High Frequency Jump Characteristics of Financial Asset Prices

Michael Stanley

Duke University
Durham, North Carolina
2007

Professor George Tauchen, Faculty Advisor

Honors thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University

---

*I would like to thank Professor George Tauchen for his immense help throughout all stages of this project. I would also like to thank Professor Tim Bollerslev for many helpful insights, and Xin Huang for providing her S&P 500 index statistics for use in Section 7.

*Please e-mail any questions or concerns to me at mhs16@duke.edu.

Abstract

There is evidence that the conventional, continuous underlying function model of financial asset pricing can be improved by adding a discontinuous jump component. Barndorff-Nielsen and Shephard (2004) propose a method of detecting these jumps that compares different estimators of quadratic variation. This paper applies that detection method to the common stock of Altria Group, Inc. and attempts to link detected jumps with news events, variance and trading volume, jumps in other assets, and previous jumps. The detected jumps are stochastic, at least to the best of our knowledge. These results imply potential advantages to modeling jumps and continuous variance independently.

1. Introduction

The assumption that financial asset prices arise from a continuous, or smooth, underlying price function is prevalent in modern finance theory. While this assumption may be mathematically convenient, an empirical survey, even a fleeting one, observing the drastic price movements inherent in many financial asset prices begs the question of whether a price model that assumes a continuous sample path is a reasonable approximation of the observed data. A possible adaptation of the continuous price model that has been considered for the last quarter-century is the use of a primarily continuous price model with the addition of a discontinuous component. This new model would imply that asset prices follow a typically smooth path with infrequent but large price "jumps" scattered stochastically over time.

If these price jumps do indeed exist, there is great interest in being able to detect when and why they occur. Soon after Black and Scholes introduced their revolutionary option pricing model, Robert Murton, (1976) realizing the potentially dramatic effect of underlying asset price jumps on option prices, reworked the Black-Scholes model to allow for discontinuous jump processes. More recently, the adverse effect of jump risk on delta hedging[1] (derived directly from its effect on option pricing) has been addressed in Cont and Voltchkova (2005). In short, the presence of jumps makes perfect delta hedging impossible even in a continuously modified portfolio, and so the risk-minimizing portfolio allocation must take into account both the general market risk and the jump risk.

[1] "Delta" in this context refers to the sensitivity of a portfolio to price changes in a particular asset. In short, delta hedging is the practice of creating a portfolio containing that particular asset and derivatives thereof such that the portfolio's delta is zero, which implies that the portfolio value is independent of the particular asset price.

Creating such an optimized portfolio would naturally require some quantifiable description of this jump risk.

Jumps appear to be stochastic in both frequency and magnitude and so are inherently unpredictable. However, if one were to remove the jump component of variance, one would expect the remaining continuous portion of the variance to be more predictable than the overall variance containing jumps. The capacity to better predict the future variance of assets is vital to pricing many financial assets, (particularly options and other derivatives) but this approach also requires that financial engineers estimate new parameters, such as the price of jump risk.

Of course, these price jumps need only be considered if they make up a significant portion of overall price variance. It has been observed in Huang and Tauchen (2006) that discontinuous jumps account for approximately seven percent of the variance of the S&P 500 index. One might also expect that assets with less trading volume and diversification than the S&P 500 index would display even larger portions of discontinuous variance. In fact, while a formal test is not presented in this paper, evidence does indeed support this expectation. These findings should provide significant incentive to analyze and properly model these jumps.

The increased availability of high frequency financial data over the past twenty years and the computational power to analyze them is a vital factor in our capacity to analyze jump characteristics, as these characteristics are often most apparent and relevant at the intraday level. Figure 1 shows two days for which the daily price change of the common stock of Altria Group, Inc. (Ticker: MO) was approximately +5%. The first of these price changes is due to a discontinuous jump and the second is due to a smooth, continuous increase. At the daily level, these days appear virtually identical, but, when observed at a higher frequency, the disparity is readily apparent.

Recently, Barndorff-Nielsen and Shephard (2004) developed the first practical approach to detecting jumps, and it employs two intraday measures of variance, namely realized variation and bi-power variation. Put simply, realized variation contains the effects of discontinuous jumps while bi-power variation does not, and so the difference between the two variance measures is a measure of the portion of variance due to jumps. A simple statistical test may be constructed to determine whether or not this difference is significant and, thus, whether a statistically significant jump has occurred. It is important to note that this detection is performed purely in hindsight and offers no direct method of forecasting.

While there is great interest in improving jump detection methods, these methods are only useful if their results tell us something about these jumps. This paper will employ the variance discrepancy detection method of Barndorff-Nielsen and Shephard to detect likely jump days, but will then focus on analyzing these results rather than considering the quality of the detection method. This analysis will begin by emulating the approach of Andersen, Bollerslev, and Diebold, (2006) in addressing the validity of the jump diffusion model, considering the connection between jumps and news announcements, and constructing forecasting models in an effort to predict future variance. Where Andersen, Bollerslev, and Diebold consider market-wide assets such as the S&P 500 index and currency futures, however, this paper considers individual common stock, in particular the common stock of Altria Group, Inc. This application to common stock has thus far not been presented in the literature.

Furthermore, this paper extends the above analyses and proposes new ones in an effort to learn more about the nature, cause, and future predictability of these jumps. This paper will examine news events connected to discontinuous jumps, but will also consider news events connected to large, smooth price changes and look for similarities and differences between the two groups. Forecasting models will be modified to predict separately the continuous and discontinuous portions of variance. A new concept of proximity effects of jumps on local price behavior will be presented and tested. A cross-asset comparison will search for correlation between jumps in indices and jumps in underlying assets, which could not be performed without the jump characteristics of the underlying assets first presented in this paper. Finally, it is a goal of this paper to begin to consider trading volume and its relation to the intraday variance and jump characteristics, as this topic has not yet been thoroughly addressed in the literature. Trading volume will be discussed and utilized in many of the aforementioned analyses and will also be compared with different estimators of variance.

The format of the rest of this paper is as follows. Section 2 sets forth the theoretical framework of the jump detection tests. Section 3 discusses the data and the necessary filtering techniques associated therewith. Section 4 provides the basic jump characteristics results and the implications of their distributions. Section 5 presents case studies of possible real-world causes of jumps and large, smooth changes in Altria Group's stock price. Section 6 analyzes the effects of jumps on variance and trading volume in the near vicinity. Section 7 considers the possibility of a correlation of jump characteristics across assets, or in particular, correlation between specific companies' stock prices and the S&P 500 futures index. Section 8 analyzes the relationship between volume and the various measures of price variance. Section 9 constructs forecasting models for the different variance components. Section 10 offers conclusions and possible avenues for future research.

2. Theoretical Framework

2.1 Price Model

The standard continuous time stochastic volatility model can be expressed as a differential equation of the form:

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \tag{1}$$

where p(t) is the logarithmic asset price, μ(t) is the price drift process, σ(t) is a strictly positive, stochastic, standard deviation process, and W(t) is a standard Brownian motion process. Thus, in this model, the time rate of change of an asset's price at any instant is a random variable with mean μ(t) and variance σ^2(t).

The above equation implies that the price function is continuous in time. If this assumption is modified, however, to allow for a discontinuous jump component, the asset price model takes the form:

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) + \kappa(t)dq(t) \tag{2}$$

where q(t) is a binary variable indicating whether or not a jump has occurred, and κ(t) describes the magnitude of those jumps. The intuition behind this model is that, if one were to draw the price function as it changes through time, the continuous process could be drawn without lifting the pencil off the paper, while the discontinuous jumps require picking up the pencil and continuing on from a different price. The notation employed above emphasizes the fact that the new model, though vastly different in its implications, is essentially a modification of the classic, continuous price diffusion model rather than an entirely new approach. This notation will also facilitate the separation of the continuous and jump components of variation that is central to the forthcoming statistical analysis.

2.2 Variation Estimators

The key insight of this statistical approach to jump detection is derived from comparing different estimators of the quadratic variation (QV) of returns. These estimators will be calculated daily, so between time t-1 and t is one trading day, and QV_t refers to the quadratic variation on day t.

