Variance Risk Premium Dynamics:
The Impact of Asset Price Jumps on Variance Risk Premia

Jackson R. Pfeiffer†

Dr. George Tauchen, Faculty Advisor

Duke University

Durham, North Carolina

2013

† Jackson Pfeiffer graduated magna cum laude from Duke University in 2013, with a Bachelors of Science in Economics with a Finance Concentration, and with Certificates in Markets and Management Studies and Arts of the Moving Image. Pfeiffer graduated with High Distinction in Economics, and is currently taking math courses at Duke University in preparation for the pursuit of a PhD in Economics, and can be contacted at jackson.pfeiffer@alumni.duke.edu.

Acknowledgements

I cannot even begin to thank Dr. George Tauchen and Dr. Tim Bollerslev adequately for their insightful comments and guidance throughout the process of researching and writing this paper. I would also like to thank my classmate Andrew Bentley for his valuable constructive criticism throughout the writing process. This paper in its current state would not be possible without the support of all those involved in this process.

Abstract

This paper utilizes the high-frequency stock price data and the corresponding daily option price data of several highly capitalized corporations in order to investigate the impact that asset price jumps of individual equities have on the equities' respective variance risk premia. The findings of this paper describe many characteristics of the variance risk premia of individual equities, supporting some expectations of the characteristics, and refuting others. In the process of investigating these characteristics, this paper proposes a simple estimator for the market price of the variance risk of an individual equity.

JEL classification: G1; G11; G19
Keywords: Variance Risk; Variance Swaps; Price Jumps

Contents

1 Introduction

All investments inherently involve risk; not even the safest of government bonds are completely risk-free. For that reason, one of the most important aspects of investing is managing risk. The most basic form of risk in the financial markets is an asset's volatility – a measure of the variability of the asset's returns. Investors, especially large institutional investors, determine both the components of their portfolio and the relative weights with which the components are represented based in large part on the volatility of the various investments available to them. Thus, it stands to reason that if the volatility of any of the assets in an investor's portfolio increased, then the asset would be riskier than it had previously been, and the investor would be less willing to hold said asset.

It has been known for some time that an asset's volatility changes over time [e.g. Fama (1965), Engle (1982), and French, Schwert, and Stambaugh (1987)]. This poses a risk for investors, therefore, that the volatility of an asset will increase, thereby making the asset less valuable. The intuition behind this is straightforward and is corroborated by evidence demonstrating the existence of the leverage effect, as discussed by Nelson (1991) and Ghysels, Harvey, and Renault (1996), whereby an asset's price movements are negatively correlated with said asset's volatility. The risk that volatility will increase is therefore known as the variance risk. As with any other risk, investors expect to be compensated in exchange for holding onto the risk. The compensation an investor receives for holding the volatility of the asset's return variance is called the variance risk premium.

Throughout the past several years, variance risk premia have become a highly researched subject in financial economics. Neurberger (1994) and Dupire (1992; 1993) are credited with initiating this research, as they both independently created techniques to study the traits of variance risk premia by using various option portfolios. Option payouts, they realized, were highly dependent on the volatility of their underlying assets, and option prices incorporated the market's expectation of the assets' ex-ante volatilities. Shortly afterwards, Carr and Madan (1998) created a formula to estimate the size of variance risk premia by creating a portfolio of options representing a cross-section of one asset's options, all with the same maturity. Since Neurberger and Dupire's work, several people have looked into how variance risk interacts with other factors [e.g. Bates (1996), Pan (2002), and Eraker (2004)]. Bakshi and Kapadia (2004b),

as well as Carr and Wu (2009), have investigated the relationship between the variance risk premia of individual equities and that of the market.

The availability of high-frequency data has made it possible to investigate how variance risk premia respond to intraday changes of an asset's volatility. Todorov (2009), for example, researched the effect that price jumps have on the market's variance risk premium. This paper aims to add to the current understanding of the dynamics of variance risk premia, extending that research to estimating the effect asset price jumps have on the variance risk premia of various individual equities rather than the market as a whole. Specifically, this paper seeks to measure how the presence, direction, and magnitude of intraday jumps impact the payout of variance swaps.

The outline of the remainder of this paper is a follows. Section 2 begins with a discussion about the role variance plays in the theoretical framework of an asset's efficient logarithmic price process. From there, the paper develops a methodology to measure the variance of an asset – first by conceptualizing a theoretical framework to describe the amount of volatility exhibited throughout a day, then by introducing a proxy for volatility that is demonstrated to be a good estimator for volatility. Section 3 examines variance risk premia. Initially, the section explains what exactly variance swaps are and how they work. Next, there is a discussion about the fair value of a variance swap and about various methodologies to estimate the fair value through option prices. These methodologies, however, prove to be infeasible to use in the estimation of the fair value of the variance swap for individual equities. Thus, this paper proposes a more computationally simple estimator. When run on the same data, the simple estimator seems to work nearly as well as the more computationally intensive estimators. Lastly, the section briefly explains the various expressions of the variance risk premium and their respective interpretations. Section 4 discusses asset price jumps. The section begins by expanding upon the model (introduced in Section 2) that describes an asset's price-path so as to allow for the existence of jumps – large price discontinuities – in the model. This paper then re-examines the consistency of the measure of variance from Section 2, given the inclusion of jumps in the price-path, and introduces a more jump robust measure of variance. Section 4 goes on to describe the common strategies that are used in attempts to locate asset price jumps. Then, the paper proceeds to create its own litmus test to determine which price movements could best be described as jumps. Continuing on, the section discusses the methods that are used to

describe and quantify each jump's respective direction and magnitude. Next, Section 5 provides an overview of the data and methodology that is used in this paper's research. The results of the research are discussed in detail in Section 6. Finally, Section 7 provides both some concluding remarks and suggestions for future research.

2 Measures of Variance

In financial economics, volatility is a measure of variability of an asset's returns, and is defined as the standard deviation of the asset's logarithmic price returns. The volatility of an asset relates directly to the path of that asset's price through time. In a standard continuous time model, an asset's price movements are described by the differential equation

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

where $p(t)$ is the logarithmic asset price, μ is the drift component of the asset price, σ is the volatility of the price returns, and $dW(t)$ is standard Brownian motion. Fama (1965) observed that stock prices follow what is commonly referred to as a random walk. If we think of an asset's price return as a random variable, μ would be the variable's expected value and σ^2 its variance. The Brownian motion component of the differential equation serves to provide randomness to the model. This random component is scaled by the asset's volatility σ.

It is important to note that this model considers volatility to remain constant. Research has shown, however, that an asset's volatility is not constant, and that models with time-varying instantaneous volatility components are more accurate than static volatility models. The differential equation describing an asset's logarithmic price movement, as presented by Merton (1971), becomes

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

An asset's price movements, however, do not always follow a continuous path. Research has shown that continuous-path models are inferior to models that incorporate asset price discontinuities, known as jumps. The differential equation can therefore be further improved by including an additional component to account for these jumps. The implications of including jumps in the model of an asset's price path will be examined during the discussion of jumps in Section 4. Furthermore, trading frictions – such as bid-ask bounce, asynchronous trading, and congestions in market making – make it impossible to observe instantaneous price movements.

