Essays on financial econometrics: cojump detection and density forecasting

Thesis submitted to Lancaster University in fulfilment of the requirements for the degree of Doctor of Philosophy in Accounting and Finance

by

Rui Fan

BBA (Hons) in Accounting and Finance (The Hong Kong Polytechnic University)

MRes in Finance (Lancaster University)

Department of Accounting and Finance

Lancaster University Management School

April 2016

Declaration

I, Rui Fan, declare that this thesis titled, "Essays on financial econometrics: cojump detection and density forecasting" is my own work. It has not been submitted in substantially the same form for the award of a higher degree elsewhere. Where any part of this thesis has previously been published, or submitted for a higher degree elsewhere, this has been clearly stated. Where I have consulted the published work of others, this has been clearly attributed. Where the thesis is the result of joint research, I have stated clearly what was done by others and what I have contributed myself.

Essays on financial econometrics: cojump detection and density forecasting

Rui Fan BBA, MRes

Thesis submitted to Lancaster University in fulfilment of the requirements for the degree of Doctor of Philosophy in Accounting and Finance

April 2016

Abstract

We choose the Andersen et al. (2007) and Lee and Mykland (2008) jump detection tests to detect intraday price jumps for ten foreign exchange rates and cojumps for six groups of two dollar exchange rates and one cross exchange rate at the one-minute frequency for five years from 2007 to 2011. We reject the null hypothesis that jumps are independent across rates as there are far more cojumps than predicted by independence for all rate combinations. We find that one dollar rate and the cross rate combination almost always has more cojumps than the two dollar rates combination. We also find some clustering of jumps and cojumps can be related to the macroeconomic news announcements affecting the exchange rates. The two selected jump detection tests find a similar number of jumps for ten foreign exchange rates.

We compare density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices for four horizons ranging from one day, one week, two weeks to one month. We use the Heston model which incorporates stochastic volatility to extract risk-neutral densities from option prices.

From historical high-frequency returns, we use the HAR-RV model to calculate realised variances and lognormal price densities. We use a nonparametric transformation to transform risk-neutral densities into real-world densities and make comparisons based on log-likelihoods. For the sixty-eight combinations from seventeen stocks for four horizons, the transformed lognormal Black-Scholes model gives the highest log-likelihoods for fifty-nine combinations. The HAR-RV model and the Heston model have similar forecast accuracy for different horizons, either before or after applying a transformation which enhances the densities. The transformed real-world densities almost always pass the Kolmogorov-Smirnov and Berkowitz tests, while the untransformed risk-neutral densities almost always fail the diagnostic tests.

Acknowledgements

I would like to show my deepest gratitude to my main supervisor, Prof. Stephen J. Taylor, who always gives me helpful guidance and advice throughout my PhD study. Without him, I could not have completed this thesis and my PhD. It is my greatest honor to be his research student. I would also like to thank my other two supervisors, Dr. Eser Arisoy and Dr. Matteo Sandri, who always help me and encourage me.

I am also grateful to fellow PhD students in the department, with whom I have had many conversations about my research. In particular, I would like to thank Xi Fu and Tristan Linke for their interesting comments. I would also like to thank Tobias Langenberg and Xiu-Ye Zhang for being good office mates.

Last but not least, I am indebted to my wife, Haiheng Yu, for her dedication, understanding and support throughout my PhD. I also want to thank my parents, they has brought me up and always supported me. I dedicate this thesis to them.

Contents

List of Tables

List of Figures

1. Introduction

The availability of high-frequency price data since the mid-1990s, typically recorded at least once every five minutes, has triggered many developments in financial econometrics during the last twenty years. This results from the additional information it contains relative to lower frequency, e.g. daily data. Subsequently, many studies focus on using this additional information. One example is nonparametric volatility modelling, or more specifically, the realised variance, which gives an accurate estimate of the quadratic variation of the underlying price process. There is a long-running debate whether continuous time processes for asset prices contain a jump component generated by a compound Poisson process besides a diffusion component driven by a Brownian motion process. A significant amount of literature, including Duffie et al. (2000), Pan (2002), Eraker et al. (2003) and Eraker (2004), has argued from low-frequency evidence that the jump component should be included.

Following this, many nonparametric tests have been proposed to detect the occurrence of jumps using high-frequency data. Barndorff-Nielsen and Shephard (2004a, 2006) initiated a method which separates the realised variance measure into a continuous component and a jump component. Subsequently, many other nonparametric tests have been proposed, which includes Jiang and Oomen (2008), Corsi et al. (2010), Podolskij and Ziggel (2010) and Andersen et al. (2012). However, all these tests only tell us on which day the jumps occur, but cannot detect the exact timing of the jumps. Andersen et al. (2007) and Lee and Mykland (2008) develop tests, on the other hand, which can detect the occurrence of jumps at the intraday level. However, nonparametric jump tests can only detect large jumps, small jumps cannot be detected and the average size of jumps is overestimated.

Many studies document the existence of jumps in high-frequency prices, typically recorded at least once every five minutes, which includes Huang and Tauchen (2005), Andersen et al. (2007), Lee and Mykland (2008), Lee and Hannig (2010) and Evans (2011). These papers have studied the equity market and all authors identify some returns which are too large to be explained by a diffusion process; the typical frequency of these large returns is one every two weeks. Lahaye et al. (2011), Dungey et al. (2009) and Dungey and Hvozdyk (2012) further present some evidence of the occurrence of jumps in foreign exchange and Treasury bond markets.

On the other hand, not so many researchers have investigated cojumps, which are simultaneous jumps in the prices of two or more assets. Dungey et al. (2009) and Dungey and Hvozdyk (2011) explore the U.S. Treasury market, Lahaye et al. (2011) investigate the U.S. equity indices, U.S. Treasury bond index and the dollar exchange rates, while Gilder et al. (2014) study the S&P 500 index and 60 of its constituent stocks. All these studies examine the timing of cojumps relative to macroeconomic news announcements, while similar macro investigations for jumps are detailed in Andersen et al. (2003, 2007).

We employ the non-parametric tests of Andersen et al. (2007) and Lee and Mykland (2008) to extract jumps and cojumps from foreign exchange rates in this thesis. In the proposed implementation, the tests compare one-minute returns with critical values dependent on a significance level, daily measures of price variation calculated from bipower and estimates of the intraday volatility pattern described in Taylor (2005). As dollar rates move together, and since cross rates are constrained by no-arbitrage equations, we expect that when one rate jumps many others will also jump. We will examine the number, sign and size of foreign exchange rate jumps, and will compare the number of cojumps with expectations derived from no-arbitrage principles.

This thesis also contains new results about density forecasts for asset prices. Density forecasts are of importance to central bankers, risk managers and other decision takers for activities such as policy-making, risk management and derivatives pricing. They can also be used to assess market beliefs about economic and political events when derived from option prices.

Volatility forecasts produce forward-looking information about the volatility of the asset price in the future, while density forecasts are more sophisticated, as they provide information about the whole distribution of the asset's future price. Since option prices reflect both historical and forward-looking information, volatility forecasters might rationally prefer implied volatilities from option prices to realised variance calculated from historical time series. We anticipate a similar preference could apply to density forecasts. There is a considerable literature comparing volatility forecasts obtained from option prices with volatility forecasts obtained from the history of asset prices. Blair et al. (2001), Jiang and Tian (2005), Giot and Laurent (2007) and Busch et al. (2011) state that option forecasts are more informative and accurate than historical forecasts of index volatility even when the historical information set includes high-frequency returns.[1] Few studies, however, make similar

[1] Further comparisons are in Poon and Granger (2003), Martens and Zein (2004) and Taylor et al. (2010).

comparisons for density forecasts. Liu et al. (2007), Shackleton et al. (2010) and Yun (2014) provide comparisons for UK and US stock indices, Hog and Tsiaras (2010) focus on crude oil prices, Ivanova and Gutierrez (2014) look at interest rates, and Trujillo-Barrera et al. (2012) investigate lean hog futures prices. These studies show option-based density forecasts outperform historical forecasts for a one-month horizon. There are no known previous results for individual stocks, so a major contribution of this thesis is to provide the first comparison for density forecasts obtained from option prices and historical intraday returns for individual stocks.

Many methods have been proposed to obtain risk-neutral densities from option prices. Parametric methods include a lognormal mixture (Ritchey, 1990; Jondeau and Rockinger, 2000), a generalised beta distribution (Anagnou-Basioudis et al., 2005; Liu et al., 2007), and a lognormal-polynomial (Madan and Milne, 1994; Jondeau and Rockinger, 2000). Other approaches include discrete probabilities (Jackwerth and Rubinstein, 1996), a nonparametric kernel regression (Ait-Sahalia and Lo, 1998; Bates, 2000), and densities obtained from implied volatility splines (Bliss and Panigirtzoglou, 2002). All these methods, however, only provide densities for horizons which match option expiry dates. We instead fit a stochastic process, to obtain densities for all horizons, following the innovative methodology of Shackleton et al. (2010).

As the implied volatility smile effect indicates that risk-neutral densities are not lognormal and volatility is not constant, some studies use a stochastic process to model volatility. Heston (1993) assumes the volatility follows a mean-reverting square-root process and gives a closed form solution for option prices. We use

Heston's model in this thesis as its parameters can be calibrated from daily option records and it also has a tractable density formula based on inverting characteristic functions. Extensions of the Heston (1993) model are in Bates (1996) who also incorporates jumps, and in Duffie et al. (2000), Eraker (2004), Eraker et al. (2003) and Pan (2002) who include a jump process in both price and volatility components. However, we do not evaluate a jump component because Bakshi et al. (2003) and Shackleton et al. (2010) both find that adding jumps does not improve their empirical results much. Furthermore, our nonparametric transformations can systematically improve mis-specified risk-neutral densities.

We compare density forecasts derived from option prices using the Heston (1993) model and forecasts obtained from historical time series using the Corsi (2009) Heterogeneous Autoregressive model of Realised Variance (HAR-RV). However, the risk-neutral density is a suboptimal forecast of the future distribution of the asset price as there is no risk premium in the risk neutral world, while in reality investors are risk-averse. Hence we need to use economic models and/or econometric methods to transform risk-neutral densities into real-world [2] densities. Pricing kernel transformations include power and/or exponential utility functions (Bakshi et al., 2003; Bliss and Panigirtzoglou, 2004; Liu et al., 2007), and the hyperbolic absolute risk aversion (HARA) function (Kang and Kim, 2006). Liu et al. (2007) use both utility and statistical calibration transformations, and they show that statistical calibration gives a higher log-likelihood than a utility transformation. Shackleton et al. (2010) compare parametric and nonparametric transformations, obtaining good results for the

[2] Similar to Liu et al. (2007) and Shackleton et al. (2010), we use "real-world" rather than other alternative adjectives, such as "risk-adjusted", "statistical", "empirical", "physical", "true", "subjective" and "objective", etc., which are all used in the literature to indicate that the price distributions incorporate risk preferences.

latter. Hence we also transform the risk-neutral densities into real-world densities using a nonparametric transformation.

Early studies including Bakshi et al. (2003), Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005) use the full dataset to make risk-transformations. The real-world densities obtained are then ex post because each forecast is made using some information from later asset prices. However it is best to apply ex ante transformations as in Shackleton et al. (2010). Thus we only use past and present asset and option prices to construct real-world densities. We investigate seventeen stocks from the Dow Jones 30 Index for four horizons ranging from one day to one month for the period from 2003 to 2012.

The rest of the thesis is structured as follows. Chapter 2 reviews the related literature on high-frequency price dynamics. It discusses variation measures, market microstructure noise, daily and intraday nonparametric jump tests, cojump tests, macroeconomic news announcements and some empirical evidence.

Chapter 3 uses the ABD and the LM jump detection tests to detect intraday price jumps for ten foreign exchange rates and cojumps for six groups of two dollar rates and one cross rate at the one-minute frequency for five years from 2007 to 2011. The null hypothesis that jumps are independent across rates is rejected as there are far more cojumps than predicted by independence for all rate combinations. Some clustering of jumps and cojumps are also detected and can be related to the macroeconomic news announcements affecting the exchange rates. The selected ABD and LM jump detection tests detect a similar number of jumps for ten foreign exchange rates.

Chapter 4 reviews the relevant literature on volatility and density forecasts, which includes high-frequency data, options information in volatility forecasts, methods to extract risk-neutral densities from option prices, transformations from risk-neutral densities into real-world densities, applications of density forecasts to estimate the risk aversion of investors, infer probabilities of future market changes and manage risk, and the methods to evaluate density forecasts.

Chapter 5 compares density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices. We use the Heston model which incorporates stochastic volatility to extract risk-neutral densities from option prices. From historical high-frequency returns, we use the HAR-RV model to calculate realised variances and lognormal price densities. We use a nonparametric transformation to transform risk-neutral densities into real-world densities and make comparisons based on log-likelihoods for four horizons ranging from one day to one month. For the sixty-eight combinations from seventeen stocks for four horizons, the transformed lognormal Black-Scholes model gives the highest log-likelihoods for fifty-nine combinations. The HAR-RV model and the Heston model have similar forecast accuracy for different horizons, either before or after applying a transformation which enhances the densities.

Finally, Chapter 6 summarises the thesis and draws some conclusions. It also points out possible directions of future research.

2. High frequency price dynamics literature

2.1 Introduction

High-frequency data research has prospered in the last 20 years after several studies were presented at the Olsen & Associates conference in 1995. New ideas have been generated as fast as data availability and computational power, which permitted the focus to shift from daily frequency to very high-frequency intraday data.

Financial markets are now known to possess significant price discontinuities, called jumps, in financial time series data. Many recent theoretical and empirical studies have confirmed the existence of jumps and their substantial influence on hedging risks and exposure to derivatives using underlying assets under certain circumstances. Faced with unpredictable jumps, there are some risks which we can no longer hedge and researchers find that jumps are empirically difficult to detect as only discrete data are available from continuous-time models.

This chapter reviewing high-frequency literature is organized as follows. Section 2.2 looks at various variation measures including realised variance (RV), bipower variation (BV), quadratic variation (QV) and integrated variance (IV). Section 2.3 investigates microstructure noise and the optimal sampling frequency. Section 2.4 explores a range of jump detection test methods and their empirical implications. Section 2.5 looks into a different area of research of cojumps and Section 2.6 studies the effects between macroeconomic news announcements and jumps. Section 2.7 summarizes and concludes.

2.2 Variation measures

A general jump diffusion process has price dynamics given as

$$dP_t = \mu_t dt + \sigma_t dW_t + \kappa_t dN_t \qquad (2.1)$$

where P_t is the logarithmic price of the asset, μ_t is the drift of the logarithmic asset price, σ_t is the volatility process and can be stochastic, κ_t is the size of any jump in P_t at time t, W_t is a standard Wiener process and N_t is a Poisson process counting the number of jumps from time 0 to t inclusive. Various variation measures are explored to better capture the dynamic properties of the high-frequency price data.

2.2.1 Integrated variance

The integrated variance IV_t for day t is the quantity stated in (2.2)

$$IV_t = \int_{t-1}^{t} \sigma^2(s) ds. \qquad (2.2)$$

IV_t equals the variation of the continuous component. The notation is for one interval of time from t-1 to t and conceptually the total continuous variation of a day can be summed up by integrating the variance through the day and is a random quantity.

2.2.2 Quadratic variation

The quadratic variation QV_t for day t includes the variation of the jump component. It is characterised as the summation of the integrated variance and the squared jumps as in (2.3)

$$QV_t = IV_t + \sum_{t-1 \leq s \leq t} \kappa^2(s) \tag{2.3}$$

QV_t is a measure of the total variation during the day. Prices move for one of the two reasons, either due to diffusion or owing to jump. Movements caused by diffusion are reflected in the first term IV_t, while movements as a result of jumps are captured in the second term of squared jumps. The difference between QV_t and IV_t isolates the jump contribution to price variation.

2.2.3 Realised variance and realised range-based variance

Realised variance measures how much the prices move over periods of time. Suppose each daily return r_t is the sum of N intraday returns, where N is a positive integer. Hence the day is subdivided into N parts and in each part there is an intraday return which is usually small. Changing the number N will change the values of intraday returns. The representation is given as:

$$r_{t,j,N} \text{ }, j = 1, 2, \dots, N$$

where t denotes the day, j indicates the intraday period and N counts the number of intraday periods. Andersen and Bollerslev (1998), Andersen et al. (2001) and Andersen et al. (2003) propose the realised variance $RV_t(N)$ for day t in (2.4)

$$RV_t(N) = \hat{\sigma}^2_{t,N} = \sum_{j=1}^{N} r^2_{t,j,N} \tag{2.4}$$

Christensen and Podolskij (2007) suggest the intraday realised range-based variance (RRV). This estimator differs from RV, which sums squared returns over intraday periods, but instead uses the difference between the highest and the lowest price during the period. It is defined as

$$RRV_t(N) = \frac{1}{\lambda_{2,N}} \sum_{j=1}^{N} s_{t,j,N}^2 \quad (2.5)$$

where $\{s_{t,j,N}\}_{i=1,\ldots,N}$ are the observed ranges of the intraday intervals of the log prices, and $\lambda_{2,N}$ is a normalizing constant. This daily range measure is preferred because it provides a more efficient estimator of daily variance than squared daily returns. Martens and van Dijk (2007) also show empirically that the realised range has a lower mean squared error than the realised variance by using S&P 500 futures and S&P 100 constituent stocks.

2.2.4 Bipower variation and other variance estimators

Barndorff-Nielsen and Shephard (2004a) show that a quantity called bipower variation is to some degree robust to rare jumps in the log-price process. They also provide the first robust method which splits quadratic variation into its continuous and jump components, without making strong parametric assumptions, and this stimulated significant discussions about bipower variation consequently. The term $BV_t(N)$ for day t is defined in (2.6) as

$$BV_t(N) = \frac{\pi N}{2(N-1)} \sum_{j=2}^{N} |r_{t,j-1,N}| \, |r_{t,j,N}| \quad (2.6)$$

where $2/\pi$ comes from a standard normal variable z, for which $E[|z|] = \sqrt{2/\pi}$; the multiplier $N/(N-1)$ is used to ensure an unbiased estimate. If a jump occurs in either one of successive absolute returns in (2.6) but not both, then a big return is multiplied by a small return and this results in a small number. Hence a jump will not have big effect on $BV_t(N)$ but will have a much larger impact on $RV_t(N)$, and the difference between the two helps to identify the jump component.

Though Barndorff-Nielsen and Shephard (2004a, 2006) show that the bipower variation $BV_t(N)$ converges to the true underlying quantity of integrated variance IV_t and the realised variance $RV_t(N)$ converges to the quadratic variation QV_t, as the sampling frequency N approaches to infinity, when microstructure noise is ignored as in (2.1). However, we need to choose N to simultaneously avoid bias from microstructure effects and measurement errors in $BV_t(N)$ and $RV_t(N)$.

Motivated by the idea of bipower variation and transformed power functions, Mancini (2004, 2009) and Jacod (2008) propose a jump-robust variance estimator, the threshold realised variance (TRV), which is given as

$$TRV_N = \sum_{j=1}^{N} \left|r_{t,j}\right|^2 1_{\{|r_{t,j}|<cN^{-\varpi}\}}, \text{ for } \varpi \in (0, 1/2). \tag{2.7}$$

The choices $\varpi = 0.47$ and $c = 6\sqrt{IV}$, and IV is estimated using BV, as recommended by Ait-Sahalia and Jacod (2009). Lowering the value of c could introduce a downward bias but will also give more robustness to jumps.

Though BV is an efficient estimator of IV asymptotically, Andersen et al. (2012) argue that the BV estimator is not robust to jumps and may cause an upward bias. They suggest the MinRV and MedRV estimators of integrated variance in (2.8)

$$MinRV_N = \frac{\pi}{\pi-2}\left(\frac{N}{N-1}\right)\sum_{j=1}^{N-1} min(|r_{t,j}|, |r_{t,j+1}|)^2$$

$$MedRV_N = \frac{\pi}{6 - 4\sqrt{3} + \pi}\left(\frac{N}{N-2}\right)\sum_{j=2}^{N-1} med\left(|r_{t,j-1}|, |r_{t,j}|, |r_{t,j+1}|\right)^2 \quad (2.8)$$

When a large jump occurs in one of the consecutive returns, the MinRV or the MedRV estimators will simply square the minimum or the median of the adjacent diffusive term.

Christensen et al. (2010) introduce the quantile-based variance (QRV) estimator, which is efficient, and robust to jumps and outliers. Their modified form estimator is also immune to microstructure noise and hence can be applied on high-frequency data. QRV is similar to RRV and we replace intraday ranges with intraday quantiles. The detected jumps may affect the extreme quantiles, but the impact can be ignored if the quantiles are not used in estimation. More than one quantile can be used in estimation to improve efficiency. Simulation results show that QRV is more robust than BV in finite samples.

2.3 Market microstructure noise

Microstructure noise results from many sources and can be categorised into two groups according to Hansen and Lunde (2006) and Ait-Sahalia et al. (2011). The discrete microstructure noise refers to the tick size, which is the minimum allowed asset price change, and the positive bid ask spread. The residual microstructure noise focuses on the trading environment and includes effects from order flow, block trades and asymmetry information. Diebold and Strasser (2013) investigate microstructure noise from an economic perspective.

If there are no microstructure effects, the realised variance is a fairly accurate measure of the price variation, and it becomes more accurate as the sampling frequency increases. But the presence of the microstructure noise creates the problem of choosing the appropriate sampling frequency when we study high-frequency data. Consequently many studies investigate how to correct microstructure bias when estimating the RV.

2.3.1 Volatility signature plot and choosing the appropriate frequency

Some studies investigate the volatility signature plot. The volatility signature plot shows the average realised variance, which is an average measure of variance across time, i.e.

$$V_N = \frac{1}{M}\sum_{t=1}^{M} \hat{\sigma}_{t,N}^2$$

The length of the return interval is $1/N$ and the database covers M days. If there was no microstructure noise, the number V_N would be approximately constant irrespective of the number M. Appropriate values of N should be found for which V_n is approximately the same for $1 \leq n \leq N$, consequently $\hat{\sigma}_{t,N}^2$ is free from microstructure bias.

The volatility signature plots of forty U.S. equities from a working paper version (dated 29th May 2007) of Bollerslev et al. (2008) present some examples. A typical U-shaped signature plot of Exxon Mobil starts high, comes down, and then becomes flat. The microstructure noise gives no extra variation when the length of the return interval exceeds twenty minutes. Bollerslev et al. (2008) state that 17.5-minute is an appropriate time between prices for large American firms.

Other researchers also use volatility signature plots to choose the optimal sampling frequency to minimize the bias from microstructure noise. Hansen and Lunde (2006) perform empirical analysis of Dow Jones Industrial Average 30 stocks and indicate that the microstructure noise can be neglected when intraday returns are sampled at lower frequencies, such as 20-minute interval. Bandi and Russell (2006) investigate S&P 100 constituent stocks and state that the 5 minute sampling frequency is optimal to maximize the accuracy of variance estimates and minimize the noise from microstructure effects. Andersen et al. (2007) indicate that two-minute is an appropriate frequency for the S&P 500 index which is a popular U.S. index. The two-minute frequency contains a lot of information and is a very high frequency.

2.3.2 Kernel-based estimators

Zhou (1996) is one of the first to suggest kernel-based estimators to correct the bias resulting from autocorrelation induced by microstructure noise. Kernel-based estimators correct the bias by including auto-covariance terms in the estimators. Hansen and Lunde (2006) also apply the kernel-based estimators and claim that the microstructure noise correlates with the price and is time-dependent.

2.3.3 Subsampled estimators

Zhang et al. (2005) introduce a new two-scale realised variance (TSRV), and they claim that it is unbiased and more efficient than the RV estimator. Zhang (2006) further improves the TSRV and proposes the multi-scale realised variance (MSRV). Ait-Sahalia et al. (2011) extend the TSRV and the MSRV under more general assumptions.

2.3.4 Pre-averaging estimators

Podolskij and Vetter (2009a, b) and Jacod et al. (2009) introduce the idea of pre-averaging, where the variance estimators are obtained from price averages over short time intervals and can reduce microstructure noise. Simulation results show that pre-averaging variance estimators have lower variance and are less biased.

2.4 Jump detection test methods

There are in general two large groups of nonparametric jump detection tests. The first category can examine the exact timing of the occurrence of the jumps, while the second group only investigates if a jump occurs or not, but cannot give the exact timing of the jump arrival. The first category can be further subdivided into daily and intraday jump tests, where the former looks at the day when the jump is present and the latter focuses on the intraday intervals. Some empirical evidence regarding jump tests in different asset markets is also presented in their detection.

2.4.1 Daily nonparametric jump tests

Barndorff-Nielsen and Shephard (2004a, 2006) (BNS) propose a daily jump detection test method based on comparing the statistical difference between the two variation measures, RV and BV, and the null hypothesis states that there is no jump on the day.

Huang and Tauchen (2005) expand the set of plausible test statistics. They study the fractional difference between the two measures of variation relative to the total, which is the relative proportion of variation attributed to jumps. They suggest comparing the relative jump measure

$$RJ_t = \frac{RV_t - BV_t}{BV_t} \tag{2.9}$$

with a consistent estimate of the standard error. The test statistic is

$$z_t = \frac{RJ_t}{\sqrt{Y_t}}\text{ , with}$$

$$Y_t = cN^{-1} \max\left(1, \frac{TP_t}{BV_t^2}\right), c \cong 0.609$$

$$TP_t = \frac{dN^2}{N-2} \sum_{j=3}^{N} \left|r_{t,j-2,N}\right|^{4/3} \left|r_{t,j-1,N}\right|^{4/3} \left|r_{t,j,N}\right|^{4/3}, d \cong 0.831^{-3}$$

and $\sqrt{Y_t}$ is the standard error. Barndorff-Nielsen and Shephard (2004a) states that TP_t is the tripower quarticity which estimates the integrated fourth power of volatility. The test statistic is compared with the standard normal distribution. As *N* increases, the null distribution of z_t converges to the standard normal when microstructure effects are ignored. The jump is detected when the value of $|z_t|$ is very large.

Christensen and Podolskij (2007) implement a jump test similar to the BNS test, but the test is based on the statistical difference between RRV and range-based BV, and the simulations indicate that the test has better power but worse size compared to the BNS test. Podolskij and Vetter (2009b) again propose a method similar to the BNS test, but using pre-averaging variance estimators. Jiang and Oomen (2008) also introduce a BNS-type jump test which is based on variance swaps.

2.4.2 Intraday nonparametric jump tests

Andersen, Bollerslev and Dobrev (2007) (ABD) apply an alternative test that determines which returns exhibit significant evidence of a jump. The idea is to inspect the total variation when the null hypothesis asserts that the price process is a diffusion. If the null is correct, BV can estimate the total variation. The total variation $BV_t(N)$ for day t is then divided equally into N parts for each part of the day. After taking the square root, we multiply the standard deviation by some constant c_N which is determined by the significance level of the hypothesis test and the cumulative distribution function (c.d.f.) of the standard normal distribution. A jump is detected from a return whenever

$$|r_{t,j,N}| > c_N \sqrt{N^{-1} BV_t(N)}. \qquad (2.10)$$

If (2.10) holds, the absolute return is large relative to what is expected over a short interval of time. The multiplier c_N also needs to be considered to determine if the returns are large or not, which is the critical value from the standard normal distribution and is determined by the daily significance level. 10^{-5} and 10^{-3} are two commonly chosen daily significance levels, because such small significance levels can help to protect against inaccurate bipower estimation, and the variation of the volatility during the trading day itself. Thus choosing small significance levels can make the results more certain by avoiding random variation and inappropriate assumptions. For example, as in ABD, if the daily significance level is 10^{-5} and N = 195 for 2-minute returns, then each return is tested under a level equal to $10^{-5}/195$ and $c_N = 5.45$. The reason for such a large number is due to the small daily significance level 10^{-5}. A number in the region of $5x10^{-8}$ is produced when it is divided by the number of intervals in a day, and this is an extremely small probability.

Lee and Mykland (2008) (LM) propose a new intraday jump detection test using high-frequency data. Unlike ABD, the test does not use BV, but refers to the distribution of the maximum values of the test statistic under the null hypothesis of a diffusion process. The choice of significance levels may ensure that only large jumps are being detected. Their Monte Carlo simulation results also show that the test has better size and power than the BNS test, where size is the probability of falsely rejecting the true null hypothesis, and is also referred as the probability of making a Type I error. While a Type II error is the failure to reject a false null hypothesis.

Lee and Hannig (2010) extend the LM test and employ a combination of the QQ test and a belief measure to identify the presence of small jumps besides large jumps in financial markets. However, similar to the LM test, no adjustment is made for the intraday volatility pattern.

Fan and Wang (2007) propose a wavelet method which removes the jumps and can estimate integrated volatility more accurately. Simulation results show that the wavelet test method has better size than the BNS test under the null hypothesis of no jumps.

2.4.3 Other nonparametric jump tests

Ait-Sahalia and Jacod (2009) suggest a new test to detect jumps by considering power variation. The estimator is given as

$$\hat{B}\left(p, \frac{1}{N}\right) = \sum_{j=1}^{N} \left|r_{t,j,N}\right|^{p} \tag{2.11}$$

where $p \geq 2$ and $\frac{1}{N}$ is the interval during which the prices are observed and goes towards zero. They then compare the measure over two different time scales, and the jump test statistic is the ratio of the two

$$\hat{S}\left(p, k, \frac{1}{N}\right) = \frac{\hat{B}\left(p, \frac{k}{N}\right)}{\hat{B}\left(p, \frac{1}{N}\right)} \tag{2.12}$$

This test statistic converges to 1 or 2, respectively in the presence or absence of jumps, when $p = 4$ and $k = 2$ as suggested by Ait-Sahalia and Jacod (2009).

Podolskij and Ziggel (2010) introduce a new jump detection test by comparing the difference between the realised power variation and a modified measure equivalent to the TRV in Mancini (2004, 2009). The estimator contains an indicator function equal to 1 when the threshold is satisfied and 0 otherwise. Their test uses the pre-averaging technique to minimize the bias attributable to microstructure effects.

Corsi et al. (2010) employ a local variance based threshold estimator to detect jumps, and the test can identify spurious jumps which are large returns and reduce the bias. Specifically, they propose the threshold bipower variation by combining the bipower variation and threshold estimation. Empirically, they investigate S&P 500 index, US individual stocks and bond yields, and show that their method can improve the forecasts of future volatility, particularly for periods after the occurrence of a jump.

2.4.4 Empirical evidence

ABD investigate two-minute returns from the S&P 500 futures contracts for the 17 years between 1988 and 2004, and they find 382 detected jumps in 4126 days (so less than one in ten days) when the daily significance level is 10^{-5}. The average bipower variation equals 95.6% of the average realised variance.

Huang and Tauchen (2005) explore the S&P 500 index between 1997 and 2002. Between 15% and 28% of the days have significant values of z at the 5% significance level; hence the fraction of the days identified with jumps is much higher than the significance level. If the significance level is pushed down to 0.1%, still a large number of days are detected as having evidence of jump effects.

Bollerslev et al. (2008) in their Figure 3 display test values of the firm Procter and Gamble (PG), of which 17 days have significant values of z at the significance level of 0.1%, from a dataset of 1246 days. Given the Type I error rate is approximately one out of a thousand for one day, there is evidence of significant jump effects. In terms of Type II errors, it is unknown how many jumps are undetected, but the jump detection tests may only find big jumps but fail to find small jumps. This can happen in a jump-diffusion model when the return is not big and the diffusion part dominates.

Lee and Mykland (2008) perform an empirical study of U.S. equity markets and collect high-frequency returns for three firms and the S&P 500 index, for a short period from 1st September to 30th November 2005. For individual stocks, jumps are related to firm-specific news releases such as scheduled earnings announcements, and other unscheduled news. For the index, jumps are associated with overall market news announcements including Federal Open Market Committee (FOMC) meetings and macroeconomic reports.

2.5 Cojumps

The previous section focuses on jumps in a univariate price process, and a natural extension is to consider the multivariate case. Few studies, however, investigate cojumps, which are simultaneous jumps in individual stocks, the index and other asset classes.

2.5.1 Cojump tests

Barndorff-Nielsen and Shephard (2004b) extend their univariate bipower approach to a multivariate case to identify cojumps between a pair of returns, but their theory is difficult to implement empirically. Gobbi and Mancini (2007) extend the TRV in Mancini (2004, 2009) to a bivariate setting. A cojump occurs if the threshold jump test identifies jumps simultaneously in both series. But they do not implement the test empirically. Jacod and Todorov (2009) extend the test in Ait-Sahalia and Jacod (2009) to a bivariate case. Two null hypotheses are tested on the days detected with jumps: the null hypothesis of common jumps must not be rejected, while the null of disjoint jumps must be rejected.

Bollerslev et al. (2008) conduct an empirical investigation based on high-frequency intraday returns for a sample of 40 large U.S. equities and the corresponding index of the same stocks over the sample period from 2001 to 2005. The index has fewer jumps than the individual stocks due to diversification of idiosyncratic jumps. They propose the mean cross-product (mcp) statistic given by the normalised sum of the intraday high-frequency returns as

$$mcp_{t,j} = \frac{2}{n(n-1)} \sum_{i=1}^{n-1} \sum_{k=i+1}^{n} r_{i,t,j} r_{k,t,j} , \qquad \text{for } j = 1, 2, \dots, M,$$

when there are M intraday returns for each of n assets. The test statistic is then studentised using daily means and standard deviations as

$$z_{mcp,t,j} = \frac{mcp_{t,j} - \overline{mcp}_t}{s_{mcp,t}}, \qquad \text{for } j = 1, 2, \dots, M, \tag{2.13}$$

where

$$\overline{mcp}_t = \frac{1}{M} mcp_t = \frac{1}{M} \sum_{j=1}^{M} mcp_{t,j}$$

and

$$s_{mcp,t} = \sqrt{\frac{1}{M-1} \sum_{j=1}^{M} (mcp_{t,j} - \overline{mcp}_t)^2}.$$

Bollerslev et al. (2008) do not identify an asymptotic distribution for this statistic but use a bootstrap to get the distribution under the null hypothesis of no jumps.

Caporin et al. (2014) introduce a novel nonparametric test for economically and statistically significant multivariate jumps for an arbitrary number of stocks. The test compares two smoothed power variations, and high values of the test statistics indicate a multi-jump among the stocks.

2.5.2 Empirical evidence

Dungey et al. (2009) implement the BNS test on 2, 5, 10 and 30-year bond prices, a cojump occurs when more than one bond price jumps on the day. The occurrence of cojumps is related to scheduled macroeconomic news announcements. Dungey and Hvozdyk (2011) employ the Jacod and Todorov (2009) cojump test on spot and futures U.S. Treasury contracts and they find that cojumps occur more frequently when the bond contracts have shorter maturities.

Lahaye et al. (2011) employ both ABD and LM tests to detect jumps and cojumps from three asset classes including stock index futures, US Treasury bond futures, and four foreign exchange rates and link them to U.S. macroeconomic news releases. They find that exchange rates and equities have frequent but small jumps, while bond prices have relatively large jumps.

Gnabo et al. (2014) extend both the LM jump test and the mcp statistic of Bollerslev et al. (2008) to a bivariate setting, and implement on S&P 500 futures, 30-year US Treasury bond futures and USD/JPY exchange rate to identify cojumps between bond and index futures, and between bond and F/X rate. In contrast to previous research, they find a positive correlation between stocks and bonds as the majority of cojumps occur for returns with the same sign. The frequency of cojumps is relatively stable but increases during the crisis period.

Gilder et al. (2014) use BNS, ABD and LM tests to detect jumps from 60 liquid US stocks and the Spyder ETF (SPY), and then use a coexceedance criterion to detect cojumps. They argue that this method has similar power to the mcp statistic in Bollerslev et al. (2008), but it cannot detect all cojumps. They also find the coexceedance based detection methods can detect systematic cojumps in the index and the underlying individual stocks. They also present evidence for an association between the systematic cojumps and the Federal Funds Target Rate announcements.

2.6 Macroeconomic news announcements

Intraday volatility and the occurrence of jumps tend to be related to macroeconomic news announcements. The average level of volatility is not constant but depends on the time of the day and has a significant intraday variation. The U-shaped curve starts high, comes down in the middle of the day and then goes up again as the day ends. Volatility also increases substantially around the times of important scheduled macroeconomic news announcements.

Evans (2011) investigates the statistically significant intraday jumps in S&P500 E-Mini, 10-Year US Treasury Bond and EUR/USD futures markets and their relation to US macroeconomic news announcements. Evans and Speight (2011) explore the association between 5-minute EUR/USD, EUR/GBP and EUR/JPY exchange rates and the scheduled macroeconomic news releases in US, UK, Japan and Eurozone.

2.7 Summary of nonparametric jump evidence

With the help of various variation measures, a range of jump detection test methods have been proposed. However, many jump identification techniques only find jumps during a minority of days. There is also a possibility that a lot of small jumps in asset prices are not detected as quite a number of methods may only have power to find large jumps. The methods have been utilised on the index as well as the individual firms, leading to the research into cojumps. Macroeconomic news announcements also have an influence on the occurrence and size of the jumps.

3. Price jumps and cojumps in the foreign exchange market

3.1 Introduction

This chapter will detect jumps in foreign exchange rates and search for jumps which appear simultaneously in many rates. Explanations for jumps, such as macroeconomic news announcements, will be explored. A deeper understanding of jumps has practical implications for high-frequency traders, designers of trading systems and risk managers. Although there are now many results for U.S. equity indices, there are few results so far for foreign exchange data.

There is a long-running debate whether continuous time processes for asset prices contain a jump component generated by a compound Poisson process besides a diffusion component driven by a Brownian motion process. A significant amount of literature, including Duffie et al. (2000), Pan (2002), Eraker et al. (2003) and Eraker (2004), has argued from low-frequency evidence that the jump component should be included.

Following this, many nonparametric tests have been proposed to detect the occurrence of jumps using high-frequency data. The breakthrough work of Barndorff-Nielsen and Shephard (2004a, 2006) employed a method which separates the realised variance measure into a continuous component and a jump component. Subsequently, many other nonparametric tests have been proposed, which includes Jiang and Oomen (2008), Corsi et al. (2010), Podolskij and Ziggel (2010) and Andersen et al. (2012). However, all these tests only detect on which day the jumps occur, but cannot tell the exact timing of the presence of the jumps. Andersen et al. (2007) and Lee and Mykland (2008) develop tests, on the other hand, which can detect the occurrence of jumps at the intraday level.

The Andersen et al. (2007) jump test can identify the exact timing of intraday jumps, hence it is particularly helpful for investigating the effects of macroeconomic news announcements on jumps and cojumps. The Lee and Mykland (2008) jump detection technique possesses a similar property and is thus employed here and the number and timing of jumps detected under the two methods are compared. Scheduled macroeconomic news announcements are among the most important factors that may cause foreign exchange rate jumps.

The existence of jumps in high-frequency prices, typically recorded at least once every five minutes, has been established in several recent papers, including Huang and Tauchen (2005), Andersen et al. (2007), Lee and Mykland (2008), Lee and Hannig (2010) and Evans (2011). These papers have studied the equity market and all authors identify some returns which are too large to be explained by a diffusion process; the typical frequency of these large returns is one every two weeks. Lahaye et al. (2011), Dungey et al. (2009) and Dungey and Hvozdyk (2012) further present some evidence of the occurrence of jumps in foreign exchange and Treasury bond markets.

However, not so many papers have investigated cojumps, which are simultaneous jumps in the prices of two or more assets. Dungey et al. (2009) and Dungey and Hvozdyk (2012) explore the U.S. Treasury market, Lahaye et al. (2011) investigate the U.S. equity indices, U.S. Treasury bond index and the dollar exchange rates, while Gilder et al. (2014) study the S&P 500 index and 60 of its constituent firms. All these studies explore the timing of cojumps relative to macroeconomic news announcements, while similar macro investigations for jumps can be found in Andersen et al. (2003, 2007).

In this study, the non-parametric ABD and LM tests are employed to extract jumps and cojumps from foreign exchange rates. In our proposed implementation, the tests compare one-minute returns with critical values dependent on a significance level, daily measures of price variation calculated from bipower and estimates of the intraday volatility pattern described in Taylor (2005). As dollar rates move together, and since cross rates are constrained by no-arbitrage equations, we anticipate that when one rate jumps many others will also jump. We will document the number, sign and size of foreign exchange jumps, and will compare the number of cojumps with expectations derived from no-arbitrage principles.

The chapter is organised as follows. Section 3.2 lays the theoretical foundation regarding the ABD and LM jump detection test methods which are employed subsequently. Section 3.3 describes the data. Section 3.4 presents empirical analysis of ten foreign exchange rates. Section 3.5 summarises the findings and concludes.

3.2 Methods

3.2.1 Variation measures

A general jump diffusion process has price dynamics similar to that of equation (2.1) in Chapter 2. Similar definitions also apply to integrated variance IV_t, quadratic variation QV_t, realised variance $RV_t(N)$ and bipower variation $BV_t(N)$ in equations (2.2) to (2.4) and (2.6) in Chapter 2.

3.2.2 ABD jump detection test

Andersen et al. (2007) have developed the test method that identifies which return displays significant evidence that there is a jump. The ABD test uses the bipower variation to estimate daily volatility under the null hypothesis of a diffusion process. We use a modified jump detection test which takes account of the intraday volatility pattern to count the number of detected jumps. A jump is detected if

$$|r_{t,j,N}| > c_N \sqrt{s_{j,N}^2 BV_t(N)}, \tag{3.1}$$

with c_N determined by the significance level of the hypothesis test. Hence we do not allocate the total variation $BV_t(N)$ of one day t equally to every short period of time, but allocate more total variation to those parts of the day with high volatility, where $s_{j,N}^2$ is the average fraction of the sample variation of all returns in period j divided by the sample variance of all intraday returns. If the variation is relatively high during some part of the day, then $s_{j,N}^2$ is greater than one, and we will get a large number on the right hand side of the equation and it may be harder to find evidence of a jump.

3.2.3 LM jump detection test

The LM test first defines the statistic $L_{t,i}$, which equals the return on day *t* during the intraday period *i*. The quantity $L_{t,i}$ scaled by an estimate of volatility for one intraday period in time increment *i* is given as

$$L_{t,i} = \frac{r_{t,i}}{\hat{\sigma}_{LM,t,i}} \tag{3.2}$$

where $r_{t,i}$ is the intraday return on day *t* in increment *i*, and the variation $\hat{\sigma}^2_{LM,t,i}$ is defined as

$$\hat{\sigma}^2_{LM,t,i} = \frac{1}{K-2} \sum_{j=i-K+2}^{i-1} |r_{t,j}|\,|r_{t,j-1}| \tag{3.3}$$

for some window size *K*. Since the variation $\hat{\sigma}^2_{LM,t,i}$ is used for the instantaneous volatility estimate in the denominator of the test statistic, the method employed is thus robust to the occurrence of jumps in prior intraday periods. We neglect the drift part as we use the high-frequency data in our study, and the drift part is mathematically negligible relative to the diffusion and the jump component. The window size *K* is determined in a way that the jumps have no effect on the volatility estimation. Lee and Mykland (2008) suggest that $K = \sqrt{M \times 252}$, where *M* is the number of increments per day and $252M$ is the number of observations in one year.

We next focus on the selection of the rejection region. The test statistics will present different behaviour depending on the presence of jumps during the testing time. On the one hand, if there is no jump from the test, the test statistics may follow an approximate normal distribution. On the other hand, if there are jumps, the test statistic will be very large, and the sample maximum converges to a Gumbel distribution. Although it is abnormal that we compare test statistics with critical values whose distribution is not normal, we simply follow Lee and Mykland (2008). We then need to decide how large the test statistic could be when there is no jump. An investigation of the asymptotic distribution of maximums of the test statistics with no jumps in increment *i* shows that a jump is detected if

$$|L_{t,i}| > \frac{\varepsilon}{c\sqrt{2\ln M}} + \frac{\sqrt{2\ln M}}{c} - \frac{(\ln 4\pi + \ln(\ln M))}{2c\sqrt{2\ln M}} \tag{3.4}$$

where $c = \sqrt{2/\pi}$, $\varepsilon = -\ln(-\ln(1-\alpha))$, and α represents the daily significance level. In other words, the criterion to choose a rejection region is that if the test statistics are not in the usual region of maximums for a set of M intraday returns, it is unlikely that the return comes from the diffusion part of the jump diffusion model.

3.3 Data

High-frequency exchange rates data are obtained from forex tester website.[3] We acquire the foreign exchange rates at the one-minute frequency for ten currency pairs, four of which are dollar rates for Euro, Pound, Yen and Swiss Franc while the other six are cross-rates. The data covers the time period from January 2007 to December 2011.

[3] www.forextester.com

Our time stamps are all for London time, which is Greenwich Mean Time (GMT) in the winter and British Summer Time (BST) in the summer. Throughout this chapter all times are for London. GMT is the local time from 1st January to the morning of the last Sunday of March, and then from the last Sunday of October to 31st December, while BST lasts from the last Sunday of March to the morning of the last Sunday of October. Foreign exchange rates operate in a 24-hour market. Our datasets contain records from 23:00 on Sunday until 21:00 on Friday, while there are no records from 21:00 on Friday, the whole of Saturday and until 23:00 on Sunday, which is also consistent with the definition of weekends in Andersen and Bollerslev (1997) and Taylor (2005). Hence we define the time period from 23:00 on Sunday to 23:00 on Monday as our 'Monday'. Similar time slot arrangements also apply for Tuesday, Wednesday and Thursday. We will then define the time from 23:00 on Thursday to 21:00 on Friday as our 'Friday'.

We delete all price records on a day when there are more than 20 missing consecutive prices, and we fill up the missing minute's price as the previous minute's price if there are fewer than 20 missing records in a day. This is a standard method and we assume no price changes and zero returns when there are missing data, because usually this happens when there is no trading. Consequently there are 1278 days in our dataset, of which 1020 days are days from Monday to Thursday, and 258 days are Fridays. The detailed list of the days that are eliminated from 2007-2011 are displayed in Table 3.1. The days deleted are usually close to holidays such as New Year's Day, Easter, Independence Day, Thanksgiving Day and Christmas. We notice that year 2010 has many more deleted days than others, which is probably due to the quality of the data (i.e., people did not record data properly). We calculate the return as the change in log

Table 3.1

List of eliminated days from 2007 to 2011.

