An Investigation into the Interdependency of the Volatility of Technology Stocks[†⊥]

Zoravar Lamba

Adviser: Prof. George Tauchen

Spring 2009, Duke University

[†] The Duke Community Standard was upheld in the completion of this report.
[⊥] This report is submitted in partial completion of the Duke University Department of Economics' requirements for graduation with distinction.

ABSTRACT

This paper examines the contemporaneous and dynamic relationships between the volatilities of the technology stocks in the S&P 100 index. Factor analysis and heterogeneous autoregressive regressions are used to examine contemporaneous and dynamic, inter-temporal relationships, respectively. Both techniques utilize high frequency data by measuring stock prices every 5 minutes from 1997-2008. We find that a strong industry effect explains the bulk of the volatility of the technology stocks and that the market's volatility has very low correlation with the stocks' volatility. Further, we find the market's volatility has insignificant predictive content for the stocks' volatility. The stocks themselves contain large quantities of unique predictive content for each other's volatilities.

ACKNOWLEDGEMENTS

- Principal adviser: Prof. George Tauchen
- Research guidance also provided by:
    - Prof. Tim Bollerslev
    - Prof. Bjorn Eraker

1. INTRODUCTION

Understanding the nature of equity volatility is of paramount importance to the financial services industry. Many derivative securities and stock options in particular are priced based upon expectations of future volatility. Furthermore, in light of the recent chaos in the markets, effective risk management has become more critical than ever before, which can only be achieved with reliable estimates of volatility. Financial institutions are required to use a standard measure called Value at Risk (VaR), the estimated possible loss to the institution over the upcoming ten days that is so severe that it is expected to occur only 1% of the time (Hull 192). This measure is updated daily and is used to decide the amount of capital that should be set aside to absorb any shocks or losses. The calculation of VaR requires accurate and reliable information on the volatility of various underlying assets and securities. Moreover, the development of financial theory also extensively utilizes asset volatility forecasts. Thus, a clearer picture of equity volatility can be very beneficial and useful. In this paper, the volatility of stocks in the technology sector will be examined.

Conventional wisdom on the components of stock returns (from which volatility is measured) suggests that there is a systematic component to returns and an idiosyncratic, stock-specific component. The systematic component indicates how the return on a stock changes due to changes in the overall market and this statistic is known as the stock's beta (Stock 122). Table 1 shows the beta for every stock analyzed in this paper – the range of the betas is around 1.00 to 1.50. This means they should all exhibit higher returns and volatility when the market's return and volatility increase, and vice versa. We test this conventional wisdom by examining the contemporaneous relationship between the stocks' volatility and the market's volatility using factor analysis. This analysis also enables us to explore how well the volatilities of different series correlate with each other. Another benefit of factor analysis is that it determines how many common factors might explain the stocks' volatilities. This would be useful in narrowing down the sources of volatility in the dataset being examined.

Further, we examine the dynamic, causal relationships among the stocks' volatilities and the market's volatility by testing how well they forecast each other. It has been well established over the last several decades that stock returns follow a random walk and are very difficult to predict (Stock 665); however, forecasting volatility is a different story altogether. Bollerslev (2006) explains that volatility appears in clusters over time and that assets do not have constant volatility (variance) over time, as was previously believed. Volatility clusters over time are summarized succinctly in Bollerslev's lecture by the following quotation from Mandelbrot (1963): "… large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes …" Thus, by examining the trends in an asset's volatility, one can expect to make reasonable and educated projections regarding the asset's volatility in the future (Stock 665).

We examine how well the stocks' volatilities and the market's volatility forecast each other by conducting heterogeneous autoregressive (HAR) regressions (Corsi 2003) and draw conclusions about each stock's predictive power by examining the results of the appropriate Granger Causality Test (Granger 1969). As will be seen, Granger Causality is different from, but related to, the normal interpretation of causality and the outcome of a Granger Causality Test shows us how much predictive power a stock's volatility has for another stock's volatility, or for the market's volatility, and vice versa.

Neither of these two techniques has been employed using high frequency data before. As described in Andersen, Bollerslev, Diebold, and Labys (2003), forecasting volatility with high-frequency data provides better results than does forecasting with lower frequency data. They also find Realized Variance, defined in Section 2, to be the best measure of volatility for making forecasts.

The rest of the paper continues as follows. Section 2 explains the model of volatility (realized variance) utilized here. Section 3 contains important notes on the data used in the paper – both the source of the data, as well as data preparation are described. Sections 4 and 5 discuss the factor analysis and HAR-RV regressions, respectively. Sections 6 and 7 contain detailed discussions and interpretations of the results of factor analysis and HAR regressions, respectively. Section 8 summarizes the paper and highlights the key results.

