

Patterns Within the Trading Day: Volatility and Jump Discontinuities in High Frequency Equity Price Series^{*†}

Peter Van Tassel[‡]

Professor George Tauchen, Faculty Advisor

^{*}The Duke Community Standard was upheld in the completion of this report.

[†]Honors Thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University.

[‡]Peter.VanTassel@Duke.edu. Box 93481. Duke University; Durham, North Carolina; 2007.

Acknowledgment

In completing this work I would like to thank the members of the Duke Econometrics and Finance Lunch Group for their continued support. I am particularly grateful for having had the opportunity to work with George Tauchen, Tim Bollerslev and Bjorn Eraker. Their comments and feedback have been an integral component of my research. I would also like to thank Tzuo Hann Law, Andrey Fradkin and Warren Davis for their helpful suggestions and willingness to collaborate.

Abstract

This paper identifies systematic patterns within the trading day by analyzing high frequency data from a market index and nine individual stocks. Empirical results expand on the previously documented *U*-shape in intraday equity volatility by implementing non-parametric statistics to test for patterns in the jump and diffusive components of volatility. Additional results indicate that a recently developed non-parametric jump detection scheme may under-report the number of returns flagged as statistically significant jumps in the middle of the day while exaggerating the number of statistically significant jumps in the early morning and late afternoon. The paper concludes by investigating whether incorporating the observed patterns into a historical forecasting model can improve performance.

1 Introduction

In financial markets there are rare events during which asset prices exhibit volatile and unexpected movements. These occurrences are of utmost concern for traders and portfolio managers. Commonly referred to as jumps, they correspond to abrupt shifts in the expectation of market participants that underscore new risks and uncertainties about the future of the economy or the future of an individual firm.

There are many alleged explanations for why jumps occur. Some suggest that jumps are information driven. Among financial practitioners it is well known that scheduled macroeconomic and idiosyncratic announcements that differ significantly from expectations can have dramatic effects on stock prices, equity indices, fixed-income securities and foreign exchange rates. Another hypothesis is suggested by the popular literature in behavioral finance. Shiller (2002) argues that market participants tend to create unsustainable situations that result in volatile and highly correlated changes in asset prices. Others believe that jumps are a product of the so-called “Black Swan” events. Taleb (2007) cites unpredictable occurrences for having a large impact on asset prices.

Recent research in financial econometrics has placed an increasing emphasis on studying jumps. In part this is a response to a long-standing criticism of the asset pricing literature. Despite the fact that papers as early as Merton (1976) motivate the importance of including jumps into models for valuing derivatives such as option prices, many of the continuous time models for asset price evolution have sample paths that do not allow for jump discontinuities. Andersen, Bollerslev and Diebold (2006) note that the assumption of a continuous sample path in theoretical models for asset pricing is clearly violated in practice. Eraker, Johannes, and Polson (2003) and Eraker (2004) expand upon this point. They detail the importance of including jumps into models for stochastic volatility and the pricing of derivatives.

The interest in jumps has also been apparent in developments made by the non-parametric literature. Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold and Labys (2001), and Barndorff-Nielsen and Shephard (2002a,b) recommend the use of high-frequency financial data to compute the non-parametric realized variance as a method for estimating, forecasting, and modeling volatility. They conclude that the realized variance maintains much of the information in intraday data while eliminating many of the complications related to working with parametric models. A series of papers by Barndorff-Nielsen and Shephard (2004, 2006) expand on this result by introducing the bi-power variation as a non-parametric statistic that provides a consistent estimator for the diffusive component in volatility. They go on to show that the difference in the realized variance and the

bi-power variation provides a consistent estimator of the jump component in volatility. Their result provides the theoretical framework to study jumps. Huang and Tauchen (2005) implement this idea to develop a new statistic that they validate through extensive Monte Carlo analysis. It tests for statistically significant trading days under the null hypothesis that there are no jump discontinuities in asset prices. Lee and Mykland (2006) introduce a new test that builds on the results from Huang and Tauchen (2005) and the framework in Barndorff-Nielsen and Shephard (2004). The Lee-Mykland statistic differentiates itself from the other jump detection schemes in that it tests individual returns for statistically significant jump discontinuities under the null hypothesis that there are no jumps in asset prices. The advantage of the Lee-Mykland statistic is that it identifies when jumps occur throughout the trading day.

This paper identifies systematic patterns that persist in equity prices within the trading day. It expands on the non-parametric literature in financial econometrics in several ways. First, it is among the earliest attempts to implement the Lee-Mykland statistic to a high frequency data set to investigate the arrival of statistically significant jump discontinuities within the trading day. It also includes the most extensive application of the Lee-Mykland statistic to individual stocks. Whereas Lee and Mykland (2006) use a three month sample for three individual stocks and the market, this paper implements the statistic over a five year sample for nine individual stocks and the market. The empirical results confirm the intuition from Lee and Mykland (2006). They indicate that statistically significant jump discontinuities are most frequent in the first hour of trading. This is particularly noted for the idiosyncratic jumps in individual stocks that appear to cluster around 10:00am.

Further analysis reveals that the results using the Lee and Mykland statistic may be spurious. There is evidence the Lee-Mykland statistic over-specifies the number of statistically significant jumps in the morning and late afternoon while under reporting the number of jumps in the middle of the day. This discussion relates to intraday patterns in volatility. Wood, McInish and Ord (1985) document a *U*-shaped pattern in intraday equity volatility where the morning corresponds to the most volatile period of the day. This paper builds upon their result. The bi-power variation developed by Barndorff-Nielsen and Shephard (2004) provides a consistent estimator for the diffusive component in volatility whereas the difference between the realized variance and the bi-power variation provides a consistent estimator for the jump component in equity volatility. These statistics are computed in the first attempt to test for intraday patterns in decomposition of equity volatility into its diffusive and jump components. The results confirm the *U*-shaped pattern observed in Wood, McInish and Ord (1985) for the realized variance and indicate that the diffusive component exhibits a similar pattern. The pattern in the jump component is more skewed. It is particularly volatile

during the first hour of trading. These results have important implications for the Lee-Mykland statistic. The Lee-Mykland statistic compares individual returns to a backward looking average of diffusive component in volatility. At the recommended window sizes for the statistic in Lee and Mykland (2006) it seems that local patterns in volatility are eliminated by taking the average of the diffusive component in volatility over a relatively long period of time. This may result in Type I errors that exaggerate the reported number of statistically significant jumps in the morning and afternoon and Type II errors that reduce the reported number of statistically significant jumps in the middle of the day.

The intraday patterns in volatility are then incorporated into a historical forecasting model to see whether they can improve performance. Andersen, Bollerslev, Diebold, and Labys (2003) report that forecasting with realized variance outperforms the well known GARCH and a variety of more complicated stochastic volatility models in out-of-sample forecasting. This paper considers the HAR-RV-CJ forecasting model introduced in Andersen, Bollerslev, and Diebold (2006). A modification to the model tests whether there is informational content in the jump component of equity volatility during the first hour of trading. The informational content is measured by its ability to improve the performance of the HAR-RV-CJ forecast over daily, weekly and monthly horizons.

The remainder of the paper will proceed as such: Section 2 presents the model for stock price evolution and the statistical measures considered throughout the subsequent discussion; Section 3 describes the data; Section 4 discusses systematic patterns in the arrival of statistically significant jump discontinuities and equity volatility; Section 5 investigates whether the observed patterns are helpful in improving historical forecasts of volatility; Section 6 concludes. All references, tables and figures can be found at the back of the paper.

2 Modeling Jump Components

2.1 Log-Price Stock Evolution

To motivate our discussion of jump discontinuities in stock prices we will begin by investigating a standard model of stock price evolution. Consider the stochastic differential equation for log-price $dp(t)$ defined as

$$dp(t) = \mu \cdot dt + \sigma \cdot dW(t). \quad (1)$$

The first term, $\mu \cdot dt$, is a deterministic drift component in the stock price. The second term, $\sigma \cdot dW(t)$, is a random component in the stock price defined as volatility multiplied by a standard Brownian motion. The model represents the limiting case of one of the most simple models for stock price evolution. In particular, the equation defined above is merely the continuous time interpretation of a Binomial model in which the stock price has a positive probability of moving up or down over small time intervals. The drift term and the standard Brownian motion are merely the mathematical consequences of adding the independent identically distributed log-returns over infinitesimally small time periods. When viewed over longer periods of time it will seem that stock prices in this model tend to drift higher with a considerable amount of random behavior or noise, an example can be found in Figure 9 in the Technical Appendix.

This model is a good starting point for investigation because it has played an important role in the financial literature for its intuitive appeal and mathematical simplicity. A notable example would be its use in deriving the well known Black-Scholes option pricing formula in Merton (1973) and Black and Scholes (1973). Of course, the disadvantage of simple models is that they often fail to encompass the complexities of the real world. One of the first criticisms of the standard model above is its failure to allow for the time varying volatility clustering found in Engle (1982) and Bollerslev (1986). The random component of the stock price merely multiplies a real number σ by a normal random variable with mean 0 and standard deviation t . A more recent criticism of the standard model is that it also fails to account for jumps discontinuities or large price moves. Again this is a mathematical consequence. Standard Brownian motions have continuous sample paths with probability one. The implication is that the model restricts stock price movements in a manner that does not seem to be aligned with the real world.

One example that comes to mind is risk-arbitrage. In many situations there are significant price movements associated with the success of deals or lawsuits. Outcomes can result in large price moves that would practically be zero probability events under the standard model described above. As a result, we introduce a new model that takes into consideration these complications and allows for more dramatic price movements. Here we define log-price evolving in continuous time as

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) + dL_J(t). \quad (2)$$

Note that both the drift and instantaneous volatility are re-defined so they are time-varying. As before $W(t)$ is a standard Brownian motion. The new term, $L_J(t)$, is a particular class of the pure jump Lévy process with increments $L_J(t) - L_J(s) = \sum_{s \leq \tau \leq t} \kappa(\tau)$. Our notation is adopted from

Basawa and Brockwell (1982). The jump process we consider is a compound Poisson process. It allows for discontinuities or jumps in stock prices with constant jump intensity λ and independent identically distributed jump size $\kappa(\tau)$. This is the model we will consider throughout the rest of our discussion. It underlies all of the work in the upcoming section where we introduce statistical measures that will allow us to study jump discontinuities or large price moves in stocks.

2.2 Realized Variance and Bipower Variation

It may seem logical that a good place to begin looking for jump discontinuities is by investigating the level of stock prices over our sample period. Figure 10 found in the Technical Appendix motivates this idea by plotting the closing price from 2001 through 2005 for Standard and Poor's Depository Receipt, a well known exchange-traded fund that tracks the performance of the S&P 500. The plot reveals that it is actually quite difficult to visually isolate days where a large price move may have occurred, at least when sufficiently large sample periods are considered. From this elementary analysis it is possible to surmise that we will need to develop more sophisticated tools for identifying jump discontinuities in stock prices.

To begin we shift our focus from investigating the level of stock prices to analyzing stock returns, or the changes in stock prices over specified periods of time. Stock returns are computed by sampling high frequency financial data at a constant frequency. On a particular day t we have prices $p(t - 1 + \frac{1}{M})$, $p(t - 1 + \frac{2}{M})$, ..., $p(t - 1 + \frac{M}{M})$ where M is the within day sampling frequency. As before $p(t)$ refers to a log-price whereas $P(t)$ will be used for the level price at some time t . The intraday geometric returns on day t is defined as,

$$r_{t,j} = p(t - 1 + \frac{j}{M}) - p(t - 1 + \frac{j-1}{M}), \quad j = 1, 2, \dots, M; \quad (3)$$

$$r_{t,j} = \ln \left(\frac{P(t - 1 + \frac{j}{M})}{P(t - 1 + \frac{j-1}{M})} \right), \quad j = 1, 2, \dots, M.$$

Figure 10 also found in the Technical Appendix plots geometric returns from the SPY sampled at a daily frequency. By looking at the changes in stock prices instead of the level it is easier to spot days when large price moves occur.