2007				
24th December	31st December			
2008				
24th December	31st December			
2009				
24th December	31st December			
2010				
21st September	4th October	5th October	6th October	7th October
21st October	25th November	30th November	27th December	28th December
29th December	30th December			
2011				
3rd July	11th September			

prices, hence we have 1440 intraday returns from Monday to Thursday when there are 24 trading hours per day (note there are no overnight returns as the foreign exchange market trades continuously 24 hours a day) and 1320 returns on Friday when we have 22 trading hours on Friday. Finally we also calculate the variance proportions separately for Monday to Thursday and Friday as there are different numbers of intraday returns for the two categories.

3.4 Empirical analysis

3.4.1 Empirical properties of returns

The mean, standard deviation, number and proportion of zero returns of the ten foreign exchange rates from 2007 to 2011 are presented in Table 3.2. In terms of returns, EUR/JPY and GBP/JPY exchange rates have relatively larger average magnitude of returns compared to the remaining eight foreign exchange rates, while EUR/GBP, EUR/USD and USD/CHF have the smallest mean returns of the ten foreign exchange rates data series. Only EUR/GBP has positive mean return over the sample period, whereas all the remaining nine foreign exchange rates have negative average return over the same period. In terms of standard deviation, EUR/JPY, GBP/JPY and CHF/JPY have comparatively larger standard deviation than the remaining seven foreign exchange rates, while EUR/GBP, EUR/USD, GBP/USD and EUR/CHF have the lowest standard deviation across all the foreign exchange rates data series. EUR/GBP has the largest number and proportion of zero returns, followed by USD/JPY, USD/CHF and EUR/CHF, which have more than 30% zero returns. All the remaining six foreign exchange rates have around 18% to 29% zero returns.

Table 3.2

Statistical properties of the exchange rate returns from 2007 to 2011.

2007	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	1.97E-07	2.35E-07	3.99E-08	1.76E-07	-2.06E-08	-6.11E-08	6.84E-08	1.03E-07	-1.31E-07	-9.10E-08
Standard deviation	0.000115	0.000108	0.000114	0.000173	0.000186	0.000166	0.000178	0.000091	0.000130	0.000123
No. 0 returns	219648	165726	136356	94014	68031	141405	128400	137966	161927	84828
% 0 returns	60.37%	45.55%	37.48%	25.84%	18.70%	38.86%	35.29%	37.92%	44.51%	23.31%
2008	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	6.83E-07	-1.69E-07	-8.49E-07	-6.22E-07	-1.31E-06	-4.55E-07	-4.02E-07	-2.20E-07	-5.28E-08	-9.04E-07
Standard deviation	0.000226	0.000232	0.000245	0.000355	0.000381	0.000290	0.000333	0.000175	0.000249	0.000263
No. 0 returns	138005	95158	92626	60031	46342	98772	85594	94217	108014	56125
% 0 returns	37.92%	26.15%	25.45%	16.49%	12.73%	27.14%	23.52%	25.89%	29.68%	15.42%
2009	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	-2.96E-07	2.93E-08	3.25E-07	1.17E-07	4.14E-07	8.98E-08	1.21E-07	-9.18E-09	-2.88E-08	2.97E-07
Standard deviation	0.000222	0.000220	0.000263	0.000305	0.000352	0.000240	0.000304	0.000143	0.000229	0.000246
No. 0 returns	118871	82916	76825	58829	50457	109078	85471	114110	103825	63304
% 0 returns	32.66%	22.78%	21.11%	16.16%	13.86%	29.97%	23.48%	31.35%	28.53%	17.39%

2010	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	-3.08E-07	-3.96E-07	-8.97E-08	-7.14E-07	-4.24E-07	-3.16E-07	-2.17E-07	-4.94E-07	-9.68E-08	-1.85E-07
Standard deviation	0.000179	0.000199	0.000195	0.000275	0.000278	0.000199	0.000272	0.000166	0.000209	0.000210
No. 0 returns	131233	90350	88697	72895	65238	130913	95523	104433	110614	67593
% 0 returns	37.40%	25.75%	25.28%	20.77%	18.59%	37.31%	27.22%	29.76%	31.52%	19.26%
2011	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	-7.69E-08	-3.22E-08	4.34E-08	-1.97E-07	-1.18E-07	-1.69E-07	-1.37E-07	-6.10E-08	-3.06E-08	1.82E-08
Standard deviation	0.000169	0.000201	0.000162	0.000244	0.000222	0.000179	0.000267	0.000250	0.000254	0.000254
No. 0 returns	139305	88142	99959	87180	86805	174169	103058	93547	123935	76064
% 0 returns	37.99%	24.04%	27.26%	23.77%	23.67%	47.49%	28.10%	25.51%	33.80%	20.74%
Across 5 years	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Mean	3.97E-08	-6.65E-08	-1.06E-07	-2.48E-07	-2.91E-07	-1.82E-07	-1.14E-07	-1.36E-07	-6.80E-08	-1.73E-07
Standard deviation	0.000182	0.000192	0.000196	0.000271	0.000284	0.000215	0.000271	0.000165	0.000214	0.000219
No. 0 returns	747062	522292	494463	372949	316873	654337	498046	544273	608315	347914
% 0 returns	41.27%	28.85%	27.31%	20.61%	17.51%	36.16%	27.52%	30.09%	33.61%	19.23%

One possible explanation for this might be that the tick sizes of foreign exchange rates are typically small but are large relative to one-minute standard deviations. There may be a large number of zero returns in the foreign exchange rates data series even when trades do occur in almost all minutes.

In terms of individual years between 2007 and 2011, GBP/USD and GBP/CHF have relatively small average magnitude of returns in 2007, 2010 and 2011, but big average magnitude of returns in 2008 and 2009. GBP/JPY has comparatively large mean returns from 2008 to 2011, while small average returns in 2007. On the contrary, USD/CHF has relatively small average returns from 2008 to 2011, but large mean returns in 2007. Only EUR/GBP has positive returns in 2008 and all foreign exchange rates have negative returns in 2010. Regarding standard deviation, GBP/JPY has the largest standard deviation between 2007 and 2010 and CHF/JPY has the largest standard deviation in 2011, while EUR/CHF has the smallest standard deviation from 2007 to 2010 and GBP/USD has the smallest standard deviation in 2011. EUR/GBP has the largest number and proportion of zero returns from 2007 to 2010 and USD/JPY has the largest number and proportion of zero returns in 2011, while GBP/JPY has the smallest number and proportion of zero returns between 2007 and 2010 and GBP/CHF has the smallest number and proportion of zero returns in 2011.

If we look across different years, years 2008 and 2009 have more volatile returns and larger standard deviation in general, while the year 2007 is the least volatile and has the smallest standard deviation. This corresponds to the world financial crisis which happened from 2008 to 2009. Years 2007 and 2011 have a larger number and proportion of zero returns.

3.4.2 Detection of jumps

Table 3.3 presents the number of jumps detected by the ABD test under the two significance levels 10^{-5} and 10^{-3} between 2007 and 2011. These are daily significance levels and equal the expected Type I errors in one day from testing all the one-minute returns. For the total number of jumps across five years, inevitably more jumps are detected under the 10^{-3} significance level than under the 10^{-5} significance level for all ten foreign exchange rates. The number of jumps found under the 10^{-3} level is between 2.1 and 2.3 times those identified under the 10^{-5} level for each foreign exchange rate. The EUR/GBP exchange rate has the smallest number of jumps detected under the 10^{-5} significance level while CHF/JPY has the least number of jumps detected under the 10^{-3} level. The USD/JPY rate has the largest number of jumps detected under both significance levels. The four Swiss Franc exchange rates CHF/JPY, EUR/CHF, USD/CHF and GBP/CHF have quite large numbers of jumps identified under both significance levels in 2011 compared to the remaining six foreign exchange rates. One possible reason for this will be discussed in section 3.4.5.

Figure 3.1 shows that for the period from 2007 to 2011, GBP/USD, USD/JPY and USD/CHF have comparatively larger number of jumps detected, while EUR/GBP and CHF/JPY have relatively smaller number of jumps detected. If we look in terms of years, we detect the largest number of jumps in 2011 while the smallest number of jumps is in 2009. The average jump detection rate, across all currency pairs and all years, is between one and two jumps per day for the 10^{-5} significance level and three jumps per day for the 10^{-3} level.

Table 3.3

Number of jumps detected from the ABD test under two significance levels between 2007 and 2011.

2007	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	249	511	683	296	325	391	218	222	494	272
10^(-3)	658	1086	1379	675	697	885	517	484	1016	611
2008	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	253	395	476	310	286	359	210	301	349	258
10^(-3)	599	859	1077	638	658	784	496	683	758	557
2009	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	271	299	365	199	249	361	186	301	393	305
10^(-3)	613	669	827	458	533	770	448	667	819	681
2010	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	255	330	302	268	249	430	250	371	378	280
10^(-3)	544	688	666	585	534	891	565	838	782	621
2011	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	254	378	329	339	306	763	419	683	655	485
10^(-3)	569	747	710	664	702	1583	884	1382	1301	1022
All years	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
10^(-5)	1282	1913	2155	1412	1415	2304	1283	1878	2269	1600
10^(-3)	2983	4049	4659	3020	3124	4913	2910	4054	4676	3492

Figure 3.1 Number of jumps detected from the ABD test under the 10^{-5} significance level from 2007 to 2011.

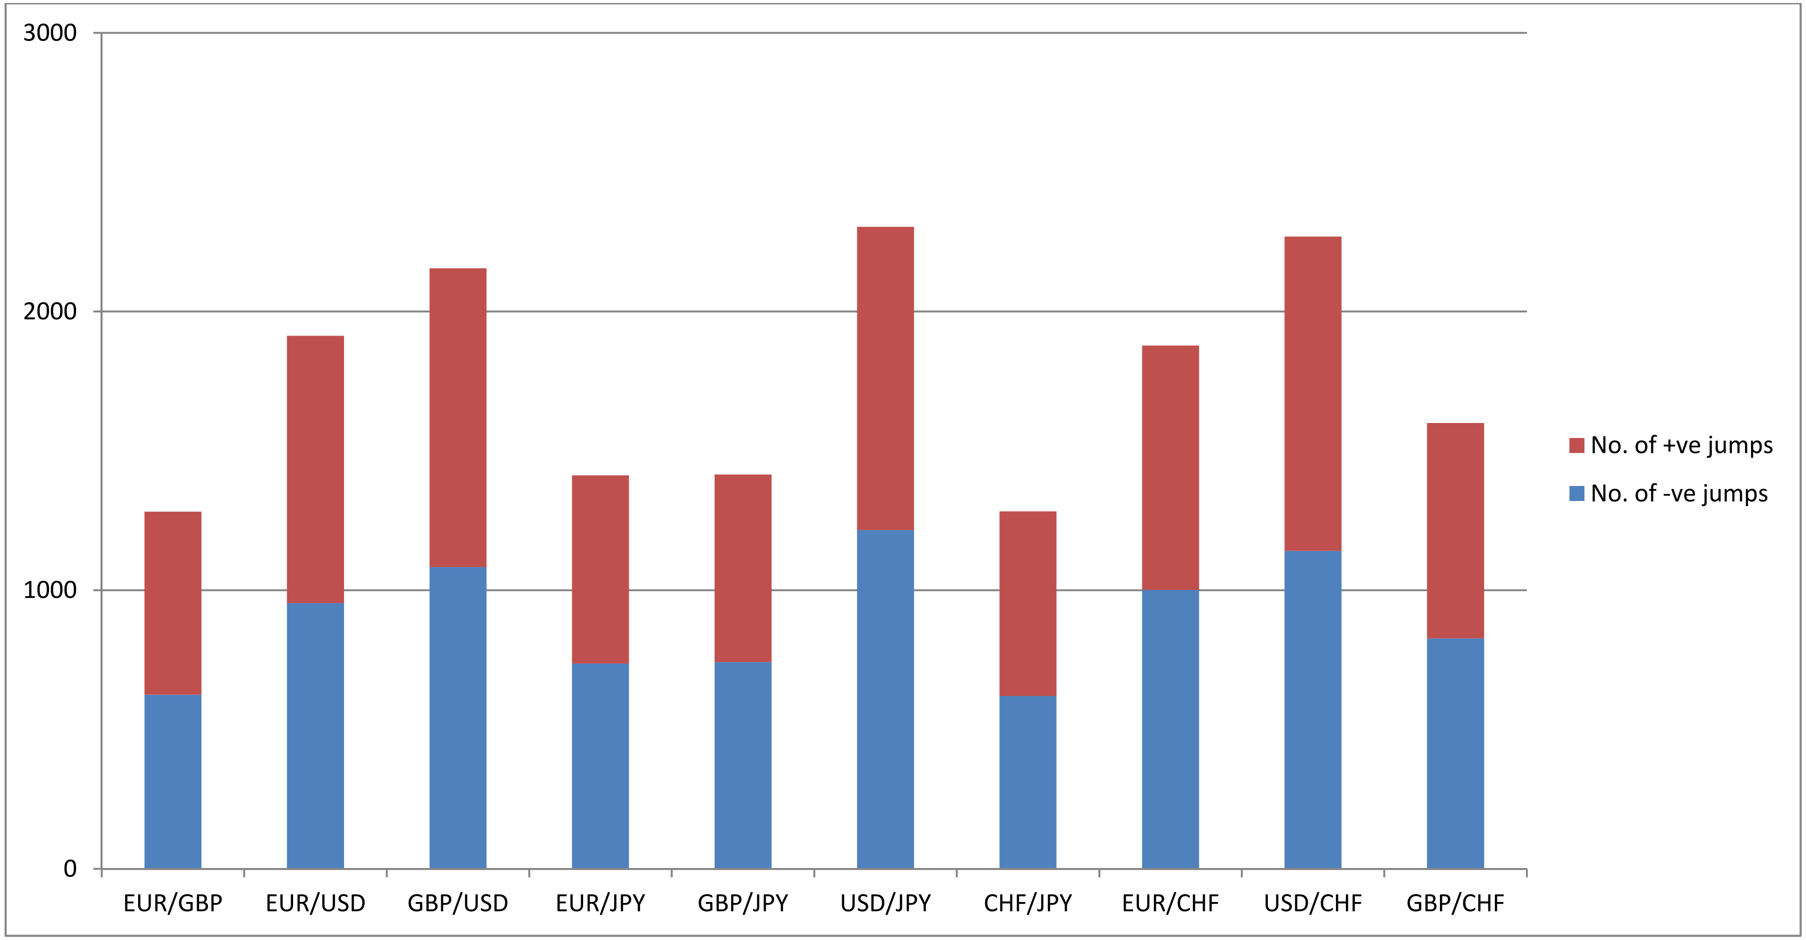

3.4.3 Statistical properties of the detected jumps

Table 3.4 details the statistical properties of the detected jumps from the ABD test for ten foreign exchange rates under the 10^{-5} significance level from 2007 to 2011. All ten foreign exchange rates have similar average magnitudes for negative and positive jumps. EUR/CHF has the smallest magnitude of mean for both negative and positive jumps (-0.000946 and 0.000934 respectively), while CHF/JPY has the largest average magnitude for negative and positive jumps (-0.001691 and 0.001640 respectively). EUR/GBP, EUR/USD and GBP/USD also have relatively smaller magnitude of average positive and negative jumps, while EUR/JPY and GBP/JPY have comparatively larger average magnitude of positive and negative jumps. This may imply that the US dollar, British pound, euro and Swiss Franc are in general more stable than the Japanese yen. GBP/USD has the smallest magnitude of the minimum negative jump (-0.008506), while CHF/JPY has the abnormal largest magnitude of minimum negative jump (-0.028183). GBP/USD has the smallest magnitude of maximum positive jump (0.007709), while EUR/USD, EUR/JPY, CHF/JPY and USD/CHF all have abnormally large magnitude of maximum jumps (0.016717, 0.017403, 0.017510 and 0.013088 respectively).

In terms of individual years from 2007 to 2011, nine and eight foreign exchange rates have slightly more negative jumps than positive jumps in 2008 and 2010, which might be related to the occurrence and aftermath effect of the world financial crisis. The years 2008, 2009 and 2011 have more large jumps with absolute value greater than 0.01 (EUR/JPY -0.010748, GBP/JPY -0.011301, EUR/JPY 0.011424 and CHF/JPY 0.010632 in 2008, USD/JPY -0.015189, USD/CHF -0.016422, GBP/CHF -0.010249,

Table 3.4

Statistical properties of jumps detected from the ABD test under the 10^{-5} significance level from 2007 to 2011.

2007	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	249	511	683	296	325	391	218	222	494	272
No. -ve	123	235	333	147	167	209	110	101	228	130
Minimum	-0.005249	-0.002721	-0.002412	-0.003902	-0.003848	-0.003736	-0.003807	-0.001623	-0.002787	-0.003204
Mean -ve	-0.000680	-0.000608	-0.000585	-0.000951	-0.001048	-0.000895	-0.001082	-0.000621	-0.000678	-0.000786
No. +ve	126	276	350	149	158	182	108	121	266	142
Maximum	0.003449	0.003153	0.006788	0.003652	0.005982	0.004209	0.002535	0.001360	0.002184	0.005489
Mean +ve	0.000737	0.000606	0.000579	0.000910	0.001066	0.000914	0.001102	0.000573	0.000691	0.000733

2008	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	253	395	476	310	286	359	210	301	349	258
No. -ve	127	228	254	167	165	188	108	167	172	136
Minimum	-0.006592	-0.004342	-0.005534	-0.010748	-0.011301	-0.007105	-0.008828	-0.003562	-0.004910	-0.005068
Mean -ve	-0.001236	-0.001168	-0.000955	-0.001470	-0.001653	-0.001367	-0.001532	-0.000953	-0.001410	-0.001357
No. +ve	126	167	222	143	121	171	102	134	177	122
Maximum	0.004955	0.005779	0.005563	0.011424	0.009503	0.005818	0.010632	0.002565	0.005317	0.004926
Mean +ve	0.001225	0.001183	0.000942	0.001505	0.001749	0.001488	0.001611	0.000907	0.001338	0.001395

2009	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	271	299	365	199	249	361	186	301	393	305
No. -ve	135	142	188	116	137	198	102	163	189	152
Minimum	-0.004529	-0.005041	-0.006771	-0.007344	-0.008869	-0.015189	-0.006928	-0.004345	-0.016422	-0.010249
Mean -ve	-0.001208	-0.001302	-0.001394	-0.001973	-0.002350	-0.001703	-0.002248	-0.000938	-0.001364	-0.001570
No. +ve	136	157	177	83	112	163	84	138	204	153
Maximum	0.010276	0.016717	0.007709	0.007069	0.006812	0.005962	0.006312	0.006150	0.007818	0.006290
Mean +ve	0.001513	0.001273	0.001331	0.001882	0.001943	0.001449	0.001930	0.000998	0.001444	0.001499

2010	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	255	330	302	268	249	430	250	371	378	280
No. -ve	120	175	151	154	125	228	129	186	193	153
Minimum	-0.009664	-0.009585	-0.005461	-0.016712	-0.006969	-0.007148	-0.009911	-0.006962	-0.004744	-0.004926
Mean -ve	-0.001254	-0.001395	-0.001339	-0.001654	-0.001804	-0.001271	-0.001732	-0.001072	-0.001371	-0.001327
No. +ve	135	155	151	114	124	202	121	185	185	127
Maximum	0.004055	0.009152	0.005086	0.006110	0.006162	0.008057	0.004875	0.008546	0.008730	0.008909
Mean +ve	0.001264	0.001369	0.001146	0.001685	0.001783	0.001296	0.001768	0.001107	0.001475	0.001542

2011	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	254	378	329	339	306	763	419	683	655	485
No. -ve	120	174	157	153	148	393	171	384	359	256
Minimum	-0.003941	-0.004667	-0.008506	-0.005004	-0.008558	-0.007394	-0.028183	-0.017507	-0.010195	-0.010149
Mean -ve	-0.001113	-0.001207	-0.000921	-0.001500	-0.001388	-0.000940	-0.001863	-0.001148	-0.001427	-0.001468
No. +ve	134	204	172	186	158	370	248	299	296	229
Maximum	0.008538	0.008963	0.005200	0.017403	0.006306	0.007108	0.017510	0.011319	0.013088	0.006091
Mean +ve	0.001258	0.001298	0.000909	0.001767	0.001621	0.001032	0.001790	0.001083	0.001306	0.001328

Across 5 years	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
Total no.	1282	1913	2155	1412	1415	2304	1283	1878	2269	1600
No. -ve	625	954	1083	737	742	1216	620	1001	1141	827
Minimum	-0.009664	-0.009585	-0.008506	-0.016712	-0.011301	-0.015189	-0.028183	-0.017507	-0.016422	-0.010249
Mean -ve	-0.001098	-0.001136	-0.001039	-0.001510	-0.001648	-0.001235	-0.001691	-0.000946	-0.001250	-0.001302
No. +ve	657	959	1072	675	673	1088	663	877	1128	773
Maximum	0.010276	0.016717	0.007709	0.017403	0.009503	0.008057	0.017510	0.011319	0.013088	0.008909
Mean +ve	0.001199	0.001146	0.000982	0.001550	0.001633	0.001236	0.001640	0.000934	0.001251	0.001299

EUR/GBP 0.010276, EUR/USD 0.016717 in 2009, CHF/JPY -0.028183, EUR/CHF -0.017507, USD/CHF -0.010195, GBP/CHF -0.010149, EUR/JPY 0.017403, CHF/JPY 0.017510, EUR/CHF 0.011319, USD/CHF 0.013088 in 2011) than year 2010 (EUR/JPY -0.016712), while the size of the jumps in 2007 is more stable. Again this may indicate that the foreign exchange rates fluctuate more in 2008 and 2009 when the financial tsunami occurred, while the market is more stable in 2007 and 2010.

3.4.4 Cojumps between the dollar rates and the corresponding cross rate

We consider it as a cojump when there are two or three corresponding jumps at the same time among the two dollar exchange rates and the parallel cross rate. Table 3.5 details the number, percentage and conditional probabilities of cojumps among the two dollar exchange rates EUR/USD, GBP/USD, and the corresponding cross rate EUR/GBP under the 10^{-5} significance level in 2011, detected from the ABD test. We consider the null hypothesis which states that the jumps are independent. If the null hypothesis is true, the probability of cojumps should be very near zero as it is just the product of the probability of jumps for the two dollar rates and/or the parallel cross rate. These individual probabilities are close to zero from previous discussions. The results in Table 3.5 lead us to reject the null hypothesis that jumps are independent, as the counts of cojumps are positive for all combinations.

We investigate the conditional probabilities of cojumps given the occurrence of jumps in either the dollar rate or the corresponding cross rate to examine the jump dependence. The dollar rate EUR/USD and the corresponding cross rate EUR/GBP

Table 3.5

Descriptive statistics of counts of EUR/GBP, EUR/USD and GBP/USD cojumps from 2007 to 2011, found using the ABD test with significance level 10^{-5}. $\mathrm{P(coj|jump)}$ is defined as follows, for example, $\mathrm{P(EUR/GBP - EUR/USD\ cojump|EUR/GBP\ jump)}$=counts of EUR/GBP and EUR/USD cojumps/counts of EUR/GBP jumps=387/1282=30.19%.

				P(coj|jump) (%)		
	No. of obs.	No. of coj	P(coj)(%)	EUR/GBP	EUR/USD	GBP/USD
EUR/GBP-EUR/USD	1809360	387	0.0214	30.19	20.23	
EUR/GBP-GBP/USD	1809360	519	0.0287	40.48		24.08
EUR/USD-GBP/USD	1809360	351	0.0194		18.35	16.29
EUR/GBP-EUR/USD-GBP/USD	1809360	74	0.0041	5.77	3.87	3.43

have the largest number of cojumps (387), which amounts to 0.02% of the 1809360 total intraday return observations. The other dollar rate GBP/USD and the parallel cross rate EUR/GBP, and the two dollar rates EUR/USD and GBP/USD have similar number of cojumps (519 and 351 respectively), amounting to 0.03% and 0.02% of the total intraday return observations respectively. There are 74 cojumps which occur at the same time among the two dollar rates EUR/USD, GBP/USD and the parallel cross rate EUR/GBP, which is 0.0041% of the total number of intraday return observations.

The conditional probabilities of a cojump in the two dollar rates or the cross rate given there is a jump in another one of these three rates range from 16.29% to 40.48%. The probabilities are the highest for the dollar rate GBP/USD and the cross rate EUR/GBP combination (40.48% and 24.08%) and the lowest for the two dollar rates EUR/USD and GBP/USD combination (18.35% and 16.29%). This may imply that the euro and the dollar are more closely related compared to the euro and the pound, as the dollar rate GBP/USD and the cross rate EUR/GBP tend to cojump more often than the two dollar rates EUR/USD and GBP/USD given there is a jump in the two dollar rates or the cross rate. The conditional probabilities of cojumps among the two dollar rates EUR/USD, GBP/USD and the cross rate EUR/GBP, given there is a jump in one of the two dollar rates or the cross rate is much lower, ranging from 3.43% to 5.77%.

The detailed cojumps in 2011 between (i) the dollar rate EUR/USD and the cross rate EUR/GBP, (ii) the dollar rate GBP/USD and the cross rate EUR/GBP, (iii) the two dollar rates EUR/USD and GBP/USD, and (iv) simultaneously among the two dollar rates and the cross rate have been investigated. The dates and the times that the cojumps occur seem to have no particular pattern. The cojumps that occur at the same time between the dollar rate EUR/USD and the cross rate EUR/GBP always have the same sign, and the two dollar rates EUR/USD and GBP/USD combination also have the same property that the cojumps are either both positive or both negative. However the dollar rate GBP/USD and the cross rate EUR/GBP have many occasions when the signs of the cojumps are opposite. For the cojumps among the two dollar rates EUR/USD, GBP/USD and the cross rate EUR/GBP, the signs of the three are always the same.

One last thing to notice is that the sum of the cross rate EUR/GBP cojump and the dollar rate GBP/USD cojump approximately equals the other corresponding dollar rate EUR/USD cojump. This is because we measure the return as the change in log prices, hence the sum of the cross rate EUR/GBP and the dollar rate GBP/USD log returns is equal to the log of the product of the two returns, and we cancel out the GBP in the product to have the dollar rate EUR/USD log return, and the equality establishes. However it is possible that the cojumps of the dollar rate GBP/USD log return and the cross rate EUR/GBP log return have opposite signs, as long as on the other side of the equality, the dollar rate EUR/USD log return has the same sign as but smaller magnitude than the dollar rate GBP/USD log return, and the equality still holds.

Figures 3.2 to 3.7 show the number of cojumps among the two dollar rates and the corresponding cross rate for six foreign exchange groups over the five-year period from 2007 to 2011 using the ABD test and the daily 10^{-5} significance level. Across five years, the EUR, CHF and USD group has the largest number of cojumps, while the USD, CHF and JPY group has the smallest number of cojumps. Between 2007

Figure 3.2: Counts of EUR/GBP, EUR/USD and GBP/USD cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

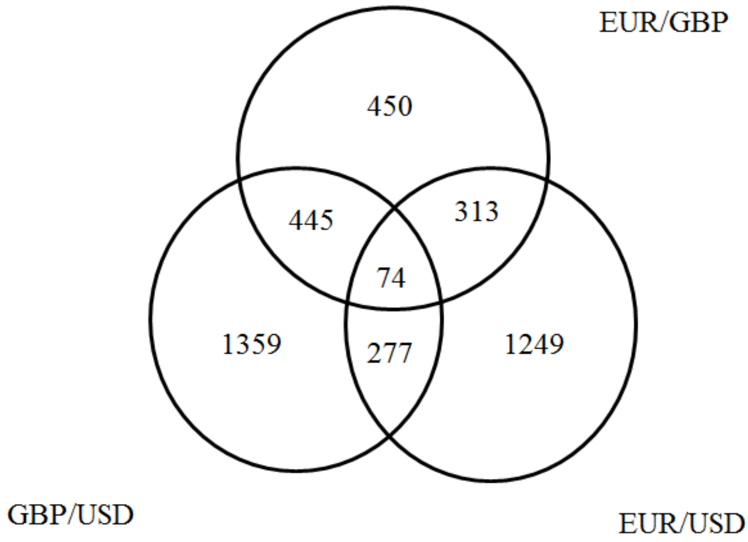

Figure 3.3 Counts of EUR/JPY, EUR/USD and USD/JPY cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

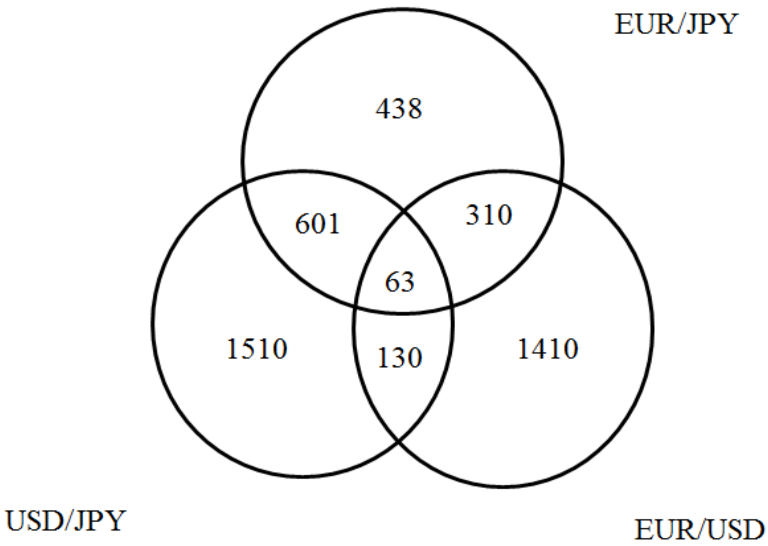

Figure 3.4 Counts of GBP/JPY, GBP/USD and USD/JPY cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

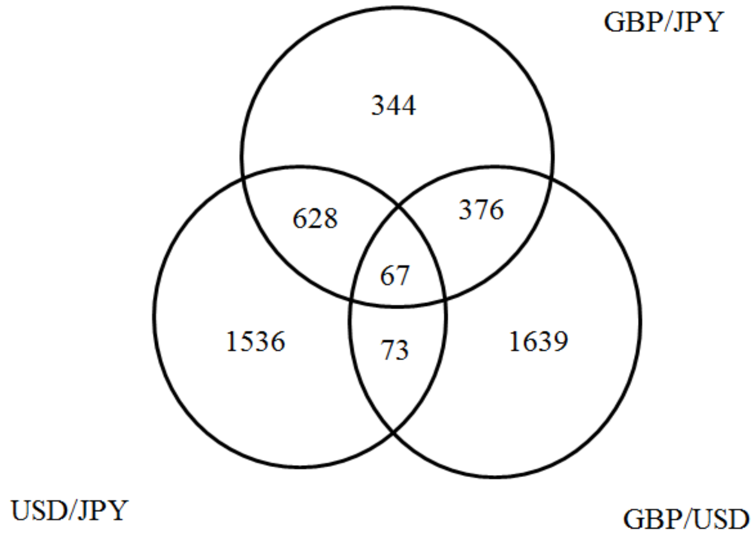

Figure 3.5 Counts of CHF/JPY, USD/JPY and USD/CHF cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

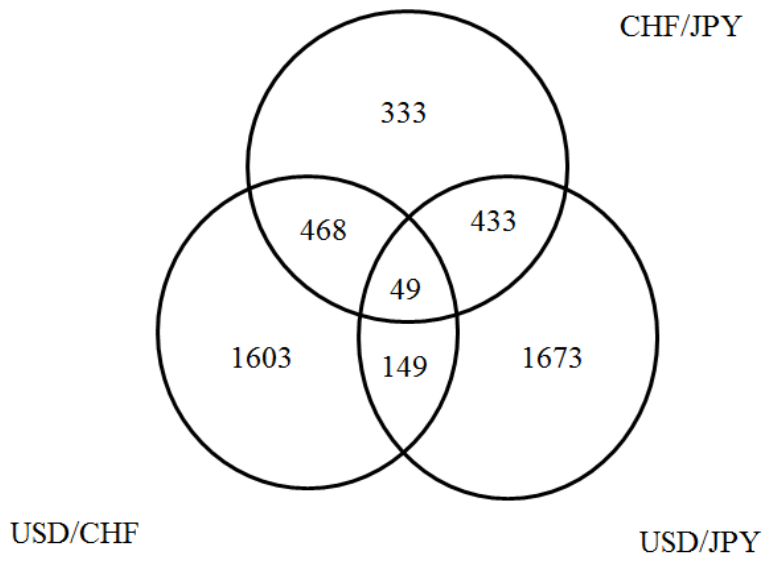

Figure 3.6 Counts of EUR/CHF, EUR/USD and USD/CHF cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

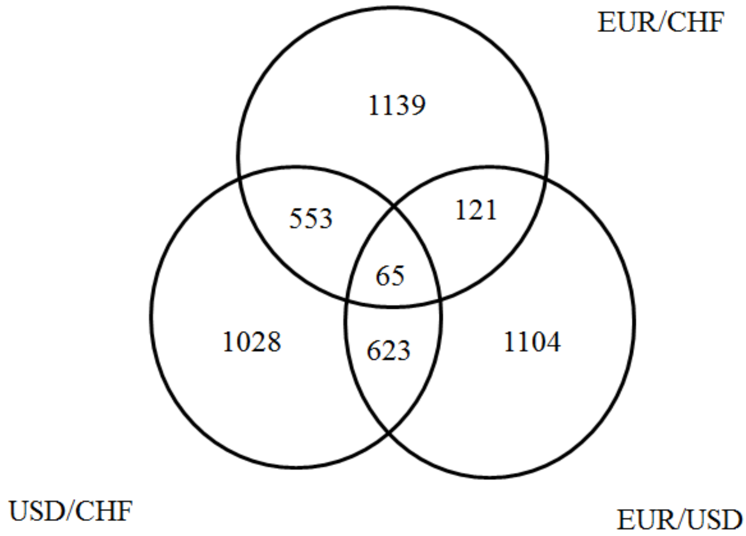

Figure 3.7 Counts of GBP/CHF, GBP/USD and USD/CHF cojumps from 2007 to 2011, ABD test, daily significance level 10^{-5}.

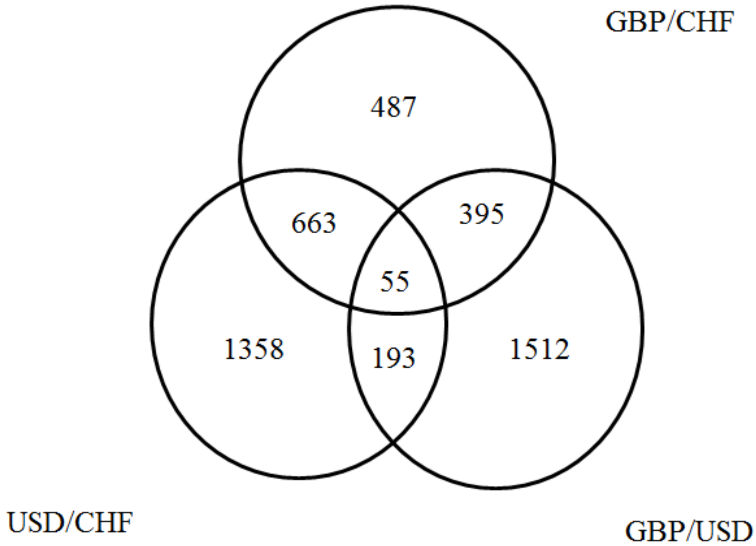

and 2009, the two dollar rates USD/CHF, EUR/USD and the cross rate EUR/CHF group almost always has the largest number of cojumps, while the two dollar rates USD/CHF, USD/JPY and the cross rate CHF/JPY and the EUR, JPY and USD group always have the least number of cojumps. For the period from 2010 to 2011, the EUR, CHF and USD group has the largest number of cojumps while the two dollar rates GBP/USD, EUR/USD and the cross rate EUR/GBP group and the GBP, USD and JPY group have the smallest number of cojumps.

Across five years, one dollar rate and the cross rate combination nearly always has more cojumps than the two dollar rates combination; only the EUR, CHF and USD group is different, either dollar rate and the cross rate combination always has fewer cojumps than the two dollar rates combination. For example, Figure 3.2 shows 277 cojumps for GBP/USD and EUR/USD, less than 313 for the cross rate and EUR/USD and 445 for the cross rate and GBP/USD. Some exceptions happen when we investigate each year separately, such as the EUR, GBP and USD group (EUR/GBP and EUR/USD has 38 (29) cojumps while EUR/USD and GBP/USD has 80 (58) cojumps in 2007 (2009), EUR/GBP and GBP/USD has 44 cojumps while GBP/USD and EUR/USD has 48 cojumps in 2011). For the different EUR, CHF and USD group, EUR/CHF and USD/CHF (231) has more cojumps than EUR/USD and USD/CHF (79) in 2011. If we look across years, year 2011 has comparatively more cojumps than the other four years. One possible explanation for this is given in section 3.4.5.

Table 3.6 summarises the times and dates when eight or nine foreign exchange rates cojump together from 2007 to 2011. There are eleven times when eight foreign exchange rates cojump together and four times when nine rates cojump together during this time period. We note that on 18th September 2007, eight foreign exchange rates cojump at the same time; and on the same day, Federal Reserve lowered target on key short-term rate for the first time in four years due to the mortgage crisis. Also on 3rd November 2010, nine foreign exchange rates cojump together; and on this day, Federal Reserve announced to pump billions of dollars to simulate the economy.

3.4.5 Plot of returns and detected jumps and some extreme returns

Plots of the returns and the detected jumps of the two dollar rates EUR/USD, GBP/USD and the cross rate EUR/GBP in 2011 under the 10^{-5} significance level are shown on Figure 3.8 to Figure 3.13. The typical size of a jump is higher towards the middle and the end of this year than the early part of the year. Jumps also tend to cluster in size through time and particularly in the middle and the end of the year. We also notice that some large positive or negative returns in a relative sense are not detected as jumps because there is more random variation during that intraday period and the period is a highly volatile period. Some detected jumps are small in magnitude because volatility is small at that time thus the returns need not be large for the jumps to be detected.

Four Japanese yen exchange rates, EUR/JPY, CHF/JPY, USD/JPY and GBP/JPY all have a large return on 31st October 2011 (0.017403, 0.017510, 0.020036 and 0.020213 respectively). The first two returns are detected as jumps while the latter two not. This maybe because all the returns for USD/JPY and GBP/JPY exchange rates are large on this day, it is hard for a large return to be detected as a jump. An

Table 3.6

Summary of dates and times when eight or nine foreign exchange rates cojump together from 2007 to 2011.

Date	Time	No. rates involved	List of foreign exchange rates cojump
18/09/2007	18:16	8	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, EUR/CHF, USD/CHF, GBP/CHF
22/01/2008	13:21	8	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, USD/CHF
04/09/2008	21:18	8	EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF
18/07/2008	10:31	8	EUR/GBP, EUR/USD, EUR/JPY, GBP/JPY, USD/JPY, EUR/CHF, USD/CHF, GBP/CHF
16/12/2009	19:17	8	EUR/USD, GBP/USD, EUR/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF, GBP/CHF
27/09/2009	23:05	9	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, USD/CHF, GBP/CHF
01/03/2010	11:40	8	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, CHF/JPY, USD/CHF, GBP/CHF
03/05/2010	0:40	8	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, CHF/JPY, USD/CHF, GBP/CHF
21/06/2010	0:18	8	EUR/GBP, EUR/USD, GBP/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF, GBP/CHF
10/08/2010	18:16	8	EUR/GBP, EUR/USD, GBP/USD, GBP/JPY, USD/JPY, CHF/JPY, USD/CHF, GBP/CHF
30/11/2010	23:01	8	EUR/GBP, EUR/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF
09/02/2010	17:42	9	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF
01/11/2010	0:01	9	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, USD/CHF, GBP/CHF
03/11/2010	18:17	9	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, EUR/CHF, USD/CHF
15/09/2011	13:01	8	EUR/GBP, EUR/USD, GBP/USD, EUR/JPY, GBP/JPY, USD/JPY, CHF/JPY, USD/CHF

Figure 3.8 Intraday EUR/GBP returns over one minute intervals in 2011.

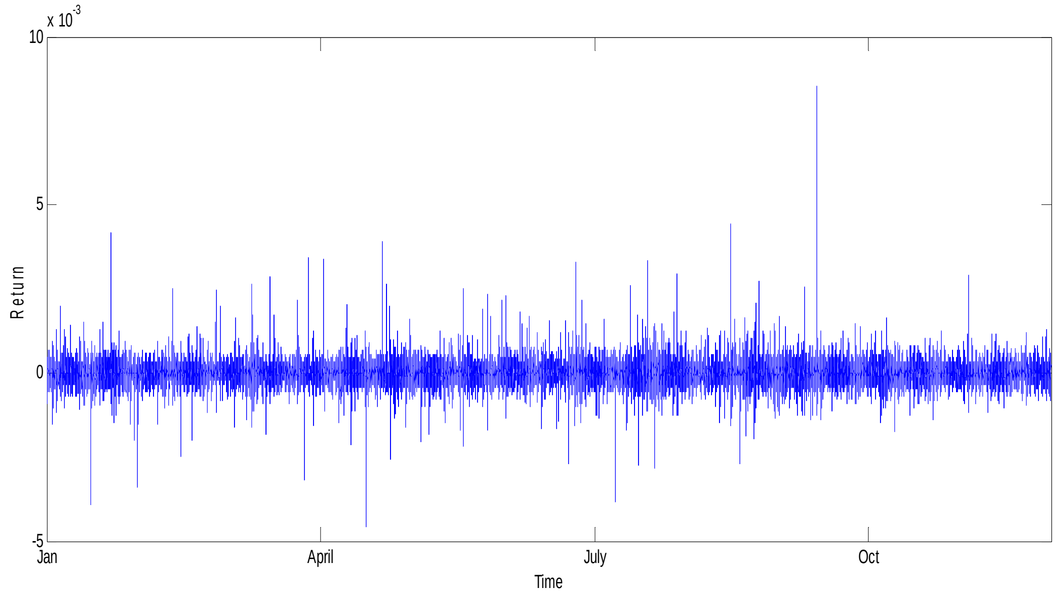

Figure 3.9 EUR/GBP detected jumps in 2011, using the ABD test and the daily 10^{-5} significance level.

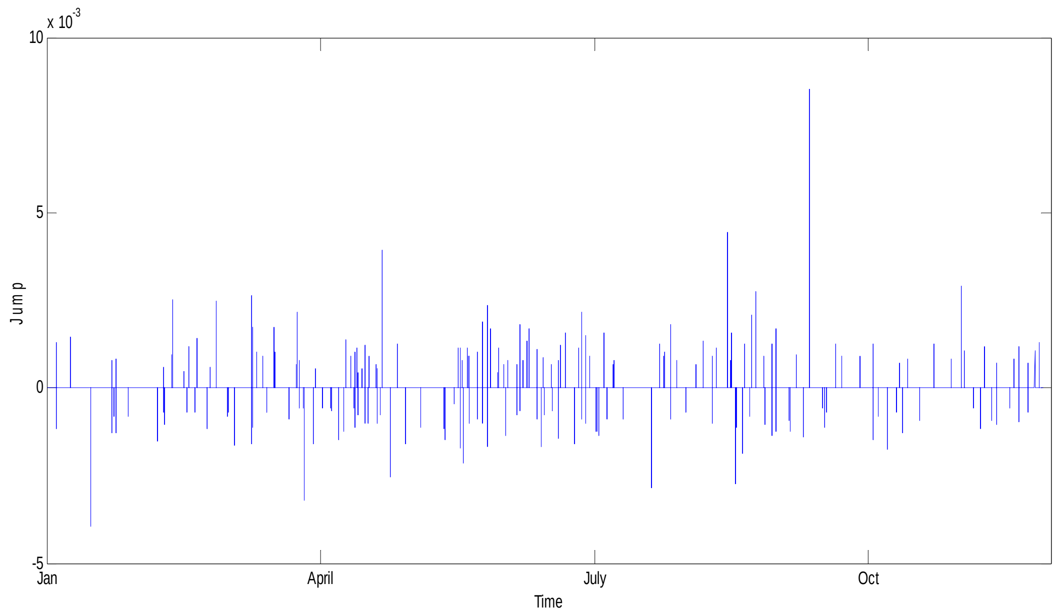

Figure 3.10 Intraday EUR/USD returns over one minute intervals in 2011.

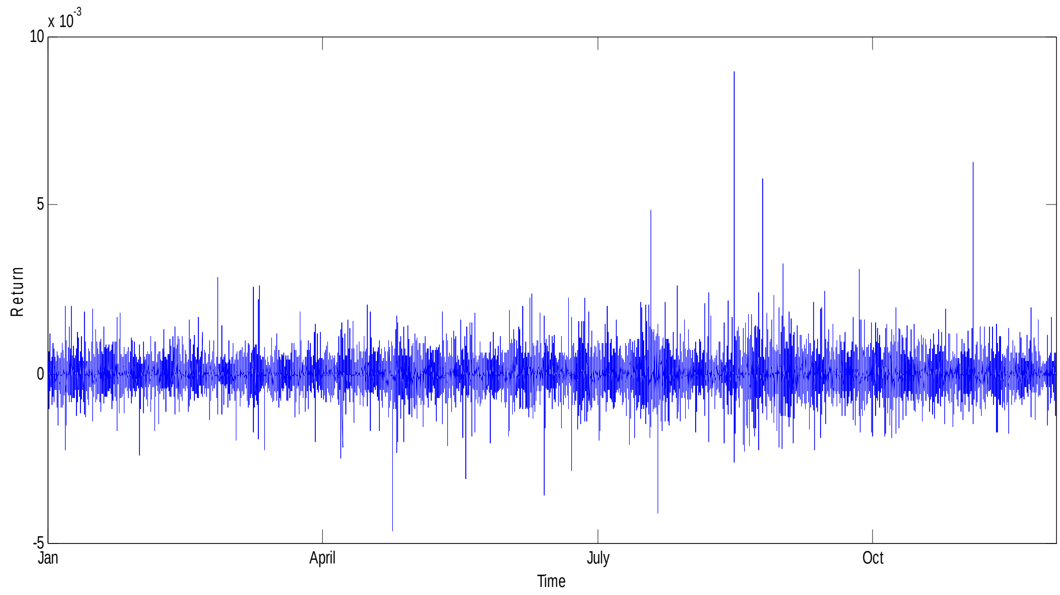

Figure 3.11 EUR/USD detected jumps in 2011, using the ABD test and the daily 10^{-5} significance level.