2. A ROBUST MODEL OF VOLATILITY

We model volatility with an approach that utilizes the benefits of high frequency data and thereby contains more information than measurements taken at a large interval, such as only once a day or once a week. Intra-day geometric returns themselves are given by:

$$r_{t,j} = p(t - 1 + \frac{j}{M}) - p(t - 1 + \frac{j - 1}{M}) \qquad (1),$$

where p is the log price, t represents the day, $j = 1, 2, \ldots, M$, and $M = 78$ is the sampling frequency. A sampling frequency of 78 corresponds to five-minute returns. The next section explains in more detail that five-minute returns are being utilized in order to minimize market microstructure noise while retaining the benefits of high frequency data. Daily realized variance is given by the formula:

$$RV_t = \sum_{j=1}^{M} r^2{}_{t,j} \qquad (2),$$

where j, M, and $r_{t,j}$ are all the same as for Equation (1) above. Thus, a day's realized variance is the sum of the squared log returns for that day. Andersen and Bollerslev (1998) illustrate that realized variance converges to the integrated variance plus the jump component as the time between observations approaches zero:

$$\lim_{M \to \infty} RV_t = \int_{t-1}^{t} \sigma^2(s) ds + \sum_{t-1 < s \leq t} \kappa^2(s) \qquad (3).$$

The first term in Equation (3) represents the integrated variance of the continuous process and the second represents squared discrete jumps. This limit is also the definition of quadratic variation. Therefore, realized variance consistently estimates quadratic variation, as it accounts for both the continuous part of variation and the discrete jump components in variation.

The next part of this section explains how market microstructure noise can cause significant distortion in estimating statistics, such as a stock's volatility.

Stock prices are frequently modeled as the discounted present value of their future earnings, with the assumption that the stock's expected Earnings Per Share and expected Dividend Per Share both grow at a constant rate (Levy 498). These also result in the share price growing at the same constant rate g. If the discount rate, i.e., the required return according to the Capital Asset Pricing Model, is represented by k, then the current price P_0 of a share is given by:

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

Since obtaining this price requires calculating the sum of an infinite series, even a small change in the market's assessment of the discount rate can result in a significant effect on the estimated price and cause it to be substantially different from the stock's fundamental price (Levy 501). The discount rate can change over time as it is highly "firm-specific" and depends upon the firm's assets and liabilities, which are not constant over time. Furthermore, the growth rate g can also cause major changes in a firm's worth, as the above equation suggests. While the assumption is that the growth rate remains constant, as with the discount rate, this is frequently not borne out in reality. In general, firms will exhibit impressive growth in their infancy, but as their size increases, it becomes increasingly difficult to sustain previous levels of growth. This illustrates how the theoretical price of a stock can differ from its observed price. Moreover, market frictions such as bid-ask bounce and order imbalance can also cause a stock's observed price to deviate from its intrinsic value. This deviation is called market microstructure noise and is represented as the error term ε_t in the following:

$$p^{*}t = p(t) + \varepsilon_t \qquad (5),$$

where $p(t)$ is the logarithm of the stock's fundamental price and $p^{*}t$ is the observed price at time t. We shall see presently how the detrimental effects of this noise can be minimized.

3. DATA

As stated in the introduction, the objective is to analyze the volatility of the stocks in the technology sector. The technology stocks with the largest market capitalizations represent a good sample to examine. They have high trade volumes and so are relatively more liquid and high-frequency data for them are readily available. All data were purchased from http://price-data.com.

The data used were available at one-minute intervals; however, sampling at that interval did not appear to be very useful, as the microstructure noise clouded the true picture. Sampling at intervals larger than five-minutes would certainly reduce the noise even further; however, it also resulted in the loss of enough of the information contained in the dataset, thereby nullifying the benefits of utilizing high frequency data. Thus, sampling returns at a frequency of five-minute intervals provides a good balance between the two contrasting forces of looking for higher frequency data to maximize the information contained in the data and looking for lower frequency data to minimize the microstructure noise.

The stocks selected to represent the technology sector, broadly defined, are as follows. Cisco (CSCO) is primarily engaged in production and distribution of networking and communication devices. Dell (DELL) focuses mostly on providing desktops and laptops to private consumers and businesses. EMC Corporation (EMC) provides servers and other solutions to the problem of storing enormous quantities of data. Google (GOOG) is the world's preeminent provider of internet search services and has also expanded into other consumer-oriented applications such as web-based email. Hewlett-Packard (HPQ) provides products and services similar to Dell, in addition to technology integration consulting and outsourcing services. IBM (IBM) used to be in the personal computing business until 2004, when it sold its PC business and brand to Lenovo. Now IBM focuses on all aspects of business computing, from servers and networking hardware and software to business process consulting and outsourcing services. Intel (INTC) is the world's largest microprocessor manufacturer. Microsoft (MSFT) is the world's largest software corporation, providing an extensive suite of products for both corporate users and the general consumer. Oracle (ORCL) is the world's second largest software corporation and is best known for its database management and enterprise resource planning products. Texas Instruments (TXN) manufactures educational technology products and microprocessors. Xerox (XRX) provides all hardware and software needed for document reproduction and transmission.

