

**Testing the Relationship between Oil Equities and Oil Futures:
A Look at Returns, Jumps, and Volatility**

Brian Northrop Jansen
Professor George Tauchen

The Duke Community Standard was upheld in the completion of this paper. Honors Thesis to be submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University.

Duke University
Durham, North Carolina
2008

1. Introduction

July 15, 2008 was a historic day for petroleum markets in the U.S. On the New York Mercantile Exchange, crude oil opened at \$146 per barrel, the highest price ever for the commodity. After opening, the price dropped \$6.44 per barrel, the largest single-day decline in over 17 years. July 15th also marked the Chicago Board Options Exchange's release of a tradable oil volatility index, the OVX, to track rapid fluctuations in the highly liquid commodity. The summer of 2008 tested new limits for oil markets--oil prices could in fact breach \$100 per barrel and prices above \$100 could be sustained over months. While oil eventually retreated from above \$100 per barrel in September 2008, the excitement around oil remains. The unpredictability of oil prices left everyone wondering what impact oil prices will have on the future of the economy.

After the oil market turmoil of high prices and extreme volatility had subsided, Exxon Mobil reported third quarter profits of \$14.83 billion, breaking its own record set the previous quarter (Larson, 2008). This result suggests that oil company earnings are highly linked to the price of oil. Most media outlets credit their increased profits to the higher prices that companies charge when crude petroleum prices increase. However, this is an over-simplification of the oil industry. This claim is supported by the fact that oil corporations own significant amounts of oil and thus face constant pumping costs per barrel. On the other hand, petroleum corporations such as Exxon Mobil, ConocoPhillips, and Chevron refine more crude petroleum than they lift out of the ground (EIA, 2008). Consequently, a portion of their input commodity must be bought through contracts that reflect changes in crude prices, increasing the cost of goods sold. These two drivers of oil company profits challenge the direct relationship that many have assumed exists between oil company earnings and crude petroleum prices. The correlation between the returns of petroleum equities and the returns of oil prices is only one element a more complex relationship between the two assets.

This paper provides a step towards explaining the connection between oil futures prices and the oil equities market in three ways. First, this paper tests the hypothesis that changes in the price of oil affect returns of petroleum companies through investigating simultaneous movements in oil future and oil equity returns. Second, simultaneous price jumps (extreme movements) in the oil equities and futures market are investigated with the hypothesis that both assets are responding to announcements that affect oil price expectations. Third, the volatilities

are compared across markets, as one might expect that volatile trading of oil futures would trigger trading of oil equities. Throughout each step of the analysis, a portfolio of companies in the S&P 100 is also considered to provide a context for the results found in the oil sector.

The main findings are that oil company asset returns are driven not only by an oil factor that describes fluctuations in crude prices, as might be expected, but also by a market factor, present across the entire portfolio of stocks included. Price jumps, on the other hand, appear to be occurring simultaneously across the oil equity sector, but discordant with jumps in the oil futures market. Similarly, the volatilities across the entire set of equity assets are highly correlated, but the oil futures market does not follow the same volatility pattern. This paper emphasizes that a strong connection between the oil futures and the oil equities market exists for returns. However, the differences in the types of assets and the motivations of traders between the equity and commodity futures market generate starkly different volatility schemes.

The current literature exploring oil markets does not examine the effects of the oil futures market on the oil equities market. Certain aspects of the oil futures market have been extensively studied. The highly volatile nature of oil futures prices led Kang (2008) and Marzo (2007) to investigate volatility forecasting in that market. Further, volatility in oil futures prices does not correlate with oil futures returns (Wang, Wu, and Yang 2008). Moosa and Al-Loughani (1994) look at the ability of oil futures to predict spot prices. Litzemberger and Rabinowitz (1995) examine the phenomenon that spot market prices are usually higher than the futures market prices at expiration.

To add to the literature on oil markets, this paper analyzes the impact of changing oil futures returns, volatility, and price jumps on individual equity returns, volatilities and price jumps. Current research exists regarding the effect of oil shocks on macroeconomic activity. Hamilton (1983, 1985) concludes that a strong causal relationship between oil prices and U.S. Gross National Product exists. Mork, Olsen and Mysen (1994) apply the strong connection between macroeconomic variables and oil prices to seven OECD countries. Jones and Kaul (1996) find a significant effect of oil price movements on aggregate stock returns by accounting for resulting changes in aggregate cash flows. Driesprong, Jacobsen and Maat (2003) determine that oil prices are economically significant in predicting stock market returns in developing countries. Sadorsky (1999) finds a significant correlation between stock index returns and oil shocks at the monthly level.

This paper is unique in that high-frequency data from specific stocks are analyzed instead of simply using aggregate indices that mimic equity markets. Further, the highly efficient oil futures market data is included, rather than using a producer price index for oil. Additionally, a study by Andersen, Bollerslev, and Diebold (2006) suggests that the continuous model of asset pricing does not hold when applying the model to real data. Thus, an analysis of simultaneous jumps in oil futures prices and stock prices should improve the understanding of oil shocks. To further identify the transfer of information between markets, the paper looks at spillovers in volatility from the two markets. Ross (1989) claims that price volatility can be an accurate measure of the rate of information flow in a financial market. If oil equities and oil futures react to similar information regarding expectations about future supply and demand of oil, their volatility patterns should also be similar.

Because one of the goals of the paper is to identify simultaneous movements across markets and the transfer of information between markets, an understanding of intraday financial metrics is required. Therefore, this paper uses a jump test developed by Lee and Mykland (2008) that is capable of identifying a jump statistic and a volatility measure for every observation throughout the trading day. Preliminary results of the Lee-Mykland test suggest that their method of jump detection may be biased towards flagging too many jumps in the morning and against identifying jumps in the afternoon. This paper offers and utilizes an adjustment to their test that appears to correct for this concern.

This paper will proceed as follows: Section 2 highlights some important practices of the oil industry that drive later results; Section 3 describes the diffusive jump model to explain asset price movement; Section 4 introduces the relevant statistics used in the financial econometrics literature, it describes the Lee-Mykland jump test, and it introduces factor analysis as a method of describing the data; Section 5 describes the data used; Section 6 explains the main results and Section 7 concludes the paper.

2. Practices of the Oil Industry

Oil is traded in three different ways: the spot market, the futures market, and through over-the-counter agreements. Spot market prices are determined through the transaction between suppliers and refiners when tankers are four days away from port and the oil is delivered immediately. The large majority of oil is traded through the futures market, with contracts

ranging in delivery date from one to 18 months in the future. This market is highly liquid and provides the most current information about the price of oil. The price of oil traded under over-the-counter contract agreements is determined by the spot and futures market prices. For the purposes of this paper, it is assumed that movements in the six-month oil futures market directly follow spot prices, although with lower volatility. According to Haubrich (2005), empirical evidence indicates that oil futures actually contain little information about the forward curve of oil. The futures price does not in fact contain any information beyond what that current spot price is, storage costs, and interest rates.

There are a couple of notable practices that connect oil company profits and the oil futures market. The first is the process of exploring for, drilling, lifting, and shipping crude oil. Increasing oil prices are a boon to all aspects of this business. As the spot market price of oil increases, petroleum corporations directly own or have long-term contracts to extract the now more highly valued resource. Exxon Mobil, for example, has 22 billion barrels of proven reserves and a presence in 200 countries that makes exotic exploration possible (Larson, 2008). Firms have the opportunity to sell their oil at spot market prices as oil is shipped into the U.S. Additionally, the valuation of an oil company is dependent on the firm's current and future assets, of which oil reserves account for a significant portion. Rising profitability across the oil industry can spark profitable asset sales for large oil companies (Martin 2006).