Essentially, since we are only able to observe an asset's price at the discrete moments in time at which the asset is traded, we must assume that these points would all lie on or near the continuous-time path of our models. We are therefore unable to observe the instantaneous volatility of an asset.

It is, however, theoretically possible to approximate an asset's volatility over a discrete period of time, such as a day, by finding the average of the asset's volatility over that time. This theoretical average is known as the integrated variance (IV), and it is defined by Anderson and Bollerslev (1998) to be

$$IV_t = \int_{\tau=0}^{1} \sigma^2(t+\tau) \cdot d\tau.$$

It is still infeasible to calculate IV_t without the ability to observe a continuous price path. How, then, is volatility measured?

By using high-frequency data, it is possible to observe discrete price movements using arbitrarily small, fixed increments of time. The interval of time between observations, otherwise known as the sampling interval (Δ), is inversely proportional to the number of intervals each day, known as the sampling frequency (M). Under this framework$\{p_{t,1}, p_{t,2}, \dots, p_{t,M+1}\}$ is the set of $M+1$ price observations on day t, and $\{r_{t,1}, r_{t,2}, \dots, r_{t,M}\}$ is the set of M logarithmic price returns on day t, such that

$$r_{t,j} = \ln(p_{t,j+1}) - \ln(p_{t,j}).$$

Andersen and Bollerslev (1998) suggest measuring volatility by using the cumulative squared intraday returns. This proxy is known as the realized variance (RV) and is defined as

$$RV_t(M) = \sum_{j=1}^{M} r_{t+j/M}^2.$$

In their paper, Andersen and Bollerslev show that, assuming there are no jumps in the price process, by the theory of quadratic variation as discussed by Karatzas and Shreve (1988),

$$\operatorname*{plim}_{M\to\infty}\left(\int_{\tau=0}^{1} \sigma^2(t+\tau) \cdot d\tau - \sum_{j=1}^{M} r_{t+j/M}^2\right) = 0.$$

By substituting the definitions of integrated variance and realized variance into the above limit and rearranging the components, we are therefore able to show that, as the sampling frequency

(M) approaches infinity and the length of each interval (Δ) approaches zero, the realized variance (RV) asymptotically approaches the integrated variance (IV).

3 Variance Risk Premium

3.1 Variance Swaps

Investors are generally thought to be risk averse. The basic concept of risk aversion is that when deciding between two independent investments, ceteris paribus, the risk-averse investor will prefer the investment that is less risky. This is not to say, however, that a risk-averse investor will only hold one investment in his portfolio. There is a vast selection of potential investments from which he can choose, each with its own level of riskiness and expected rate of return. The investor chooses several different investments, such that he is holding his desired level of risk in his portfolio. Generally, the riskiness of an asset is considered to be the volatility of its returns.

All else equal, investors prefer assets with less volatility. Volatility can, as previously discussed, change over time. If an asset's volatility increases, it becomes less desirable to investors. This means that investors are at risk of the volatility of an asset in their portfolios increasing. Investors seeking to reduce their exposure to any such increases in volatility want a product that has a positive payout if the asset's volatility increases and a negative payout if that volatility decreases. Demand for such a product drove the creation of variance swaps.

A variance swap[1] is a type of financial instrument that enables investors to exchange cash flows based on a variable rate, the realized variance, for cash flows based on a fixed rate, the swap rate. The contracts are structured so that the buyer is paid the difference between the realized variance (RV) of some underlying product over a period of time (usually 30 days) and the contract's agreed upon fixed swap rate (SW), such that

$$\text{Payout} = N \times (RV - SW),$$

where N is the notional amount of the contract, and $(RV - SW)$ is the payout of a contract with a notional amount of \$1.

[1] For a detailed explanation of swaps, see Appendix D.

For example, let us assume that firm A buys a variance swap from firm B with a swap rate of 18% and a notional amount of \$100,000. If, after 30 days, the realized variance is calculated to be 23%, then firm A's payout would be

$$\begin{aligned} \text{Payout} &= \$100{,}000 \times (0.23 - 0.18) \\ &= \$5{,}000 \qquad , \end{aligned}$$

meaning that firm B would pay firm A \$5,000. If, on the other hand, the realized variance is calculated to be 15%, then firm A's payout would be

$$\begin{aligned} \text{Payout} &= \$100{,}000 \times (0.15 - 0.18) \\ &= -\$3{,}000 \qquad . \end{aligned}$$

The negative payout for firm A indicates that firm A would actually pay firm B \$3,000. Variance swaps can therefore be used to effectively hedge out the risk of the asset's volatility changing.

It is important to note that while there is a limit to how low any asset's realized variance can become, with a theoretical minimum of zero, there is no upper bound to how high an asset's volatility can spike. While spikes in volatility are somewhat rare, the party that sells a variance swap risks losing a large amount of money if such a spike occurs. Continuing the example above, if the realized variance is calculated to be 45%, then firm A's payout would be

$$\begin{aligned} \text{Payout} &= \$100{,}000 \times (0.45 - 0.18) \\ &= \$27{,}000 \qquad , \end{aligned}$$

meaning that firm B would have to pay firm A \$27,000.

The above examples assumed the contract's swap rate to be 18%, but how is this rate determined? This leads to the economic question of the fair value of a financial instrument with the payout of a variance swap.

3.2 Variance Swap Rate

In order to compensate the sellers of variance swaps for the risk they are taking, the swap rates are set such that buying a variance swap has a small negative expected return. The small negative return for buyers of variance swaps can be interpreted as a small positive expected return for the sellers of the swaps. This small reward is referred to as the variance risk premium. The fair value of a variance swap should therefore be slightly higher than the market's expectation of future variance.

Variance swaps are over-the-counter instruments, meaning that they are private contracts between two parties and are not traded on any exchange. As such, the details about the contract

terms are never publicly reported. Therefore, the swap rates cannot be observed directly. How then can we estimate the swap rate? We must instead infer the fair swap rate of a variance swap from some other financial instruments that are traded on an open exchange and the prices of which reflect the market's expectation of future variance.

In order to get a view of volatility, we look at a type of financial instrument called option contracts. An option, which is traded on the Chicago Board of Options Exchange (CBOE), is a derivative, meaning that its fair value is based on, or derived from (hence the term derivative), some underlying asset. A call option gives its buyer the right, but not the obligation, to buy the underlying asset on the option's expiration date at the option's strike price. Conversely, a put option gives its buyer the right, but not the obligation, to sell the underlying asset on the option's expiration date at the option's strike price. For example, if a person had purchased a single call option on XYZ stock with a May 17 expiration date and a strike price of $100, then on May 17 that person would have the right to buy 100 shares of XYZ stock for a price of $100 per share. If the spot price of XYZ stock is $110 on the expiration date of the option, then the basket of options is effectively worth $1,000. If, however, the spot price of XYZ stock is lower than $100 on the expiration date, then the owner of the call option would be better suited just buying the stock rather than executing the option, making the option worthless.