These represent all the technology stocks in the S&P 100, with the exception of Apple (AAPL), as reliable data for it could not be obtained. Data for Google are only available from the time of its Initial Public Offering in 2004, but for all the other stocks, data were available from April 1997 to January 2008. Thus, we first consider the data for all the stocks other than Google, then later incorporate Google data from the point possible and see whether including Google changes any results or conclusions.

The S&P 500 futures data, which is used to represent the market, contained some shortened trading days (i.e., days with fewer than the required 78 five-minute returns). A tolerance limit of 70 was set for the required number of returns. Days with fewer than 70 five-minute returns were removed from the entire dataset, including from the history of each stock, prior to any calculations. In addition, for the factor analysis section, instead of the daily realized variations, their logs have been used. This scales the results – as shown in Equation (3), returns are squared to measure realized variation, which can really exaggerate outliers and skew the data.

4. METHOD ONE: FACTOR ANALYSIS

This section is divided into three subsections – Section 4.1 provides some historical context and describes the situations in which one might chose conduct factor analysis; Section 4.2 explains how factor analysis actually works and how its results are interpreted; finally, Section 4.3 describes some methods to rearrange the results of a regular factor analysis, which can make interpreting the results of a factor analysis easier.

4.1 The purpose of Factor Analysis

The technique of factor analysis, while initially pioneered for the domain of psychology in 1905 by Spearman, has a wide range of applications in economics and finance. Gorsuch (1976) states that there are three major purposes of employing factor analysis: first, that "the number of variables for further research can be minimized while also maximizing the amount of information in the analysis"; second, that it enables us determine manageable hypotheses when faced with a huge amount of raw data; and third, that it enables us to test preexisting hypotheses about some data, i.e., whether or not it possesses certain qualities or distinctions. Thus, it would be a good fit for examining the parallel variables being studied here, particularly since we do not know what causal relationships are at play here.

4.2 Theory, Technique, and Interpretations

Next is an examination of exactly what factor analysis is, precisely what results it provides, and how they are to be interpreted. Kim and Mueller (1978a) explain that factor analysis involves recreating the variables to be analyzed by taking linear combinations of the fewest possible common factors, while minimizing the information lost with this data reduction. The following equation is instructive:

$$Y_i = \beta_1 Z_{1i} + \ldots + \beta_n Z_{ni} + \varepsilon \qquad (6).$$

Here, Y_i represents a variable being examined by the factor analysis, Z_{ni} represents a factor generated by the analysis, β_i represents the loading on the factor (i.e., the correlation between the dependent variable Y_i and the factor Z_{ni}), and ε is the error term.

In a regression, the loadings would be interpreted as the regression coefficients on each of the explanatory variables; however, in a factor analysis, the explanatory variables are rescaled as necessary to enable the interpretation of loadings as correlation coefficients. A factor loading of 0.7 or greater is considered to be a benchmark figure for judging whether a variable and a factor are strongly correlated (StataCorp 2005).

Equation (6) appears to be similar to a regression; however, the key difference is that the factors, unlike the explanatory variables in a regression, are not directly observed. This is why it is often a challenge to produce appropriate and meaningful interpretations from a factor analysis. In addition, results can be further confounded, as factors could have significant correlation, which is not consistent with the base model assumption that there is no inter-factor correlation. Garson explains that such an analysis can help in deciphering the "latent structure" of the variables being analyzed. In other words, it provides information about the structure and make-up of a set of variables that would not otherwise be apparent – when analyzing a large quantity of data, it is useful to see how many different factors explain most of the variance within the dataset.

Each factor is associated with an eigenvalue that indicates how much of the total overall variance that particular factor explains. Factors with negative eigenvalues can be completely disregarded and in general, the greater its eigenvalue, the greater the variance explained by a particular factor (Garson). Furthermore, factors with eigenvalues very close to zero may also be ignored. In addition, it is critical to understand the concepts of *Communality* and *Uniqueness*. They are related as indicated below:

$$Uniqueness = (1 - Communality) \qquad (7).$$

Uniqueness indicates the amount of a variable's variation that is not well explained by the common factors generated by the factor analysis. Communality, as Equation (5) suggests, is the measure of the percent of a variable's variance that is well explained by the common factors. Naturally, high communality corresponds to low uniqueness, and vice versa. If a variable's uniqueness is below `0.6` (making the communality greater than `0.4`), then we may say that its variance is well explained by the common factors (StataCorp 2005).