The quadratic variation is defined as the sum of squared returns calculated over infinitesimally small time periods, expressed symbolically as:

$$QV = \lim_{M \to \infty} \sum_{i=1}^{M} r_i^2 \tag{3}$$

where r_i is the logarithmic return and M is the number of time periods. When applied to the continuous time model of equation (1) over the time period from t-1 to t, the quadratic variation becomes:

$$QV_t = \lim_{M \to \infty} \sum_{i=1}^{M} r_i^2 \to \int_0^t \sigma^2(s)\,ds \tag{4}$$

where σ is the stochastic, standard deviation process from above and the integrated term is called the integrated volatility (IV). When applied to the continuous-time jump diffusion process described in equation (2), the quadratic variation over the same time period is known to be:

$$QV_t = \int_0^t \sigma^2(s)\,ds + \sum_{0<s\le t} \kappa^2(s) \tag{5}$$

The first term in this equation represents the volatility due to the continuous process and is equal to the integrated volatility found above. The second term is the volatility due to jumps and is equal to the summation of the squared magnitudes of all the jumps between time t-1 and time t. One should note that, if $\kappa(s)=0$ from time t-1 to time t, equation (5) reduces to the integrated volatility and is equal to equation (4).

The first measure of variance employed in jump detection is the realized variation, RV, defined as:

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2 \tag{6}$$

where $r_{t,j} = p(t-1+j/M) - p(t-1+(j-1)/M)$ and M is the number of samples in period t. If period t is one trading day, then M is the number of samples taken daily and $r_{t,j}$ denotes the j^{th} sample on day t. One should note that the realized variance function is very similar to the quadratic variation function except that it considers returns over finite time periods rather than the infinitesimally small time periods of quadratic variation. Thus, it makes intuitive sense that the asymptotic value of realized variance in continuous time, as discussed in Barndorff-Nielsen and Shephard (2002) and elsewhere, is:

$$\lim_{M \to \infty} RV_t = \int_{t-1}^{t} \sigma^2(s)\,ds + \sum_{t-1<s\le t} \kappa^2(s) \tag{7}$$

which means that it is a consistent estimator of the aforementioned quadratic variation.

The second variance measure employed is the bi-power variation, BV, defined as:

$$BV_t = \mu_1^{-2}\left(\frac{M}{M-2}\right)\sum_{j=3}^{M} \left|r_{t,j-2}\right|\left|r_{t,j}\right| \tag{8}$$

where $\mu_1 = \sqrt{\frac{2}{\pi}}$. While the summed terms for both RV and BV are quadratic with respect to returns, the BV terms are staggered, multiplying together returns that are separated by a time period, whereas RV simply squares each return. The constants outside of the BV summation are used to compensate for the degrees of freedom lost from the lag of returns so that the magnitudes of RV and BV may be nominally compared. Interestingly, the asymptotic theory described by Barndorff-Nielsen and Shephard (2004) shows that, under certain assumptions:

$$\lim_{M \to \infty} BV_t = \int_{t-1}^{t} \sigma^2(s)\,ds \tag{9}$$

which estimates the continuous portion of the quadratic variation of equation (4).

The bi-power variation is a measure of volatility that is robust to jumps. In determining RV_t, the large, jump returns are squared and so their effect on the variance measure is magnified. If one considers equation (8) and assumes that jumps are scarce, it is unlikely that both $r_{t,j-2}$ and $r_{t,j}$ are jumps, and so the effect of the large, jump return on the daily BV_t will be washed out when the jump return is multiplied by a small, non-jump return. As the size of the sampling periods approaches zero, the effect of jumps becomes negligible in the BV_t measure but remains significant in the RV_t measure. Even in the finite sampling case, large returns will have a much smaller impact on lagged variance measures than on non-lagged measures.

Since RV_t estimates the entire quadratic variation while BV_t estimates only the continuous portion, RV_t-BV_t is a measure of the jump component of the quadratic variation:

$$\lim_{M \to \infty} (RV_t - BV_t) = \sum_{t-1<s<t} \kappa^2(s) \tag{10}$$

The fraction of the total variation attributable to this jump component can be measured by simply comparing the above difference to RV_t. This statistic is known as the relative jump (RJ) and is defined as:

$$RJ_t = \frac{RV_t - BV_t}{RV_t}. \tag{11}$$

2.3 Hypothesis Test for Jumps

The asymptotic distributions described above can be employed to formulate a statistical test to address the validity of adding this jump component to the standard price model. Since, in the absence of jumps, RV_t and BV_t have the same expected value, it is natural to construct a hypothesis test with the null hypothesis that there are no jumps. The two alternative hypotheses are RV_t being larger than BV_t and vice versa, but, since the jump component κ^2 is strictly positive, the test only lends evidence to jumps if the hypothesis that RV_t is larger than BV_t is found to be true. Thus, the hypothesis test is one-sided and can be represented algebraically as:

$$\begin{aligned} H_0 &: E(RV_t) - E(BV_t) = 0 \\ H_A &: E(RV_t) - E(BV_t) > 0 \end{aligned}$$

To test this hypothesis, the RV_t-BV_t must be studentized by subtracting off its mean and dividing by its standard deviation. The mean, under the null hypothesis, is zero. Estimating the standard deviation is far from trivial, but since RV and BV both asymptotically approach the integrated volatility in the absence of jumps, the variance of their difference is proportional to $\int_{t-1}^{t} \sigma^4(s)ds$. This value is known as the integrated quarticity and can be estimated with a statistic known as the quad-power quarticity (QP) defined in Barndorff-Nielsen and Shephard (2004) as:

$$QP = M\mu_1^{-4}(\tfrac{M}{M-3})\sum_{j=4}^{M}|r_{t,j-3}||r_{t,j-2}||r_{t,j-1}||r_{t,j}| \tag{12}$$

The constants outside the summation are necessary to correct for the degrees of freedom lost from the lagged returns. A similar statistic called the tri-power quarticity (TP) is employed by Andersen, Bollerslev, and Diebold (2004) to estimate the integrated quarticity and offers nearly identical results.

2.4 Daily Statistics

With QP_t, a consistent estimator of the integrated quarticity, a normalized statistic can be created to test the null hypothesis. The simplest form of this z-statistic is:

$$z_{QP,t} = \frac{RV_t - BV_t}{\sqrt{[(\frac{\pi}{2})^2 + \pi - 5](\frac{1}{M})QP_t}} \tag{13}$$

The bracketed constant in the denominator is a characteristic of the joint distribution of RV_t and BV_t, as shown in Barndorff-Nielsen and Shephard (2004). In hopes of improving finite sample performance, statistics can also be created using the logarithms of the volatility measures, as suggested in Barndorff-Nielsen and Shephard (2002). This statistic is of the form:

$$z_{QP,l,t} = \frac{\log(RV)_t - \log(BV_t)}{\sqrt{[(\frac{\pi}{2})^2 + \pi - 5](\frac{1}{M})(\frac{QP_t}{BV_t^2})}} \tag{14}$$

One should note that, since the logarithm renders the numerator dimensionless, the denominator is normalized by dividing QP_t by BV_t^2. This statistic can be further modified, as in Barndorff-Nielsen and Shepherd (2004), by making a maximum adjustment, which replaces $\frac{QP_t}{BV_t^2}$ with $\max(1, \frac{QP_t}{BV_t^2})$. Theoretically, $\frac{QP_t}{BV_t^2}$ is always less than one, but in practice this is not always true, and this truncation improves the performance of the statistics.

This paper will use a z-statistic that employs the relative jump measure in the numerator. Like the logarithm form, RJ_t is dimensionless and so the denominator must again be normalized. The statistic used here is of the form:

$$z_{QP,rm,t} = \frac{RJ_t}{\sqrt{[(\frac{\pi}{2})^2 + \pi - 5](\frac{1}{M})\max(1, \frac{QP_t}{BV_t^2})}}. \tag{15}$$

In the absence of jumps, $z_{QP,rm,t} \rightarrow N(0,1)$ as $M \rightarrow \infty$. Thus, a simple t-test on the z-statistics can be used as a jump detector. Intuitively, on a day when RV_t is large relative to BV_t and QP_t, there is a high probability that $\sum_{t-1<s<t} \kappa^2(s)ds$ is non-zero, which implies that one or more jumps have occurred on that day. Consequently, the z-statistic for that day will be large, or, more precisely, significant at a high confidence level. For example, a z-statistic greater than 3.09 indicates a jump day at the 99.9% confidence level. It is important to note that these are *daily* statistics, and so the test indicates days on which a jump occurs rather than the jumps themselves.

