What Causes the Volatility Asymmetry?
– A Clue to the Volatilities Puzzle

Woo Chang Kim[1]
John M. Mulvey[2]

Department of Operations Research & Financial Engineering
Princeton University

July 29, 2007

Abstract

The skewness of the return distribution causes the asymmetric movements of volatilities to the contemporaneous returns. Employing the skew normal distribution, we propose a model that estimates the measures for the volatility asymmetry from the return distribution skewness. We find that the sign and the magnitude of the asymmetry measures can be determined by the skewness. Empirical evidence of DJIA from 1935 to 2006 supports the hypothesis. This relationship between the skewness and the volatility asymmetry provides a clue to the unsettled volatility puzzle: the market risk premium and the conditional market risk have a positive relation. The empirical evidence contradicting the theoretical prediction may have been caused by the negatively skewed return distribution of the market portfolio.

I. Introduction

In finance, it is widely accepted wisdom that the volatility changes over time. There are many theoretical volatility models that try to capture its various aspects. Prominent examples include the continuous time model by Hull and White (1987) and the GARCH model by Bollerslev (1986) for the discrete time setting. One of interesting characteristics of the equity return volatility is the asymmetric movements. For instance, Harvey (1989), Turner et al. (1989) and Duffee (1995) report a positive contemporaneous relation between volatilities and returns, while French et al. (1987), Campbell (1988), Nelson (1991), Campbell and Hentschel (1992) and Glosten et al. (1993) find evidence that they posses a negative relation[3]. In response to such empirical works, Nelson (1991) proposes the EGARCH model which utilizes the weighted innovation process to reproduce the asymmetry. Also, Heston (1993) constructs a continuous-time stochastic volatility model which replicates the asymmetry by allowing the Wiener processes for the SDE's of the return and the volatility correlated.

---

[1] Ph.D. Candidate, ORFE, Princeton University, contact: wookim@princeton.edu
[2] Professor, ORFE, Princeton University, contact: mulvey@princeton.edu
[3] See Figure 1 for the example of the negative relation of returns and volatilities. It depicts the prices and both the realized and implied volatilities of the Dow Jones Industrial Average index (DJIA) and the S&P 500 index from 1998 to 2005. It is evident that both realized and implied volatilities are negatively correlated with the price (thus, return) processes.

Figure 1
Price, Realized Volatility and Implied Volatility of DJIA (Left) and SP500 (Right)
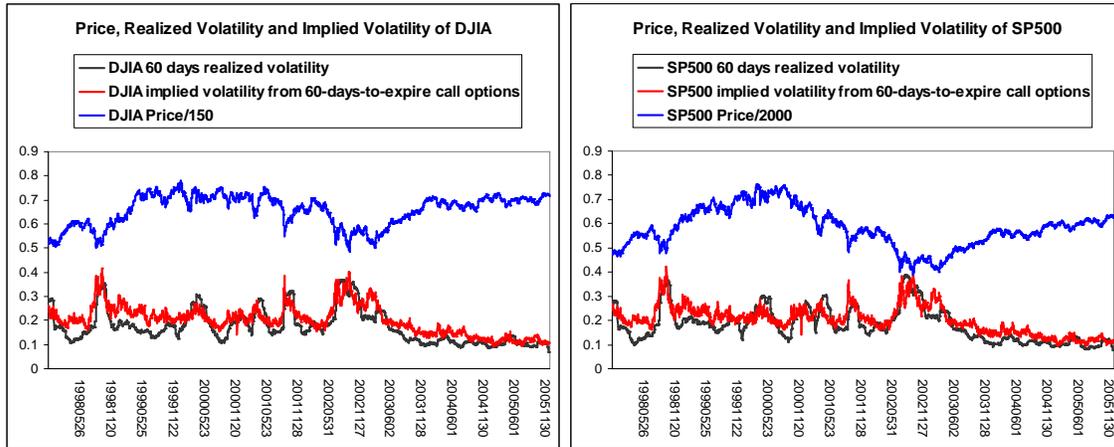
Efforts regarding the relationship between the contemporaneous returns and the realized conditional volatilities, which is often called the volatility feedback effect, have been mainly focused on identifying and calibrating proper stochastic volatility models. For instance, Scruggs (1998) employs a conditional two-factor ICAPM model which includes the long-term government bond returns as the second factor, under the assumption that the volatility process follows the EGARCH model. Bollerslev and Zhou (2006) explain the relationship with the Heston Model. Also, Smith (2006) proposes a modified stochastic volatility model that includes the volatility feedback effect as its stylized feature. However, in spite of the immense attention from both academics and practitioners, no conclusive explanation on what causes such an asymmetry of the volatility movement with respect to the contemporaneous returns has yet been made.

