−0.61)	(−0.75)	(15.62)	(−25.72)	(−3.45)
D_{12}	0.0080	0.0079	0.0001	0.0023	536.8	−41.3	−0.094
	(9.84)	(10.29)	(0.27)	(0.17)	(9.69)	(−34.06)	(−5.43)
R^2 Adj	0.09	0.10	0.01	0.00	0.81	0.95	0.32
F Stat	32.46	35.20	1.66	0.78	297.89	370.35	144.12

Table 3.8A: Regression analysis is used to test for differences in five-minute intraday volatilities on Wednesdays from May 2003 to December 2007 for crude oil and heating oil. During this period, the petroleum report was released on Wednesdays at 10:30. OLS is used to estimate the regression model,

$$\mathrm{Vol}_{t,k} = \beta_0 + \sum_{i \neq l} \beta_i \mathrm{D}_i + \epsilon_{t,k},$$

where $\mathrm{Vol}_{t,k}$ denotes the intraday volatility for the kth five-minute intraday interval, D_i is a dummy variable set to one for the ith five-minute interval and zero otherwise, and l denotes the five-minute interval 10:10 - 10:15 which is used as a basis for comparison. The intraday volatility is given by the absolute value of the difference between the logarithm of closing and opening prices per interval. The reported results are confined to the trading time 10:00 - 11:35. The interval 10:30 - 10:35 matches the announcement time of the petroleum report. The columns present estimates for announcement days with and without significant jumps, respectively. The significance level α to test for jumps is 0.99. t-statistics are reported in parentheses.

	Crude Oil		Heating Oil	
	Jump	No Jump	Jump	No Jump
Intercept	0.0015	0.0014	0.0009	0.0011
	(4.23)	(12.54)	(1.15)	(4.23)
10:00 - 10:05	−0.0004	−0.0003	−0.0003	0.0000
	(−0.49)	(−1.20)	(−0.29)	(0.08)
10:05 - 10:10	0.0000	−0.0001	0.0009	0.0006
	(0.09)	(−0.82)	(1.13)	(2.09)
10:15 - 10:20	−0.0006	−0.0002	0.0004	0.0006
	(−1.28)	(−1.19)	(0.44)	(2.04)
10:20 - 10:25	−0.0005	−0.0001	0.0003	0.0004
	(−0.97)	(−0.65)	(0.41)	(1.37)
10:25 - 10:30	0.0000	0.0003	0.0008	0.0009
	(0.10)	(1.86)	(0.94)	(3.10)
10:30 - 10:35	0.0116	0.0040	0.0139	0.0042
	(23.62)	(24.73)	(17.13)	(15.06)
10:35 - 10:40	0.0022	0.0014	0.0031	0.0019
	(4.38)	(8.93)	(3.77)	(6.90)
10:40 - 10:45	0.0005	0.0016	0.0011	0.0020
	(1.11)	(10.18)	(1.36)	(7.30)
10:45 - 10:50	0.0010	0.0014	0.0011	0.0021
	(2.15)	(8.59)	(1.33)	(7.44)
10:50 - 10:55	0.0005	0.0010	0.0014	0.0014
	(1.07)	(6.49)	(1.75)	(5.05)
10:55 - 11:00	0.0015	0.0010	0.0018	0.0015
	(3.08)	(6.27)	(2.28)	(5.26)
11:00 - 11:05	0.0008	0.0009	0.0020	0.0010
	(1.73)	(5.33)	(2.46)	(3.62)
11:05 - 11:10	0.0003	0.0007	0.0012	0.0012
	(0.66)	(4.41)	(1.46)	(4.22)
11:10 - 11:15	0.0002	0.0004	0.0009	0.0009
	(0.44)	(2.45)	(1.11)	(3.30)
11:15 - 11:20	0.0009	0.0006	0.0016	0.0008
	(1.90)	(3.73)	(2.03)	(3.04)
11:20 - 11:25	0.0004	0.0005	0.0011	0.0009
	(0.76)	(3.19)	(1.38)	(3.26)
11:25 - 11:30	0.0003	0.0005	0.0008	0.0008
	(0.68)	(3.15)	(0.97)	(3.00)
11:30 - 11:35	0.0003	0.0004	0.0013	0.0007
	(0.69)	(2.26)	(1.55)	(2.48)
R^2 Adj	0.46	0.13	0.47	0.12
F Stat	20.11	30.26	28.96	24.42

Table 3.8B: Regression analysis is used to test for differences in five-minute intraday volatilities for natural gas on Wednesdays from March 2000 to May 2002 when the storage reports was released around 14:00 (Panel A) and on Thursdays from June 2002 to December 2007 when the report was released around 10:30 (Panel B). The reported results are limited to the trading time 12:55 - 14:30 on Wednesdays and 10:00 - 12:10 on Thursdays. The five-minute interval 13:05 - 13:10 is used as the basis on Wednesdays and 10:05 - 10:10 on Thursdays. The intervals 14:00 - 14:05 and 10:30 - 10:35 match the announcement time of the natural gas storage report on Wednesdays and Thursdays, respectively. The columns present estimates for announcement days with and without significant jumps, respectively. The significance level α to test for jumps is 0.99. t-statistics are reported in parentheses.

	Jump	No Jump
Panel A: Wednesday		
Intercept	0.0014	0.0024
	(2.15)	(5.78)
12:55 - 13:00	0.0007	−0.0006
	(0.76)	(−0.95)
13:00 - 13:05	−0.0001	−0.0004
	(−0.06)	(−0.67)
13:10 - 13:15	−0.0002	0.0000
	(−0.18)	(0.02)
13:15 - 13:20	0.0001	−0.0005
	(0.10)	(−0.82)
13:20 - 13:25	−0.0001	−0.0003
	(−0.07)	(−0.46)
13:25 - 13:30	0.0002	−0.0004
	(0.22)	(−0.67)
13:30 - 13:35	−0.0002	−0.0005
	(−0.20)	(−0.79)
13:35 - 13:40	−0.0001	−0.0006
	(−0.08)	(−1.08)
13:40 - 13:45	0.0004	−0.0005
	(0.40)	(−0.89)
13:45 - 13:50	0.0003	0.0000
	(0.34)	(0.00)
13:50 - 13:55	0.0002	0.0002
	(0.25)	(0.44)
13:55 - 14:00	0.0019	0.0015
	(2.16)	(2.64)
14:00 - 14:05	0.0258	0.0096
	(28.77)	(16.76)
14:05 - 14:10	0.0059	0.0039
	(6.48)	(6.74)
14:10 - 14:15	0.0042	0.0044
	(4.60)	(7.61)
14:15 - 14:20	0.0035	0.0034
	(3.91)	(5.94)
14:20 - 14:25	0.0027	0.0036
	(2.97)	(6.30)
14:25 - 14:30	0.0044	0.0027
	(4.87)	(4.69)
R^2 Adj	0.47	0.20
F Stat	29.31	17.33

Table 3.8B continue

	Jump	No Jump
Panel B: Thursday		
Intercept	0.0023	0.0022
	(5.58)	(11.70)
9:55 - 10:00	−0.0017	−0.0008
	(−1.31)	(−1.86)
10:00 - 10:05	−0.0001	−0.0008
	(−0.09)	(−1.86)
10:10 - 10:15	−0.0003	−0.0003
	(−0.59)	(−1.31)
10:15 - 10:20	−0.0000	−0.0003
	(−0.01)	(−1.14)
10:20 - 10:25	−0.0007	−0.0003
	(−1.13)	(−1.02)
10:25 - 10:30	−0.0001	0.0001
	(−0.23)	(0.46)
10:30 - 10:35	0.0227	0.0062
	(38.60)	(24.07)
10:35 - 10:40	0.0033	0.0034
	(5.63)	(12.96)
10:40 - 10:45	0.0013	0.0027
	(2.19)	(10.41)
10:45 - 10:50	0.0014	0.0018
	(2.31)	(7.03)
10:50 - 10:55	0.0019	0.0019
	(3.20)	(7.36)
10:55 - 11:00	0.0016	0.0010
	(2.81)	(3.99)
11:00 - 11:05	0.0013	0.0013
	(2.25)	(4.86)
11:05 - 11:10	0.0008	0.0010
	(1.30)	(3.85)
11:10 - 11:15	0.0008	0.0013
	(1.32)	(4.96)
11:15 - 11:20	0.0002	0.0009
	(0.28)	(3.58)
11:20 - 11:25	0.0005	0.0008
	(0.85)	(2.96)
11:25 - 11:30	0.0006	0.0010
	(0.95)	(3.82)
11:30 - 11:35	0.0007	0.0007
	(1.19)	(2.65)
11:35 - 11:40	0.0009	0.0007
	(1.53)	(2.80)
11:40 - 11:45	0.0005	0.0005
	(0.83)	(2.02)
11:45 - 11:50	0.0005	0.0007
	(0.78)	(2.71)
11:50 - 11:55	0.0002	0.0004
	(0.38)	(1.37)
11:55 - 12:00	0.0003	0.0004
	(0.49)	(1.63)
12:00 - 12:05	0.0000	0.0006
	(0.05)	(2.47)
12:05 - 12:10	−0.0006	0.0004
	(−1.06)	(1.54)
R^2 Adj	0.48	0.13
F Stat	51.94	30.53