The second practice of an oil company is refining, encompassing all processes involved with transforming crude oil into a delivered, usable end-product. This function of an oil company is negatively impacted during oil price increases. The primary reason is that oil is an input cost for the business and the higher cost cannot be completely transferred to consumers. The highly competitive oil industry results in declining profit margins per gallon from refining activity as the price per gallon increases (EIA). However, there are several factors that mitigate shrinking refining margins. The primary mitigating factor is that refining operations only purchase a portion of their oil at market rates. For example, approximately 55% of the oil refined by Exxon is purchased on the spot or futures market subject to Nymex prices. The remaining 45% has been lifted from Exxon's own oil reserves (EIA). Thus, oil companies are partially hedged against rising oil prices when buying on the spot or futures market. Another mitigating factor is the strength of the consumer market. Cooper (2003) found that the price elasticity of demand for crude oil was -0.05 in the U.S., indicating the demand curve's

insensitivity to price changes. This has been true even during the historically high crude prices of 2007 and 2008 (EIA).

Although the refining component of an oil company is negatively affected by increasing oil prices, the positive relationship between lifting activity and oil prices outweighs that effect and leads to overall higher profits. This translates to higher valuations of oil companies in equity markets and a positive relationship between the expected future price of oil and oil equity returns. While this is a simplistic model of an oil company, it accounts for the main drivers of revenue and costs. Equity analysts confirm this connection between crude oil prices and expected company earnings. Recent windfall profits by Exxon Mobil, ConocoPhillips, and Chevron empirically indicate this connection during the 2005-2007 period of steadily increasing oil prices.

It is important to note that oil companies do not participate in oil markets for financial arbitrage. Instead, they are active members of the markets only to secure accurate delivery and pricing of oil as it impacts their refining operations.

3. Returns Model

Before analyzing real data, it is necessary to understand the theoretical model of asset price movements and market microstructure noise. Section 3.1 introduces the jump diffusive model of asset pricing on which the jump and volatility measures described in Section 4 are based. Section 3.2 explains how the presence of market microstructure noise in high-frequency stock quotes influences the sampling interval used in jump statistics and volatility metrics.

3.1 Model of Returns and Volatility

The foundational model for this paper is a stochastic differential equation that describes the evolution of stock price. The model to represent asset log-price movement was discussed by Merton (1971) and is described below:

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

In equation (1), the stock log-price movement, $dp(t)$, is explained in terms of a time-varying drift component $\mu(t) \cdot dt$ and a time-varying volatility component $\delta(t) \cdot dW(t)$, where $dW(t)$ introduces randomization through standard Brownian motion. Including a Brownian motion

component is the result of summing independent identically distributed log-returns over infinitesimally small time periods. When the log-price summations are viewed over a longer time period, the model allows the stock price to take on a drift-like movement. This model accounts for the continuous component of stock price movement, because standard Brownian motions follow continuous sample paths with probability one.

Recent financial literature, by Barndorff-Nielsen and Shephard (2004, 2006) in particular, suggests that the continuous model does not fully capture the dynamics of asset returns and that discontinuities represent an essential element in stock price evolution. It is important to note that identifying jumps in stock prices is a difficult conceptual and practical task. The argument could be made that all price movement is discontinuous as assets are quoted at discrete prices. However, allowing prices to be modeled continuously provides many useful results. The problem remains that when there are sharp price movements, the continuous model is not appropriate. Therefore, Merton (1976) developed a jump diffusion component to account for such discontinuities. The log-price $dp(t)$ is now defined as

$$dp(t) = \mu(t) \cdot dt + \delta(t) \cdot dW(t) + \kappa(t) \cdot dq(t). \quad (2)$$

The third term in the log-price model (2), $\kappa(t) \cdot dq(t)$, accounts for potential discontinuities or jumps in the asset price movement, where $\kappa(t)$ represents the magnitude of the jump and $dq(t)$ is a counting process. This model provides the foundation for the study of price jumps throughout the remainder of the paper and underlies the statistics described in Section 4.

3.2 Market Microstructure Noise

When analyzing asset prices at high frequencies, special attention must be paid to the effects of sampling the data at different intervals. There are two potential complications of assigning the proper price to any given asset traded in a market. The first concern addresses the problem of determining an asset price given the cyclical, speculative, and random-walk nature of financial products. The second complication in asset pricing is the microstructure noise of observed price quotes. This problem has deeper implications for the paper and can be defined as deviance of observed prices from their true values. As sampling frequencies increase, microstructure noise has an increasing effect on the reliability of observed prices. This paper

maximizes the use of accurate information by reducing sampling frequencies from their highest levels and therefore reducing microstructure noise's effect on the data.

The market microstructure noise can be explained through summing the discounted value of future dividend payments and thus determining the theoretical price of a stock. Given the constant dividend payout ratio assumed by the constant dividend growth model outlined by Levy (2005), the expected value of the dividend payout in year i is

$$E(D_i) = D_0(1 + g)^i = E_0d(1 + g)^i, \quad (3)$$

where g is the constant yearly growth rate, E_0 is the current annual dividend earnings per share and d is the equity's constant dividend payout ratio. Once the equity's future earnings are determined, the "fundamental" price is

$$P_0 = \sum_{i=0}^{\infty} \frac{E(D_i)}{k^i} = \frac{D_0(1 + g)}{k - g} \quad (4)$$

where k is the required rate of return as determined by the Capital Asset Pricing Model (CAPM) (Levy 2005). The CAPM is a method of pricing an asset based on its exposure to the systemic risk of the market and the risk-free rate of return offered by the market. The value P_0 is the theoretical true price of the equity and it follows the changes in dividends, growth rates, and required rate of return.

Microstructure noise confounds the observation of an equity's true price. When accounting for microstructure noise, the observed log-price at time t becomes

$$\ln[P^*(t)] = \ln[P(t)] + \varepsilon_t \quad (5)$$

where $P(t)$ is the fundamental price as described in equation 4 and ε_t is microstructure noise. The errors in equation 5 are uncorrelated, have unit variance, and mean zero. The noise in the observed price comes primarily from the bid-ask bounce and various other trading anomalies at the high-frequency tick level. The slight deviations from the true price must be accounted for in determining sampling frequencies and window sizes for the jump statistics explained in Section 4.