The main contribution of this paper is to identify the fundamental aspect of the asymmetric volatility movements. We advocate that the source of the asymmetric volatility movements is the skewness of the return distribution. Evidence of the skewed return distribution in the equity market has been reported repeatedly (Singleton and Wingender (1986), Harvey and Siddique (1999, 2000), Brannas and Nordman (2003), Rydberg and Shephard (2003) and Cappuccioa et al. (2006)). We find that the skewness of the return distribution determines the sign and the magnitude of the return-volatility relationship: when the return distribution is positively (negatively) skewed, the correlation between the realized volatilities and the returns are positive (negative). Furthermore, the magnitude of the asymmetry increases as the absolute value of the skewness increases. By adopting the skew normal distribution discussed in Azzalini and Dalla Valle (1996), a model that explicitly yields the asymmetry from the skewness of the underlying return distribution is proposed. Also, an empirical analysis for the model is conducted on the daily return data of DJIA from 1935 to 2006.

This framework provides a clue to one of the long unsettled puzzles: the market returns and the contemporaneous volatilities relation. Many theoretical asset pricing models

advocate a positive contemporaneous return-risk relation of the market portfolio. For instance, under the assumption of the single factor capital asset pricing model (CAPM), Merton (1980) and Pindyck (1984) have argued that the risk premium of the market portfolio has a positive linear relation with the market risk, which is measured by the market variance. However, as mentioned previously, the relationship between the market returns and the volatilities is insignificant or even significantly negative (French et al. (1987), Campbell (1988), Nelson and Glosten et al. (1993)). We claim that the opposing empirical results to the theoretical framework might have been caused by the skewness of the return distribution: Since the measure of the market risk, the variance is an intrinsic value, it cannot be directly measured. The conventional measure of the realized variances is in fact conditioned on the realized returns. Thus, when the return distribution is negatively skewed, the negative variance-return relation may be observed, even if the true risk has a positive relation to the return.

The remainder of the paper is organized as follows. The following section explains the relationship of the skewness to the volatility asymmetry. Section III provides the results from the empirical test. Section IV discusses the market volatility puzzle in the context of the skewed return distribution. Section V concludes the paper.

II. Asymmetries in the Volatility Movements and the Skewness of the Return Distribution

In many articles, the asymmetric movements of volatilities (σ_t) to the contemporaneous returns (r_t) are measured by the correlation between the return and the volatility (corr(r_t, σ_t)) or the slope from the simple regression (β in $\sigma_t = \alpha + \beta \cdot r_t$). By convention, we call the correlation and the slope the asymmetry correlation (AC) and the asymmetry slope (AS), respectably.

Asymmetry Correlation (AC) = corr(r_t, σ_t)

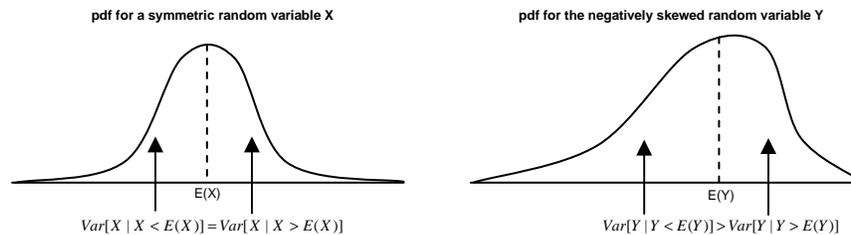Asymmetry Slope (AS) = β from the simple regression equation $\sigma_t = \alpha + \beta \cdot r_t$

The direction of the asymmetry is determined by the sign of AC and AS, and the magnitude is measured by their absolute values: when the sign is positive, the return and the volatility have a positive relation, and their relation is considered to be strong if the magnitude is large.

In order to have some intuition, let's consider simple examples that illustrate how the skewness of the return distribution affects the relationship between the volatilities and the contemporaneous returns. First, suppose a random variable X has a symmetric distribution as illustrated on the left of Figure 2. By the symmetry, it is obvious that the variance conditioned on the left half of the distribution, $Var[X \mid X < E(X)]$ is the same as the one conditioned on the right tail, $Var[X \mid X > E(X)]$. Interpreting the random

variable X as the random return, the equality of the conditional variances implies that the upside volatility is equivalent to the downside one.