Table 3.9A: The table reports OLS estimates of the HAR-RV model (equation (3.30) on page 111) for daily ($h = 1$), weekly ($h = 5$), and monthly ($h = 22$) realized variance, RV_t, for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). Sine and cosine terms are added to capture seasonal variation. The model is also computed for two nonlinear forms of the realized variation: standard deviation form, $\text{RV}_t^{1/2}$, and logarithmic form, $\log(\text{RV}_t)$. The t-statistics in parentheses are computed based on Newey-West/Bartlett adjusted standard errors (Newey and West (1987)). The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

	$\text{RV}_{t,t+h}$			$(\text{RV}_{t,t+h})^{1/2}$			$\log(\text{RV}_{t,t+h})$		
h	1	5	22	1	5	22	1	5	22
				Panel A: Crude Oil					
β_0	0.000	0.000	0.000	0.003	0.004	0.007	−1.460	−1.869	−3.018
	(4.79)	(4.49)	(4.35)	(6.27)	(6.15)	(5.90)	(−9.00)	(−8.32)	(−7.65)
β_D	0.232	0.162	0.061	0.234	0.184	0.075	0.228	0.182	0.073
	(5.11)	(5.65)	(4.37)	(7.15)	(8.47)	(6.73)	(7.86)	(9.87)	(6.84)
β_W	0.356	0.226	0.159	0.352	0.215	0.178	0.364	0.218	0.188
	(7.59)	(3.62)	(4.12)	(6.83)	(3.33)	(4.37)	(8.31)	(4.06)	(4.90)
β_M	0.189	0.295	0.275	0.211	0.339	0.327	0.246	0.379	0.376
	(3.02)	(3.44)	(3.55)	(4.70)	(4.75)	(4.43)	(7.07)	(6.70)	(5.56)
AdjR^2	0.35	0.47	0.50	0.45	0.57	0.59	0.52	0.64	0.66
F	200.63	332.03	374.10	302.86	495.07	537.51	400.11	671.00	703.42
				Panel B: Heating Oil					
β_0	0.000	0.000	0.000	0.003	0.004	0.007	−1.430	−1.818	−3.013
	(3.88)	(3.94)	(4.13)	(5.54)	(5.79)	(6.21)	(−7.40)	(−6.95)	(−7.32)
β_D	0.246	0.151	0.059	0.211	0.146	0.066	0.189	0.132	0.062
	(3.96)	(4.75)	(5.55)	(5.21)	(5.78)	(6.27)	(6.50)	(7.48)	(6.45)
β_W	0.277	0.171	0.171	0.306	0.217	0.207	0.331	0.247	0.211
	(3.41)	(3.01)	(3.48)	(5.82)	(4.78)	(5.14)	(8.40)	(6.62)	(6.04)
β_M	0.292	0.407	0.330	0.303	0.400	0.335	0.317	0.402	0.356
	(6.23)	(7.47)	(5.05)	(7.67)	(7.90)	(5.69)	(8.36)	(8.47)	(6.36)
AdjR^2	0.37	0.50	0.56	0.44	0.59	0.62	0.47	0.64	0.66
F	205.75	360.61	447.11	275.43	517.19	586.81	319.99	633.94	696.84
				Panel C: Natural Gas					
β_0	0.000	0.000	0.000	0.003	0.003	0.007	−1.355	−1.889	−3.787
	(0.16)	(0.23)	(0.38)	(3.53)	(2.90)	(3.20)	(−9.18)	(−8.98)	(−9.95)
β_D	0.221	0.140	0.055	0.212	0.167	0.087	0.191	0.159	0.096
	(2.48)	(2.65)	(3.46)	(5.17)	(5.59)	(7.08)	(6.80)	(10.67)	(9.40)
β_W	0.360	0.262	0.229	0.416	0.350	0.331	0.470	0.400	0.407
	(4.94)	(3.03)	(3.08)	(11.57)	(7.75)	(5.86)	(12.37)	(11.08)	(7.88)
β_M	0.172	0.244	0.111	0.165	0.214	0.075	0.184	0.217	0.040
	(3.74)	(6.00)	(1.98)	(4.55)	(5.69)	(1.22)	(5.66)	(5.31)	(0.61)
AdjR^2	0.27	0.37	0.33	0.40	0.51	0.42	0.47	0.58	0.47
F	127.12	201.03	165.47	225.93	348.03	240.82	298.83	474.63	306.11

Table 3.9B: The table reports OLS estimates of the HAR-RV-J model (equation (3.31) on page 111) for daily ($h = 1$), weekly ($h = 5$), and monthly ($h = 22$) realized variance, RV_t, for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). Sine and cosine terms are added to capture seasonal variation. The model is also estimated for two nonlinear forms of the realized variation: standard deviation form, $\text{RV}_t^{1/2}$, and logarithmic form, $\log(\text{RV}_t)$. The t-statistics in parentheses are computed based on Newey-West/Bartlett adjusted standard errors (Newey and West (1987)). The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

	$\text{RV}_{t,t+h}$			$(\text{RV}_{t,t+h})^{1/2}$			$\log(\text{RV}_{t,t+h})$		
h	1	5	22	1	5	22	1	5	22
				Panel A: Crude Oil					
β_0	0.000	0.000	0.000	0.003	0.004	0.006	−1.067	−1.333	−2.116
	(3.24)	(3.96)	(4.66)	(5.74)	(6.54)	(6.80)	(−7.95)	(−6.65)	(−5.46)
β_D	0.579	0.373	0.206	0.343	0.248	0.121	0.300	0.222	0.099
	(5.85)	(5.16)	(4.28)	(7.15)	(7.70)	(5.27)	(8.01)	(9.99)	(7.33)
β_W	0.054	0.097	0.035	0.202	0.140	0.132	0.311	0.189	0.206
	(1.10)	(1.96)	(0.98)	(2.93)	(2.42)	(2.88)	(7.16)	(4.19)	(6.40)
β_M	0.296	0.323	0.373	0.286	0.413	0.443	0.268	0.428	0.432
	(4.38)	(3.27)	(3.73)	(5.96)	(6.04)	(6.02)	(9.13)	(8.79)	(7.07)
β_J	−1.161	−0.779	−0.426	−0.160	−0.140	−0.071	−93.090	−76.571	−47.977
	(−5.77)	(−5.07)	(−4.45)	(−2.11)	(−1.90)	(−1.79)	(−6.87)	(−6.82)	(−3.83)
AdjR^2	0.11	0.28	0.43	0.46	0.56	0.61	0.61	0.70	0.71
F	51.51	160.07	297.55	338.09	501.39	614.83	624.62	946.64	974.48
				Panel B: Heating Oil					
β_0	0.000	0.000	0.000	0.003	0.004	0.007	−1.285	−1.687	−2.891
	(4.07)	(4.05)	(4.19)	(5.58)	(5.82)	(6.24)	(−6.66)	(−6.58)	(−6.97)
β_D	0.330	0.206	0.108	0.234	0.166	0.087	0.213	0.154	0.083
	(4.13)	(4.79)	(4.65)	(5.57)	(5.45)	(5.04)	(7.35)	(8.03)	(6.19)
β_W	0.259	0.160	0.161	0.301	0.212	0.202	0.329	0.245	0.210
	(3.30)	(2.90)	(3.45)	(5.81)	(4.72)	(5.12)	(8.51)	(6.69)	(6.07)
β_M	0.274	0.396	0.319	0.298	0.396	0.331	0.310	0.395	0.350
	(5.77)	(7.39)	(5.05)	(7.51)	(7.80)	(5.64)	(8.23)	(8.36)	(6.27)
β_J	−0.409	−0.266	−0.234	−0.052	−0.046	−0.047	−406.496	−367.955	−343.616
	(−3.85)	(−3.43)	(−3.65)	(−2.83)	(−2.69)	(−2.40)	(−3.27)	(−4.46)	(−4.33)
AdjR^2	0.37	0.51	0.56	0.44	0.59	0.63	0.47	0.64	0.66
F	191.78	332.48	413.99	249.05	468.20	532.55	289.65	575.03	632.84
				Panel C: Natural Gas					
β_0	0.000	0.000	0.000	0.003	0.003	0.007	−1.200	−1.767	−3.670
	(0.04)	(0.15)	(0.35)	(3.40)	(2.84)	(3.18)	(−7.12)	(−8.14)	(−9.53)
β_D	0.273	0.189	0.087	0.264	0.206	0.109	0.228	0.189	0.124
	(2.24)	(3.15)	(5.39)	(6.50)	(6.66)	(7.41)	(7.75)	(9.34)	(8.36)
β_W	0.337	0.240	0.215	0.393	0.333	0.321	0.462	0.394	0.401
	(4.04)	(2.71)	(2.90)	(11.68)	(7.39)	(5.68)	(12.69)	(10.88)	(7.72)
β_M	0.171	0.242	0.110	0.163	0.212	0.074	0.175	0.209	0.033
	(3.88)	(6.08)	(1.97)	(4.55)	(5.66)	(1.21)	(5.40)	(5.12)	(0.50)
β_J	−0.185	−0.177	−0.112	−0.097	−0.073	−0.040	−164.304	−128.756	−122.877
	(−0.73)	(−1.92)	(−2.78)	(−2.55)	(−2.56)	(−2.20)	(−2.03)	(−2.77)	(−3.81)
AdjR^2	0.27	0.38	0.33	0.40	0.51	0.42	0.47	0.59	0.48
F	115.73	184.23	150.86	206.00	316.44	217.63	272.32	431.98	278.98