3. Data and Data Filtering Techniques

This paper will consider primarily the common stock of Altria Group, Inc., formerly Phillip Morris. Altria Group has a market cap of approximately $150 billion and is a member of both the DJIA and the S&P 500. While there is no particular uniqueness that motivated the use of Altria Group stock in the following analyses, the company does have a very fortuitous attribute: more than two-thirds of its annual revenue comes from the tobacco industry, an industry that has received considerable public scrutiny in the last few years. The legal ramifications of this attribute are extremely relevant to the jump characteristics of Altria Group, Inc.

The primary data is common stock price data for Altria Group, Inc. (Ticker: MO) from 1999-2005 as acquired from the WRDS TAQ database. These data include the price, share volume, and exchange on which every trade occurred for the 1,760 trading days. This data was filtered so as to only include NYSE trades and was then sampled at 5-minute intervals throughout the standard trading hours, (9:30am to 4:00pm EST) creating 78 price values and 77 return values per day. Share volume on the NYSE was summed up over 5-minute intervals as well. Overnight returns (price changes outside of trading hours) are not considered in this paper.

The sampling program used a backtracking method, whereby the transaction price occurring immediately before each 5-minute tick was taken as the price at that tick. While this backtracking may cause complications in assets with small trading volume when the magnitude of this backtrack may be a few minutes on average, Altria Group common stock had an average daily share volume of about six million over the data set and this backtrack was only a few seconds on average. On the occasions in which an entire 5-minute interval contained no trades, the most recent trade from the interval before was taken as the price. This technique is reasonable since, if there have been no trades, one may assume that the price has not changed a significant amount.

Filters were also put in place to correct or remove erroneous data from the sample. The most common errors found in the data were time stamp errors (attributing a trade to the wrong time) and simple typos. These errors, if they occur on a sampled trade, will show up as a spike in the price, changing drastically over one interval and returning to near the original price level in the next interval. Unfortunately, the jump statistics employed in this paper are quite susceptible to these errors, and will often generate a large z-statistic because of a price spike. Filters were set up to order the prices by time stamp and to remove price spikes by eliminating large returns that are immediately followed by large returns in the opposite direction. These filters eliminated the vast majority of errors in the data set.

The S&P 500 futures data[2] that is utilized in Section 7 contain daily z-statistics and RJ_t values from 1999-2002. This data set comprises 980 trading days.

Price data for five other assets, namely the S&P 500 spider index fund (Ticker: SPY) and the common stock of Procter and Gamble Co., (Ticker: PG) American International Group, Inc., (Ticker: AIG) Wal-Mart Stores, (Ticker: WMT) and International Business Machines Corp. (Ticker: IBM) were acquired for January 3, 2001 from the WRDS TAQ database. The common stock data was limited to NYSE trades, while the spider index, which is not traded on the NYSE, was not filtered for a particular market.

4. Jump Statistic Results

The four panels of Figure 2 show, in descending order, RV_t, BV_t, RJ_t, and the daily statistic $z_{QP,rm,t}$ for Altria Group from 1999-2005. Two RV_t values exceed the chart dimensions and are approximately 104 and 74. The horizontal line on the z-statistic plot at $z = 3.09$ corresponds to the 99.9% significance level.

[2] The S&P 500 futures data was generously provided by Xin Huang and originally described in Huang and Tauchen (2006).

Over the 1,760 day sample period, the z-statistic detected 190 jumps at the 99.9% significance level, which is 10.8% of the total sample days. In the absence of jumps, the z-statistic should be approximately N(0,1) distributed (exactly in continuous time) and so the test would be expected to only detect 1.8 jumps at this significance level. The fact that the number of jumps actually present is two orders of magnitude larger than this expected value is incredibly strong evidence for the existence of jumps.

Figure 3 compares the distribution of the 1760 daily z-statistics in the sample to the normal distribution predicted by the null hypothesis. The histogram is the actual distribution of returns, while the associated curve has been smoothed with a kernel estimator. The mean z-statistic is clearly greater than zero and the distribution features more outliers than would be expected. Both of these properties are further evidence of the existence of jumps.

Table 1 provides basic statistics for the considered variance and jump measures, as well as 5-minute share volume. One should note that RV_t, BV_t, and daily volume have standard deviations larger than their means, which implies that they are all highly volatile. Also, since trading volume cannot be negative, its distribution must be significantly skewed. Finally, it is particularly important to note that the average RJ_t value is 14.2%, which means that discontinuous jumps account for more than 14% of the total price variation, which is twice as large as the portion accounted for in the S&P 500 index, as stated in Huang and Tauchen (2006).

5. Causes of Price Jumps and Smooth Changes

Prices can make large movements in one of two ways: discontinuous jumps and continuous smooth movements. This study considers four days with price changes of at least 4% over the day, two of which contain a jump and two of which are smooth price changes.

Figure 4 shows the price and trading volume of Altria Group's common stock for two significant jump days, February 4, 2005 and August 18, 2005. One should recall that this contains only NYSE trades sampled at 5-minute intervals, and the volume is the total number of shares traded during each interval.

On February 4, 2005, a US federal court ruled that tobacco companies could not be forced to pay the proposed $280 billion of past earnings that they were being sued for. As Phillip Morris accounts for about half of the domestic cigarette market, this decision was very good news and the stock price jumped about 4%, a market cap rise of about $5 billion, between 2:00pm and 2:05pm. The z-statistic on February 4 was 4.81 which is significant at virtually any level. Also, the 5-minute volume jumped to approximately twenty times its sample average during and immediately after the jump.

On August 18, 2005, the Illinois Supreme Court reversed a $1 billion judgment against State Farm Insurance Co. The relevance of this event for Altria Group is that the case was a class-action lawsuit similar to the pending $10.1 billion class-action case against Altria Group's Phillip Morris USA and was viewed as an indication of how that case would be ruled on as well. This decision was announced at approximately 10:00am EST and, as is apparent in Figure 4, caused Altria Group's stock price to immediately jump between 10:05am and 10:10am. The z-statistic on August 18 was 5.37 and so very significant and the 5-minute volume during and after the jump was about fifteen times the sample average.

In contrast with the above findings, Figure 5 shows price and trading volume of Altria Group for two large, smooth movement days. Both days have returns over the day of over 5% and z-statistics below 2.

On October 21, 2002, Philip Morris announced that it was considering launching a new, discount cigarette brand to compete with the small, discount brands that had been eroding its market share during the recent recession. The market considered this positive news and the stock price increased 6.1%, but it did so smoothly over the first four hours of trading and the z-statistic for the day was -0.44. While one may suggest many reasons why the price increase was more drawn out than the jumps described above, it was likely because 1) the announcement was made before trading began, 2) the announcement was indecisive and made no promises, or because 3) the market was unsure how to value this new project.

On May 24, 2004, a federal judge refused to reduce the $280 billion lawsuit described above (this refusal occurred nine months before the tobacco industry won the lawsuit). The stock price dropped 8.9%, but did so smoothly over a few hours, and the z-statistic was 0.64. One could imagine that, since this ruling was not final and the expected value of the ensuing lawsuit still had to be estimated probabilistically, it took the market longer to value the effects of this ruling.

This case study provides some new insight into the causes of jumps for specific companies. While jumps in market indices like the S&P 500 and currency exchange rates are most often caused by macroeconomic announcements, as observed in Andersen, Bollerslev, and Diebold (2006), the most significant jumps in the stock price of Altria Group were due to company-specific, or at least industry-specific, announcements and, in both these cases, more specifically, lawsuits. Since earnings announcements are typically made outside of trading hours and so have virtually no effect on price volatility during standard trading hours, one could expect other announcements, such as lawsuits and executive resignations to be the primary catalysts of individual company price jumps.