Nonetheless, a visual selection of jump discontinuities is neither rigorous nor convincing. In what follows stock returns are used to compute two measures of integrated volatility that allow for

the study jumps in a more formal manner. Both measures are discussed in Barndorff-Nielsen and Shephard (2004) to analyze the realized or historical price variance.

The first measure will be referred to as the realized variance,

$$RV_t = \sum_{j=1}^M r_{t,j}^2. \quad (4)$$

It sums the squared geometric returns over intraday periods. This provides a consistent estimator for the integrated variance. Andersen, Bollerslev, Diebold (2002) expands upon this idea by noting that the realized variance for a particular day t satisfies

$$\lim_{m \rightarrow \infty} RV_t = \int_{t-1}^t \sigma^2(s) ds + \sum_{j=1}^{N_t} k_{t,j}^2 \quad (5)$$

where $\int_{t-1}^t \sigma^2(s) ds$ is a consistent estimator of the integrated variance, N_t is the number of jumps, and $k_{t,j}$ is the jump size. The second measure to be considered is the realized bipower variation. It is defined as

$$BV_t = \mu_1^{-2} \left(\frac{M}{M-1} \right) \sum_{j=2}^M |r_{t,j-1}| |r_{t,j}| = \frac{\pi}{2} \left(\frac{M}{M-1} \right) \sum_{j=2}^M |r_{t,j-1}| |r_{t,j}| \quad (6)$$

where $\mu_a = E(|Z|^a)$, $Z \sim N(0,1)$, $a > 0$.¹ This statistic sums the absolute value of the intraday return at time j multiplied by the absolute value of the preceding return at time $j-1$ over the course of the trading day. The $\frac{\pi}{2}$ adjustment is a technical consideration so the statistics that will be defined later are distributed $N(0,1)$. The sum is multiplied by $\frac{M}{M-1}$ so it is possible to directly compare the bipower variation which has $M-1$ terms and the realized variance which has M terms for each trading day.

The rationale for computing both of these non-parametric statistics is to use their difference as a consistent estimator for the jump component in equity volatility. As was previously noted, the realized variance is a consistent estimator for the integrated variance plus the jump contribution. The bipower variance provides a consistent estimator of the integrated variance unaffected by jumps. Under what Huang and Tauchen (2005) describe as reasonable assumptions about Equation 2, Barndorff-Nielsen and Shephard (2004, 2006) imply that the difference $RV_t - BV_t$ can be used as a consistent estimator of the pure jump component. To be specific,

¹ $\mu_p \equiv 2^{\frac{p}{2}} \frac{\Gamma(\frac{1}{2}(p+1))}{\Gamma(\frac{1}{2})} = E(|Z|^p)$. We use μ_1^{-2} so the bipower variation and realized variance are directly comparable. The term $\left(\frac{M}{M-1}\right)$ is a degrees of freedom adjustment in bipower variation. Further explanations can be found in Barndorff-Nielsen and Shephard (2004).

$$\lim_{m \rightarrow \infty} BV_t = \int_{t-1}^t \sigma^2(s) ds; \quad (7)$$

$$\lim_{m \rightarrow \infty} (RV_t - BV_t) = \sum_{j=1}^{N_t} k_{t,j}^2. \quad (8)$$

This result is the basis for all of our subsequent work. In the upcoming section it will be used to discuss two non-parametric jump detection schemes that will be used to compute statistically significant jump discontinuities in stock prices.

2.3 Jump Test Statistics

The theoretical framework presented by Barndorff-Nielsen and Shephard (2004) has been implemented to study jump discontinuities and create detection schemes that will flag statistically significant jumps. One example can be found in Andersen, Bollerslev, Diebold (2006). They use realized variance and bipower variation to investigate jumps in equity, fixed income, and foreign exchange markets. Their research suggests that the standard continuous-time stochastic volatility model,

$$dp(t) = \mu(t)dt + \sigma(t)dt,$$

fails to explain the number of large intraday price moves. Huang and Tauchen (2005) expand upon these results by performing extensive Monte Carlo analysis on a series of non-parametric test statistics. They conclude that the average contribution of jumps to total daily price variance is 4.5% to 7.0%. Their preferred statistic will be considered in this paper as a starting point for flagging jumps. It is defined as,

$$z_t = \frac{RJ_t}{\sqrt{(\nu_{bb} - \nu_{qq}) \frac{1}{M} \max(1, \frac{TP_t}{BV_t^2})}}; \quad (9)$$

$$RJ_t = \max\left(\frac{RV_t - BV_t}{RV_t}, 0\right);$$

$$TP_t = M \mu_{4/3}^{-3} \left(\frac{M}{M-2} \right) \sum_{j=3}^M |r_{t,j-2}|^{\frac{4}{3}} |r_{t,j-1}|^{\frac{4}{3}} |r_{t,j}|^{\frac{4}{3}};$$

$$\nu_{qq} = 2, \nu_{bb} = \left(\frac{\pi}{2}\right)^2 + \pi - 3.$$

The numerator or relative jump component, RJ_t , makes use of the difference between the realized variance and bipower variation to test for large contributions in the jump component of equity returns over the course of a trading day. The denominator studentizes the statistic. Tri-Power Quarticity, TP_t , is a jump robust estimator for the integrated quarticity $\int_{t-1}^t \sigma^4(s)ds$ that Huang and Tauchen (2005) conclude is preferred to other estimators for integrated quarticity such as the Quad-Power Quarticity. The modifications to the statistic are made for technical considerations. To be specific, the max adjustments correct for the fact that asymptotically $TP_t \geq BV_t^2$ and $RV_t \geq BV_t$. Finite sampling frequencies do not guarantee that these conditions will always be met. The coefficients in the denominator assure that $z_t \sim N(0,1)$ as $M \rightarrow \infty$ under the null hypothesis that there are no jumps in equity prices. As a result, the interpretation for statistically significant z_t is evidence for at least one jump occurring over the course of the trading day.

The z_t statistic is a good starting point for investigating jumps because it is easy to compute and it has been rigorously validated. Huang and Tauchen (2005) find that the statistic computed on a daily basis does “an outstanding job of identifying the days on which jumps occur.” If we are thus convinced that jumps in stock prices do occur and they account for a non-trivial 4.5% to 7.0% of total price variation, it is logical to delve deeper into researching when jumps occur.

Lee and Mykland (2006) elaborate on this idea by presenting a new jump detection scheme. Using the same principle of comparing realized variance and bipower variation they define a statistic that flags individual returns as statistically significant events under the null hypothesis that there are no jump discontinuities in stock prices. As the sampling frequency is increased to compute returns over smaller time periods the Lee-Mykland statistic presents the opportunity to test when statistically significant jumps occur within the trading day. This differs from the z_t statistic which flags entire trading days as statistically significant events in which a jump may have occurred. The Lee-Mykland statistic will be defined as

$$\ell_{t,j} = \frac{r_{t,j}}{\sqrt{BV_{t,j}}}; \quad (10)$$

$$BV_{t,j} = \frac{1}{k} \left(\frac{\pi}{2} \right) \sum_{i=j-k+1}^j |r_{t,i-1}| |r_{t,i}|.$$

By dividing individual returns by a trailing average of the bipower variation the $\ell_{t,j}$ statistic introduces a test that will flag individual returns as statistically significant jumps.² The notation is

²Note that negative values for the index i will refer to previous days. In particular, $r_{t,0} = r_{t-1,M}$, $r_{t,-1} = r_{t-1,M-1}$, etc.

analogous to the z_t statistic. The subscript t refers to a particular trading day while j refers to the within day return. As before the bipower variation is multiplied by $\frac{\pi}{2}$ so $\ell_{t,j} \sim N(0,1)$. The only difference from the previous definition is that $\ell_{t,j}$ and $BV_{t,j}$ have subscripts for both t and j . This is the case because the $\ell_{t,j}$ statistic is computed for individual returns. The modification to the bipower variation is necessary because the $\ell_{t,j}$ statistic computes a local average for the bipower variation where the variable k determines the degree to which the average of the bipower variation is backward looking. From here on the variable k will be referred to as the window size. Lee and Mykland recommend window sizes of 7 weeks, 16 days, 12 days, 8.5 days, 6 days, and 3.5 days for sampling intervals of 1 week, 1 day, 1 hour, 30 minutes, 15 minutes, and 5 minutes, respectively. Both the choice of window size and the perceived advantage that the $\ell_{t,j}$ statistic can flag individual returns as jumps will be discussed later in the paper.

The next section will briefly digress to discuss the data set. It is important to note that the data is an integral component of the analysis performed later in this paper. The statistical measures defined in the previous paragraphs are intended for high frequency financial data from heavily traded stocks. The objective to study when jump discontinuities occur within the trading day would not be possible without an extensive price series that includes intraday data over a substantial amount of time. This will be accomplished by computing RV_t , BV_t , z_t , and $\ell_{t,j}$ using the high frequency data to be described. Any additional modifications to the statistical measures will be introduced as necessary.

3 Data

High frequency financial data from the New York Stock Exchange will be considered in the subsequent sections of this paper. The focus will be placed on analyzing the Standard and Poor's Depository Receipt or SPDR. The SPDR is an exchange-traded fund that launched in 1993 to track the performance of the S&P 500. It is an extremely liquid instrument that allows individual traders to buy or sell a basket of stocks. From here on the SPDR will serve as our proxy for the market portfolio and it will be referred to by its ticker symbol: SPY. Nine individual stocks will also be considered from the consumer goods, services, basic materials, healthcare, financial, and technology industries. The specific stocks and their ticker symbols are included in Table 1.

The high frequency data was obtained from the Trade and Quote Database (TAQ) via Wharton Research Data Services (WRDS). It was then formatted in Law (2007) into a thirty second price

series using an adapted version of the previous tick method from Dacorogna, Gencay, Müller, Olsen, and Pictet (2001). The final price series excludes non-full trading days and the first five minutes of each trading day in order to ensure uniformity of trading and information arrival.³ The result for each stock is a data set with 771 prices from 9:35am to 4:00pm over 1241 days between January 1, 2001 and December 31, 2005.⁴

The selection of stocks is highly motivated by liquidity. Each stock is among 40 of the most heavily traded stocks on the NYSE as defined by their 10-day trading volume. This is an important criterion for the statistics implemented in Section 4. The primary concern with illiquid stocks and high frequency price sampling is the effect of market microstructure noise. The literature on market microstructure noise dates back to Black (1976) and discusses a variety of sources that bias prices when sampling at a high frequency, including trading mechanisms and discrete prices. Some of these concerns have been mitigated in recent years. During the late 1990s the trading frequency increased significantly at the NYSE and by 2001 almost all of the stock prices were converted from fractional to decimal trading.

Figure 12 in the Technical Appendix includes the volatility signature plots recommended in Andersen, Bollerslev, Diebold, and Labys (2000). Volatility signature plots provide a visual tool for understanding the effect of the sampling frequency on the market-microstructure noise. For the SPY and each of the individual stocks the volatility signature plots indicate the average realized variance in units of annualized volatility when computed at various sampling frequencies.⁵ The expectation is that the average of the realized variance will increase in response to market-microstructure noise as prices are sampled at progressively smaller intervals. This is noted in all of the stocks except MO and LOW. Their behavior is unexpected and for the purposes of this discussion will not be explored further. Rather, the volatility signature plots are used as a means to balance the preference for a fine sampling frequency with the visible effects of market-microstructure noise. Based on these grounds a 17.5 minute sampling frequency is selected and is used in subsequent analysis.

The final consideration that needs to be made with respect to the data is the potential for errors. First, the data is checked for outliers by plotting the price series and daily returns for each stock. No data points are removed by this method. Afterward a simple algorithm is applied to the data to detect for suspicious behavior. In particular, sequential returns are set equal to zero if they are

³Excluding the prices from 9:30 to 9:35am is also notable because Lee and Mykland (2006) report the majority of the statistically significant jumps in their data set during this time period.

⁴A more detailed discussion of how the data was formatted into a thirty second price series can be found in Law (2007).