Table 3.10: Regression results of five models based on equation (3.37) where the dependent variable, the realized variation, is in standard deviation form, $\text{RV}_{t+1}^{1/2}$ $(h = 1)$. The dash (-) denotes that the corresponding variable is excluded. elws_t is the extreme low temperature surprise; invs_t is the inventory surprise; and spread_t^- denotes the interest-rate adjusted negative spread. The Fourier terms are estimated but not reported to conserve space. The weather and inventory estimates are scaled by 10^5. The t-statistics are obtained based on the procedure by Newey and West (1987) and reported in parentheses.

	1	2	3	4	5
			Panel A: Crude Oil		
β_0	0.003	0.003	0.003	0.003	0.003
	(6.27)	(5.74)	(6.10)	(5.93)	(5.77)
β_D	0.234	0.343	0.259	0.260	0.260
	(7.15)	(7.15)	(7.56)	(7.56)	(7.57)
β_W	0.352	0.202	0.339	0.343	0.340
	(6.83)	(2.93)	(6.37)	(6.70)	(6.43)
β_M	0.211	0.286	0.203	0.205	0.200
	(4.70)	(5.96)	(4.95)	(4.69)	(4.94)
$\lvert\text{elws}_{t+1}\rvert$	–	–	–	0.130	0.100
				(0.54)	(0.43)
$\lvert\text{elws}_t\rvert$	–	–	–	−0.400	−0.400
				(−1.39)	(−1.49)
$\lvert\text{elws}_{t-1}\rvert$	–	–	–	0.760	0.740
				(2.15)	(2.09)
$\lvert\text{invs}_{t+1}\rvert$	–	–	–	0.000	0.000
				(0.14)	(0.20)
$\lvert\text{spread}_t^-\rvert$	–	–	0.000	–	0.000
			(1.66)		(1.61)
β_J	–	−0.160	−0.060	−0.058	−0.061
		(−2.11)	(−2.30)	(−2.18)	(−2.30)
AdjR^2	0.45	0.46	0.45	0.45	0.45
F	302.86	338.09	250.04	196.40	184.07
			Panel B: Heating Oil		
β_0	0.003	0.003	0.004	0.003	0.004
	(5.54)	(5.58)	(6.02)	(5.53)	(6.53)
β_D	0.211	0.234	0.220	0.233	0.220
	(5.21)	(5.57)	(5.81)	(5.53)	(5.76)
β_W	0.306	0.301	0.276	0.297	0.276
	(5.82)	(5.81)	(4.72)	(5.60)	(4.69)
β_M	0.303	0.298	0.236	0.289	0.236
	(7.67)	(7.51)	(5.83)	(7.37)	(5.94)
$\lvert\text{elws}_{t+1}\rvert$	–	–	–	0.470	0.280
				(1.82)	(1.10)
$\lvert\text{elws}_t\rvert$	–	–	–	0.046	−0.100
				(0.13)	(−0.42)
$\lvert\text{elws}_{t-1}\rvert$	–	–	–	0.760	0.530
				(2.29)	(1.57)
$\lvert\text{invs}_{t+1}\rvert$	–	–	–	0.002	0.001
				(1.30)	(0.60)
$\lvert\text{spread}_t^-\rvert$	–	–	0.040	–	0.037
			(5.00)		(4.76)
β_J	–	−0.052	−0.049	−0.052	−0.049
		(−2.83)	(−2.44)	(−2.83)	(−2.46)
AdjR^2	0.44	0.44	0.45	0.44	0.45
F	275.43	249.05	235.75	179.58	173.31
			Panel C: Natural Gas		
β_0	0.003	0.003	0.004	0.003	0.004
	(3.53)	(3.40)	(4.08)	(3.48)	(4.55)
β_D	0.212	0.264	0.246	0.264	0.246
	(5.17)	(6.50)	(7.13)	(6.53)	(7.15)
β_W	0.416	0.393	0.356	0.393	0.354
	(11.57)	(11.68)	(7.06)	(11.58)	(6.88)
β_M	0.165	0.163	0.153	0.163	0.160
	(4.55)	(4.55)	(4.69)	(4.07)	(4.33)
$\lvert\text{elws}_{t+1}\rvert$	–	–	–	0.660	0.480
				(1.54)	(1.14)
$\lvert\text{elws}_t\rvert$	–	–	–	−0.700	−0.900
				(−1.29)	(−1.59)
$\lvert\text{elws}_{t-1}\rvert$	–	–	–	−0.500	−0.700
				(−1.16)	(−1.56)
$\lvert\text{invs}_{t+1}\rvert$	–	–	–	0.038	−0.015
				(0.54)	(−0.23)
$\lvert\text{spread}_t^-\rvert$	–	–	0.008	–	0.008
			(4.62)		(4.52)
β_J	–	−0.097	−0.086	−0.098	−0.086
		(−2.55)	(−2.10)	(−2.55)	(−2.12)
AdjR^2	0.40	0.40	0.41	0.40	0.41
F	225.93	206.00	193.25	147.66	142.46

3.7 Figures

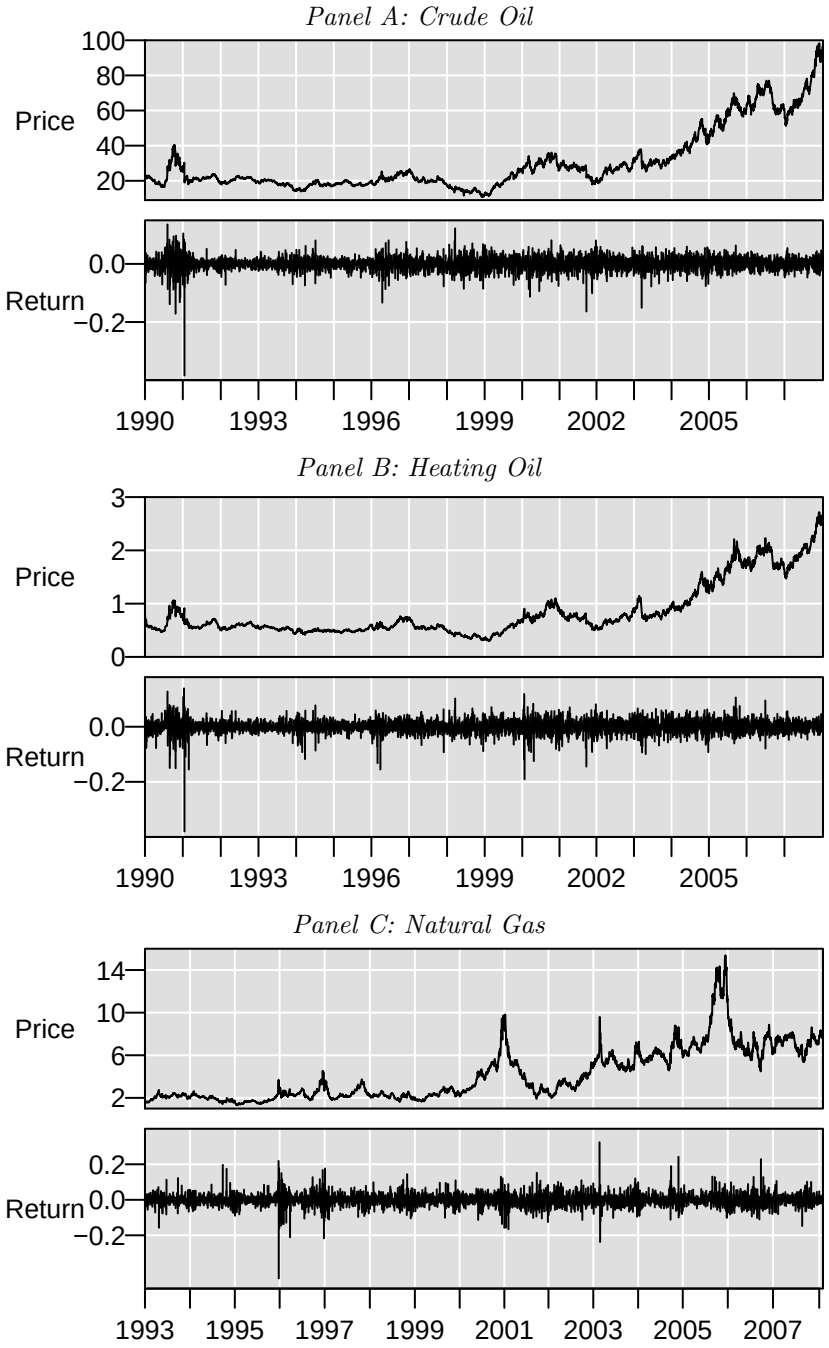

Figure 3.1: The figure graphs closing prices (top panel) and daily returns (bottom panel) for futures contracts on crude oil, heating oil and natural gas. The returns are computed as the logarithmic price difference between the last and first transactions per day.