Interestingly, some company specific news announcements were accompanied by smooth price changes. The evidence of this case study suggests that jumps are caused by decisive and easily quantifiable announcements. The jump day rulings were conclusive and involved lawsuits of a precise amount, while the smooth day announcements were less absolute and would require significantly more modeling to value. It is reasonable that the market would come to a consensus on the value of the first two announcements almost instantaneously but take a few hours to agree on the value of the latter two announcements. Also, announcements made outside of trading hours, while unable to cause jumps, might cause smooth price changes during trading hours if the less active after-hours markets cannot provide the necessary liquidity. These announcements would still have virtually no effect on the daily jump statistics.

No clear evidence was found of macroeconomic announcements causing jumps in the stock price of Altria Group, although such evidence for other companies is presented in Section 7. The apparent effect of company and industry related announcements suggests that estimating the jump characteristics of individual firm-based assets would require company and industry specific parameters, in addition to those of the general market.

One must note that this analysis can only be performed in hindsight, since announcements are inherently unpredictable to some extent and cannot typically be categorized before they occur. If an announcement were perfectly predicted, then it would already be accounted for in the asset price. In summary, while there is significant evidence that many jumps are associated with real world events and even that certain types of events are more likely to cause jumps than others, these events and their relevant properties are not much easier to forecast than are the asset prices themselves.

6. Proximity Effects of Jump Days

A qualitative analysis of the significant jump days shown in Figure 4 suggests that volatility and volume may differ before and after jumps occur. Since jumps often represent an abrupt change in the underlying asset or the market's valuation thereof, it seems reasonable that different traders and shareholders would have different perceptions of this change and would cause higher volatility and trading volume in striving to augment their positions. These potential increases in volatility and volume after a jump day will henceforth be referred to as proximity effects.

To substantiate this informal observation, a comparison of BV_t, RJ_t, and daily volume for time periods immediately before and immediately after significant jump days was performed. One should note that the continuous (BV_t) and jump (RJ_t) components of volatility are addressed separately. The method of comparison was to calculate the difference between the measured values before and after jump days. The comparisons were made for two time periods; one day and one week (5 trading days) before and after each jump. The results of this test are provided in Table 2.

Contrary to the initial qualitative prediction, the results of this test provide virtually no evidence for any proximity effects of jump days. None of the variance measures or volume displays a statistically significant difference around jump days, although this is partly due to the fact that both volatility and volume have large standard deviations. Interestingly, while the effects on the single day before and after might be considered relevant, (daily volume showed a mean difference of 8%) the effects measured at the week time period are almost completely undetectable. This suggests that, if there are any proximity effects associated with significant jumps, they are of short duration. This makes qualitative sense if one considers Figure 4 and notes the rapid return to normality after the jump, most evident in the volume plot.

While it would be interesting to test this theory at the intraday level, such an experiment poses a few technical dilemmas. First, in order to compare volatility or volume within the jump day before and after it occurs, one must be able to determine when exactly the jump occurs. The jump detection tests as they are currently constructed offer no information as to when within a day the jump occurs, and so a new method of detection would have to be created. While something as simple as determining the largest absolute 5-minute return in the jump day might be reasonably effective, it has not been considered in this paper. Second, if one can indeed determine the exact time when the jump occurs, one then needs a measure of variance that either compares a different number of returns (jumps don't have to happen in the middle of the day) or combines returns from multiple days into one statistic. If the latter is used, it would doubtless have to consider overnight returns, which are not typically considered in intraday frequency data analysis and would add many complications.

The absence of any real proximity effects supports the discontinuous jump diffusion price model. Since jump days have virtually no effect on the smooth, continuous days around them, the jump volatility and continuous volatility can be considered independent of one another and may be addressed separately. Also, these results suggest that predicting jumps based on any measure of volatility or volume will be ineffective, since neither offers any indicator that a jump is forthcoming. This idea will be addressed formally in Section 9.

7. Cross-asset Comparisons of Jump Characteristics

One might expect that a stock with a CAPM beta close to one would have a price jump when the S&P 500 index has a price jump. According to several authoritative financial websites (Reuters.com and Yahoo! Finance, among others), Altria Group has a beta between 0.86 and 1.08. Plotting the daily z-statistics of the S&P 500 against those of Altria Group presents a simple way of determining whether or not any correlation exists between the jump characteristics of the two assets. A linear regression would be more formal, but a glance at Figure 6 reveals why a regression is, in this case, unnecessary.

The daily z-statistics for Altria Group common stock and S&P 500 futures will be compared for the four years 1999-2002. This data set contains 980 trading days. Over this time period, Altria Group experienced 137 jump days significant at the 99.9% level while the S&P 500 futures index experienced 22 significant jumps. As one might expect, the average jump component for the S&P 500 was significantly smaller (7%) than that of Altria Group (16%).

Figure 6 plots the z-statistics against each other for every day in the data set. Surprisingly, there appears to be no correlation whatsoever between the two sets of statistics. On only one occasion do both assets experience a significant jump. Though this paper only considers Altria Group and the S&P 500, similar comparisons have been performed in contemporaneous work for over 30 common stocks, and with comparable results.

These results demand a closer consideration of the dynamics of a price jump in an asset index. Under the assumption of no arbitrage, an index price can only jump when the value of its composing assets jumps. In continuous time, the index would jump if and only if at least one underlying asset itself jumped and was not cancelled out by another asset jump in the opposite direction. However, in discrete time the dynamics are more subtle, and the index may not jump when one or more underlying assets do jump if the magnitude of these underlying jumps becomes insignificant in the aggregate index portfolio. For example, if Altria Group's price jumps 5% and the prices of the rest of the S&P 500 stocks do not change, the index will probably not show a statistically significant jump. As is the case with individual assets, the determination of jumps in discrete time is predominantly a statistical issue.

There are also circumstances under which the index can jump in the absence of statistically significant underlying asset jumps. The S&P 500 index is less volatile than individual assets and so, since the jump statistic compares the jump volatility (RV-BV) to overall volatility (RV), a price change of a given magnitude may be a statistically significant jump for the index and not for its underlying assets. This will only occur, however, if a large number of underlying assets experience simultaneous price changes in the same direction. At this point it becomes clear that index jumps are determined not only by statistically significant jumps in underlying assets, but also by the covariance of the underlying asset prices. Thus, in order for an event to cause a jump in the S&P 500, it must either cause a small group of stocks to jump a large magnitude, or a larger number of stocks to jump a smaller magnitude. It seems reasonable that the lawsuits mentioned previously that caused Altria Group's stock price to jump did not have enough of a widespread effect to cause the S&P 500 index to jump. Also, one could conceive of events that would only moderately affect Altria Group but would also moderately affect lots of other companies, thus causing a jump in the index.

A case study of a particularly significant jump day for the S&P 500 provides some insight into the properties of these more widespread events. January 3, 2001 is the most significant jump day in the data set for the S&P 500 (z=6.43) and coincides with the Federal Reserve unexpectedly lowering the federal funds rate 50 basis points. This announcement occurred at approximately 1:15pm EST and caused the stock market to run-up drastically.

Figure 7 plots the price of six assets on January 3, 2001. Though the S&P 500 futures data was unavailable, SPY is an exchange traded fund that traces the S&P 500 index. In addition to Altria Group, price charts are provided for Procter and Gamble Co., American International Group, Inc., Wal-Mart Stores, and International Business Machines Corp. All five companies are members of the S&P 500 and the DJIA and have CAPM betas close to 1.

Aside from the enigmatic price drop for Altria Group, the other four stocks experienced a drastic jump in price at exactly 1:15pm, the same time that SPY experienced a jump. Furthermore, three of the five stocks have a significant z-statistic for January 3, which is strong evidence for the fact that, at least in very drastic situations, macroeconomic announcements can cause discontinuous jumps in individual companies as well as market-wide indices. Interestingly, the z-statistic for the index was larger than all the underlying stocks except IBM, which supports the suggestion that the index could indeed experience a statistically significant jump in the absence of underlying asset jumps. While much of the jump volatility of individual stocks is eliminated in a diversified index, the index itself has additional jump characteristics associated with the covariances of the underlying assets.