4. Statistical Methods

4.1 Barndorff-Nielsen Shepard Statistics

The measures of variance used in this paper are heavily influenced by Barndorff-Nielsen and Shephard's (2004, 2006) development of two volatility measures: Realized Variance and

Bipower Variation. Realized Variance converges to the daily volatility of both the continuous and jump components of asset price variation. Bipower Variation is robust to jumps and converges to the volatility of the only the continuous price movement. The paper defines geometric returns $R(t_i)$ at time t_i for an asset with price $P(t_i)$ as

$$R(t_i) = \log \left[\frac{P(t_i)}{P(t_{i-1})} \right]. \quad (6)$$

The volatility statistics and their asymptotic properties are defined below:

$$RV_t = \sum_{i=2}^n |R(t_i)|^2 \xrightarrow{n \rightarrow \infty} \int_0^T \sigma^2(v) dv + \sum_{i=1}^n \kappa^2(t_i) q(t_i) \quad (7)$$

$$BV_t = \frac{\pi}{2} \left(\frac{n}{n-1} \right) \sum_{i=3}^n |R(t_i)| |R(t_{i-1})| \xrightarrow{n \rightarrow \infty} \int_0^T \sigma^2(v) dv. \quad (8)$$

The limits in equations (7) and (8) indicate that as n approaches infinity, the interval between the returns approaches zero but the time period between the first and last observation is held constant. In equation (7), Realized Variance converges to the daily variance that includes both a continuous component and a jump component, while the Bipower Variation limit in equation (8) only describes an asset's continuous daily variance. The reason for this is that Bipower Variation multiplies adjacent returns and the effects of a single large return $|R(t_i)|$ are mitigated by the smaller returns $|R(t_{i-1})|$ and $|R(t_{i+1})|$. The jump detection test that Barndorff-Nielsen and Shephard developed utilizes RV and BV to isolate the variance due to jumps in prices and compares the size of the jump to the continuous component of daily variance (2005). This paper does not utilize the Barndorff-Nielsen Shephard jump test because the test can only identify days that contain jumps, not individual observations that are jumps.

4.2 Lee-Mykland Test Statistics

The Lee-Mykland (2008) jump test uses a similar methodology for jump detection as the Barndorff-Nielsen Shephard test, but has the additional capability of identifying specific returns that can be classified as jumps. After knowing the exact time of a jump, it is possible to identify simultaneous jumps occurring in the equities and oil futures market. The Lee-Mykland jump test creates a standardized return value $L(i)$ at every observation with time t_i ,

$$L(i) = \frac{R(t_i)}{\sigma(t_i)} \quad (9)$$

$$\sigma^2(t_i) = \frac{1}{K-2} \sum_{j=i-K+2}^{i-1} |R(t_j)| |R(t_{j-1})|. \quad (10)$$

The $\sigma(t_i)$ statistic is a scaled version of the Bipower Variation used in the Barndorff-Nielsen and Shephard literature. Lee and Mykland recommend using 5-minute data to reduce the effect of microstructure noise on high-frequency price quotes. By using K equal to 270, the window size of the test is 3.5 days when using 5-minute sampling of the data. That is, the test generates a local volatility measure by calculating the Bipower Variation of the previous 3.5 days of returns. Therefore, the $L(i)$ statistic is a ratio of the return scaled by the local, continuous volatility $\sigma(t_i)$ that is robust to discontinuous or “jumpy” behavior. Lee and Mykland then identify the following statistic as having an exponential distribution:

$$\Lambda_i = \frac{\max_{i \in M} |L(i)| - C_n}{S_n} \rightarrow \xi, \quad M = (1, 2, 3, \dots, n) \quad (11)$$

where

$$P(\xi \leq x) = \exp(-e^{-x}) \quad (12)$$

and

$$C_n = \frac{(2 * \log n)^{1/2}}{c} - \frac{\log(\pi) + \log(\log n)}{2 * c(2 * \log n)^{1/2}} \quad (13)$$

$$S_n = \frac{1}{c(2 * \log n)^{1/2}} \quad (14)$$

$$c = \sqrt{2/\pi}. \quad (15)$$

The maximum absolute value of the adjusted return $L(i)$ over the observations in a day is shifted by C_n and scaled by S_n , both determined under the normal distribution of the $L(i)$ statistic (Lee and Mykland, 2008). The resulting statistic Λ_i is exponentially distributed. Further, the test is run at the 1% significance level for rejecting the hypothesis that there is no jump at time t_i .

4.3 Lee-Mykland Test Adjustment

After performing the jump tests on three oil stocks, it was evident that the Lee and Mykland jump test method may be biased, as the results presented in Section 6.1 confirm. The histograms presented in Figure 1 display the intraday jump behavior of three oil stocks. These figures indicate that jumps occur disproportionately during the time period from 9:35 to 10:35. Van Tassel found the same intraday pattern in the distribution of jumps using the slightly altered Lee-Mykland test proposed in 2006 (Van Tassel, 2008). To investigate this phenomenon further,

the average Bipower Variation for each observation time across the entire dataset is calculated. This statistic is defined as

$$\overline{BV}_h = \frac{1}{N} * \sum_{b=1}^N \sqrt{R_{b,h-1}} * |R_{b,h}| * \sqrt{R_{b,h+1}}, \quad (16)$$

where $R_{b,h}$ is the return of an observation for the b^{th} day and the h^{th} time within that day and N is the number of days in the sample. The U-shaped pattern of \overline{BV}_h for both the futures and the equities is noticeable in Figure 2. This pattern is cause for concern over the validity of the Lee-Mykland test. The instantaneous volatility $\sigma(t_i)$ is assumed to be flat throughout the window size, which spans 3.5 days when using 5-minute data. However, with a morning volatility 2.5 times greater than the afternoon volatility, the Lee-Mykland test is biased towards incorrectly finding more jumps in the morning and failing to flag jumps in the afternoon.

To adjust for the biased jump test, the following method is proposed: to scale the individual returns by their corresponding \overline{BV}_h value and then re-perform the Lee-Mykland test on the adjusted returns

$$\hat{R}_{b,h} = \frac{R_{b,h}}{\overline{BV}_h}. \quad (17)$$

In this method, the instantaneous volatility reflects the local volatility, but also compensates for the intraday volatility pattern. When performing this test, first and last returns of the day are omitted, as there are no \overline{BV}_h calculations for these returns. The results from this Lee-Mykland correction are presented in detail in Section 5.1. The rest of the paper utilizes the following

statistics: adjusted returns $L(b,h) = \frac{\log(\hat{R}_{b,h})}{\sigma_{b,h}}$, the jump statistic $|L(b,h)|$, and instantaneous

$$\text{volatility } \sigma_{b,h}^2 = \frac{1}{K-2} * \sum_{j=i-K+2}^{i-1} |\hat{R}_{b,h}| * |\hat{R}_{b,h-1}|.$$

4.4 Factor Analysis

The primary method of data analysis presented in this paper is factor analysis. In contrast to using regression tools, factor analysis does not presuppose a single dependent variable and a group of explanatory variables. Instead, factor analysis extracts a number of unobserved factors that explain the common variance across concurrent observations of a set of variables (Gorsuch

1983). If markets are efficient and a connection between oil futures prices and the stock market exists, then changes in price, volatility, and jumps should occur simultaneously on average. If this did not hold, after accounting for the bid-ask transaction cost, then an arbitrage situation would arise. Therefore, the analytic structure of this study requires factor analysis to examine contemporaneous movements in stock returns, jumps, and volatilities.