⁵To annualize the realized variance the computation $\sqrt{\left(\frac{252}{1241}\right) \sum_{t=1}^{1241} RV_t}$ is made where RV_t is calculated at different sampling frequencies for each day in the five year sample.

more than 1.5% in opposite directions over the span of one minute. A likely explanation for this phenomenon is data entry error. Of course, what appears to be irrational behavior is not always the case. One curious example is found in the thirty second price series of Pepsi on April 25th 2005 and is included in Figure 11 in the Technical Appendix. At approximately 10:00am there is a trade that appears to be unreasonably low. It is the lowest price that Pepsi trades at for the entire day and it seems to be out of sync with the other trades around 10:00am. However, further inspection reveals that the volume on the trade is actually 37 times greater than the average volume over the 5 year sample. A rational explanation for the phenomenon is an investor who agrees to sell a significant number of shares at a discounted price to compensate the buyer for taking on the risk of a large order. This example is helpful in that it alludes to the advantage and disadvantage of our data set. The ability to zoom in on financial markets using high frequency data adds to the complexity of research at the same time it allows us to better decipher the on-goings of the real world.

4 Intraday Jump Components

4.1 Patterns in the Lee-Mykland Statistic

One of the underlying goals in using the high frequency data is to better understand how financial markets behave. To address this objective the Lee-Mykland statistic is computed using the data described in the previous section. The analysis focuses on the fundamental advantage of the Lee-Mykland statistic over the statistic proposed by Huang and Tauchen. Namely, the $\ell_{t,j}$ statistic affords the ability to flag individual returns as statistically significant jump discontinuities. This is a natural extension in the literature on non-parametric jump statistics. It provides the opportunity to investigate when large price moves occur within the trading day and detect whether there are patterns in the arrival of jump discontinuities. The intuition is that jumps in the SPY are more likely to occur near the market's open than at other times of the day. This hypothesis is supported by the fact that information flow is high in the morning. Ederington and Lee (1993) suggest that a large volume of information is released after the market closes that is subsequently priced into securities at the start of the subsequent trading day. An additional source of new information can be found in macroeconomic data releases. Chaboud et al. (2004) note that foreign exchange volatility and volume increase significantly for 8:30am data releases such as Nonfarm Payrolls and the Consumer Price Index. Admittedly it seems these announcements would have a more direct impact on over-the-counter products that trade before the equity markets open. Of course, there is

a second round of macroeconomic data releases at 10:00am which occurs after the New York Stock Exchange opens at 9:30am. Furthermore, the observations from Chaboud et al. (2004) underscore the high information flow that traders face in the early hours of the day.

In addition to information flow, the hypothesis that there are more jumps in the morning is motivated by the empirical results of Lee and Mykland (2006). Their paper analyzes high frequency data obtained from the Trade and Quote database for three stocks and the SPY over the period from September 1, 2005 to November 30, 2005 using a five minute sampling frequency. They find that statistically significant jumps in the market portfolio and individual stocks are most frequent in the morning, particularly between 9:30am and 9:35am. Their conclusion is that statistically significant jumps in equity price series often correspond with scheduled announcements. In the case of the market portfolio they observe that flagged jumps tend to coincide with macroeconomic announcements in the morning and the Federal Open Market Committee announcements in the early afternoon. This differs from individual stocks. Jumps in individual stocks tend to correspond with unannounced idiosyncratic news releases that tend to occur early in the morning.

This paper expands upon the results of Lee and Mykland (2006) by computing the $\ell_{t,j}$ statistic over a five year data set for the SPY and nine individual stocks. Trades from 9:30am to 9:35am are excluded from the analysis to ensure uniformity of trading. This consideration is made to account for the findings such as Ederington and Lee (1993) that suggest overnight information may be priced into stocks during the first minutes of trading. If the period from 9:30am to 9:35am was to be included in the analysis it would serve to exaggerate the findings, yielding similar results to those reported in Lee and Mykland (2006). As a result, excluding the first five minutes of trading from the analysis can be viewed as a conservative decision that strengthens the results reported later.

Statistically significant or flagged jumps are defined as $|\ell_{t,j}| > \Phi^{-1}(.999)$. In other words, the null hypothesis that there are no jumps in equity prices is rejected when the absolute value of the statistic $\ell_{t,j} \sim N(0,1)$ is greater than $\Phi^{-1}(.999) \approx 3.09$. The number of flagged jumps are then counted over the course of the five year sample for each intraday period. In particular, *Jump* is defined below

$$Jump(j, k) = \sum_{t=1}^{Days} I[|\ell_{t,j}| \geq \Phi^{-1}(.999)], \quad (11)$$

$$j \in \{1, 2, \dots, 22\}, \quad k \in \{1, 2, \dots, 150\}. \quad (12)$$

Jump is a function of the window size k and the intraday interval j that reports the number of flagged

jumps during each intraday period. $I[\cdot]$ is the indicator function. Figure 1 reports the results for the SPY data set. The independent axes correspond to the window size and the intraday interval. The window size takes on integer values in the set $\{1, 2, \dots, 150\}$ which permits the trailing average of the bipower variation to be as far as 6.7 days backward looking. From the plot it is clear that the statistic stabilizes for a sufficiently high window size. The plot also highlights the total number of flagged jumps between particular times during the trading day. At the 17.5 minute sampling frequency there are 22 intraday returns. As an example, when $j = 1$ the variable $Jump(1, k)$ reports the number of statistically significant returns between 9:35am and 9:52:30am for each day in the five year period computed using the window size k . The intraday pattern is striking. Across all window sizes the $\ell_{t,j}$ statistic flags more returns in the morning than later parts of the day. For a window size of 4.5 days approximately 41% of the total flagged jumps for the SPY occur before 11:00am. This period includes the 10:00am macroeconomic announcements including Consumer Confidence, Factory Orders, the ISM Index, and New Existing Home Sales. The number of flagged jumps also begins to increase later in the trading day, a potential response to the Federal Open Market Committee statements that are released at approximately 2:15pm. Figure 13 in the Technical Appendix plots the same data but collapses the curves corresponding to different window sizes onto two axis. From this plot it is easy to see the pattern in flagged jumps converges to a “smile” or “smirk.” As the window size increases the local average of volatility becomes progressively more backward looking and the number of flagged jumps stabilizes.

Figure 16 plots the number of flagged jumps computed using a window size of 4.5 days for each of the individual stocks. The first observation is that all of the individual stocks exhibit a similar pattern to the SPY. The vast majority of flagged idiosyncratic jumps occur before 11:00am. The results are distinct from the SPY in that the number of flagged idiosyncratic jumps does not noticeably increase during the afternoon. As an example, 67% of the flagged jumps for Lowes Companies occur during the first 1.5 hours of trading while only 10% of the flagged jumps occur during the last 1.5 hours of trading. The specific number of flagged jumps over each intraday period is reported in Table 2 using a 4.5 day window size for the nine individual stocks and the SPY. More empirical analysis is included in the Technical Appendix. Figures 14 and 15 show the Quantile-Quantile plots and a non-parametric kernel density estimation for the probability density function of $\ell_{t,j}$ computed using the SPY data across various window sizes.

4.2 Patterns in Realized Variance and Bipower Variation

In this section the focus shifts from investigating patterns in the arrival of flagged jump discontinuities by the $\ell_{t,j}$ statistic to analyzing whether intraday patterns persist in the jump and diffusive components of equity volatility. As a brief aside, it is important to note that the literature on intraday patterns in volatility dates back to a paper from Wood, McInish, and Ord (1985) documenting U -shaped patterns in both intraday volatility and volume for equities. There is also an extensive literature that investigates intraday patterns in foreign exchange volatility and volume, including studies such as Chaboud et al. (2004) and Andersen and Bollerslev (1998). They document increases in volatility and volume around data releases at 8:30am. This paper expands on the literature by testing for intraday patterns in the decomposition of equity volatility into its diffusive and jump components. To the author's knowledge, this is the first attempt to test for intraday patterns in the decomposition of equity volatility. The intuition is that the diffusive and jump components will follow a U -shape. This hypothesis is supported by the empirical results from computing the $\ell_{t,j}$ statistic in the previous section. Further confirmation is found by rigorously testing the realized variance and bipower variation introduced earlier. Ultimately, the objective in separating the pattern between the diffusive and jump components in equity volatility is to better understand the evolution of stock prices.

To begin the analysis the realized variance and bipower variation are redefined from their previous definition in Section 2.2. In particular, both measures are computed over seven intraday periods instead of using the entire trading day. Each period within the day is approximately 52.5 minutes and is defined in the following manner⁶

$$RV_{t,i}^* = \sum_{j=3i-2}^{3i} r_{t,j}^2; \quad (13)$$

$$BV_{t,i}^* = \left(\frac{\pi}{2}\right) \sum_{j=3i-1}^{3i} |r_{t,j-1}| |r_{t,j}|. \quad (14)$$

The decision to consider seven intraday intervals implies that $i \in \{1, 2, \dots, 7\}$. The jump component is then defined as the difference between the realized variance and bipower variation or $J_{t,i}^* = RV_{t,i}^* - BV_{t,i}^*$. Empirical results are recorded in a series of Tables and Figures in the back. Figure 2 provides an interesting series of plots to begin the discussion. It compares weekly averages of the

⁶Because a 17.5 minute sampling frequency results in computing 22 returns per day we do not have each data point correspond to exactly 52.5 minutes. In particular, the last data point corresponds to the last 70 minutes of each trading day. It is defined as $RV_{t,7}^* = \frac{3}{4} \sum_{j=19}^{22} r_{t,j}^2$ and $BV_{t,7}^* = \frac{3}{4} \sum_{j=20}^{22} |r_{t,j-1}| |r_{t,j}|$. As a result, we scale each term appropriately so it can be compared to the other intraday periods. We also report all computations in terms of annualized volatility for ease of comparison.

volatility, the diffusive component of volatility and the jump component of volatility during two intraday periods for the SPY. One can note from the three plots that all measures of volatility tend to be higher between 9:35am and 10:27am than 12:13pm and 13:05pm. Figure 17 in the Technical Appendix reports the same plots for Citigroup where a similar conclusion can be reached. The results indicate that the total level of volatility, the diffusive component of volatility and the jump component of volatility are higher in the morning than the middle of the day.

Table 2 reports the averages of the intraday realized variance, bipower variation and jump component across the sample. The numbers are annualized and reported in terms of standard deviation. For example, the value of 20.49 from 9:35am to 10:27am for the SPY intraday realized variance can be interpreted as 20.49% annualized volatility between 9:35am and 10:27am. The data for the SPY is plotted in Figure 3. The results indicate that the average of the realized variance, bipower variation and jump component for the SPY are all higher in the morning than the middle of the day. To be specific, the average for the realized variance between 9:35am and 10:27am is 20.49% annually whereas the the average for realized variance between 12:30pm and 1:05pm is 11.81%. This breaks down into 17.00% and 3.49% for the bipower variation and jump component between 9:35am and 10:27am compared to 9.65% and 2.16% for the bipower variation and jump component between 12:30pm and 1:05pm. This difference is certainly significant. It constitutes over a 33% decrease in the realized variance from morning to midday. Figure 3 also indicates that the realized variance and bipower variation exhibit what are commonly referred to as “smiles” or “smirks” in the literature. They tend to have their highest values in the morning, lowest values in the middle of the day, and intermediate values near the close. The jump component is distinct in that it falls during the last hour of trading and it appears to be even more skewed toward the start of the trading day. In fact, the drop in the average jump component of volatility between 9:35am and 10:27am to 12:30pm and 1:05pm is 38%. Figure 18 plots the data for some of the individual stocks. One can observe virtually identical patterns to those noted in the market portfolio. For all stocks the realized variance, bipower variation and jump component take on their highest value in the first hour of the trading day. With the exception of the last hour of trading where the jump component tends to decrease, it generally constitutes about 20% to 25% of the realized variance as measured by $J_{t,i}^*/RV_{t,i}^*$.