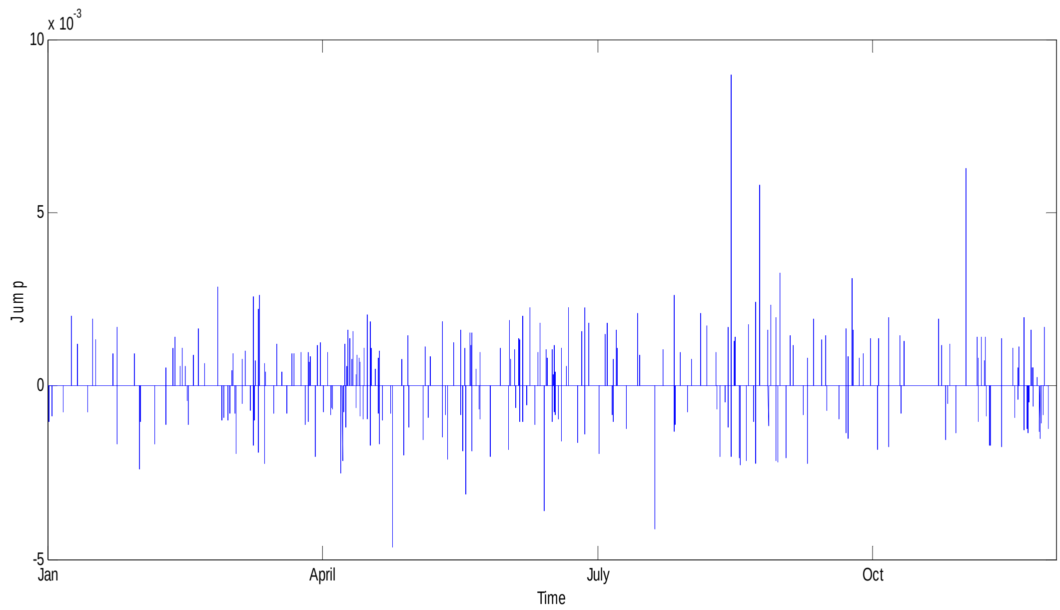

Figure 3.12 Intraday GBP/USD returns over one minute intervals in 2011.

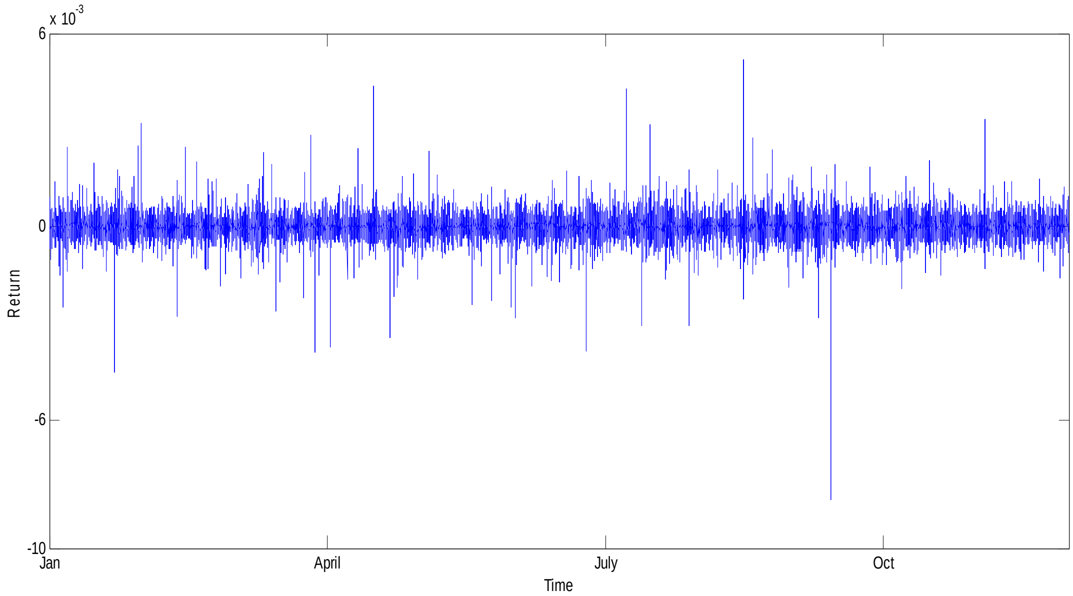

Figure 3.13 GBP/USD detected jumps in 2011, using the ABD test and the daily 10^{-5} significance level.

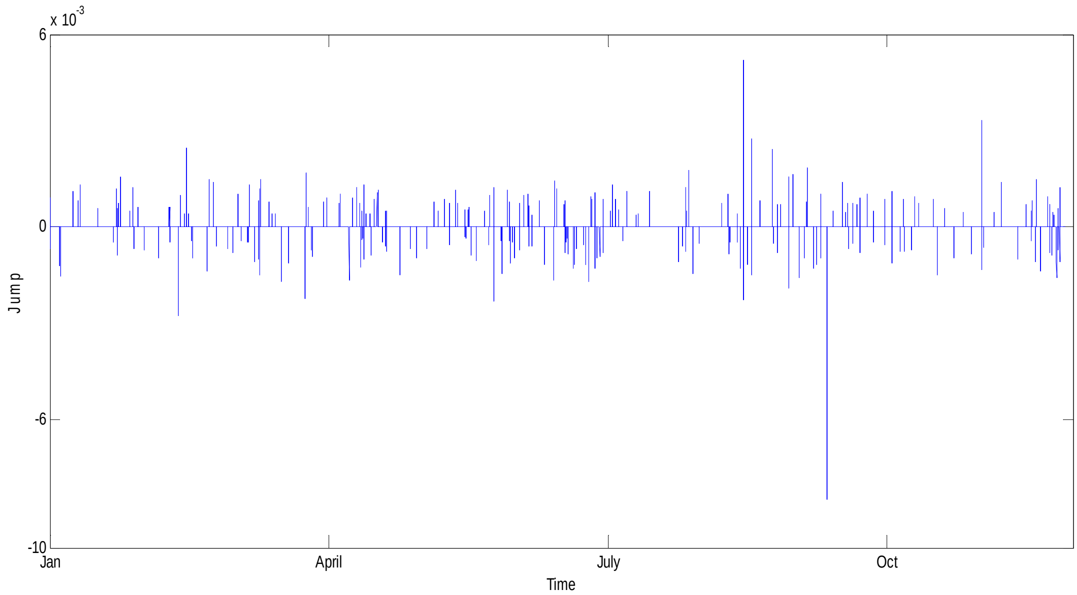

investigation reveals that this arises from the Japanese government intervention to weaken the yen against the dollar rate after it has reached the highest level since World War II. The Japanese government fears that the strengthened yen will impact the country's economy as it relies heavily on exports. The country is still slowly recovering from the destructive consequences of the 11th March earthquake and tsunami so that such appreciation cannot be tolerated.

Four Swiss Franc rates CHF/JPY, EUR/CHF, USD/CHF and GBP/CHF all have an unusually large jump on 20th September 2011 (0.010627, -0.017507, -0.010195 and -0.010149 respectively). An investigation shows that this is because the Swiss National Bank announced to set a minimum level of the Swiss Franc against the Euro. The Swiss Franc dropped by 9% against other currencies in fifteen minutes as the Swiss National Bank did not allow one Swiss Franc to be worth more than 0.83 Euro. The reason for this action is that the investors considered the Swiss Franc a haven in the European debt crisis, and the Swiss companies worried that the exporters of the country will be less competitive in the market abroad. The action from the Swiss National Bank further depreciated the Swiss Franc in subsequent weeks and months and investors also tried to look for alternative investments. This incident may also help to explain the cluster in the jump size and the frequency of jumps for the three Swiss Franc exchange rates towards the end of the year.

3.4.6 Jumps and macroeconomic news announcements

Figures 3.14, 3.16 and 3.18 present the number of jumps against the time of day for the two dollar rates EUR/USD, GBP/USD and the cross rate EUR/GBP. We choose

Figure 3.14 Counts of EUR/GBP jumps against time, using the ABD test and the daily 10^{-5} significance level.

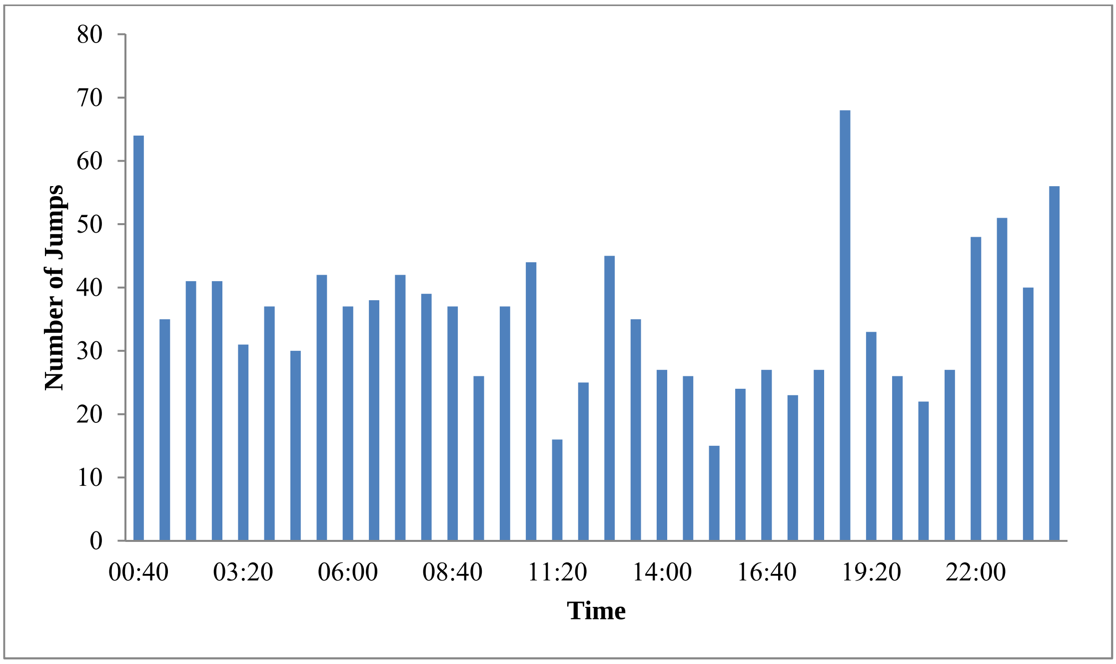

Figure 3.15 EUR/GBP variance proportion plot from 2007 to 2011.

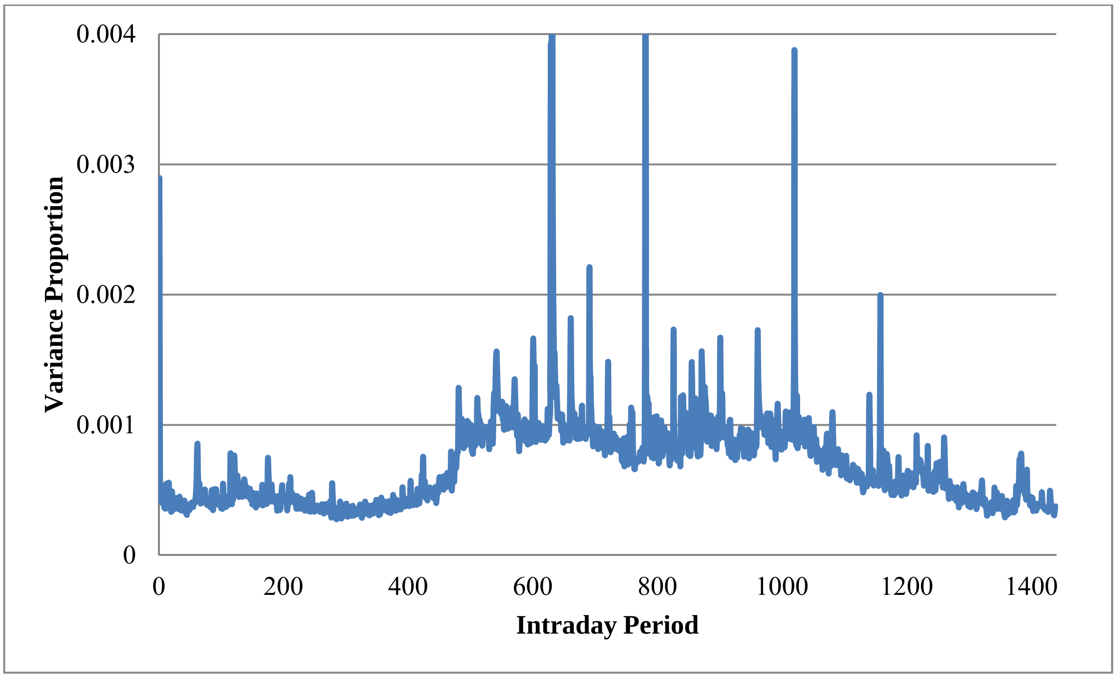

Figure 3.16 Counts of EUR/USD jumps against time, using the ABD test and the daily 10^{-5} significance level.

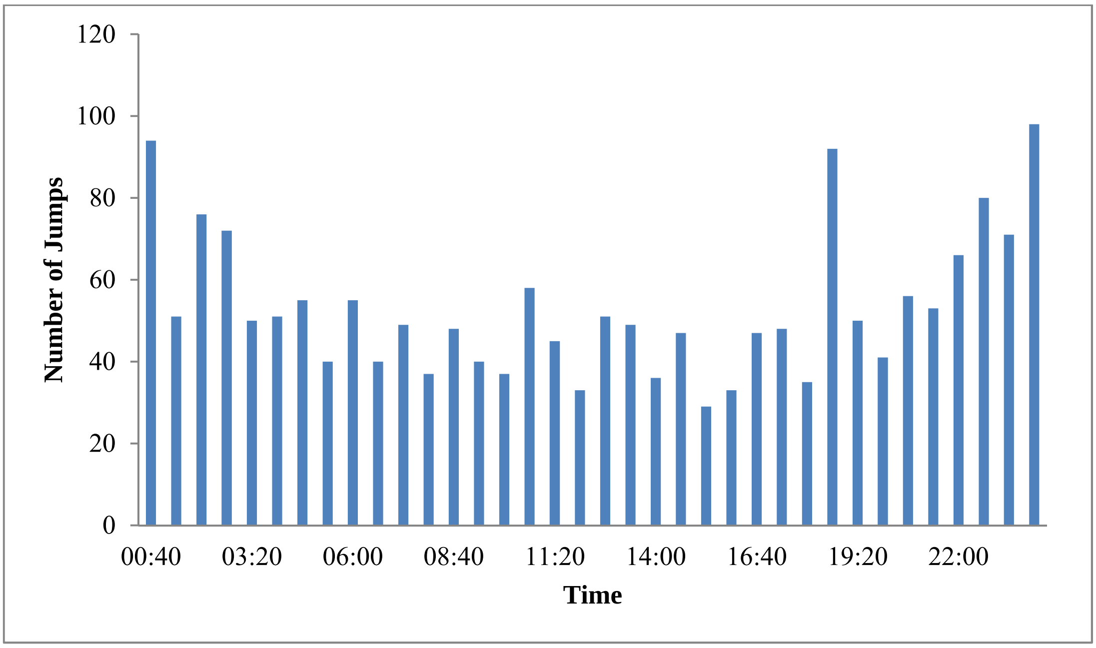

Figure 3.17 EUR/USD variance proportion plot from 2007 to 2011.

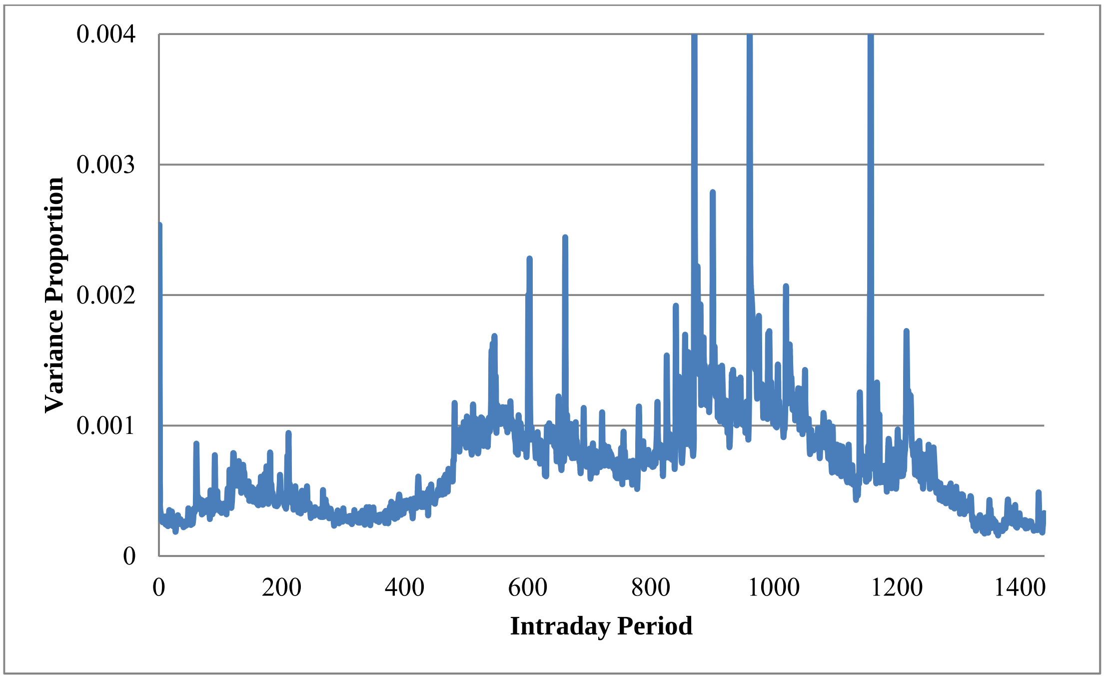

Figure 3.18 Counts of GBP/USD jumps against time, using the ABD test and the daily 10^{-5} significance level.

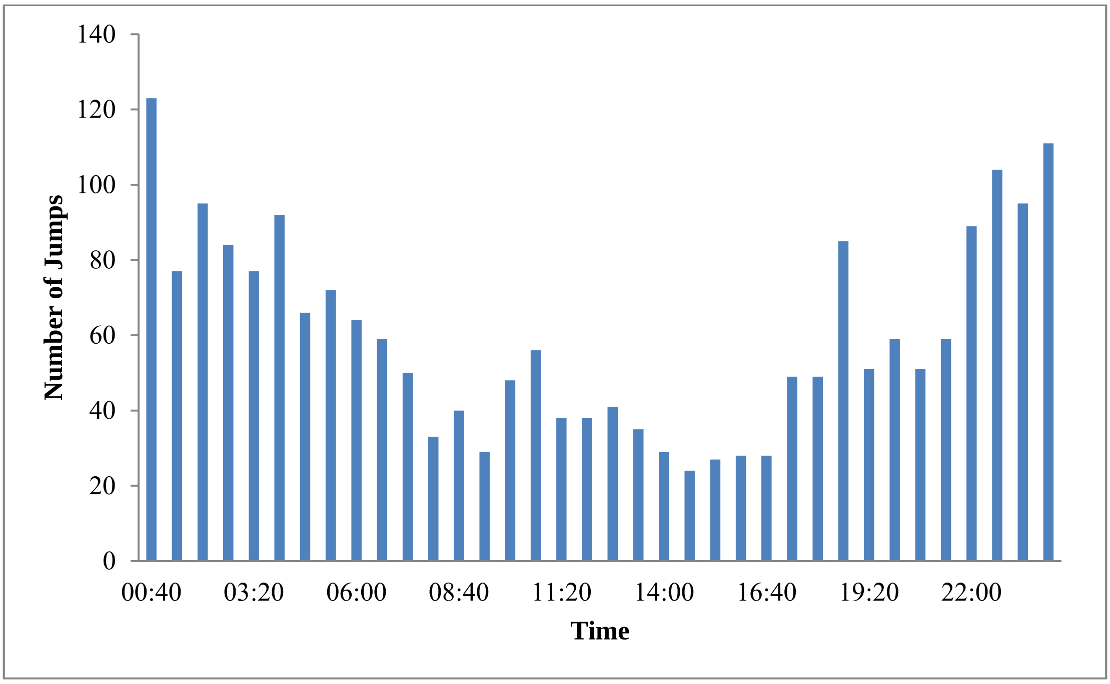

Figure 3.19 GBP/USD variance proportion plot from 2007 to 2011.

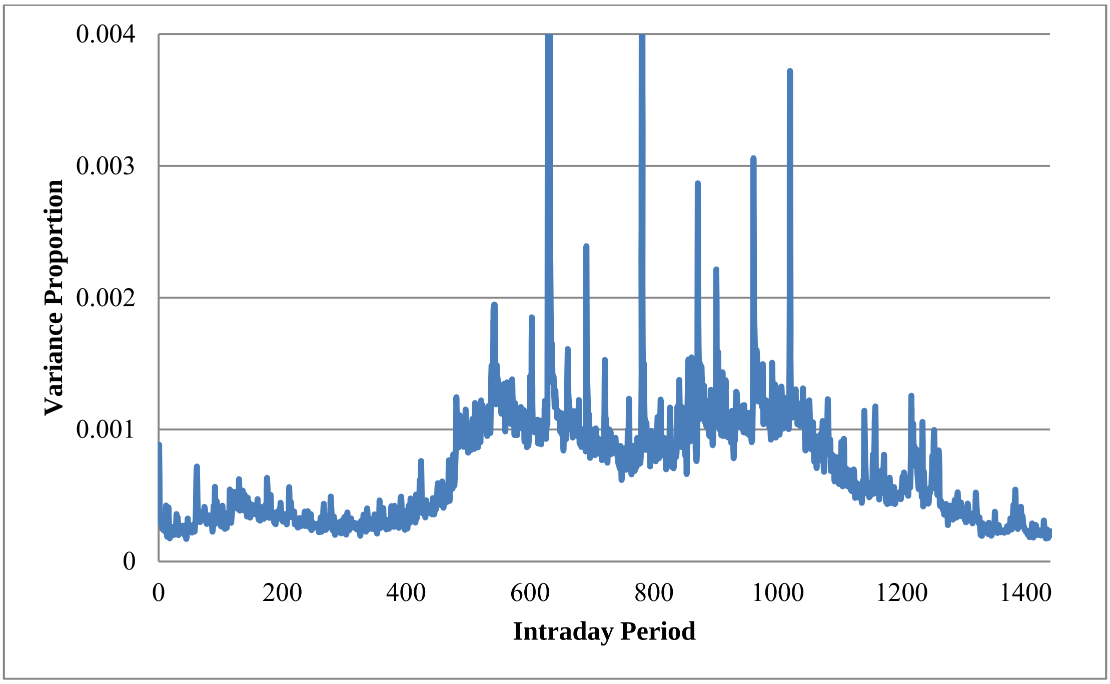

the 40-minute interval, and given the 24-hour foreign exchange market there are 36 intervals per day. We aggregate all the jumps in the same time interval, but on different days from 2007 to 2011, and then try to link them to the time of macroeconomic news announcements. A glimpse at the figures indicates that the two dollar rates EUR/USD and GBP/USD have a spike around 19:00, which may correspond to macroeconomic news announcements in the U.S. afternoon.[4] The four Japanese yen rates EUR/JPY, GBP/JPY, USD/JPY and CHF/JPY are all detected for a large number of jumps around GMT midnight, which may be related to the early morning news announcements in Japan and other Asian countries.[5] The intraday volatility plots for the two dollar rates EUR/USD, GBP/USD and the cross rate EUR/GBP are shown alongside the detected jump pattern in Figures 3.15, 3.17 and 3.19. In general, the intraday volatility is low at the start and the end of the day, but high during the middle of the day. There are more jumps detected when the volatility is low but fewer jumps detected when the volatility is high.

3.4.7 Cojumps and macroeconomic news announcements

Figures 3.20 to 3.23 plot the number of cojumps against the time between the dollar rate EUR/USD and the cross rate EUR/GBP, the dollar rate GBP/USD and the cross rate EUR/GBP, two dollar rates EUR/USD and GBP/USD, and among the two dollar rates and the cross rate. We choose the 20-minute interval, given the 24-hour foreign exchange market, there are 72 intervals per day. We aggregate all the cojumps in the same time interval but on different days from 2007 to 2011, and then try to link them to the time of macroeconomic news announcements. The spike of cojumps is similar

[4] U.S. announces government fiscal surplus or deficit at Eastern Standard Time (EST) 14:00 monthly and federal funds target at EST 14:15 every six weeks.

[5] Japan announces macroeconomic news at 8:50 local time.

to that of jumps. There are spikes at approximately 19:00 and midnight. The 19:00 spike may correspond to macroeconomic news announcements in the U.S. afternoon, while the midnight clustering of cojumps can be linked to the early morning news announcements in Japan and other Asian countries. Similar to jumps, there are more cojumps detected when the volatility is low but fewer cojumps detected when the volatility is high.

3.4.8 Comparison between ABD and LM jump detection tests

A comparison between the number of jumps detected under the ABD test and the number of jumps detected under the LM test is provided in Table 3.7. The LM test detects more jumps than the ABD test, and the ratio ranges between 1.26 to 2.01, where GBP/USD has the largest ratio and EUR/CHF has the smallest ratio.

3.5 Conclusions

We investigate one-minute returns of ten foreign exchange rates for five years from 2007 to 2011. We use the ABD and LM jump detection tests to detect intraday price jumps for ten rates and cojumps for six groups of two dollar rates and one cross rate. We reject the null hypothesis that jumps are independent across rates, as there are far more cojumps than predicted by independence for all rate combinations. We also find that some clustering of jumps and cojumps can be related to the macroeconomic news announcements affecting the exchange rates. The chosen ABD and LM jump detection tests find a similar number of jumps for the foreign exchange rates.

Figure 3.20 EUR/GBP and EUR/USD cojumps from 2007 to 2011, using the ABD test and the daily 10^{-5} significance level.

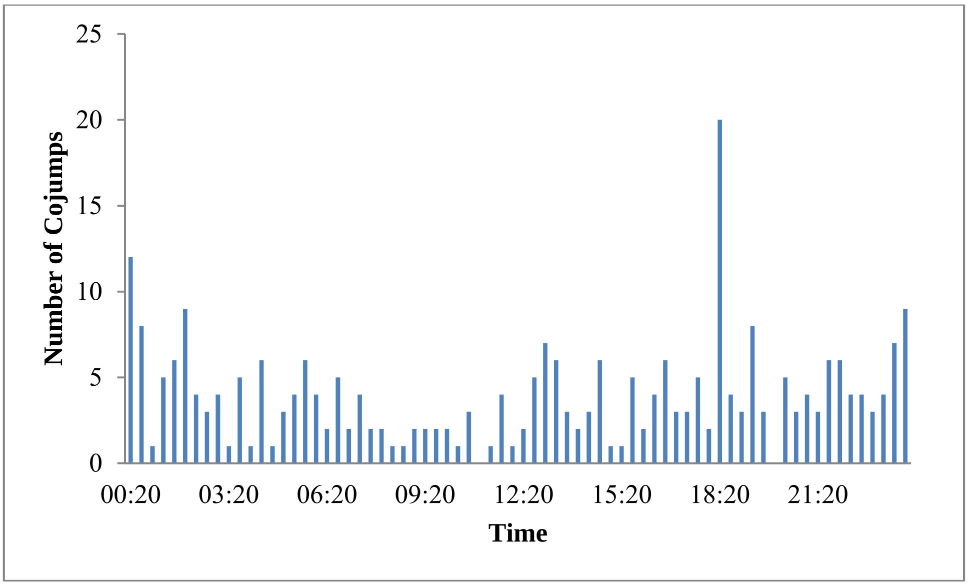

Figure 3.21 EUR/GBP and GBP/USD cojumps from 2007 to 2011, using the ABD test and the daily 10^{-5} significance level.

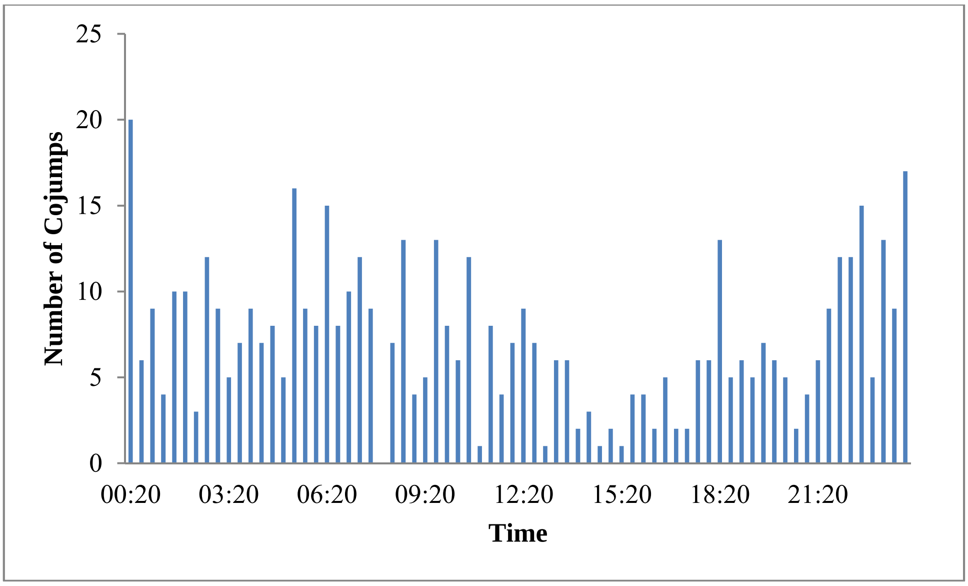

Figure 3.22 EUR/USD and GBP/USD cojumps from 2007 to 2011, using the ABD test and the daily 10^{-5} significance level.

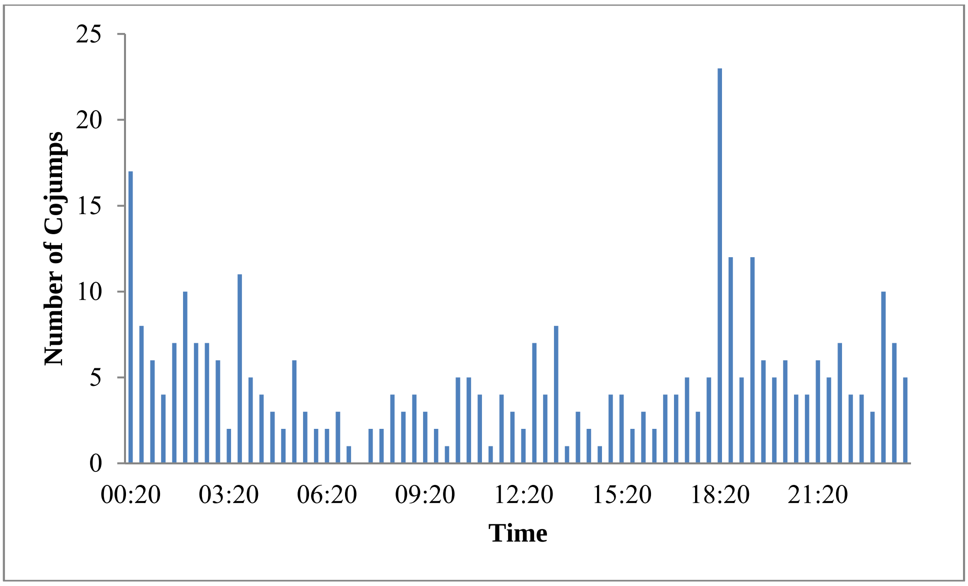

Figure 3.23 EUR/GBP, EUR/USD and GBP/USD cojumps from 2007 to 2011, using the ABD test and the daily 10^{-5} significance level.

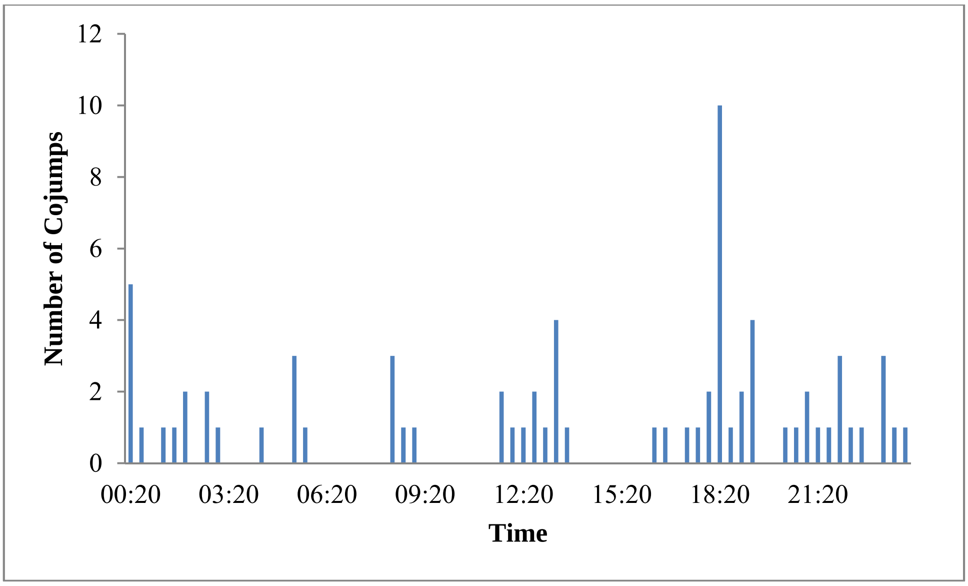

Table 3.7

Number of jumps detected from the ABD and the LM tests under the 10^{-5} level from 2007 to 2011.

2007	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	249	511	683	296	325	391	218	222	494	272
LM	515	1194	1411	398	498	590	356	408	1170	550
2008	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	253	395	476	310	286	359	210	301	349	258
LM	456	746	939	326	451	519	242	388	769	489
2009	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	271	299	365	199	249	361	186	301	393	305
LM	511	464	671	372	442	574	363	368	605	564
2010	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	255	330	302	268	249	430	250	371	378	280
LM	525	544	640	402	442	582	423	547	655	557
2011	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	254	378	329	339	306	763	419	683	655	485
LM	483	575	679	475	517	756	502	647	783	644
Across 5 years	EUR/GBP	EUR/USD	GBP/USD	EUR/JPY	GBP/JPY	USD/JPY	CHF/JPY	EUR/CHF	USD/CHF	GBP/CHF
ABD	1282	1913	2155	1412	1415	2304	1283	1878	2269	1600
LM	2490	3523	4340	1973	2350	3021	1886	2358	3982	2804

Foreign exchange rates have frequent and relatively small jumps as they are usually affected by two sources of news and they have more liquidity shocks during the continuously traded 24-hour market. Some groups of foreign exchange rates jump and cojump more than other groups, this may either be due to some exchange rates are highly correlated, or it is easy to simultaneously trade some exchange rates. For example, the U.S. scheduled macroeconomic news announcements may affect all dollar exchange rates, and some European news may affect both euro and pound exchange rates.

Previous studies including Lahaye et al. (2011) usually only focus on dollar rates, our study investigates more currencies by examining six groups of two dollar rates and one cross rate at the more frequent and informative one minute level for ten years. We find that one dollar rate and the cross rate combination nearly always has more cojumps than the two dollar rates combination.

The limitation of this study is that the chosen test methods may only find jumps during a minority of days when the daily significance level is low. There are possibilities that a lot of small jumps in foreign exchange rates are not detected as the nonparametric jump detection tests can only detect large jumps, where a large jump is large relative to the volatility of the diffusion component of the asset prices. Also the power of these tests in detecting jumps should be compared.

There is always the debate whether continuous time processes for asset prices contain a jump component generated by a compound Poisson process besides a diffusion component driven by a Brownian motion process. Recent theoretical and empirical evidence have confirmed that jumps exist in financial time series data, and it is important to understand their big impact on hedging risks and trading derivatives. For instance, the extreme comovements during the recent financial crisis may have caused large jumps and cojumps in financial asset prices, while some less extreme news may only create small jumps and cojumps. Hence it is important to understand these comovements in the financial market and hedge against their risks. Cojumps contain useful information to understand asset price dynamics, and can help to model and forecast volatility and covariance matrix. Cojumps are also important for risk managers and option traders.

4. Volatility and density forecasting literature

4.1 Volatility forecasting

4.1.1 High-frequency information

High-frequency data records prices more than once during a day, typically at frequency of one or a few minutes. Many research studies have shown that high-frequency data provides much more useful information than daily data for volatility forecasting. Taylor and Xu (1997) state that five-minute DM/dollar returns contain incremental information to option prices when forecasting volatility one hour ahead. Andersen and Bollerslev (1998) also find that using high-frequency intraday data can improve the foreign currency volatility estimates obtained from GARCH models. Blair et al. (2001) confirm this finding for a U.S. equity index. They claim that the realised volatility estimates obtained from high-frequency intraday returns are superior to one day or multi-day forecasts by using ARCH models.

Andersen, Bollerslev, Diebold and Ebens (2001), and Andersen, Bollerslev, Diebold and Labys (2001) find that a long memory process can better model the realised volatility. Andersen et al. (2003) also find that the combination of a long memory process and the use of high-frequency returns provide better volatility forecasts for foreign exchange rates. But Pong et al. (2004) argue that the better accuracy of volatility forecasts comes from the use of high-frequency data, not the long memory model. They show that the performance of volatility forecasts of foreign exchange rates using a long memory model is similar to using a short memory model for different horizons.

4.1.2 Options information

Volatility forecasts provide information about the volatility of the asset price in the future, while density forecasts are more complicated, as they produce information about the whole distribution of the asset's future price. Since option prices not only reflect historical information, but also contain forward-looking information about the future distribution of the asset price, volatility forecasters might rationally prefer implied volatilities inferred from option prices to realised variance calculated from historical time series.

There is a considerable literature comparing volatility forecasts obtained from option prices with volatility forecasts calculated from the history of asset prices. Xu and Taylor (1995) find that the historical volatility estimates are superior to implied volatility estimates for four foreign exchange rates for the period from 1985 to 1991. Blair et al. (2001) compare the information content of intraday returns and implied volatilities when forecasting index volatility, and they state that in both in-sample estimates and out-of-sample forecasting, implied volatilities estimated by the "old" VIX index perform the best. Martens and Zein (2004) state that the volatility estimates inferred from options are superior to volatility forecasts calculated from historical daily returns for S&P 500 index, YEN/USD exchange rate and light, sweet crude oil. They also find that the forecasts can be improved by using high-frequency data and a long memory process. Jiang and Tian (2005) extend the model-free implied volatility model proposed by Britten-Jones and Neuberger (2000), and find that the model-free implied volatility estimates outperform both the Black-Scholes implied volatility and past realised volatility for S&P 500 index. Giot and Laurent (2007) further confirm that implied volatility contains incremental information to past realised volatility for two stock indices, even when they separate the diffusion and jump components. Busch et al. (2011) similarly show that implied volatility is superior to realised volatility when forecasting future volatility in foreign exchange, stock and bond markets. The past realised volatility is decomposed into continuous and jump components and the forecast is made using a vector HAR model. These studies all state that option forecasts are more informative and accurate than historical forecasts of index volatility even when the historical information set includes high-frequency returns. We therefore anticipate a similar preference could apply to density forecasts.

However, we must note that some studies compare forecasts obtained from option prices and intraday returns and rank intraday index information highly. Bali and Weinbaum (2007) use both high-frequency data and implied volatility models to estimate S&P 100 index volatility and find that the forecasts obtained using intraday returns are superior to the forecasts obtained using daily option data, for one day and twenty days forecast horizons. Becker et al. (2007) find that VIX does not contain incremental information for forecasting volatility compared to model based forecasts. Martin et al. (2009) find that the spot-based volatility forecasts are superior to the options-based forecasts for three Dow Jones Industrial Average stocks during the period from 2001 to 2006.

4.2 Risk-neutral densities

4.2.1 Theoretical setup

Black and Scholes (1973) initiate the approach to price options under a no-arbitrage assumption. In the Black-Scholes model, we assume the price of the underlying asset follows a stochastic process, geometric Brownian motion

$$dS/S = \mu dt + \sigma dW \tag{4.1}$$

where μ is the expected return per annum, and is equal to the risk free rate r minus the dividend yield q, and plus the asset's risk premium.

Under the real world measure P, the distribution of stock price S_T is lognormal, then the distribution of $\log(S_T)$ is normal with mean $\log(S_0) + \mu T - \frac{1}{2}\sigma^2 T$ and variance $\sigma^2 T$, denoted:

$$log(S_T) \sim N\left(log(S_0) + \mu T - \frac{1}{2}\sigma^2 T, \sigma^2 T\right).$$

Under the risk-neutral Q-distribution, the risk-neutrality assumption requires a drift rate r-q instead of μ, and there is no risk premium, hence we have

$$dS/S = (r - q)dt + \sigma dW \tag{4.2}$$

$$log(S_T) \sim N\left(log(S_0) + (r - q)T - \frac{1}{2}\sigma^2 T, \sigma^2 T\right).$$

Let $\psi\left(x \middle| F_{0,T}, \sigma, T\right)$ be the density for S_T given by equation (4.2), with futures price $F_{0,T} = S_0 e^{(r-q)T}$. Then Black-Scholes call option prices are given by

$$c(x) = e^{-rT} \int_{0}^{\infty} max(x - X, 0)\, \psi\big(x \big| F_{0,T}, \sigma, T\big) dx \qquad (4.3)$$

Breeden and Litzenberger (1978) first obtain a general risk-neutral density (RND) from theoretical option prices by taking the second derivative of option prices with respect to strike prices. For a general RND $g_Q(x)$, fair call option prices are

$$c(X) = e^{-rT} \int_{X}^{\infty} (x - X) g_Q(x) dx. \qquad (4.4)$$

The RND is then obtained from option prices as

$$g_Q(X) = e^{rT} \frac{\partial^2 c}{\partial X^2} \qquad (4.5)$$

We use the example of call option, call spread and butterfly spread to illustrate the intuition behind. We construct the bull call spread by buying one call option with a lower strike price *X* and selling another call option with a higher strike price *X*+a as shown in Figure 4.1. Since the first derivative of a function *f(x)* is given as

$$f'(x) = \lim_{h \to 0} \frac{f(x + h) - f(x)}{h}.$$

Then the first difference of the call option is

$$\frac{C(S, X + a) - C(S, X)}{a}.$$

A long butterfly spread position is constructed by buying one call option with a strike price of *X*-*a*, selling two call options with a strike price of *X*, and buying one call option with a strike price of *X*+a as shown in Figure 4.2. Since the second derivative of a function *f(x)* is

Figure 4.1 Payoff of bull spread using call options.

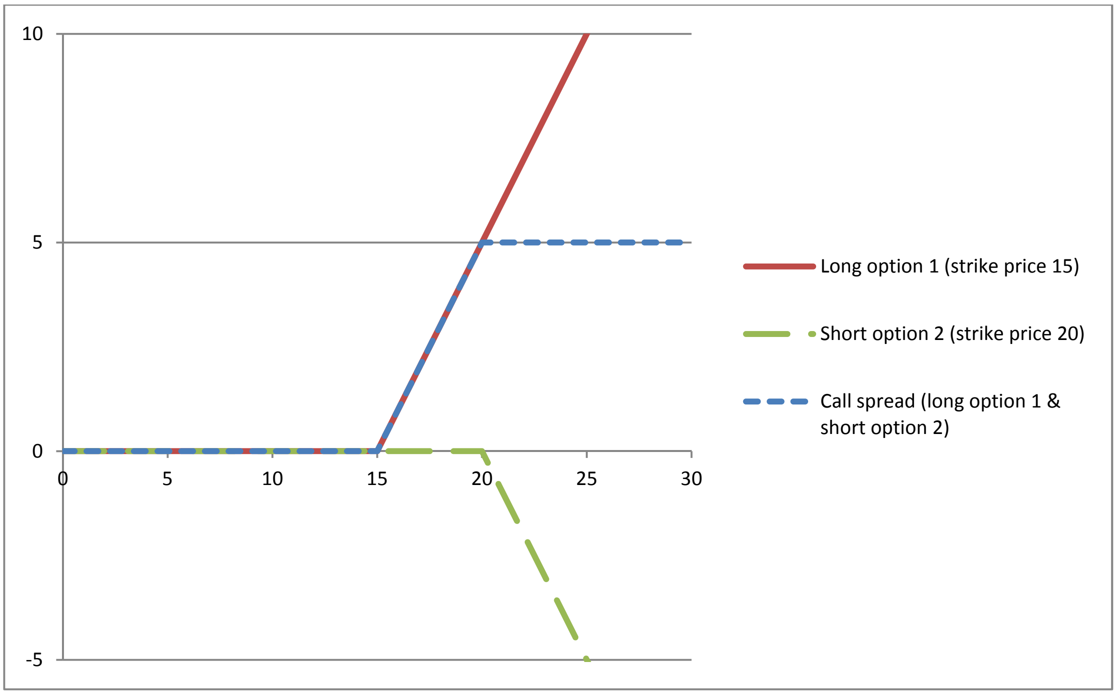

Figure 4.2 Payoff of butterfly spread using call options.