The asymmetric payout structure of options makes the fair value of an option highly dependent on the market's expectation of the underlying asset's future volatility. Indeed, options traders often refer to options not by their nominal price, but instead by the amount of expected volatility that they are pricing into the options. This paper follows the intuition that the market's expected volatility that is priced into options should be close to the market's expected volatility that would be used as the swap rate. Unfortunately, while option prices depend greatly upon the expected level of volatility, the options themselves must officially be traded in nominal terms. We therefore cannot directly observe the expected volatility that is priced into options. How then can we estimate the market's expected volatility? In 1973, Fischer Black and Myron Scholes released a paper entitled "The Pricing of Options and Corporate Liabilities." The paper introduced a formula that was designed to determine the appropriate price of options. According to this formula, known as the Black-Scholes formula, the price of an option depends upon several variables: the current (or spot) price of the asset, the strike price of the option, the time

until the option's expiration date, the risk free interest rate, and the expected volatility of the asset's price over the life of the option.

All but one of these variables, the expected volatility, can be directly observed from the financial markets. It is possible, therefore, to determine the amount of expected volatility that must be entered into the Black-Scholes formula for it to output the price of the option that is actually observed. This level of expected volatility found by backing out of the formula is known as the implied volatility, as it is the volatility implied by the current price of the option. It should be noted, however, that for any given maturity of an asset's options, the various options each have different implied volatilities. This poses a problem: the market should only have one expectation of volatility, not several. We must therefore reconcile these differing levels of expected volatility in order to estimate the one level of expected volatility that would fairly value a variance swap, the swap rate.

It is theoretically possible to synthesize an estimator of the swap rate from the prices of a cross-section of the asset's call and put options for a given maturity. The estimator of the variance swap rate on day t, as considered by Carr and Wu (2009), is

$$SW_{t,T} = \mathbb{E}_t^*\left[RV_{t,T}\right] = \frac{2}{T-t}\int_0^\infty \frac{\Theta_t(K,T)}{B_t(T)K^2}\,dK,$$

where $\mathbb{E}_t^*\left[RV_{t,T}\right]$ equals the time t conditional expectation operator the realized variance over the life of the swap under a set of risk-adjusted probabilities[2], T equals the maturity date of the option, K equals the strike price of the option, $\Theta_t(K,T)$ equals the time t value of out-of-the-money option with $K > 0$ and $T \geq t$, and $B_t(T)$ equals the time t price of a risk free bond paying one dollar at time T.

This equation to estimate the variance swap of a given asset, however, is purely theoretical. The equation integrates across an infinite number of option strike prices, ranging from zero to infinity. In actuality, options are only actively traded if their strike prices fall within a relatively narrow range of the equity's spot price. Furthermore, options are only made

[2] Risk adjusted probabilities slightly skew the probabilities of each outcome so that, for the purposes of the calculation, bad outcomes are calculated as having higher probabilities of occurring than they actually have and good outcomes are calculated as having lower probabilities of occurring than they actually have. This reflects the human behavior of giving bad outcomes more weight and good outcomes less weight when evaluating opportunities.

available at discrete strike prices. The limited availability of options renders using the equation infeasible for estimating the variance swap rate.

For this reason, we look to the CBOE's Volatility Index (VIX), which aims to estimate the market's expected volatility of the S&P 500 index (SPX). In calculating the VIX, the CBOE employs an estimator of the market's expectation of 30-day volatility that is computationally similar to the theoretical estimator described above. The CBOE uses a weighted average of the prices of SPX index options across various strike prices.

Unfortunately, the available data on the option prices of individual equities is not as comprehensive as the data on the prices of SPX index options that is used by the CBOE. It is therefore infeasible to use the CBOE's methodology to estimate the variance swap rate for individual equities. We can, however, use the VIX as a benchmark against which to compare alternative estimators of the variance swap rate. It is possible to test the accuracy of an alternative estimator by using said alternative to estimate the expected 30-day volatility of the SPX index and comparing the alternative's estimate to the VIX. If, when performed on co-temporal SPX option data, the alternative estimator's results are similar to the VIX, we can assume that, when performed on the option data of individual equities, the alternative estimator would continue to produce estimates that are reasonably close to those that would be found by using the CBOE's estimator.

One such alternative estimator is the simple arithmetic average of the 26 implied volatilities given by Option Metrics' 30-day volatility data. The Option Metrics data represents what would theoretically be the implied volatilities of a cross-section of options with 30 days to maturity. It is comprised of the implied volatilities of the 13 out-of-the-money[3] calls and 13 out-of-the-money puts with strike prices nearest the equity's spot price. For each strike price, the 30-day implied volatility is calculated through a linear interpolation of the implied volatilities of the two options whose strike dates are closest to the date 30 days in the future. As seen in Figure 1, this simple estimator tracks the VIX quite well when applied to SPX options. This simple alternative estimator will therefore be used in this paper as the estimator of the variance swap rate.

[3] An out-of-the-money call is one for which the option's strike price is above the equity's spot price, and an out-of-the-money put it one for which the strike price is below the equity's spot price.

3.3 Variance Risk Premium

There are two different expressions for the realized variance risk premium, each with its own interpretation. The standard variance risk premium is defined by the equation

$$VRP_t = RV_{t,30} - SW_{t,30}.$$

As entering into a swap contract entails no upfront costs, VRP_t represents the nominal payout on day $t + 30$ of entering into a \$1 notational variance swap contract on day t. The Log Variance Risk Premium ($LVRP_t$), however, is defined as

$$\begin{aligned} LVRP_t &= \ln\left(RV_{t,30}\right) - \ln\left(SW_{t,30}\right) \\ &= \ln\left(RV_{t,30}/SW_{t,30}\right) \end{aligned}.$$

If we were to consider a variance swap to be a forward contract and the swap rate $SW_{t,30}$ to be the cost of the forward, then $\left(RV_{t,30}/SW_{t,30} - 1\right)$ represents the excess return from the investment in a variance swap. By conceptualizing the variance risk premium in this way, $LVRP_t$ acts as the continuously compounded excess return of entering into a variance swap on day t and holding it to maturity. For the purposes of this paper, we will be using $LVRP_t$ to express the ex-post variance risk premium because we are able to remain scale-free by working in logarithms.