The process of extracting factors for an individual variable can be described by the following equation:

$$z_j = a_{j,1} \cdot F_1 + a_{j,2} \cdot F_2 + \dots + a_{j,m} \cdot F_m + \varepsilon_i, \quad (j=1,2,\dots,n) \quad (18)$$

where the observed data for variable j is represented by z_j . Further, z_j is described as a linear combination of the common factors F_m and an error term ε_i . The coefficients of the factors, $a_{j,m}$, are referred to as the loadings. The common factors account for the correlations among the variables and the error term captures the remaining variance of that variable (Harman 1967). Additionally, the extracted factors are orthogonal and uncorrelated. To allow for multiple observations, equation (18) has been generalized to:

$$z_{j,i} = \sum_{p=1}^m a_{j,p} \cdot F_{p,i} + \varepsilon_{j,i} \quad (i=1,2,\dots,N; j=1,2,\dots,n) \quad (19)$$

In equation 19, $z_{j,i}$ represents the company variable being explained, either returns, jumps, or volatility, while $a_{j,p}$ and $F_{p,i}$ are the factor loadings and factors. The coefficients $a_{j,p}$ are transformed to be able to represent the correlation between $F_{p,i}$ and $z_{j,i}$ and thus do not have unit variance. The errors are uncorrelated with unit variance and mean zero. Equation (19) is capable of explaining observations across multiple companies. Factor analysis is only meaningful when the number of factors m is low. That is, the researcher can explain a large quantity of data with only a few key factors.

To simplify the interpretation of the factor analysis results, a rotated factor matrix can be created. In this process, the new rotated factors are linear combinations of the previously preserved factors. Thus, the variables are described by the linear combination of the old factors. The rotated factor-loading matrix forces the loadings to be close to either zero, one, or negative one, as each company variable is intended to be either correlated highly or not at all with the factor combination. With rotated factor loadings, the same number of factors is retained and the same degree of explanatory power is preserved (Gorsuch 1983).

There are two key tests of validity for the extracted factors: statistical significance and practical significance. First, the retained factors themselves need to be statistically significant and explain a meaningful portion of the variance. There are several competing ideas about the correct number of factors that should be retained. Kaiser (1960) developed a criterion where only the factors with an eigenvalue, a scalar multiplier of an eigenvector that indicates magnitude, greater than one should be retained. Further research by Cattell (1966) indicates that factors should be retained until any additional factors would only explain a sharply reduced percentage of the variance. For this study, only the factors with an eigenvalue above 0.2 are retained for the factor-loading matrix. Any non-included additional factors explained a significantly smaller portion of the variance. As will be discussed in Section 6, two extracted factors were found to be significant in the factor analysis results.

To determine the statistical significance of specific factor loadings, Stevens (2003) describes testing the loadings using standard errors calculated by doubling the standard error for a standard Pearson correlation, $\frac{1}{\sqrt{n-1}}$, where n is the number of observations in the sample and the test is performed at the $\alpha = .01$ level from standard normal distribution. In practice, a sample size greater than 1,000 must have a correlation above 0.162 to be determined significant. As this dataset contains approximately 65,000 observations, factor loadings, or correlations, above 0.162 are statistically different from zero. To determine practical significance, Hair et al. (1998) suggest the correlation thresholds of $\pm 0.3, \pm 0.4, \pm 0.5$ to determine minimally, moderate, and significant practical interpretation. These thresholds were chosen through their historical work—using their field expertise to determine what the expected number of practically significant factors should be and comparing the correlation. Intuitively, a correlation of ± 0.3 indicates that 9% of the variation in the dependent variable is explained by the factor. A correlation smaller than ± 0.3 would not explain enough of the variation to be determined practically significant.

4.5 Co-Jump Correlation

The method of factor analysis is not the only way to identify common jumps or jump processes across the oil and equity markets. Roeber (1993) argues that the oil futures market responds very rapidly to changes in price expectations. Market participants include oil companies, refining companies, electric utilities, traders, and investors, creating a well-

functioning and efficient market. However, following news announcements regarding oil or oil expectations, the oil futures market may respond faster to the information than the stock price of an oil company might. Therefore, to account for the potential lag time in the equity market response, this study uses a correlation-type statistic, comparing days that contain a jump rather than specific returns that contain a jump, giving the equity market the remaining time in the trading day to react. This statistic is defined as

$$\rho_{a,b} = \frac{C_{a,b}}{\sqrt{J_a * J_b}} \quad (20)$$

where $C_{a,b}$ is the number of common days between stocks a and b that contain at least one jump as determined by the adjusted Lee-Mykland test and J_a is the number of days that contain at least one jump for stock a .

5. Data

The high-frequency data for the eleven equities traded on the New York Stock Exchange were obtained from the commercial data source price-data.com. Price quotes from the oil futures traded on the New York Mercantile Exchange (Nymex) were obtained from tickdata.com. The paper refers to the oil futures price with the shorthand OIL and the equities are quoted by their ticker symbol.

As mentioned in Section 3.1, to test and adjust the Lee-Mykland jump detection method, three oil stocks are analyzed, ExxonMobil (XOM), ConocoPhillips (COP), and Chevron (CVX), over the period of 9/3/2002 to 1/24/2008, for a total of 1343 trading days. The dataset has 385 observations per-day, recorded at the 1-minute frequency, with each day beginning at 9:35 AM and ending at 3:59 PM. For the Lee-Mykland test, the data was sampled at the 5-minute frequency, creating 77 observations per day. However, the first observation of the day was not included in my test as it contained an overnight return that was not a part of the intraday price movement.

The dataset used in the factor analysis section includes OIL (oil futures), the three refining oriented oil company stocks used above, an upstream drilling company, and 7 other equities with different relationships to petroleum. The results from the factor analysis are grouped to reflect oil's potential relationship on the company's performance. The first group

regards the companies directly related to lifting and selling oil: XOM, COP, CVX, and Baker Hughes (BHI). BHI is an oilfield services company dealing with lifting and drilling services. The second group contains companies with services or products related to the price of oil, including Entergy (ETR), Ford (F), and FedEx (FDX). ETR is a consumer utilities company with primary operations in electricity delivery originating from nuclear power plants. ETR's product is potentially a substitute for electricity originating from an oil-fueled plant. Ford products are heavily reliant upon consumer petroleum prices as they concentrate in pickup and industrial truck production and have been slow to develop smaller or hybrid powered cars. With the largest air cargo fleet in the world and a formidable fleet of ground trucks, FedEx is an enormous consumer of oil. The final grouping includes companies whose products are less related to the price of oil: Goldman Sachs (GS), Proctor & Gamble (PG), Boeing (BA), and Dell (DELL).

Because oil futures trade on the Nymex between 10:00 a.m. and 2:30 p.m., the equities data was truncated to this time interval. The Nymex light, sweet crude oil futures data contains 54 observations per-day when sampled at the 5-minute frequency, with an open at 10:00 a.m. and close at 2:30 p.m., for the majority of the days dating. The OIL and equities datasets contained a small number of incongruous days, where the trading time was either shortened or non-existent. The time period of the data spans from 9/3/2002 to 1/24/2008, with 1284 complete trading days 1284 and 65,484 individual observations. The adjusted returns $L(b,h)$, the jump statistic $|L(b,h)|$, and instantaneous volatility $\sigma_{b,h}^2$ are constructed from this formatted data for each observation in accordance with the theory described in Sections 3 and 4. The oil contracts are six-month-forward futures contracts for light, sweet crude. Light, sweet crude is the highest grade of unrefined oil and comes primarily from gulf oil reserves such as the Ghawar field in Saudi Arabia via tankers.

6. Results

The first portion of this section describes the results of the adjusted Lee-Mykland jump test and potential explanations for the pattern of intraday volatility. The second portion of the section applies the adjusted Lee-Mykland test, and the statistics it utilizes, to the much larger dataset, described above. Sections 6.3-6.5 use factor analysis to examine the patterns in adjusted returns, jump statistics, and instantaneous volatilities for OIL and the eleven equities. Section 6.6 presents the Co-Jump Matrix of common jump days across OIL and the eleven equities.