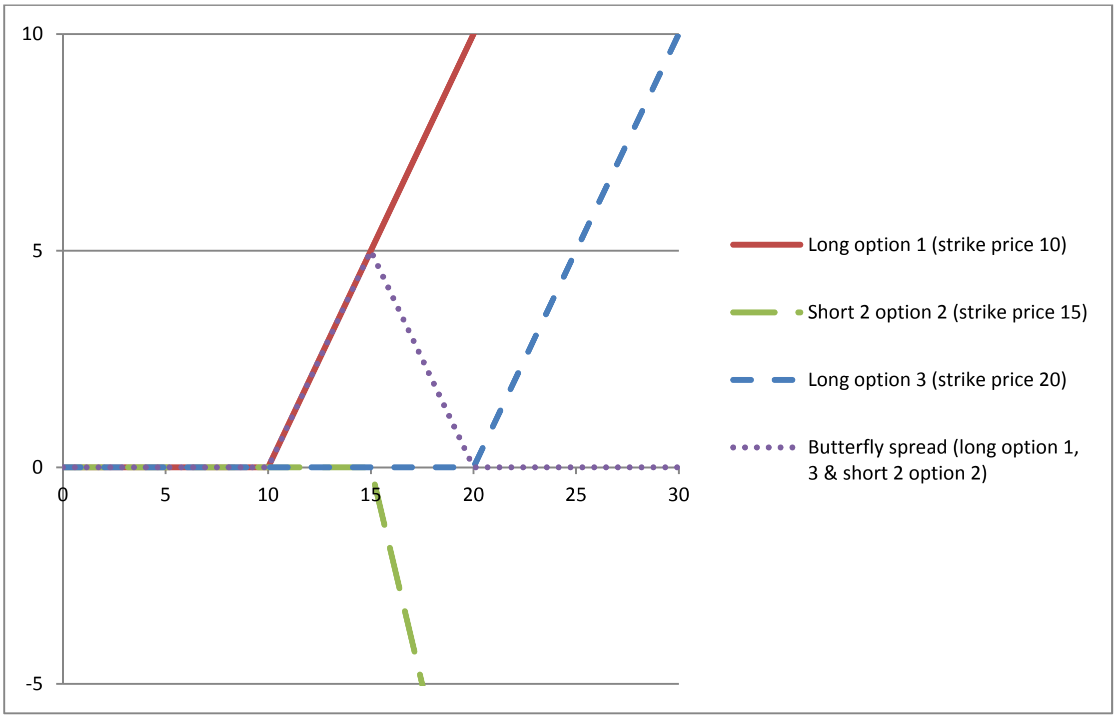

$$f''(x) = \lim_{h \to 0} \frac{\frac{f(x+h) - f(x)}{h} - \frac{f(x) - f(x-h)}{h}}{h}$$

$$= \lim_{h \to 0} \frac{f(x+h) - 2f(x) + f(x-h)}{h^2}.$$

Then the second difference of the call option is

$$\frac{C(S, X+a) - 2C(S, X) + C(S, X-a)}{a^2}.$$

As a approaches zero, the above function is Dirac and can represent the density. For an Arrow Debreu security, it pays one unit of numeraire under a particular state and zero under all other states.

4.2.2 Methods to extract risk-neutral densities

The lognormal Black-Scholes model is one of the most common ways to extract a risk-neutral density. As the implied volatility smile effect proposed by Rubinstein (1994) indicates that risk-neutral densities are not lognormal and volatility is not constant, hence some other methods rather than the lognormal Black-Scholes model should be employed to model the risk neutral process. Empirically the RND can be obtained by fitting the market option prices to the theoretical option prices across different strikes. Different methods have been proposed to obtain risk-neutral densities from option prices. These methods can be grouped into different categories including parametric methods, nonparametric methods, implied volatility spline methods and price dynamics related methods.

4.2.2.1 Parametric methods

The lognormal Black-Scholes model contains only one parameter, the implied volatility, hence the model is not flexible. In order to obtain a more accurate and flexible model, some researchers use a mixture of lognormal densities. Ritchey (1990) first introduces the mixture of two lognormal densities, which is a weighted combination of two lognormal densities. This model has five parameters, which are two implied volatilities, two futures prices, and the weight parameter. Hence the model is easy to apply and can always give non-negative risk-neutral densities.

The lognormal mixture has been applied on exchange rates by Jondeau and Rockinger (2000), and on equity indices by Bliss and Panigirtzoglou (2002), Anagnou-Basioudis et al. (2005) and Liu et al. (2007). Melick and Thomas (1997) extend this model to a mixture of three lognormal densities, with eight parameters. They study the risk-neutral densities of crude oil during the period of the first Gulf War.

Bookstaber and McDonald (1987) first introduce GB2, the generalised beta distribution of the second kind. The GB2 method can always give a non-negative risk-neutral density, and can reflect a flexible shape of tails of the distribution. But this method does not rely on a strong theoretical foundation. Anagnou-Basioudis et al. (2005) apply it to estimate risk-neutral densities for the S&P 500 index and the GBP/USD exchange rate, while Liu et al. (2007) use it to extract risk-neutral densities for the FTSE 100 index.

The Black-Scholes model assumes a standard normal distribution for the standardized returns. Madan and Milne (1994) modify this assumption by proposing the lognormal-polynomial method. They assume that the density of standardized returns equals the lognormal density multiplying a polynomial. This method has a sound theoretical foundation but may give negative empirical densities. This is because there are finite number of strikes which are not continuous. Hence it cannot capture the flexible shape of the tail distribution. Madan and Milne (1994) and Jondeau and Rockinger (2000) apply the lognormal-polynomial density functions for the S&P 500 index and exchange rates respectively.

4.2.2.2 Nonparametric methods

Since the parametric methods in general cannot reflect a very flexible shape of tails of the distribution, two main types of nonparametric methods are proposed to capture a more flexible distribution, which are flexible discrete distributions and kernel regression methods.

For the flexible discrete distributions method, usually the probabilities are obtained by minimizing some objective functions, which either measures the match between the observed and fitted option prices, or the smoothness of the risk-neutral densities. Jackwerth and Rubinstein (1996) examine the fit for S&P 500 index option prices, while Jackwerth (2000) investigates the smoothness of the implied volatility function. However, this method can give negative probabilities due to the discreteness of the dataset.

Ait-Sahalia and Lo (1998, 2000) employ a nonparametric kernel regression method to extract risk-neutral densities from S&P 500 option prices. This method does not need to presume any dynamics of the distribution of the asset prices, but assumes a non-linear relationship and then regresses the option price or the implied volatility against stock and strike prices, time to expiration, dividend yield and risk-free rate. The density is then obtained by taking the second differential of the estimated option prices. However, this method assumes the option prices and the implied volatilities have time-invariant distributions and also requires a large amount of data.

4.2.2.3 Implied volatility method

While the parametric and nonparametric methods focus on fitting the distribution of option prices, another method examines the fit of implied volatilities. The option prices are then obtained from converting the estimated implied volatilities, and the risk-neutral densities are extracted by taking the second derivative of the option prices relative to strike prices. Bates (2000) fits the observed S&P 500 futures option prices, while Bliss and Panigirtzoglou (2002) use a cubic smoothing spline and focus on delta, rather than the implied volatilities, to extract the risk-neutral densities. This method does not require intensive data, but could give negative probabilities.

4.2.2.4 Price dynamics methods

Apart from the methods focusing on the distribution of the prices and fitting the implied volatilities, the last type of method centres on price dynamics. The risk-neutral densities can be extracted if the underlying assets are assumed to have specific risk-neutral dynamics, such as the geometric Brownian motion for the Black-Scholes model which leads to lognormal distributions of the underlying asset prices.

Option pricing with jump diffusion, stochastic volatility model and stochastic volatility model with jump diffusion can be generalised using the affine class of option pricing models. The affine jump diffusion process assumes that the drift vector, covariance matrix and the jump intensities all have affine dependence on the state vector (Duffie et al., 2000). One property of the affine class of option pricing models is that the state vector has closed form conditional characteristic function and analytically tractable solution, which can be evaluated by the Fourier inversion transformation.

The implied volatility smile effect indicates that geometric Brownian motion is not a proper asset price process as the risk-neutral densities are not lognormal and volatility is not constant. Merton (1976) extends geometric Brownian motion to a jump diffusion process. Bates (1991) confirms that an option pricing formula incorporating a jump diffusion asset price process fits market prices better. Some studies use a stochastic process to model volatility. Hull and White (1987) first propose option pricing model with stochastic volatility process. Heston (1993) assumes the volatility follows a mean-reverting square-root process and has a closed form solution for option prices and densities based on inverting characteristic functions. Hence the stochastic volatility process of Heston (1993) is a desirable choice as it considers both the volatility smile and the term structure effects. Extensions of the Heston (1993) model are in Bates (1996) who also incorporates jumps, and in Duffie et al. (2000), Eraker (2004), Eraker et al. (2003) and Pan (2002) who include a jump process in both price and volatility components.

The use of the Heston (1993) model enables calculation of density forecasts for all horizons, which can be a considerable advantage. Many studies can only estimate densities for horizons identical to option expiry dates such as Jackwerth and Rubinstein (1996), Melick and Thomas (1997) and Bliss and Panigirtzoglou (2002), while Shackleton et al. (2010) use the Heston model to compare density forecasts for multiple horizons.

4.2.3 Comparisons among estimation methods

Many researchers conduct comparisons among different methods to extract risk-neutral densities based on the accuracy of the estimates. Jondeau and Rockinger (2000) compare the mixture of lognormal, the Heston stochastic volatility with jumps and the lognormal polynomial methods for FF/DM exchange rate. They find that the mixture of lognormal is the best for short time-to-expiry options, while the jump diffusion model performs the best for longer maturities. Bliss and Panigirtzoglou (2002) compare the mixture of lognormal and the smoothed implied volatility spline methods to forecast risk-neutral densities on sterling interest futures options and FTSE 100 index options and find that the smoothed implied volatility smile method outperforms the lognormal mixture method. Liu et al. (2007) use parametric methods including the lognormal mixture, the smooth spline and the GB2 methods to forecast risk-neutral densities for the FTSE 100 index and find that the mixture of lognormal and the GB2 methods give higher log-likelihoods than the spline smoothing method.

4.3 Transformations from risk-neutral densities into real-world densities

Both Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005) argue that the risk-neutral densities are not good forecasts of the future distribution of asset prices. The risk-neutral density is a suboptimal forecast of the future distribution of the asset price as there is no risk premium in the risk-neutral world, while in reality investors are risk-averse. Hence we need to use economic models and/or econometric methods to transform risk-neutral densities into real-world densities (RWDs).

4.3.1 Economic models to transform densities

4.3.1.1 Utility method

Economic models define the pricing kernel, which is the stochastic discount factor derived from risk-neutral and real-world densities, respectively g_Q and g_P , as follows:

$$m(S_T) = e^{-rT} \frac{g_Q(S_T)}{g_P(S_T)} \tag{4.6}$$

The pricing kernel is used to transform the risk-neutral densities into the real-world densities, and is proportional to the marginal utility of the representative agent given appropriate assumptions, hence the focus is on the choice of utility function in

$$m(x) = \lambda \frac{du}{dx} \tag{4.7}$$

where $u(x)$ is the utility function and λ is a positive constant.

Many researchers employ power and/or exponential utility functions to transform the risk-neutral densities into the real-world densities, and the power utility assumes a constant risk aversion while the exponential specification states that the risk aversion is varying. Some studies also examine more general cases by assuming wider classes of utility functions.

Anagnou-Basioudis et al. (2005) use the power utility function to transform the risk-neutral densities into the real-world densities for sterling exchange rates and the S&P 500 index, and state that the null hypothesis that the RWD is an efficient estimate of the real densities cannot be rejected. Liu et al. (2007) also employ the power utility function to transform the risk-neutral densities into the real-world densities for FTSE 100 index and make comparisons based on the log-likelihood criterion; they show that the latter outperforms the former. Bliss and Panigirtzoglou (2004) use both power and exponential utility functions to transform the RNDs into the RWDs for S&P 500 and FTSE 100 indices at different horizons, and state that the estimated RWDs are all reasonable. Kang and Kim (2006) extend the analysis to more generality by using the hyperbolic absolute risk aversion (HARA) function, the log plus power, and the linear plus exponential utility. They examine the FTSE 100 index and conclude that the more flexible utility functions provide more forecasting power.

4.3.1.2 Drift correction method

Drift transformations are possible for specific price dynamics. For example, suppose the continuous-time risk-neutral price dynamics for the stock, which incorporate the stochastic variance V, follows a square-root process, as follows:

$$dS/S = (r - q)dt + \sqrt{V}dW_1 \quad (4.8)$$

$$dV = \kappa(\theta - V)dt + \sigma\sqrt{V}dW_2 \quad (4.9)$$

Thus we can include linear risk premium terms in both the price and the volatility components to define an affine real-world diffusion process

$$dS/S = (aV + r - q)dt + \sqrt{V}dW_1 \quad (4.10)$$

$$dV = [bV + \kappa(\theta - V)]dt + \sigma\sqrt{V}dW_2 \quad (4.11)$$

The drift adjustment terms are here assumed to be linear in V, the property of the affine class of option pricing models ensures that an analytical solution can be obtained for the real-world characteristic functions. The real-world densities are then given by equation (5.14) and depend on the drift rates a and b.

4.3.2 Econometric methods to transform densities

The econometric approach is based on Rosenblatt (1952), which states that, if the forecasted density is correct, then the forecasted cumulative probability is uniform i.i.d..

We can use a parametric method to transform the RNDs into the RWDs. At time 0, we let $g_{Q,T}(x)$ and $G_{Q,T}(x)$ define the risk-neutral density and the cumulative distribution function (c.d.f.) of the random variable S_T. We denote u_T=$G_{Q,T}(S_T)$. We then follow Bunn (1984), Dawid (1984), and Diebold et al. (1999) to denote the calibration function $C_T(u)$, which is the real-world c.d.f. of the random variable u_T. The calibration function depends on the forecast horizon T. We now consider the real world c.d.f. of S_T, with Pr standing for the real world probabilities. The real-world c.d.f. of S_T is

$$\Pr(S_T \leq x) = \Pr\left(G_{Q,T}(S_T) \leq G_{Q,T}(x)\right) = \Pr\left(u_T \leq G_{Q,T}(x)\right)$$
$$= C_T\left(G_{Q,T}(x)\right) \quad (4.12)$$

Hence the real-world c.d.f. of S_T is

$$G_{P,T}(x) = C_T\left(G_{Q,T}(x)\right) \quad (4.13)$$

The real-world density of S_T is

$$g_{P,T}(x) = \frac{d}{dx} C_T\left(G_{Q,T}(x)\right) = \frac{du}{dx}\frac{d}{du} C_T(u) = g_{Q,T}(x) c_T(u). \quad (4.14)$$

where $c_T(u)$ is the density of u_T.

It is necessary to assume the calibration function $C_T(u)$ is invariant over time, and it is standard to assume the parametric calibration function is the c.d.f. of the Beta distribution. Fackler and King (1990) first use the equation to transform the risk-neutral densities into the real-world densities for corn, soybeans, live cattle and hogs option prices, and Liu et al. (2007) also employ it to transform the densities. The calibration density is

$$c_T(u) = u^{c-1}(1-u)^{d-1}/B(c,d), \qquad 0 \leq u \leq 1 \quad (4.15)$$

and the constant $B(c, d)=\Gamma(c)\,\Gamma(d)/\,\Gamma(c+d)$. The two calibration parameters c and d depend on the horizon T. The special case that c=d=1 denotes a uniform distribution and the RNDs and the RWDs are identical. The real-world density is

$$g_{P,T}(x) = \frac{G_{Q,T}(x)^{c-1}(1 - G_{Q,T}(x))^{d-1}}{B(c,d)} g_{Q,T}(x) \quad (4.16)$$

Liu et al. (2007) use both utility and statistical calibration transformations, and show that a statistical calibration gives a higher log-likelihood than a utility transformation. Shackleton et al. (2010) compare parametric and nonparametric transformations, obtaining better results for the latter. Hence we also transform the risk-neutral densities into the real-world densities using a nonparametric transformation explained in chapter 5.

4.4 Density forecast applications

Density forecasts have been applied in many areas. They can be used to estimate the risk aversion of investors. They can also be employed to infer probabilities of future market changes for different asset classes. Furthermore, they can be used to assess market beliefs about future economic and political events when derived from option prices. Last but not least, density forecasts are important in risk management, particular for the estimation of Value-at-Risk (VaR). Hence density forecasts are of importance to central bankers and other decision takers for activities such as policy-making, risk management and derivatives pricing.

4.4.1 Estimated risk aversion

We can assess the rationality of estimated risk-neutral densities by referring to their associated risk aversion estimates. The utility function has first derivative under the representative agent model given by

$$u'(x) = \frac{e^{-rT} g_Q(x)}{\lambda g_P(x)} \quad (4.17)$$

where λ is a positive constant. The second derivative is negative for a rational utility function for all values of x. We can assess the rationality of risk-neutral densities by estimating the risk aversion function implied by the first and second derivatives of the utility function as

$$RA(x) = -\frac{u''(x)}{u'(x)} = \frac{g_P'(x)}{g_P(x)} - \frac{g_Q'(x)}{g_Q(x)} \tag{4.18}$$

The risk aversion function must be positive for all x if the utility function is rational.

Jackwerth (2000) estimates risk aversion for the S&P 500 index around the 1987 market crash. Before the crash, the risk aversion function is positive and consistent with the economic theory, while after the crash the risk aversion function has negative values and increases with wealth, which contradicts the assumptions. Jackwerth (2000) argues that mispriced options is the most likely reason. Ait-Sahalia and Lo (2000) estimate risk aversion for S&P 500 index options for 1993. They find the risk aversion function is positive, but has an irregular U-shape. Bliss and Panigirtzoglou (2004) infer the relative risk aversion (RRA) function for FTSE 100 and S&P 500 index options for multiple horizons. They state that all their estimates are reasonable, and the RRA declines as the forecast horizon increases, and it is lower when the market volatility is high.

4.4.2 Infer future market change

Density forecasts have been employed to estimate probabilities of future market changes for different asset classes including for stock indices Shackleton et al. (2010) and Yun (2014), for interest rates Ivanova and Gutierrez (2014), for exchange rates Sarno and Valente (2004), for commodities Hog and Tsiaras (2010) and for lean hog futures Trujillo-Barrera et al. (2012).

Melick and Thomas (1997) employ the mixture of the lognormal to estimate density for crude oil during the first Persian Gulf crisis. The mixture of lognormal method to extract densities can clearly show the change of investor expectation in the market, as the single lognormal model would overestimate the market's assessment of the probability of a major disruption and underestimate the effect on prices of such a disruption.

4.4.3 Assess market beliefs

Density forecasts can be employed to assess market beliefs about future economic and political events when derived from option prices due to its forward-looking property. The ex-ante analysis infers the possible outcome of the market due to the event, while the ex post analysis checks if the market reacts to the event as expected.

Early studies including Bakshi et al. (2003), Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005) use the full dataset to make risk-transformations. The real-world densities obtained are then ex post because each forecast is made using some information from later asset prices. However it is best to apply ex ante transformations as in Shackleton et al. (2010). Thus we should only use past and present asset and option prices to construct real-world densities as is done in chapter 5.

4.4.4 Estimate Value-at-Risk

Density forecasts play an important role in risk management, especially for the estimate of VaR, which measures how much one can lose at a pre-defined confidence interval over different horizons. Many institutions, such as J.P. Morgan and the Bank of England, periodically publish their density estimates, which enable investors to access risk for their investment portfolios. Hence density forecasts are important to central bankers and other decision takers for activities such as policy-making, risk management and derivatives pricing.

4.5 Density forecast evaluation

Volatility forecast evaluation can be problematic because volatility is latent, and density forecast evaluation faces a similar problem. Blair et al. (2001) use the squared daily returns and Martens and Zein (2004) use the realised variance as the ex post proxy, however, no similar proxy exists for density forecasts. Some studies including Jondeau and Rockinger (2000) and Bliss and Panigirtzoglou (2004) evaluate density forecasts based on option pricing. But this method only works for risk-neutral densities which have closed-form solutions for options, but not the real-world densities. Many researchers focus on the time series properties of density forecasts and use the probability integral transform (PIT), while some other people prefer log-likelihoods.

4.5.1 Diagnostic tests

Many studies use the properties of time series of density forecast. Rosenblatt (1952) employs the probability integral transform, and states that the PIT is i.i.d. uniform if density forecasts are correctly specified. Diebold et al. (1998) initiated the idea of using PIT to evaluate density forecasts. They employ a graphical tool to check the i.i.d. of the PIT. Some researchers extend the method in Diebold et al. (1998) to formal diagnostic tests. One to mention is the Kolmogorov and Smirnov (KS) test. The KS test checks the maximum difference between the empirical and theoretical cumulative functions, so that we can evaluate if the values of a variable are compatible with a certain distribution. The KS test is applied widely as it is simple to implement. However, one needs to be careful when interpreting the test results, as the KS test tests the uniformity under the i.i.d. assumption rather than checks the i.i.d. and the uniformity jointly.

Many studies question the power of the KS test when evaluating density forecast. Berkowitz (2001) proposes the BK test, which states that if the PIT is i.i.d. uniform, then the normal inverse cumulative function of the PIT is i.i.d. normal. The benefit of the BK test is that it can test the independence and the uniformity jointly. The BK test has been employed widely in studies including Clements and Smith (2000) and Shackleton et al. (2010). Clements (2004) uses PIT to evaluate the UK Monetary Policy Committee's inflation density forecasts. Some studies, including Bliss and Panigirtzoglou (2004), Anagnou-Basioudis et al. (2005), and Kang and Kim (2006), by minimising the BK test statistics, estimate parameters of utility functions to transform the RNDs into the RWDs.

Some studies also doubt the power of the BK test. Clements et al. (2003) conduct density forecasts using linear autoregressive (AR) and self-exciting threshold autoregressive (SETAR) models. They implement a Monte Carlo simulation incorporating a two-regime SETAR process to generate data. They find that the BK test produces high *p* value and is unable to reject the mis-specified linear model. Guidolin and Timmermann (2005) investigate the economic implications of 'bull' and 'bear' regimes in UK stock and bond returns. They find the BK test gives a low rejection rate and can hardly show the single-state model to be deficient, while the standard Jarque-Bera test is more powerful to reject the mis-specified single-state model.

4.5.2 Maximum log-likelihood

Apart from diagnostic tests, researchers also use the log-likelihood metric to evaluate density forecasts. One shortcoming of the BK test is that models cannot be compared if they are all accepted or rejected. A comparison of log-likelihood among different models can solve this problem, as employed by Bao et al. (2007), Liu et al. (2007) and Shackleton et al. (2010).

Amisano and Giacomini (2007) use an out-of-sample "weighted likelihood ratio" test to compare density forecasts. The forecasts can be made based on different models, including parametric (nested or non-nested), semiparametric and nonparametric models and Bayesian estimation techniques. They employ the test to evaluate density forecasts of U.S. inflation and state that the Markov-switching, Phillips curve model obtained by maximum likelihood gives the best density forecasts of U.S. inflation.

4.6 Density forecasting comparisons

Only a few studies have compared density forecasts. Liu et al. (2007) examine FTSE 100 index futures and options prices from 1993 to 2003. They extract risk-neutral densities using three methods, including a mixture of two lognormals, a generalised beta and a flexible density defined by spline functions. They transform the RNDs (defined in the first paragraph on page 74) into the RWDs (defined in the second paragraph on page 80) using both a utility function and a statistical calibration. They find that densities obtained from option prices are superior to historical densities based on the log-likelihood criterion, and a combination of parametric, real-world and historical densities produces the best density forecasts.

Shackleton et al. (2010) compare density forecasts of the S&P 500 index from 1991 to 2004, using both daily option prices and five-minute index returns. They use the GJR model to obtain densities from historical returns and employ the Heston (1993) model which incorporates stochastic volatility to extract RNDs from option prices. They use three methods to transform the RNDs into the RWDs, including a drift correction method, a parametric and a nonparametric method. They conduct ex ante density forecasts for multiple horizons ranging from one day to twelve weeks and obtain mixed findings. ARCH densities are more informative for the one day and one week horizons, because an accurate forecast of tomorrow’s variance is obtainable from high frequency returns summarized by the daily measures of the realised volatility. RWDs provided by option prices perform better for two weeks and four weeks horizons.

They argue that this can be attributable to the forward looking property of option prices as they only use option prices for the contracts with medium-term maturities longer than one week, while the short term one day and one week RNDs are extrapolations that are not supported by trading options for such short maturities.

Kostakis et al. (2011) use monthly closing prices for S&P 500 futures options from 1986 to 2009 and extract implied distributions and transform them into the corresponding risk-adjusted ones. They then investigate, from a portfolio allocation perspective, combining investment in a risky and a risk-free asset and state that the risk-adjusted implied distributions perform better than the historical returns' distributions even when they consider transaction costs.

Yun (2014) studies the S&P 500 stock index and options from 1987 to 2000 and conducts out-of-sample density forecasts of the affine jump diffusion models. They find that the time-varying jump risk premia models are superior for density forecasts than the other models based on the log-likelihood criterion.

Hog and Tsiaras (2010) focus on crude oil prices for the period from 1994 to 2006. They extract risk-neutral densities from crude oil option prices and compare with the standard ARCH type models. They transform the RNDs into the RWDs using both parametric and nonparametric calibration. They evaluate density forecasts using the goodness-of-fit tests and out-of-sample likelihood comparisons, and state that nonparametric calibration is superior to parametric transformation and option prices are more informative than historical returns.

Ivanova and Gutierrez (2014) examine interest rate, Euribor futures options daily observations, from 1999 to 2012. They extract risk-neutral densities from option prices using the spline method proposed by Bliss and Panigirtzoglou (2002), and transform the RNDs into the RWDs using parametric and nonparametric calibrations following Fackler and King (1990). They obtain density forecasts four weeks prior to option expiry and conclude that the RWDs, not the RNDs, can generate reliable forecasts.

5. Density forecast comparisons for stock prices, obtained from high-frequency returns and daily option prices

5.1 Introduction

Density forecasts have been used to infer the probabilities of future market changes for different asset classes including stock prices, interest rates, exchange rates and commodities. They can also be used to assess market beliefs about future economic and political events when derived from option prices. Also, density forecasts are important in risk management, particular for the estimation of Value-at-Risk, which measures how much one can lose at a pre-determined confidence interval over different horizons. Hence density forecasts are of importance to central bankers and other decision takers for activities such as policy-making, risk management and derivatives pricing.

Volatility forecasts produce forward-looking information about the volatility of the asset price in the future, while density forecasts are more sophisticated as they provide information about the whole distribution of the asset's future price. Since option prices reflect both historical and forward-looking information, volatility forecasters might rationally prefer implied volatilities from option prices to realised variance calculated from historical time series. We anticipate a similar preference could apply to density forecasts. We compare density forecasts derived from option prices using the Heston (1993) model and forecasts obtained from historical time series using the Corsi (2009) Heterogeneous Autoregressive model of Realised Variance (HAR-RV). We also transform the risk-neutral densities into real-world densities using a nonparametric transformation. There are no known previous results for individual stocks, so our contribution is to provide the first comparison for density forecasts obtained from option prices and historical intraday returns for individual stocks. We investigate seventeen stocks from the Dow Jones 30 Index for four horizons ranging from one day to one month for the period from 2003 to 2012.

This chapter is structured as follows. Section 5.2 covers methodology, including the density forecasting methods, namely Heston model for densities inferred from option prices and HAR-RV model for density forecasts obtained from historical high-frequency returns. It also includes the econometric methods used to obtain ex-ante parameters and evaluate density forecasts. Section 5.3 describes the Dow Jones 30 stock and option prices data employed in the study. Section 5.4 focuses on the empirical analysis. Section 5.5 summarises the findings and concludes.

5.2 Methodology

5.2.1 Option pricing with stochastic volatility

We want to extract the risk-neutral density for the underlying asset from option prices, and a realistic process for an individual stock must incorporate a stochastic volatility component, whose increments are correlated with the price increments. We need to calculate an enormous number of theoretical option prices, so fast calculations are essential. The stochastic volatility process of Heston (1993) meets all our requirements as it has closed-form densities and option prices.

The risk-neutral price dynamics for the stock price S, which incorporate the stochastic variance V, is defined as below

$$\frac{dS}{S} = (r - q)dt + \sqrt{V}dW_1 \tag{5.1}$$

where r is the risk-free interest rate, q is the dividend yield, and W_1 is a Wiener process. For the variance, we have the familiar square-root process of Cox et al. (1985) written as

$$dV = \kappa(\theta - V)dt + \sigma\sqrt{V}dW_2 \tag{5.2}$$

We let ρ denote the correlation between the two Wiener processes W_1 and W_2, while θ is the level towards which the stochastic variance V reverts, and κ denotes the rate of reversion of V towards θ. The volatility of volatility parameter σ controls the kurtosis of the returns. More complicated affine jump-diffusion processes which have closed-form solutions are described by Duffie et al. (2000). We do not consider these, noting that Shackleton et al. (2010) obtained no benefits from including price jump in their study.

Heston (1993) assumes q=0 and also makes some assumptions about the price of volatility risk, by referring to Black and Scholes (1973) and Merton (1973), where the value of any asset denoted u should satisfy the partial differential equation (PDE) (5.3) as below

$$\frac{1}{2}VS^2\frac{\partial^2 U}{\partial S^2} + \rho\sigma VS\frac{\partial^2 U}{\partial S\,\partial V} + \frac{1}{2}\sigma^2 V\frac{\partial^2 U}{\partial V^2} + rS\frac{\partial U}{\partial S}$$

$$+\{\kappa[\theta - V] - \lambda(S,V,t)\}\frac{\partial U}{\partial V} - rU + \frac{\partial U}{\partial t} = 0, \tag{5.3}$$

where the term $\lambda(S,V,t)$ denotes the price of the volatility risk and Heston (1993) assumed $\lambda(S,V,t) = \lambda V$.

A European call option with strike price K which matures at time T and satisfies equation (5.3) satisfies the following boundary conditions in equations (5.4), for $0 \le t \le T$.

$$U(S,V,T) = Max(0, S-K)$$

$$U(0,V,t) = 0$$

$$\frac{\partial U}{\partial S}(\infty, V, t) = 1$$

$$rS\frac{\partial U}{\partial S}(S,0,t) + \kappa\theta\frac{\partial U}{\partial V}(S,0,t) - rU(S,0,t) + U_t(S,0,t) = 0$$

$$U(S,\infty,t) = S \tag{5.4}$$

Similar to the Black-Scholes formula, at time 0 the Heston call price formula is derived by assuming

$$C(S_0, V_0, 0) = S_0P_1 - KP(0,T)P_2. \tag{5.5}$$

The first term S_0 is the current value of the spot price, while the second term $KP(0,T)$ is the present value of the strike price K. The terms P_1 and P_2 are functions of S_0, V_0 and the parameters in (5.1) and (5.2).

We also use x to denote the logarithm of the spot price as defined in equation (5.6)

$$x_t = log(S_t). \tag{5.6}$$

Heston (1993) substituted equation (5.5) into the PDE equation (5.3) to show that P_1 and P_2 must both satisfy the following PDEs in equation (5.7)

$$\frac{1}{2}V\frac{\partial^2 P_j}{\partial x^2} + \rho\sigma V\frac{\partial^2 P_j}{\partial x\,\partial V} + \frac{1}{2}\sigma^2 V\frac{\partial^2 P_j}{\partial V^2} + \left(r + u_jV\right)\frac{\partial P_j}{\partial x}$$

$$+\left(a_j - b_j V\right)\frac{\partial P_j}{\partial V} + \frac{\partial P_j}{\partial t} = 0 \tag{5.7}$$

for $j = 1, 2,$ and

$$u_1 = \frac{1}{2}, u_2 = -\frac{1}{2}, a_1 = a_2 = \kappa\theta, b_1 = \kappa + \lambda - \rho\sigma, b_2 = \kappa + \lambda \tag{5.8}$$

with $\lambda = 0$ when the other parameters are for risk-neutral dynamics.

When x follows the stochastic process, the relevant price dynamics are given by equations (5.9)

$$dx = \left[r - q + u_j V\right]dt + \sqrt{V}dW_1,$$

$$dV = \left(a_j - b_j V\right)dt + \sigma\sqrt{V}dW_2. \tag{5.9}$$

where the parameters u_j, a_j and b_j are denoted as before. Each P_j in equation (5.5) is a conditional probability that the call option expires in-the-money. The term P_2 is derived from the characteristic function of S_T under the risk-neutral measure Q, while P_1 is derived from the characteristic function of S_T under a related measure Q^* for different drift rates in equation (5.9).

Probabilities are obtained from the conditional characteristic function of $log(S_T)$, which is denoted by $g(\Phi)$ and defined for all real numbers Φ, with $i = \sqrt{-1}$, as

$$g(\Phi) = E\left(e^{i\Phi log(S_T)} \middle| S_0, V_0\right). \tag{5.10}$$

This is a complex-valued function. Heston (1993) solves the PDEs to get the characteristic function solution

$$g(\Phi) = e^{\,C + DV_0 + i\Phi\, log(S_0)} \tag{5.11}$$

The terms C and D are calculated from long equations which can be written as follows in equations (5.12), for a selected value of *i*:

$$C = r\Phi Ti + \frac{a_j}{\sigma^2}\{(b_j - \rho\sigma\Phi i + d)T - 2\,log[\frac{1 - ke^{dT}}{1 - k}]\}$$

$$D = \frac{b_j - \rho\sigma\Phi i + d}{\sigma^2}[\frac{1 - e^{dT}}{1 - ke^{dT}}]$$

and

$$k = \frac{b_j - \rho\sigma\Phi i + d}{b_j - \rho\sigma\Phi i - d}$$

$$d = -\sqrt{(\rho\sigma\Phi i - b_j)^2 - \sigma^2(2u_j\Phi i - \Phi^2)}. \tag{5.12}$$

When the asset pays continuous dividends, so $q>0$. S_0 is replaced by S_0e^{-qT} in (5.5) and (5.11). For options on futures, $q=r$. Each desired probability can be obtained by inverting the characteristic function, which is given as

$$P(S_T \geq K \mid S_0, V_0) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{-i\Phi\,log(K)}\,g(\Phi)}{i\Phi}\right]d\Phi \tag{5.13}$$

where $Re[.]$ is the real part of a complex number (Kendall et al. 1987). This integral can be evaluated rapidly and accurately by numerical methods.

Two methods can be implemented to evaluate this integral. One method is to use Adaptive Simpson's Rule, Matlab function quad (@fun, a, b) uses Adaptive Simpson's Rule on the function @fun over the interval [a, b]. The other method employs Gauss Lobatto Rule, Matlab function quadl (@fun, a, b) uses Gauss Lobatto Rule to integrate @fun over [a, b] numerically. Matlab defines the functions, quad and quadl, as lower and higher order quadrature rules, we hence expect quadl to be superior and employ it in our study.

The integral also provides the conditional cumulative distribution function of S_T, therefore

$$F(y) = P(S_T \leq y \mid S_0, V_0) = 1 - P(log(S_T) \geq log(y) \mid S_0, V_0).$$

From routine calculations, the conditional risk-neutral density for positive values of y is hence

$$f(y) = \frac{dF}{dy} = \frac{1}{\pi y} \int_0^\infty Re\left[e^{-i\Phi\, log(y)} g(\Phi)\right] d\Phi. \tag{5.14}$$

However, several studies including Kahl and Jackel (2006) and Shackleton et al. (2010) point out that using the positive root solution for *d* in equations (5.12) can cause a discontinuity problem in the integrand in equation (5.13), and an investigation shows that this actually arises from the complex logarithm in function *C* in equation (5.11). The problem occurs if software chooses values of the complex logarithm which do not guarantee a continuous characteristic function. We emphasise the problem because many researchers appear to be unaware of it. Several methods have been proposed to solve this problem, and we follow Shackleton et al. (2010) to take the negative root solution *d* in equations (5.12). Gatheral (2006) asserts that using the negative root does not lead to false option prices, based upon extensive empirical experience.

5.2.2 High-frequency HAR methods

The HAR-RV model of Corsi (2009) is a simple AR-type model for the realised volatility which combines different volatility components calculated over different time horizons, and has been applied in Andersen et al. (2007) and Busch et al. (2011).

The HAR-RV model states that the multiperiod realised variance is the average of the corresponding one-period measures denoted as

$$RV_{t,t+h} = h^{-1}[RV_{t+1} + RV_{t+2} + \cdots + RV_{t+h}] \tag{5.15}$$

where h=1, 2, …, by definition $RV_{t,t+h} \equiv RV_{t+h}$ and we use h=5 and h=22 to represent the weekly and monthly realised volatility. Here the time period for predictions is from t to $t+h$, both counting trading days. In contrast, our options notation is a time period from 0 to T, both measured in years. The HAR model is a regression model which is standard and unbiased. There are literature about the GARCH model and heavy tailed time series using variance targeting as a means of reducing estimation bias, but there is no literature about the HAR model.

The HAR-RV model of Corsi (2009) is stated as a regression of the next RV on today's RV and the average RVs over the latest week and month:

$$RV_{t,t+1} = \beta_0 + \beta_D RV_{t-1,t} + \beta_W RV_{t-5,t} + \beta_M RV_{t-22,t} + \varepsilon_{t,t+1}.$$

To make predictions for the next h-day period, the regression specification is simply:

$$RV_{t,t+h} = \beta_{0,h} + \beta_{D,h} RV_{t-1,t} + \beta_{W,h} RV_{t-5,t} + \beta_{M,h} RV_{t-22,t} + \varepsilon_{t,t+h}. \tag{5.16}$$

Some volatility forecast models also employ standard deviations as opposed to variances. Andersen et al. (2007) present the standard deviation form of HAR-RV model as

$$(RV_{t,t+h})^{1/2} = \beta_{0,h} + \beta_{D,h} \left(RV_{t-1,t}\right)^{1/2} + \beta_{W,h} (RV_{t-5,t})^{1/2} + \beta_{M,h} (RV_{t-22,t})^{1/2} + \varepsilon_{t,t+h} \tag{5.17}$$

Given the logarithmic daily realised volatilities are approximately unconditionally normally distributed (Andersen et al., 2001), Andersen et al. (2007) also predict the realised variance in logarithmic form as

$$log\left(RV_{t,t+h}\right) = \beta_{0,h} + \beta_{D,h}\, log\left(RV_{t-1,t}\right) + \beta_{W,h}\, log\left(RV_{t-5,t}\right) + \beta_{M,h}\, log\left(RV_{t-22,t}\right) + \varepsilon_{t,t+h} \tag{5.18}$$

We also use the logarithmic form of realised variance in our study. However, Pong et al. (2004) state that we cannot simply take the exponential of a forecast of logarithmic volatility to get a forecast of the variance, as the forecasts obtained will be biased. We thus follow Granger and Newbold (1976) to get the volatility forecast. In their notation,

$$x_{n+h} = f_{n+h} + e_{n+h}^{(x)} \tag{5.19}$$

where $e_{n+h}^{(x)}$ is the h-step forecast error of x_{n+h} and f_{n+h} is the optimal forecast of x_{n+h} made at time *n*. Using $I_n = \{x_{n-j}, j \geq 0\}$, we define $S^2(h)$ to be the variance of the h-step forecast error of x_{n+h}:

$$S^2(h) = var\left(e_{n+h}^{(x)}\right). \tag{5.20}$$

The optimal forecast of $exp(x_{n+h})$ using I_n is then given by

$$g_{n+h}^{(x)} = \exp\left(f_{n+h} + \frac{1}{2}S^2(h)\right) \tag{5.21}$$

assuming $\{x_n\}$ is a Gaussian process. This is a standard assumption for $log(RV_t)$. The reason for this is that the logarithmic returns are biased, and we reduce the bias in the variance. While in (5.25) we reduce the bias in the expectation of the lognormal prices.

5.2.3 Lognormal densities, from the Black-Scholes model and HAR-RV forecasts

In the Black-Scholes model, we assume the prices follow geometric Brownian motion

$$dS/S = \mu dt + \sigma dW \tag{5.22}$$

where μ is the expected return per annum, and is equal to the risk free rate plus the asset's risk premium and minus the dividend yield.

Since the distribution of stock price S_T is then lognormal, the distribution of $log(S_T)$ is normal:

$$log(S_T) \sim N(log(S_0) + \mu T - \frac{1}{2}\sigma^2 T, \sigma^2 T)$$

Under the risk-neutral or the Q-distribution, the risk-neutrality assumption requires a drift rate r-q instead of μ, and hence we have

$$log(S_T) \sim N(log(S_0) + (r-q)T - \frac{1}{2}\sigma^2 T, \sigma^2 T)$$

and

$$E^Q[S_T] = S_0 e^{(r-q)T} = F_{0,T} \tag{5.23}$$

where $F_{0,T}$ is the no-arbitrage, futures price at time 0 for a contract to exchange at time T.

The risk-neutral density of S_T then depends on three parameters $(F_{0,T}, \sigma, T)$ and is given by the lognormal density

$$\psi(x|F_{0,T}, \sigma, T) = \frac{1}{x\sigma\sqrt{2\pi T}} e^{-\frac{1}{2}\left(\frac{log(x) - \left[log(F_{0,T}) - \frac{1}{2}\sigma^2 T\right]}{\sigma\sqrt{T}}\right)^2}. \tag{5.24}$$

Similarly, a risk-neutral, lognormal density from the HAR-RV model can be given by replacing $\sigma\sqrt{T}$ by a term $\widehat{RV}_T$ to give:

$$\psi\left(x \middle| F_{0,T}, \widehat{RV}_{t,t+h}\right) = \frac{1}{x\sqrt{2\pi\widehat{RV}_{t,t+h}}} e^{-\frac{1}{2}\left(\frac{log(x)-\left[log(F_{0,T})-\frac{1}{2}\widehat{RV}_{t,t+h}\right]}{\sqrt{\widehat{RV}_{t,t+h}}}\right)^2} \quad (5.25)$$

The quantity $\widehat{RV}_{t,t+h}$ is calculated from (5.18) and (5.21) with the horizon h (measured in trading days) and maturity T (measured in years).

5.2.4 Nonparametric transformations

The risk-neutral, Q-densities are not satisfactory specifications of the real-world densities. One reason is that Q-variance obtained from option prices is usually higher than the real-world variance, because there is a negative volatility risk premium (Carr and Wu, 2009). Consequently there are fewer observations than predicted in the tails of the Q-densities. A second reason is that the equity risk premium is, by definition, absent from all the risk-neutral densities. Hence it is necessary to use some technique to transform risk-neutral densities into real-world densities.

We consider the nonparametric calibration method in this study. Nonparametric calibration functions are re-estimated for each period t. At time t (which counts trading days), the nonparametric transformation for a selected horizon h is determined by a set of t-h+1 cumulative, risk-neutral probabilities

$$u_{s+h} = G_{Q,s,T}(S_{s+h}|\Theta_s), \qquad 0 \le s \le t-h, \quad (5.26)$$

with T (years) matching h (trading days), s a time before t-h+1, $G_{Q,s,T}$ the cumulative distribution function of the price S_{s+h}, and with Θ_s a vector of density parameters. We assume the observations u_{s+1} are i.i.d. and their c.d.f. is given by the calibration function $C_T(u)$.

The values of the variables u for the Heston model are given by (5.13). The variables u for the HAR-RV model can be derived in the following way. For the risk-neutral dynamics,

$$log(S_{s+h}) \sim N(log(F_{s,s+h}) - \frac{1}{2}\widehat{RV}_{s,s+h}, \widehat{RV}_{s,s+h})$$

with $F_{s,\,s+h}$ the futures price at time s for a transaction at time $s+h$ and with $\widehat{RV}_{s,s+h}$ the forecast of RV for the period from time s to $s+h$ inclusive. From the outcome $\hat{S}_{s+h}$ we calculate

$$\begin{aligned} u_{s+h} &= G_{Q,s,T}\left(\hat{S}_{s+h}\middle|\Theta_s\right) \\ &= \Phi\left(\frac{log(\tilde{S}_{s+h}) - (log(F_{s,s+h}) - \frac{1}{2}\widehat{RV}_{s,s+h})}{\sqrt{\widehat{RV}_{s,s+h}}}\right). \end{aligned} \quad (5.27)$$

The values of the variables u for the Black-Scholes model are given in a similar way[6]

$$u_{s+h} = \Phi\left(\frac{log(\tilde{S}_{s+h}) - (log(\tilde{F}_{s,s+h}) - \frac{1}{2}\sigma^2 T)}{\sigma\sqrt{T}}\right) \quad (5.28)$$

We use $\varphi()$ and $\Phi()$ to represent the density and the c.d.f. of the standard normal distribution. We then transform the observations u_i, whose domain is from 0 to 1, to new variables $y_i=\Phi^{-1}(u_i)$, and then fit a nonparametric kernel c.d.f. to the set $\{y_1, y_2, ..., y_{t-h+1}\}$. We use a normal kernel with bandwidth B to obtain the kernel density and c.d.f.:

6 When calculating densities and variables u, we use forward prices on day s for future transactions at time $s+h$.

$$\hat{h}_T(y) = \frac{1}{(t-h+1)B} \sum_{i=1}^{t-h+1} \varphi\left(\frac{y-y_i}{B}\right),$$

$$\hat{H}_T(y) = \frac{1}{t-h+1} \sum_{i=1}^{t-h+1} \Phi\left(\frac{y-y_i}{B}\right). \tag{5.29}$$

The bandwidth B decreases as t increases. We apply the standard formula of Silverman (1986), where $B=0.9\sigma_y/t^{0.2}$ and σ_y is the standard deviation of the terms y_i.

The empirical calibration function is then

$$\hat{C}_T(u) = \hat{H}_T\left(\Phi^{-1}(u)\right) \tag{5.30}$$

which is calculated at time *t*. At the same time, we let $g_{Q,T}(x)$ and $G_{Q,T}(x)$ denote the risk-neutral density and the cumulative distribution function of the random variable S_T. We define $u_T = G_{Q,T}(S_T)$. We follow Bunn (1984) and denote the calibration function $C_T(u)$, which is the real-world c.d.f. of the random variable u_T. Now we consider the real world c.d.f. of S_T, with *Pr* referring to the real world probabilities. The c.d.f. is

$$\Pr(S_T \le x) = \Pr\left(G_{Q,T}(S_T) \le G_{Q,T}(x)\right) = \Pr\left(u_T \le G_{Q,T}(x)\right)$$
$$= C_T\left(G_{Q,T}(x)\right) \tag{5.31}$$

Consequently replacing $C_T(.)$ by $\hat{C}_T(.)$, the predictive real-world c.d.f. of S_T is

$$G_{P,T}(x) = \hat{C}_T\left(G_{Q,T}(x)\right) \tag{5.32}$$

The real-world density is

$$g_{P,T}(x) = \frac{d}{dx}\widehat{H}_T\left(\Phi^{-1}\left(G_{Q,T}(x)\right)\right) = \frac{d}{dx}\widehat{H}_T(y) = \frac{dy}{dx}\frac{d\widehat{H}_T(y)}{dy}$$

$$= \frac{du}{dx}\frac{dy}{du}\hat{h}_T(y) = \frac{g_{Q,T}(x)\hat{h}_T(y)}{\varphi(y)}. \tag{5.33}$$

Also the nonparametric calibration density is

$$\hat{c}_T(u) = \frac{d}{du}\hat{C}_T(u) = \frac{d}{du}\widehat{H}_T(y) = \frac{d\widehat{H}_T(y)}{dy}\frac{dy}{du} = \frac{\hat{h}_T(y)}{\varphi(y)}. \tag{5.34}$$

5.2.5 Parameter estimation

The densities are all evaluated out-of-sample and thus the parameter values are obtained ex ante, i.e. the values at time t are estimated based on the information available at time t. For the HAR variances we estimate all parameters from regressions over five-year windows. For Black-Scholes lognormal densities, we use the nearest-the-money, nearest-to-expiry option implied volatility.