8. Volume-Variance Regressions

The relationship between trading volume and volatility has been studied extensively (Tauchen and Pitts (1983) and Shephard (1996), among others), but little work has been done using new high-frequency measures of variance. The following regressions are all univariate OLS regressions of the form:

$$Volume = \alpha + \beta * Variance \tag{16}$$

where α and β are the OLS coefficients, and volume denotes daily share volume. The regression will be performed with four different measures of variance: daily absolute returns, RV_t, BV_t, and RJ_t. Table 3 presents the results for this regression and the logarithmic regression:

$$Log(Volume) = \alpha + \beta * Log(Variance) \tag{17}$$

A few insights may be made from these results. First, the fact that RV_t and BV_t exhibit far more explanatory power (as evidenced by the larger R^2 values) than the absolute daily returns lends credibility to the use of high-frequency data analysis. Whether or not one considers jump characteristics, there is much to be learned about an asset's price dynamics by sampling at intraday frequencies.

Second, and more pertinent to the topics of this paper, daily volume displays almost no correlation with RJ_t, which implies that volume is virtually independent of jump characteristics. This conclusion is further supported by the slightly larger R^2 value for the BV_t regression as compared with the RV_t regression.

This result can perhaps be better explained through the volatility-volume model of Tauchen and Pitts (1983). In this model, an asset's price at time i, P_i, is an average of the reservation prices, P_{ij}^*, of J traders such that:

$$P_i = \tfrac{1}{J}\sum_{j=1}^{J} P_{ij}^* \tag{18}$$

The change in price P_{ij}^* from time i-1 to time i is modeled as the sum of a universal component, φ_i, and a trader-specific component, ψ_{ij}, and so:

$$\Delta P_{ij}^* = P_{ij}^* - P_{i-1,j}^* = \varphi_i + \psi_{ij} \tag{19}$$

$$\text{and} \quad \Delta P_i = \tfrac{1}{J}\sum_{j=1}^{J} \Delta P_{ij}^* = \varphi_i + \overline{\psi}_i \tag{20}$$

where $\overline{\psi}_i$ is the mean of the trader-specific components at time i. Trading volume, V_i, is modeled as a factor α multiplied by the difference between the change in trader reservation prices and the change in the market price, or:

$$V_i = \tfrac{\alpha}{2}\sum_{j=1}^{J} \left| \Delta P_{ij}^* - \Delta P_i \right| \Rightarrow \tfrac{\alpha}{2}\sum_{j=1}^{J} \left| \psi_{ij} - \overline{\psi}_i \right| \tag{21}$$

The central insight in this model is that trading volume does not depend on the universal component of price change, φ_i, and so if the legitimacy of a jump in an asset's price is generally agreed upon among traders, it need not have any more effect on trading volume than a smooth price change of the same magnitude. In general, it is reasonable that traders' reservation prices will differ from the market price just as much during a large, smooth price change as they will over a jump, and so the experimental finding that daily volume is virtually independent of the jump volatility component, RJ_t, agrees with the model's prediction.

9. Volatility Forecasting Models

It is well documented that volatility displays significant autocorrelation (Baillie and Bollerslev (1996)) and, more recently, this relationship has been confirmed by Corsi (2004) using high frequency volatility measures. In the presence of jumps, however, this concept becomes more subtle. One would expect that the jump volatility component might display less autocorrelation than the continuous component because jumps have proven to be virtually unpredictable. For this reason, forecasting models that separate these two volatility components may provide new insight into volatility forecasting.

The following forecasting models will be logarithmic in both variables, and so it is easier to model the jump component with JM_t such that:

$$JM_t = \frac{RV_t}{BV_t} \tag{21}$$

rather than RJ_t, which is often zero (or negative if it is not truncated). The forecasting models used here regress JM_t, BV_t, and RV_t on the lagged variables JM, BV, and daily volume, and are of the form:

$$\begin{aligned}\log(JM) &= \beta_{dj}\log(JM_d) + \beta_{wj}\log(JM_w) + \beta_{mj}\log(JM_m) + \beta_{db}\log(BV_d) \\ &+ \beta_{wb}\log(BV_w) + \beta_{mb}\log(BV_m) + \beta_{dv}\log(Vol_d) + \beta_{wv}\log(Vol_w) + \beta_{mv}\log(Vol_m)\end{aligned} \tag{22}$$

where subscripts d, w and m denote day, week and month variable lags, which are defined as:

$$JM_{t,t+h} = \tfrac{1}{h}[JM_{t+1} + JM_{t+2} + ... + JM_{t+h}] \tag{23}$$

and h=1, 5, and 22 for day, week, and month lags, respectively. The BV and volume lag variables are generated with the same equation. These regressions were performed in the form of Equation (22), but two similar regressions were performed as well. The volatility measures were regressed on just lagged JM and BV variables, and also just on BV or JM in a typical autocorrelation regression. These regressions for JM are as follows:

$$\begin{aligned} \log(JM) &= \beta_{dj}\log(JM_d) + \beta_{wj}\log(JM_w) + \beta_{mj}\log(JM_m) + \beta_{db}\log(BV_d) \\ &+ \beta_{wb}\log(BV_w) + \beta_{mb}\log(BV_m) \end{aligned} \tag{24}$$

$$\log(JM) = \beta_{dj}\log(JM_d) + \beta_{wj}\log(JM_w) + \beta_{mj}\log(JM_m) \tag{25}$$

The results of these regressions are provided in Tables 4a-c.

Considering the JM regressions of Table 4a reveals a great deal about jump forecasting, or the impossibility thereof. First, it is readily apparent that jumps display virtually no autocorrelation, as the first regression yields an R^2 value of 0.0086. This agrees with the results from Section 6 that showed jump days having insignificant proximity effects on the volume and volatility of the surrounding days. Interestingly, the day and month lagged variables of BV and daily volume are all statistically significant and provide more predictive power than do the lagged JM variables. It should be noted, however, that the day and month lagged variables are of comparable magnitude and opposing signs, which implies that the jump component has short-memory dependence on these variables and little long-memory dependence. Despite these statistically significant variables, none of the JM regressions yields an R^2 value greater than 0.087, which agrees with previous findings that suggest that jumps are stochastic.

The BV regressions reveal that the continuous volatility component is far more predictable than the jump component. BV is highly autocorrelated, as all three lagged BV variables are significant in all three regressions and generate R^2 just below 0.6. None of the lagged JM variables are ever statistically significant, which supports the conclusion that knowledge of past jump characteristics provides no help predicting continuous volatility. The day and week lagged volume variables are statistically significant, but again are of opposite sign and provide little additional predictive power.

The RV regressions mimic the BV regressions almost exactly, but are provided to show that the power to forecast volatility is entirely encapsulated in the forecasting of the continuous component.

These regressions provide a more formal argument for the earlier conclusion that jumps are stochastic. The continuous volatility component, BV, is highly autocorrelated and predictable, but no more so than the entire volatility, RV. Thus, while these regressions of decomposed volatility components has provided significant insight into the mechanics of volatility forecasting, it has little or no inherent forecasting power.

10. Conclusions and Future Considerations

The application of the variance difference jump detection method to the common stock of Altria Group generated results comparable to those found when the method was applied to market indices and currency exchange markets. The number of significant jumps detected was more than two orders of magnitude greater than one would expect if the price function was continuous, which is strong evidence for the presence of jumps. Additionally, the distribution of daily z-statistics, which should be N(0,1) in the absence of jumps, has a mean well above zero and is not normally distributed. The magnitude of the jump component of variation was found to be even larger for this individual stock than for the indices and exchange rates, with an average of over 14%.

The news-related causes of discontinuous jumps in Altria Group's stock price were found to be mostly company specific announcements rather than the macroeconomic announcements found to affect the larger market assets. Though many jumps can be undoubtedly linked to these news announcements, similar announcements can cause large, smooth price movements and it is very difficult to predict whether or not a particular event will cause a discontinuous, price jump or a smooth price movement. Any proximity effects of significant jumps on the surrounding days was found to be either insignificant or of very short duration. Curiously, there was found to be no correlation between the daily z-statistics of Altria Group stock and the S&P 500 futures index, but a case study of the effects of an unexpected Federal Reserve announcement on the price of the S&P 500 index and five common stocks showed convincingly that a monumental macroeconomic announcement can, indeed, cause indices and common stocks to jump simultaneously. Regressions of daily volume versus variance found that high-frequency variance measures are much more strongly correlated with volume than are absolute daily returns, but volume appears uncorrelated with jump characteristics. The forecasting regressions confirmed that, while the continuous component of volatility is highly autocorrelated and exhibits long memory characteristics, the jump component is stochastic and so virtually unpredictable. Also, daily volume was found to be a significant variable in forecasting both the continuous and jump components of volatility (though with miniscule explanatory power), although the jump component exhibited shorter memory with respect to volume.