These results confirm the intuition in regard to the diffusive and jump component of equity volatility. By looking at intraday averages across our sample in a similar manner to the Wood, McInish, and Ord (1985) it is clear that the diffusive component of volatility exhibits behavior that resembles the previously documented *U*-shaped patterns in equity volatility. The jump component is distinct from the realized variance and bipower variation in that it does not increase at the end

of the trading day. Rather, it is highest at the start of the day and generally decreases during the last hour of trading. The next section investigates the implications of these patterns for the $\ell_{t,j}$ statistic. It finds that the U -shaped pattern in the diffusive component of volatility has an impact on the ability of $\ell_{t,j}$ statistic to properly identify jumps.

4.3 Implications for the Lee-Mykland Statistic

The previous section suggests that patterns seem to persist in the bipower variation and jump component of equity returns within the trading day. This has important implications for the $\ell_{t,j}$ statistic. Namely, the $\ell_{t,j}$ statistic defined in Equation 18 divides individual returns by a trailing average of the bipower variation. As the window size for computing the average bipower variation is increased the intraday patterns in local volatility are eliminated. A potential result is that individual returns in the morning may be compared to an average of the bipower variation that is actually lower than the true bipower variation during the morning, with the consequence of overstating the number of flagged jumps in the morning and afternoon and understating the number of flagged jumps in the middle of the day. To pursue this idea further a theoretical example is considered and later compared to empirical results.

The theoretical example is derived from Figure 3. The primary concern for the $\ell_{t,j}$ statistic is whether the bipower variation included in the denominator will behave properly. To that end, a simple scenario is investigated. In a situation without stochastic volatility one might argue that the bipower variation behaves as a periodic function over the trading day. When taking the average over sufficiently large window sizes one would eliminate all of the periodicity seen in true bipower variation potentially leading to Type I errors in the morning and afternoon and Type II errors in the middle of the day. To explore this idea from a theoretical standpoint Figure 4 plots the average bipower variation and a polynomial fit to the data. In the context of the theoretical example the polynomial represents the true value for the bipower variation,

$$True_{t,j} = -.000171 \cdot j^4 + .00243 \cdot j^3 - .00456 \cdot j^2 - .0381 \cdot j + .0211. \quad (15)$$

As usual t refers to the day and j refers to 1 of 22 intraday periods. The next step is to compute the lagged values for the bipower variation over various window sizes. They are calculated in the following manner

$$Lag_{t,j} = \frac{1}{n} \sum_{i=j-n+1}^j True_{t,i}. \quad (16)$$

Figure 5 plots the lagged values for window sizes of 1 hour, 3 hours, 4.5 hours, and 1 day. In this particular case it is clear that any intraday patterns are eliminated after the window sizes reaches 1 day. Of course, the window size used for calculating $\ell_{t,j}$ in the previous section is 4.5 days. If the theoretical example being considered is representative of the sample period than one might expect to find some of the errors alluded to above.

One way to investigate the degree to which the $\ell_{t,j}$ statistic misidentifies jumps in the real data is to compare how the original statistic compares to a new statistic where the bipower variation in the denominator is centered at the current time. In the context of our theoretical example this means computing lagged values defined as

$$Lag_{t,j}^c = \frac{1}{n+1} \sum_{i=j-\frac{n}{2}}^{j+\frac{n}{2}} True_{t,i}. \quad (17)$$

⁷ Figure 6 notes that the newly defined $Lag_{t,j}^c$ does a better job of approximating the true value for low window sizes. For example, at the 4.5 hour window size it is clear that $Lag_{t,j} < Lag_{t,j}^c$ in the morning whereas $Lag_{t,j} > Lag_{t,j}^c$ in the middle of the day. This can be seen in Figure 7 which plots the difference $Lag_{t,j} - Lag_{t,j}^c$. Another note from Figure 6 is that $Lag_{t,j}^c$ and $Lag_{t,j}$ approach the average of the periodic function $True_{t,j}$ as the window size is increased to a sufficiently high level. This concept is applied to the real data by re-defining the Lee-Mykland statistic such that the original statistic can be compared to a new statistic where the average of local volatility is centered at the current time. Formally the new statistic $\ell_{t,j}^c$ is defined as

$$\ell_{t,j}^c = \frac{r_{t,j}}{\sqrt{BV_{t,j}^c}}; \quad (18)$$

$$BV_{t,j}^c = \frac{1}{K+1} \left(\frac{\pi}{2}\right) \sum_{i=j-K/2}^{j+K/2} |r_{t,i-1}| |r_{t,i}|.$$

Both $\ell_{t,j}$ and $\ell_{t,j}^c$ are computed for the SPY and PEP data. The difference in the number of flagged jump discontinuities reveals an interesting difference between the two statistics.

Figures 8 plots the data for the SPY while Figure 19 in the Technical Appendix plots the data for

⁷When n is odd $Lag_{t,j}^c$ is defined as $Lag_{t,j}^c = \frac{1}{n} \sum_{i=j-\frac{n}{2}+\frac{1}{2}}^{j+\frac{n}{2}-\frac{1}{2}} True_{t,i}$.

PEP. Table 3 reports the number of flagged jumps for the SPY data at the 3 and 4.5 hour window sizes. To begin analyzing the data consider the start of the trading day. At small window sizes like 4.5 hours the theoretical example in Figure 7 suggests that the average of the bipower variation centered at the current time will provide a better approximation to the bipower variation than the backward looking average. This results in $BV_{t,j} < BV_{t,j}^c$ for small window sizes and suggests that $\ell_{t,j}$ would flag more jumps than $\ell_{t,j}^c$. This is observed in the data. Table 3 shows that $\ell_{t,j}$ statistic flags more large price moves than the $\ell_{t,j}^c$ statistic in the early hours of trading when the window size is 3 or 4.5 hours. This is particularly notable for the flagged jumps computed between 9:52 and 10:10am. During this period the $\ell_{t,j}$ statistic flags 46 jumps compared to the $\ell_{t,j}^c$ statistic which flags 24 jumps at a 3 hour window size using the SPY data. As the window size is increased beyond one day the number of statistically significant jumps flagged by $\ell_{t,j}$ and $\ell_{t,j}^c$ tends to converge. This is what one would expect from the theoretical example. A sufficiently high window size will eliminate intraday patterns in the local volatility in both statistics. This suggests both statistics will flag the same number of jumps. The minor differences that are observed in the number of flagged jumps by $\ell_{t,j}^c$ and $\ell_{t,j}$ can be explained by stochastic volatility. The $\ell_{t,j}^c$ statistic is both trailing and forward looking so it will pick up changes in the level of realized variance earlier than the $\ell_{t,j}$ statistic, providing an explanation for the small discrepancies in the number of statistics they flag at high window sizes like 6.7 days. During the early hours of trading the conclusion is that intraday patterns in the bipower variation may result in Type I errors when using the $\ell_{t,j}$ statistic.

The expected effect of intraday patterns in the bipower variation is also observed in the middle and end of the trading day. For example, at a 3 hour window size one would expect that the trailing average of the bipower variation to over estimate the average of the bipower variation centered at the current time, $BV_{t,j} > BV_{t,j}^c$, resulting in the $\ell_{t,j}^c$ flagging more statistically significant jumps than the $\ell_{t,j}$ statistic. At the end of the day one would expect the relationship to reverse such that the $\ell_{t,j}$ statistic flags more statistically significant jumps than the $\ell_{t,j}^c$ statistic. This is observed in the SPY data. Table 2 confirms the intuition that the statistic $\ell_{t,j}^c$ will flag more jumps in the middle of the day and less jumps at the end of the day. The additional observation can be made that the difference between $\ell_{t,j}$ and $\ell_{t,j}^c$ is more pronounced in the morning and the late afternoon. Continuing with the SPY data at a 3 hour window size as the example, one notes that between 9:52am and 10:10am the $\ell_{t,j}$ statistic flags 22 more jumps than the $\ell_{t,j}^c$ statistic, between 12:47 and 13:05pm the $\ell_{t,j}^c$ statistic flag 2 more jumps than the $\ell_{t,j}$ statistic and between 15:42 and 16:00 the $\ell_{t,j}$ statistic flag 28 more jumps than the $\ell_{t,j}^c$ statistic.

A plausible explanation for why intraday patterns in the bipower variation have a more visible

effect on the number of flagged jumps in the morning and late afternoon is that large price movements seem to be more frequent in the morning and late afternoon. As noted in the previous section, both the diffusive component and jump component of equity volatility are lowest in the middle of the trading day. This may indicate that there aren't as many large price movements for the $\ell_{t,j}^c$ statistic to flag as significant during the middle of the day. It is also important to note that this observation does not contradict the previous results. In fact, it actually strengthens the hypothesis that large price movements are not as frequent in the middle of the trading day. Regardless of how the bipower variation is defined there are only a few returns flagged as statistically significant jump discontinuities in the middle of the trading day. The additional conclusion from this section is that the number of jumps flagged by the $\ell_{t,j}$ statistic may be spurious. While it is clear from the results in Section 4.1 that flagged jumps are more frequently in the morning and late afternoon, it appears that taking a trailing average of the bipower variation will lead to Type I errors that exaggerate the number of statistically significant jumps in the morning and late afternoon and Type II errors that result in understating the number of statistically significant jumps in the middle of the day.

5 Forecasting

After observing patterns in both intraday volatility and the $\ell_{t,j}$ statistic a natural question to ask is whether the empirical results are useful in forecasting the future volatility of equity returns. It is feasible that informational content from the patterns observed in high frequency financial data could improve the efficiency and performance of historical models from the financial econometrics literature. This is an important question to resolve regardless of the outcome. In addition to the research interest in the academic community, forecasting and risk management have become increasingly important for investors and major corporations as evidenced by the growing demand for derivative products based on equity volatility.

Recent papers from Müller et al. (1997) and Corsi (2003) develop heterogeneous autoregressive realized variance (HAR-RV) models that are particularly relevant. These models use exclusively historical data that would be readily available to a trader or an academic researcher, for example. They signify the expectation of future volatility for a representative agent who ignores all information except for the historical data used to compute the forecast. Empirical tests performed by Andersen, Bollerslev, Diebold, and Labys (2003) and Andersen, Bollerslev, and Meddahi (2004) use the HAR-RV class of volatility models to show that simple linear forecasts can outperform the more

complicated GARCH and stochastic volatility models that incorporate long-memory dependence of financial markets.

HAR models are included in this paper to forecast the average future value of realized variance over different horizons using lagged averages of realized volatility. The notation is borrowed from Andersen, Bollerslev, and Diebold (2006). The average future value of realized variance, what ABD (2006) refers to as the multi-period normalized realized variation, is defined as

$$RV_{t,t+h} = h^{-1}[RV_{t+1} + RV_{t+2} + \dots + RV_{t+h}] \quad (19)$$

where $h = 1$, $h = 5$ and $h = 22$ correspond to averages over daily, weekly and monthly time intervals. The HAR-RV model introduced in Corsi (2003) to forecast daily realized variance can then be denoted as

$$RV_{t,t+1} = \beta_o + \beta_d RV_{t-1,t} + \beta_w RV_{t-5,t} + \beta_m RV_{t-22,t} \quad (20)$$

where the horizons for the lagged explanatory variables are chosen for their intuitive appeal and natural economic interpretation. ABD (2006) extends this model to forecast the realized variance over longer horizons and presents one of the first attempts to decompose the realized variance into its diffusive and jump components for forecasting purposes. The manner they divide the realized variance into its diffusive and jump component relates to this paper because the non-parametric statistic they implement is the z_t statistic introduced in Section 2 and the bi-power variation investigated in Section 4. The question at hand is whether the performance of their regression can be enhanced by incorporating the observed intraday patterns in volatility into the proposed model for dividing the realized variance into its continuous sample path variability and jump variation.