For the Heston model, we estimate the risk-neutral parameters of the asset price dynamics every day. On each day, we estimate the initial variance V_t, the rate of reversion κ_t, the unconditional expectation of stochastic variance θ_t, the volatility of volatility σ_t, and the correlation ρ_t between the two Wiener processes. Assume there are N_t European, call[7] option contracts traded on day t, denoted by $i=1, \ldots, N_t$, and the market prices are $c_{t,i}$, for strike prices $K_{t,i}$, and expiry times $T_{t,i}$. We also assume $p_{t,i}$ is the futures price for the asset, calculated for a synthetic futures contract which expires in $T_{t,i}$ years. Then we calibrate the five risk-neutral Heston parameters by minimising

[7] We use put-call parity to obtain the equivalent European call prices from the put prices, and then apply them to (5.5), this is also discussed in section 5.3.2.

the total squared errors in

$$\sum_{i=1}^{N_t}(c_{t,i} - c(p_{t,i}, K_{t,i}, T_{t,i}, V_t, \kappa_t, \theta_t, \sigma_t, \rho_t))^2 \tag{5.35}$$

with $c(.)$ the solution for the European call option price from the Heston model given in (5.5).[8]

Christoffersen and Jacobs (2004) argue that the loss function used in parameter estimation and model evaluation should be the same for any given model, and the estimation loss function should be identical when comparing across models. Different loss functions are used in the estimation and evaluation stages in the literature. Bakshi et al. (1997) use mean-squared absolute option pricing errors in estimation, but both mean-squared absolute and relative option pricing errors in evaluation. Rosenberg and Engle (2002) employ mean-squared absolute option pricing errors in estimation, but relative hedging errors in the evaluation stage. Pan (2002) estimates parameters using generalised method of moments (GMM) loss function and evaluate models using implied volatility mean squared errors. Chernov and Ghysels (2000) use efficient method of moments (EMM) in estimation, and both mean-squared absolute and relative option pricing errors in evaluation. Benzoni (2002) employs EMM and mean-squared absolute option pricing errors in estimation, and mean-squared absolute option pricing errors in the evaluation stage.

[8] Christoffersen and Jacobs (2004) conclude that it is a "good general-purpose loss function in option valuation applications". Christoffersen et al. (2010) also employed it in the study of S&P 500 dynamics.

5.2.6 Econometric methods

5.2.6.1 Maximum log-likelihood

There are several ways to evaluate density forecasts, and we will use the standard log-likelihood criterion previously employed by Bao et. al (2007), Liu et. al (2007) and Shackleton et al. (2010). For a given horizon h, assuming method m gives densities $g_{m,t}(x)$ at times $i, .., j$ for the asset price at times $i+h, ..., j+h$. Our goal is to find the method which maximises the out-of-sample log-likelihood of observed asset prices, and this log-likelihood for method m is given by

$$L_m = \sum_{t=i}^{j} log\left(g_{m,t}(S_{t+h})\right) \tag{5.36}$$

To compare two methods we apply a version of the log-likelihood ratio test in Amisano and Giacomini (2007). The null hypothesis states that two different density forecasting methods m and n have equal expected log-likelihood. The test is based on the log-likelihood differences

$$d_t = log\left(g_{m,t}(S_{t+h})\right) - log\left(g_{n,t}(S_{t+h})\right), \qquad i \leq t \leq j. \tag{5.37}$$

Amisano and Giacomini (2007) follow Diebold and Mariano (1995) and add the assumption that the differences are uncorrelated and ignore all covariance terms in the estimator. Hence the AG test statistic is

$$t_{i,j} = \frac{\bar{d}}{s_d/\sqrt{(j-i+1)}} = \frac{L_m - L_n}{s_d\sqrt{(j-i+1)}} \tag{5.38}$$

This statistic follows a standard normal distribution, where $\bar{d}$ is the mean and s_d is the standard deviation of the terms d_t.

When h>1 the forecasts overlap and it is plausible to expect some autocorrelation in the differences. A Newey-West adjustment should then be made when estimating the variance of $\bar{d}$. Assuming the terms d_t are stationary,

$$var(\bar{d}) = var\left(\frac{d_1 + d_2 + \cdots + d_n}{n}\right)$$

$$= \frac{1}{n^2}[nvar(d_1) + 2(n-1)cov(d_1, d_2) + \cdots + 2cov(d_1, d_n)]$$

$$= \frac{var(d_1)}{n}\left[1 + 2\left(\frac{n-1}{n}\right)\rho_1 + 2\left(\frac{n-2}{n}\right)\rho_2 + \cdots + 2\left(\frac{1}{n}\right)\rho_{n-1}\right]$$

where the autocorrelations are $\rho_\tau = cor(d_t, d_{t+\tau})$. The typical estimate of the variance of $\bar{d}$ is

$$\frac{s_d^2}{n}[1 + 2\omega_1\hat{\rho}_1 + \cdots + 2\omega_k\hat{\rho}_k]$$

and a standard set of weights for k estimated autocorrelations is $\omega_\tau = \frac{k+1-\tau}{k+1}$, $1 \leq \tau \leq k$.

5.2.6.2 Diagnostic tests

Appropriate diagnostic tests use properties of time series derived from density forecasts. Rosenblatt (1952) introduces the probability integral transform, and states that the PIT values are i.i.d. uniform for known densities. Diebold et al. (1998) initiated the idea of using PIT values to evaluate density forecasts. Following this and Shackleton et al. (2010), we also employ a series of observed cumulative probabilities to check the accuracy of the forecasts. For a given method *m* the PIT probabilities are given by

$$u_{t+1} = \int_0^{S_{t+1}} g_{m,t}(x)dx, \tag{5.39}$$

for prices S_{t+1} matched with densities $g_{m,t}(x)$.

We then evaluate if the values of u are compatible with i.i.d. observations from the uniform distribution. We can employ the Kolmogorov and Smirnov test. The KS test checks the maximum difference between the empirical and theoretical cumulative functions. For forecasts made at times $i \le t \le j$, the sample c.d.f. of $\{u_{i+1}, \ldots, u_{j+1}\}$, evaluated at u, is the proportion of values less than or equal to u, i.e.

$$\tilde{C}(u) = \frac{1}{j-i+1} \sum_{t=i+1}^{j+1} S(u-u_t) \tag{5.40}$$

with $S(x)$=1 if $x \ge 0$, and $S(x)$=0 if $x < 0$. The test statistic is given by

$$KS = \sup_{0 \le u \le 1} \left| \tilde{C}(u) - u \right|. \tag{5.41}$$

The KS test is widely applied because it is easy to implement. However, one needs to be cautious when interpreting the test results, as the KS test checks for uniformity under the i.i.d. assumption rather than tests i.i.d. and uniformity jointly.

Some researchers doubt the power of the KS test when evaluating density forecasts. Berkowitz (2001) invented the BK test, which states that if the PIT is i.i.d. uniform, then the normal inverse cumulative function of the PIT is i.i.d. normal. The advantage of the BK test is that it can test independence and uniformity jointly. The BK test has been applied in Clements and Smith (2000), Clements (2004), Guidolin and Timmermann (2005) and Shackleton et al. (2010).

The BK method transforms the observations u_i to new variables $y_i = \Phi^{-1}(u_i)$, with $\Phi()$ the c.d.f. of the standard normal distribution. The null hypothesis of the test is that the values of y are i.i.d. and follow a standard normal distribution, against the alternative hypothesis that y is a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. Let

$$y_t - \mu = \rho(y_{t-1} - \mu) + \varepsilon_t. \tag{5.42}$$

Then the null hypothesis is that $\mu = 0$, $\rho = 0$, and $var(\varepsilon_t) = 1$. The log-likelihood for T observations from (5.42) is

$$-\frac{T}{2}log(2\pi) - \frac{1}{2}log[\sigma^2/(1-\rho^2)] - \frac{(y_1 - \mu/(1-\rho))^2}{2\sigma^2/(1-\rho^2)} - \frac{T-1}{2}log(\sigma^2)$$

$$-\sum_{t=2}^{T}\left(\frac{(y_t - \mu - \rho y_{t-1})^2}{2\sigma^2}\right) \tag{5.43}$$

Here σ^2 is the variance of ε_t. The log-likelihood is written as a function of the unknown parameters of the model, $L(\mu, \sigma^2, \rho)$. The log-likelihood ratio test (LR3) is

$$LR_3 = -2(L_0 - L_1) = -2\big(L(0,1,0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})\big). \tag{5.44}$$

Here hats denote maximum-likelihood values, L_0 and L_1 are the maximum log-likelihoods for the null and alternative hypotheses, and the test statistic has an asymptotic x_3^2 distribution. One disadvantage of the BK test is that models cannot be easily compared if they are all accepted or rejected. The AG test, which we discussed before, compares the log-likelihoods between models and solves this problem.

5.3 Data

5.3.1 Option data

We investigate a majority of the Dow Jones Industrial Average (DJIA) 30 Index stocks for 10 years from 1st January 2003 to 31st December 2012. The calculations are time consuming and consequently we report results for only 17 stocks which we find are sufficient to obtain clear conclusions. Table 5.1 lists the stocks studied, which were all DJIA constituents at the end of our sample period.

The option data are obtained from Ivy DB OptionMetrics, which includes price information for all U.S. listed equity options, based on daily closing quotes at the CBOE. The OptionMetrics database also includes information about end-of-day security prices and zero-coupon interest rate curves. The security price file provides the closing price for each security on each day from CRSP.

5.3.2 Option prices

In terms of filtering option price records, we follow the criteria of Carr and Wu (2003, 2009 and 2010) and Huang and Wu (2004). We delete an option record when the bid price is zero or negative. We also delete an option record when the bid price is greater than the ask price. As do Carr and Wu (2009), we eliminate all the options which have maturity equal to or more than one year. Following Carr and Wu (2003), Huang and Wu (2004), Shackleton et al. (2010) and Taylor et al. (2010), we delete all data for options with maturity equal to or less than seven calendar or five business days.

Table 5.1

List of 17 DJIA constituent stocks studied.

Number	Company	Exchange	Symbol	Industry	Date Added
1	Alcoa	NYSE	AA	Aluminum	1959/6/1
2	American Express	NYSE	AXP	Consumer finance	1982/8/30
3	AT&T	NYSE	T	Telecommunication	1999/11/1
4	Boeing	NYSE	BA	Aerospace and defense	1987/3/12
5	Cisco Systems	NASDAQ	CSCO	Computer networking	2009/6/8
6	General Electric	NYSE	GE	Conglomerate	1907/11/7
7	Hewlett-Packard	NYSE	HPQ	Computers & technology	1997/3/17
8	The Home Depot	NYSE	HD	Home improvement retailer	1999/11/1
9	Intel	NASDAQ	INTC	Semiconductors	1999/11/1
10	IBM	NYSE	IBM	Computers & technology	1979/6/29
11	Johnson & Johnson	NYSE	JNJ	Pharmaceuticals	1997/3/17
12	JPMorgan Chase	NYSE	JPM	Banking	1991/5/6
13	McDonald's	NYSE	MCD	Fast Food	1985/10/30
14	Merck	NYSE	MRK	Pharmaceuticals	1979/6/29
15	Pfizer	NYSE	PFE	Pharmaceuticals	2004/4/8
16	Wal-Mart	NYSE	WMT	Retail	1997/3/17
17	Walt Disney	NYSE	DIS	Broadcasting and entertainment	1991/5/6

All the equity options are American. OptionMetrics provides implied volatilities, calculated from binomial trees which incorporate dividends and permit early exercise. We use equivalent European option prices defined by assuming the European and American implied volatilities are equal. This method assumes the early exercise premium can be obtained from constant volatility pricing models. The assumption is particularly reasonable for out-of-the-money options which have small early exercise premia.

European call and put prices for the same strike and maturity theoretically contain the same information. Either the call option or the put option will be out-of-the-money (OTM), or under rare circumstances both are at-the-money (ATM). Options are ATM when the strike price equals the stock price (S=K), calls are OTM when $S<K$ and puts are OTM when $S>K$; they are nearest-the-money if $|S-K|$ is nearer zero than for all other contemporaneous strikes. We choose to use the information given by the prices of OTM and ATM options only, because in-the-money (ITM) options are less liquid and have higher early exercise premia. We use put-call parity to obtain equivalent European call prices from the European OTM put prices.

5.3.3 Interest rates

We follow Taylor et al. (2010) to get the interest rate corresponding to each option's expiry by linear interpolation of the two closest zero-coupon rates supplied by Ivy DB.

5.3.4 IBM example

We use IBM to illustrate our data and results. A total of 109,111 option prices are investigated in our sample period for IBM stock. The average number of option prices used per day is 44, consisting of 19 OTM calls and 25 OTM puts. Table 5.2 presents the quantity, moneyness and maturity of the option contracts used in this paper.

5.3.5 Futures prices

We calculate synthetic futures prices, which have the same expiry dates as the options, as the future value of the current spot price minus the present value of all the dividends expected during the life of the futures contract until the option expiry time T, i.e.

$$F_{0,T} = e^{rT}\big(S - PV(dividends)\big) \tag{5.45}$$

We use the actual dividends amount in the distribution file from OptionMetrics.

5.3.6 High-frequency stock prices

We use the transaction prices of DJIA 30 Index stocks for ten years during the period between 1st January 1998 and 31st December 2012. The data are obtained from pricedata.com. The prices provided are the last prices in one-minute intervals. After an inspection of the high-frequency data, we find a number of problematic days which do not have complete trading records. We set the price equal to that for the previous minute when there is a missing record, and we delete a day when there are more than

Table 5.2

Summary statistics for IBM option data. Information about out-of-the-money (OTM) and at-the-money (ATM) options on IBM stock from 2003 to 2012.

	Total	Average per day	Maximum per day	Minimum per day
Calls	47709	19	46	6
Puts	61402	25	74	5
Total	109111	44	115	12

	Maturity	<1 month	Between 1 and 6 months	>6 months	Subtotal
Moneyness	S/K				
Deep OTM put	>1.05	6462	30100	13596	50158
		(5.92%)	(27.59%)	(12.46%)	(45.97%)
OTM put	1.01-1.05	2040	5123	1839	9002
		(1.87%)	(4.70%)	(1.69%)	(8.25%)
At/near the money	0.99-1.01	1049	2641	973	4663
		(0.96%)	(2.42%)	(0.89%)	(4.27%)
OTM call	0.95-0.99	2278	5733	2330	10341
		(2.09%)	(5.25%)	(2.14%)	(9.48%)
Deep OTM call	<0.95	3168	20393	11386	34947
		(2.90%)	(18.69%)	(10.44%)	(32.03%)
Subtotal		14997	63990	30124	109111
		(13.74%)	(58.65%)	(27.61%)	(100.00%)

40 consecutive missing prices. The days deleted are usually close to holidays such as New Year's Day, Easter, Independence Day, Thanksgiving Day and Christmas.

Between 2003 and 2012, 17 days are deleted because of missing high-frequency prices and these days usually only have prices for half a day. There are also 8 days with unsatisfactory option price data. All 25 days are deleted from the high-frequency and option files leaving a sample of 2488 days for each firm for the ten-year period ending on 31st December 2012.

The stocks are traded for six-and-a-half-hours, from 9:30 EST to 16:00 EST. We calculate realised variances from 5-minute returns because Bandi and Russell (2006) state that the 5-minute frequency provides a satisfactory trade-off between maximising the accuracy of volatility estimates and minimising the bias from microstructure effects. As usual, returns are changes in log prices. We have 77 5-minute intraday returns for each day after deleting the data in the first five minutes to avoid any opening effects. The realised variance for day t is the sum of the squares of the 5-minute returns $r_{t,i}$:

$$RV_t = \sum_{i=1}^{77} r_{t,i}^2. \tag{5.46}$$

However, this calculation of realised variance is downward biased as a measurement of close-to-close volatility over a 24-hour period. This is because we only include the information during the trading period when we calculate the realised variance for a day, so the variation overnight (from close-to-open) is excluded. We thus need to scale the realised variance up. We multiply forecasts from the HAR-RV model by a scaling factor. The denominator of the scaling factor is the sum of the squares of the 5-minute returns representing the open market period, while the numerator of the scaling factor is the sum of the squares of the daily returns representing open and closed market periods. We use a rolling window for the scaling factor, hence if we forecast the realised variance on day t, then we use the information about returns up to and including day t to calculate

$$\widehat{RV}_{t,t+h}\left(\frac{\sum_{i=1}^{t} r_t^2}{\sum_{i=1}^{t}\sum_{j=1}^{77} r_{t,j}^2}\right).$$

This quantity replaces $\widehat{RV}_{t,t+h}$ in (5.25) when the high-frequency, lognormal densities are evaluated.

5.4 Empirical results

5.4.1 Heston risk-neutral parameters

Table 5.3 shows the summary statistics for risk-neutral parameters calibrated for IBM and across all stocks for each day in our sample period. The risk-neutral parameters minimise the mean squared error (MSE) of option prices on each day.

For IBM, our median estimate of the stochastic variance θ is 0.3457, equivalent to an annualized volatility level of 58.80%. The mean estimate of the rate of reversion κ is 1.6861, for which the half-life parameter of the variance process is then about 5 months. The median estimate of the volatility of volatility parameter σ which controls the kurtosis of returns is 0.8617. Also the median estimate of the correlation

Table 5.3

Summary statistics for risk-neutral calibrated parameters for IBM and across all stocks. Estimates are summarised for the risk-neutral dynamics (5.2). The parameters are estimated each day from 2003 to 2012, from the OTM and ATM options, through minimising the MSE of the fitted option prices. We apply the constraints $0 \leq \kappa \leq 36$, $0 \leq \theta \leq 1$, $\sigma \geq 0$, $-1 \leq \rho \leq 1$, $0 \leq v_0 \leq 1$.

	κ	θ	σ	ρ	v_0
IBM					
Mean	1.6861	0.5042	1.2038	-0.6723	0.0653
Median	0.1661	0.3457	0.8617	-0.6652	0.0444
Standard deviation	3.6779	0.4201	2.1596	0.1051	0.0726
Averages across all firms					
Mean	3.0401	0.4037	1.9675	-0.6331	0.1081
Median	1.1136	0.2308	1.0267	-0.6305	0.0692
Standard deviation	5.2434	0.3594	5.6694	0.1462	0.1206

Figure 5.1 Plot of IBM Heston parameter κ from 2003 to 2012.

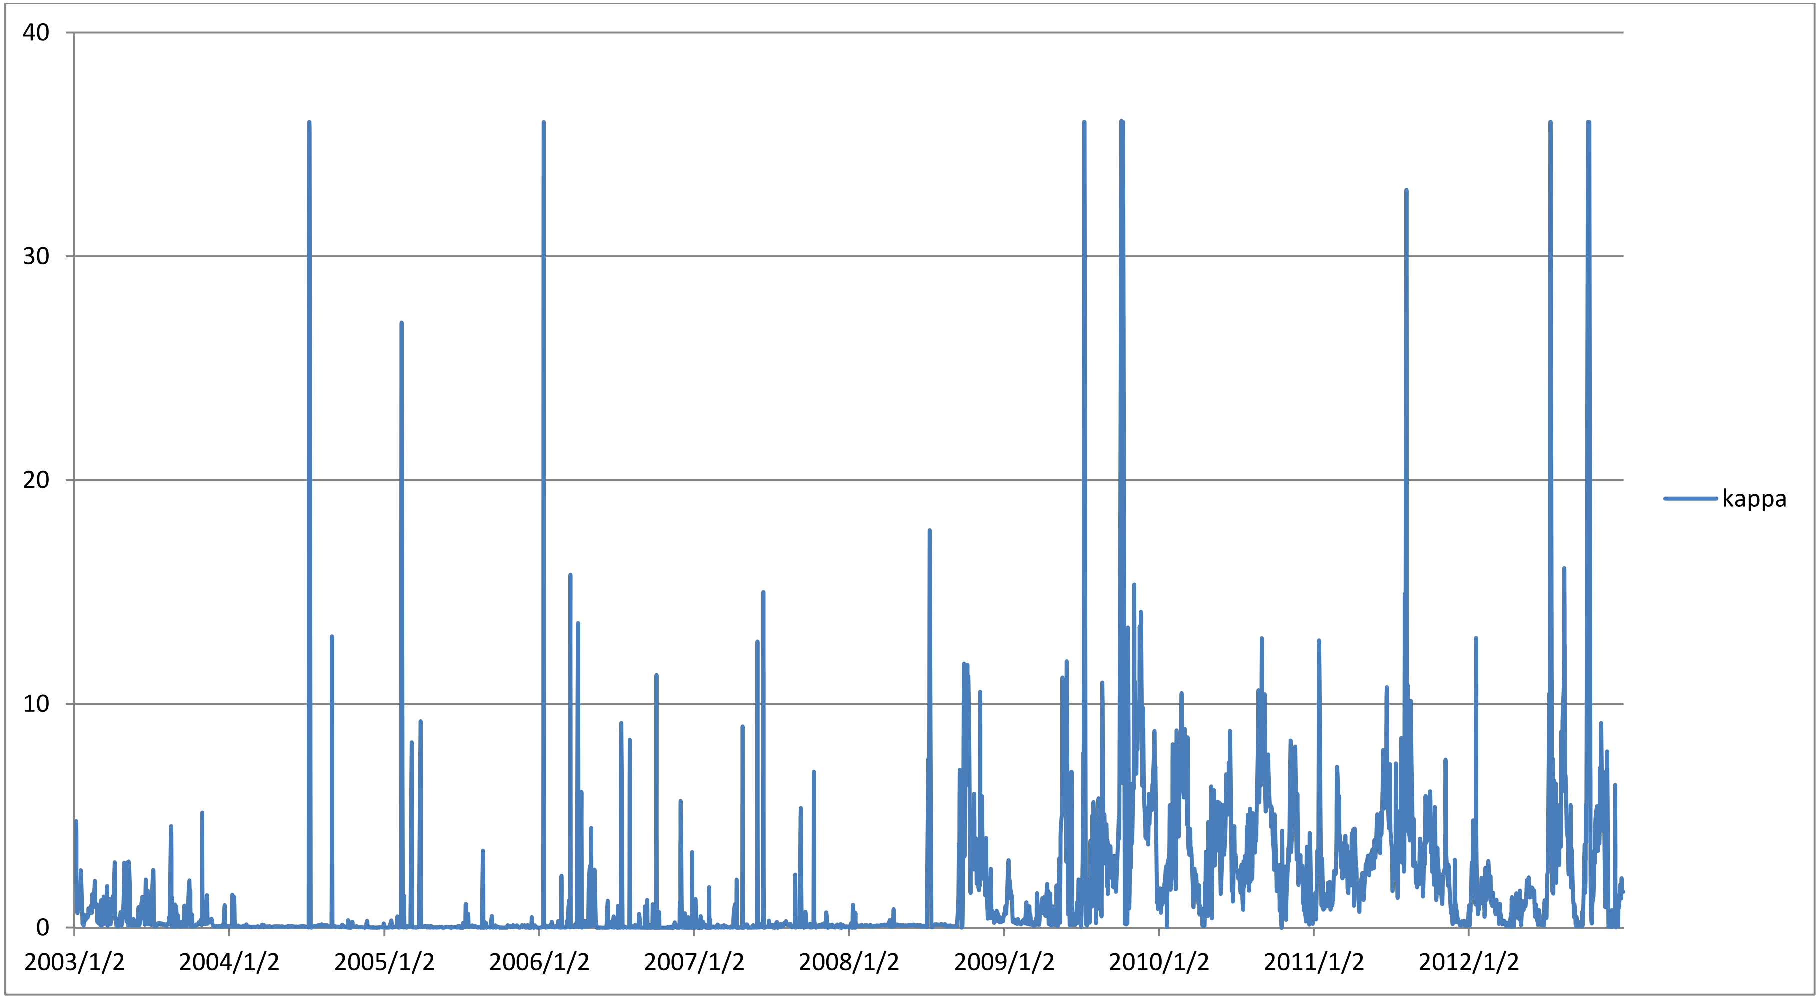

Figure 5.2 Plot of IBM Heston parameters θ and v_0 from 2003 to 2012.

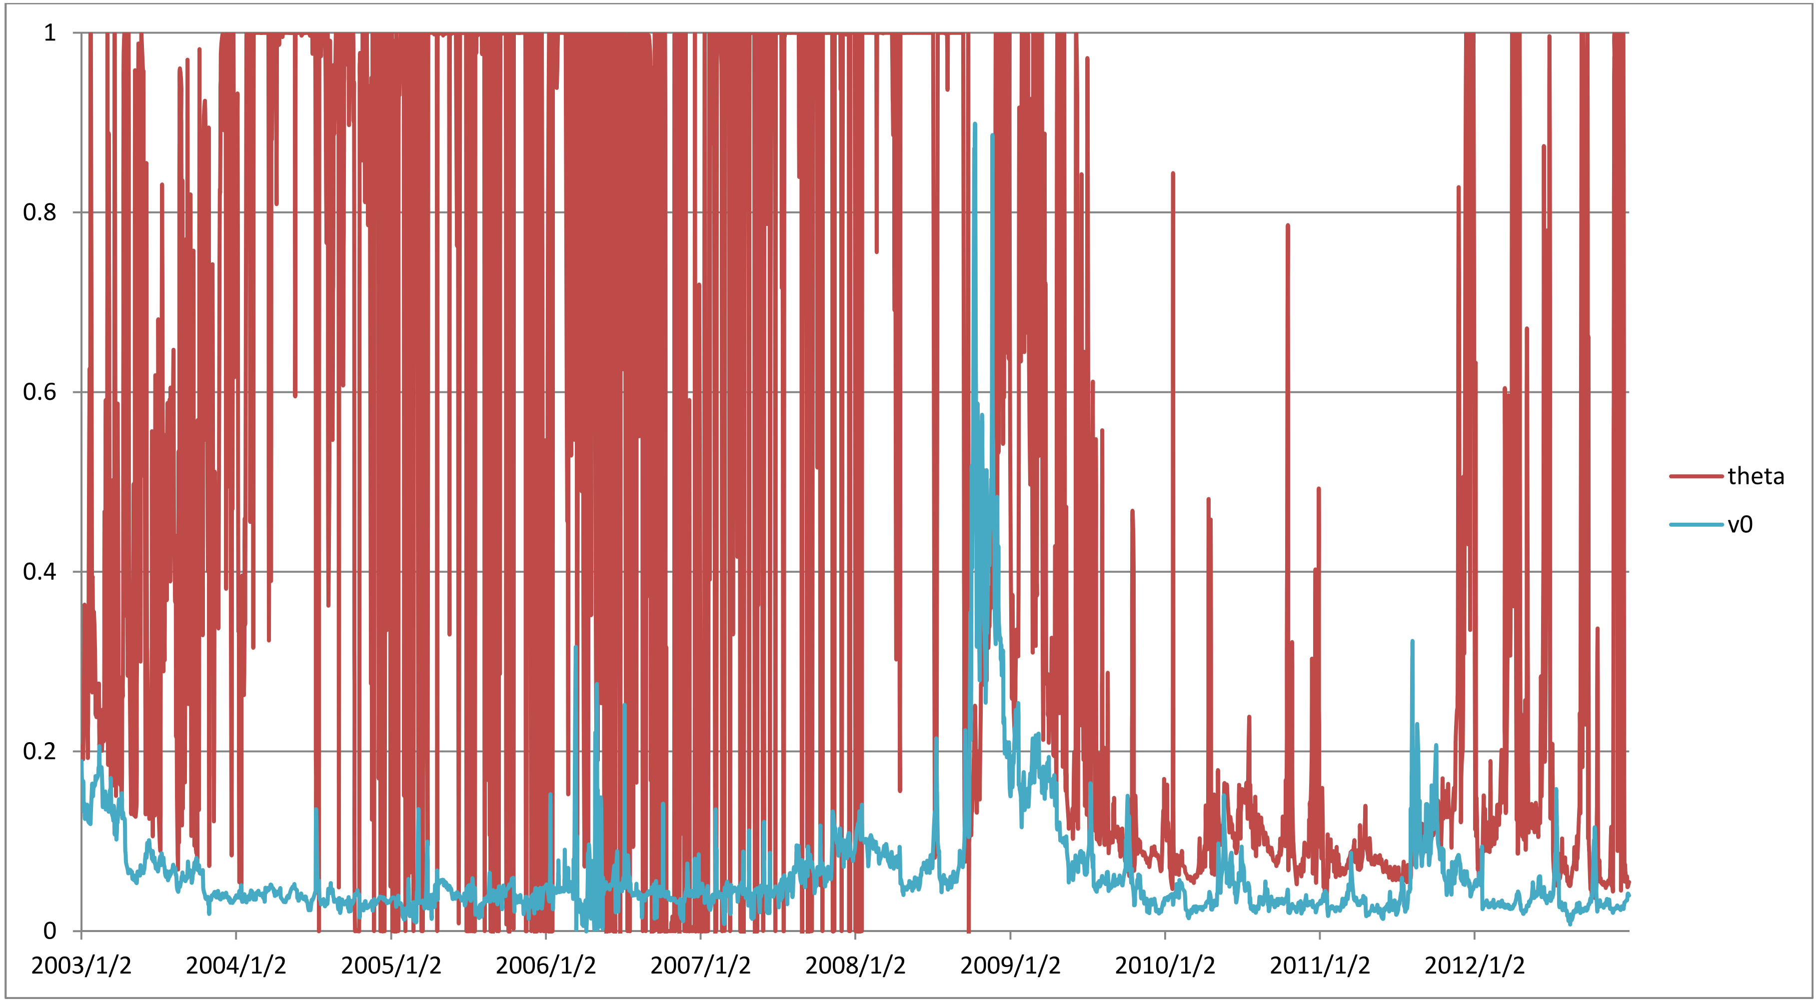

Figure 5.3 Plot of IBM Heston parameter σ from 2003 and 2012.

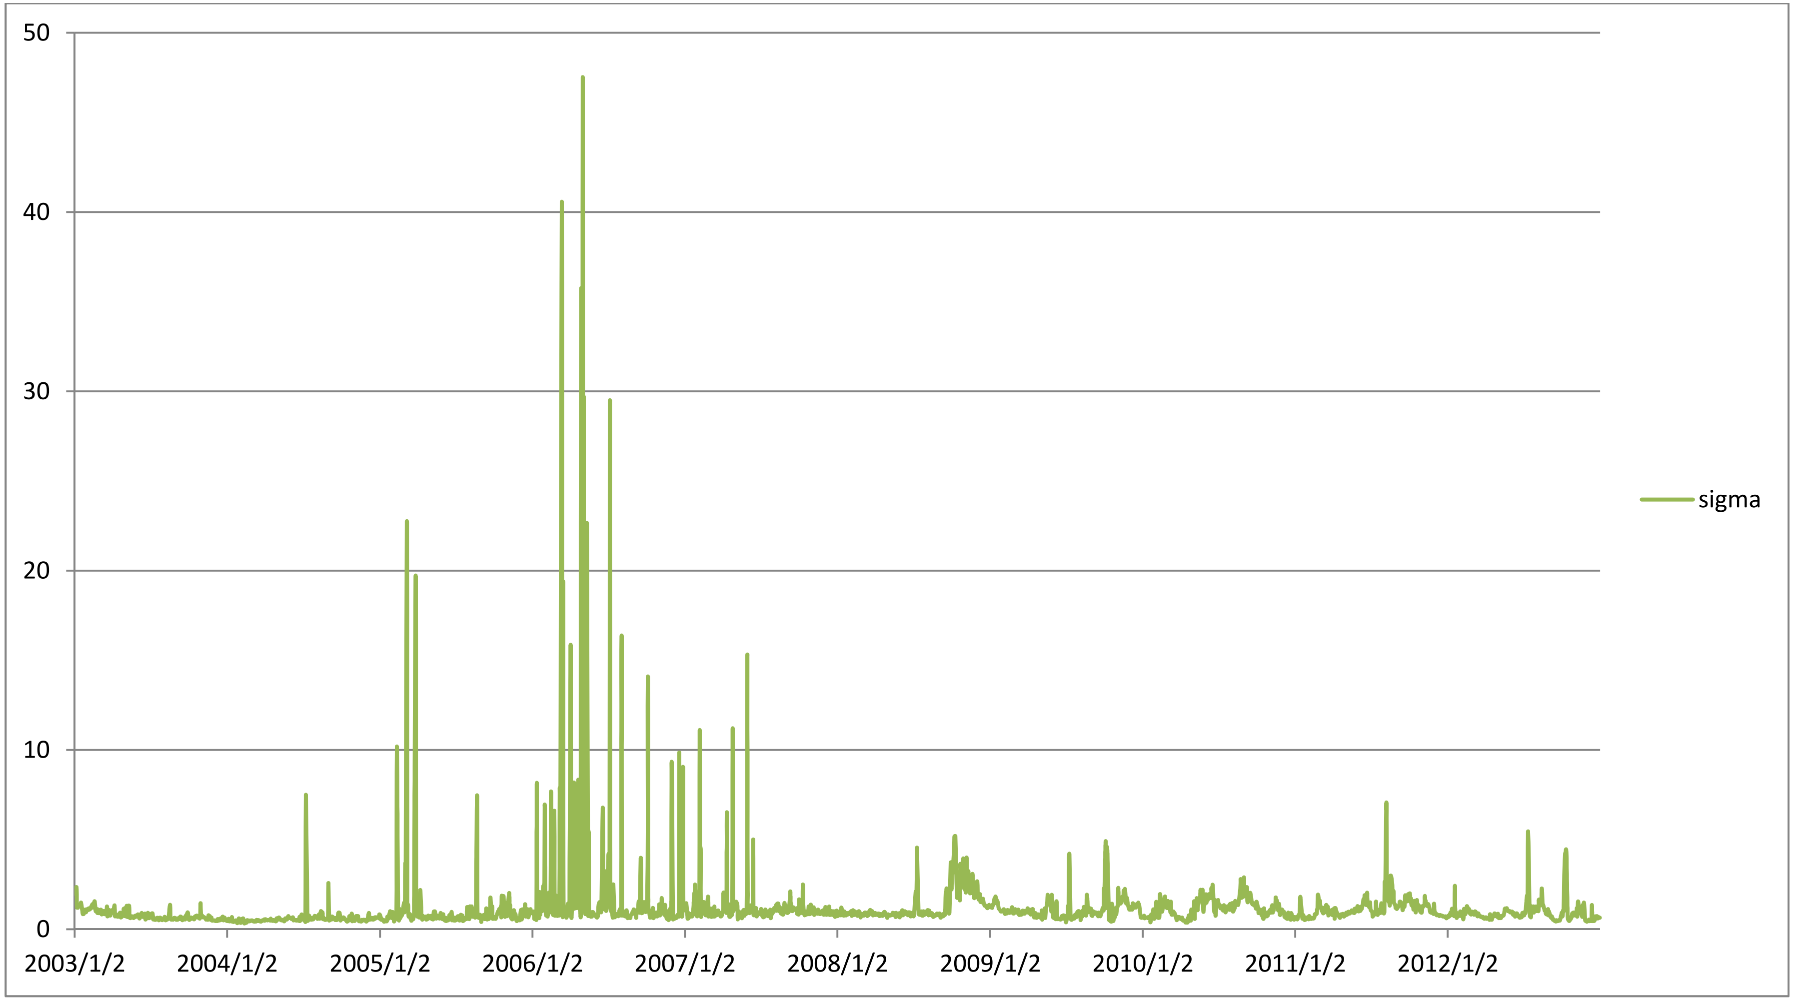

Figure 5.4 Plot of IBM Heston parameter ρ from 2003 to 2012.

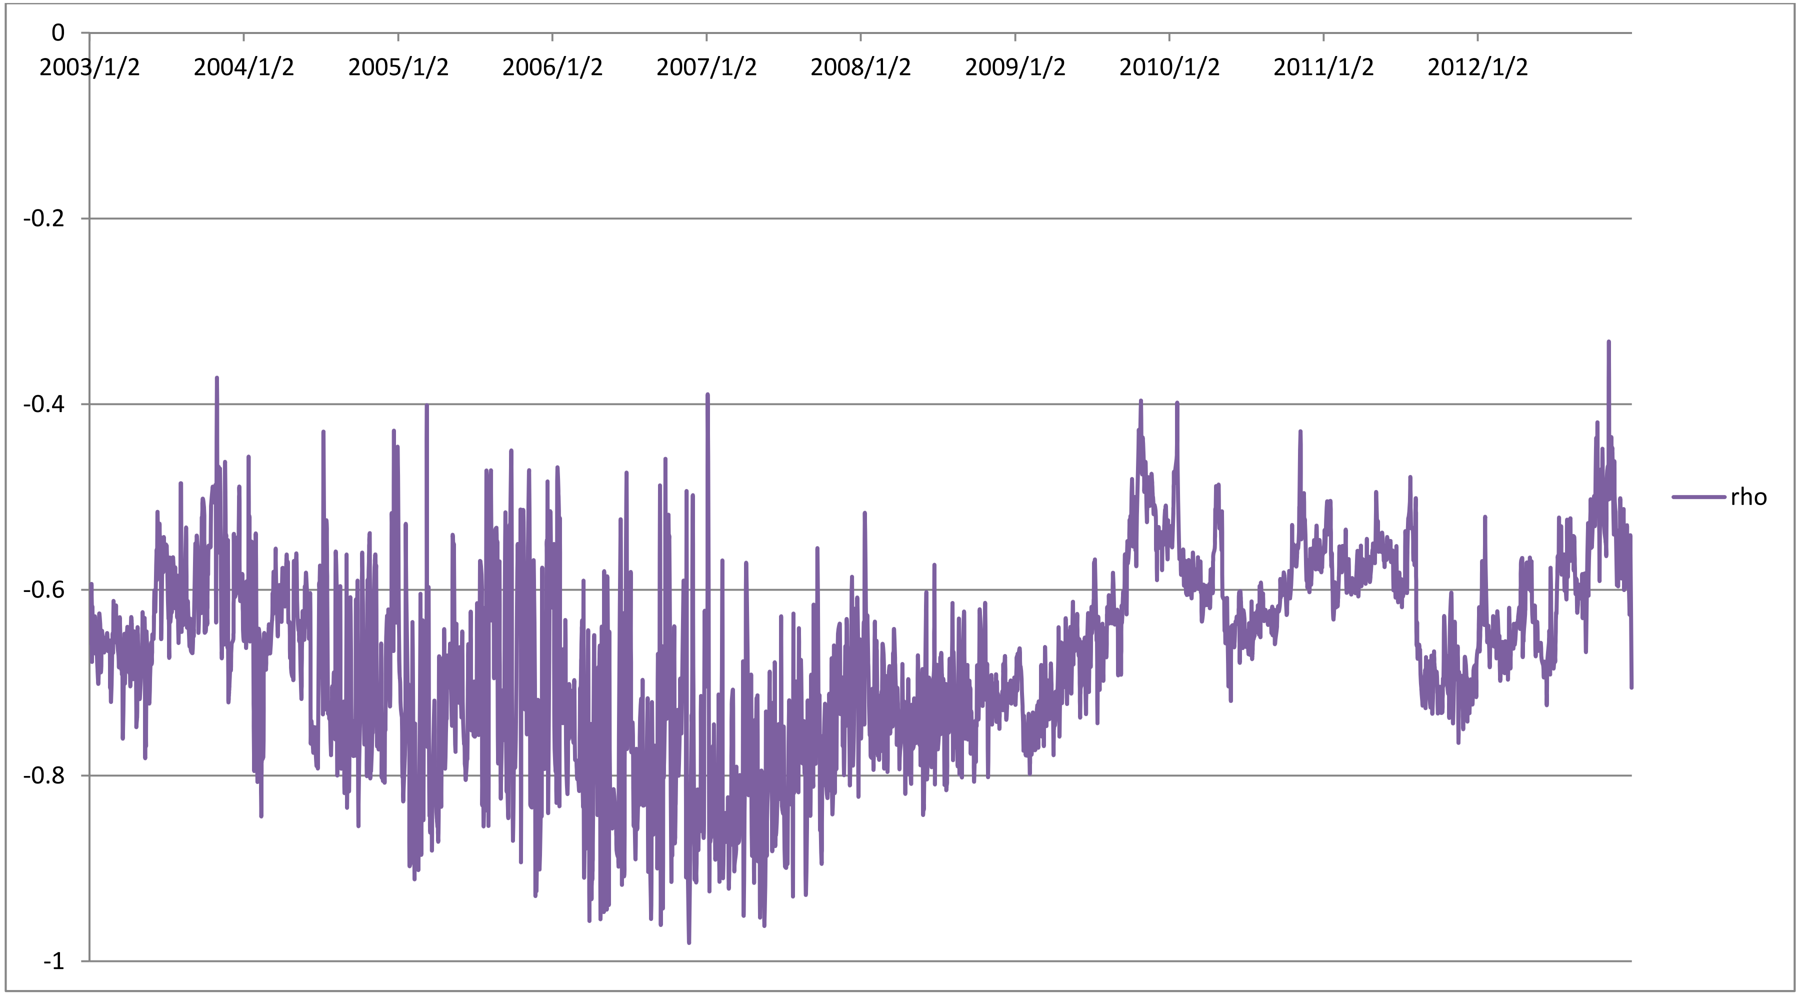

Table 5.4

Initial and calibrated parameters for IBM, estimated on five days from 2003 to 2012, from five different initial values.

κ	θ	σ	ρ	v_0	MSE
Initial Parameter					
4.9292	0.0505	0.9296	-0.6590	0.1898	
4.1528	0.0452	0.7925	-0.6624	0.1787	
3.8748	0.0421	0.6977	-0.6788	0.1877	
3.0920	0.0347	0.6400	-0.6795	0.1651	
2.0000	0.0100	0.1000	0.0000	0.0100	
Calibrated Parameter					
Day 1					
3.3211	0.2155	1.9444	-0.6689	0.1857	0.0238
3.3201	0.2155	1.9446	-0.6689	0.1858	0.0238
3.3224	0.2154	1.9440	-0.6690	0.1858	0.0238
3.3207	0.2154	1.9445	-0.6688	0.1858	0.0238
3.3214	0.2154	1.9443	-0.6690	0.1858	0.0238
Day 50					
0.0774	1.0000	0.8654	-0.6433	0.1342	0.0256
0.0774	1.0000	0.8654	-0.6433	0.1342	0.0256
0.0775	1.0000	0.8655	-0.6432	0.1342	0.0256
0.0774	1.0000	0.8655	-0.6432	0.1342	0.0256
0.0775	0.9999	0.8654	-0.6433	0.1342	0.0256
Day 100					
0.1050	1.0000	0.7849	-0.6612	0.0670	0.0135
0.1050	1.0000	0.7849	-0.6612	0.0670	0.0135
0.1052	0.9982	0.7850	-0.6612	0.0670	0.0135
0.1049	1.0000	0.7849	-0.6612	0.0670	0.0135
0.1050	0.9999	0.7848	-0.6612	0.0670	0.0135

Day 150					
κ	θ	σ	ρ	v_0	MSE
0.0518	0.9935	0.5303	-0.5889	0.0731	0.0068
0.0576	0.9005	0.5306	-0.5890	0.0731	0.0068
0.0623	0.8377	0.5309	-0.5890	0.0731	0.0068
0.0707	0.7478	0.5317	-0.5888	0.0731	0.0068
0.0531	0.9686	0.5299	-0.5893	0.0731	0.0068
Day 200					
0.1051	1.0000	0.7621	-0.5203	0.0461	0.0142
0.1058	0.9936	0.7620	-0.5203	0.0461	0.0142
0.1051	1.0000	0.7619	-0.5203	0.0461	0.0142
0.1056	0.9956	0.7620	-0.5203	0.0461	0.0142
0.1052	1.0000	0.7621	-0.5203	0.0461	0.0142

ρ is -0.6652, consistent with estimates in the literature. Time series plots of the five calibrated parameters are shown from Figures 5.1 to 5.4.

Given that considerable time variation is noted for the estimated Heston parameters, some possible reasons are discussed below. Firstly, the market may not use the Heston model to reflect stochastic volatility, hence it cannot explain why the parameters change. Additionally, the Matlab software may not be able to find the true minimum, thus it has more variation than it should. Moreover, the quality of the estimates can also be related to how many option records we have on each day, and we do not have fixed number of options during the sample period. We check our Matlab code against the Heston (1993) paper to make sure it is reliable. However, the potential non convergence of the search algorithm around the parameter space to a global minimum of the loss function may not be the true reason. As shown in Table 5.4, we start with five different initial values on five different days in our dataset, and we get similar calibrated parameters and the same MSE.

5.4.2 Examples of density forecasts

The one-day ahead risk-neutral Heston, lognormal and HAR densities for IBM calculated on January 2nd 2003 are shown in Figure 5.5. The Heston density is negatively skewed while the lognormal density is slightly positively skewed. The HAR density is seen to have less variance than the Heston and the lognormal densities. The one-month ahead risk-neutral Heston, lognormal and HAR densities for IBM calculated on January 2nd 2003 are shown in Figure 5.6 display similar properties. These densities are all risk-neutral because the expectation is equal to the futures.