The central insight of this paper is the somewhat dismal fact that price jumps are both relevant and stochastic. They explain significant portions of price volatility, but exhibit no autocorrelation, cannot be predicted with a knowledge of past trading volume, and have no measurable effect on the volatility and volume proceeding or following their occurrence. Though many jumps can be linked to news events, smooth price changes can be linked to similar news events. Jumps in indices do not even correlate with jumps in the underlying assets. While these may appear to be frustrating findings, they indicate that the common practice of modeling price functions continuously is flawed and may be improved in the future.

The most obvious extension of the findings of this paper is to consider a larger body of common stocks. The legal ramifications for Altria Group as the leading producer in the cigarette industry potentially creates a unique jump profile, as evidenced by the real-world jump causes discussed previously. Extending the analysis to other common stocks would lend credibility to the results of this case study, as well as allow for a more thorough analysis of cross-asset jump characteristics. Also, a study of different indices could be performed to examine the relationship between an index's jump characteristics and the covariance of its comprising assets.

Extending the proximity effects experiment to include intraday measures might give a more definitive answer to whether or not the proximity effects of jumps are significant over any time period.

Finally, the most useful results that can come from this high-frequency data analysis are improved forecasting models, particularly for future volatility. If significant improvements can be made on models that only employ daily data by the use of high-frequency data and a consideration of jump characteristics, then many aspects of finance that depend greatly on volatility forecasting, notably derivative pricing, can be improved.

References

Andersen, Torben G. (1996). Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility. *Journal of Finance,* 51(1), 169-204.

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold. (2002). Parametric and Nonparametric Volatility Measurement. In Yacine Ait-Sahalia and Lars Peter Hansen (eds.), *Handbook of Financial Econometrics,* Amsterdam, North Holland.

Andersen, Torben G., Tim Bollerslev, and Francis 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, Inc.

Ballie, Richard T. and Tim Bollerslev. (2000). The Forward Premium Anomaly Is Not As Bad As You Think. *Journal of International Money and Finance,* 19(4), 471-488.

Barndorff-Nielsen, Ole and Neil Shephard. (2002). Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society,* 64(2), 253-280.

Barndorff-Nielsen, Ole and Neil Shephard. (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics,* 2(1), 1-37.

Cont, Rama and Ekaterina Voltchkova. (2005). A Finite Difference Scheme for Option Pricing in Jump Diffusion and Exponential Levy Models. *SIAM Journal of Numerical Analysis,* 43(4), 1596-1626.

Huang, Xin, and George Tauchen. (2006). The Relative Contribution of Jumps to Total Price Variance. *Journal of Financial Econometrics,* 3(4), 456-499.

Hull, John C. (2000). *Options, Futures, & Other Derivatives, 4th Edition*. New Jersey: Prentice-Hall.

Merton, Robert C. (1976). Option Pricing When Underlying Stock Returns Are Discontinuous. *Journal of Financial Economics,* 3(2), 125-144.

Tauchen, George and Mark Pitts. (1983). The Price Variability-Volume Relationship on Speculative Markets. *Econometrica,* 51(2), 485-506.

Table 1: Summary Statistics for Altria Group common stock for 1999-2005.

	RV	BV	RJ	5-minute Volume (in thousands)
Mean	3.262	2.676	0.142	77.42
St. Dev.	5.252	3.806	0.131	122.74
Min.	0.112	0.073	0	0
Max	103.4	58.7	0.792	22118.1

Table 2: Proximity Effects of Jump Days.

	BV	RJ	Daily Vol (Mill)
Mean			
Overall	*2.654*	*0.14*	*5.915*
Before Jump	3.35	0.174	6.06
After Jump	3.43	0.163	6.5
Median			
Overall	*1.633*	*0.118*	*4.991*
Before Jump	2.41	0.157	5.25
After Jump	2.49	0.149	5.34
Standard Deviation			
Overall	*3.833*	*0.129*	*3.583*
Before Jump	5.27	0.134	3.36
After Jump	3.96	0.1342	3.57
Proximity Effect (Aft - Bef)			
Day	0.0819	-0.0113	0.436
Week	0.1197	-0.012	0.445
Proximity Effect (% of Mean)			
Day	3.10%	8.10%	7.40%
Week	0.90%	1.70%	1.50%

Table 3: Volume vs. Variance Regressions.

$$Volume = \alpha + \beta * Variance$$

	Standard Model (Vol in Millions)			Logarithmic Model		
Variance Measure	Beta	C	R-squared	Beta	C	R-squared
abs(Ret)	154	4.17	0.289	0.184	16.39	0.209
SE	(12.1)	(.134)		(.010)	(.050)	
RV	4160	4.68	0.334	0.300	18.06	0.412
SE	(719)	(.211)		(.010)	(.085)	
BV	6330	4.35	0.405	0.313	18.22	0.419
SE	(745)	(.177)		(.010)	(.090)	
RJ (JM for Log Model)	2.03	5.74	0.005	0.1886	15.46	0.006
SE	(.912)	(.160)		(.0584)	(.0143)	

Table 4a: Volatility Forecasting: Jump Component, JM.

	$\log(JM_t)$		
α	0.1468	0.53590	1.8907
	(0.0082)	(0.04960)	(0.5117)
$\beta_{BV\text{-}d}$		-0.0486	-0.076
		(0.0113)	(0.0115)
$\beta_{BV\text{-}w}$		0.0350	0.0488
		(0.0142)	(0.0196)
$\beta_{BV\text{-}m}$		0.0610	0.0971
		(0.0137)	(0.0183)
$\beta_{JM\text{-}d}$	-0.0762	-0.0325	-0.0436
	(0.0367)	(0.0344)	(0.0334)
$\beta_{JM\text{-}w}$	0.1373	0.1285	0.1210
	(0.0435)	(0.0414)	(0.0397)
$\beta_{JM\text{-}m}$	-0.0008	0.0077	-0.0038
	(0.0409)	(0.0375)	(0.0374)
$\beta_{Vol\text{-}d}$			0.0799
			(0.0195)
$\beta_{Vol\text{-}w}$			-0.0444
			(0.0283)
$\beta_{Vol\text{-}m}$			-0.1098
			(0.0324)
R^2	0.0085	0.0672	0.0871

Table 4b: Volatility Forecasting: Continuous Component, BV.

	Log(BV_t)		
α	-0.8506	-0.8264	0.9122
	(0.1680)	(0.1699)	(1.5540)
$\beta_{BV\text{-}d}$	0.3080	0.3065	0.2062
	(0.0294)	(0.0293)	(0.0331)
$\beta_{BV\text{-}w}$	0.2511	0.2532	0.3854
	(0.0469)	(0.0470)	(0.0581)
$\beta_{BV\text{-}m}$	0.3550	0.3555	0.3516
	(0.0434)	(0.0435)	(0.0554)
$\beta_{JM\text{-}d}$		0.1992	0.1779
		(0.1140)	(0.1120)
$\beta_{JM\text{-}w}$		-0.0560	-0.0698
		(0.1232)	(0.1252)
$\beta_{JM\text{-}m}$		-0.2293	-0.2391
		(0.1117)	(0.1127)
$\beta_{Vol\text{-}d}$			0.2874
			(0.0565)
$\beta_{Vol\text{-}w}$			-0.3532
			(0.0822)
$\beta_{Vol\text{-}m}$			-0.0301
			(0.1010)
R^2	0.5911	0.5923	0.5994

Table 4c: Volatility Forecasting: Combined Volatility, RV

	Log(RV_t)		
α	-0.2690	-0.2490	3.2840
	(0.1642)	(0.1728)	(1.5820)
$\beta_{BV\text{-}d}$	0.3087	0.3074	0.2083
	(0.0287)	(0.0290)	(0.0327)
$\beta_{BV\text{-}w}$	0.2560	0.2582	0.3889
	(0.0459)	(0.0470)	(0.0593)
$\beta_{BV\text{-}m}$	0.3992	0.4002	0.4259
	(0.0413)	(0.0440)	(0.0564)
$\beta_{JM\text{-}d}$		0.1965	0.1616
		(0.1182)	(0.1149)
$\beta_{JM\text{-}w}$		-0.0428	-0.0690
		(0.1220)	(0.1229)
$\beta_{JM\text{-}m}$		-0.1732	-0.2027
		(0.1114)	(0.1116)
$\beta_{Vol\text{-}d}$			0.2794
			(0.0571)
$\beta_{Vol\text{-}w}$			-0.3408
			(0.0841)
$\beta_{Vol\text{-}m}$			-0.1333
			(0.1023)
R^2	0.6099	0.6107	0.618

Figure 1: Discontinuous Jump vs. Smooth Increase: Altria Group stock price on February 4, 2005 (top) and November 1, 2005 (bottom).