4 Asset Price Jumps

As previously discussed, an asset's volatility can change over time. Observers of financial markets understand that high levels of variance in an asset's price movements tend to come in waves. This effect, whose first observation was noted by Benoit Mandelbrot (1963), is known as "volatility clustering." His observation, that "large changes tend to be followed by large changes – of either sign – and small changes tend to be followed by small changes," provided intuition to more accurate autoregressive models forecasting volatility in financial time series, notably Engle's (1982) Autoregressive Conditional Heteroskedasticity (ARCH) model and Bollerslev's (1986) Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model. These models break up volatility into two components, one of which represents the asset's natural base level of volatility, and one that describes the more fleeting levels of current volatility, both of which are impacted by recent price movements. Jumps can be viewed essentially as price movements much larger than one would expect given the current level of volatility. For this

reason we would expect jumps to have a significant impact on the component of volatility that relates to the recent price movements. The relative size and duration of a jump's impact on its respective asset's level of volatility has several implications on what the fair value of a variance swap should be.

In order to discuss jumps in detail, we must first expand upon the discussion of modeling the path of an asset's logarithmic price from Section 2. Recall the differential equation

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

By adding $\kappa(t) \cdot dq(t)$ to the model as the jump component of the price process, as suggested by Merton (1976), the differential equation becomes

$$dp(t) = \mu \cdot dt + \sigma(t) \cdot dW(t) + \kappa(t) \cdot dq(t),$$

where $\kappa(t)$ represents the magnitude of the jump and $q(t)$ is a binary variable indicating the presence of a jump. By using this notation, the model considers the jump component to be independent of the rest of the price process, allowing us to study the effect of a jump on volatility separately from the continuous price movements.

Andersen, Bollerslev, and Diebold (2002) demonstrated that when jumps are included in the model of price movements, as the sampling interval approaches zero, the limit of realized variance converges to the integrated variance plus the sum of the squares of the magnitudes of all the jumps that occur throughout the day, such that

$$\operatorname*{plim}_{\Delta \to 0} RV_t = IV_t + \sum_{\tau=0}^{1} \kappa^2(t+\tau).$$

Barndorff-Nielsen and Shepherd (2004b) suggested their own proxy for volatility, which is known as the bi-power variation (BV), and is defined in this paper as

$$BV_t(M) = \frac{\pi}{2} \times \frac{M}{M-1} \sum_{j=2}^{M} |r_{t,j}||r_{t,j-1}|.$$

Furthermore, Barndorff-Nielsen and Shepherd (2006) showed that their bi-power variation measure is quite robust and that, even with the presence of jumps, the bi-power variation converges to the integrated variance as the sampling interval approaches zero. Thus, bi-power variation is a consistent estimator for integrated variance. The characteristics of the bi-power variation theoretically allow us to isolate the jump component of the price process, for as the sampling interval approaches zero, the difference between the realized variance and the bi-power

variation converges to the cumulative sum of the squared magnitudes of jumps that occurred throughout the day, such that

$$\operatorname*{plim}_{\Delta \to 0}(RV_t - BV_t) = \sum_{\tau=0}^{1} \kappa^2(t+\tau).$$

Noting that the difference between realized variance and bi-power variation estimates the amount of realized variance that is contributed by jumps, several researchers have used this intuition as the basis for tests attempting to determine whether or not a jump occurred on any given day, notably Barndorff-Nielsen and Shepherd (2006) and Huang and Tauchen (2005). These tests, however, can only tell us if a jump occurred on any given day; the tests are unable to attribute jumps to the individual price movements during which they occur. Taking a different approach to the problem, a large amount of literature is devoted to creating tests that attempt to determine whether individual price movements can be classified as jumps. Often, these tests involve some threshold that a price movement must surpass in order to be classified as a jump [e.g., Lee and Mykland (2008), Mancini (2009), and Corsi, Pirino, and Renò (2010)].

4.1 Locating Jumps

In considering a method to classify jumps, this paper takes its intuition from Cecilia Mancini's (2009) threshold estimator and the idea of truncated variance, which is defined as

$$TV_t = \sum_{j=1}^{M} \left|r_{t,j}\right|^2 \times I\left[r_{t,j}\right],$$

$$\text{where } I\left[r_{t,j}\right] = \begin{cases} 1 \text{ if } r_{t,j} \leq \text{cutoff}_t \\ 0 \text{ if } r_{t,j} > \text{cutoff}_t. \end{cases}$$

In this way, $I[\cdot]$ is a binary indicator of whether a jump has occurred. Truncated variance classifies a price movement as a jump if the logarithmic return is greater than some cutoff value. In order to account for the latent amount of volatility present, the cutoff value is set each day based on a measure of the previous day's volatility. To minimize the effect that jumps have on that measure of volatility, this paper uses the bi-power variation as the measure of volatility. We can then define the cutoff value as

$$\text{cutoff}_t = \gamma \times \sqrt{\frac{1}{M} BV_{t-1}},$$

where γ equals the minimum magnitude of return relative to BV_{t-1} to be classified as a jump.

This paper adopts the intuition of a cutoff value determining the presence of jumps, and compares price returns to the cutoff value in order to determine the location of jumps. Furthermore, recall that, because trading frictions restrict us from viewing the true continuous-time price path of any asset, we must observe the price level at discrete time intervals. It is therefore much harder to estimate the jump component of the price model than the other components, as it is difficult to be certain whether a change in price is due to a jump or to a large continuous movement throughout the interval. Therefore, for the purposes of this paper, γ is set equal to eight (rather than the common range of three to five) so that we can be confident that each movement that is classified as a jump truly is best described as a jump.

4.2 Measuring Jumps

For those returns classified as jumps, this paper considers the size of the jumps to be the absolute value of the ratio between the jump's log-difference return and the previous day's bi-power variation, such that

$$\text{jump\_return}_t = \max_{j=1\to M} \left| \frac{r_{t,j}}{\sqrt{\frac{1}{M} BV_{t-1}}} \right|.$$

The maximizing function, therefore, serves to ensure that, if multiple jumps occur on the same trading day, the day's largest jump is considered to be its only jump.

4.3 Direction of Jumps

The direction of a jump, positive or negative, depends simply on the direction of the jump's underlying price movement, upward or downward. If multiple jumps occur on the same trading day, the direction of the day's jump is considered to be the direction of the price movement underlying the day's largest jump in magnitude. It should be noted that we would expect the presence of multiple jumps in a single day to have a larger impact on the future variability of some equity's price movements than the presence of a single jump. This represents a potential weakness in this paper's methodology for classifying jumps, as days are sorted into only two categories: days in which at least one jump occurs, and days in which no jumps occur.