5.4.3 Examples of cumulative probabilities and nonparametric transformations

The one-day ahead risk-neutral densities give the cumulative distribution functions $G_{Q,t}(x)$ for the next stock price p_{t+1}, and the observed risk-neutral probabilities $u_{t+1}=G_{Q,t}(p_{t+1})$ are not consistent with uniform probabilities, as expected. The sample cumulative probabilities $\tilde{C}(u)$ are calculated using (5.40), and the deviations between the sample c.d.f. and a uniform c.d.f., namely $\tilde{C}(u) - u$, are plotted in Figure 5.7 for IBM, for one-day-ahead forecasts obtained from the Heston model. We can observe from the figure that there are few observations u close to either zero or one; only 7.3% of the variables u are below 0.1 and only 5.1% of them are above 0.9. The KS test statistic is the maximum value of $|\tilde{C}(u) - u|$, which is equal to 7.1%, hence the null hypothesis of a uniform distribution is rejected at the 0.01% significance level. The shape of the curve may be explained by the fact that the historical volatility is lower than the risk-neutral volatility, hence the risk-neutral probabilities of large price changes exceed the real-world probabilities. The corresponding plot for IBM for one-day-ahead forecasts obtained from Black-Scholes and HAR lognormal densities are shown in Figures 5.8 and 5.9.

The nonparametric transformation of the probabilities u_{t+1} used in the calculation of the real-world density is calculated from (5.34). The calibration densities $\hat{c}(u)$, for one-day ahead HAR, Black-Scholes and Heston lognormal forecasts are shown in Figure 5.10; these densities use the values of u for all 10 years from 2003 to 2012. The purpose of the calibration is to create real-world densities which have uniformly distributed observed probabilities u_{t+1}. These calibration densities are not smooth at

Figure 5.5 Heston, lognormal and HAR one-day ahead risk-neutral density forecasts for IBM on January 2nd 2003.

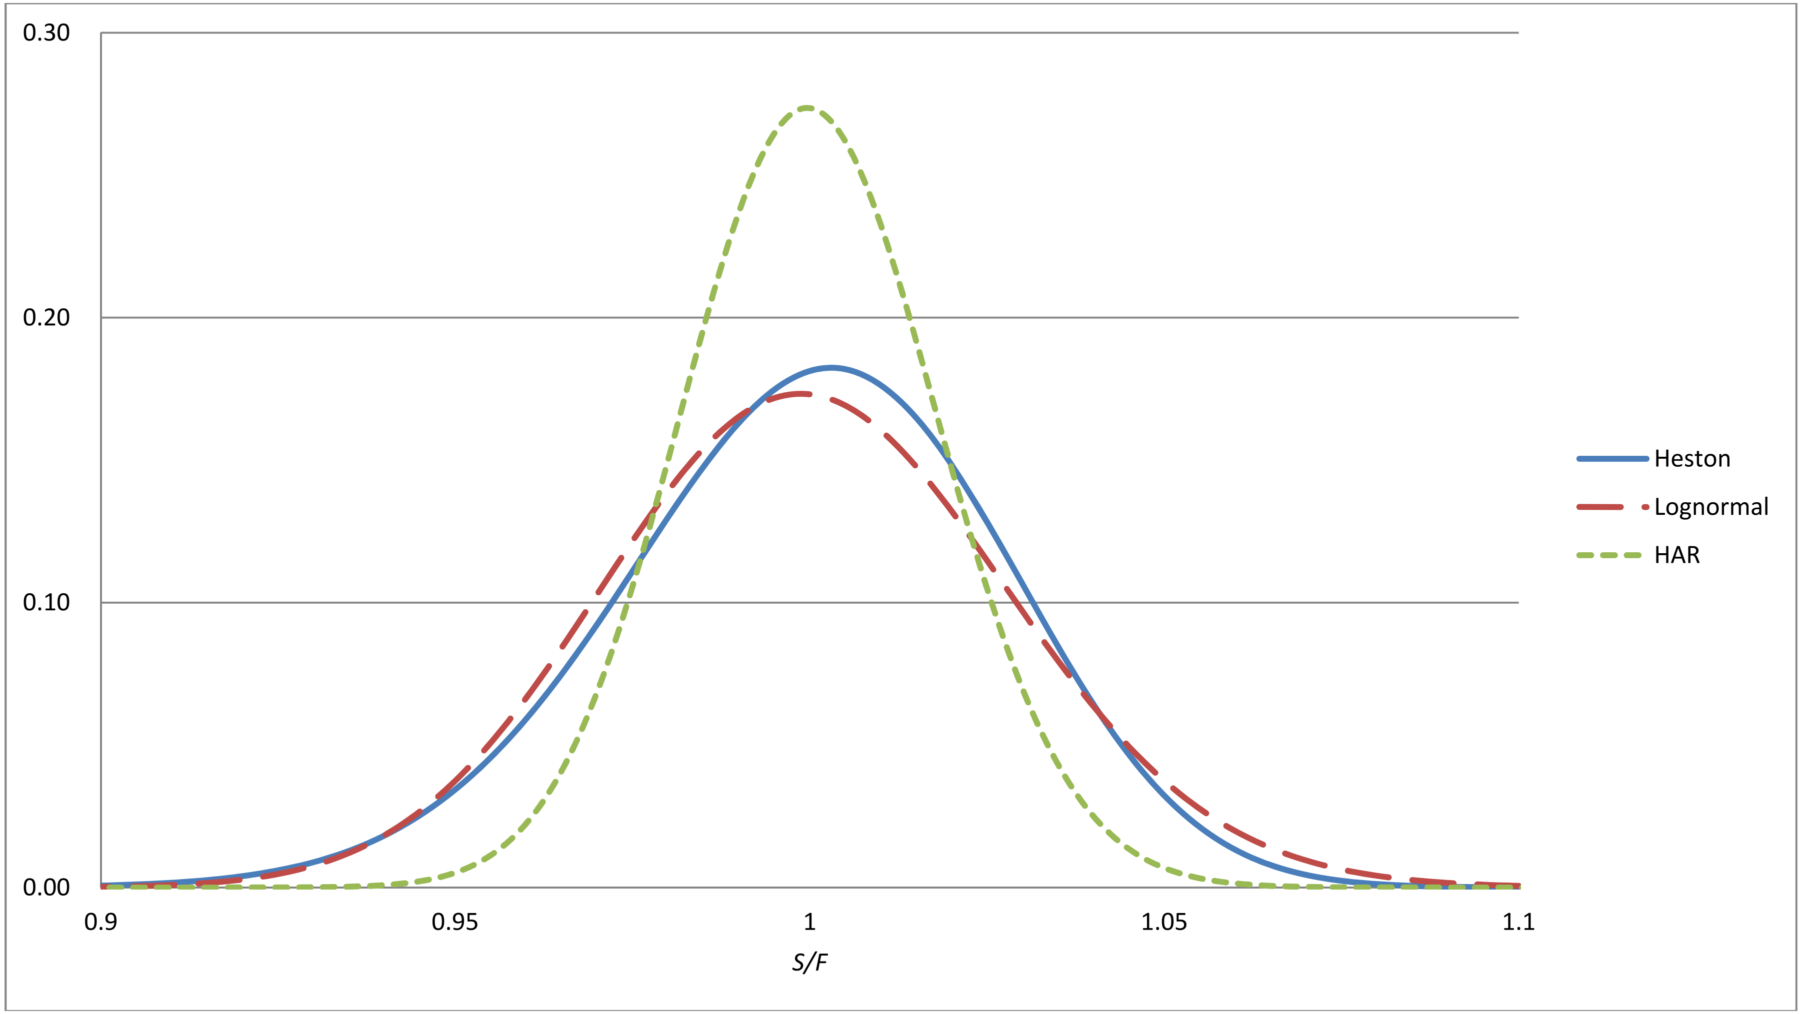

Figure 5.6 Heston, lognormal and HAR one-month ahead risk-neutral density forecasts for IBM on January 2nd 2003.

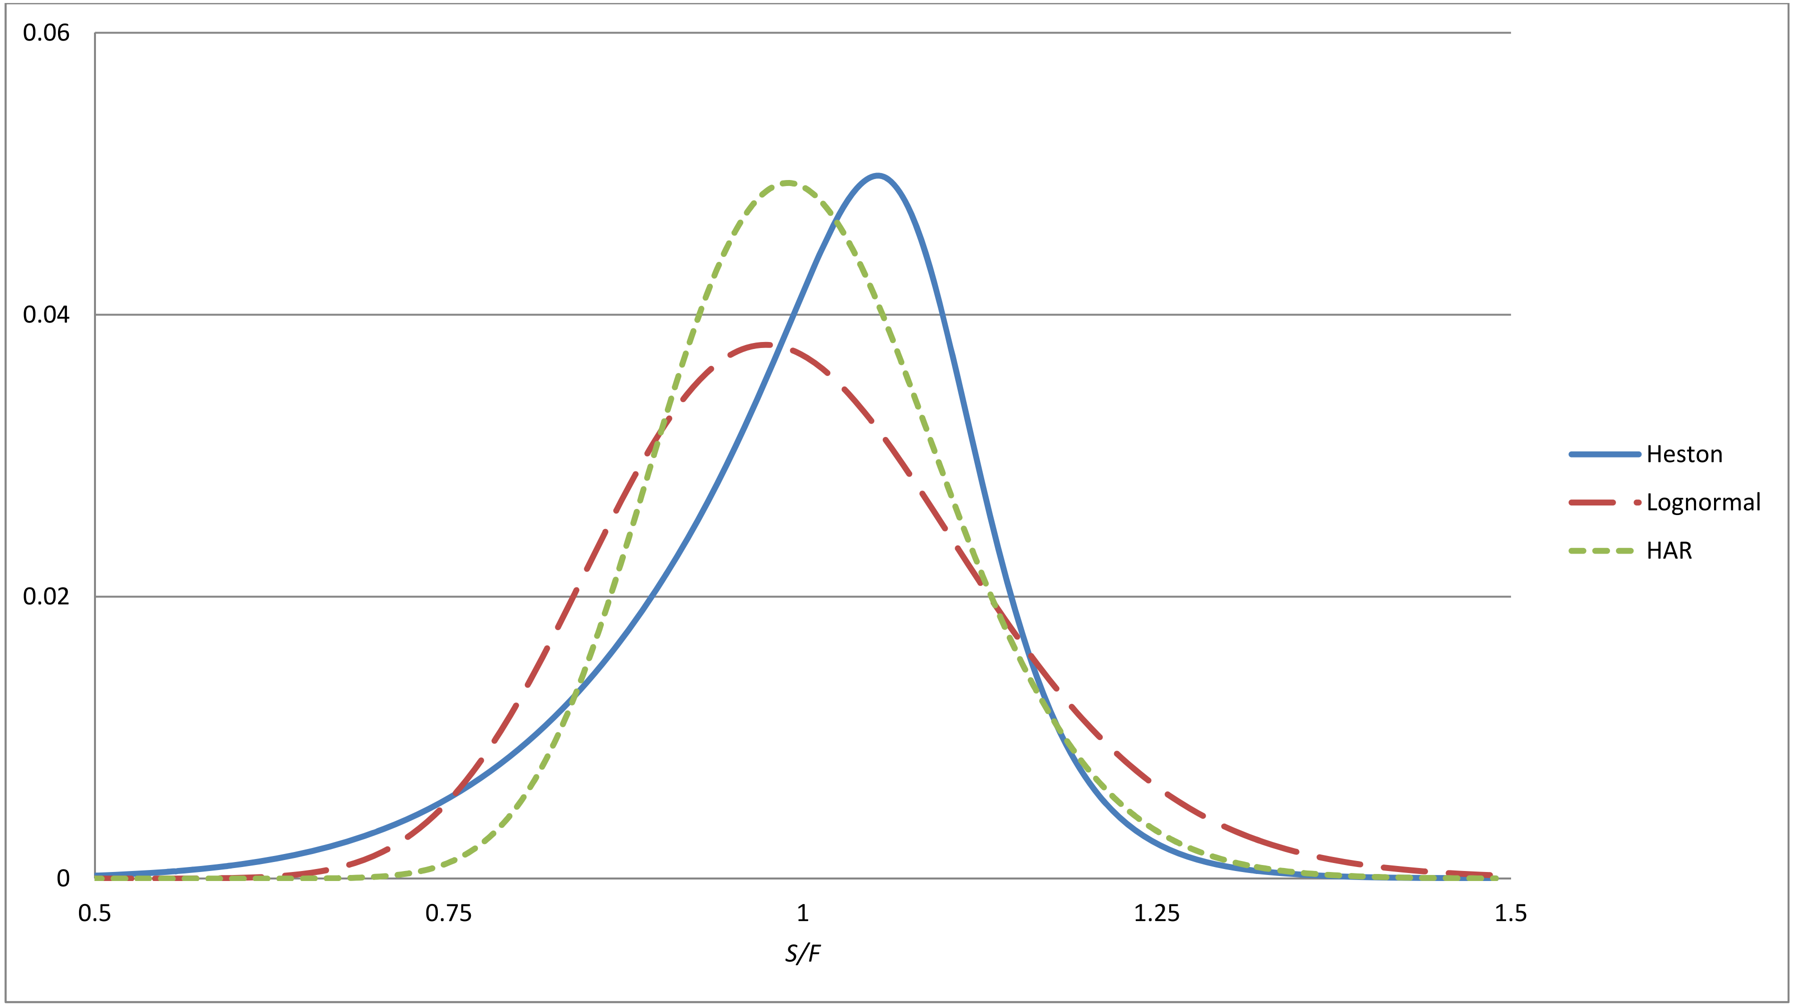

Figure 5.7 Function $\tilde{C}(u) - u$ for one-day ahead forecasts from the Heston model and a nonparametric transformation for IBM.

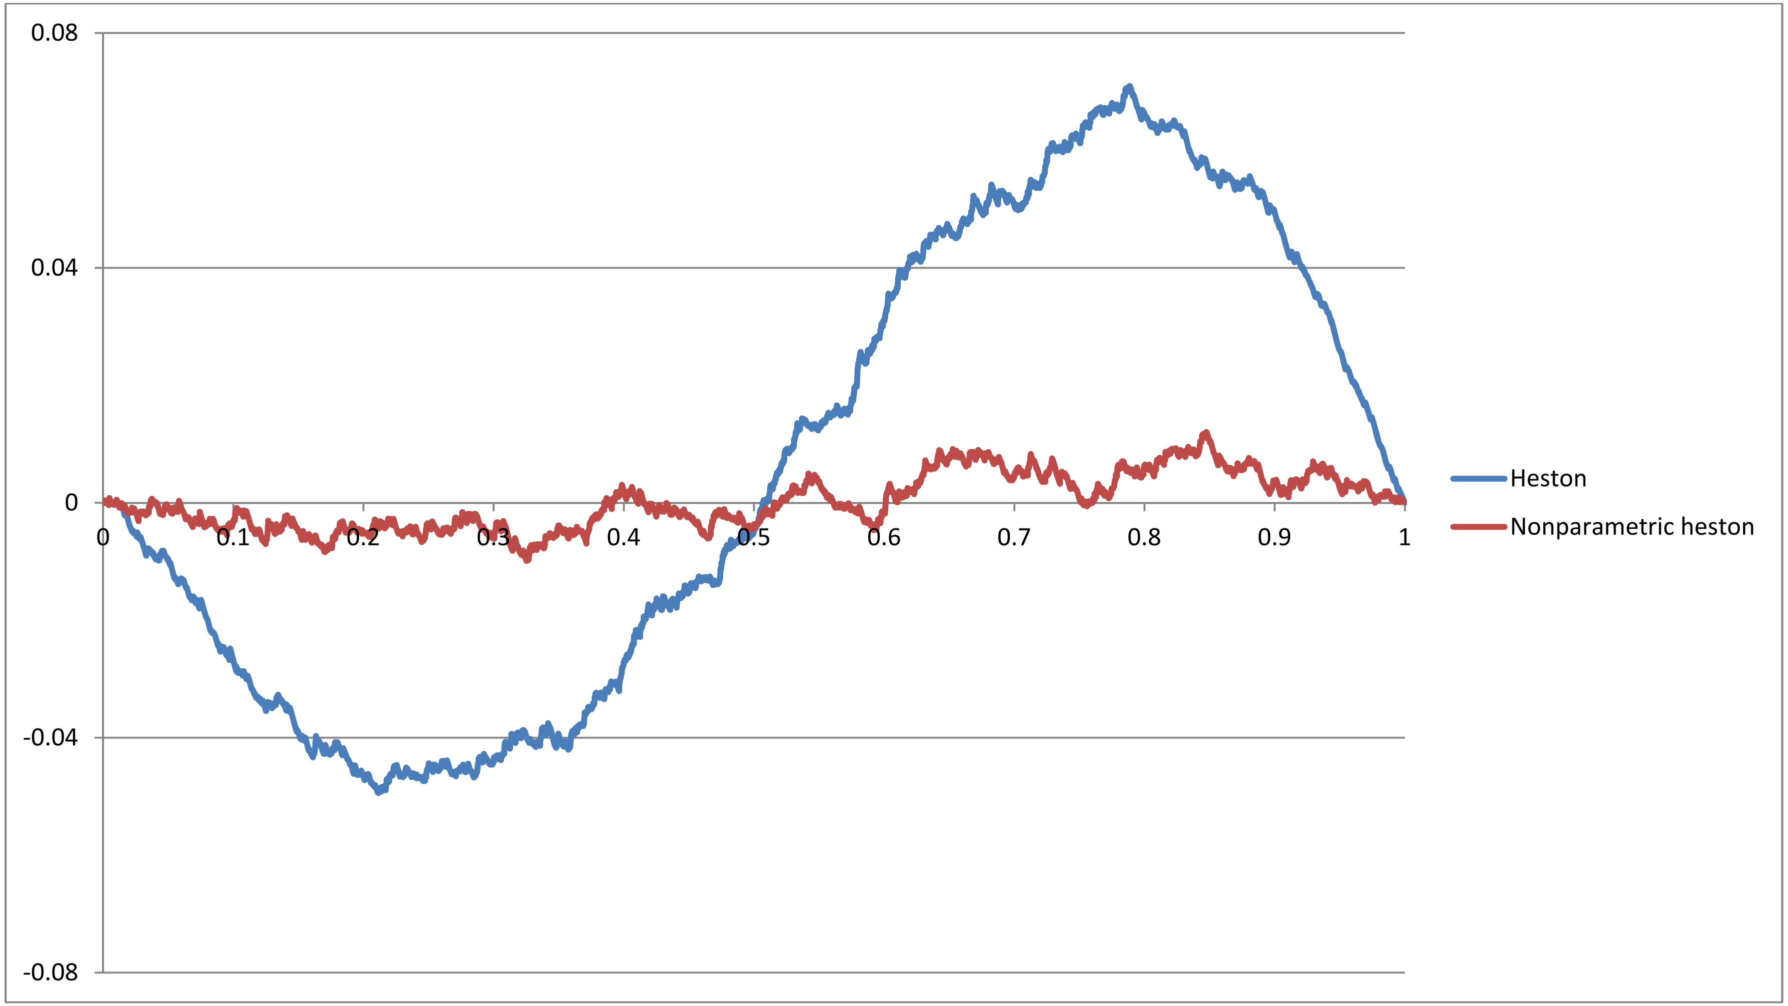

Figure 5.8 Function $\tilde{C}(u) - u$ for one-day ahead forecasts from the Black-Scholes model and a nonparametric transformation for IBM.

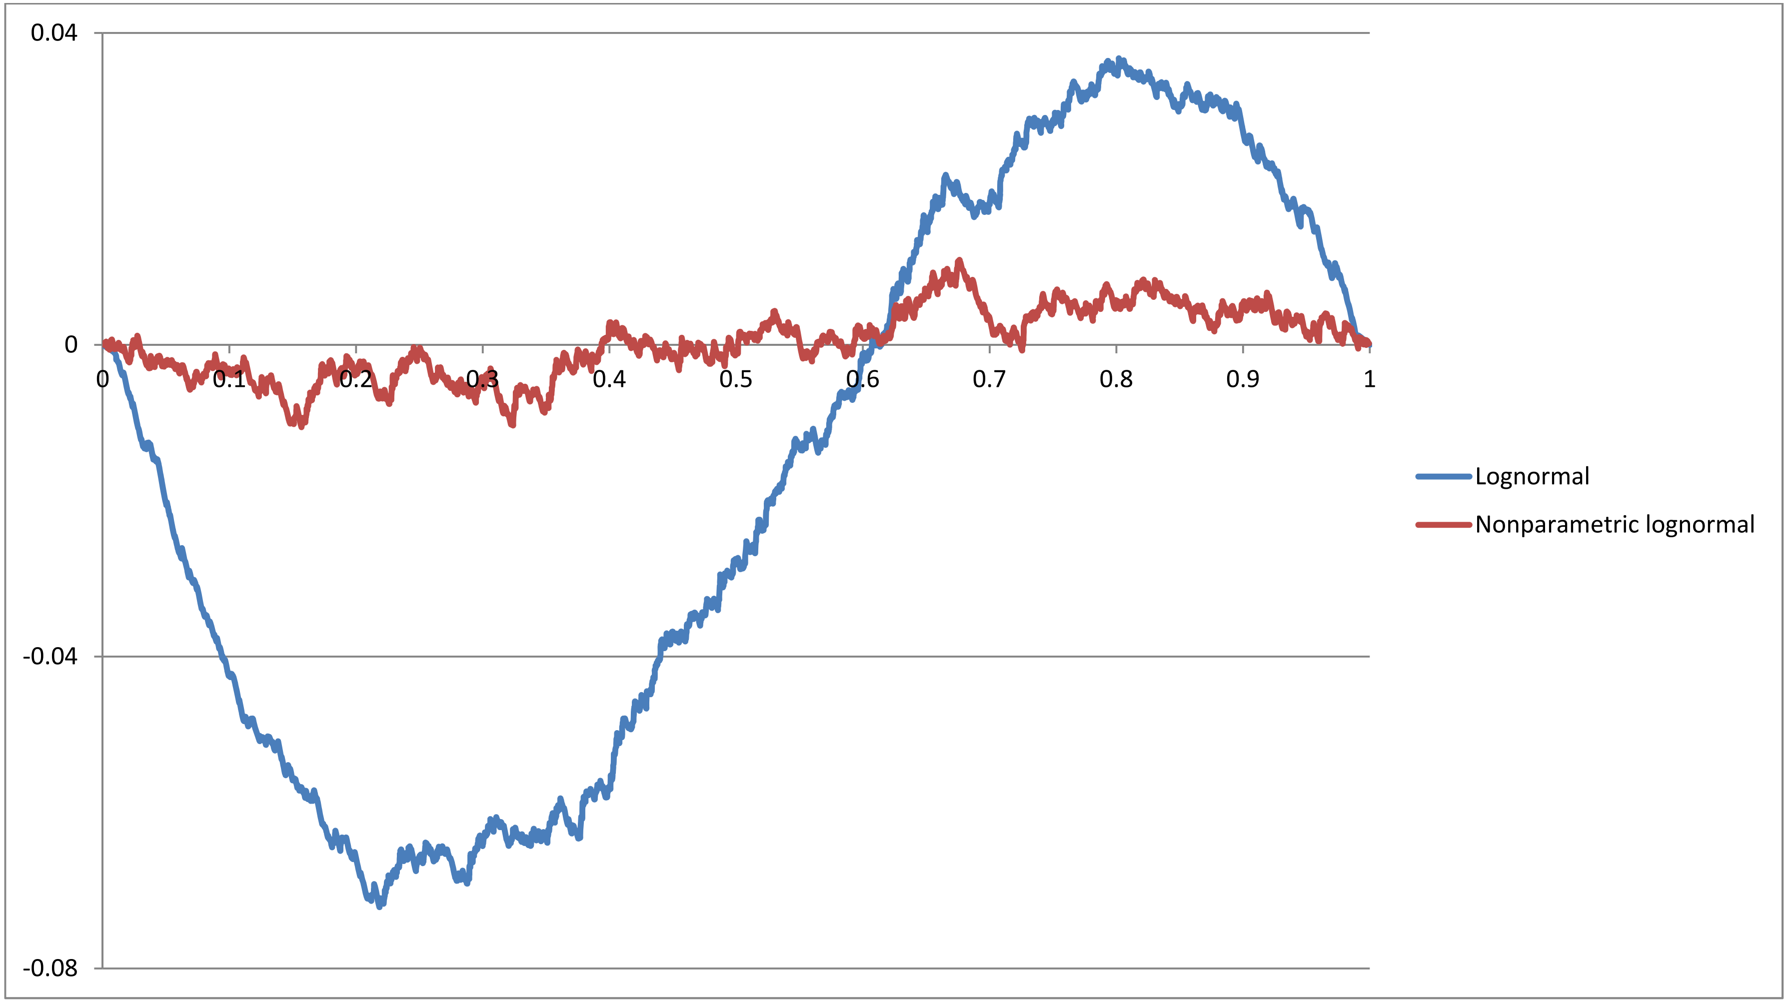

Figure 5.9 Function $\tilde{C}(u) - u$ for one-day ahead forecasts from the HAR model and a nonparametric transformation for IBM.

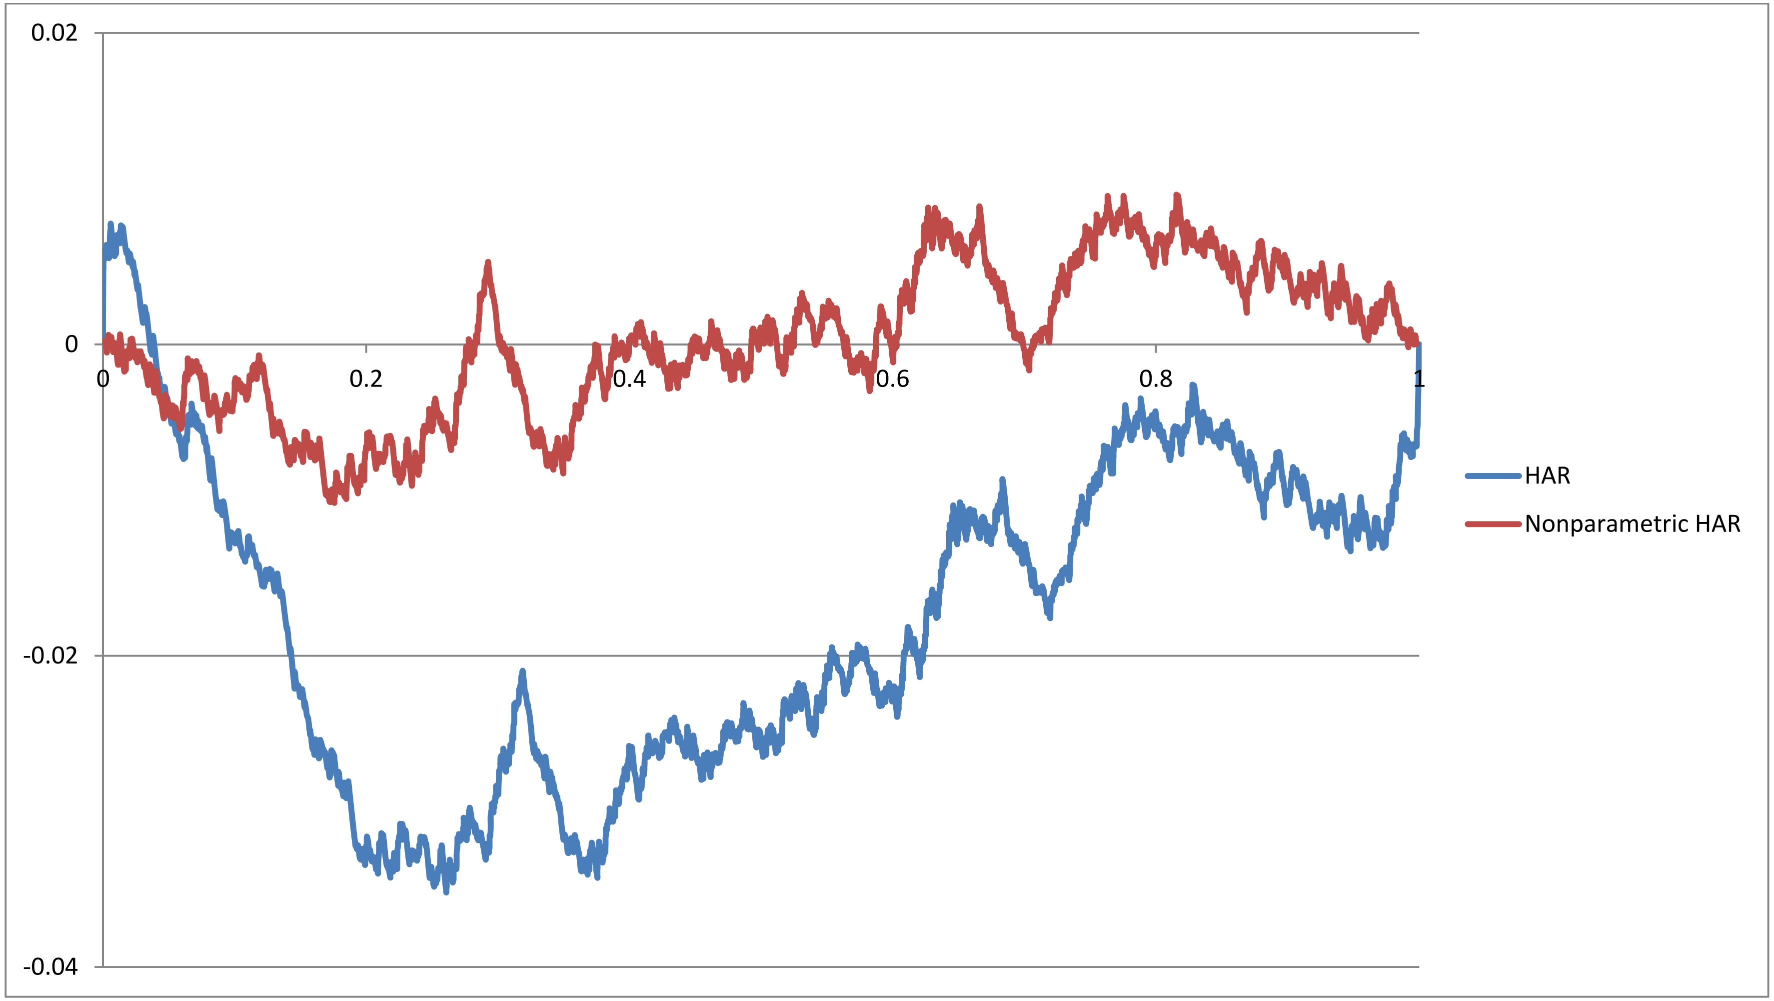

Figure 5.10 Nonparametric calibration densities $\hat{c}(u)$ from one-day ahead HAR, Lognormal and Heston forecasts for IBM.

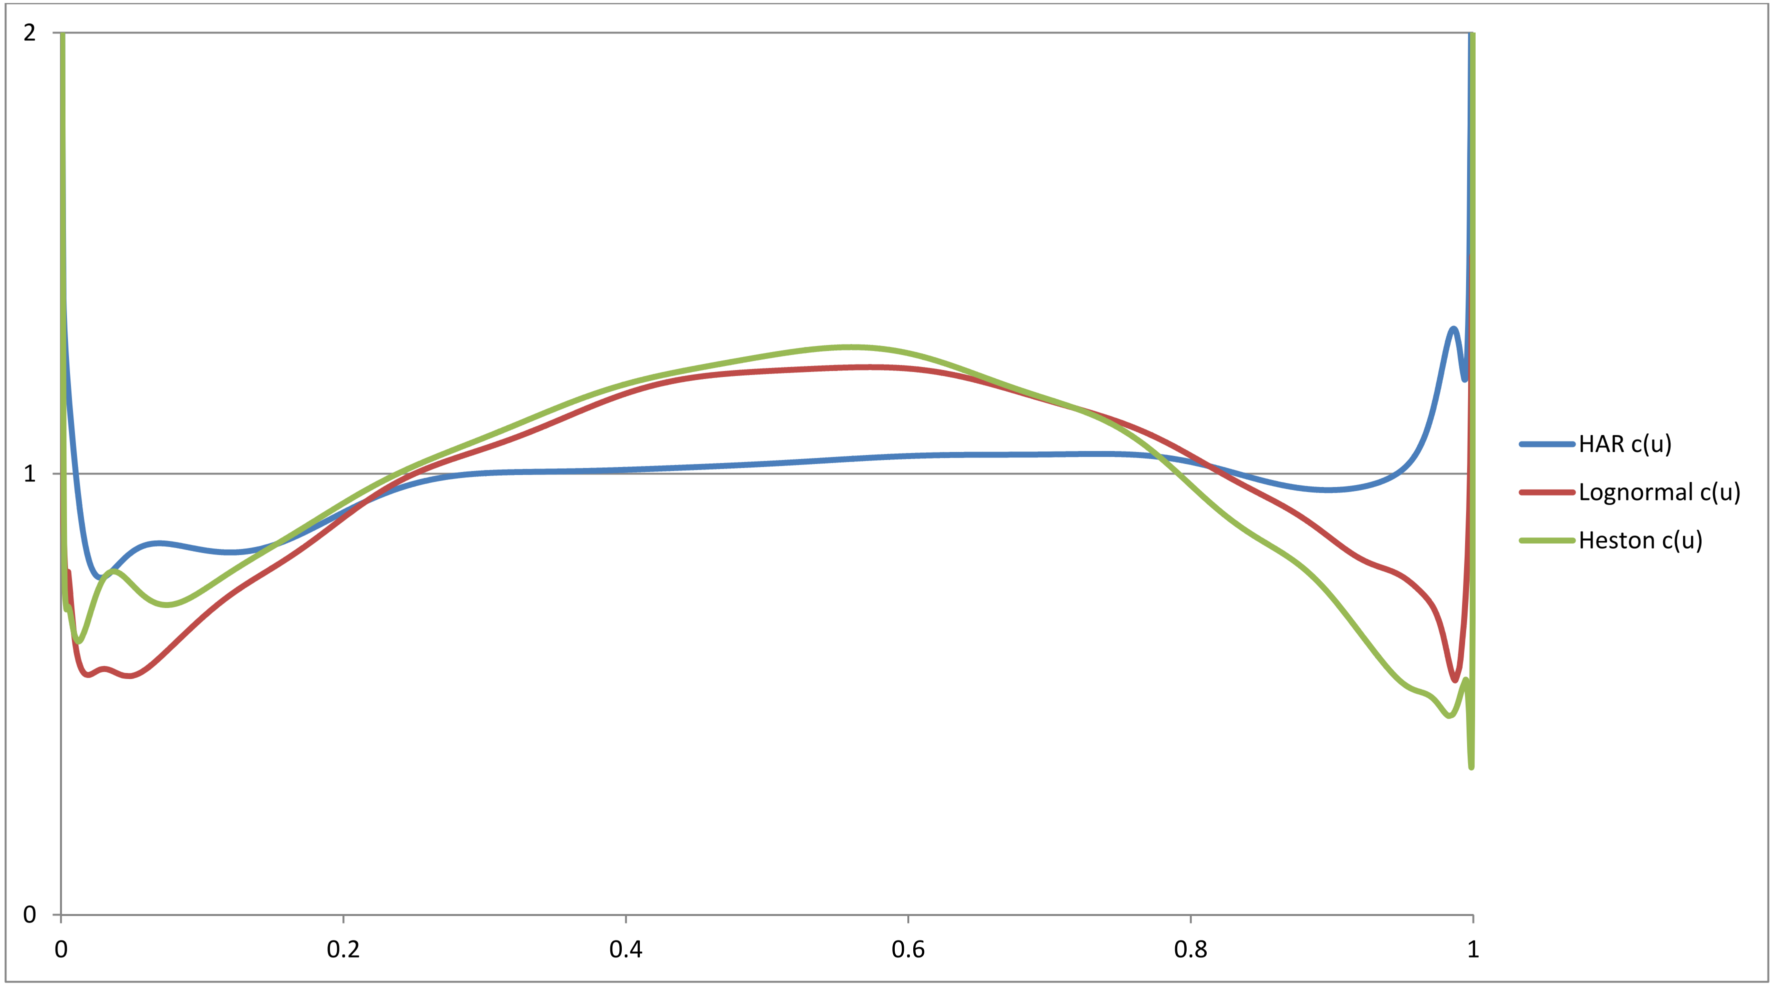

the tails because they are based on samples which have few observations at the extremes. The differences $\tilde{C}(u) - u$ after applying the nonparametric calibration method for one-day ahead forecasts from Heston, Black-Scholes and HAR lognormal densities are shown in Figures 5.7 to 5.9. The differences are much nearer zero compared to the risk-neutral densities. Comparable figures and results are obtained for longer horizon density forecasts.

5.4.4 Log-likelihood comparison

Table 5.5 gives the log-likelihoods for IBM, another sixteen stocks and the average across the seventeen stocks from 2003 to 2012, for six forecasting methods. The density forecasts are overlapping for four horizons, namely one day, one week (5 trading days), two weeks (10) and one month (22).[9] Overlapping forecasts are evaluated for horizons exceeding one day. The log-likelihood of the untransformed HAR model is defined as the benchmark level, the log-likelihoods of the other five density forecasting methods exceeding the benchmark are summarised in the table. For IBM stock, the lognormal Black-Scholes model gives the highest log-likelihoods for all four horizons ranging from one day to one month, for both risk-neutral and transformed real-world densities. The HAR model and the Heston model give similar likelihoods for all four horizons after applying transformations. The log-likelihoods for nonparametric transformation are always higher than those under risk-neutral measure for all methods and horizons, and the differences range from 66.3 to 192.8.

Similarly, for the average across seventeen stocks, the lognormal Black-Scholes

[9] For a horizon *h* trading days, we set $T=h/252$ to calculate option implied densities.

Table 5.5

Log-likelihoods for overlapping forecasts. The numbers shown are the log-likelihoods of the HAR untransformed density forecasts (0 for the average across 17 stocks) and the log-likelihoods of the other forecasts in excess of the HAR benchmark values. The letter Q defines untransformed and risk-neutral densities, while the letter P denotes nonparametric transformation of the Q densities defined by (5.33). The numbers in bold in each row refer to the best method, which has the highest log-likelihood for the selected forecast horizon.

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	P	Q	P	Q	P
IBM							
1 day	2487	-4312.5	124.1	33.0	**128.5**	-9.3	113.2
1 week	2483	-6419.1	157.3	100.1	**217.4**	100.9	167.2
2 weeks	2478	-7222.1	189.3	78.1	**270.9**	76.1	176.2
1 month	2466	-8232.5	179.9	77.2	**257.6**	65.7	151.1
Alcoa							
1 day	2487	-1616.5	81.3	68.2	**117.1**	6.9	77.7
1 week	2483	-3687.9	60.9	77.8	**107.6**	7.0	82.9
2 weeks	2478	-4575.0	131.0	108.8	**161.8**	-5.1	111.1
1 month	2466	-5693.6	305.0	202.5	**332.0**	-17.7	239.7
Boeing							
1 day	2487	-3706.5	168.5	178.7	**208.5**	119.5	174.6
1 week	2483	-5644.3	110.8	115.3	**150.9**	44.4	134.8
2 weeks	2478	-6439.4	100.4	72.3	**119.8**	-57.5	109.2
1 month	2466	-7387.4	**158.9**	57.3	147.5	-186.9	113.4
Cisco							
1 day	2487	-1235.1	260.2	185.3	**269.8**	129.3	242.8
1 week	2483	-3218.6	161.0	223.4	**266.9**	92.9	226.9
2 weeks	2478	-3966.3	109.0	130.3	**189.3**	-29.9	115.3
1 month	2466	-4904.4	81.3	68.1	**133.3**	-181.3	50.0

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Disney							
1 day	2487	-1787.1	163.6	166.3	**231.4**	98.9	200.7
1 week	2483	-3569.7	69.3	85.0	**131.8**	-38.9	93.0
2 weeks	2478	-4368.9	105.3	76.2	**191.5**	-117.6	118.9
1 month	2466	-5342.1	169.2	43.7	**237.0**	-255.5	204.5
General Electric							
1 day	2487	-2636.3	**185.2**	-150.7	-8.6	-702.0	-97.3
1 week	2483	-3330.3	208.8	347.9	**385.7**	135.2	267.3
2 weeks	2478	-3910.0	62.0	75.6	**109.2**	-166.0	-2.3
1 month	2466	-5220.7	160.3	396.3	**453.1**	36.2	335.6
Home Depot							
1 day	2487	-2009.2	78.3	40	**98.5**	-222.6	-59.9
1 week	2483	-4014.8	54.2	77.6	**110.8**	-261.3	-151.3
2 weeks	2478	-4815.7	72.8	46.9	**117.4**	-238.3	-162.7
1 month	2466	-5821.9	92.6	26.2	**136.7**	-321.8	-221.1
Hewlett Packard							
1 day	2487	-2395.3	356.2	238.3	**401.3**	257.6	386.3
1 week	2483	-4299.4	255.6	193.4	**316.8**	248.5	311.5
2 weeks	2478	-5035.9	200	127.8	**245.1**	180	232.9
1 month	2466	-6095.3	280.6	136.7	**332.6**	244.1	302.2
Intel							
1 day	2487	-2395.3	**85.5**	7.6	77.2	-1.9	71
1 week	2483	-4299.4	75.4	65.9	**104.8**	52.1	91.6
2 weeks	2478	-5035.9	**83.6**	26.7	80.3	16	71.8
1 month	2466	-6101	**86.8**	-34.2	51.3	5.5	38.7

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Johnson & Johnson							
1 day	2487	-2395.3	**171.9**	62.6	163.7	-58.3	106.1
1 week	2483	-4299.4	146.5	59.3	**174.8**	-113.7	46.2
2 weeks	2478	-5035.9	105.4	-4.5	**145.9**	-208.4	-31.4
1 month	2466	-6101.0	104.6	-37.5	**112.7**	-286.4	-137.6
JP Morgan Chase							
1 day	2487	-2395.3	**101.7**	15.1	95.6	5.5	87.6
1 week	2483	-4299.4	62.5	28.4	**93.7**	-2.3	61.3
2 weeks	2478	-5035.9	53.1	9.7	**106.6**	-50.4	39.8
1 month	2466	-6101.0	43.5	-33.3	**79.5**	-112.9	-4.6
McDonald's							
1 day	2487	-2395.3	135.9	92.7	**167.1**	57.2	148.5
1 week	2483	-4299.4	207	429.6	**516.1**	384.1	453.7
2 weeks	2478	-5035.9	78.7	-31.7	**85.8**	-72.6	19.7
1 month	2466	-6101	**153.8**	-42.3	121.8	-65.4	58.9
Merck							
1 day	2487	-2395.3	790.4	235.6	**850.8**	570.2	751.0
1 week	2483	-4299.4	553.8	137.1	**609.6**	353.0	492.6
2 weeks	2478	-5035.9	582.3	102.2	**648.8**	459.7	586.7
1 month	2466	-6101.0	431.9	-20.5	**464.6**	266.2	404.9
Pfizer							
1 day	2487	-2395.3	197.6	91.5	**222.3**	1.5	180.2
1 week	2483	-4299.4	82.5	57.9	**113.1**	3.5	71.0
2 weeks	2478	-5035.9	64.9	30.3	**96.2**	-33.7	48.3
1 month	2466	-6101.0	49.1	12.9	**101.3**	-123.7	13.7

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
AT&T							
1 day	2487	-2395.3	85.1	74.0	**121.9**	-597.4	-284.7
1 week	2483	-4299.4	57.0	59.6	**105.7**	-754.4	-494.7
2 weeks	2478	-5035.9	79.0	86.5	**139.0**	-703.0	-435.1
1 month	2466	-6101.0	116.6	109.4	**171.9**	-786.2	-491.3
Walmart							
1 day	2487	-2395.3	183.2	127.4	**213.7**	83.6	183.5
1 week	2483	-4299.4	69.7	56.2	**102.0**	-25.0	50.7
2 weeks	2478	-5035.9	40.5	5.4	**71.1**	-98.0	-7.9
1 month	2466	-6101.0	29.3	-36.8	**48.9**	-140.6	-31.5
American Express							
1 day	2487	-3046.9	271.7	154.4	267.2	53.1	**305**
1 week	2483	-4829.3	117.9	89	**153.4**	-75.1	12.1
2 weeks	2478	-5608.6	94.9	52.3	**144.5**	-40.5	10.5
1 month	2466	-6456.2	92.7	33.7	**145.5**	-91.9	9.9
Average							
1 day	2487	0	202.0	95.3	**213.3**	-12.2	152.1
1 week	2483	0	144.1	129.6	**215.4**	8.9	112.8
2 weeks	2478	0	126.6	58.4	**172.0**	-64.1	58.9
1 month	2466	0	149.0	56.5	**195.8**	-114.5	61.4

Table 5.6

Best methods. Each count is the frequency that the method has the highest log-likelihood for the selected forecast horizon across 17 stocks. Separate counts are shown for risk-neutral (Q) and transformed (P) densities. The log-likelihood always increases after transforming from Q to P, for all stocks, horizons and methods.

Forecast horizon	No. of obs.	Q			P		
		HAR	Lognormal	Heston	HAR	Lognormal	Heston
1 day	2487	1	14	2	4	12	1
1 week	2483	0	14	3	0	17	0
2 weeks	2478	2	13	2	1	16	0
1 month	2466	4	10	3	3	14	0
Total		7	51	10	8	59	1

Figure 5.11 Nonparametric HAR and Lognormal Black-Scholes log-likelihoods for 17 stocks, relative to untransformed HAR model.

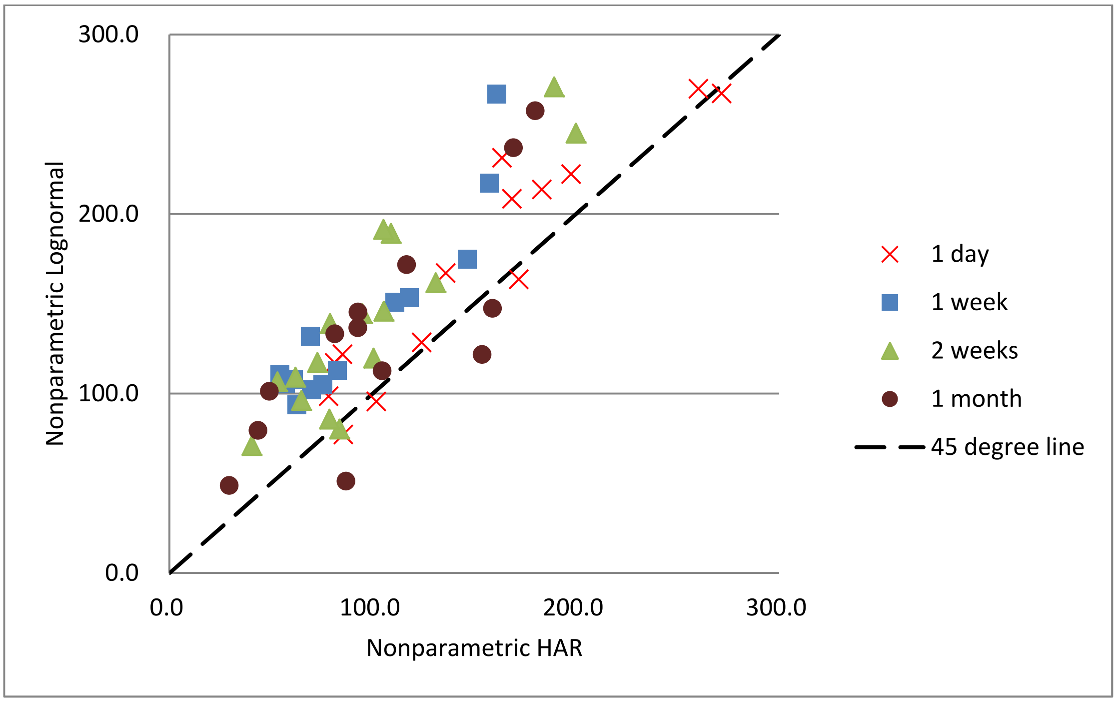

Figure 5.12 Nonparametric Lognormal Black-Scholes and Heston log-likelihoods for 17 stocks, relative to untransformed HAR model.