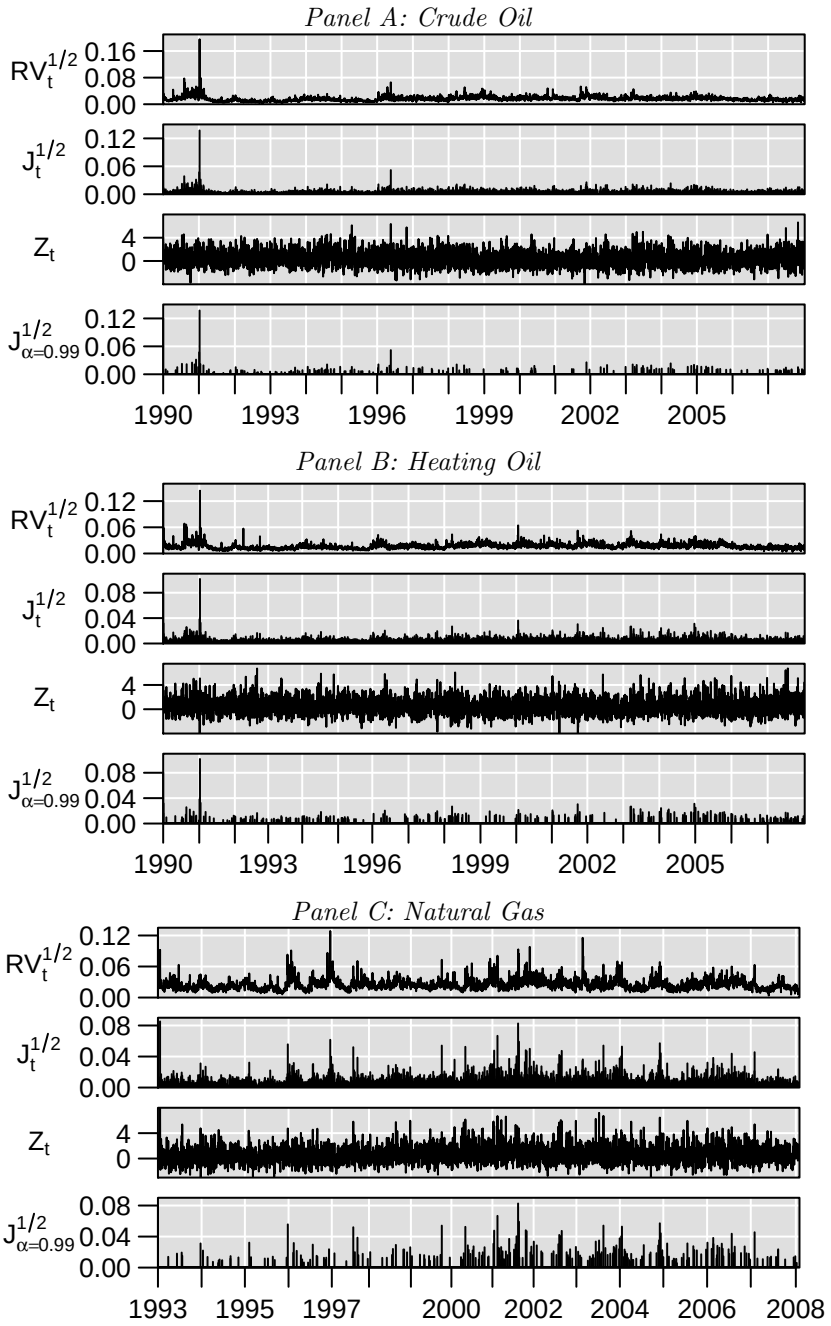

Figure 3.2: The figure graphs time-series of the realized volatility and jump component for futures contracts on crude oil, heating oil and natural gas. The top panel for respective contract graphs the daily realized volatility, $\text{RV}_t^{1/2}$ (equation (3.2)); the second panel plots the jump component $J_t^{1/2}$ (equation (3.21)); the third panel shows the jump statistic, Z_t (equation (3.12)); and the bottom panel plots the significant jump component, $J_{t,\alpha=0.99}^{1/2}$ (equation (3.22)). The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

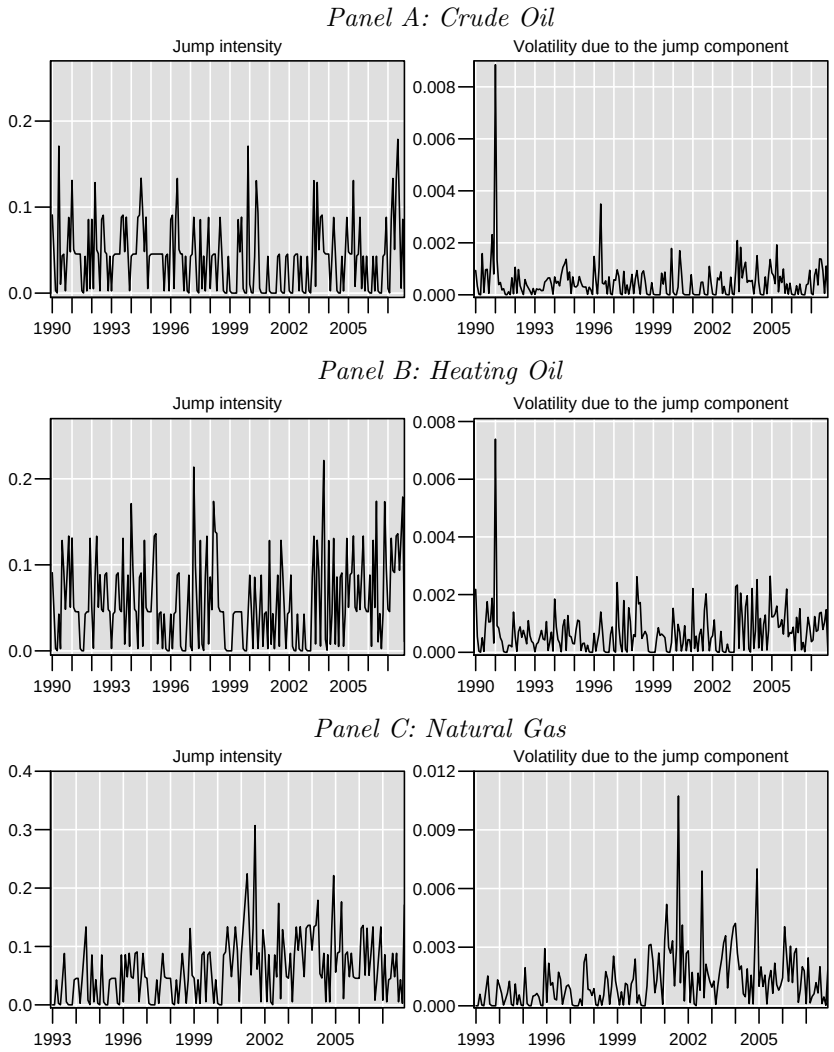

Figure 3.3: The figure graphs the exponentially smoothed (with smoothing parameter 0.94) monthly mean of the intensity (left column) and size (right column) of significant jumps, $J_{t,\alpha=0.99}^{1/2}$ (equation (3.22)) for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). The realized variations are computed based on five-minute returns and staggered returns with one lag offset.