5 Empirical Analysis

5.1 Data

The calculations in this paper employ two main sets of time series data, applying them to 10 diverse, highly capitalized corporations: Alcoa (AA), Apple (AAPL), Amazon (AMZN), Boeing (BA), Bank of America (BAC), Caterpillar (CAT), Dell Computers (DELL), General Electric (GE), Google (GOOG), and Goldman Sachs (GS). The first set of time series data is composed of each stock's intraday prices, reported in one-minute increments from 9:35am to 4:00pm, for every full trading day from April of 1997 through December 30th, 2010. The second set of time series data for each corporation is Option Metrics' 30-day volatility data, which reports daily levels of volatility implied by the closing prices of options. Every date is composed of the implied volatilities of 13 calls and 13 puts of various strike prices, each of which is interpolated to a maturity of 30 days.

It should be noted that the trading frictions discussed earlier in this paper cause a well-documented pattern of short-term deviations of the asset's observed price from its fundamental value, such that

$$p_{Obs}(\tau) = p(\tau) + \varepsilon_{\tau},$$

where $p_{Obs}(\tau)$ is the observed price at moment τ, $p(\tau)$ is the asset's fundamental price at moment τ, and ε_{τ} is the short-term deviation caused by the trading frictions. Some approaches to compensate for this noise attempt to estimate the fundamental price at every observation, such as the strategy proposed by Barndorff-Nielson and Shephard (2005) of averaging the bid quote and the ask quote at the time of each observation to create a set of mid-quote date. Other literature focuses on limiting the impact of the noise by optimizing the length of the interval between samples, such as Aït-Sahalia, Mykland, and Zhang (2005). This paper follows the approach of optimizing the duration of sample intervals, using the common sampling interval of five minutes. This interval represents a good middle ground between maximizing the sampling frequency and minimizing the impact of the trading frictions on the measurement. The first set of time-series data is therefore sampled into five-minute increments, which in turn sets the sampling interval M equal to 77.

5.2 Methodology

The statistical framework of this paper is based upon the ordinary least squares regression of several contemporaneous factors: the change of the current day's ex-post log variance risk premium from that of the previous day, the occurrence of a jump in the current trading day, and the direction and relative magnitude of any such jump.

For the purposes of this paper's statistical analysis, the dependent variable of each of the regressions is set as the nominal difference between the ex-post Log Variance Risk Premiums of variance swaps entered into at the close of day t and at the close of day $t-1$ respectively. This paper refers to the change in Log Variance Risk Premium from day $t-1$ to day t, $\Delta LVRP_t$, as

$$\begin{aligned}\Delta LVRP_t &= LVRP_t - LVRP_{t-1} \\ &= \ln\left(\frac{RV_{t,t+30}}{SW_{t,t+30}}\right) - \ln\left(\frac{RV_{t-1,t+29}}{SW_{t-1,t+29}}\right),\end{aligned}$$

where the 30-day variance swap rate for day t, $SW_{t,t+30}$, is taken to be the arithmetic average of the 26 implied volatilities given by Option Metrics 30-day volatility data for day t.

By the conventions of the financial markets, in computing the payout of a variance swap, the realized variance of an asset over a 30-day period that begins at the close of day t ($RV_{t,t+30}$) is equal to the sum of the log differences of the asset's daily closing prices throughout that period. Therefore, the *ex post RV* corresponding to a 30-day variance swap created at the close of day t is defined as

$$RV_{t,t+30} = \sum_{i=1}^{30} (\ln(P_{t+i}) - \ln(P_{t+i-1}))^2,$$

where P_n equals the closing price of the asset on day n. $RV_{t,t+30}$ looks forward 30 days, and is therefore not known until day t+30. Furthermore, while the New York Stock Exchange (NYSE) is not open every day, in the interest of ensuring contemporaneous regressions, this paper will consider the closing price on any day the market is closed to be the closing price from the previous trading day. In this way, we can ensure both that each $RV_{t,30}$ includes all of the daily returns for its respective 30-day period, and that there is a corresponding $RV_{t,30}$ for days in which the CBOE is open and the NYSE is closed.

Furthermore, the presence of a jump is a dummy variable in some regressions. If any returns are greater than the cutoff described earlier, there is deemed to have been a jump.

Otherwise, there is deemed not to have been any jumps. If there is a day in which the NYSE is open while the CBOE is closed, any jumps that occur that day are "rolled over" into the next day.

6 Findings

6.1 Presence of a Jump

Insofar as the participants in the financial markets include the most recent intraday price movements in their respective forecasts of volatility, we would expect the financial firms selling variance swaps to attempt to correct their volatility forecasts for any jump that occurs throughout day. This should increase the swap rate of any variance swap made after the jump. If the jump occurred on day t, then the information regarding the presence of a jump is priced into a variance swap created at the end of day t, but is obviously not priced into such a swap created at the end of day $t-1$, before knowledge of the jump ever existed. While the swap rates are different, the two variance swaps have payouts based on the realized variance over most of the same days. For this reason, we expect the *ex post* payout of buying a variance swap after a jump on day t to be less than that of buying variance swap on day $t-1$. In order to test this expectation, we perform the regression

$$\Delta \widehat{LVRP}_t = \alpha + \beta \times jump\_binary_t,$$

where $jump\_binary_t$ is a dummy variable for presence of a jump on day t. The results of this regression are in Table 1.

In testing the null hypothesis $H_0: \alpha = 0$ against the alternative $H_1: \alpha \neq 0$, we cannot reject the null hypothesis at the 95% confidence interval for any of the 10 equities tested. This result is expected, for the Log Variance Risk Premium should not change significantly between two days without the presence of a jump.

When we tested the null hypothesis $H_0: \beta = 0$ against the alternative $H_1: \beta \neq 0$, we found that we can reject the null hypothesis at the 95% confidence level for each of the ten equities included in the regression. This means that the presence of a jump on day t has a statistically significant effect on the change of $LVRP$ from day $t-1$ to day t. Furthermore, the coefficient found for β in each of the ten equities' respective regression is negative, suggesting that the presence of a jump on day t causes the $LVRP$ to decrease on day t. As $LVRP$ represents the continuously compounded excess return of entering into a variance

swap, the presence of a jump on day t causes the return from buying a variance swap on day t to be lower than it otherwise would have been.

6.2 Positive Versus Negative Jumps

Working on the assumption that the market participants are risk averse, and are therefore judging risk based on an autoregressive forecast of an asset's price return variance rather than on past price performance, we would assume that positive asset price jumps are priced into variance swaps with the same weight as negative asset price jumps. To test our assumption, we first change the previous regression by separating the presence of positive and negative asset price jumps into two mutually exclusive exogenous variables. The regression becomes

$$\Delta L\widehat{VRP}_t = \alpha + \beta_1 \times neg\_jump\_binary_t + \beta_2 \times pos\_jump\_binary_t,$$

where $neg\_jump\_binary_t$ is a dummy variable for the presence of a negative jump on day t, and $pos\_jump\_binary_t$ is a dummy variable for the presence of a positive jump on day t. The results of this regression are in Table 2.