ABD (2006) introduces the heterogeneous autoregressive realized variance continuous jump model, hereafter referred to as HAR-RV-CJ, to decompose the realized variance into its continuous and sample path variability. The model defines $RV_{t,\alpha}$ as

$$RV_{t,\alpha} = C_{t,\alpha} + J_{t,\alpha}, \quad (21)$$

$$C_{t,\alpha} = I[z_t \leq \Phi^{-1}(\alpha)] \cdot [RV_t] + I[z_t > \Phi^{-1}(\alpha)] \cdot [BV_t], \quad (22)$$

$$J_{t,\alpha} = I[z_t > \Phi^{-1}(\alpha)] \cdot [RV_t - BV_t]. \quad (23)$$

This is a natural separation of the realized variance using the indicator function $I[\cdot]$. Note that RV_t and BV_t are the same non-parametric statistics that are introduced in Section 2 and later studied in Section 4. The first component, $C_{t,\alpha}$, corresponds to the continuous sample path variability which is computed as RV_t on any trading day where the z_t statistic from Huang and Tauchen (2005) does not reject the null hypothesis of no jumps at the $1 - \alpha$ significance level. On days that are flagged as statistically significant jumps the continuous component $C_{t,\alpha}$ corresponds to the bipower variation BV_t and the jump component $J_{t,\alpha}$ is defined as the difference between realized variance and bi-power variation $RV_{t,\alpha} - BV_{t,\alpha}$. The value of α is set equal to .999 as in ABD (2006) and is omitted in the subsequent references to C_t and J_t .

The explanatory variables in HAR-RV-CJ model are introduced in an analogous manner to the HAR model. Over different horizons the continuous and jump sample path variability are defined as

$$C_{t,t+h} = h^{-1}[C_{t+1} + C_{t+2} + \dots + C_{t+h}] \quad (24)$$

and

$$J_{t,t+h} = h^{-1}[J_{t+1} + J_{t+2} + \dots + J_{t+h}], \quad (25)$$

resulting in the definition of the HAR-RV-CJ model as

$$RV_{t,t+h} = \beta_o + \beta_{CD}C_{t-1,t} + \beta_{CW}C_{t-5,t} + \beta_{CM}C_{t-22,t} + \beta_{JD}J_{t-1,t} + \beta_{JW}J_{t-5,t} + \beta_{JM}J_{t-22,t} + \epsilon_{t,t+h}. \quad (26)$$

A consideration needs to be made for the error term which will generally be serially correlated at least to the order $h - 1$. While this will not affect the consistency of regression coefficient estimates, the estimates for standard errors are adjusted by using the Newey-West heteroskedasticity covariance matrix estimator with a lag of 60 days. The results from the regression are reported in Tables 5 and 6. Coefficient estimates for the explanatory jump variables are negative in all the statistically significant cases. One can interpret this as signifying that flagged jumps in the z_t statistic forecast lower levels of future volatility over daily, weekly and monthly horizons for the S&P 500.

To incorporate the intraday patterns in the jump component of equity volatility into the HAR-RV-CJ model a modification is introduced for the definition of the explanatory jump variable. The modification is fairly analogous to J_t . An indicator function defines the jump component to be zero on days that are not flagged as statistically significant by the z_t statistic. On days that are flagged

as statistically significant the jump component is the difference in the realized variance and the bipower variation during the first hour of trading. In other words,

$$J_{t,\alpha}^1 = I[z_t > \Phi^{-1}(\alpha)] \cdot [RV_{t,1}^* - BV_{t,1}^*]. \quad (27)$$

Results are reported in Table 5. The first observation is that none of the coefficient estimates for the newly defined explanatory variable J_t^1 are statistically significant from zero. This contrasts from using J_t where several of the coefficient estimates for the explanatory jump variables are negative and statistically significant. However, the additional observation can be made that the performance of the regressions as measured by R^2 and $adj - R^2$ is practically the same both cases. It is notable that the performance is highest using J_t as an explanatory variable. This signifies that the informational content relevant for forecasting the average future value of realized volatility over different horizons is higher in the explanatory variable J_t than in the variable J_t^1 . However, the difference in informational content between these variables is marginal and not particularly important from the perspective of a financial practitioner or forecaster.

6 Conclusion

The empirical results in this paper expand on the non-parametric literature in financial econometrics. The hypothesis is that large moves in stock prices occur most frequently in the morning. This conjecture is tested by using high frequency data from nine individual stocks and the SPY to investigate whether there are patterns in jump discontinuities and volatility throughout the trading day.

Lee and Mykland (2006) introduce a new statistic that flags individual returns as statistically significant jump discontinuities under the null hypothesis that there are no jumps in asset prices. They report that jumps are most frequent in the first hour of trading using a three month sample of high frequency data for the SPY. This paper expands on their result by using a five year sample of high frequency data. It is also the most extensive attempt to apply the Lee-Mykland statistic to individual stocks and the first paper to analyze the number of statistically significant jumps across different window sizes of the Lee-Mykland statistic. The results are included in Section 4.1. They confirm the intuition from Lee and Mykland (2006). Figure 1 shows that the statistic

stabilizes for a sufficiently high window size and Table 2 indicates that statistically significant jumps in both individual stocks and the market portfolio are most frequent in the morning. This pattern is particularly apparent in the idiosyncratic jumps of individual stocks, with a large number of the flagged jumps occurring at approximately 10:00am.

The discussion that follows is an investigation of intraday patterns in volatility. In one sense this is an extension of the previous work. The Lee-Mykland statistic is based on the theoretical results from Barndorff-Nielsen and Shephard (2004). It implicitly makes a comparison between the realized variance and the bi-power variation in determining which returns are statistically significant jumps. Wood, McInish and Ord (1985) document that equity volatility exhibits a *U*-shaped intraday pattern. Section 4.2 expands on their result by testing for similar patterns in the decomposition of equity volatility into its jump and diffusive components. The realized variance and the bi-power variation are computed over intraday intervals and averaged over the sample as in Wood et al. (1985). The empirical results indicate that both equity volatility and the diffusive component in volatility exhibit *U*-shaped patterns within the trading day. The jump component in volatility is highest in the morning but it does not increase later in the day. Evidence for this can be found in Table 3, Figure 3 and Figure 16 in the Technical Appendix. Similar results seem to hold for the individual stocks and the market.

The subsequent section investigates the affect of intraday patterns in volatility on the Lee-Mykland statistic. It compares the Lee-Mykland statistic which has a backward looking average of local volatility with a modification to the Lee-Mykland statistic that defines the average of the local volatility to be centered at the current time. The conclusion from Figure 8 and Table 4 is that the Lee and Mykland statistic may over specify the number of statistically significant jumps in the morning and late afternoon while under reporting the number of jumps in the middle of the trading day. The intuition comes from Figures 5 and 7. It seems that the patterns observed in the local volatility are eliminated at the recommended window sizes for computing the statistic. The result is that returns in the morning and late afternoon are compared to an average of the diffusive component in volatility that may be lower than the true diffusive component in volatility, whereas returns in the middle of the trading day are compared an average of the diffusive component in volatility that may be higher than the true diffusive component in volatility.

The last section investigates whether the observed patterns in the realized variance and bi-power variation can be used to improve the performance of the HAR-RV-CJ forecasting model proposed in Andersen, Bollerslev and Diebold (2006). The results reported in Table 5 are similar to those

in Andersen, Bollerslev and Diebold (2006). They indicate that only the continuous component of the realized variance has predictive power in forecasting. The coefficient estimates for the jump explanatory variables are generally insignificant using Newey-West standard errors. In fact, over all horizons the jump explanatory variables based on the first hour of trading have insignificant coefficient estimates.

These results suggest several interesting extensions for future research. First, the discussion related to the Lee-Mykland statistic could be enhanced by performing Monte Carlo simulations to compare the performance of the statistic when there are intraday patterns in the diffusive component of volatility with a simulation where there are no intraday patterns in the diffusive component of volatility. This would be most convincing if performed in a model where the volatility was stochastic in addition to exhibiting intraday patterns.

Another extension could be a new definition for the Lee-Mykland statistic that attempts to take into consideration the patterns in intraday volatility. As noted previously, the Lee-Mykland statistic has an advantage over the other non-parametric jump detection schemes in that it tests individual returns under the null hypothesis that there are no jumps in asset prices. If the estimate of local volatility in the denominator could be corrected to take into account the effects of intraday volatility it would provide a very powerful tool for identifying when statistically significant jumps occurred in historical data.

Finally, it would be interesting to look further into what drives jumps in asset prices. The empirical results suggest that there are more jumps in the morning than any other time of day. The morning is also the most volatile period of the trading day. One interpretation of this result is that the morning may serve as a period of price discovery. From this perspective the morning may be viewed as an opportunity for traders to interact with each other to decide the price of assets after they encode new information into their expectations about the future. This idea could be extended to a comparative analysis of volatility and jumps in the morning for the SPY versus the individual stocks. There is evidence alluded to throughout the paper that indicates something different may be driving the market and the individual stocks to jump. Explaining this difference in behavior suggests an interesting source of future research.

References

- [1] Amihud, Y. and H. Mendelson (1987). Trading mechanisms and stock returns: An empirical investigation. *Journal of Finance*, 42, 533-553.
- [2] Andersen, T.G., and T. Bollerslev (1998). Answering the Skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39, 885-905.
- [3] Andersen, T.G., T. Bollerslev and F.X. Diebold (2000). Great Realizations. *Risk Magazine*, 13, 105-108.
- [4] Andersen, T.G., T. Bollerslev, and F.X. Diebold (2006). Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility. Working Paper, National Bureau of Economic Research.
- [5] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2001). The distribution of realized exchange rate volatility. *Journal of American statistical Association*, 96, 42-55.
- [6] Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, 71, 579-626.
- [7] Andersen, T.G., T. Bollerslev, and N. Meddahi (2004). Analytic Valuation of Volatility Forecasts. *International Economic Review*, 45, 1079-1110.
- [8] Barndorff-Nielsen, O.E., and N. Shephard (2002a). Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models. *Journal of Royal Statistical Society, B* 63, 253-280.
- [9] Barndorff-Nielsen, O.E. and N. Shephard (2002b). Estimating Quadratic Variation Using Realised Variance. *Journal of Applied Econometrics*, 17, 457-477.
- [10] Barndorff-Nielsen, O.E., and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics*, 2, 1-37.
- [11] Barndorff-Nielsen, O.E., and N. Shephard (2006). Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation. *Journal of Financial Econometrics*, 4, 1-30.
- [12] Bollerslev, T. (1986). Generalized autoregressive heteroskedasticity. *Journal of Econometrics*, 31, 307-327.
- [13] Basawa, I. and P. Brockwell (1982). Non-parametric estimation for non-decreasing lévy processes. *Journal of the Royal Statistical Society Series, B* 44, 262-269.
- [14] Black, F. and M. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, 637-654.
- [15] Black, F. (1986). Noise. *Journal of Finance*, 41, 529-543.
- [16] Chaboud, A., S. Chernenko, H. Edward, R.S.K. Iyer., D. Liu and J.H. Wright (2004). The High-Frequency Effects of U.S. Macroeconomic Data Releases on Prices and Trading Activity in the Global Interdealer Foreign Exchange Market. *FRB International Finance Discussion Paper No. 823*.
- [17] Corsi, F. (2004). A Simple Long Memory Model of Realized Volatility. *Manuscript*, University of Southern Switzerland.
- [18] Dacorogna, M., R. Gencay, U. Müller, R. Olsen, and O. Pitchet (2001). *An Introduction to High Frequency Finance*. San Diego: Academic Press.
- [19] Ederington, L.H. and J.H. Lee (1993). How Markets Process Information: New Releases and Volatility. *Journal of Finance*, 48, 1161-1191.