6.1 Lee-Mykland Test Results

By adjusting for the intraday pattern in Bipower Variation, the Lee-Mykland test used in this paper flags an even number of jumps throughout the trading day. The results discussed in Section 4.3 explore the shortcomings of the Lee-Mykland test and a potential correction to the test. The most important finding was that the Lee-Mykland test assumed that intraday Bipower Variation was flat, while the data indicates that volatility is 2.5 times higher in the morning than in the afternoon. The uncorrected, intraday average Bipower Variation \overline{BV}_h results are shown in Figure 2, with OIL as a solid line and the average of XOM, COP, and CVX as an asterisk line. The diamond line is the \overline{RV}_h average of XOM, COP, and CVX and is constant throughout the day. The graph indicates that the intraday volatility is not constant, but U-shaped. There is a clear spike in the daily average Bipower Variation starting at 10:30 and remaining in the 10:35 and 10:40 averages. The \overline{BV}_h then trails off after 10:40, until it rises slightly again in the last half hour of the trading day. The most likely explanation for the sharp 10:30 increase is that the Department of Energy's Weekly Petroleum Status Report is released at 10:30 every Wednesday. This report contains all the major oil measures including world prices, refinery activity, net imports, oil reserves, and products supplied. Van Tassel (2007) found that within the equity market, intraday Bipower Variation was higher in the morning and lower in the afternoon due to the higher frequency of morning news announcements and higher volume of market transactions. This paper confirms that result. The oil industry's intraday volatility pattern is most likely subject to similar processes, but is additionally affected by the importance of the Department of Energy announcements on oil expectations and their corresponding effect on oil futures and companies.

The problematic jump detection method of the original Lee-Mykland test appears to be corrected. The results of the adjusted Lee-Mykland test are displayed in a histogram in Figure 3 and present jumps evenly distributed throughout the day. Figure 4 applies a kernel density to the combined, adjusted jumps. The histogram in Figure 4 indicates that jumps are evenly distributed throughout the day. However, the Kernel density identifies the subdued U-shaped intraday pattern. This remains an intuitive result as morning announcements should still cause more discontinuities around 10:30 a.m. and 11:00 a.m. The new \overline{BV}_h graph is not included, but the

volatility has variance zero and a mean of one throughout the day, as it should by construction. After adjusting the returns to reflect their sample average Bipower Variation, the Lee-Mykland test correctly assumes that instantaneous volatility is constant throughout the window size of 2.5 days. Therefore, there is no longer a bias towards over-identifying jumps in the morning and under-identifying jumps in the afternoon. The test continues to reflect global periods of higher or lower volatility. The corrected test flags an average of 8.01% fewer days as containing at least one jump when compared to the uncorrected Lee-Mykland test.

6.3 Adjusted Returns

To test the connection between equity returns and the oil futures market, the first variable investigated was the adjusted returns. The factor loading matrix is described in Table 1. Only two of the factors were determined to have eigenvalues above 0.2 and significantly contribute to the explanation of the variance of the adjusted returns. For a visual demonstration of the results, the factor loadings are plotted in Figure 4.3, where the correlations with factor one and two form the coordinates. There are three clear groupings in Figure 4.3: OIL, the oil companies, which includes XOM, COP, CVX, and BHI, and the remaining, non-oil companies.

The first extracted factor can be interpreted to be the “market” factor, describing the overall trend in the S&P 100. It is highly correlated with all of the equity stock returns in the sample as seen in the second column of Table 1. As is indicated by a correlation of only .09, factor one does not meaningfully describe OIL return variance. This is expected as the price of oil is more related to political events and supply-shocks than macroeconomic indicators (Chua 1994). The first factor describes 37% of the total variance. This factor would be related to general employment statistics, the strength of consumer spending, inflation impacts, and the macro-economic cycle. Another explanation of the market factor could be the systemic risk defined by the Capital Asset Pricing Model. Under the CAPM framework, the common variance across the different equity returns should indicate the risk-return profile for the stock market.

The second factor can be interpreted to be the “oil” factor, as oil futures and the oil equities are significantly and positively correlated with factor two while the non-oil companies were negatively and less significantly correlated with factor two. Therefore, this factor could simply be the expected future price of oil, where an increase in the expected oil price would lead to an increase in the adjusted returns of oil futures and oil business sector equities. Further, an

expected increase in oil prices should decrease the returns of non-oil related companies as petroleum is an input in their products and may place a damper on overall consumer spending. The importance of oil on the returns of the U.S. economy, as studied by Hamilton (1983) and Mork, Olsen, and Mysen (1994), is apparent even at the 5-minute return level when looking at individual equity returns. Finally, together the two factors are able to explain an average of 53% of the variance of the oil company stocks, while only 25% of the non-oil equities. By including four stocks from the same oil-related sector, factor analysis is able to isolate the oil factor and explain an additional portion of the sector's returns beyond the return profile of the broad market.

These two factors were present during the first half of 2008. While it was a record year for profits across the oil sector due to high crude prices, many petroleum companies were trading down approximately 20% through October 2008. However, the overall market has declined an average of 34% through the same period. Experts suggest that the poor health of the economy has depressed the consumer demand for petroleum products and led to decreased expectations for oil equities. Fluctuations in oil futures prices affect oil equity returns at the 5-minute return level. However, petroleum companies are not immune to the macroeconomic cycles and outlooks affecting stock markets, also seen at the return level.

6.4 Jump Statistics

The second test to investigate the connection between the futures and equities market is to analyze whether price jumps occur simultaneously. The analytic result of examining simultaneous jump statistics is very similar to the returns model of the previous section: a market and an oil factor drive the data. When examining the jump statistics data, the rotated factor-loading matrix yielded more interpretable results. The results of the adjusted returns above suggest there is correlation between the two factors as oil prices can affect the macro economy in the form of oil shocks (Chua 1994). For this reason, an oblique rotation is used for interpretation that allows for correlation between factors. The output is presented in Table 2.

As expected, the jump results are similar to the results of the standardized returns. However, the factors selected capture less of the variance of the jump component than they did for the adjusted returns. This is reflected in the lower communality for each company's counterpart in the adjusted return factor analysis presented previously in Table 1. For example,

60% of Exxon Mobil's return variance is explained by the two factors while only 39% of the jump variance is explained.

Factor one is highly correlated with oil companies and slightly less so with oil futures and non-oil companies. As was the case in analyzing the adjusted returns, this appears to be a "market" factor that captures the systematic likelihood for all equities in the market to jump. The second factor is the oil factor—positively correlated with oil companies and negatively correlated with the non-oil companies. Oil company stock prices are likely jumping contemporaneously as they respond similarly to oil announcements and crude oil price fluctuations while non-oil companies are not likely to be jumping. Further, the low levels of oil futures correlation, as evidenced by the top row, indicate that the oil futures market is reacting to similar, yet unique events that determine the price of oil. The presence of non-simultaneous jumping of oil futures and equities is puzzling. If oil equity returns were closely tied to oil futures returns, one might expect efficient markets would create contemporaneous jumps. The incredible access to information and the use of new technology in electronic trading should make jumps across related markets simultaneous. A possible explanation is that unexpected, random jumps affect the more volatile futures market first and then later reverberate to the equity markets. It is also possible that markets are inefficient and an arbitrage opportunity exists.