On the other hand, when the distribution is skewed, the equivalence of the volatilities conditioned on the left and the right tail disappears. To see this, consider another random variable Y, whose distribution is negatively skewed as illustrated on the right of Figure 2. Since the negative skewness means that the left tail is longer than the right, it is easy to see that the variance conditioned on the downside movements, $Var[Y\,|\,Y < E(Y)]$ is larger than the upside conditional variance $Var[Y\,|\,Y > E(Y)]$. Therefore, if the return distribution is negatively skewed, the volatilities measured during the periods of the stock price dropping are more likely to be higher than the ones during the upturns. Also, it is evident that the relationship between returns and volatilities would be reversed if the distribution is positively skewed.

Figure 2
Simple Illustrations of the Effects of the Return Distribution Skewness to the Conditional Variances

pdf for a symmetric random variable X

pdf for the negatively skewed random variable Y

E(X)

E(Y)

$Var[X\,|\,X < E(X)] = Var[X\,|\,X > E(X)]$

$Var[Y\,|\,Y < E(Y)] > Var[Y\,|\,Y > E(Y)]$

Although the illustration above provides an intuitive example for the effect of the skewness, it does not fully reflect the actual mechanism; we are interested in the connections between the volatility (σ_t) and the *contemporary* return (r_t), not the volatility conditioned on the specific state of the underlying return. The main difficulty of modeling the real situation is that the volatility is not directly observable. Since the volatility is an intrinsic value, it only can be measured via directly observable data such as the returns and the prices of the derivatives. For instance, the implied volatility is estimated from the option prices, by inverting the Black-Scholes option pricing formula. Alternatively, one can employ the realized volatility. For a given return, it is obtained by calculating the sample standard deviation of the intra-returns. For example, the realized volatility of the monthly return in July of 2007 is the standard deviation of the daily returns *within* the month. If the high frequency data is available, one can acquire the value from the minute-by-minute data, or even from the tick-by-tick data. In other words, the realized volatility for a given period is the sample standard deviation of the observations within the time frame. So, we can formalize the circumstance with the following terms. For simplicity, we assume the unit time length is 1. Then, the return

during the *t*-th interval (r_t) is the sum of *n* returns ($r_{t,1}, r_{t,2}, ..., r_{t,n}$), which represent the returns for the *n* evenly divided sub-intervals of the *t*-th interval. That is,

$$r_t = \sum_{i=1}^{n} r_{t,i} \text{ for } t \in \{1,2,...,T\}, \text{ where}$$

$$r_t = \ln[P_t / P_{t-1}],$$

$$r_{t,i} = \ln[P_{t+i/n} / P_{t+(i-1)/n}] \text{ and}$$

P_s is the stock price observed at $s \in \{0,1,2,...T\}$.

Then, the corresponding volatility

$$\sigma_t = \tilde{\sigma}(r_{t,1}, r_{t,2}, ..., r_{t,n}) \text{ for } t \in \{1,2,...,T\}, \text{ where}$$

$\tilde{\sigma}(\cdot)$ is the sample standard deviation operator.

Therefore,

AC = the sample correlation of the pair { r_t, σ_t } for $t \in \{1,2,...,T\}$,

AS = β from the linear regression equation $\sigma_t = \alpha + \beta \cdot r_t + \varepsilon_t$ for $t \in \{1,2,...,T\}$.

Next, we introduce the skew normal (SN) random variable (Azzalini and Dalla Valle (1996), Azzalini and Capitanio (1999)) to model the stock returns with a skewed distribution. A random variable $X \sim SN(\alpha)$, if

$$f_X(x;\alpha) = 2\phi(x)\Phi(\alpha x) \text{ for } \forall x \in R, \text{ where}$$

$\phi(x), \Phi(x)$ are the pdf and the cdf of $N(0,1)$, and $\alpha \in R$ is the shape parameter.

The shape parameter α determines the skewness of the distribution. Figure 3 illustrates the SN density functions for three different values of α.

Figure 3
Densities of the SN distribution (α =5, 0 and -5)

5
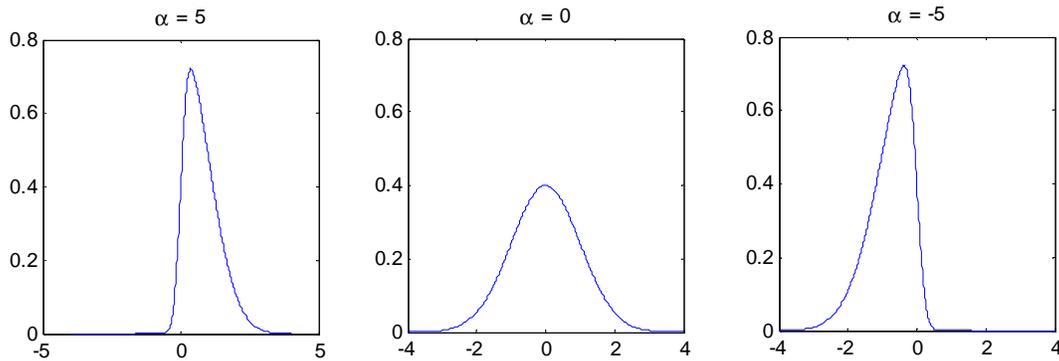
The followings are the properties of the skew normal distribution that will be exploited in the model.