- [20] Engle, R.F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of Variance of United Kingdom Inflation. *Econometrica*, 50, 987-1008.
- [21] Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance*, 59, 1367-1403.
- [22] Eraker, B., M. Johannes, and N. Polson (2003). The Impact of Jumps in Volatility and Returns. *Journal of Finance*, 58, 1269-1300.
- [23] Harris, L. (1990). Estimation of stock variance and serial covariance from discrete observations. *Journal of Financial and Quantitative Analysis*, 25, 291-306.
- [24] Harris, L. (1991). Stock price clustering and discreteness. *Review of Financial Studies*, 4, 389-415.
- [25] Huang, X., and G. Tauchen (2005). The Relative Contributions of Jumps to Total Variance. *Journal of Financial Econometrics*, 3, 456-499.
- [26] Law, T.H. (2007). The Elusiveness of Systematic Jumps. *Duke Journal of Economics*, 19, Honors Thesis.
- [27] Lee, S. S., and P. Mykland (2006). Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics. Working Paper. Georgia Institute of Technology and University of Chicago.
- [28] Merton, R.C. (1976). The Impact on Option Pricing of Specification Error in the Underlying Stock Price Returns. *Journal of Finance*, 31, 333-350.
- [29] Merton, R.C. (1973). Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science*, 4, 141-183.
- [30] Müller, U., M. Dacorogna, R. Dav, R. Olsen, and J. von Weizsacker (1997). Volatility of different time resolutions - analysing the dynamics of market components. *Journal of Empirical Finance*, 4, 213-239.
- [31] Shiller, R. (2005). *Irrational exuberance*. New York: Currency.
- [32] Taleb, N. (2007). *The black swan: the impact of the highly improbable*. New York: Random House.
- [33] Wood, R. A., T. H. McInish, and J. K. Ord (1985). An investigation of transaction data for NYSE stocks. *Journal of Finance*, 25, 723-739.

7 Tables

Stock Names and Ticker Symbols

Ticker Symbol	Actual Name
SPY	Standard and Poor's Depository Receipt
C	Citigroup
LOW	Lowe's
MO	Altria Group
MRK	Merck
PEP	Pepsi
UPS	United Parcel Service
VZ	Verizon Communications
WMT	Wal-Mart Stores
XOM	Exxon Mobile

Table 1: Reports the ticker symbols for the nine individual stocks and the market portfolio considered throughout the paper.

Statistically Significant Jump Discontinuities During Intraday Periods

Time at NYSE	SPY	C	LOW	MO	MRK	PEP	UPS	VZ	WMT	XOM
9:52:30	14	9	3	17	15	8	26	10	8	9
10:10	27	79	137	117	121	104	107	110	101	97
10:27:30	42	55	89	70	55	55	61	68	71	46
10:45	21	35	34	35	29	33	40	25	24	16
11:02:30	14	20	22	32	23	22	17	18	16	22
11:20	7	13	19	19	12	14	23	14	17	12
11:37:30	12	16	9	14	7	16	13	13	10	14
11:55	8	5	8	18	11	7	5	9	2	6
12:12:30	5	2	6	6	7	6	9	13	4	5
12:30	4	0	6	8	11	2	11	6	2	7
12:47:30	4	2	2	5	4	7	7	2	5	0
1:05	3	1	3	11	4	5	6	1	2	3
1:22:30	2	3	2	5	2	2	1	3	1	4
1:40	4	4	3	5	3	3	4	2	2	0
1:57:30	5	2	5	5	4	3	2	2	5	3
2:15	4	1	2	6	8	6	5	6	2	2
2:32:30	14	7	6	9	8	3	8	8	6	3
2:50	19	15	2	14	10	5	9	6	8	7
3:07:30	10	6	8	6	5	10	11	6	9	14
3:25	22	13	9	8	9	4	16	19	4	12
3:42:30	13	5	8	11	12	10	7	10	3	8
4:00	28	7	7	17	11	21	17	15	10	15
Total	282	300	390	438	371	346	485	366	312	305

Table 2: Returns are flagged as statistically significant jumps when $\ell_{t,j} > \Phi^{-1}(.999)$ using a 4.5 day window size. The first column corresponds to the end of each intraday period using a 17.5 minute sampling frequency throughout the trading day. The data is plotted in Figure 16.

Intraday Realized Variance $RV_{t,i}^$*

Intraday Period i	SPY	C	LOW	MO	MRK	PEP	UPS	VZ	WMT	XOM
9:35 to 10:27	20.49	36.21	42.72	31.84	33.77	28.71	23.00	35.50	33.81	28.89
10:27 to 11:20	16.05	26.20	30.04	23.01	23.88	21.21	16.05	26.59	24.49	21.74
11:20 to 12:12	13.26	22.35	24.62	19.58	20.71	17.19	13.16	22.37	19.27	17.84
12:12 to 13:05	11.81	19.65	22.00	17.00	17.97	15.44	11.92	19.28	17.70	16.02
13:05 to 13:57	12.80	20.38	22.53	16.73	18.40	15.90	12.94	20.24	19.02	16.14
13:57 to 14:50	14.73	23.04	24.43	18.73	20.12	17.22	13.85	22.13	20.72	18.47
14:50 to 16:00	16.68	25.39	24.89	20.42	23.12	19.92	15.60	25.15	22.90	21.67

Intraday Diffusive Component $BV_{t,i}^$*

Intraday Period i	SPY	C	LOW	MO	MRK	PEP	UPS	VZ	WMT	XOM
9:35 to 10:27	17.00	29.61	34.67	25.66	27.38	23.16	18.07	28.44	27.43	23.15
10:27 to 11:20	13.29	20.89	24.56	18.33	19.35	17.22	12.93	21.89	19.81	17.39
11:20 to 12:12	10.68	18.31	19.58	15.68	16.48	13.70	10.30	18.01	15.67	14.40
12:12 to 13:05	9.65	15.86	17.89	13.60	14.55	12.29	9.65	15.68	14.13	13.24
13:05 to 13:57	10.31	16.61	17.76	13.22	14.96	12.80	10.69	16.20	15.73	13.09
13:57 to 14:50	12.03	18.51	19.52	15.07	16.36	13.80	11.04	17.99	16.63	15.20
14:50 to 16:00	14.26	21.84	21.51	17.52	19.60	17.18	12.94	21.75	19.99	18.68

Intraday Jump Component $RV_{t,i}^ - BV_{t,i}^*$*

Intraday Period i	SPY	C	LOW	MO	MRK	PEP	UPS	VZ	WMT	XOM
9:35 to 10:27	3.49	6.61	8.05	6.18	6.39	5.54	4.93	7.06	6.37	5.75
10:27 to 11:20	2.76	5.32	5.48	4.68	4.53	3.99	3.12	4.70	4.67	4.35
11:20 to 12:12	2.58	4.04	5.04	3.90	4.23	3.49	2.86	4.36	3.60	3.43
12:12 to 13:05	2.16	3.79	4.11	3.40	3.42	3.15	2.27	3.59	3.56	2.79
13:05 to 13:57	2.49	4.77	4.76	3.51	3.44	3.10	2.25	4.04	3.28	3.05
13:57 to 14:50	2.70	4.53	4.91	3.66	3.77	3.41	2.81	4.14	4.09	3.27
14:50 to 16:00	2.42	3.55	3.39	2.91	3.52	2.74	2.65	3.40	2.90	2.99

Table 3: The averages of the intraday periods $i = \{1, \dots, 7\}$ for the realized variance, diffusive component, and jump component as defined in Section 4.2 are computed for the SPY and nine individual stocks. The reported numbers are then annualized and expressed in terms of volatility to make them easier to interpret. For example, the computation for realized variance is $\sqrt{\left(\frac{1}{Days}\right) \sum_{t=1}^{Days} RV_{t,i}^* \left(\frac{252 \cdot M}{3}\right)}$ for each intraday period i . This data corresponds to Figures 3 and 18.

$\ell_{t,j}$ and $\ell_{t,j}^c$: Statistically Significant Returns in the SPY Data

Window Size	3 Hours	3 Hours	3 Hours	4.5 Hours	4.5 Hours	4.5 Hours
Time at NYSE	$\ell_{t,j}$	$\ell_{t,j}^c$	$\ell_{t,j} - \ell_{t,j}^c$	$\ell_{t,j}$	$\ell_{t,j}^c$	$\ell_{t,j} - \ell_{t,j}^c$
9:52:30	12	4	8	14	8	6
10:10	48	8	40	53	11	42
10:27:30	46	24	22	56	30	26
10:45	31	9	22	34	16	18
11:02:30	13	4	9	16	12	4
11:20	13	9	4	11	9	2
11:37:30	8	4	4	10	10	0
11:55	2	5	-3	1	2	-1
12:12:30	3	5	-2	4	5	-1
12:30	5	6	-1	6	4	2
12:47:30	5	7	-2	6	3	3
1:05	3	6	-3	2	5	-3
1:22:30	6	5	1	4	4	0
1:40	9	5	4	4	0	4
1:57:30	11	8	3	2	6	-4
2:15	19	12	7	7	8	-1
2:32:30	34	12	22	21	13	8
2:50	42	12	30	32	8	24
3:07:30	25	4	21	16	5	11
3:25	26	12	14	25	11	14
3:42:30	22	3	19	21	4	17
4:00	37	9	28	42	13	29

Table 4: The first column plots the end of each 17.5 minute intraday period. Notably, the $\ell_{t,j}$ and $\ell_{t,j}^c$ statistics behave as we expect from Figure 20. At the window sizes of 3 and 4.5 hours $\ell_{t,j}$ flags more jumps in the morning than $\ell_{t,j}^c$ while the $\ell_{t,j}^c$ flags more jumps in the middle of the day than $\ell_{t,j}$. Furthermore, it may be misleading to compare the difference $\ell_{t,j} - \ell_{t,j}^c$ across different intraday periods. Figures 1, 8 and 15 indicate that statistically significant jumps are most frequent in the morning for all window sizes and for both the $\ell_{t,j}$ and $\ell_{t,j}^c$ statistics, which may led to a more pronounced difference between the statistics at different intraday periods.

Summary Statistics

Variables	RV_t	C_t	J_t	J_t^1
Mean	1.03	1.01	.014	.026
St. Dev.	1.30	1.24	.234	.276
Skewness	4.88	4.27	32.1	14.4
Kurtosis	36.9	27.4	1090	233
Min.	0.05	0.05	0	0
Max.	16.3	12.9	7.94	5.46

Correlation Matrix

Variables	RV_t	C_t	J_t	J_t^1
RV_t	1	.989	.272	.006
C_t	.989	1	.126	-0.02
J_t	.272	.126	1	.182
J_t^1	.006	-0.02	.182	1

Table 5: The top panel reports the sample mean, standard deviation, skewness, kurtosis, minimum and maximum.. The bottom panel reports the correlation matrix. The realized variance RV_t and the explanatory variables C_t , J_t , and J_t^1 are defined in Section 5 and are computed using the SPY data over the five year sample.

S&P 500 HAR-RV-CJ Regressions

$$RV_{t,t+h} = \beta_o + \beta_{CD}C_{t-1,t} + \beta_{CW}C_{t-5,t} + \beta_{CM}C_{t-22,t} + \beta_{JD}J_{t-1,t}^i + \beta_{JW}J_{t-5,t}^i + \beta_{JM}J_{t-22,t}^i + \epsilon_{t,t+h}$$

J_t, J_t^1, J_t^2	$h = 1$	$h = 5$	$h = 22$	$h = 1$	$h = 5$	$h = 22$
β_{CD}	0.329	0.208***	0.148***	0.377	0.201**	0.147**
β_{CW}	0.400*	0.447***	0.328***	0.401*	0.456***	0.336**
β_{CM}	0.202*	0.210**	0.212***	0.184*	0.194**	0.201*
β_{JD}	-0.249	-0.177**	-0.135***	0.131	0.126	0.070
β_{JW}	-.541**	-0.617***	-0.424**	0.086	-0.054	-0.243
β_{JM}	-.287	-0.167	0.046	-0.347	-0.336	-0.066
β_o	0.097**	0.159**	0.326**	0.101**	0.167**	0.333*
R^2	0.522	0.640	0.521	0.517	0.634	0.517
$Adj R^2$	0.519	0.639	0.518	0.514	0.632	0.515

Table 6: The table reports the ordinary least squares estimates for the HAR-RV-CJ forecasting model over daily $h = 1$, weekly $h = 5$ and monthly $h = 22$ horizons. The first three columns correspond to the jump component defined as J as in Andersen, Bollerslev and Diebold (2006). The next three columns correspond to J^1 where J^1 is the modified jump components defined in Section 5. The last rows include the R^2 and the Adjusted R^2 . Newey-West Standard errors are also computed where *implies $p < .05$, ** implies $p < .01$ and *** implies $p < .001$.