6.5 Instantaneous Volatility

While the adjusted return and jump statistic data yield similar analytic results, an investigation of the contemporaneous volatility across the futures and equity market does not fit the same pattern. Table 3 displays the results of the unrotated factor-loading matrix. The most interesting aspect of the analysis is that on average, 85% of the oil company price volatility is explained by the common three factors. Approximately 75% of the non-oil equity price volatility is explained by the three retained factors. This is not a surprising result as the volatilities of any two equities traded in the same market are highly correlated. However, only 26 percent of the variance in oil futures volatility is explained by the three factors. The first factor is highly correlated with all of the equity variables, yet not at all with the oil futures volatility. The extremely volatile nature of the oil futures market most likely excludes it from sharing correlation with factor one. The presence of factor one suggests that there exists a common element of volatility across the entire equity market, regardless of business sector.

Additionally, the second factor contributes uniquely to the oil-company stocks and OIL, accounting for the higher communality across those sector-specific stocks. The negative relationship between oil futures volatility and factor two, and thus the oil-equity sector, is puzzling. One might expect that an increase in volatility of oil futures would trigger trading of the oil equities, but this is not supported by the data. Wang, Wu, and Yang (2008) find that oil futures volatility is symmetric across positive and negative returns. Although increased volatility does not indicate changes in expected returns, risk-averse equity investors could be avoiding the uncertainty of buying and selling during periods of highly volatile oil prices. The low communality of oil futures volatility indicates that neither factor explains oil futures volatility well. Thus, there is still a possibility for spurious correlation between OIL volatility and factor two. In contrast to the results from analyzing the adjusted returns and jump statistics data, oil price volatility and equity volatility appear to have no connection when assessed at the 5-minute observation level.

6.6 Co-Jump Matrix

The jump-day correlation-type matrix, using the statistic defined in Section 4.6, is displayed in Table 4. The results of the matrix suggest that there are “market” and “oil” factors similar to the factor analysis results above. The equity market factor appears to contribute approximately 0.20 to the statistic, as any two given stocks will have several common jump days. This could be due to market-wide jump-inducing announcements like a major interest rate change, or it could be a spurious correlation indicating nothing more than the probability of two random processes sharing a common jump day. However, the correlation-type relationship between XOM, CVX, and COP is not sufficiently explained by the “market” alone. The average correlation between the three stocks is 0.40, double the average correlation between any two stocks chosen at random. By examining the existence of jumps over a given day, this method compensates for a possible lagged response of the equity market. The correlation between OIL and the three stocks remains around 0.20. This indicates that oil-company prices are not jumping in response to jumps in the oil futures price. Further, by examining the OIL column, it appears that oil price jumps simply do not occur simultaneously with jumps in the equity market beyond what can be accounted for by spurious jumping. The fact that OIL and equities are traded in

markets with completely different dynamics, investors, and price determinants could be limiting the ability to compare the two markets when applying a jump test.


7. Conclusion

This paper looks at simultaneous returns, jumps, and volatilities of oil futures, oil equities, and other equities in the S&P 100. After adjusting returns for the natural intraday volatility pattern that exists in the data, factor analysis is used to extract the essential components that drive the variation in the three dependent variables. Through this method, a market factor is found to affect the overall level of returns across the equities and the likelihood that two given equities to jump simultaneously. This component likely describes the macroeconomic conditions that drive equity returns regardless of specific industry. A second factor is found to affecting the returns and jumps that uniquely describes the variation in the oil equity and futures data. This characteristic supports the intuitive connection between oil futures and oil equity results discussed in the paper. Volatility in oil futures and equities, however, is not found to have a common root due to the differences in types and motivations of traders.

One future extension of this paper is to study the potential lag-lead relationships between the two markets. The value of volatility spillovers on determining the rate of informational flow that Ross (1989) outlines suggests that the volatility relationship between the markets could be better explained by lead-lag relations rather than a contemporaneous relationship. A GARCH framework could elucidate the potential leading role that oil futures volatility plays in imputing oil equity volatility. Of course, the rise of instant information and computerized trading could have vastly changed in the 20-year old model that Ross (1989) introduced. Another recommendation of further research the inclusion of an equity index, such as the S&P 500, in determining market returns, jumps, and volatility. The true activity of equity markets may not be fully captured through looking at individual stocks.