a. The skewness of the distribution increases as the shape parameter (α) increases. Especially, when $\alpha = 0$, the distribution becomes standard normal.

b. $E(X) = \sqrt{2/\pi} d$, $Var(X) = 1 - 2d^2/\pi$ and $\gamma_1 = (4-\pi)E(X)^3 / 2Var(X)^{3/2}$, where γ_1 is the skewness factor and $d = \alpha/\sqrt{1+\alpha^2}$.

c. The linear transform $Y = m + wX$ is a skew normal random variable, and denoted by $Y \sim SN(m, w^2, \alpha)$.

d. For Y in 3), $E(Y) = m + wE(X)$, $Var(Y) = w^2 Var(X)$, and the skewness is γ_1.

e. The sum of the skew normal random variables is again skew normal.

f. $\gamma_1 \in (-0.995, 0.995)$.

To reveal the relationship between the volatility asymmetry and the skewness, we conduct simulations based on the skewed normal distribution, instead of deriving analytic solutions for AC and AS. The idea is as follows: 1) first, generate a set of sample returns from the SN distributions with various skewness levels, 2) then, evaluate the asymmetry measures from the samples. We make an assumption that the returns are i.i.d. skew normal for the simplicity of the tests. The simulation procedures can be summarized as follows.

1) For a given skewness level (γ_1), calculate α.

2) Since the mean and the variance of the basic skew normal distribution ($SN(\alpha)$) changes as α changes (property b), compute the location (m) and the scale parameter (w) so that the linear transformation via m and w ($Y=m+wX$), which follows $SN(m,w,\alpha)$, will have the unconditional mean 0 and the unconditional standard deviation 1. i.e. $E(Y)=0$ and $Var(Y)=1$. Note that when the skewness is set to 0, Y becomes the standard normal distribution ($N(0,1)$).

3) Construct T i.i.d. n-by-1 vectors, whose entries are i.i.d $SN(m,w,\alpha)$. T is the total number of intervals and n is the number of sub-intervals within an interval. We choose 1,000 and 10,000 for n and T, respectably. Here, each entry of every n-by-1 vector ($r_{t,1}, r_{t,2}, ..., r_{t,n})^T$ represents the return for a sub-interval. The sum of all entries of the t-th vector represents the return for the t-th interval. Namely, $r_t = \sum_{i=1}^{n} r_{t,i}$ for $t \in \{1,2,...,T\}$. By the property e, r_t is skew normal, too.

4) Evaluate the contemporaneous realized volatility at the t-th interval with the sample standard deviation of $\{ r_{t,1}, r_{t,2}, ..., r_{t,n} \}$. That is, $\sigma_t = \tilde{\sigma}(r_{t,1}, r_{t,2}, ..., r_{t,n})$ for $t \in \{1,2,...,T\}$.

5) Estimate AC and AS from r_t and σ_t obtained from 3) and 4).

6) Repeat 1) ~ 5) for various skewness levels.

Figure 4
Simulation Results (γ_1 =-0.995 (left), 0 (center) and .995 (right))
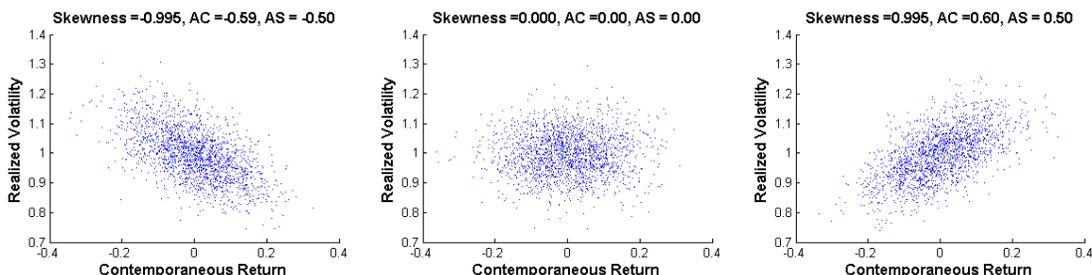

Figure 4 illustrates the results from the simulation for γ_1 =-0.995, 0 and 0.995. As we expected, both AC and AS for the symmetric case (γ_1 =0) are 0. However, when the return distribution is negatively skewed (γ_1 =-0.995), the test yields AC = -0.59 and AS = -0.5. Also, from the simulation where the samples are drawn from the positively skewed distribution (γ_1 =0.995), the asymmetry measures have the same magnitudes as the negatively skewed ones with different signs (AC=0.6 and AS=0.5). Furthermore, from Figure 5, which depicts the values of AC and AS from the tests with eleven skewness levels, it is found that the asymmetry measures are proportionate to the skewness and their relationship is almost linear.