8 Figures

Statistically Significant Jump Discontinuities in the Market Portfolio




Figure 1: The z-axis plots the quantity of statistically significant jump discontinuities for the SPY where flagged jumps are defined as $\ell_{t,j} > \Phi^{-1}(.999)$. The window size k is plotted on an independent axis where it ranges from 35 minutes to 6.7 days. The intraday time intervals j are plotted on the other independent axis where j takes on integer values from 0 to 22 corresponding to the 22 returns computed at a 17.5 minute sampling frequency on a daily basis. One observation is that the number of flagged jumps appears to stabilize for a sufficiently high window size. There are also considerably more flagged jumps in the morning and late afternoon than other parts of the trading day.

Morning and Midday Volatility

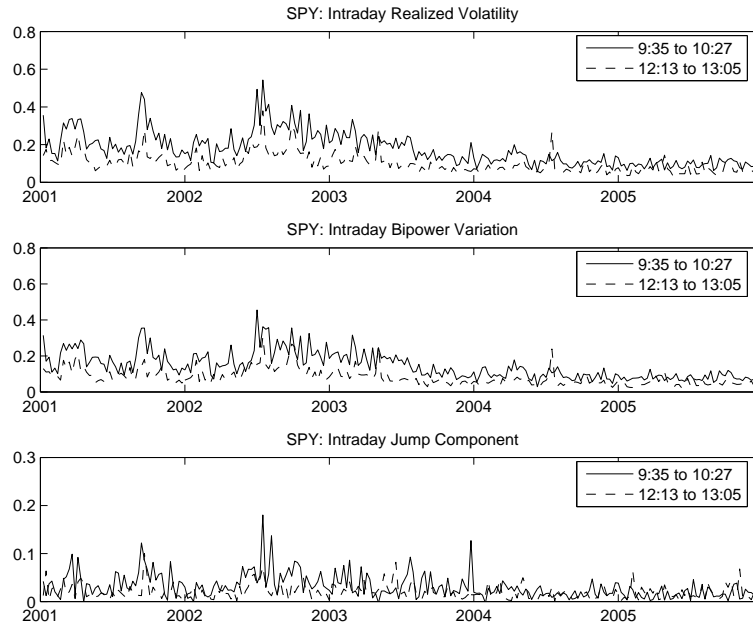


Figure 2: Each plot compares the morning and midday annualized weekly averages of the intraday realized volatility, bipower variation, and jump component for the SPY data over the five year sample period. The x-axis plots time while the y-axis plots $RV_{t,i}^*$, $BV_{t,i}^*$, and $J_{t,i}^*$ as defined in Section 4.2. The observation is made that in nearly every week of the five year sample the realized volatility, bipower variation, and jump component are higher between 9:35 to 10:27 than 12:13 to 13:05.

Intraday Patterns in the Volatility of the S&P 500




Figure 3: The averages of the intraday periods $i = \{1, \dots, 7\}$ for the realized variance, diffusive component, and jump component as defined in Section 4.2 are computed for the SPY. The reported numbers are then annualized and expressed in terms of volatility to make them easier to interpret. For example, the computation for realized variance is $\sqrt{\left(\frac{1}{Days}\right) \sum_{t=1}^{Days} RV_{t,i}^* \left(\frac{252 \cdot M}{3}\right)}$ for each intraday period i . The top figure plots $RV_{t,i}^*$ and $BV_{t,i}^*$, the middle figure plots $J_{t,i}^*$ and the bottom figure plots the relative jump component $J_{t,i}^*/RV_{t,i}^*$. From the three figures it is clear that the first hour of trading is the most volatile intraday period. Additionally, it is evident that the averages of $RV_{t,i}^*$ and $BV_{t,i}^*$ follow an intraday U -shaped pattern that is skewed to the left. This contrasts with the jump component $J_{t,i}^*$ where the volatility in the last period is lower than the preceding period.




Figure 4: The x-axis plots the time of day. The y-axis plots a fourth degree polynomial that is fit to the intraday averages of the bipower variation $BV_{t,i}^*$ for the SPY data from Figure 3.

Implications of Intraday Patterns in Bipower Variation for $\ell_{t,j}$




Figure 5: The x-axis plots the time of day. The y-axis plots various curves that represent the denominator in the $\ell_{t,j}$ statistic during the trading day under the assumption that an intraday U -shaped pattern persists in the bipower variation. The No Lag curve corresponds to the fourth-degree polynomial fit to the S&P 500 data from Figure 4. The additional curves correspond to averages of the No Lag curve over various window sizes. They are computed as $Lag_{t,j}$ introduced in Section 4.3.

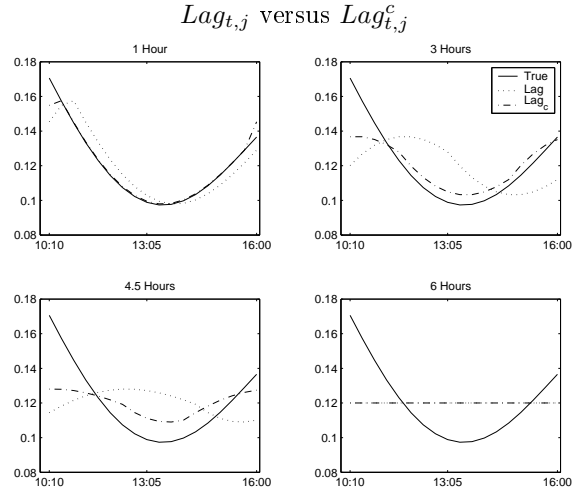


Figure 6: The x-axis plots the time of day. The y-axis plots $Lag_{t,j}$ and $Lag_{t,j}^c$ as defined in Section 4.3. As is visually evident, $Lag_{t,j}^c$ provides a better approximation for the intraday pattern in bipower variation at small window sizes because it is centered at the current time whereas $Lag_{t,j}$ is backward looking. Once the window sizes increases to 6 hours both $Lag_{t,j}$ and $Lag_{t,j}^c$ converge to the average of bipower variation over the trading day.

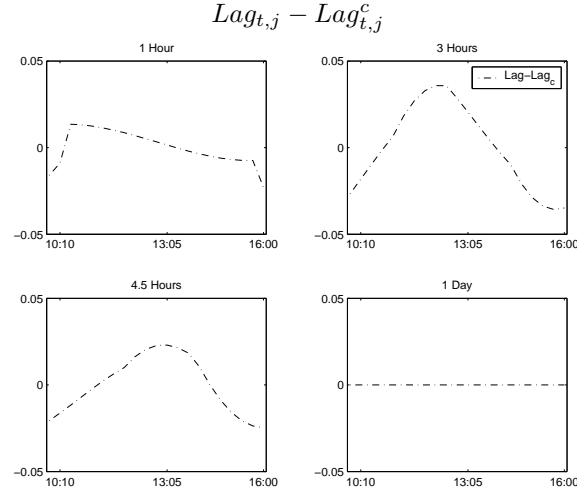


Figure 7: The x-axis plots the time of day while the y-axis plots $Lag_{t,j} - Lag_{t,j}^c$ as defined in Section 4.3. The most pronounced difference occurs at the 3 and 4.5 hour window sizes where $Lag_{t,j} > Lag_{t,j}^c$ in the middle of the day and $Lag_{t,j} < Lag_{t,j}^c$ in the morning and late afternoon.

$\ell_{t,j}$ versus $\ell_{t,j}^c$ for the S&P 500

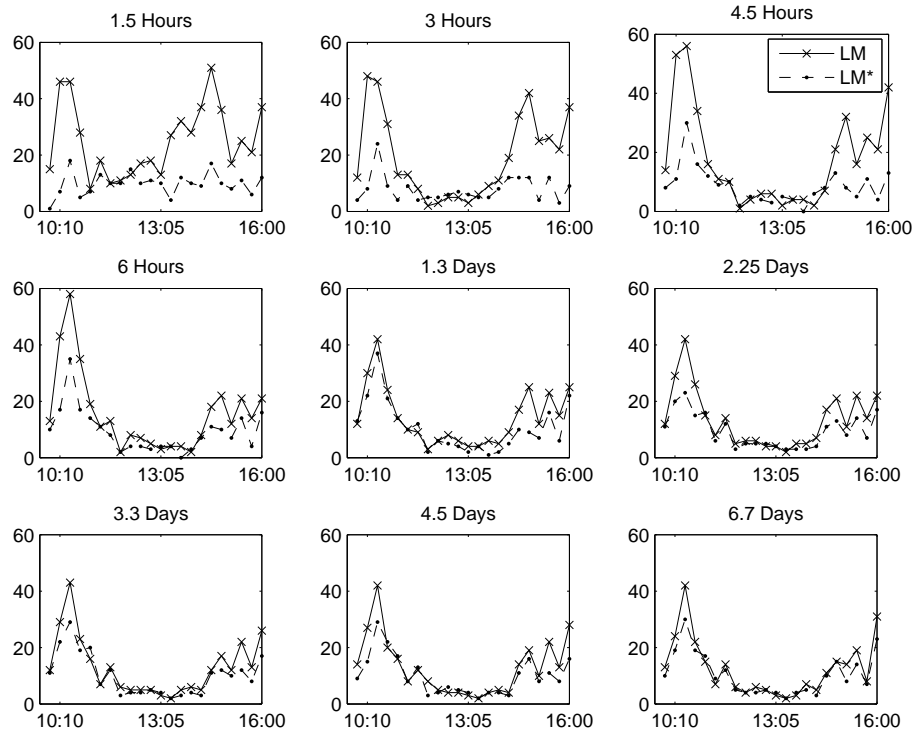


Figure 8: In the legend LM corresponds to $\ell_{t,j}$ and LM^* corresponds to $\ell_{t,j}^c$. Statistically significant jumps are defined as $\ell_{t,j} > \Phi^{-1}(.999)$ and $\ell_{t,j}^c > \Phi^{-1}(.999)$ from Section 4.3. Notably, the statistics behave as expected. At small window sizes $\ell_{t,j} > \ell_{t,j}^c$ in the morning and afternoon while this relationship is reversed in the middle of the day. As the window size increases the number of jumps flagged by the two statistics converge for each of the intraday periods. The exact number of jumps at the 3 and 4.5 hour window sizes are reported in Table 4.

9 Technical Appendix

Model for Stock Price Evolution

Included below is a Monte Carlo simulation following the standard model $dp(t) = \mu \cdot dt + \sigma \cdot dW(t)$ described in Section 2. Stock price evolution modeled as a standard geometric Brownian motion is defined as $S(t) = S(0) \cdot e^{\mu t + \sigma W(t)}$ where $W(t) \sim N(0, t)$. In the example below the annualized drift and volatility are $\mu = 10\%$ and $\sigma = 25\%$. The starting stock price is $S(0) = 100$.

Standard Model: Simulated Stock Price Evolution




Figure 9: Stock price in dollars is plotted on the y-axis versus the x-axis which plots time in years.

Data

High frequency financial data from the TAQ database is analyzed for the empirical research in this paper. In total the data set contains 1241 days with 771 intraday prices for the SPY and nine stocks from the NYSE. At the outset it is reassuring to check that the high frequency data agrees with the price series data that is readily available online. Below two plots show the closing price and log-returns for the SPY sampled at a daily frequency. Visually they confirm that the high frequency data from 2001 to 2005 behaves as expected. Similar plots can be generated by using price series data from Yahoo Finance.





Figure 10: Plots the closing price and daily returns of the SPY for each of the 1241 days in our sample. The dates on the x-axis correspond to the end of each year. The y-axis plots the dollar price for the index and the geometric returns from the closing prices on each of the 1241 days.

One concern in computing the non-parametric statistics from Section 2 is the choice of sampling frequency. Below two plots illustrate the difference between intraday sampling at a 30 second versus a 5 minute frequency. The price series corresponds to PepsiCo. on April 25th 2005. This particular day was investigated at the outset of the empirical research because it is flagged as statistically significant by the z_t statistic. Note how the 5 minute sampling actually removes 30 second prices from the plot, smoothing a significant amount of the noise in the 30 second price series. An interesting omission from the 5 minute price series occurs slightly before 10:00am. In the 30 second price series there is a trade made at \$54.95 that seems to be out of sync with the rest of the prices. Further investigation reveals that the volume on this trade was 181,700 shares compared to a five year average of 4,869 shares per trade. A plausible explanation for why the

trade occurred at a lower price is compensation for the additional risk of taking on a large order. More importantly, this example demonstrates how high frequency data can shed light on how rational markets actually operate.