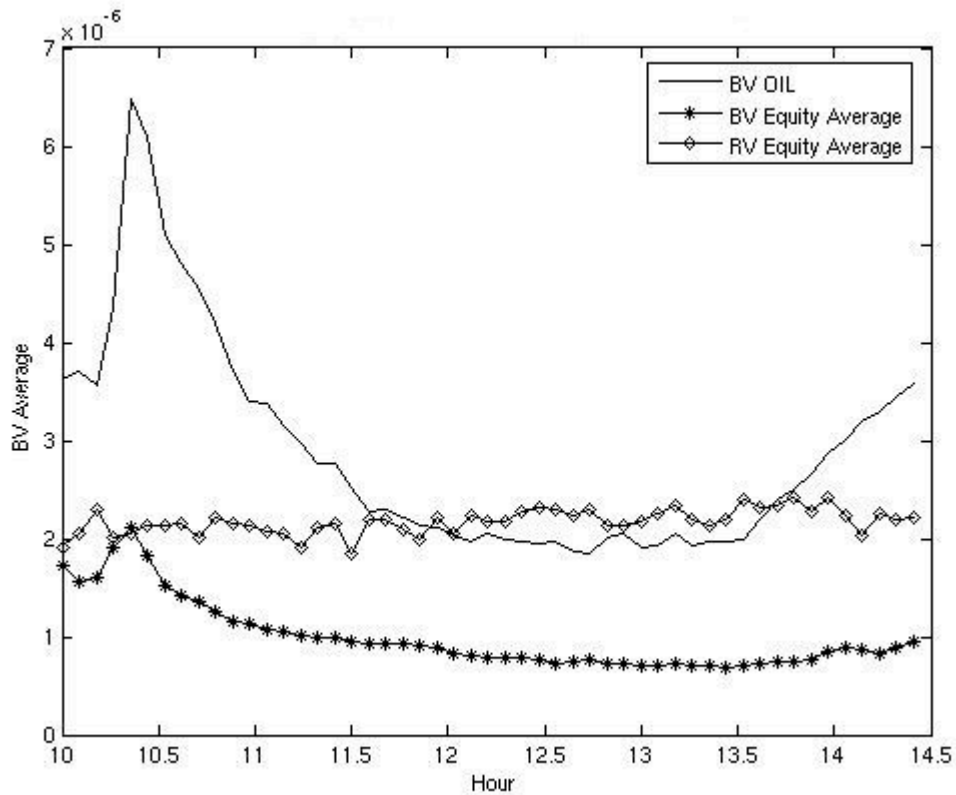
A. Figures and Tables

Figure 1:
Histograms of Intraday Unadjusted Jump Pattern



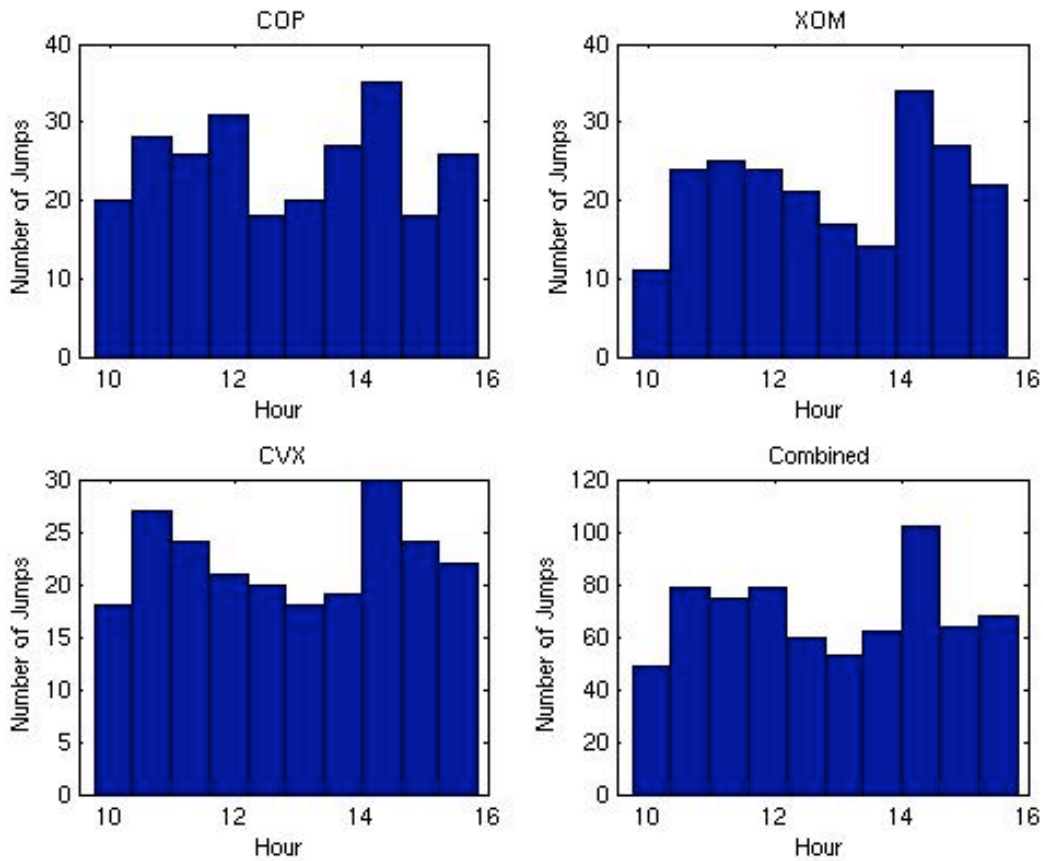
The histograms are for COP, XOM, CVX and Combined, which displays the total number from the three stocks. The histograms display the total number of flagged jumps at every observation time over the period from September 3rd, 2002 to January 24, 2008 for a total of 1,343 trading days. The jumps detection method is the unadjusted Lee-Mykland (2008) test. The Lee-Mykland test uses observations starting at 9:40 a.m. and ending at 4 p.m.

Figure 2:
Intraday Bipower Variation for Oil Equities and Futures




The solid line displays the intraday Bipower Variation as defined in equation (16) for oil futures prices. The asterisk line averages the intraday Bipower Variation, from equation (16), across XOM, COP, and CVX. The diamond line averages the Realized Variance, as defined in equation (7), across XOM, COP, and CVX. All plots use data over the period from September 3rd, 2002 to January 24, 2008 for a total of 1,343 trading days. The BV and RV calculations use observations starting at 10 a.m. and ending at 2:30 p.m.

**Figure 3:
Intraday Adjusted Jump Pattern**



The histograms are for COP, XOM, CVX and Combined, which displays the total number from the three stocks. The histograms display the total number of flagged jumps at every observation time over the period from September 3rd, 2002 to January 24, 2008 for a total of 1,343 trading days. The jumps detection method is the adjusted Lee-Mykland test, as defined in equation (17). The adjusted Lee-Mykland test uses observations starting at 9:40 a.m. and ending at 4 p.m.

Figure 4:
Combined Jumps with Kernel Density




The histogram displays the total number from three stocks: XOM, CVX, and COP. The histogram displays the total number of flagged jumps at every observation time over the period from September 3rd, 2002 to January 24, 2008 for a total of 1,343 trading days. The jumps detection method is the adjusted Lee-Mykland test, as defined in equation (17). The adjusted Lee-Mykland test uses observations starting at 9:40 a.m. and ending at 4 p.m. The diamond line displays the kernel density for the histogram of combined jumps.

**Table 1:
Adjusted Returns Factor Loadings**

	F1 - Market	F2 - Oil	Communality
rOIL	0.0985	0.3428	0.1272
rCOP	0.653	0.3128	0.5243
rCVX	0.7242	0.2753	0.6003
rXOM	0.7479	0.2073	0.6023
rBHI	0.5413	0.3196	0.3951
rETR	0.4018	-0.2064	0.204
rF	0.2977	-0.1567	0.1132
rFDX	0.4608	-0.2785	0.2899
rGS	0.5245	-0.2801	0.3535
rPG	0.4588	-0.2768	0.2871
rBA	0.4553	-0.266	0.278
rDELL	0.4524	-0.2486	0.2665

The table represents the factor analysis results from the adjusted returns data. The data is described in detail in section 5 and the adjusted returns are derived in equation (17). Each row represents a different financial asset. The first column is the correlation between the ‘Market’ factor and the asset’s returns, the second column is the correlation between the ‘Oil’ factor and the asset’s returns, and the third column provides the percentage of total return variance explained by the two factors. See Section 4.4 for an explanation of factor analysis.

Figure 4.3:
Plot of Adjusted Returns Factor Loadings



The factor loadings from the adjusted returns factor analysis output in Table 1 are plotted. Three distinct groups are displayed: oil equities returns in the lower right corner, oil futures returns in the lower left corner, and non-oil equities in the upper middle.

**Table 2:
Factor Loadings of Jump Statistic**

	F1 - Market	F2 - Oil	Communality
jOIL	0.1358	0.0866	0.0259
jCOP	0.5341	0.2113	0.3299
jCVX	0.6098	0.1871	0.4068
jXOM	0.6137	0.1341	0.3946
jBHI	0.4076	0.1836	0.1998
jETR	0.261	-0.1714	0.0975
jF	0.1663	-0.1172	0.0414
jFDX	0.2991	-0.2158	0.136
jGS	0.3523	-0.2324	0.1781
jPG	0.3097	-0.2196	0.1442
jBA	0.2873	-0.2093	0.1263
jDELL	0.2889	-0.1973	0.1224

The table represents the factor analysis results from the jump statistic data. The data is described in detail in section 5 and the jumps statistic is derived in equation (9). Each row represents a different financial asset. The first column is the correlation between the ‘Market’ factor and the asset’s jump statistic, the second column is the correlation between the ‘Oil’ factor and the asset’s jump statistic, and the third column provides the percentage of total jump variance explained by the two factors. See Section 4.4 for a more complete explanation of factor analysis.

See Section 4 for an explanation of the jump statistic used.