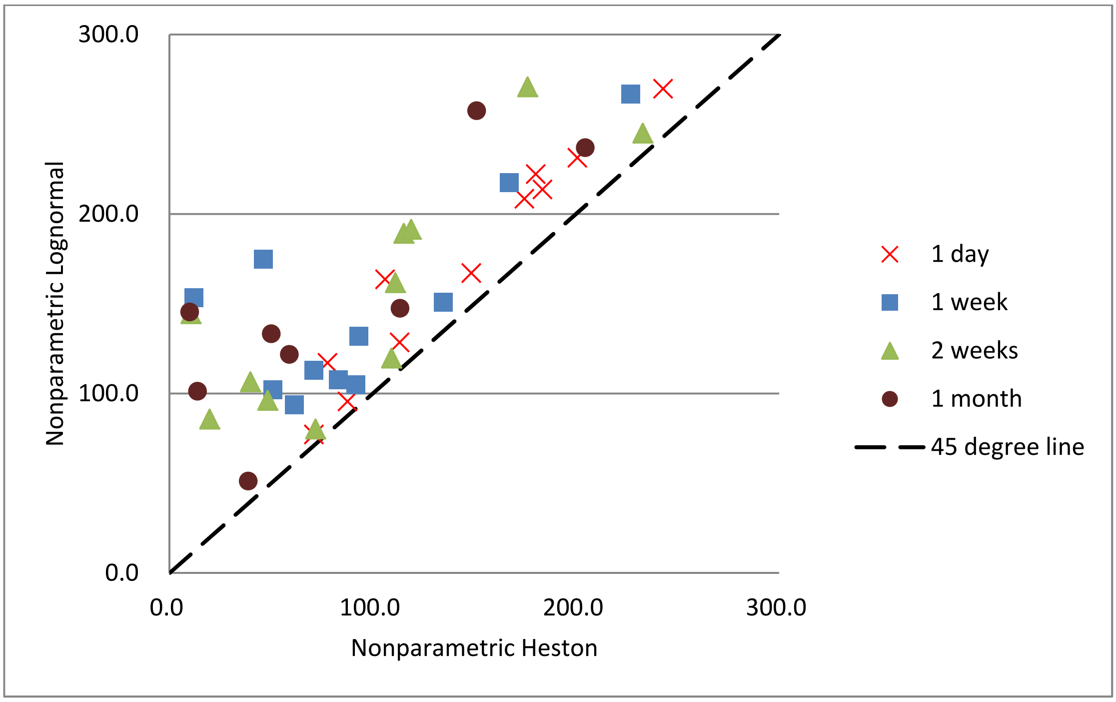

model gives the highest log-likelihoods for all four horizons, and for both untransformed risk-neutral and transformed real-world densities. The HAR model produces higher log-likelihoods than Heston model for almost all horizons both before and after applying transformations, with the exception of the risk-neutral density for one week horizon. The log-likelihoods from nonparametric transformation are always higher than those under risk-neutral measure for all methods and horizons, and the average differences vary between 85.8 and 202.0.

Table 5.6 gives the number of times that the respective method has the highest log-likelihoods for the selected forecast horizon across seventeen stocks. For transformed real-world densities, the lognormal Black-Scholes model gives the highest log-likelihoods for fifty-nine out of sixty-eight combinations from seventeen stocks and four horizons. Figures 5.11 and 5.12 show graphically that the nonparametric lognormal Black-Scholes model gives higher log-likelihoods than the nonparametric HAR and the nonparametric Heston models. (some points are outside the plotted range) The lognormal Black-Scholes model also gives the highest log-likelihoods fifty-one times for untransformed risk-neutral densities. The HAR model and the Heston model give the highest log-likelihoods for a similar number of times for risk-neutral densities, while the HAR model gets the highest log-likelihoods more times than the Heston model for transformed real-world densities.

5.4.5 Diagnostic tests

The KS statistic tests if the densities are correctly specified under the i.i.d. assumption. Table 5.7 summarises the *p*-values for the KS test for six density forecasting methods for four horizons for IBM and another sixteen stocks. Since the null hypothesis is rejected at the α significance level when $p<\alpha$, for IBM stock all the risk-neutral measure p-values reject the null hypothesis at the 5% significance level, which might be due to the mis-specified risk-neutral densities which have higher variance than real-world densities. The untransformed HAR densities are also mis-specified, as they are conditionally normal. All nonparametric transformations have satisfactory p-values greater than 50%.

Table 5.8 gives the number of times that the null hypothesis is rejected at the 5% significance level for the KS test across seventeen stocks. All the nonparametric transformations pass the KS test while the null hypothesis is rejected for almost all risk neutral and untransformed cases at the 5% significance level.

The Berkowitz LR3 statistic tests the null hypothesis that the variables $y_i=\Phi^{-1}(u_i)$ are i.i.d. and follow a standard normal distribution, against the alternative hypothesis of a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. Table 5.9 presents the LR3 test statistic, and the estimates of the variance and AR parameters for six density forecasting methods and four horizons for IBM and another sixteen stocks.

For IBM stock, the MLEs of the autoregressive parameters are between -0.01 and 0.01 for the one-day horizon, hence there is no significant evidence of time-series dependence. However, the MLEs for the one-week horizon range between -0.04 and -0.08, thus four of them reject the null hypothesis that the autoregressive parameter is 0 at the 5% significance level. The longer two-weeks and one-month horizons also

Table 5.7

KS test results for overlapping forecasts. The numbers are the percentage *p*-values of the KS test for the null hypothesis that the terms u_t are uniformly distributed. The letter *Q* defines untransformed and risk-neutral densities, while the letter *P* denotes nonparametric transformation of the real-world densities defined by (5.33). * indicates that the *p*-values are greater than 50%. The null hypothesis is rejected at the α significance level when $p<\alpha$.

Forecast horizon	No. of obs.	HAR (%)		Lognormal (%)		Heston (%)	
		Q	*P*	*Q*	*P*	*Q*	*P*
IBM							
1 day	2487	0.42	*	0.00	*	0.00	*
1 week	2483	0.01	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
Alcoa							
1 day	2487	49.13	*	0.73	*	0.00	*
1 week	2483	*	*	7.53	*	1.97	*
2 weeks	2478	30.38	*	0.55	*	2.24	*
1 month	2466	1.86	*	0.17	*	0.01	*
Boeing							
1 day	2487	0.01	*	0.15	*	0.02	*
1 week	2483	0.01	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
Cisco							
1 day	2487	0.02	*	0.00	*	0.00	*
1 week	2483	2.22	*	0.60	*	0.00	*
2 weeks	2478	0.01	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*

Forecast horizon	No. of obs.	HAR (%)		Lognormal (%)		Heston (%)	
		Q	*P*	*Q*	*P*	*Q*	*P*
Disney							
1 day	2487	0.54	*	0.62	*	0.00	*
1 week	2483	0.14	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	41.18
General Electric							
1 day	2487	9.38	*	1.68	*	0.00	*
1 week	2483	1.91	*	0.12	*	0.00	*
2 weeks	2478	24.41	*	0.01	*	0.00	*
1 month	2466	0.03	*	0.00	*	0.00	44.16
Home Depot							
1 day	2487	9.62	*	0.00	*	2.21	*
1 week	2483	0.01	*	0.00	*	15.83	*
2 weeks	2478	0.00	*	0.00	*	8.38	*
1 month	2466	0.00	*	0.00	*	0.00	*
Hewlett Packard							
1 day	2487	6.08	*	0.00	*	0.00	*
1 week	2483	0.93	*	0.00	*	0.01	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	48.90	0.08	*
Intel							
1 day	2487	1.53	*	0.00	*	0.01	*
1 week	2483	4.19	*	0.69	*	1.36	*
2 weeks	2478	0.03	*	0.00	*	0.01	*
1 month	2466	0.01	*	0.00	*	11.13	*

Forecast horizon	No. of obs.	HAR (%)		Lognormal (%)		Heston (%)	
		Q	*P*	*Q*	*P*	*Q*	*P*
Johnson & Johnson							
1 day	2487	0.06	*	0.00	*	0.00	*
1 week	2483	0.11	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
JP Morgan Chase							
1 day	2487	2.80	*	0.00	*	0.05	*
1 week	2483	0.00	*	0.00	*	0.51	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
McDonald's							
1 day	2487	0.00	*	0.00	*	0.00	*
1 week	2483	0.00	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
Merck							
1 day	2487	0.00	*	0.00	*	0.00	*
1 week	2483	0.01	*	0.00	*	0.47	*
2 weeks	2478	0.02	*	0.00	*	0.14	*
1 month	2466	0.03	*	0.09	*	0.00	*
Pfizer							
1 day	2487	0.62	*	0.00	*	0.00	*
1 week	2483	7.72	*	0.00	*	0.00	*
2 weeks	2478	12.95	*	0.00	*	0.00	*
1 month	2466	3.35	*	0.00	*	0.00	*

Forecast horizon	No. of obs.	HAR (%)		Lognormal (%)		Heston (%)	
		Q	*P*	*Q*	*P*	*Q*	*P*
AT&T							
1 day	2487	0.00	*	0.00	*	0.00	*
1 week	2483	0.00	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
Walmart							
1 day	2487	0.05	*	0.00	*	0.00	*
1 week	2483	0.33	*	0.00	*	0.00	*
2 weeks	2478	0.16	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*
American Express							
1 day	2487	0.48	*	0.00	*	0.07	*
1 week	2483	0.23	*	0.00	*	1.64	*
2 weeks	2478	0.00	*	0.00	*	0.49	*
1 month	2466	0.00	*	0.00	*	0.00	*

Table 5.8

KS test results for overlapping forecasts. The numbers are the times that the null hypothesis is rejected at the 5% significance level for 17 stocks.

Forecast horizon	HAR		Lognormal		Heston	
	Q	*P*	*Q*	*P*	*Q*	*P*
1 day	13	0	17	0	17	0
1 week	16	0	16	0	16	0
2 weeks	14	0	17	0	16	0
1 month	17	0	17	0	16	0

Table 5.9

Berkowitz test results for overlapping forecasts. The null hypothesis that the variables $y_i=\Phi^{-1}(u_i)$ are i.i.d. and follow a standard normal distribution is tested against the alternative hypothesis of a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. The numbers are the LR3 test statistic, and the estimates of the variance and AR parameters. * indicates that the null hypothesis is rejected at 5% significance level when LR3>7.81.

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
IBM							
1 day	AR	-0.01	-0.01	0.01	0.00	0.01	0.00
	Variance	1.17	0.97	0.79	0.97	0.78	0.97
	LR3	42.19*	1.74	74.23*	1.36	75.42*	1.47
1 week	AR	-0.04	-0.07	-0.06	-0.08	-0.05	-0.06
	Variance	1.18	0.96	0.86	0.96	0.84	0.96
	LR3	50.06*	15.07*	44.06*	19.08*	44.69*	11.08*
2 weeks	AR	0.01	0.00	0.01	0.00	0.01	0.01
	Variance	1.11	0.96	0.82	0.96	0.81	0.96
	LR3	30.61*	2.42	67.22*	2.51	56.32*	2.54
1 month	AR	0.01	-0.02	0.01	-0.02	-0.02	-0.01
	Variance	1.12	0.96	0.86	0.96	0.90	0.96
	LR3	44.77*	3.42	62.91*	4.12	23.80*	2.64

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Alcoa							
1 day	AR	0.03	0.03	0.03	0.03	0.04	0.04
	Variance	1.18	0.97	0.94	0.97	0.91	0.97
	LR3	38.04*	3.55	6.78	3.32	15.38*	5.06
1 week	AR	-0.01	-0.02	-0.02	-0.02	0.00	0.00
	Variance	1.11	0.97	1.04	0.97	1.01	0.96
	LR3	15.70*	2.69	2.58	2.37	1.48	2.21
2 weeks	AR	0.02	-0.01	0.00	-0.01	0.03	0.01
	Variance	1.09	0.96	1.01	0.96	1.01	0.95
	LR3	11.57*	2.92	2.92	2.49	3.49	3.66
1 month	AR	0.07	-0.01	0.02	-0.01	0.06	0.02
	Variance	1.17	0.95	1.07	0.96	1.12	0.94
	LR3	49.71*	3.69	8.32*	2.84	27.69*	7.11
Boeing							
1 day	AR	0.03	0.02	0.02	0.02	0.02	0.02
	Variance	1.51	0.97	0.95	0.97	0.97	0.97
	LR3	254.41*	2.79	9.95*	2.66	14.53*	2.50
1 week	AR	-0.02	-0.02	-0.02	-0.02	-0.01	-0.02
	Variance	1.36	0.97	1.00	0.97	0.94	0.96
	LR3	144.11*	2.87	18.19*	2.55	12.59*	2.88
2 weeks	AR	-0.01	-0.02	-0.01	-0.01	0.02	0.01
	Variance	1.27	0.96	0.95	0.96	0.92	0.95
	LR3	94.77*	2.87	37.40*	2.32	29.18*	3.63
1 month	AR	-0.04	-0.05	-0.06	-0.06	-0.04	-0.04
	Variance	1.21	0.94	0.93	0.95	0.95	0.93
	LR3	80.46*	11.79*	87.29*	13.48*	50.36*	12.58*

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Cisco							
1 day	AR	0.01	0.00	-0.01	-0.01	-0.01	-0.01
	Variance	1.18	0.96	0.89	0.97	0.86	0.97
	LR3	44.79*	1.81	22.55*	1.53	42.32*	1.77
1 week	AR	-0.07	-0.07	-0.07	-0.07	-0.06	-0.04
	Variance	1.09	0.96	0.88	0.96	0.91	0.96
	LR3	22.12*	15.76*	29.54*	12.91*	26.05*	7.30
2 weeks	AR	-0.03	-0.04	-0.02	-0.02	0.01	0.01
	Variance	0.98	0.96	0.85	0.97	0.89	0.96
	LR3	3.81	4.89	38.37*	2.74	26.66*	2.89
1 month	AR	-0.05	-0.04	-0.04	-0.04	-0.03	-0.01
	Variance	0.92	0.97	0.88	0.97	0.98	0.95
	LR3	15.90*	6.97	34.24*	5.31	10.08*	5.39
Disney							
1 day	AR	0.01	0.02	0.02	0.03	0.02	0.03
	Variance	1.44	0.97	0.90	0.97	0.86	0.97
	LR3	183.84*	2.25	14.36*	3.25	35.39*	3.07
1 week	AR	-0.05	-0.06	-0.07	-0.07	-0.05	-0.06
	Variance	1.16	0.96	0.83	0.96	0.84	0.96
	LR3	40.18*	10.03*	59.04*	14.59*	59.73*	11.86*
2 weeks	AR	-0.01	-0.03	-0.02	-0.03	0.00	-0.02
	Variance	1.06	0.97	0.76	0.96	0.82	0.95
	LR3	18.85*	3.65	110.89*	4.27	90.37*	4.62
1 month	AR	-0.02	-0.03	-0.03	-0.03	-0.03	-0.03
	Variance	1.00	0.96	0.78	0.95	0.89	0.92
	LR3	35.55*	5.63	130.83*	6.09	95.00*	12.85*

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
General Electric							
1 day	AR	0.06	0.06	0.05	0.05	0.07	0.08
	Variance	1.32	0.96	0.93	0.97	0.96	0.96
	LR3	118.50*	10.08*	12.54*	7.29	15.85*	16.01*
1 week	AR	-0.01	-0.01	-0.01	-0.01	0.00	0.01
	Variance	1.34	0.97	0.95	0.97	1.06	0.96
	LR3	116.36*	1.88	5.58	1.67	4.68	2.68
2 weeks	AR	-0.01	-0.01	-0.01	-0.01	0.01	0.01
	Variance	1.16	0.96	0.92	0.96	1.07	0.95
	LR3	29.72*	2.33	16.05*	2.45	9.27*	4.39
1 month	AR	0.03	0.01	-0.01	-0.01	0.04	0.04
	Variance	1.28	0.97	0.94	0.96	1.14	0.92
	LR3	80.81*	2.07	15.26*	3.59	31.18*	13.50*
Home Depot							
1 day	AR	0.03	0.03	0.02	0.02	0.02	0.02
	Variance	1.20	0.96	0.85	0.96	1.19	0.96
	LR3	60.24*	3.14	34.33*	2.67	47.95*	2.10
1 week	AR	-0.04	-0.05	-0.05	-0.06	-0.08	-0.08
	Variance	1.13	0.96	0.88	0.96	1.25	0.96
	LR3	31.36*	7.09	38.10*	9.51*	90.39*	17.01*
2 weeks	AR	0.02	0.03	0.03	0.03	-0.01	0.01
	Variance	1.05	0.97	0.83	0.97	1.17	0.96
	LR3	17.37*	4.13	64.31*	4.47	37.53*	2.09
1 month	AR	0.03	0.03	0.04	0.05	-0.01	0.01
	Variance	1.04	0.96	0.88	0.96	1.29	0.96
	LR3	30.00*	5.37	65.97*	8.69*	101.27*	2.53

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Hewlett Packard							
1 day	AR	0.00	-0.01	-0.01	-0.02	-0.01	-0.02
	Variance	1.54	0.96	0.94	0.96	0.84	0.97
	LR3	272.00*	2.20	6.01	2.36	35.41*	2.16
1 week	AR	-0.06	-0.08	-0.08	-0.08	-0.08	-0.08
	Variance	1.38	0.96	0.94	0.96	0.83	0.96
	LR3	155.17*	17.76*	18.49*	18.89*	51.72*	17.68*
2 weeks	AR	0.02	-0.01	0.00	-0.02	-0.01	-0.02
	Variance	1.25	0.96	0.88	0.96	0.78	0.97
	LR3	68.49*	2.33	20.58*	2.91	70.93*	2.44
1 month	AR	-0.02	-0.03	-0.03	-0.03	-0.04	-0.04
	Variance	1.29	0.95	0.96	0.95	0.82	0.96
	LR3	91.00*	6.62	4.73	6.79	57.88*	7.63
Intel							
1 day	AR	0.02	0.01	0.01	0.01	0.02	0.01
	Variance	1.03	0.97	0.91	0.97	0.87	0.97
	LR3	4.77	1.68	14.61*	1.58	26.62*	1.69
1 week	AR	-0.04	-0.05	-0.04	-0.04	-0.03	-0.04
	Variance	0.99	0.96	0.93	0.97	0.90	0.97
	LR3	4.19	7.81	11.58*	5.46	15.53*	4.98
2 weeks	AR	0.03	0.02	0.03	0.03	0.02	0.03
	Variance	0.92	0.97	0.89	0.97	0.86	0.97
	LR3	11.78*	3.18	21.89*	3.58	28.22*	3.70
1 month	AR	0.01	-0.01	0.01	-0.01	-0.01	0.00
	Variance	0.92	0.97	0.98	0.97	0.90	0.97
	LR3	11.30*	2.27	2.12	2.01	16.82*	2.18

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Johnson & Johnson							
1 day	AR	0.01	0.01	0.01	0.00	0.04	0.02
	Variance	1.35	0.97	0.81	0.97	0.86	0.96
	LR3	149.16*	1.57	65.97*	1.33	33.90*	2.56
1 week	AR	-0.06	-0.06	-0.04	-0.05	-0.02	-0.03
	Variance	1.18	0.96	0.78	0.97	0.89	0.96
	LR3	4.19	7.81	11.58*	5.46	15.53*	4.98
2 weeks	AR	-0.04	-0.04	-0.01	-0.02	0.01	0.00
	Variance	1.07	0.96	0.74	0.96	0.90	0.96
	LR3	19.99*	6.10	115.44*	3.42	14.16*	2.81
1 month	AR	-0.06	-0.07	-0.06	-0.07	-0.06	-0.04
	Variance	1.01	0.96	0.75	0.96	0.88	0.96
	LR3	21.39*	13.86*	110.62*	12.91*	37.97*	6.42
JP Morgan Chase							
1 day	AR	-0.04	-0.05	-0.05	-0.05	-0.04	-0.05
	Variance	1.31	0.97	0.97	0.96	0.91	0.97
	LR3	107.55*	7.08	8.21*	8.10*	14.07*	6.80
1 week	AR	-0.04	-0.05	-0.04	-0.04	-0.03	-0.04
	Variance	1.13	0.96	0.91	0.96	0.91	0.96
	LR3	27.27*	7.45	14.94*	6.57	14.10*	5.89
2 weeks	AR	-0.03	-0.03	-0.03	-0.03	-0.03	-0.02
	Variance	1.04	0.96	0.85	0.96	0.90	0.96
	LR3	10.70*	4.36	36.20*	4.52	15.77*	3.83
1 month	AR	0.00	0.00	0.01	0.01	0.00	0.00
	Variance	0.94	0.96	0.83	0.95	0.91	0.96
	LR3	21.01*	3.40	51.31*	4.55	12.68*	3.70

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
McDonald's							
1 day	AR	-0.02	-0.02	-0.02	-0.02	-0.01	-0.01
	Variance	1.27	0.97	0.87	0.97	0.84	0.97
	LR3	104.87*	2.21	42.66*	2.24	47.85*	1.78
1 week	AR	-0.07	-0.07	-0.05	-0.06	-0.05	-0.05
	Variance	0.71	0.96	0.82	0.96	0.84	0.96
	LR3	181.30*	12.98*	97.32*	9.19*	64.87*	8.86*
2 weeks	AR	-0.04	-0.05	-0.05	-0.05	-0.05	-0.04
	Variance	0.85	0.95	0.75	0.95	0.83	0.95
	LR3	113.76*	9.39*	179.73*	10.18*	88.46*	8.08*
1 month	AR	-0.02	-0.01	-0.01	0.00	0.01	0.01
	Variance	0.76	0.94	0.70	0.94	0.84	0.94
	LR3	257.23*	6.56	308.69*	5.87	125.15*	6.18
Merck							
1 day	AR	0.06	0.08	0.05	0.07	0.07	0.07
	Variance	1.82	0.95	1.41	0.95	1.05	0.96
	LR3	594.85*	17.91*	192.10*	16.73*	36.75*	14.52*
1 week	AR	0.03	0.00	0.03	0.00	0.01	0.00
	Variance	1.63	0.96	1.43	0.96	1.16	0.97
	LR3	350.23*	2.38	179.70*	2.44	31.37*	1.58
2 weeks	AR	0.02	0.02	0.02	0.02	0.03	0.02
	Variance	1.63	0.95	1.48	0.95	1.13	0.96
	LR3	347.35*	4.23	214.50*	5.01	19.98*	3.15
1 month	AR	0.06	0.02	0.05	0.01	0.01	0.00
	Variance	1.51	0.96	1.52	0.96	1.19	0.96
	LR3	243.04*	4.74	244.28*	3.97	36.47*	2.94

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Pfizer							
1 day	AR	0.03	0.02	0.02	0.02	0.02	0.02
	Variance	1.37	0.97	0.88	0.97	0.82	0.97
	LR3	151.72*	2.82	26.35*	2.33	48.88*	2.27
1 week	AR	-0.05	-0.05	-0.05	-0.05	-0.05	-0.05
	Variance	1.18	0.97	0.86	0.97	0.83	0.96
	LR3	44.17*	7.17	34.23*	7.90*	46.27*	7.52
2 weeks	AR	-0.04	-0.04	-0.05	-0.05	-0.05	-0.05
	Variance	1.08	0.97	0.79	0.97	0.78	0.96
	LR3	12.68*	6.38	68.36*	7.07	82.74*	7.29
1 month	AR	0.08	0.08	0.08	0.08	0.10	0.08
	Variance	0.99	0.96	0.76	0.96	0.79	0.95
	LR3	18.08*	18.75*	98.79*	17.58*	94.23*	20.11*
AT&T							
1 day	AR	0.12	0.11	0.12	0.11	0.11	0.10
	Variance	1.12	0.96	0.99	0.96	1.11	0.95
	LR3	100.19*	29.29*	71.13*	30.45*	64.83*	25.06*
1 week	AR	-0.02	-0.02	-0.03	-0.03	0.01	0.01
	Variance	0.91	0.97	0.95	0.97	1.10	0.96
	LR3	26.82*	2.92	19.79*	3.71	12.14*	2.91
2 weeks	AR	-0.03	-0.04	-0.05	-0.05	-0.02	0.01
	Variance	0.81	0.95	0.85	0.95	1.03	0.94
	LR3	75.83*	7.48	56.20*	11.25*	6.00	4.94
1 month	AR	0.00	0.00	0.01	0.01	0.06	0.08
	Variance	0.74	0.94	0.80	0.94	1.05	0.93
	LR3	138.63*	6.21	90.57*	6.32	17.90*	24.04*

Forecast horizon		HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Walmart							
1 day	AR	-0.01	-0.02	-0.01	-0.01	-0.01	-0.01
	variance	1.44	0.97	0.86	0.97	0.84	0.97
	LR3	199.90*	2.14	34.59*	1.80	39.83*	1.77
1 week	AR	-0.05	-0.05	-0.06	-0.06	-0.06	-0.06
	variance	1.23	0.97	0.84	0.97	0.85	0.96
	LR3	69.06*	7.79	45.08*	9.71*	36.78*	9.21*
2 weeks	AR	-0.05	-0.05	-0.07	-0.07	-0.05	-0.06
	variance	1.13	0.96	0.78	0.96	0.82	0.96
	LR3	28.78*	8.03*	87.93*	13.46*	55.36*	10.27*
1 month	AR	-0.08	-0.09	-0.09	-0.10	-0.07	-0.08
	variance	1.00	0.96	0.72	0.96	0.77	0.96
	LR3	24.63*	19.51*	147.79*	24.77*	94.43*	18.27*
American Express							
1 day	AR	-0.05	-0.06	-0.06	-0.06	-0.03	-0.05
	variance	1.52	0.96	0.95	0.96	1.19	0.96
	LR3	259.78*	9.63*	13.62*	11.54*	44.80*	9.25*
1 week	AR	-0.04	-0.05	-0.05	-0.05	-0.05	-0.05
	variance	1.26	0.96	0.88	0.96	1.00	0.96
	LR3	75.86*	8.26*	31.00*	8.84*	6.29	6.94
2 weeks	AR	-0.08	-0.09	-0.09	-0.09	-0.10	-0.10
	variance	1.15	0.95	0.95	0.81	0.88	0.95
	LR3	47.05*	22.85*	23.10*	79.78*	42.49*	25.68*
1 month	AR	-0.05	-0.03	-0.05	-0.04	-0.01	-0.02
	variance	1.01	0.96	0.74	0.96	0.87	0.96
	LR3	14.41*	5.40	135.46*	6.53	27.68*	4.51

Table 5.10

Berkowitz LR3 test results for overlapping forecasts. The numbers are the frequencies that the null hypothesis is rejected at the 5% significance level for 17 stocks.

Forecast horizon	HAR		Lognormal		Heston	
	Q	*P*	*Q*	*P*	*Q*	*P*
1 day	16	4	15	4	17	4
1 week	15	6	15	9	14	6
2 weeks	16	3	16	4	15	3
1 month	17	4	15	5	17	6

Figure 5.13 Untransformed HAR and Lognormal Black-Scholes Berkowitz LR3 statistic for 17 stocks.

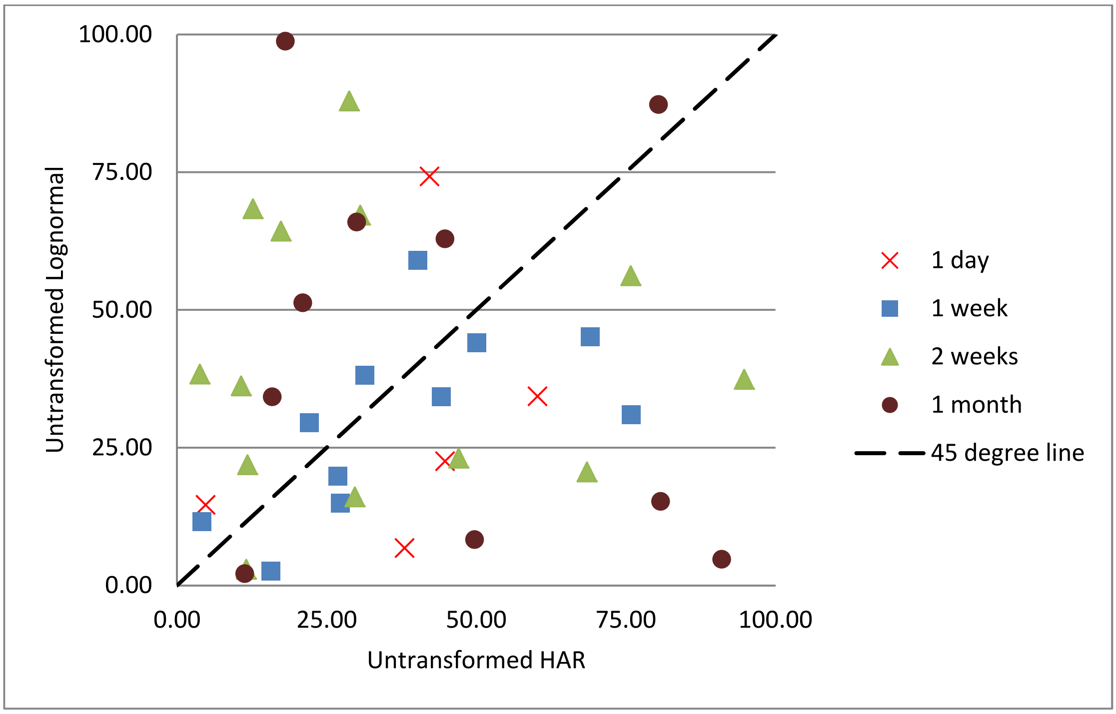

Figure 5.14 Untransformed Lognormal Black-Scholes and Heston Berkowitz LR3 statistic for 17 stocks.

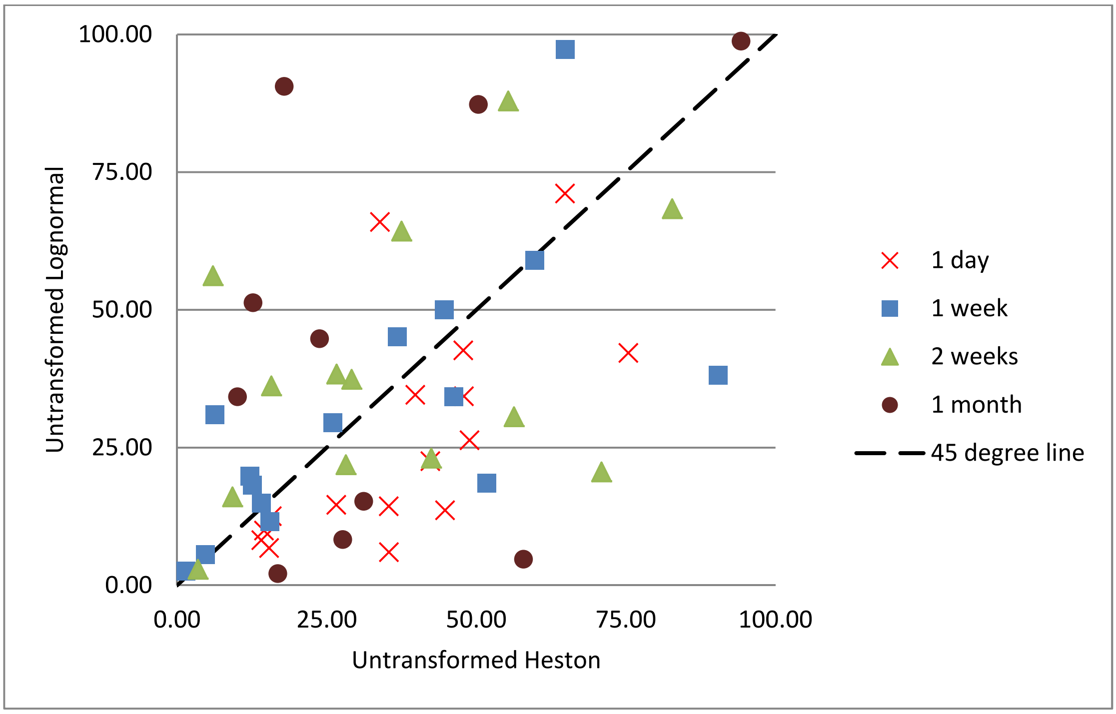

Figure 5.15 Nonparametric HAR and Lognormal Black-Scholes Berkowitz LR3 statistic for 17 stocks.

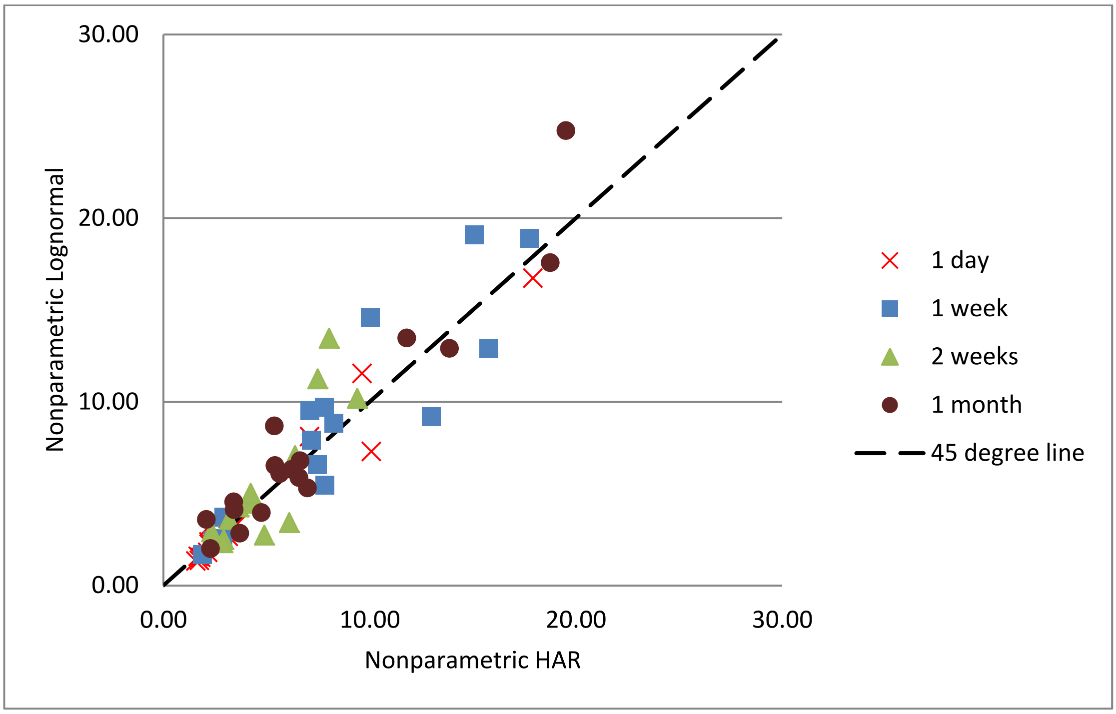

Figure 5.16 Nonparametric Lognormal Black-Scholes and Heston Berkowitz LR3 statistic for 17 stocks.

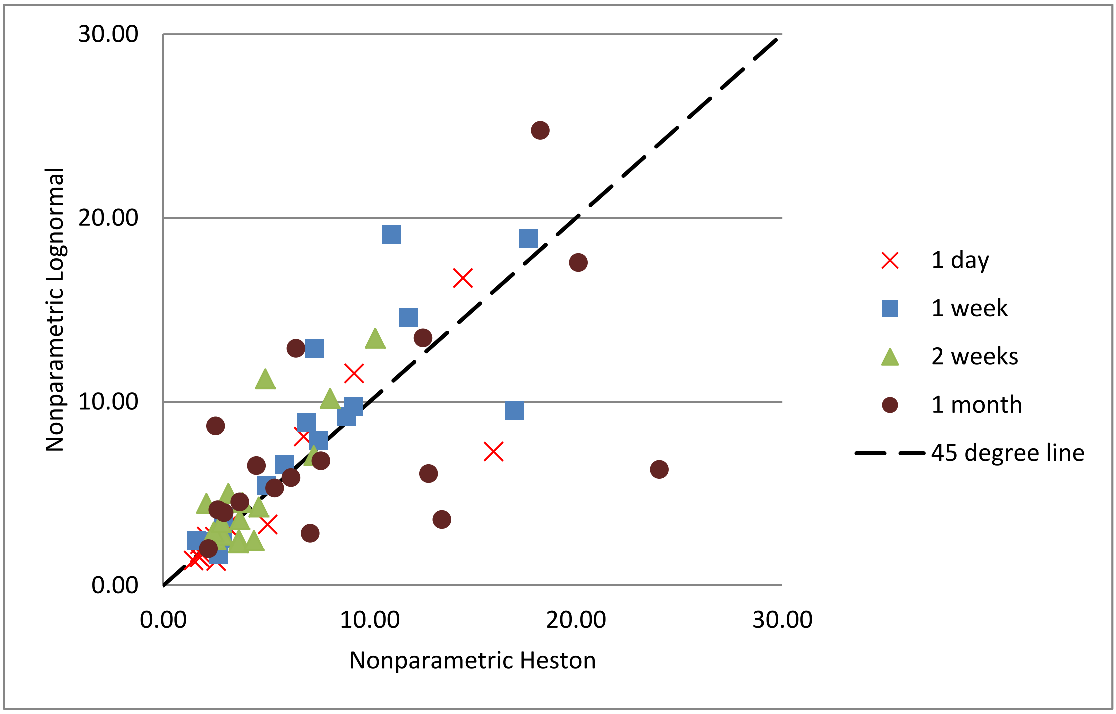

provide no evidence of dependent observations. The MLEs of the variance parameter are near one for correctly specified densities. The low estimates for one-day lognormal and Heston forecasts under Q measure can be explained by the fact that the risk-neutral standard deviations are on average are higher than the historical standard deviations.

The LR3 test statistic is significant at the 5% level when it exceeds 7.81. Table 5.9 indicates that, for IBM stock, the null hypothesis is rejected for all risk-neutral forecasts and all one-week forecasts. The null hypothesis is accepted for all real-world forecasts for one day, two-weeks and one-month horizons. The significant values of the LR3 test statistic might be attributed to the negative estimates of the AR parameter for the one-week horizon, or the mis-specified risk-neutral density which has higher variance than the real-world level.

Table 5.10 shows the number of times that the null hypothesis is rejected at the 5% significance level for all seventeen stocks for the LR3 test. Figures 5.13 and 5.14 show that the null hypothesis is rejected for almost all risk neutral measures at the 5% significance level as most Berkowitz LR3 statistics are greater than 7.81. (some points are outside the plotted range) While figures 5.15 and 5.16 show that the majority of the nonparametric transformations pass the LR3 test at the 5% significance level as most Berkowitz LR3 statistics are smaller than 7.81. (some points are outside the plotted range) The number of times that the null is rejected at the 5% level for seventeen stocks are similar across different horizons.

Table 5.11 shows the number of times that the row method provides statistically better forecasts than the column method at the 5% significance level for all seventeen stocks for the Amisano and Giacomini (AG) test. For all four horizons, the nonparametric lognormal method has the largest number of times to be statistically better than the other five density forecasting methods. For one day and one week horizons, the HAR method has the least number of times to be significantly better than the other five methods, while for the longer two weeks and one month horizons, the Heston method gets the least number. The number of times that each method is statistically better than the remaining methods are similar across the four horizons, and the nonparametric methods have more times to be significantly better than the parametric methods.

Table 5.12 gives the number of times that the row method is statistically better than the column method at the 5% level for all seventeen stocks for the AG test when the Newey-West adjustment is made to the estimated variance of $\bar{d}$ and 20 autocorrelations are used. The results are similar for the one day horizon, but the number of times that each method is statistically better than the remaining methods decreases as the forecast horizon increases to one week, two weeks and one month. As the Newey-West adjustments are necessary, we should rely on Table 5.12 rather than on Table 5.11.

Table 5.13 summarises the test statistics for the AG test for six density forecasting methods and four horizons for IBM. At the one day horizon, two of the AG test statistics are insignificant at the 5% level when the best method, nonparametric lognormal, is compared to the five alternatives; the AG test statistics equal -0.37 and 1.27 for tests against nonparametric HAR and nonparametric Heston methods.

Table 5.11

AG test results for overlapping forecasts. The numbers are the times that the row method is statistically better than the column method at the 5% level for 17 stocks.

1 day	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	2	1
HAR-*P*	16	/	12	0	16	4
Lognormal-*Q*	10	0	/	0	12	2
Lognormal-*P*	15	7	16	/	17	9
Heston-*Q*	3	0	0	0	/	0
Heston-*P*	13	1	9	0	17	/

1 week	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	3	2
HAR-*P*	17	/	5	0	14	5
Lognormal-*Q*	15	3	/	0	12	4
Lognormal-*P*	17	17	17	/	17	14
Heston-*Q*	7	1	1	0	/	0
Heston-*P*	13	4	5	0	17	/

2 weeks	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	1	0	8	2
HAR-*P*	17	/	10	0	16	7
Lognormal-*Q*	11	0	/	0	13	3
Lognormal-*P*	17	14	17	/	17	14
Heston-*Q*	2	0	2	0	/	0
Heston-*P*	8	1	7	0	17	/

1 month	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	4	0	11	3
HAR-*P*	17	/	14	2	17	11
Lognormal-*Q*	9	1	/	0	12	4
Lognormal-*P*	17	13	17	/	17	16
Heston-*Q*	2	0	3	0	/	0
Heston-*P*	10	2	7	0	17	/

Table 5.12

AG test results for overlapping forecasts when the Newey-West adjustment is made and 20 autocorrelations are used. The numbers are the times that the row method is statistically better than the column method at the 5% level for 17 stocks.

1 day	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	2	1
HAR-*P*	15	/	11	0	15	3
Lognormal-*Q*	10	0	/	0	10	2
Lognormal-*P*	14	5	16	/	17	7
Heston-*Q*	3	0	0	0	/	0
Heston-*P*	13	1	7	0	16	/

1 week	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	1	1
HAR-*P*	3	/	1	0	3	2
Lognormal-*Q*	1	1	/	0	5	1
Lognormal-*P*	4	6	7	/	11	5
Heston-*Q*	0	0	1	0	/	0
Heston-*P*	0	2	1	0	9	/

2 weeks	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	1	1
HAR-*P*	0	/	0	0	3	2
Lognormal-*Q*	0	0	/	0	2	1
Lognormal-*P*	1	2	1	/	9	3
Heston-*Q*	0	0	0	0	/	0
Heston-*P*	0	0	0	0	6	/

1 month	HAR-*Q*	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	/	0	0	0	0	0
HAR-*P*	0	/	0	0	0	0
Lognormal-*Q*	0	1	/	0	0	0
Lognormal-*P*	0	1	0	/	0	0
Heston-*Q*	0	0	0	0	/	0
Heston-*P*	0	0	0	0	1	/

Table 5.13

AG test results for IBM overlapping forecasts. The null hypothesis states that two different density forecasting methods have equal expected log-likelihood. The numbers are the test statistics. * indicates that the null hypothesis is rejected at the 5% significance level when $|t|$>1.96.

IBM 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.28*	-0.92	-3.09*	0.26	-2.81*
HAR-*P*		4.20*	-0.37	6.07*	0.72
Lognormal-*Q*			-5.31*	4.30*	-4.27*
Lognormal-*P*				6.99*	1.27
Heston-*Q*					-6.70*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.56*	-2.97*	-4.63*	-2.22*	-3.61*
HAR-*P*		2.27*	-4.65*	3.11*	-0.69
Lognormal-*Q*			-5.58*	-0.04	-2.90*
Lognormal-*P*				7.34*	4.24*
Heston-*Q*					-6.45*

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-4.10*	-3.03*	-5.64*	-1.59	-3.68*
HAR-*P*		2.91*	-6.68*	5.83*	0.85
Lognormal-*Q*			-5.38*	0.06	-2.65*
Lognormal-*P*				10.91*	6.94*
Heston-*Q*					-8.42*

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-6.13*	-3.69*	-7.80*	-1.74	-4.33*
HAR-*P*		4.22*	-6.05*	5.43*	1.59
Lognormal-*Q*			-8.33*	0.40	-2.82*
Lognormal-*P*				9.79*	6.48*
Heston-*Q*					-7.32*

Table 5.14

AG test results for overlapping forecasts when the Newey-West adjustment is made to the estimated variance of $\bar{d}$ and 20 autocorrelations are used. The null hypothesis states that two different density forecasting methods have equal expected log-likelihood. The numbers are the test statistics. * indicates that the null hypothesis is rejected at the 5% significance level when $|t|>1.96$.