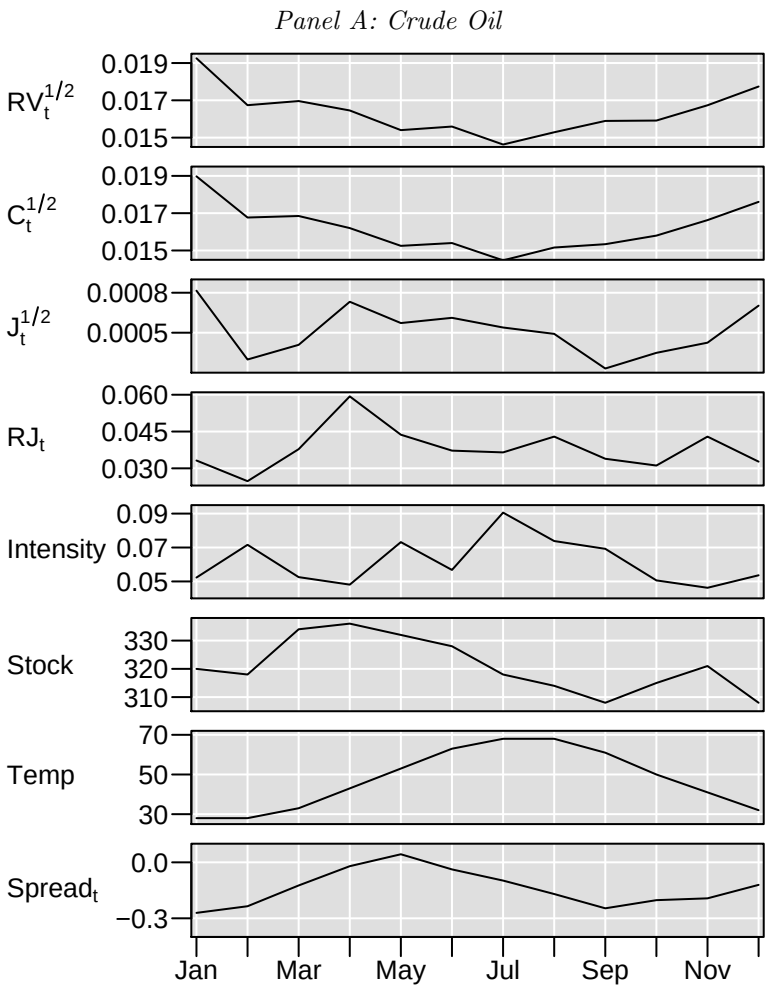

Figure 3.4: The figure graphs the monthly mean of the daily realized volatility, $\mathrm{RV}_t^{1/2}$ (equation (3.2)), the smooth component, $\mathrm{C}_{t,\alpha=0.99}^{1/2}$ (equation (3.23)), the significant jump component, $J_{t,\alpha=0.99}^{1/2}$ (equation (3.22)), and the relative contribution, RJ_t (equation (3.13)), for futures contracts on crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). Intensity denotes the fraction of days with a significant jump over the number of trading days per month. Stock denotes the inventory level which is scaled by 10^{-3}. Temp denotes the monthly average of daily minimum temperature in New York City for crude oil (Panel A) and heating oil (Panel B), and Chicago for natural gas (Panel C). Spread denotes the monthly average of daily spreads (equation (3.34)). The realized variations are computed based on five-minute intraday returns and staggered returns with one lag offset.