Because we know that $\beta_1 \times neg\_jump\_binary_t + \beta_2 \times pos\_jump\_binary_t$ from this regression is a linear combination of $\beta \times jump\_binary_t$ from the first regression, we know that the set of α values is shared among both regressions, and therefore that in testing the null hypothesis $H_0: \alpha = 0$ against the alternative $H_1: \alpha \neq 0$, we cannot reject the null hypothesis at the 95% confidence interval for any of the ten equities tested. Again this result assures us that the Log Variance Risk Premiums did not change significantly between two days without the presence of a jump.

In testing the null hypothesis $H_0: \beta_1 = \beta_2$ against the alternative hypothesis $H_1: \beta_1 \neq \beta_2$, we are unable to reject the null hypothesis at the 95% confidence level for nine of the ten equities we included in the regression, meaning that positive and negative jumps do not affect the $LVRP$ differently than each other in a statistically significant way. This result is especially interesting because the well-documented leverage effect – the positive correlation between negative return surprises and volatility [e.g., Nelson (1991), and Ghysels, Harvey, and Renault (1996)] – suggests that this should not be the case. According to the leverage effect, a negative price jump should increase the volatility of the asset by a greater amount than a positive price jump of the same magnitude.

6.3 Relative Size of Jumps

It stands to reason that a relatively larger asset price jump occurring on a given day should be weighted into the asset's volatility forecasts more heavily than a relatively smaller price jump occurring on another day. To test this logic, we run the regression

$$\Delta \widehat{LVRP_t} = \alpha + \beta \times jump\_relative_t,$$

where jump_relative$_t$ is the magnitude of the largest jump on day t relative to the previous day's bi-power variation, BV_{t-1}, such that

$$\text{jump\_relative}_t = \max_{j=1\to M} \left| \frac{r_{t,j}}{\sqrt{\frac{1}{M} BV_{t-1}}} \right|.$$

The results of this regression are in Table 3.

Once again, when we test the null hypothesis $H_0: \alpha = 0$ against the alternative $H_1: \alpha \neq 0$, we find that we cannot reject the null hypothesis for any of the ten equities, reassuring us that LVRP is behaving as we would expect.

When we tested the null hypothesis $H_0: \beta = 0$ against the alternative $H_1: \beta \neq 0$, we found that we can reject the null hypothesis at the 95% confidence level for each of the ten equities included in the data set. Furthermore, seven out of the ten equities had higher test-statistic values in this regression than their respective test-statistic values in the regression of the presence of jumps alone. This suggests that the size of jumps has a greater effect on $LVRP$ than the presence of jumps alone.

7 Conclusion

This paper contributes to the current understanding of variance risk premia in two main ways. First, it proposes a quick and easy method to estimate the fair value of a variance swap. This new method, which appears to be quite accurate in its estimation, makes the investigation of the variance risk premia of individual equities feasible. Whereas few stocks have options with enough liquidity to use the CBOE's estimator, the requisite data for the simple estimator proposed herein is readily available for most, if not all, stocks traded on the NYSE. This should greatly enhance the ability of other researchers to test the characteristics of variance risk premia. Secondly, the findings of this paper confirm many of the expectations held about how asset price

jumps affect the variance risk premia of individual equities. Specifically, we have presented evidence that the variance risk premia of individual equities respond to price jumps of their respective equities in much the same way that the variance risk premia of the market index respond to index price jumps. Furthermore, this result suggests that the variance risk premium of an individual equity is not solely dependent upon the equity's proportion of the variance risk of the market as a whole. This inference, drawn from the evidence in this paper, contradicts the conclusion of Carr and Wu (2009) that much of the variance risk premia of individual equities is a result of the market's variance risk.

It is important, however, to keep the results of this paper in context. First of all, there were several days during which multiple jumps occurred. For the purpose of this paper, however, only the largest jump of the day was recorded. It is possible that the presence of multiple jumps in a single day could have affected the variance risk premium much more than the presence of a single jump would, which would in turn make the effect of the jump appear to be much more statistically significant. It may therefore be beneficial to include a factor in future regressions to compensate for the presence of multiple jumps in the same day. Also, while this paper used high-frequency intraday asset price information, the only available Option Metrics volatility data was based on the daily closing prices of options. If an asset price jump occurred early enough in the day, it is possible that the jump's effect on the asset's volatility began to wane, which would result in the jump not getting enough credit. If intra-daily option data becomes available, further research should be done to investigate how long a jump's effects on variance risk premia last.

These caveats aside, the research presented in this paper provides strong evidence that asset price jumps have significant effects on the variance risk premia of individual equities. Furthermore, the methodologically simple estimator proposed within this paper has thus far proved to be an efficient tool for accurately estimating the variance risk premium of an individual equity.

A References

Aït-Sahalia, Y., Mykland, P., and Zhang, L. (2005). "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data." *Journal of the American Statistical Association* 100: 1394-1411.

Andersen, T., and Bollerslev, T. (1998). "Answering the Skeptics: Yes Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review* 39.4: 885-905.

Andersen, T., Bollerslev, T., and Diebold, F. (2002). "Parametric and Nonparametric Volatility Measurement." In Y. Aït-Sahalia and L. P. Hansen (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland.

Bakshi, G, and Kapadia, N. (2003a). "Delta-Hedged Gains and the Negative Market Volatility Risk Premium." *The Review of Financial Studies* 16.2: 527-566.

Bakshi, G, and Kapadia, N. (2003b). "Volatility Risk Premiums Embedded in Individual Equity Options: Some New Insights." *The Journal of Derivatives* 11.1: 45-54.

Bakshi, G., and Madan, D. (2006). "A Theory of Volatility Spreads." *Management Science* 52.12: 1945-1956.

Barndorff-Nielsen, O., Graversen, S., Jacod, J., and Shephard, N. (2006). "Limit Theorems for Bipower Variation in Financial Econometrics." *Econometric Theory* 22.4: 677-719.

Barndorff-Nielsen, O. and Shephard, N. (2003). "Realized Power Variation and Stochastic Volatility Models." *Bernoulli* 9.2: 243-265.

Barndorff-Nielsen, O. and Shephard, N. (2004a). "Econometric Analysis of Realized Covariation: High Frequency Covariance, Regression and Correlation in Financial Economics." *Econometrica* 72: 885–925.

Barndorff-Nielsen, O. and Shephard, N. (2004b). "Power and Bipower Variation with Stochastic Volatility and Jumps." *Journal of Financial Econometrics* 2.1: 1-37.

Barndorff-Nielsen, O. and Shephard, N. (2005). "Variation, Jumps, Market Frictions, and High-Frequency Data in Financial Econometrics." *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress (Econometric Society Monographs)*. Cambridge University Press, 2007. 328-372.

Barndorff-Nielsen, O. and Shephard, N. (2006). "Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation." *Journal of Financial Econometrics* 4 .1: 1-30.

Barndorff-Nielsen, O., Shephard, N., and Winkel, M. (2006). "Limit Theorems for Multipower Variation in the Presence of Jumps in Financial Econometrics." *Stochastic Processes and their Applications* 116: 796-806.