The implication is now evident: since the volatilities can be only estimated via observed returns, and since the observations are random, the realized volatilities should appear to be random, *even when the unconditional volatility is constant*. Furthermore, since the skewness of the return distribution drives the differences in the frequencies between the left and the right tail events, it causes the volatility asymmetry. The sign of the asymmetry is the same as the sign of the skewness and its magnitude is proportionate the scale of the skewness.

Figure 5
Values of AC and AS from the Simulation
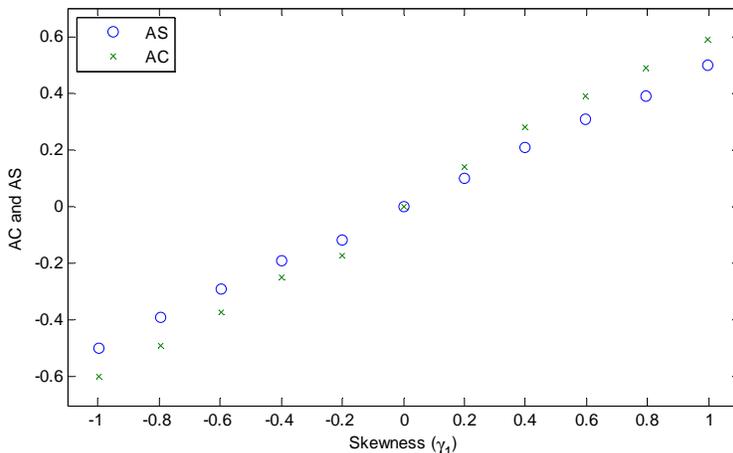
III. Volatility Asymmetries and Return Distribution Skewness: Evidence from DJIA

As illustrated in Figure 6, which depicts the sample paths of prices and volatilities generated from the skew normal distributions, the volatility asymmetry is caused by the skewness of the return distribution. In this section, we apply the model proposed in the previous section to the daily data of Dow Jones Industrial Average Index (DJIA) from 1935 to 2006.

For the empirical test, the seventy-two years long sample period is divided into thirty-six two-year long sub-periods and the asymmetry measures (AC and AS) are estimated from both the historical data and the model for each sub-period. The following describes the detailed test procedure.

1) For each two-year long sub-period, estimate the unconditional mean, the unconditional volatility and the skewness from the historical daily returns of DJIA.
2) Calculate the asymmetry measures (AC and AS) by comparing the monthly returns and the monthly volatilities ($n = 20$): the monthly return is the sum of 20 daily log-returns and the monthly volatility is the sample standard deviation of the daily returns.
3) Generate random vectors from a skew normal distribution whose unconditional mean, unconditional volatility and skewness are the same as those from the historical data: each element of the vector, the sum of all elements and the sample standard deviation of all elements of a vector correspond to the daily return, the monthly return and the monthly volatility, respectably.
4) From the randomly generated vectors, estimate AC and AS.
5) Compare the asymmetry measures from the historical data and the model.

Figure 6
Sample Paths of Log-Prices and Volatilities Generated from SN Distributions
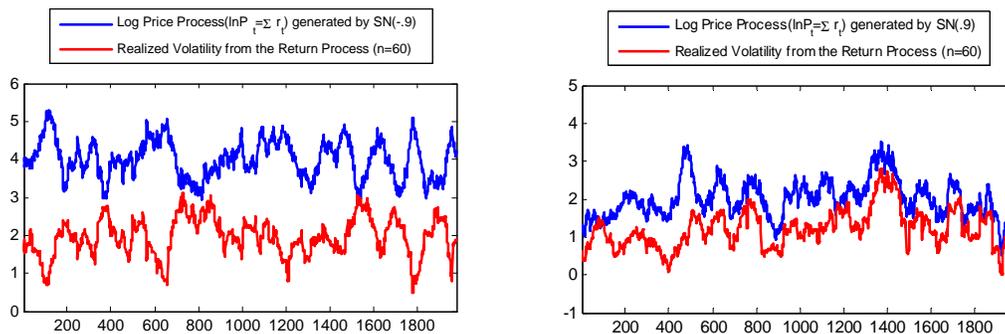**(Left: Skewness = -0.9, Right: Skewness = 0.9)**