Figure 3.4 continue

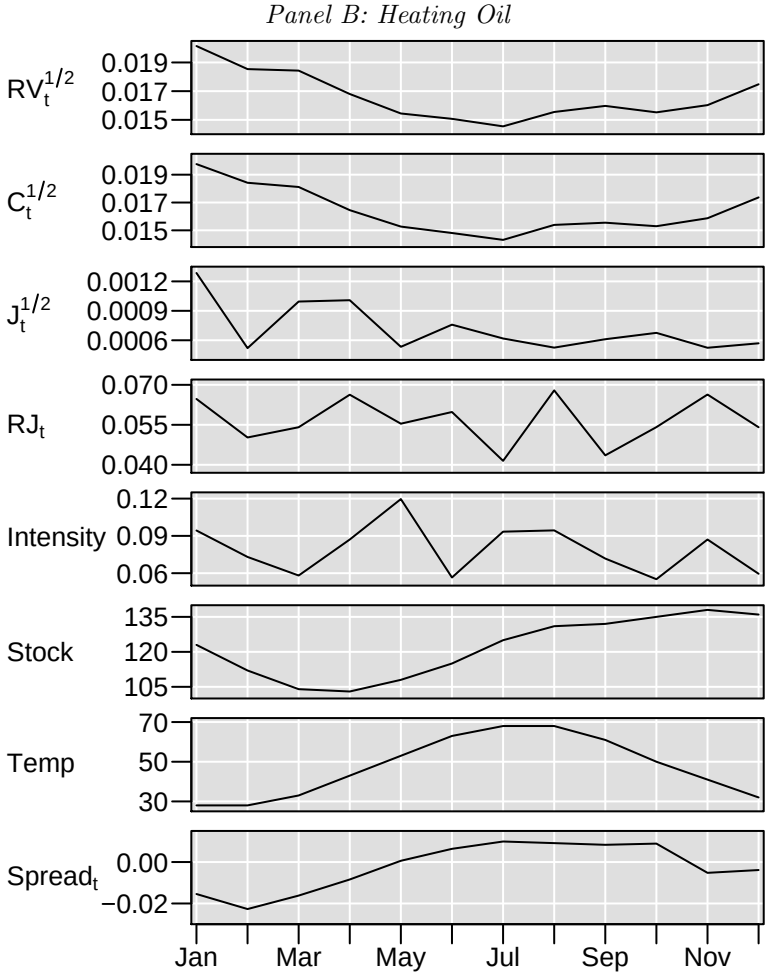

Figure 3.4 continue

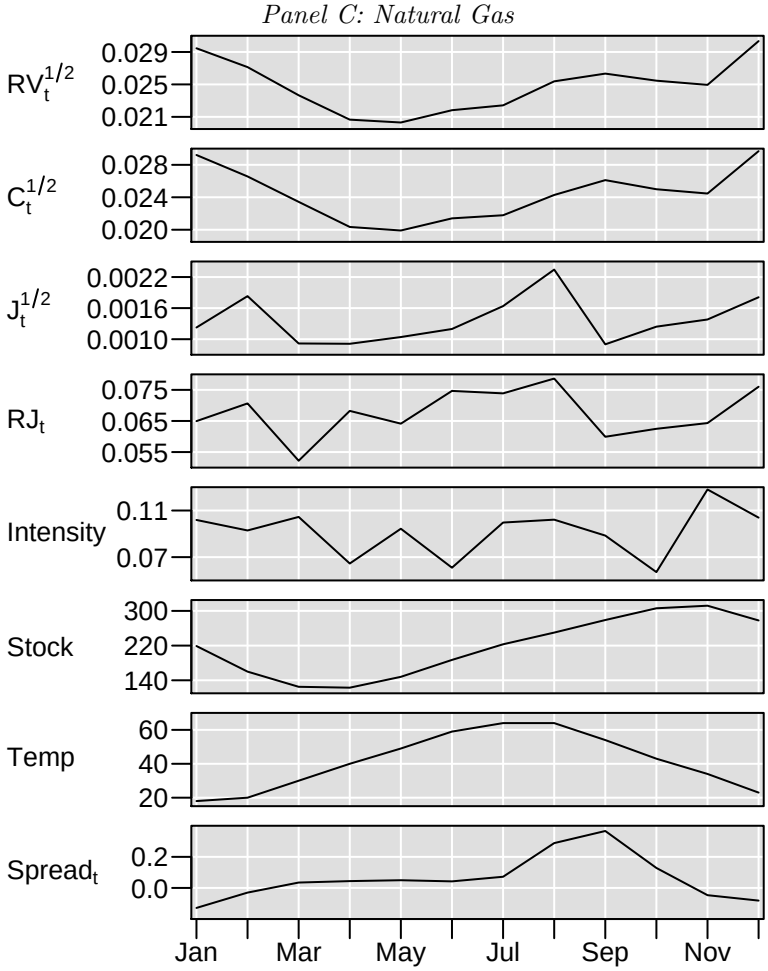

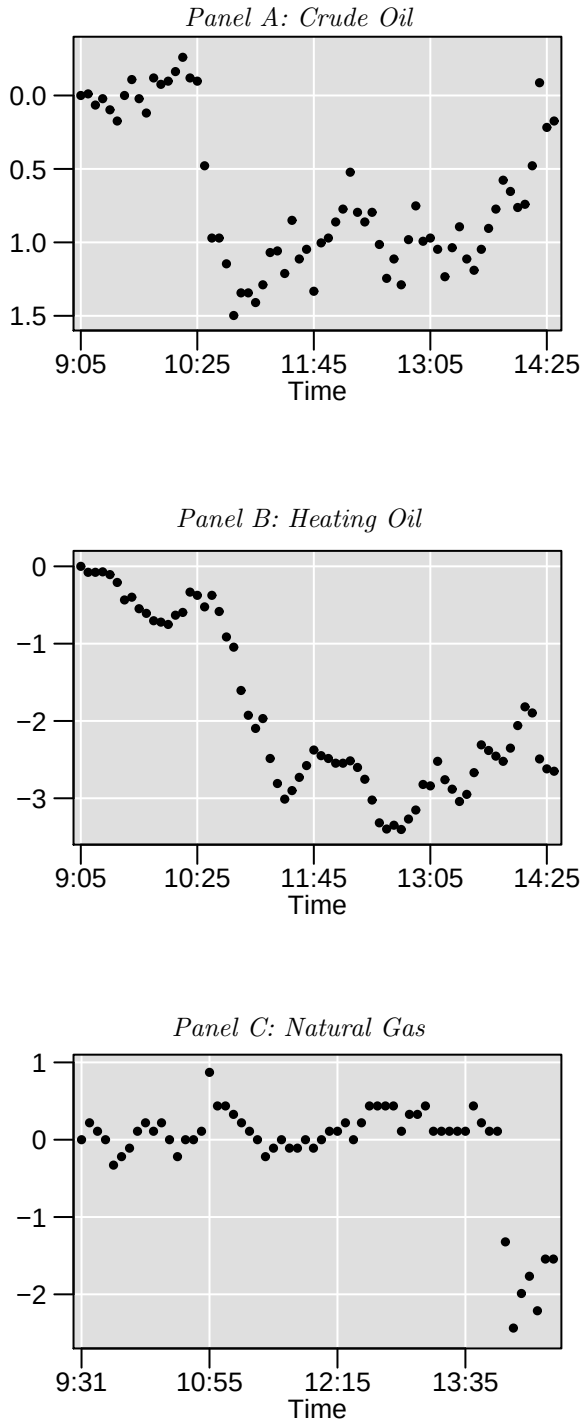

Figure 3.5: The figure graphs cumulative five-minute intraday price returns in percentage for crude oil on November 15, 2007, heating oil on February 14, 2007, and natural gas on May 2, 2001,

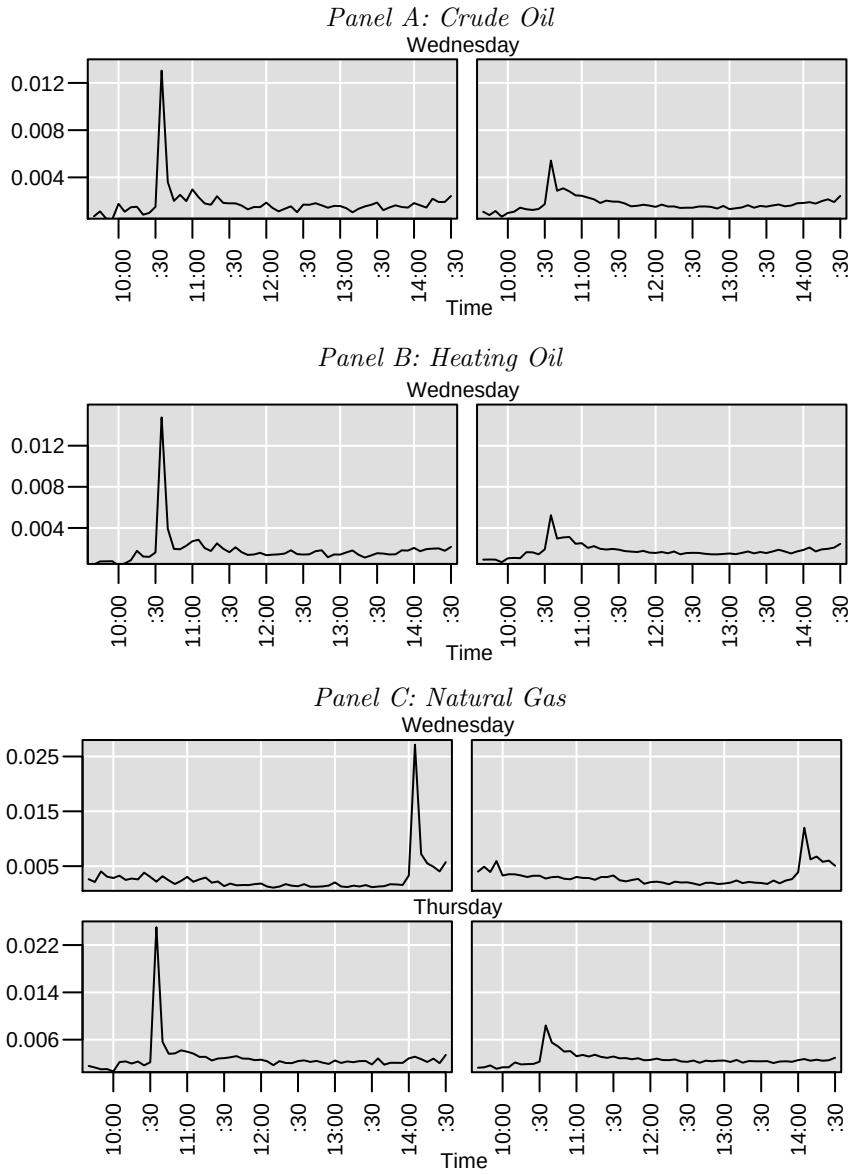

Figure 3.6: The figure graphs the mean of five-minute intraday volatilities, $\text{Vol}_{t,k}$ (equation (3.2)), for crude oil (Panel A), heating oil (Panel B) and natural gas (Panel C). The intraday volatility is given by the absolute value of the difference between the logarithm of closing and opening prices per interval. For crude and heating oil, the values are computed for Wednesdays from May 2003 to December 2007 when the petroleum report was released around 10:30. Two days are considered for natural gas; the storage report was released around 14:00 on Wednesdays from March 2000 to May 2002, and on Thursdays around 10:30 from June 2002 to December 2007. The left (right) column plots volatilities for announcement days with (without) significant jumps ($\alpha = 0.99$).

A Appendix

A.I Contract Specifications

Table A.1: Key features of contract specifications for crude oil, heating oil and natural gas.

Panel A: Light, Sweet Crude Oil Futures
Trading Unit
1000 U.S. barrels (42000 gallons)
Price Quotation
U.S. dollars and cents per barrel
Trading Hours
Open outcry trading is conducted from 9:00 AM until 2:30 PM.
Trading Months
Crude oil futures are listed nine years forward using the following listing schedule: consecutive months are listed for the current year and the next five years; in addition, the June and December contract months are listed beyond the sixth year.
Minimum Price Flucuation
\$0.01 (1¢) per barrel (\$10.00 per contract).
Maximum Daily Price Flucuation
\$10.00 per barrel (\$10,000 per contract) for all months.
Last Trading Day
Trading terminates at the close of business on the third business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th calendar day of the month is a non-business day, trading shall cease on the third business day prior to the business day preceding the 25th calendar day.
Settlement Type
Physical
Delivery
F.O.B. seller's facility, Cushing, Oklahoma, at any pipeline or storage facility with pipeline access to TEPPCO, Cushing storage, or Equilon Pipeline Co., by in-tank transfer, in-line transfer, book-out, or inter-facility transfer (pumpover).
Trading Symbol
CL
Panel B: Heating Oil Futures
Trading Unit
42000 U.S. gallons (1000 barrels)
Price Quotation
U.S. dollars and cents per barrel
Trading Hours
Open outcry trading is conducted from 9:00 AM until 2:30 PM.
Trading Months
36 consecutive months
Minimum Price Flucuation
\$0.0001 (0.01¢) per gallon (\$4.20 per contract).
Maximum Daily Price Flucuation
\$0.25 per gallon (\$10,500 per contract) for all months.
Last Trading Day
Trading terminates at the close of business on the last business day of the month preceding the delivery month.
Settlement Type
Physical
Delivery
F.O.B. seller's facility in New York harbor, ex-shore. All duties, entitlements, taxes, fees, and other charges paid. Requirements for seller's shore facility: capability to deliver into barges. Buyer may request delivery by truck, if available at the seller's facility, and pays a surcharge for truck delivery. Delivery may also be completed by pipeline, tanker, book transfer, or inter- or intra-facility transfer. Delivery must be made in accordance with applicable federal, state, and local licensing and tax laws.
Trading Symbol
HO

Panel C: Henry Hub Natural Gas Futures
Trading Unit
10,000 million British thermal units (mmBtu).
Price Quotation
U.S. dollars and cents per barrel
Trading Hours
Open outcry trading is conducted from 9:00 AM until 2:30 PM.
Trading Months
The current year plus the next twelve years through December 2020. A new calendar year will be added following the termination of trading in the December contract of the current year.
Minimum Price Flucuation
\$0.001 (0.1¢) per mmBtu (\$10.00 per contract).
Maximum Daily Price Flucuation
\$3.00 per barrel (\$30,000 per contract) for all months.
Last Trading Day
Trading terminates three business days prior to the first calendar day of the delivery month.
Settlement Type
Physical
Delivery
The Sabine Pipe Line Co. Henry Hub in Louisiana. Seller is responsible for the movement of the gas through the Hub; the buyer, from the Hub. The Hub fee will be paid by seller.
Trading Symbol
NG