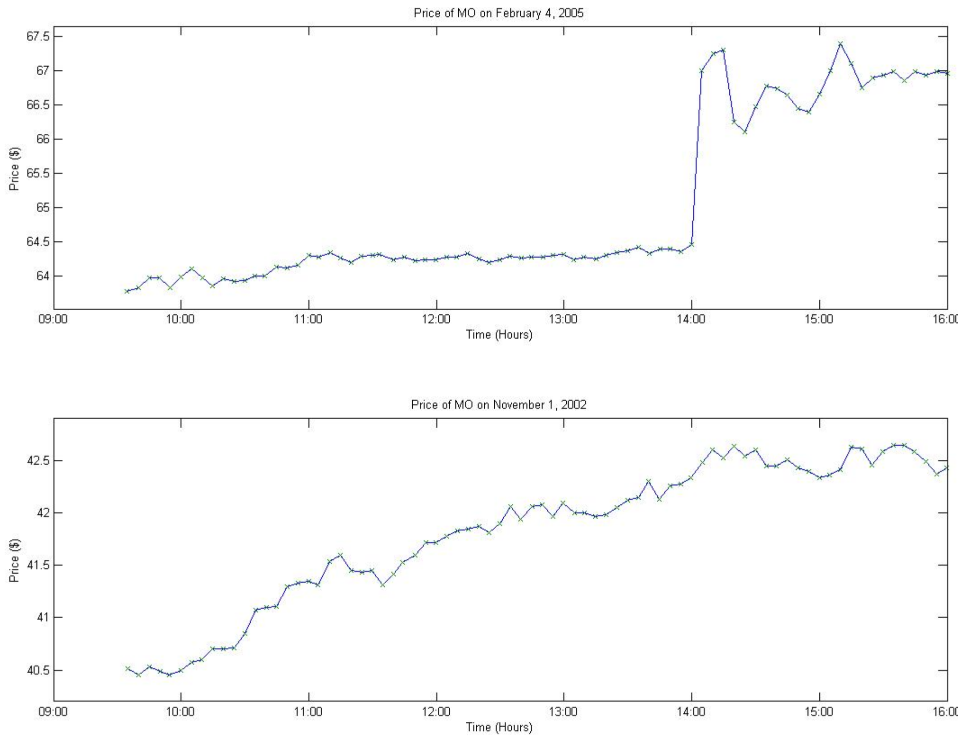

Figure 2: Jump Statistics for Altria Group common stock for 1999-2005

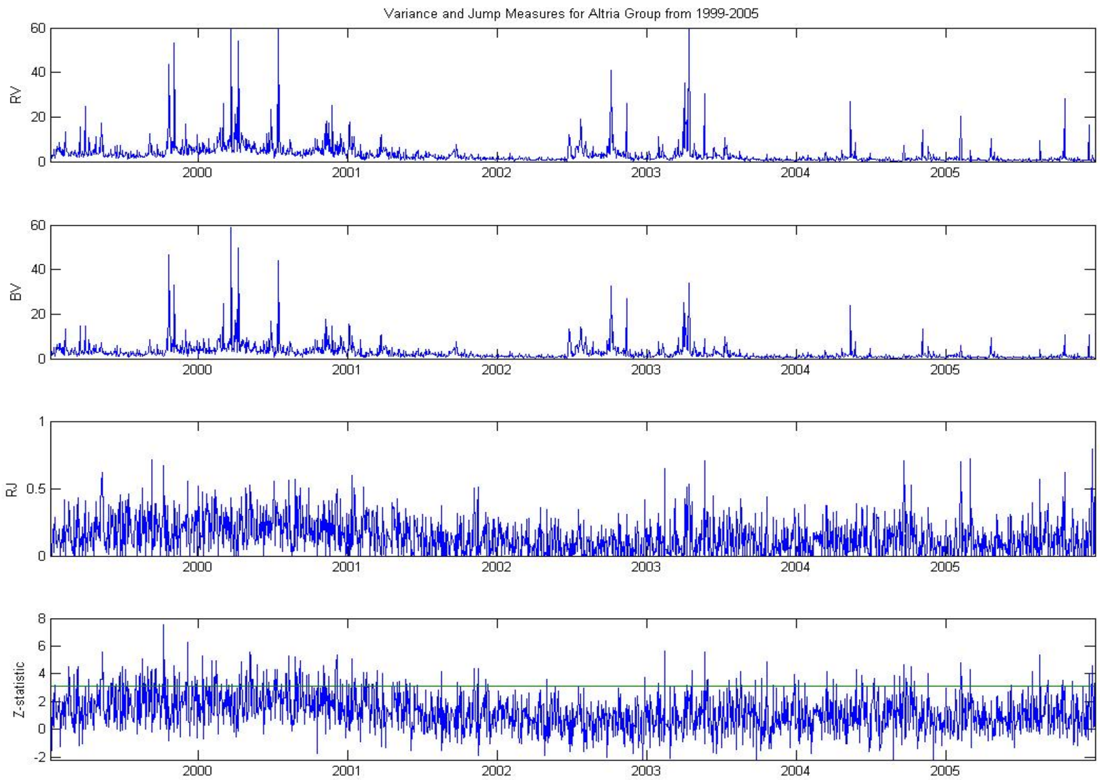

Figure 3: Distribution of daily z-statistics compared to the expected N(0,1) distribution.

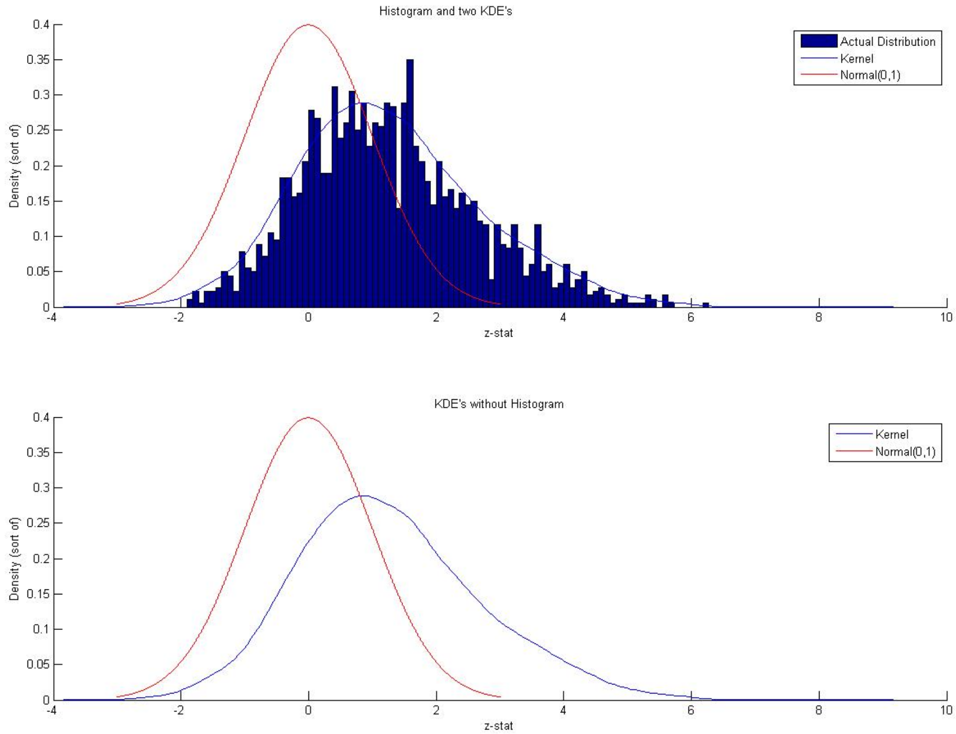

Figure 4: Jump Days: Altria Group Stock Price and Share Volume on February 4, 2005 (left) and August 18, 2005 (right).