Figure 7
Comparisons of Asymmetry Measures from Historical Data and the Model
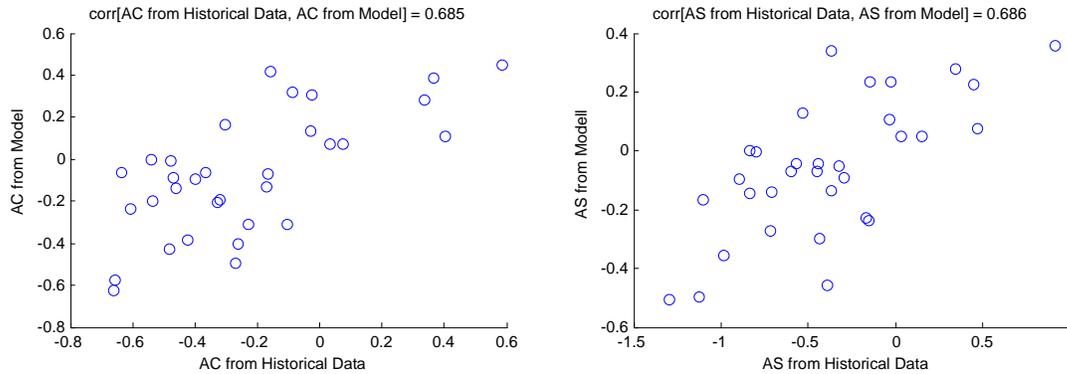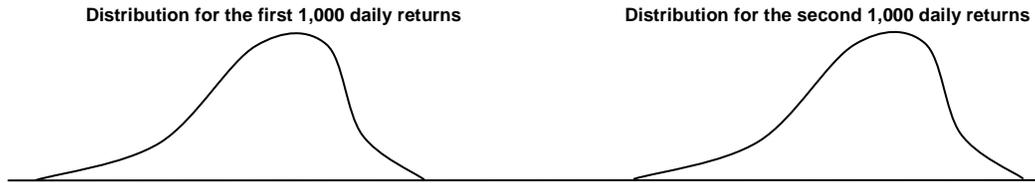**(Left: AC, Right: AS)**

Figure 7 illustrates the test results. The correlations between the theoretical and the empirical asymmetry measures are approximately 0.69, which indicates that the skewness of the return distribution plays a major role in the volatility asymmetry. However, unlike the test for AC, the magnitude of AS from the model and the data are significantly different: AS for the historical data ranges from -1.5 to 1, while the one from the model varies only from -0.6 to 0.4. This can be explained by the difference in the fatness of the tail distribution of the skew normal distribution and the empirical return distribution. It is well known that the distributions of stock returns are leptokurtic: their tails are fatter than that of the normal distribution. But since the tail of the skew normal distribution converges to 0 at the similar rate to the normal distribution ($\exp(-x^2)$), the extreme events such as "Black Monday" are not well reflected in the model. Therefore, extreme values are less likely to be drawn from the skew normal distribution, compared to the "real" distribution, so the conditional volatilities should be less that the empirical ones. Although further analysis is required regarding the magnitude of AS, considering that it does not take into account of the stylized features of the stock returns such as the serial dependence of the return series, seasonality and possible changes in the first and the second moments (expected return, unconditional volatility), the empirical tests strongly back up the hypothesis.

We close this section with issues related to estimating the skewness. Just like the volatility, the skewness cannot be observed directly. It is usually estimated from the directly observable data such as the returns, so the estimated skewness depends on the sample. Thus, as we have argued about the volatility, the estimation may not fully reflect the true value, if such a thing exists. For instance, the skewness of the daily log-returns of DJIA from 1933 to 2006[4] is -0.974. However, if the return on the day of "Black Monday" [5], October 19th 1987, is excluded, the skewness increases to -0.036. But, such an extreme event is very rare, so it may be impossible to estimate the skewness close to the true one, if the length of the sample period is not sufficiently large.

---

[4] Total number of trading days is 18,583.
[5] The log-return of DJIA on this day was -25.6%.