IBM 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.17*	-1.01	-3.27*	0.26	-2.96*
HAR-*P*		2.98*	-0.33	3.68*	0.56
Lognormal-*Q*			-4.05*	2.75*	-3.32*
Lognormal-*P*				4.40*	0.99
Heston-*Q*					-4.34*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.96	-0.83	-1.25	-0.61	-1.01
HAR-*P*		0.73	-1.73	1.06	-0.29
Lognormal-*Q*			-1.69	-0.01	-1.00
Lognormal-*P*				2.56*	1.83
Heston-*Q*					-2.38*

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.68	-0.54	-0.93	-0.26	-0.64
HAR-*P*		0.60	-1.58	1.61	0.24
Lognormal-*Q*			-1.10	0.01	-0.60
Lognormal-*P*				2.87*	2.10*
Heston-*Q*					-2.38*

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.53	-0.40	-0.69	-0.17	-0.41
HAR-*P*		0.46	-0.95	0.93	0.22
Lognormal-*Q*			-0.82	0.04	-0.33
Lognormal-*P*				1.73	1.08
Heston-*Q*					-1.34

Alcoa 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.01*	-3.03*	-3.89*	-0.32	-2.67*
HAR-*P*		0.88	-3.40*	2.95*	0.31
Lognormal-*Q*			-3.73*	3.20*	-0.57
Lognormal-*P*				4.01*	2.99*
Heston-*Q*					-3.41*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.82	-0.82	-1.00	-0.12	-0.94
HAR-*P*		-0.43	-1.23	0.93	-0.55
Lognormal-*Q*			-1.31	1.17	-0.14
Lognormal-*P*				1.37	0.69
Heston-*Q*					-1.69

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.62	-0.54	-0.61	0.06	-0.55
HAR-*P*		0.42	-0.61	0.97	0.41
Lognormal-*Q*			-0.86	0.89	-0.04
Lognormal-*P*				0.90	0.80
Heston-*Q*					-0.99

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.37	-0.38	-0.38	0.06	-0.33
HAR-*P*		0.41	-0.40	0.69	0.54
Lognormal-*Q*			-0.44	1.01	-0.18
Lognormal-*P*				0.67	0.60
Heston-*Q*					-0.66

Boeing 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-4.35*	-4.39*	-4.66*	-2.63*	-3.81*
HAR-*P*		-0.70	-3.36*	1.76	-0.31
Lognormal-*Q*			-3.12*	2.13*	0.19
Lognormal-*P*				3.05*	1.69
Heston-*Q*					-3.51*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.70	-1.67	-1.99*	-0.54	-1.67
HAR-*P*		-0.25	-1.78	-0.93	-0.65
Lognormal-*Q*			-1.77	1.29	-0.55
Lognormal-*P*				1.68	0.50
Heston-*Q*					-2.16*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.13	-0.70	-1.14	0.55	-0.98
HAR-*P*		0.48	-0.49	1.61	-0.14
Lognormal-*Q*			-1.09	1.59	-0.55
Lognormal-*P*				1.79	0.24
Heston-*Q*					-2.04*
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.86	-0.26	-0.65	1.08	-0.50
HAR-*P*		0.64	0.12	1.73	0.44
Lognormal-*Q*			-0.71	1.61	-0.32
Lognormal-*P*				1.57	0.38
Heston-*Q*					-1.54

Cisco 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.74*	-2.60*	-2.97*	-1.65	-2.62*
HAR-*P*		2.46*	-0.67	4.09*	0.84
Lognormal-*Q*			-3.48*	2.58*	-1.84
Lognormal-*P*				4.50*	1.29
Heston-*Q*					-4.44*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.96	-1.33	-1.45	-0.76	-1.31
HAR-*P*		-1.27	-2.66*	0.90	-2.11*
Lognormal-*Q*			-1.55	2.19*	-0.09
Lognormal-*P*				2.18*	1.48
Heston-*Q*					-1.84

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.01	-1.07	-1.40	0.33	-1.00
HAR-*P*		-0.37	-1.58	1.32	-0.11
Lognormal-*Q*			-1.87	1.78	0.24
Lognormal-*P*				1.94	1.31
Heston-*Q*					-1.57

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.94	-0.86	-1.52	1.36	-0.40
HAR-*P*		0.14	-1.04	1.41	0.30
Lognormal-*Q*			-0.86	1.55	0.13
Lognormal-*P*				1.67	0.80
Heston-*Q*					-1.47

Disney 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.18*	-3.18*	-3.34*	-1.65	-2.89*
HAR-*P*		-0.10	-1.88	1.74	-0.97
Lognormal-*Q*			-2.72*	2.61*	-1.15
Lognormal-*P*				3.88*	1.32
Heston-*Q*					-4.33*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.12	-0.96	-1.50	0.40	-1.01
HAR-*P*		-0.59	-2.36*	1.63	-0.53
Lognormal-*Q*			-2.20*	2.23*	-0.19
Lognormal-*P*				2.49*	0.98
Heston-*Q*					-2.60*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.96	-0.62	-1.39	0.95	-0.87
HAR-*P*		0.53	-2.25*	1.74	-0.21
Lognormal-*Q*			-1.90	1.98*	-0.53
Lognormal-*P*				2.19*	1.16
Heston-*Q*					-2.05*
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.76	-0.32	-0.99	1.78	-0.81
HAR-*P*		0.90	-1.58	1.56	-0.41
Lognormal-*Q*			-1.23	1.74	-0.83
Lognormal-*P*				1.70	0.34
Heston-*Q*					-1.60

General Electric					
1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.00*	-1.59	-1.91	0.68	-0.72
HAR-*P*		2.33*	0.67	3.31*	2.05*
Lognormal-*Q*			-2.56*	2.77*	1.00
Lognormal-*P*				3.29*	2.01*
Heston-*Q*					-3.94*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.40	-1.81	-1.95	-0.85	-1.46
HAR-*P*		-2.07*	-2.69*	0.71	-0.87
Lognormal-*Q*			-2.02*	2.14*	1.36
Lognormal-*P*				2.27*	2.00*
Heston-*Q*					-1.79
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.86	-0.87	-1.19	0.76	0.02
HAR-*P*		-0.39	-1.62	1.52	0.77
Lognormal-*Q*			-1.76	1.71	0.91
Lognormal-*P*				1.99*	1.31
Heston-*Q*					-0.90
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.64	-1.17	-1.28	-0.14	-1.01
HAR-*P*		-1.99*	-2.20*	0.60	-0.99
Lognormal-*Q*			-1.08	1.56	0.39
Lognormal-*P*				1.66	0.88
Heston-*Q*					-1.60

Home Depot 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.74*	-1.51	-3.08*	3.70*	1.19
HAR-*P*		1.52	-1.49	5.02*	3.22*
Lognormal-*Q*			-3.06*	4.78*	2.24*
Lognormal-*P*				5.18*	3.42*
Heston-*Q*					-3.72*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.96	-0.90	-1.38	1.70	1.31
HAR-*P*		-0.49	-2.10*	2.36*	2.46*
Lognormal-*Q*			-1.90	2.38	2.63
Lognormal-*P*				2.65*	3.02*
Heston-*Q*					-1.49

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.79	-0.51	-1.02	1.10	0.95
HAR-*P*		0.46	-1.18	1.76	2.00*
Lognormal-*Q*			-1.19	1.42	1.49
Lognormal-*P*				1.92	2.21*
Heston-*Q*					-0.88

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.94	-0.48	-1.20	1.11	0.91
HAR-*P*		0.65	-1.08	1.47	1.54
Lognormal-*Q*			-0.94	1.19	1.03
Lognormal-*P*				1.62	1.75
Heston-*Q*					-0.81

Hewlett Packard					
1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.36*	-2.67*	-3.42*	-2.44*	-3.30*
HAR-*P*		2.98*	-1.40	2.22*	-0.88
Lognormal-*Q*			-3.74*	-0.76	-3.45*
Lognormal-*P*				4.00*	1.39
Heston-*Q*					-3.95*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.43	-1.40	-1.66	-1.34	-1.65
HAR-*P*		1.02	-2.09*	0.20	-2.20*
Lognormal-*Q*			-1.98*	-0.96	-1.82
Lognormal-*P*				2.24*	0.31
Heston-*Q*					-2.25*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.33	-0.78	-1.42	-1.03	-1.45
HAR-*P*		0.92	-0.86	0.32	-0.67
Lognormal-*Q*			-1.73	-0.96	-1.57
Lognormal-*P*				1.40	0.45
Heston-*Q*					-1.32
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.73	-0.41	-0.76	-0.58	-0.71
HAR-*P*		1.09	-0.61	0.54	-0.28
Lognormal-*Q*			-1.31	-0.83	-1.14
Lognormal-*P*				1.38	0.60
Heston-*Q*					-1.42

Intel 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.36*	-0.37	-2.32*	0.07	-2.32*
HAR-*P*		2.56*	0.56	2.88*	0.91
Lognormal-*Q*			-3.05*	0.53	-2.86*
Lognormal-*P*				3.31*	0.63
Heston-*Q*					-3.83*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.03	-0.95	-1.17	-0.71	-1.25
HAR-*P*		0.31	-0.96	0.57	-0.71
Lognormal-*Q*			-1.92	0.60	-1.43
Lognormal-*P*				1.53	0.71
Heston-*Q*					-1.70
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.12	-0.30	-0.75	-0.16	-0.85
HAR-*P*		0.99	0.05	1.06	0.26
Lognormal-*Q*			-1.96	0.29	-1.50
Lognormal-*P*				1.77	0.32
Heston-*Q*					-1.78
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.89	0.50	-0.39	-0.04	-0.30
HAR-*P*		1.38	0.61	0.94	0.79
Lognormal-*Q*			-0.90	-0.33	-0.77
Lognormal-*P*				0.99	0.36
Heston-*Q*					-1.03

Johnson &Johnson					
1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.36*	-0.37	-2.32*	0.07	-2.32*
HAR-*P*		2.56*	0.56	2.88*	0.91
Lognormal-*Q*			-3.05*	0.53	-2.86*
Lognormal-*P*				3.31*	0.63
Heston-*Q*					-3.83*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.43	-0.59	-1.63	1.06	-0.40
HAR-*P*		2.36*	-1.47	3.00*	1.75
Lognormal-*Q*			-3.50*	2.19*	0.20
Lognormal-*P*				3.22*	2.12*
Heston-*Q*					-2.29*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.28	0.03	-1.54	1.39	0.21
HAR-*P*		1.47	-1.33	2.13*	1.34
Lognormal-*Q*			-2.42*	1.53	0.23
Lognormal-*P*				2.33*	1.70
Heston-*Q*					-1.95
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.51	0.29	-1.37	1.13	0.73
HAR-*P*		1.48	-0.18	1.54	1.42
Lognormal-*Q*			-1.71	1.20	0.65
Lognormal-*P*				1.74	1.61
Heston-*Q*					-1.33

JP Morgan Chase					
1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-4.20*	-0.45	-3.52*	-0.22	-3.33*
HAR-*P*		2.23*	0.40	3.71*	0.85
Lognormal-*Q*			-3.11*	0.30	-2.41*
Lognormal-*P*				3.94*	0.73
Heston-*Q*					-4.18*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.10*	-0.80	-2.31*	0.01	-1.53
HAR-*P*		0.84	-1.21	1.48	0.03
Lognormal-*Q*			-2.18*	0.65	-0.81
Lognormal-*P*				2.43*	1.15
Heston-*Q*					-1.97*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.06	-0.26	-1.81	0.53	-0.75
HAR-*P*		0.45	-1.64	1.53	0.34
Lognormal-*Q*			-1.05	0.55	-0.35
Lognormal-*P*				2.43*	1.55
Heston-*Q*					-2.27*
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.63	0.30	-0.96	0.59	0.16
HAR-*P*		0.55	-0.77	1.26	0.81
Lognormal-*Q*			-0.89	0.36	-0.04
Lognormal-*P*				1.48	1.15
Heston-*Q*					-1.35

McDonald's 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-4.02*	-3.22*	-4.51*	-1.72	-3.91*
HAR-*P*		2.21*	-2.95*	3.02*	-1.02
Lognormal-*Q*			-4.39*	2.45*	-3.09*
Lognormal-*P*				4.39*	1.75
Heston-*Q*					-4.38*
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.10*	-0.80	-2.31*	0.01	-1.53
HAR-*P*		0.84	-1.21	1.48	0.03
Lognormal-*Q*			-2.18*	0.65	-0.81
Lognormal-*P*				2.43*	1.15
Heston-*Q*					-1.97*
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.29	0.55	-1.26	0.96	-0.36
HAR-*P*		1.22	-0.21	2.17*	1.40
Lognormal-*Q*			-1.39	0.71	-0.60
Lognormal-*P*				2.69*	1.82
Heston-*Q*					-1.95
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.97	0.43	-0.80	0.31	-0.45
HAR-*P*		1.00	0.66	1.40	1.25
Lognormal-*Q*			-0.87	0.17	-0.60
Lognormal-*P*				1.47	1.21
Heston-*Q*					-1.08

Merck 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.57	-3.08*	-1.68	-1.18	-1.53
HAR-*P*		1.27	-1.90	3.78*	1.38
Lognormal-*Q*			-1.40	-0.80	-1.21
Lognormal-*P*				4.05*	3.21*
Heston-*Q*					-3.96*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.73	-2.08*	-0.80	-0.57	-0.71
HAR-*P*		0.55	-2.76*	1.24	0.90
Lognormal-*Q*			-0.62	-0.35	-0.52
Lognormal-*P*				1.52	1.58
Heston-*Q*					-1.41

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.48	-1.48	-0.53	-0.43	-0.49
HAR-*P*		0.39	-2.03*	0.64	-0.09
Lognormal-*Q*			-0.44	-0.33	-0.40
Lognormal-*P*				1.02	1.60
Heston-*Q*					-0.79

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.35	0.17	-0.37	-0.26	-0.34
HAR-*P*		0.32	-0.71	0.61	0.27
Lognormal-*Q*			-0.34	-0.24	-0.31
Lognormal-*P*				0.78	0.74
Heston-*Q*					-0.74

Pfizer 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.87*	-2.41*	-3.13*	-1.13	-2.70*
HAR-*P*		2.03*	-1.86	3.63*	0.92
Lognormal-*Q*			-2.53*	0.84	-1.92
Lognormal-*P*				4.60*	2.65*
Heston-*Q*					-4.34*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.17	-0.68	-1.26	-0.04	-0.77
HAR-*P*		0.50	-0.93	1.41	0.27
Lognormal-*Q*			-1.69	1.00	-0.32
Lognormal-*P*				2.37*	1.62
Heston-*Q*					-2.58*

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.96	-0.27	-0.87	0.25	-0.41
HAR-*P*		0.65	-0.74	1.23	0.27
Lognormal-*Q*			-1.85	0.90	-0.34
Lognormal-*P*				1.98*	1.25
Heston-*Q*					-1.82

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.98	-0.09	-0.91	0.57	-0.07
HAR-*P*		0.33	-0.61	0.89	0.25
Lognormal-*Q*			-1.66	0.78	-0.01
Lognormal-*P*				1.45	0.88
Heston-*Q*					-1.82

AT & T 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.93*	-2.60*	-3.13*	2.10*	2.29*
HAR-*P*		0.58	-2.25*	2.36*	2.99*
Lognormal-*Q*			-2.92*	2.30*	2.81*
Lognormal-*P*				2.39*	3.05*
Heston-*Q*					-2.82*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-2.17*	-1.39	-2.40*	2.10*	2.46*
HAR-*P*		-0.03	-1.49	2.17*	2.58*
Lognormal-*Q*			-1.68	2.09*	2.43*
Lognormal-*P*				2.19*	2.59*
Heston-*Q*					-3.10*

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.51	-1.56	-2.01*	2.03*	2.29*
HAR-*P*		-0.10	-1.38	2.09*	2.44*
Lognormal-*Q*			-1.40	2.02*	2.32*
Lognormal-*P*				2.12*	2.50*
Heston-*Q*					-2.67*

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.16	-1.45	-1.54	1.73	1.77
HAR-*P*		0.11	-0.91	1.83	1.91
Lognormal-*Q*			-0.82	1.71	1.77
Lognormal-*P*				1.82	1.92
Heston-*Q*					-2.18*

Walmart 1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-3.72*	-3.69*	-4.34*	-1.97*	-3.72*
HAR-*P*		2.02*	-3.12*	3.21*	-2.84*
Lognormal-*Q*			-3.49*	1.74	-2.02*
Lognormal-*P*				4.30*	3.10*
Heston-*Q*					-3.22*

1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.56	-0.94	-1.79	0.29	-0.95
HAR-*P*		0.44	-1.31	1.30	0.67
Lognormal-*Q*			-2.15*	1.11	0.17
Lognormal-*P*				1.74	2.04*
Heston-*Q*					-1.34

2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.95	-0.07	-1.05	0.90	0.09
HAR-*P*		0.58	-0.93	1.72	1.03
Lognormal-*Q*			-1.44	1.23	0.21
Lognormal-*P*				2.09*	1.69
Heston-*Q*					-2.05*

1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.72	0.25	-0.64	1.01	0.35
HAR-*P*		0.51	-0.36	1.48	0.92
Lognormal-*Q*			-0.83	0.76	-0.05
Lognormal-*P*				1.75	1.31
Heston-*Q*					-1.34

American Express					
1 day	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-1.82	-1.47	-1.80	-0.40	-2.04*
HAR-*P*		2.34*	0.54	2.64*	-0.41
Lognormal-*Q*			-2.29*	2.03*	-1.78
Lognormal-*P*				2.55*	-0.49
Heston-*Q*					-1.69
1 week	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.79	-0.78	-1.01	0.35	-0.07
HAR-*P*		0.61	-1.45	1.41	1.28
Lognormal-*Q*			-1.52	1.20	0.86
Lognormal-*P*				1.77	1.86
Heston-*Q*					-1.57
2 weeks	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.55	-0.32	-0.75	0.16	-0.04
HAR-*P*		0.87	-1.49	1.20	0.87
Lognormal-*Q*			-1.86	0.73	0.37
Lognormal-*P*				1.70	1.48
Heston-*Q*					-1.60
1 month	HAR-*P*	Lognormal-*Q*	Lognormal-*P*	Heston-*Q*	Heston-*P*
HAR-*Q*	-0.57	-0.15	-0.68	0.29	-0.03
HAR-*P*		0.48	-0.82	0.99	0.44
Lognormal-*Q*			-0.94	0.72	0.14
Lognormal-*P*				1.57	0.92
Heston-*Q*					-1.58

Table 5.14 summarises the test statistics for the AG test for six methods and four horizons for IBM and another sixteen stocks when the Newey-West adjustment is made to the estimated variance of $\bar{d}$ and 20 autocorrelations are used. For IBM stock the AG test has similar test values and the same conclusions. The insignificant values become -0.33 and 0.99 when twenty autocorrelations are considered. The AG test results show that the best method for one week horizon is significantly better than two of the remaining five methods at the 5% level, and the best method is statistically better than one method at the 5% level for two weeks horizon, while the best method is statistically indifferent to the other methods at the longest, one month horizon, when the Newey-West adjustment is employed.

5.5 Conclusions

We compare density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices by using Heston, lognormal Black-Scholes, lognormal HAR-RV and transformed densities. Our comparison criterion is the log-likelihood of observed stock prices. For the sixty-eight combinations from seventeen stocks for four horizons, the transformed, lognormal Black-Scholes model gives the highest log-likelihoods for fifty-nine combinations. The HAR-RV model and the Heston model have similar forecast accuracy for different horizons, either before or after applying a transformation which enhances the densities.

Jiang and Tian (2005) suggest that daily option prices are more informative than daily and intraday index returns when forecasting the volatility of the S&P 500 index over horizons from one to six months. Shackleton et al. (2010) similarly imply that option prices are more informative when based on mid-term forecast horizons due to the forward-looking nature of option prices. They only use option prices for the contracts with maturities of more than one week, hence the short horizons of one day and one week density forecasts are extrapolations which are not backed by active trading. They state that the historical density is best for the one day horizon as we can forecast the volatility for tomorrow accurately by calculating the realised variance from recent high-frequency returns.

Most density research only focuses on either risk-neutral densities or ex post real-world density forecasts for horizons matching option expiry dates, while we generate ex ante real-world densities for different forecast horizons. We use a nonparametric transformation to transform the risk-neutral density into real-world density. The log-likelihoods for the nonparametric transformation are always higher than those under the risk-neutral measure for all methods and horizons. The nonparametric transformation also gives better diagnostic test results. Hence central banks, risk managers and other decision takers should not merely focus on risk-neutral densities, but should also obtain more accurate predictions by using risk transformations applied to risk-neutral densities. The relatively unsatisfactory performance of the Heston model for individual firms might be attributed to the illiquidity of their out-of-the-money options. Compared to the index, the individual firm stocks options have fewer strikes that are traded.

Density forecasts can be applied in many areas. They can be used to estimate the risk aversion of investors. They can also be employed to infer probabilities of future market changes for different asset classes including stock indices, interest rates, exchange rates and commodities. In particular, it can help to analyse the impact of a market crash, whether it is anticipated by investors, and whether it is a temporary phenomenon or it results from market failure. Furthermore, they can be used to assess market beliefs about future economic and political events when derived from option prices due to its forward looking nature. The ex-ante analysis infers the possible outcome of the market due to the event, while the ex posts analysis checks if the market reacts to the event as expected. Last but not least, density forecasts are important in risk management, particular for the estimation of Value-at-Risk, which measures how much one can lose at a pre-defined confidence interval over different horizons. Many institutions, such as investment banks and central banks, periodically publish their density estimates, which enable investors to assess risk for their investment portfolios. Hence density forecasts are of importance to central bankers and other decision takers for activities such as policy-making, risk management and derivatives pricing. Concerning the current study, investors should use options information rather than stock prices when choosing their pricing and forecasting models. The simple Black-Scholes lognormal model performs better than the stochastic volatility Heston model. And a nonparametric transformation from risk-neutral densities to real-world densities always give more accurate forecasts.

Appendix. Assumptions about prices, dividends and options

Stock prices jump when dividends are assigned. We apply the Heston dynamics to futures prices which do not jump. We also need to assume all synthetic futures prices have the same dynamics. We assume futures and options contracts expire at time T_1, and there is a dividend at time τ_1 between time 0 and time T_1. The second expiry time for futures and options is T_2 and there is another dividend at τ_2 between time T_1 and T_2. We can use the same dynamics for all futures from simple dividend assumptions; this is easy for continuous dividends but harder for discrete dividends. We denote the futures price at t for delivery at T to be $F_{t,T}$. Our discussion below refers to dividend constants $c_1, c_2, \ldots$, which do not need to be calculated.

We assume, at time t before time τ_i, that the expected dividends are

$$\begin{aligned} E_t[dividend\ at\ \tau_i] &= c_1 e^{r(\tau_1 - t)} S_t && \tau = 1, t < \tau_1 \\ &= c_2(1 - c_1) e^{r(\tau_2 - t)} S_t && \tau = 2, t < \tau_2 \\ &= c_3(1 - c_1)(1 - c_2) e^{r(\tau_3 - t)} S_t && \tau = 3, t < \tau_3 \end{aligned}$$

etc. We assume futures prices are set by no-arbitrage conditions, so

$$F_{t,T} = e^{r(T-t)}[S_t - PV(expected\ dividends\ from\ t\ to\ T)].$$

Then for the first contract

$$\begin{aligned} F_{t,T_1} &= e^{r(T_1 - t)}\left[S_t - e^{-r(\tau_1 - t)} c_1 e^{r(\tau_1 - t)} S_t\right] \\ &= (1 - c_1) e^{r(T_1 - t)} S_t && 0 \le t < \tau_1, \\ &= e^{r(T_1 - t)} S_t && \tau_1 \le t \le T_1. \end{aligned}$$

Then we have

$$ln\left(F_{t,T_1}/S_t\right) = ln(1-c_1) + r(T_1 - t) \qquad 0 \le t < \tau_1,$$

$$= r(T_1 - t) \qquad \tau_1 \le t \le T_1.$$

Thus

$$d\left(lnF_{t,T_1}\right) = d(lnS_t) - rdt \qquad t \neq \tau_1$$

Also S_t jumps down by $c_1 S_{\tau_1}$ at time $t=\tau_1$, but F_{t,T_1} does not jump at $t=\tau_1$.

Similarly, for the second contract

$$F_{t,T_2} = e^{r(T_2-t)}\left[S_t - e^{-r(\tau_1-t)}c_1 e^{r(\tau_1-t)}S_t - e^{-r(\tau_2-t)}c_2(1-c_1)e^{r(\tau_2-t)}S_t\right]$$

$$= e^{r(T_2-t)}(1-c_1)(1-c_2)S_t \qquad 0 \le t < \tau_1,$$

$$= e^{r(T_2-t)}(1-c_2)S_t \qquad \tau_1 \le t < \tau_2,$$

$$= e^{r(T_2-t)}S_t \qquad \tau_2 \le t \le T_2.$$

Hence we have

$$d\left(lnF_{t,T_2}\right) = d(lnS_t) - rdt \qquad t \neq \tau_1, \tau_2,$$

$$= d\left(lnF_{t,T_1}\right) \qquad 0 \le t \le T_1.$$

And we also have

$$\frac{F_{t,T_2}}{F_{t,T_1}} = e^{r(T_2-T_1)}(1-c_2) \qquad 0 \le t \le T_1.$$

We estimate the Heston parameters from the prices of European options which expire at $T_1, T_2, \ldots, T_N$, and strike prices are available as $K_{i,j}$, with $1 \le i \le N$ and $1 \le j \le n_i$. At time 0 we have Black-Scholes implied volatilities $\sigma_{i,j}$, these give market prices from the standard formula for options on futures,

$$c_{i,j} = c_B\left(F_{0,T_i}, T_i, K_{i,j}, r, \sigma_{i,j}\right).$$

Here we have

$$F_{0,T_i} = e^{rT_i}[S_0 - PV(expected\ dividends\ from\ zero\ to\ T_i]$$

and S_0 is the spot price.

Our target is to estimate the Heston parameters θ as:

$$\hat{\theta} = arg\,\min_{\theta} \sum_i \sum_j \left[c_{ij} - c_{Heston}\left(F_{0,T_i}, T_i, K_{ij}, r, \theta\right)\right]^2$$

At time 0 and for any future time τ, we can obtain the density of $S_\tau = F_{\tau,\tau}$ by evaluating the Heston-density with initial price $F_{0,\tau}$ and parameters $\hat{\theta}$.

6. Conclusions

One-minute returns of ten foreign exchange rates are investigated for five years from 2007 to 2011. We employ the ABD and LM jump detection tests to detect intraday price jumps for ten rates and cojumps for six groups of two dollar rates and one cross rate. The null hypothesis that jumps are independent is rejected, as there are far more cojumps than predicted by independence for all rate combinations. Some clustering of jumps and cojumps are also detected and can be related to the macroeconomic news announcements affecting the exchange rates. The selected ABD and LM jump detection tests detect a similar number of jumps for ten foreign exchange rates.

Foreign exchange rates contain frequent and relatively small jumps as they are usually affected by two sources of news and they have more liquidity shocks during the continuously traded 24-hour market. Some foreign exchange rates jump and cojump more than others, this is because some exchange rates are closely correlated, or it is easy to simultaneously trade some exchange rates. For example, the U.S. scheduled macroeconomic news announcements may affect all dollar exchange rates, and some European news may affect both euro and pound exchange rates.

Previous studies such as Lahaye et al. (2011) only investigate dollar rates, while we examine more currencies through checking six groups of two dollar rates and one cross rate at the higher one minute frequency for ten years. We find that one dollar rate and the cross rate combination almost always has more cojumps than the two dollar rates combination.

We compare density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices by using Heston, lognormal Black-Scholes, lognormal HAR-RV and transformed densities. We base comparisons on the log-likelihood of observed stock prices. For the sixty-eight combinations from seventeen stocks for four horizons, the transformed lognormal Black-Scholes model gives the highest log-likelihoods for fifty-nine combinations. The HAR-RV model and the Heston model gives the highest log-likelihood for a similar number of times, either before or after applying a nonparametric transformation.

Jiang and Tian (2005) argue that daily option prices are more informative than daily and intraday index returns when forecasting the volatility of the S&P 500 index over horizons from one to six months. Shackleton et al. (2010) also imply that option prices contain more information when based on mid-term forecast horizons due to the forward-looking nature of option prices. They only use option prices for the contracts with maturities of more than seven calendar days, hence the short horizons of one day and one week density forecasts are extrapolations which are not supported by frequent trading. They state that the historical density is best for the one day horizon as we can forecast the volatility for tomorrow accurately by calculating the realised variance from recent high-frequency returns.

Most density research only focuses on either risk-neutral densities or ex post real-world densities for horizons matching option expiry dates, while we generate ex ante real-world densities for different forecast horizons. We use a nonparametric transformation to transform the risk-neutral densities into the real-world densities. The log-likelihoods for the nonparametric transformation are always higher than those under the risk-neutral measure for all methods and horizons. The nonparametric transformation also provides better diagnostic test results. Hence central banks, risk managers and other investors should not only look at risk-neutral densities, but also obtain more accurate predictions by using risk transformations applied to risk-neutral densities. The relatively unsatisfactory performance of the Heston model for individual firms might be attributed to the illiquidity of their out-of-the-money options. Compared to the index, the individual firm stocks options have fewer strikes that are traded.

There are several possible directions of future research to point out. Concerning the foreign exchange rates jump and cojump study, it should be more helpful if we can use more detailed sources of macroeconomic news announcements and employ some models to formally assess the effect of macroeconomic news announcements on jumps and cojumps. Additionally, we can conduct a Monte Carlo simulation to compare the size and power of the ABD and the LM jump detection tests. Regarding the density forecast study, since Pong et al. (2004) state that the better accuracy of volatility forecasts comes from the high-frequency data, but not necessarily from a long memory specification. We could check if long memory models (e.g. ARFIMA) for realised variance can improve density forecasts obtained using high-frequency data. Moreover, we evaluate density forecasts using log-likelihoods and diagnostic tests. It might be interesting if we can also make comparisons based on some risk management application, such as the value-at-risk. Last but not least, Shackleton et al. (2010) focus on a U.S. stock index and get different findings, hence we could extend the analysis to other asset classes, such as currencies, commodities and interest rates, to see what findings we can get.

References

Ait-Sahalia, Y. and J. Jacod, 2009. Testing for jumps in a discretely observed process. *Annals of Statistics* 37, 184-222.

Ait-Sahalia, Y. and A.W. Lo, 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499-547.

Ait-Sahalia, Y. and A.W. Lo, 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9-51.

Ait-Sahalia, Y., P.A. Mykland and L. Zhang, 2011. Ultra-high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics* 160, 160-175.

Amisano, G.G. and R. Giacomini, 2007. Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business and Economic Statistics* 25, 177-190.

Anagnou-Basioudis, I., M. Bedendo, S.D. Hodges and R. Tompkins, 2005. Forecasting accuracy of implied and GARCH-based probability density functions. *Review of Futures Markets* 11, 41-66.

Andersen, T.G. and T. Bollerslev, 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115-158.

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

Andersen, T.G., T. Bollerslev and F.X. Diebold, 2007. Roughing it up: including jump components in measuring, modeling and forecasting asset return volatility. *Review of Economics and Statistics* 89, 701-720.

Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens, 2001. The distribution of stock return volatility. *Journal of Financial Economics* 61, 43-67.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2001. The distribution of realised exchange rate volatility. *Journal of American Statistical Association* 96, 42-55.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003. Modelling and forecasting realised volatility. *Econometrica* 71, 579-625.

Andersen, T.G., T. Bollerslev, F.X. Diebold and C. Vega, 2003. Micro effects of macro announcements: real-time price discovery in foreign exchange. *American Economic Review* 93, 38–62.

Andersen, T.G., T. Bollerslev, F.X. Diebold and C. Vega, 2007. Real-time price discovery in stock, bond and foreign exchange markets. *Journal of International Economics* 73, 251–277.

Andersen, T.G., T. Bollerslev and D. Dobrev, 2007. No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: theory and testable distributional implications. *Journal of Econometrics* 138, 125-180.

Andersen, T.G., D. Dobrev and E. Schaumburg, 2012. Robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics* 169, 75-93.

Bakshi, G., C. Cao and Z. Chen, 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003-2049.

Bakshi, G., N. Kapadia and D.B. Madan, 2003. Stock return characteristics, skew laws and the differential pricing of individual equity options. *Review of Financial Studies* 16, 101-143.

Bali, T.G. and D. Weinbaum, 2007. A conditional extreme value volatility estimator based on high-frequency returns. *Journal of Economic and Dynamic Control* 31, 361-397.

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

Bao, Y., T.H. Lee and B. Saltoglu, 2007. Comparing density forecast models. *Journal of Forecasting* 26, 203-225.

Barndorff-Nielsen, O.E. and N. Shephard, 2004a. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1–37.

Barndorff-Nielsen, O.E. and N. Shephard, 2004b. Measuring the impact of jumps in multivariate price processes using bipower variation. *Working Paper, University of Oxford*.

Barndorff-Nielsen, O.E. and N. Shephard, 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4, 1-60.

Bates, D., 1991. The crash of '87: was it expected? The evidence from the options markets. *Journal of Finance* 46, 1009-1044.

Bates, D., 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. *Review of Financial Studies* 9, 69-107.

Bates, D., 2000. Post-'87 crash fears in S&P 500 futures options. *Journal of Econometrics* 94, 181-238.

Becker, R., A.E. Clements and S.I. White. 2007. Does implied volatility provide any information beyond that captured in model-based volatility forecasts? *Journal of Banking and Finance* 31, 2535-2549.

Benzoni, L., 2002. Pricing options under stochastic volatility: an empirical investigation. *Working Pape*r, *University of Minnesota*.

Berkowitz, J., 2001. Testing density forecasts, with applications to risk management. *Journal of Business and Economic Statistics* 19, 465-474.

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

Blair, B.J., S.H. Poon and S.J. Taylor, 2001. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high frequency index returns. *Journal of Econometrics* 105, 5-26.

Bliss, R. and N. Panigirtzoglou, 2002. Testing the stability of implied probability density functions. *Journal of Banking and Finance* 26, 381-422.

Bliss, R. and N. Panigirtzoglou, 2004. Option-implied risk aversion estimates. *Journal of Finance* 59, 407-446.

Bollerslev, T., T.H. Law and G. Tauchen, 2008. Risk, jumps, and diversification. *Journal of Econometrics* 144, 234-256.

Bookstaber, R.M. and J.B. McDonald, 1987. A general distribution for describing security price returns. *Journal of Business* 60, 401-424.

Breeden, D. and R. Litzenberger, 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621-651.

Britten-Jones, M. and A. Neuberger, 2000. Option prices, implied prices processes and stochastic volatility. *Journal of Finance* 55, 839-866.

Bunn, D.W., 1984. Applied Decision Analysis. McGraw-Hill.

Busch, T., B.J. Christensen and M.Ø. Nielsen, 2011. The role of implied volatility in forecasting future realised volatility and jumps in foreign exchange, stock, and bond markets. *Journal of Econometrics* 160, 48–57.

Caporin, M., A. Kolokolov and R. Reno, 2014. Muti-jumps. *Working Paper, University of Siena*.

Carr, P. and L. Wu, 2003. Finite moment log stable process and option pricing. *Journal of Finance* 58, 753-777.

Carr, P. and L. Wu, 2009. Variance risk premiums. *Review of Financial Studies* 22, 1311-1341.

Carr, P. and L. Wu, 2010. Stock options and credit default swaps: a joint framework for valuation and estimation. *Journal of Financial Econometrics* 8, 409-449.

Chernov, M. and E. Ghysels, 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of option valuation. *Journal of Financial Economics* 56, 407-458.

Christensen, K., R. Oomen and M. Podolskij, 2010. Realised quantile-based estimation of the integrated variance. *Journal of Econometrics* 159, 74-98.

Christensen, K. and M. Podolskij, 2007. Realised range-based estimation of integrated variance. *Journal of Econometrics* 141, 323-349.

Christoffersen, P. and K. Jacobs, 2004. The importance of the loss function in option valuation. *Journal of Financial Economics* 72, 291-318.

Christoffersen, P., K. Jacobs, and K. Mimouni, 2010. Models for S&P 500 dynamics: evidence from realised volatility, daily returns and option prices. *Review of Financial Studies* 23, 3141-3189.

Clements, M.P., 2004. Evaluating the Bank of England density forecasts of inflation. *Economic Journal* 114, 855-877.

Clements, M.P., P.H. Franses, J. Smith and D. van Dijk, 2003. On SETAR non-linearity and forecasting. *Journal of Forecasting* 23, 359-375.

Clements, M.P. and J. Smith, 2000. Evaluating density forecasts of linear and non-linear models: applications to output growth and unemployment. *Journal of Forecasting* 19, 255-276.

Corsi, F., 2009. A simple approximate long-memory model of realised volatility. *Journal of Financial Econometrics* 7, 174-196.

Corsi, F., D. Pirino and R. Reno, 2010. Threshold bipower variation and the impact of jumps on volatility forecasting. *Journal of Econometrics* 159, 276-288.

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

Dawid, A.P., 1984. Statistical theory, the prequential approach. *Journal of the Royal Statistical Society*, Series A 147, 278-290.

Diebold, F.X., A. Gunther and A.S. Tay, 1998. Evaluating density forecasts with applications to financial risk management. *International Economic Review* 39, 863-883.

Diebold, F.X., J. Hahn and A.S. Tay, 1999. Multivariate density forecast evaluation and calibration in financial risk management: high-frequency returns on foreign exchange. *Review of Economics and Statistics* 81, 661-673.

Diebold, F.X. and R.S. Mariano, 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.

Diebold, F.X. and G. Strasser, 2013. On the correlation structure of microstructure noise: a financial economic approach. *Review of Economic Studies* 80, 1304-1337.

Duffie, D., J. Pan, and K.J. Singleton, 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343-1476.

Dungey, M. and L. Hvozdyk, 2012. Cojumping: Evidence from the US treasury bond and futures markets. *Journal of Banking and Finance* 36, 1563-1575.

Dungey, M., M. McKenzie and V. Smith, 2009. Empirical evidence on jumps in the term structure of the US treasury market. *Journal of Empirical Finance* 16, 430–445.

Eraker, B., 2004. Do equity prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59, 1367-1403.

Eraker, B., M.S. Johannes, and N.G. Polson, 2003. The impact of jumps in returns and volatility. *Journal of Finance* 53, 1269-1300.

Evans, K.P., 2011. Intraday jumps and US macroeconomic news announcements. *Journal of Banking and Finance* 35, 2511-2527.

Evans, K.P. and A.H. Speight, 2011. Intraday euro exchange rates and international macroeconomic announcements. *European Journal of Finance* 17, 83-110.

Fackler, P.L. and R.P. King, 1990. The calibration of option-based probability assessments in agricultural commodity markets. *American Journal of Agricultural Economics* 72, 73-83.

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

Gatheral, J., 2006. The volatility surface, a practitioner's guide. John Wiley and Sons, Inc.

Gilder, D., M.B. Shackleton and S.J. Taylor, 2014. Cojumps in stocks: empirical evidence. *Journal of Banking and Finance* 40, 443-459.

Giot, P. and S. Laurent, 2007. The information content of implied volatility in light of the jump/continuous decomposition of realised volatility. *Journal of Futures Markets* 27, 337-359.

Gnabo, J.Y., L. Hvozdyk and J. Lahaye, 2014. System-wide tail comovements: A bootstrap test for cojump identification on the S&P 500, US bonds and currencies. *Journal of International Money and Finance* 48, 147-174.

Gobbi, F. and C. Mancini, 2007. Diffusion covariation and co-jumps in bidimensional asset price processes with stochastic volatility and infinite activity levy jumps. *Working Paper, University of Florence*.

Granger, C.W.J. and P. Newbold, 1976. Forecasting economic time series. Academic Press, New York.

Guidolin, M. and A. Timmermann, 2005. Economic implication of bull and bear regimes in UK stock and bond returns. *Economic Journal* 115, 111-143.

Hansen, P.H. and A. Lunde, 2006. Realised variance and market microstructure noise. *Journal of Business and Economic Statistics* 24, 127-161.

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

Hog, E. and L. Tsiaras, 2010. Density forecasts of crude-oil prices using option implied and arch-type models. *Journal of Futures Markets* 31, 727-754.

Huang, X. and G. Tauchen, 2005. The relative contribution of jumps to total price variation. *Journal of Financial Econometrics* 4, 456-499.

Huang, J. and L. Wu, 2004. Specification analysis of option pricing models based on time-changed levy processes. *Journal of Finance* 59, 1405-1439.

Ivanova, V. and J.M. Gutierrez, 2014. Interest rate forecasts, state price densities and risk premium from Euribor options. *Journal of Banking and Finance* 48, 210-223.

Jackwerth, J.C., 2000. Recovering risk aversion from option prices and realised returns. *Review of Financial Studies* 13, 433-451.

Jackwerth, J.C. and M. Rubinstein, 1996. Recovering probability distributions from option prices. *Journal of Finance* 51, 1611-1631.

Jacod, J., 2008. Asymptotic properties of realised power variations and related functionals of semimartingales. *Stochastic Processes and their Applications* 118, 517-559.

Jacod, J., Y. Li, P. Mykland, M. Podolskij and M. Vetter, 2009. Microstructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and their Applications* 119, 2249-2276.

Jacod, J. and V. Todorov, 2009. Testing for common arrival of jumps in discretely-observed multidimensional processes. *Annuals of Statistics* 37, 1792-1838.

Jiang, G.J. and R.A. Oomen, 2008. Testing for jumps when asset prices are observed with noise-a ‘swap variance’ approach. *Journal of Econometrics* 144, 352-370.

Jiang, G.J. and Y.S. Tian, 2005. The model-free implied volatility and its information content. *Review of Financial Studies* 13, 433-451.

Jondeau, E. and M. Rockinger, 2000. Reading the smile: the message conveyed by methods which infer risk neutral densities. *Journal of International Money and Finance* 19, 885-915.

Kahl, C. and P. Jackel, 2006. Not-so-complex logarithms in the Heston model. *Working paper, ABN AMRO*.

Kang, B.J. and T.S. Kim, 2006. Option-implied risk preferences: an extension to wider classes of utility functions. *Journal of Financial Markets* 9, 180-198.

Kendall, M.G., A. Stuart and J.K. Ord, 1983. The advanced theory of statistics, vol.3, 4th edition. London: Charles Griffin.

Kostakis, A., N. Panigirtzoglou and G. Skiadopoulos, 2011. Market Timing with Option-Implied Distributions: A Forward-Looking Approach, *Management Science* 57, 1231-1249.

Lahaye, J., S. Laurent and C.J. Neely, 2011. Jumps, cojumps and macro announcements. *Journal of Applied Econometrics* 26, 893-921.

Lee, S.S. and J. Hannig, 2010. Detecting jumps from levy jump diffusion processes. *Journal of Financial Economics* 96, 271-290.

Lee, S.S. and P. Mykland, 2008. Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial Studies* 21, 2535-2563.

Liu, X., M.B. Shackleton, S.J. Taylor and X. Xu, 2007, Closed-form transformations from risk-neutral to real-world densities. *Journal of Banking and Finance* 31, 1501-1520.

Madan, D.B. and F. Milne, 1994. Contingent claims valued and hedged by pricing and investing in a basis. *Mathematical Finance* 4, 223-245.

Mancini, C., 2004. Estimating the integrated volatility in stochastic volatility models with levy type jumps. *Working Paper, University of Florence*.

Mancini, C., 2009. Nonparametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics* 36, 270-296.

Martens, M. and D. van Dijk, 2007. Measuring volatility with the realised range. *Journal of Econometrics* 138, 181-207.

Martens, M. and J. Zein, 2004. Predicting financial volatility: High-frequency time-series forecast vis-à-vis implied volatility. *Journal of Futures Markets* 24, 1005-1028.

Martin, G.M., A. Reidy and J. Wright, 2009. Does the option market produce superior forecasts of noise-corrected volatility measures? *Journal of Applied Econometrics* 24, 77-104.

Melick, W.R. and C.P. Thomas, 1997. Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91-115.

Merton, R.C., 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141-183.

Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125-144.

Pan, J., 2002. The jump risk premium implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3-50.

Podolskij, M. and M. Vetter, 2009a. Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps. *Bernoulli* 15, 634-658.

Podolskij, M. and M. Vetter, 2009b. Bipower-type estimation in a noisy diffusion setting. *Stochastic Process and their Applications* 119, 2803-2831.

Podolskij, M. and D. Ziggel, 2010. New tests for jumps in semimartingale models. *Statistical Inference for Stochastic Processes* 13, 15-41.

Pong, S.H., M.B. Shackleton, S.J. Taylor and X. Xu, 2004. Forecasting currency volatility a comparison of implied volatilities and ar(fi)ma models. *Journal of Banking and Finance* 28, 2541-2563.

Poon, S.H. and C.W.J. Granger, 2003. Forecasting volatility in financial markets: a review. *Journal of Economic Literature* 41, 478-539.

Ritchey, R.J., 1990. Call option valuation for discrete normal mixtures. *Journal of Financial Research* 13, 285-296.

Rosenberg, J. and R.F. Engle, 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 341-372.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. *Annuals of Mathematical Statistics* 23, 470-472.

Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 771-818.

Sarno, L. and G. Valente, 2004. Comparing the accuracy of density forecasts from competing models. *Journal of Forecasting* 23, 541-557.

Shackleton, M.B., S.J. Taylor and P. Yu, 2010. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. *Journal of Banking and Finance* 34, 2678-2693.

Silverman, B.W., 1986. Density estimation for statistical and data analysis. Chapman Hall.

Taylor, S.J., 2005. Asset price dynamics, volatility and prediction, Princeton University Press.

Taylor, S.J. and X. Xu, 1997. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4, 317-340.

Taylor, S.J., P.K. Yadav and Y. Zhang, 2010. The information content of implied volatilities and model-free volatility expectations: evidence from options written on individual stocks. *Journal of Banking and Finance* 34, 871-881.

Trujillo-Barrera, A., P. Garcia and M. Mallory, 2012. Density forecasts of lean hog futures prices. Proceedings of the NCCC-134 conference on applied commodity price analysis, forecasting, and market risk management. St. Louis, MO. [http://www.farmdoc.illinois.edu/nccc134].

Xu, X. and S.J. Taylor, 1995. Conditional volatility and informational efficiency pf the PHLX currency options market. *Journal of Banking and Finance* 19, 803-821.

Yun, J., 2014. Out-of-sample density forecasts with affine jump diffusion models. *Journal of Banking and Finance* 47, 74-87.

Zhang, L., 2006. Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach. *Bernoulli* 12, 1019-1043.

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

Zhou, B., 1996. High-frequency data and volatility in foreign exchange rates. *Journal of Business and Economic Statistics* 14, 45-52.