Bates, D. (1996). "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies* 9.1: 69-107.

Black, F. and Scholes, M. (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81.3: 637-54.

Bollerslev, T. (1986). "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31.3: 307-27.

Carr, P., Lee, R., and Wu, L. (2011). "Variance Swaps On Time-Changed Lèvy Processes." *Finance and Stochastics, Springer* 9.4: 453-475.

Carr, P. and Madan, D. (1998). "Towards a Theory of Volatility Trading." In Jarrow, R. (ed.) *Volatility*, 417-427. Risk Publications, London.

Carr, P. and Wu, L. (2006). "A Tale of Two Indices." *Journal of Derivatives* 13.3:13-29

Carr, P. and Wu, L. (2009). “Variance Risk Premiums.” *The Review of Financial Studies* 22.3: 1311-1341.

Corsi, F., Pirino, D., and Renò, R. (2010). “Threshold Bipower Variation and the Impact of Jumps on Volatility Forecasting.” *Journal of Econometrics* 159.2: 276-288.

Dupire, B. (1992). *Arbitrage Pricing with Stochastic Volatility*. Société Générale, Paris

Dupire, B. (1993). “Model art.” *Risk* 6.9, 118-124.

Engle, R. (1982). "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica (pre-1986)* 50.4: 987-1007.

Eraker, B. (2004). “Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices.” *Journal of Finance* 59.3: 1367-1404.

Fama, E. (1965). “The Behavior of Stock-Market Prices.” *Journal of Business* 38.1: 34-105.

French, K., Schwert, G., and Stambaugh, R. (1987). “Expected Stock Returns and Volatility.” *Journal of Financial Economics* 19: 3-29.

Ghysels, E., A. C. Harvey, and E. Renault. (1996). ‘‘Stochastic Volatility.’’ In C. R. Rao and G. S. Maddala (eds.), *Statistical Methods in Finance*, Amsterdam: North-Holland.

Huang, X., and Tauchen, G. (2005). “The Relative Contribution of Jumps to Total Price Variance.” *Journal of Financial Econometrics* 3.4: 456-499.

Karatzas, I. and Shreve, S. (1988). *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag.

Lee, S., and Mykland, P. (2008). “Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics.” *The Review of Financial Studies* 21.6: 2535-2563.

Mancini, C. (2009). “Non-parametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps.” *Scandinavian Journal of Statistics* 36.2: 270-296.

Mandelbrot, B. (1963). "The Variation of Certain Speculative Prices." *The Journal of Business* 36.4: 394-419.

Merton, R. (1971). “Optimum Consumption and Portfolio Rules in a Continuous-Time Model.” *Journal of Economic Theory*, 3: 373-413.

Nelson, D. (1991). ‘‘Conditional Heteroskedasticity in Asset Returns: A New Approach.’’ *Econometrica* 59: 347–370.

Neuberger, A. (1994). “The Log Contract.” *Journal of Portfolio Management* 20: 74-80.

Pan, J. (2002). “The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study.” *Journal of Financial Economics* 63.1: 3-50

Todorov, V. (2010). “Variance Risk-Premium Dynamics: The Role of Jumps.” *Review of Financial Studies* 23.1: 345-383.

B Tables

Table 1 – Results from regressing the presence of an asset price jump on day t on the difference in payouts of variance swaps purchased at the end of days t and t-1.

	Intercept			jump_boolean			
Stock	α	Std. Err.	t	β	Std. Err.	t	R^2
AA	0.0030	0.0026	1.14	-0.1260	0.0169	-7.45	0.0159
AAPL	0.0012	0.0030	0.39	-0.0621	0.0183	-3.39	0.0034
AMZN	0.0033	0.0035	0.95	-0.1504	0.0225	-6.70	0.0135
BA	0.0027	0.0026	1.03	-0.1174	0.0168	-7.00	0.0141
BAC	0.0021	0.0028	0.76	-0.1427	0.0227	-6.29	0.0114
CAT	0.0015	0.0027	0.55	-0.1313	0.0222	-5.92	0.0101
DELL	0.0006	0.0027	0.23	-0.0649	0.0238	-2.72	0.0022
GE	0.0012	0.0027	0.45	-0.0971	0.0219	-4.44	0.0057
GOOG	0.0007	0.0045	0.15	-0.0518	0.0251	-2.06	0.0027
GS	0.0031	0.0029	1.07	-0.1924	0.0218	-8.82	0.0267

Table 2 – Results from regressing the presence of either a positive or a negative asset price jump on day t on the difference in payouts of variance swaps purchased at the end of days t and t-1.

	Intercept			neg_jump_boolean			pos_jump_boolean			
Stock	α	Std. Err.	t	B_1	Std. Err.	t	B_2	Std. Err.	t	Adj R^2
AA	0.0030	0.0026	1.14	-0.1167	0.0254	-4.60	-0.1333	0.0225	-5.93	0.0154
AAPL	0.0012	0.0030	0.39	-0.0517	0.0243	-2.13	-0.0753	0.0274	-2.75	0.0029
AMZN	0.0033	0.0035	0.95	-0.1698	0.0358	-4.74	-0.1381	0.0286	-4.83	0.0130
BA	0.0027	0.0026	1.03	-0.1260	0.0237	-5.31	-0.1090	0.0234	-4.65	0.0136
BAC	0.0021	0.0028	0.76	-0.2093	0.0305	-6.86	-0.0622	0.0335	-1.86	0.0139
CAT	0.0016	0.0027	0.58	-0.1404	0.0326	-4.31	-0.1312	0.0295	-4.44	0.0108
DELL	0.0006	0.0027	0.23	-0.1378	0.0365	-3.77	-0.0116	0.0313	-0.37	0.0036
GE	0.0012	0.0027	0.45	-0.0780	0.0321	-2.43	-0.1134	0.0297	-3.82	0.0053
GOOG	0.0007	0.0045	0.15	-0.0465	0.0331	-1.40	-0.0588	0.0379	-1.55	0.0015
GS	0.0031	0.0029	1.07	-0.2272	0.0323	-7.03	-0.1639	0.0293	-5.59	0.0267

Table 3 – Results from regressing the magnitude of an asset price jump on day t (relative to BV_{t-1}, the bipower variation of the previous day) on the difference in payouts of variance swaps purchased at the end of days t and t-1.