30 Second Sampling: PEP Prices on April 25th 2005



5 Minute Sampling: PEP Prices on April 25th 2005




Figure 11: The y-axis plots the level price in dollars for PepsiCo. on April 25th 2005 versus the time of day at the New York Stock Exchange for two different sampling frequencies.

The debate about the appropriate sampling frequency is also related to the literature on market-microstructure noise. Black (1976), Amihud and Mendelson (1987), and Harris (1990, 1991) conclude that trading mechanisms and discrete prices can bias high frequency data. Andersen, Bollerslev, Diebold and Labys (2000) continue this discussion in the context of computing the realized variance. They observe that market-microstructure noise that begins to bias prices at very fine sampling frequencies has the effect of increasing the volatility. To select an appropriate sampling frequency their recommendation is followed to compute the average realized variance over our entire data set at various frequencies and make a selection that balances the preference for a fine frequency with the visible effects of market-microstructure noise. The 17.5 minute sampling frequency is selected by this method and is used in all subsequent analysis unless otherwise denoted.

Volatility Signature Plots

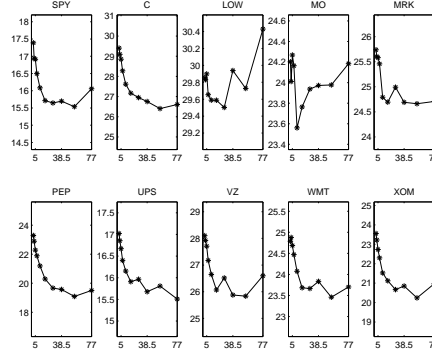


Figure 12: The x-axis plots the sampling frequency in minutes. The frequencies considered are 2.5, 3.5, 5, 7, 11, 17.5, 27.5, 38.5, 55, and 77 minutes. The y-axis plots the annualized volatility in percentage terms, $\sqrt{\left(\frac{252}{1241}\right) \sum_{t=1}^{1241} RV_t}$.

Kernel Density Estimation and Q-Q Plots

In Section 4 the $\ell_{t,j}$ statistic is computed over the five year data set for nine individual stocks and the SPY. The statistically significant jump discontinuities in the SPY data are then reported in Figure 1. Below a similar plot is included using the same data as in Figure 1. It merely collapses the curves for various window sizes from the three dimensional plot onto two dimensions. The two dimensional graph provides another visual representation for how the number of flagged jumps converges to an intraday *U*-shaped pattern as the window size increases, where a large number of returns are flagged as significant jumps in the morning and late afternoon. The pattern could also be said to resemble a “smile” or “smirk.”

Statistically Significant Jumps Within the Trading Day: SPY Data




Figure 13: The various curves correspond to the number of statistically significant jump discontinuities flagged by the $\ell_{t,j}$ statistic for different intraday periods. Each curve corresponds to a different window size.

Non-parametric kernel density estimation and Quantile-Quantile plots are included below. The Q-Q plots compare the quantiles of a standard normal distribution versus the order statistics of $\ell_{t,j}$ computed at various window sizes for the SPY data. Deviations from linearity indicate that there is evidence to potentially reject the null hypothesis of no jump discontinuities in stock prices or $\ell_{t,j} \sim N(0, 1)$. The kernel density estimates provide a non-parametric estimation of the probability distribution of $\ell_{t,j}$ at various window sizes.

Q-Q Plots for $\ell_{t,j}$




Figure 14: The x-axis plots the quantiles of the normal distribution versus the y-axis which plots the order statistics for $\ell_{t,j}$ calculated using the SPY data. Titles of the subplots indicate the window size for $\ell_{t,j}$.

Kernel Density Estimation

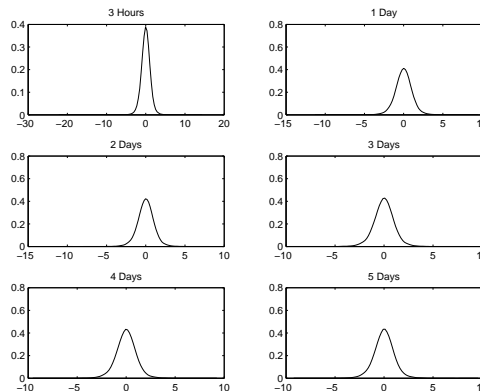


Figure 15: The non-parametric kernel density estimates provide an estimation for the probability distribution of $\ell_{t,j}$. Titles of the subplots indicate the window size.

Individual Stocks

The empirical results for the individual stocks are similar to the results observed in the SPY. One observation that can be made from the data is that flagged jump discontinuities by the $\ell_{t,j}$ statistic occur most frequently in the morning. This trend is particularly accentuated for idiosyncratic jumps. Over the five year sample the plots below indicate that for many of the stocks there were upwards of 100 flagged jumps between 9:50am and 10:10am with less than 10 jumps being flagged over intervals of the same length corresponding to the middle of the day or afternoon.

Statistically Significant Jump Discontinuities: Individual Stocks




Figure 16: The y-axis plots the number of statistically significant jump discontinuities defined as $\ell_{t,j} \geq \Phi^{-1}(.999)$. The window size used for the computation is 4.5 days as it is seen to be sufficiently long for the statistic to stabilize in Figure 1.

Another similarity to the SPY is the tendency for the volatility to be high in the morning. For Citigroup's stock the plots below annualize the realized variance, bipower variation, and jump component for each week in the sample. Over practically the entire sample one can observe that all three measures of the volatility tend to be higher in the first hour of trading than the middle of the day.

Morning and Midday Volatility

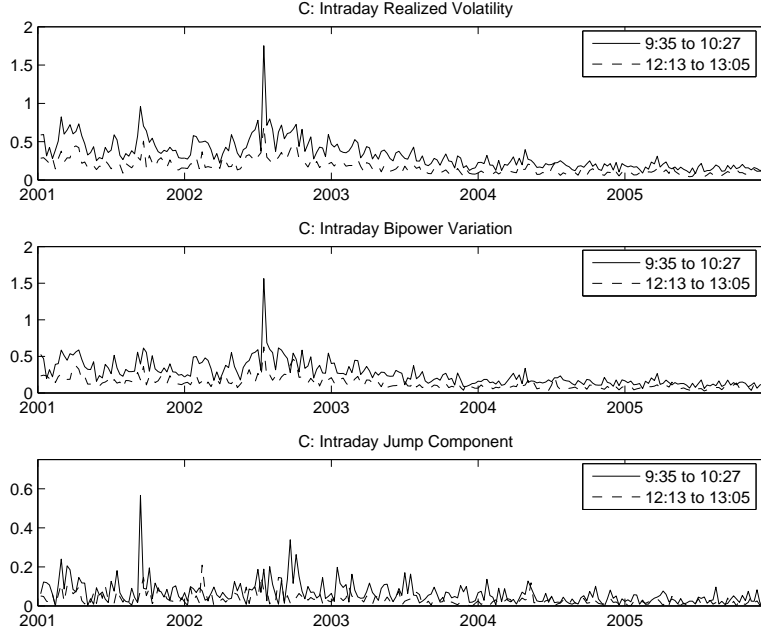


Figure 17: The x-axis plots time versus weekly averages of $RV_{t,i}^*$, $BV_{t,i}^*$ and $\max(RV_{t,i}^* - BV_{t,i}^*, 0)$ as defined in Section 4.2, or the intraday realized variance, bipower variation and jump component respectively.

On the next page the averages for the intraday realized variance, bipower variation and jump component are included for six of the individual stocks. As before they are annualized and reported in terms of volatility, so .5 would correspond to 50% annual volatility. Similar to the S&P 500, the realized variance and bi-power variation exhibit the now familiar *U*-shaped pattern. The highest period of volatility as measured by all of the non-parametric statistics is the first hour of trading. An additional observation from the data is that individual stocks are more volatile than the SPY. This is expected. Perhaps more interesting is the fact that the jump component is higher in the individual stocks than it is for the SPY. The relative jump component comprises approximately 20% of the realized variance throughout the trading day, with the exception of the last hour where it falls noticeably for both stocks and the market.

Intraday Patterns in the Volatility of Individual Stocks

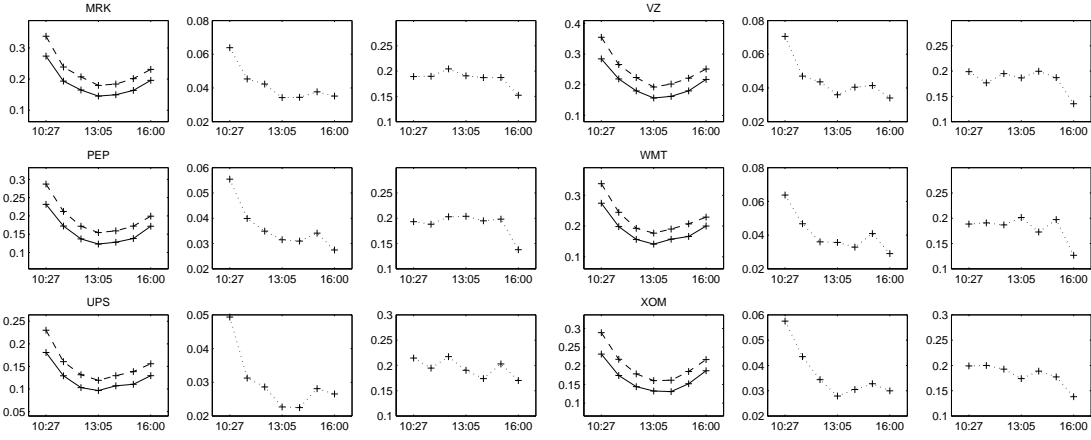


Figure 18: Averages of $RV_{t,i}^*$, $BV_{t,i}^*$ and $\max(RV_{t,i}^* - BV_{t,i}^*, 0)$ are annualized and plotted in terms of volatility. For example, the average of the realized variance for MRK is .3377 in the first hour of trading which would correspond to an annualized volatility of 33.77%. The third plot for each stock is the relative jump component or $(RV_{t,i}^* - BV_{t,i}^*)/RV_{t,i}^*$.

The final figure compares the number of statistically significant jump discontinuities for PepsiCo. defined as $\ell_{t,j} > \Phi^{-1}(.999)$ and $\ell_{t,j}^c > \Phi^{-1}(.999)$ from Section 4.3. The statistics behave as expected. At small window sizes there are more jumps flagged by $\ell_{t,j}$ than $\ell_{t,j}^c$ in the morning and afternoon. This relationship is reversed in the middle of the day. It provides even more striking evidence of intraday patterns that have an effect on the number of jumps flagged by $\ell_{t,j}$ than in the SPY. By increasing the window size to a stable level the number of flagged jumps by both statistics converge very closely and the number of jumps flagged by $\ell_{t,j}^c$ more than doubles in the period from 9:50 to 10:10am.

$\ell_{t,j}$ versus $\ell_{t,j}^c$ for PepsiCo

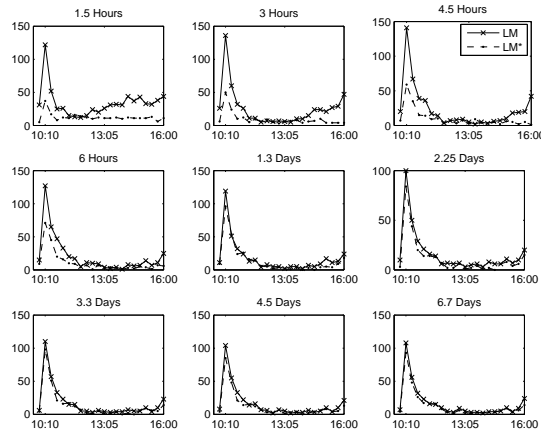


Figure 19: In the legend LM corresponds to $\ell_{t,j}$ and LM^* corresponds to $\ell_{t,j}^c$.