Figure 8
Effects of Changes in the First Moment to the Estimation of the Skewness

Distribution for the first 1,000 daily returns   **Distribution for the second 1,000 daily returns**

On the other hand, the dependence of the estimated skewness to the other moments limits the length of the sample period. To see this, consider the following extreme example: for given 2,000 daily returns, the first and the second thousands are generated from different distributions, where two distributions are equally skewed, but their first moments (expected returns) are so far from each other that no interval can have a positive probability for both of the distributions, as illustrated in Figure 8. In this case, the unconditional skewness is 0, although both distributions are negatively skewed. Since there is no clear cut for the changes in the true distribution in real life, it is almost impossible to distinguish the first 1,000 data points from the second, which may lead to biased estimations. Thus, the length of the sample period should not be too long, for it is irrational to assume that the true expected return remains constant over time.

Even worse, the skewness itself may change over time. Similar to the previous case, suppose the first 1,000 daily returns follow a skew normal distribution with the skewness of -0.5, while the second 1,000 are drawn from another skew normal with skewness of 0.5. If the first two moments of the distributions are the same, the unconditional skewness of 2,000 returns would be 0. In reality, it is hard to tell whether the returns follow a symmetric distribution or they are generated from two different distributions.

One of possible ways to decrease the bias regarding issues above is to estimate the skewness from the high frequency data with a relatively short sample length. The high frequency sampling makes it more likely to catch the extreme events, while the short sample period may reduce the possibility of changes in the distribution. Further analysis is required.

IV. A Clue to the Volatility Puzzle

The effect of the skewness to the volatility asymmetry provides an interesting insight to one of the unsolved volatility puzzles: market returns and contemporaneous volatilities relation. Accepting the CAPM is true, the market risk premium and the conditional market risk are predicted to have a positive relation. However, many articles provide evidence that the proxies of the market portfolio posses a negative risk-return relation. For instance, Campbell (1988) finds that the simple linear regression analysis yields a significantly negative risk-return relation for the CRSP value weighted index. Glosten et al. (1993) also find the similar results from the CRSP index.

The negative relation of the market return and the contemporaneous market volatility may well be caused by the flaws in the theoretical model, but it may be caused by the skewness of the index returns. In most empirical studies, the proxy of the market risk is measured by the variance of excess returns of the market portfolio. However, as we have seen in the previous sections, even when the unconditional volatility is constant over time, the conditional volatilities are negatively correlated with the returns, if the return distribution is negatively skewed. Therefore, if the returns of the market portfolio are negatively skewed, the conventional approaches estimating the volatility process may fail to reflect the true market risk-return relation. Indeed, the daily log-returns of the CRSP equal weighted index for the New York Stock Exchange from 1926 to 2006 yield a skewness of -0.519. Even after excluding the return on "Black Monday", the skewness remains considerably negative with the value of -0.15. With the skewness of -0.519, the estimated AC from the model is -0.34, which indicates that even though the market risk premium is positively correlated with the "true" market risk, the stylized feature of its proxy may have caused the negative relation in the empirical tests. Further evidence regarding the bias due to the definition of the market risk can be found in Harvey (2001) and Ghysels et al. (2005). They show that the relationship between the returns and the volatilities differs by the way to estimate the volatility.

Again, we are advocating neither the theoretical settings for the positive risk-return relation are wrong, nor the volatility is an inappropriate measure for the market risk. Rather, we suspect that the conventional approaches may be misled by the return skewness. In order to overcome this problem, one can compensate the negative relation due to the skewness from the observations over time. Another approach is to construct a new measure for the risk which is free of asymmetry caused by the return distribution skewness. For instance, instead of employing the standard deviation, one can fit the return series to the skewed parametric models such as the skew normal, the skew t (Azzalini and Capitanio (2003)) or the skew Cauchy distribution (Arnold and Beaver (2000), Behboodian et al. (2006)), and utilize the scale factor from the fitted model as the measure of the market risk.

V. Conclusion and Future Research

Since the volatilities are intrinsic values, the realized volatilities are affected by the shape of the distribution. When the return distribution is skewed, the frequencies for "good events" and "bad events" are different from each other, so the volatility shows an asymmetric response to the realized return. We find that the negative relation of the realized volatility and the contemporaneous return is caused by the negative skewness of the return distribution. The empirical evidence from DJIA supports the hypothesis. This finding gives a hint to the market risk-return puzzle. Accepting the CAPM is true, the market risk premium and the conditional market risk are predicted to have a positive relation, but empirical evidence indicates that it may not be the case. However, the test results may have been biased due to the volatility asymmetry caused by the negative skewness. Daily returns of the CRSP equal weighted index from 1926 to 2006 yield a

skewness of -0.519, which may have caused the contradicting empirical results to the theoretical prediction.