**Table 3:
Factor Loadings of Instantaneous Volatility**

	F1	F2	F3	Communality
ivOIL	-0.079	0.290	0.413	0.261
ivCOP	0.736	-0.576	0.176	0.904
ivCVX	0.837	-0.492	0.053	0.945
ivXOM	0.921	-0.241	0.021	0.908
ivBHI	0.795	-0.225	0.101	0.693
ivETR	0.860	0.097	0.041	0.750
ivF	0.777	0.310	-0.048	0.702
ivFDX	0.841	0.164	-0.122	0.749
ivGS	0.816	0.020	-0.231	0.719
ivPG	0.751	0.305	0.277	0.734
ivBA	0.712	0.535	0.105	0.804
ivDELL	0.755	0.227	-0.305	0.714

The table represents the factor analysis results from the instantaneous volatility data. The data is described in detail in section 5 and the volatility is derived in equation (10). Each row represents a different financial asset. Columns 1-3 present the correlation between instantaneous volatility and the three extracted factors across the twelve different assets. See Section 4.4 for a more complete explanation of factor analysis.

Table 4:
Co-Jump Correlation Matrix

	OIL	COP	CVX	XOM	BHI	ETR	F	FDX	GS	PG	BA	DELL
OIL	1.000											
COP	0.253	1.000										
CVX	0.146	0.420	1.000									
XOM	0.287	0.324	0.458	1.000								
BHI	0.195	0.283	0.185	0.238	1.000							
ETR	0.199	0.212	0.247	0.187	0.125	1.000						
F	0.196	0.185	0.190	0.159	0.112	0.239	1.000					
FDX	0.178	0.214	0.220	0.235	0.150	0.223	0.186	1.000				
GS	0.142	0.226	0.286	0.334	0.195	0.218	0.215	0.194	1.000			
PG	0.141	0.229	0.223	0.217	0.145	0.223	0.225	0.260	0.299	1.000		
BA	0.136	0.219	0.299	0.219	0.183	0.243	0.149	0.187	0.256	0.239	1.000	
DELL	0.147	0.264	0.267	0.233	0.134	0.233	0.223	0.233	0.262	0.299	0.241	1.000

This matrix presents the correlation type statistic that describes the correlation between the jump days of two assets. The correlation statistic is defined as

$$\rho_{a,b} = \frac{C_{a,b}}{\sqrt{J_a * J_b}},$$

defined in more detail in Section 4.5. The jump detection method is the adjusted Lee-Mykland test, as defined in equation (17). All twelve financial assets are presented, as described in Section 5. The data spans the period from September 3rd, 2002 to January 24, 2008 for a total of 1,343 trading days.

B. References

- Andersen, T., T. Bollerslev, and F. Diebold (2004). Some Like it Smooth, and Some Like It Rough: Untangling Continuous and Jump Components in Measuring, Modeling, and Forecasting Asset Return Volatility. *Working Paper, Duke University*.
- Barndorff-Nielsen, O. and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics*, 2, 1-48.
- Barndorff-Nielsen, O. and N. Shephard (2006). Econometrics of Testing for Jumps in Financial Econometrics Using Bipower Variation. *Journal of Financial Econometrics*, 1, 1-30.
- Cattell, R. B. (1966). The Scree Test for the Number of Factors. *Multivariate Behavioral Research*, 2, 245-276.
- Chua, J. and R. Woodward (1994). Financial Performance of the U.S. Oil and Gas Industry: 1980-1990. *Financial Markets, Institutions, and Instruments*, 3.1, 1-56.
- Cooper, J.C.B. (2003). Price elasticity of demand for crude oil: estimates for 23 countries. *OPEC Review*, 27, 1-8.
- Driesprong, G., B. Jacobsen, and B. Maat (2003). Striking Oil: Another Puzzle. *Working Paper, Erasmus University Rotterdam*.
- Energy Information Administration (2007, June 18). *Oil Market Basics*. Retrieved November 15, 2008, from <http://www.eia.doe.gov/pub/oil_gas/petroleum/analysis_publications/oil_market_basics/default.htm>
- Energy Information Administration (2008, October 31). *Official Energy Statistics from the U.S. Government*. Retrieved November 15, 2008, from <<http://www.eia.doe.gov/>>
- Gorsuch, R. (1983). *Factor Analysis*. Philadelphia: Saunders.
- Hair, J.F., B. Black, B. Babin, R. Anderson and R.L. Tatham (1998). *Multivariate Data Analysis*. New Jersey: Prentice Hall.
- Hamilton, J.D. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy*, 91, 228-248.
- Harman, H. (1967). *Modern Factor Analysis*. Chicago: University of Chicago.
- Haubrich, J.G. (2005). Oil Prices: Backward to the Future. *Federal Reserve Bank of Cleveland Economic Commentary*.
- Huang, R.D., R.W. Masulis, and H.R. Stoll (1996). Energy Shocks and Financial Markets. *Journal of Futures Markets*, 1, 1-27.

- Huang, X. and G. Tauchen (2005). The Relative Contribution of Jumps to Total Price Variance. *Journal of Financial Econometrics*, 4, 456-499.
- Johnston, D. and D. Johnston (2006). *Introduction to Oil Company Financial Analysis*. Tulsa: PennWell.
- Jones, C. and G. Kaul (1996). Oil and the Stock Markets. *Journal of Finance*, 2, 463-491.
- Kaiser, H.F. (1960). An analytical rotation criterion for factor analysis. *Review of Educational Research*, 5, 440-486.
- Kang, S.H., S.M. Kang and S.M. Yoon (2008). Forecasting Volatility of Crude Oil Markets. *Energy Economics*, In Press September 2008.
- Larson, P. (2008, October 31). XOM, COP, CVX Analyst Reports. *Morningstar Equity Analyst Reports*, Retrieved November 15, 2008.
- Lee, S and P. Mykland (2008). Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics. *The Review of Financial Studies*, Forthcoming, 1-29.
- Levy, H. and T. Post (2005). *Investments*. Essex: Pearson.
- Litzenberger, R. and N. Rabinowitz (1995). Backwardation in Oil Futures Markets: Theory and Empirical Evidence. *Journal of Finance*, 5, 1517-1545.
- Martin, R. and W. Leffler (2006). *Oil and Gas Production in Nontechnical Language*. Tulsa: PennWell.
- Marzo, M. and P. Zagaglia (2007). Volatility Forecasting for Crude Oil Futures. *Working Paper*, Johns Hopkins.
- Merton, R.C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3, 373-413.
- Merton, R.C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3, 125-144.
- Moosa, I.A. and N.E. Al-Loughani (1994). Unbiasedness and Time-varying Risk Premia in the Crude Oil Futures Market. *Energy Economics*, 2, 99-105.
- Mork, K.A., O. Olsen, and H.T. Mysen (1994). Macroeconomic Response to Oil Price Increases and Decreases in Seven OECD Countries. *The Energy Journal*, 5, 19-35.
- Roerber, J. (1993). *The evolution of oil markets: trading instruments and their role in price*

formation. London: Royal Institute of International Affairs.

Ross, S. (1989). Information and Volatility: The No-Arbitrage Martingale Approach to Timing and Resolution Irrelevancy. *Journal of Finance*, 44, 1-17.

Sadorsky, P. (1999). Oil Price Shocks and Stock Market Activity. *Energy Economics*, 5, 449-469.

Stevens, J. (2002). *Applied Multivariate Statistics for the Social Sciences*. New Jersey: Lawrence Erlbaum.

Van Tassel, P. (2008). Patterns Within the Trading Day: Volatility and Jump Discontinuities in High Frequency Equity Price Series. *Duke Journal of Economics, Undergraduate Research Symposium Spring 2008*.

Wang, T., J. Wu, and J. Yang (2008). Realized Volatility and Correlation in Energy Futures Markets. *Journal of Futures Markets, Forthcoming*.