A.II Modeling Daily Temperatures

I model daily low temperatures by the conditional mean temperature model proposed by Campbell and Diebold (2005),

$$\text{Tmin}_t = \text{Trend}_t + \text{Seasonal}_t + \sum_{l=1}^{L} \beta_{t-l}\text{Tmin}_{t-l} + \sigma_t\epsilon_t, \tag{A.1}$$

where,

$$\begin{aligned}
\text{Trend}_t &= \sum_{m=1}^{M} \xi_m t^m \\
\text{Seasonal}_t &= \sum_{p=1}^{P} \left(\sigma_{c,p} \sin\left(\frac{2\pi p d(t)}{365}\right) + \sigma_{s,p} \cos\left(\frac{2\pi p d(t)}{365}\right)\right) \\
\sigma_t^2 &= \sum_{q=1}^{Q} \left(\sigma_{c,q} \sin\left(\frac{2\pi q d(t)}{365}\right) + \sigma_{s,q} \cos\left(\frac{2\pi q d(t)}{365}\right)\right) \\
&\quad + \sum_{r=1}^{R} \alpha_r(\sigma_{t-r}\epsilon_{t-r})^2 + \sum_{s=1}^{S} \beta_s \sigma_{t-s}^2 \\
\epsilon_t &\sim \text{iid}(0,1),
\end{aligned} \tag{A.2}$$

where $d(t)$ cycles through $1, \ldots, 365$. I follow Campbell and Diebold (2005) and set $L = 25$, $M = 1$, $P = 3$, $Q = 3$, $R = 1$, and $S = 1$.

A.III Small Sample Properties of a Combined Statistic

This section reports a small Monte Carlo study on the combined statistic described in Section 3.2.1. I show that the test remains valid with reasonable power for price processes with market microstructure noise. The setup follows Huang and Tauchen (2005), who consider a one-factor stochastic volatility jump-diffusion model written as,

$$\begin{aligned} \mathrm{d}X_t &= \mu \mathrm{d}t + \mathrm{e}^{\beta_0 + \beta_1 v_t} \mathrm{d}w_{p,t} + \kappa_t \mathrm{d}q_t, \\ \mathrm{d}v_t &= \alpha_v v_t \mathrm{d}t + \mathrm{d}w_{v,t}, \end{aligned} \tag{A.3}$$

where v_t is a stochastic volatility factor; α_v is the mean reversion parameter; and $\mathrm{d}w_p$ and $\mathrm{d}w_v$ are standard Brownian motions with correlation, ρ. q_t is a discontinuous jump process where jumps occur at a rate denoted by λ. κ_t is the size of the jumps. In the following, I refer to the model defined in equation (A.3) as SV1F for $\lambda_t = 0$, that is, when no jumps are simulated, and SV1FJ otherwise.

Table A.2 presents values of the parameters in the data-generating processes that I consider. The values are obtained from Huang and Tauchen (2005), who select values based on empirical studies on observed market data reported in literature.

Table A.2: The experimental design for SV1F and SV1FJ (equation (A.3)) where the jump rate, λ, is set to zero for SV1F.

Parameter	Value
μ	0.030
β_0	0.000
β_1	0.125
α_v	-1.386
ρ	-0.620
λ	0.014
σ_{jmp}	1.500

I simulate observed prices per second from the stochastic differential equation following the Euler scheme. The number of simulated prices per interval t is equivalent to six hours and a half of trading, that is, t corresponds to a typical trading day. I compute intraday price returns for time intervals ranging from one to thirty minutes. I assume that the number of jumps in the SV1FJ model has a

Poisson distribution; hence, the interarrival times have an exponential distribution with parameter λ. The size of the jumps, κ, has a normal distribution with zero mean and variance, σ^2_{mn}.

Section 3.2.3 discusses the sources and impact of market microstructure noise. I report results based on an iid normal noise process where the noisy logarithmic prices, Y_{t_i}, are given by,

$$Y_{t_i} = X_{t_i} + \epsilon_{t_i}, \tag{A.4}$$

where ϵ_{t_i} denotes the microstructure noise.[22] I report results based on applying staggered returns to reduce the impact of noise, see Section 3.2.3.

Table A.3: The size of the combined statistic based on Z_{TPRM} and Z_{SWVR} is tabulated based on 10000 days simulated from the SV1F model as defined in equation (A.3). The experimental design is described in Table A.2. A noise process is added to the simulated prices, where the noise is assumed to be iid $\text{N}(0, \sigma^2_{mn})$; σ_{mn} is set to 0.000, 0.027, 0.040, 0.052, 0.065, 0.080. The panel label i denotes the staggered offset. The return horizons are one, three, five and thirty minutes. The test size is $\alpha = 0.01$.

	$\{\sigma_{mn}\}$					
Interval	0.000	0.027	0.040	0.052	0.065	0.080
			$(i = 0)$			
1 minutes	0.000	0.000	0.000	0.000	0.000	0.000
3 minutes	0.002	0.002	0.001	0.000	0.000	0.000
5 minutes	0.002	0.002	0.002	0.002	0.001	0.001
30 minutes	0.011	0.011	0.010	0.011	0.011	0.012
			$(i = 1)$			
1 minutes	0.001	0.000	0.000	0.000	0.000	0.000
3 minutes	0.002	0.002	0.001	0.001	0.001	0.001
5 minutes	0.002	0.003	0.002	0.002	0.002	0.001
30 minutes	0.016	0.016	0.015	0.016	0.016	0.016

Table A.3 presents rejection rates under the null hypothesis for different standard deviations of the noise component, σ_{mn}, ranging from 0.000 to 0.080 based on 10000 realizations. The test is highly conservative except for the longest sampling intervals irrespective of whether staggered returns are applied.

Table A.4 shows rejection rates under the alternative data-generating return process. The labels, NJ and J, denote days without and with a jump, respectively. The rows represent true events while the columns denote statistical inferences. Hence, the rows for the 2×2 matrices add up to one where the 1×1 element is the fraction of correct non-rejections of the null (no-jump) hypothesis and the 1×2 element is the false rejection rate. Meanwhile, the 2×1 element is the false non-rejection of

[22] I obtain analogous results for an AR(1) noise process (Aït-Sahalia et al. (2006)) and by enforcing a minimum tick size (Li and Mykland (2007)).

the null hypothesis and the 2×2 element is the correct rejection.

The power stays reasonably high even though the tests are highly conservative at all sampling rates except for the 30-minute return horizon. For example, the power drops from 0.638(0.0048) to 0.493(0.0050) (with standard deviation in parentheses) as the variance of the noise component increases from 0.000 to 0.080 when sampling every five minutes and applying staggered returns with one lag (see panel ($i = 1$)). The equivalent rates based on the Z_{TPRM} statistic alone range from 0.638 (0.0048) to 0.507 (0.0050) and for Z_{SWVR} range from 0.725 (0.0045) to 0.616 (0.0049).[23] The type I error for both Z_{TPRM} and Z_{SWVR} are greater than 0.01, respectively, while the errors for the combined test range from 0.001 to 0.003. The results suggests that the noise impact the statistics differently under the null hypothesis, thus leading to false rejections on different days. Hence, the combined statistic continues to reject the same false null hypotheses but reduces the type I error.

Table A.4: Confusion matrices are presented based on 10000 days simulated from the SV1FJ model as defined in equation (A.3). The experimental design is described in Table A.2. A noise process are added to the simulated prices, where the noise is assumed to be IID $\text{N}(0, \sigma_{mn})$; σ^2_{mn} takes the values 0.000, 0.027, 0.052, 0.080. Results are presented for four return horizons, which are one, three, five and thirty minutes. The panel label i denotes the staggered offset. The labels, NJ and J, denote days without and with a jump, respectively. The rows correspond to the actual event of a jump or no jump while the columns denote the statistical inference. The test size is $\alpha = 0.01$.