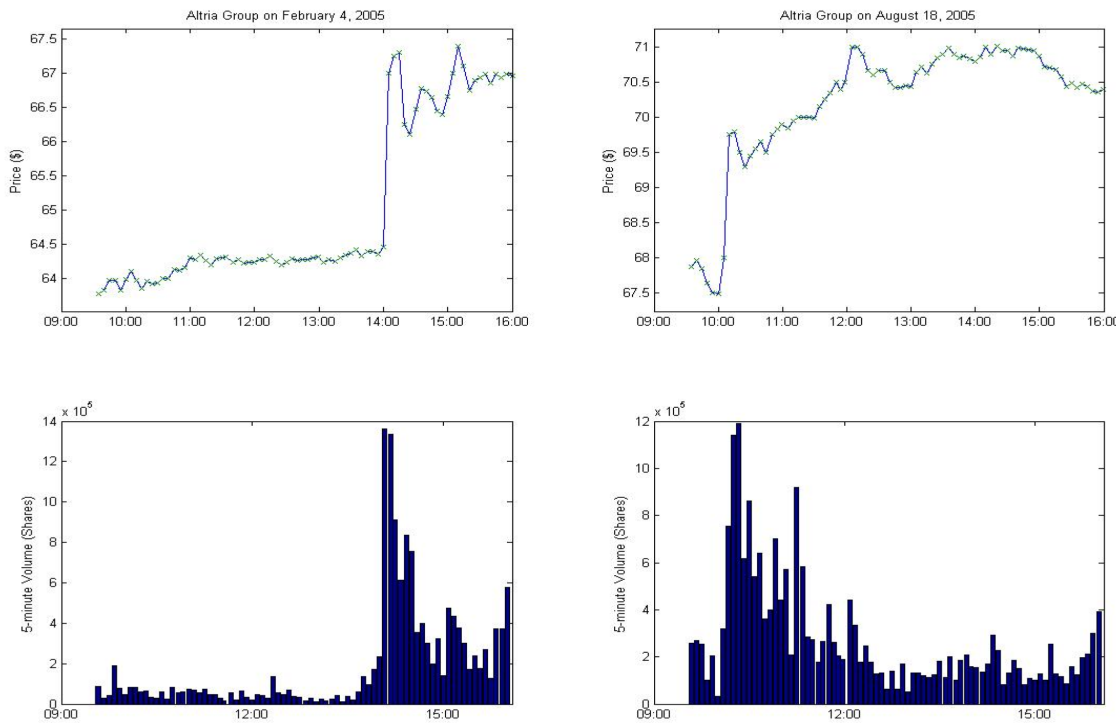

Figure 5: Smooth Days: Altria Group Stock Price and Share Volume on October 21, 2002 (left) and May 24, 2004 (right).

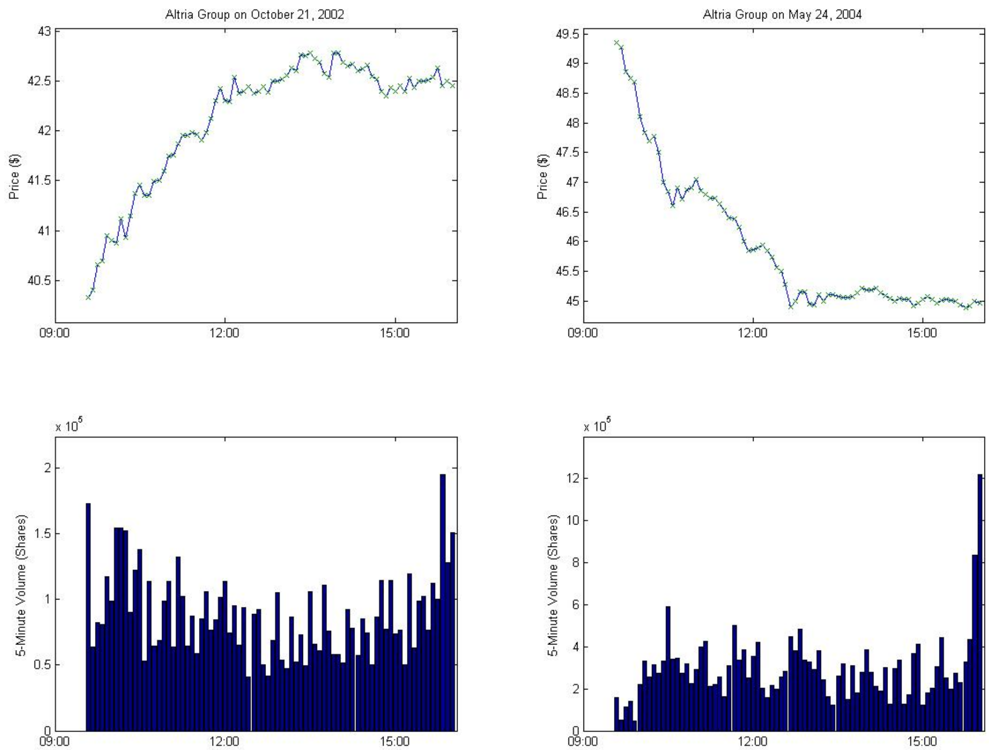

Figure 6: Daily z-statistics of Altria Group vs. daily z-statistics of S&P 500 futures index for 1999-2002.

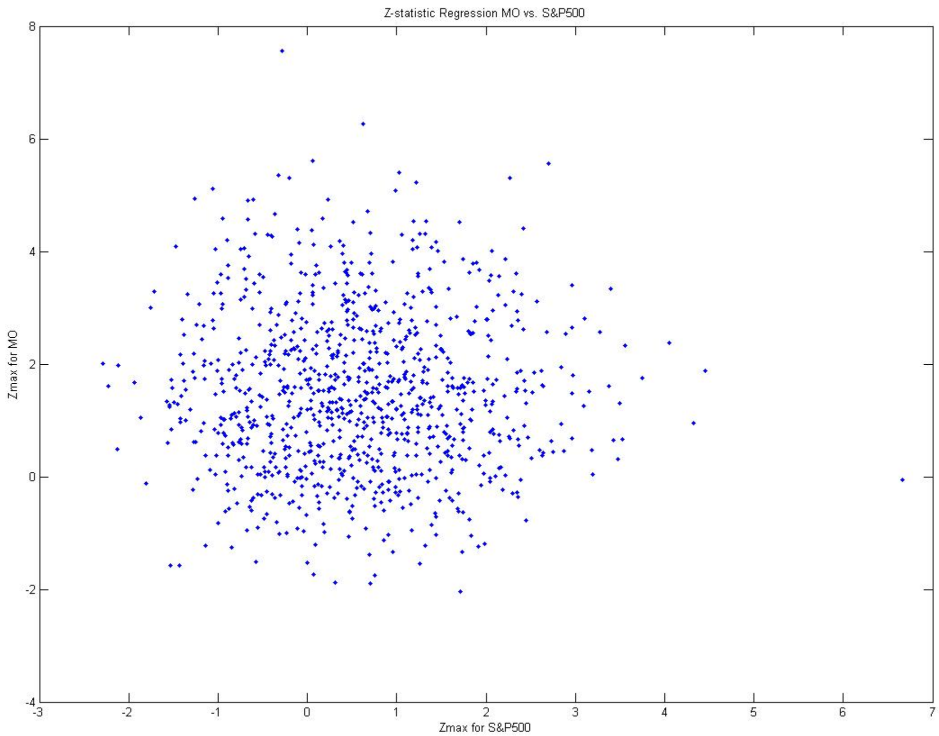

Figure 7: Price charts for SPY, AIG, WMT, MO, IBM, and PG on January 3, 2001.