	Intercept			Jump_return			
Stock	α	Std. Err.	t	β	Std. Err.	t	R^2
AA	0.0021	0.0026	0.79	-5.1573	0.8188	-6.30	0.0114
AAPL	0.0016	0.0030	0.53	-3.4143	0.5725	-5.96	0.0103
AMZN	0.0031	0.0035	0.88	-6.5859	0.9043	-7.28	0.0159
BA	0.0030	0.0026	1.16	-8.8996	0.9856	-9.03	0.0232
BAC	0.0015	0.0028	0.55	-5.1351	0.8574	-5.99	0.0103
CAT	0.0014	0.0027	0.53	-7.8448	1.2677	-6.19	0.0110
DELL	0.0005	0.0027	0.18	-2.7560	0.9168	-3.01	0.0026
GE	0.0008	0.0027	0.31	-4.2497	1.0399	-4.09	0.0048
GOOG	0.0008	0.0045	0.17	-4.1239	1.4374	-2.87	0.0052
GS	0.0029	0.0029	1.01	-9.0112	0.7945	-11.34	0.0433

C Figures

Figure 1 – Comparison of the levels found using this paper's methodology for estimating the market's predicted variance over the next 30 days and the levels founds using the CBOE's methodology for estimating the market's predicted variance, as reported in the VIX index.

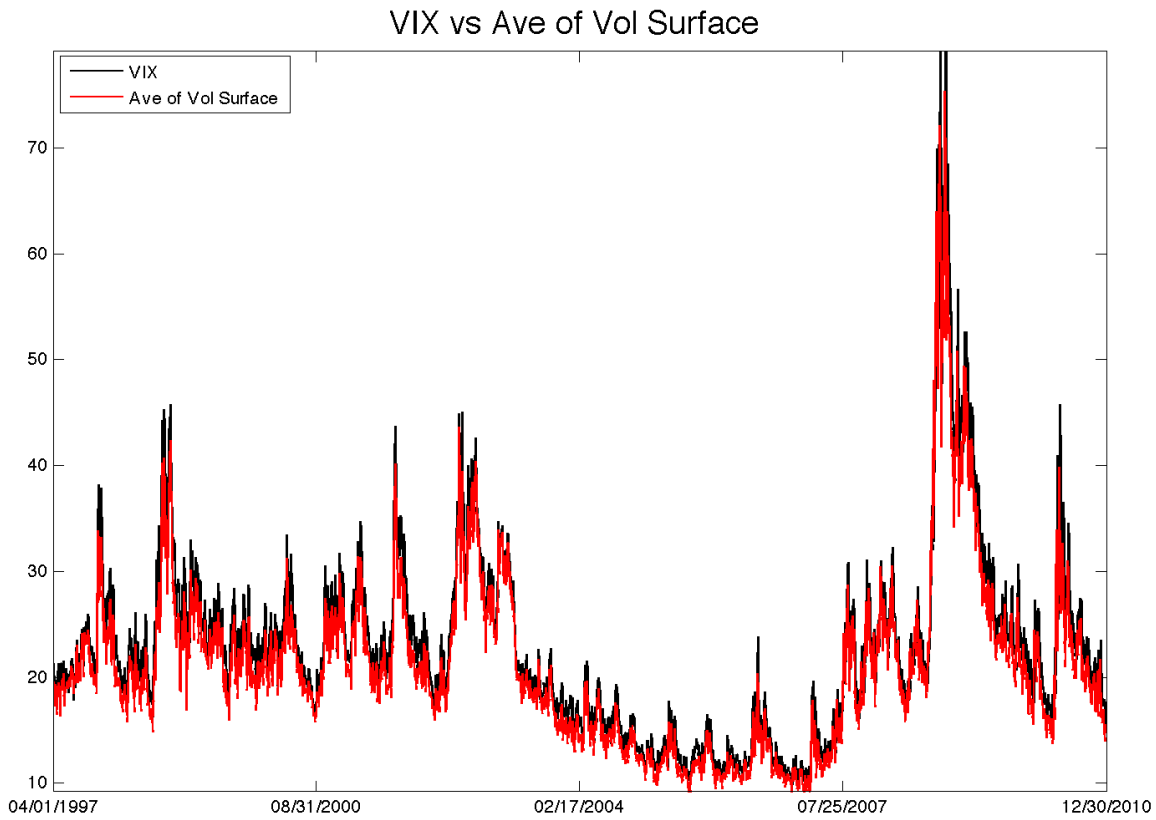

D Explanation of Swaps

A swap is a type of over-the-counter financial instrument; it is a contractual agreement between two parties to exchange a sequence of cash flows over a set period of time. In a basic swap, one of the two cash flows is fixed and the other is based on an uncertain variable, such as interest rates, currency exchange rates, commodity prices, etcetera. At the time of each payment, the two cash flows net out so that the party whose payment is greatest pays the other party the difference between the two payments.

To illustrate how swaps work, let us assume two companies, firm A and firm B, enter into an interest rate swap on January 1, 20X1 with the following terms:

- Firm A is to pay firm B the equivalent of a 5% interest payment on a notional amount of $1 million every year for five years.
- Firm B is to pay firm A the equivalent of an interest payment equal to the London Inter-Bank Offer Rate (LIBOR) one-year interest rate, a rate which varies to reflect the rate at which banks are willing to lend to each other for one year, present on January 1 on a notional amount of $1 million every year for five years.

Assuming the one-year LIBOR was equal to 4.5% on January 1, 20X1, then on January 1, 20X2 firm A would owe firm B $50,000 and firm B would owe firm A $45,000. These payments would net each other out, such that firm A would pay firm B $5,000 and firm B would not pay anything to firm A. Furthermore, assuming the one-year LIBOR on January 1, 20X2 were equal to 6%, then on January 1, 20X3 firm B would pay firm A $10,000 and firm A would not pay anything to firm B. This process would continue for the life of the swap.

Corporations often have a large amount of debt with a variable interest rate. This gives managers at those companies big headaches as they attempt to create an operating budget for their company; the variable interest rate means that the company's interest burden is subject to change from month to month, making it difficult to plan for how much money the company needs to set aside to pay its lenders. Furthermore, those companies face the risk that the interest rate will increase significantly enough in the future that, much like the homeowners in the recent mortgage crisis who had adjustable rate mortgages and could not afford their mortgage payments after interest rates increased, they will be unable to afford to pay their lenders and will be forced into bankruptcy. Banks, on the other hand, often borrow money at a fixed rate but would prefer

to have variable rate debt (lest interest rates they are able to charge drop below the rate at which they borrowed their money). For these reasons, banks and corporations occasionally enter into interest rate swaps, through which they essentially "swap" the fixed rate for the variable rate and vice versa.

In 1981, IBM and the World Bank made history when they entered into the first ever interest rate swap. At that time, IBM had a large amount of debt denominated in Swiss Francs and Deutsche Marks. The US based company's debt being based in foreign currencies meant that its interest burden was dependent upon the exchange rates between those currencies and the US Dollar, putting the company at risk of the US Dollar depreciating in value. The World Bank had the opposite problem; most of the bank's debt was denominated in US Dollars while the loans it made were to foreign nations, putting the World Bank at risk of the US Dollar appreciating in value. The World Bank therefore "swapped" its US debt for IBM's foreign debt.