		$\{\sigma_{mn}\}$							
		0.000		0.027		0.052		0.080	
		(NJ)	(J)	(NJ)	(J)	(NJ)	(J)	(NJ)	(J)
					$(i = 0)$				
1 minutes	(NJ)	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000
	(J)	0.246	0.754	0.304	0.696	0.493	0.507	0.623	0.377
3 minutes	(NJ)	0.998	0.002	0.998	0.002	1.000	0.000	1.000	0.000
	(J)	0.341	0.659	0.406	0.594	0.493	0.507	0.580	0.420
5 minutes	(NJ)	0.998	0.002	0.998	0.002	0.998	0.002	0.999	0.001
	(J)	0.377	0.623	0.384	0.616	0.486	0.514	0.543	0.457
30 minutes	(NJ)	0.989	0.011	0.989	0.011	0.989	0.011	0.988	0.012
	(J)	0.754	0.246	0.783	0.217	0.768	0.232	0.775	0.225
					$(i = 1)$				
1 minutes	(NJ)	0.999	0.001	1.000	0.000	0.999	0.001	1.000	0.000
	(J)	0.261	0.739	0.319	0.681	0.435	0.565	0.522	0.478
3 minutes	(NJ)	0.998	0.002	0.998	0.002	0.999	0.001	0.999	0.001
	(J)	0.355	0.645	0.362	0.638	0.464	0.536	0.551	0.449
5 minutes	(NJ)	0.998	0.002	0.997	0.003	0.998	0.002	0.999	0.001
	(J)	0.362	0.638	0.406	0.594	0.457	0.543	0.507	0.493
30 minutes	(NJ)	0.984	0.016	0.984	0.016	0.984	0.016	0.983	0.017
	(J)	0.768	0.232	0.768	0.232	0.761	0.239	0.790	0.210

[23] Notice that I do not tabulate the power based on the individual statistics.

Chapter 4: Summary and Future Work

This dissertation contributes to the literature in two related areas. First, I evaluate the small sample properties of a nonparametric method to identify jumps in asset prices. Second, I apply an extension of this method to data from U.S. energy futures markets to document and examine jump processes and their relative contribution to the total volatility.

Detecting Jumps in Asset Prices Using Bipower Variation

In this chapter, I examine small sample properties of nonparametric statistics developed by Barndorff-Nielsen and Shephard (2004, 2006) that are applied to test for jumps in asset prices. I primarily focus on the implication of adding noise to the observed prices and recent methods to contend with such market frictions.

My work has produced several important empirical findings. The jump test statistics converge to the limiting normal distribution with zero mean and unit variance as the sampling interval approaches zero for efficient (noise-free) prices. These statistics have converged at a one-minute sampling interval. The convergence results are highly influenced by noise, however. Literature has previously established that market microstructure noise biases the statistics against identifying jumps. I find that the limiting distribution remains normal but the mean and variance estimates become negatively biased. As a result, the rejection rate decreases since the test is based on the right tail. The optimal sampling methods by Bandi and Russell (2006) and Zhang et al. (2005) reduce the bias and increase the power of the test statistics. These methods perform similarly to applying staggered returns. I propose a modified version of

the method by Bandi and Russell (2006) and show that it corrects the jump statistics from being slightly anti-conservative to become valid. Furthermore, it increases the power. I find that the size and power are similar for three different models of market microstructure noise. Specifically, adding serial correlation to the error process and introducing rounding errors do not have a significant impact beyond the effects of an iid normal noise process. Alternative jump distributions do not alter the small sample properties of the jump statistics considerably.

A number of future research directions has emerged that I have not pursued in this work. The experimental design generates prices per second throughout the trading day. In practice, there are longer and irregular gaps in between trades. Since the statistics rely on intraday data, this approach can only be applied to liquid markets; nonetheless, the impact of missing observations is an important extension. Second, I determine the optimal sampling rate per day, which assumes that the market is sufficiently liquid. Future research may determine whether it is beneficial to estimate the rate over a longer time horizon and, as a result, obtain more robust estimates of the optimal sampling rates. Third, more recent nonparametric methods based on high-frequency data have been proposed in the literature, such as Fan and Wang (2007), Jiang and Oomen (2008), and Sen (2008). The finite sample properties of these and the statistics that I consider in this study may be compared in the future.

Volatility and Jump Dynamics in U.S. Energy Futures Markets

This essay applies a nonparametric method based on realized and bipower variations calculated from intraday high-frequency data to identify jumps in prices from U.S. energy futures markets. The futures contracts are crude oil, heating oil and natural gas, which are traded on NYMEX. The sample period of the high-frequency intraday data spans from January 1990 to January 2008. I apply alternative methods such as

staggered returns and optimal sampling frequency methods to remove the effects of microstructure noise which biases the tests against detecting jumps.

I document several interesting and important findings. The natural gas market is the most volatile among these price series. There are upward trends in volatility for the three series during the sample period; for natural gas the increase is primarily due to the jump component while the smooth component dominates the increase in the crude oil and heating oil markets. There are significant jumps (price spikes) in all price series. I document that the total realized volatility and smooth sample component for natural gas and heating oil are higher in the winter months than during the summer months. These results are consistent with the general hypothesis that when short run demand for natural gas and heating oil are suddenly shifted higher due to extreme cold weather during the winter, the short run supply is inelastic due to low inventories at this time of the year. In an intraday analysis, I find that the volatility is higher during inventory news announcement periods and that many jumps are associated with these announcement dates. Furthermore, it is interesting to observe that for all markets, that volatility returns to preannouncement levels faster when there is a jump in the futures price changes than when there is no jump. This suggest that jumps may be efficient pricing of new information.

I find that including the jump component as an explanatory variable improves the performance of a simple realized volatility forecasting model. The coefficient of the jump component attains the largest value at the daily lag and decreases for corresponding weekly and monthly regression estimates. Moreover, all of the coefficients of jumps are negative and most are significant. The above two results indicate that the jump component in the price process produces transitory surges in volatility and that there is a strong reversal in the volatility on the subsequent days of a jump. In addition, I add cold weather and inventory surprises to the model and find that

these lead to an increase in volatility in natural gas and heating oil markets. Furthermore, I compute the lagged interested-rate adjusted spread and document that the spread can serve as a proxy for negative inventory periods since the significance of the weather and inventory variables drops while the spread remains highly significant when including all three variables.

The empirical results have several important implications for participants in energy futures markets. For option traders, option pricing models with jumps are preferred over the Black option model in these three energy futures contracts during the winter months. Market participants may prefer to employ an optimal hedging ratio with jumps to hedge their exposure to energy price risk. Market risk managers should be aware that the shapes of the return distributions of energy futures prices will change over time due to the presence of significant jumps in these markets.

Several possible future research directions emerge from this work. For example, the empirical results suggest that volatility returns to preannouncement levels after an inventory report release faster on days with jumps compared to days with no jumps. A more detailed study of this finding may follow the approach used by Ederington and Lee (1993, 1995). I show that including jump and seasonal components as well as the interest-rate adjusted spread improve a model of the realized volatility. A study on whether this leads to improved forecasts of volatility compared with other models may be of interest, see for example Andersen et al. (2003). The considerable contribution of jump components to the total volatility may have important implications on Value-at-Risk (VaR) calculations in these markets. Feng and Brooks (2002) and Sullivan et al. (2002) provide a framework for VaR calculation and hedging of futures contracts and present empirical results for natural gas futures contracts. Kruse (2006) investigate if forecasting realized volatility improves VaR forecasts. They do not consider jumps, however.

Bibliography

Bibliography

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625.

Bandi, F. M. and Russell, J. R. (2006). Separating microstructure noise from volatility. *Journal of Financial Economics*, 79:655–692.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2:1–48.

Barndorff-Nielsen, O. E. and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics*, 4(1):1–30.

Ederington, L. H. and Lee, J. H. (1993). How markets process information: News releases and volatility. *Journal of Finance*, 48(4):1161–1191.

Ederington, L. H. and Lee, J. H. (1995). The short-run dynamics of the price adjustment to new information. *Journal of Financial and Quantitative Analysis*, 30(1):117–134.

Fan, J. and Wang, Y. (2007). Multi-scale jump and volatility analysis for high-frequency financial data. *Journal of the American Statistical Association*, 102(480):1349–1362.

Feng, W. and Brooks, R. (2002). Risk management of natural gas exposure. Working Paper No. 01-10-02.

Jiang, G. J. and Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise - A "Swap Variance" approach. *Journal of Econometrics*, 144(2):352–370.

Kruse, R. (2006). Can realized volatility improve the accuracy of value-at-risk forecasts? Working Paper.

Sen, R. (2008). Jumps and microstructure noise in stock price volatility. In Gregoriou, G. N., editor, *Stock Market Volatility*. Chapman Hall-CRC/Taylor and Francis.

Sullivan, J. H., Brooks, R., and Stoumbos, Z. G. (2002). Structured assessment of risk systems and value at risk (VaR). Working Paper No. 02-06-01.

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*, 100:1394–1411.

Curriculum Vitae

Johan Bjursell received a Bachelor of Science in Computing Science from the University of Gothenburg, Göteborg, Sweden, in August 2000, and a Master of Science in Computer Science from University of Southern MS in May 2002. He spent the summer of 2007 as a Visiting Ph.D. Student at Smarts Group Pty Ltd, Sydney, Australia, doing research on real-time systems for market surveillance. He was an intern at the Commodity Futures Trading Commission, Washington DC, from June 2004 to December 2005 primarily working on empirical market microstructure in futures markets. He has worked as a Graduate Research and Teaching Assistant for the Volgenau School of Information Technology and Engineering, George Mason University, since January 2004.