There are several possible extensions. First, since the skew normal distribution does not reflect the leptokurtic property of the stock returns, so we will apply skewed distributions with fat tails, such as skew t and skew Cauchy to the model. Second, in order to find out whether the volatility puzzle is caused by the skewness, it is required to construct a method to decompose the market risk-return relation into the volatility asymmetry and the true risk-return relation. Last, considering the importance of the skewness, the search for better estimation methods is strongly needed.

References

Arnold, B. C. and Beaver, R. J. (2000) The Skew-Cauchy Distribution. *Statisitics & Probability Letter*, **49**,285–290

Azzalini, A. and Capitanio, A. (1999) Statistical Applications of the Multivariate Skew Normal Distributions. *Journal of Royal Statistics Society*, **ser.B 61**, 579–602

Azzalini, A. and Capitanio, A. (2003) Distributions Generated by Perturbation of Symmetry with Emphasis on a Multivariate Skew t Distribution. *Journal of Royal Statistics Society*, **ser. B 65**, 367–389

Azzalini, A. and Dalla Valle, A. (1996) The Multivariate Skew-Normal Distribution. *Biometrika*, **83**, 715–726

Behboodian, J., Jamalizadeh, A., and Balakrishnan, N. (2006) A New Class of Skew-Cauchy distributions. *Statisitics & Probability Letter*, **76**, 1488–149

Bollerslev, T. (1986) Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, **31**, 307-327

Bollerslev, T. and Zhou, H. (2006) Volatility Puzzles: A Unified Framework for Gauging Return-Volatility Regressions. *Journal of Econometrics*, **131(1)**, 123-150

Brannas, K. and Nordman, N. (2003) An Alternative Conditional Asymmetry Specification for Stocks Returns. *Applied Financial Economics*, **13**, 537–41

Campbell, J.Y. (1988) Stock Returns and the Term Structure. *Journal of Financial Economics*, **18**, 373–399

Campbell, J., and Hentschel, L. (1992) No News is Good News. *Journal of Financial Economics*, **31**, 281–318

Cappuccioa, N., Lubian, D. and Raggi, D, (2006) Investigating Asymmetry in US Stock Market Indexes: Evidence from a Stochastic Volatility Model. *Applied Financial Economics*, **16**, 479–490

Duffee, G.R. (1995) Stock Returns and Volatility. *Journal of Financial Economics*, **37**, 399–420

French, K., Schwert, G.W., and Stambaugh, R. (1987) Expected Stock Returns and Volatility. *Journal of Financial Economics*, **19**, 3–30

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2005) There is a Risk-Return Tradeoff After All. *Journal of Financial Economics*, **76**, 509–548

Glosten, L., Jagannathan, R., and Runkle, D. (1993) On The Relation Between Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, **48**, 1779–1801

Harvey, C.R. (1989), Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics*, **24**, 289–317

Harvey, C. R. (2001) The Specification of Conditional Expectations. *Journal of Empirical Finance*, **8**,573–638

Harvey, C.R. and Siddique, A. (1999) Autoregressive Conditional Skewness. *Journal of Financial and Quantitative Analysis*, **34**, 465–88

Harvey, C.R. and Siddique, A. (2000) Conditional Skewness in Asset Pricing Tests, *Journal of Finance*, **55**, 1263–95

Heston, S. (1993) A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, **6(2)**, 327-343

Hull, J., and White, A. (1987) The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, **42(2)**, 281-300

Merton, R.C. (1980) On Estimating the Expected Return on the Market: An Exploratory Investigation. *Journal of Financial Economics*, **8**, 323–361

Nelson, D. (1991) Conditional Heteroscedasticity in Stock Returns: A New Approach. *Econometrica*, **59**, 347–370

Pindyck, R. (1984) Risk, Inflation and the Stock Market. *American Economic Review*, **74**, 335-351

Scruggs, J.T. (1998) Resolving the Puzzling Intertemporal Relation between the Market Risk Premium and Conditional Market Variance: A Two-Factor Approach. *The Journal of Finance*, **53(2)**, 575-603

Shephard, N. and Rydberg, T.H. (2003) Dynamics of Trade-by-Trade Price Movements: Decomposition and Models. *Journal of Financial Econometrics*, **1**, 2-25

Singleton, J.C. and Wingender, J. (1986) Skewness Persistence in Stock Returns. *Journal of Financial and Quantitative Analysis*, **21**, 335–341

Smith, D.R. (2006) Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters! *Working Paper*, Available at SSRN: http://ssrn.com/abstract=889263

Turner, C.M., Startz, R. and Nelson C.R. (1989) A Markov Model of Heteroskedasticity, Risk, and Learning in the Stock Market. *Journal of Financial Economics*, **25**, 3–22