4.3 Factor Rotations

Factor Rotations can make interpreting factor analysis results easier. There are two primary techniques for rotating factors – orthogonal and oblique. An orthogonal rotation entails rotating all the factors by a fixed angle, thereby keeping them uncorrelated, as before (Abdi 2003). The goal is to obtain factors with either really high or really low loadings. This would theoretically make it easy to identify which factors associate most closely with which variables. For an oblique rotation, the factors may be rotated by different angles within the factor space, which usually leads to correlated factors. An oblique rotation is open to meaningful interpretation only if the common factors that it produces have low to medium correlation (Garson). If the correlation between two factors were too high, then they would be "better interpreted as only one factor" (Abdi 2003). In many cases, an oblique rotation provides new insight into the data because allowing different factors to be rotated by different angles permits combinations that would not be possible under the rigidity of an orthogonal rotation.

This concludes the discussion of the technique to be utilized for examining contemporaneous relationships between the stocks' realized variances. The next section explains the technique for studying the dynamic relationships between the stocks' realized variances.

5. METHOD TWO: HAR-RV REGRESSIONS

This section is divided into 3 subsections – Section 5.1 gives a brief introduction to time series analysis; Section 5.2 explains heterogeneous autoregressive realized variance regressions, and finally, Section 5.3 provides the rationale behind significance tests to be conducted on these regressions, along with a detailed explanation of how these tests are actually conducted and their results interpreted.

5.1 An Introduction to Time Series and Lags

A time series measures the values of a particular entity at different points in time. Time series regressions are particularly helpful in examining the causal impact of one entity on another. Prior to details of these regressions, however, it is critical to explain the concept of a lag, represented as follows: Y_{t-j}. If Y_t is the value of the series Y in the present time period, then its j^{th} lag Y_{t-j} is the value of Y j periods ago. Working with a series' lags enables the testing of its utility in forecasting that series itself (autoregression), or another series, as desired.

5.2 The HAR-RV Model

Andersen, Bollerslev, Diebold and Labys (2003) show that autoregression models exhibit better results than other models when it comes to forecasting variance. We use heterogeneous autoregressive realized variance (HAR-RV) type models to forecast variance and then perform Granger Causality Tests in order to test the regressors' predictive power.

Corsi (2003) and Müller et. al (1997) developed the HAR-RV models. They are known to be simple to work with, while also capturing and relaying information about variance through lagged past values effectively. These models seek to utilize lagged averages of realized variation to forecast future realized variation. We use lags averaged over 1, 5, and 22 days in order to capture data from the preceding day, week, and month, respectively. 5 and 22 were chosen for lagging by a week and a month since those are the number of trading days in a week and a month, respectively. Specifically, the average lags are calculated as follows:

$$\overline{RV}_{t-k} = \frac{RV_{t-1} + RV_{t-2} + \ldots + RV_{t-k}}{k} \tag{8}$$

where RV_{t-k} is the k^{th} average lag of the series' realized variation. The HAR-RV model may then be expressed as:

$$RV_t = \beta_0 + \beta_D RV_{t-1} + \beta_W \overline{RV}_{t-5} + \beta_M \overline{RV}_{t-22} + \varepsilon_t \tag{9}$$

where t indicates the day and the betas are the coefficients in the regression. The four betas correspond to the constant term, the 1-day daily lag, the 5-day average weekly lag, and the 22-day average monthly lag respectively, and ε_t is the error term.

5.3 Granger Causality

We now discuss Granger Causality (Granger 1969), a very important result in time-series regression. Granger Causality is different from the standard idea of causality. If a series X has a causal impact on the series Y, then we would expect Y to be a direct consequence of X; however, if the series X Granger-causes the series Y, then the interpretation is that X contains unique information useful for predicting Y above and beyond every other series in the regression. From this point onward, all references to causality should be taken to mean Granger-causality.

The Granger Causality Test is conducted by employing an F-Test that tests the joint significance of multiple regressors in a regression. This F-Test is based upon standard errors, and it is critical to use Newey-West covariance matrix estimators to get accurate standard errors in a time-lagged regression, with a lag of 60 days to ensure heteroskedasticity robustness. This is critical, as stock volatility varies over time. Specifically, Newey-West covariance matrix estimators account for correlation amongst different lags of a times series and place more emphasis on correlated observations that appear closer together.

In order to conduct a Granger Causality Test, all the lags of a particular series in the regression must be jointly subject to an F-Test. The null hypothesis is that the coefficients on each of the lags of a particular series are zero. If the null hypothesis is rejected, then at least one of the coefficients being tested is significant and the series being tested has unique predictive content in its past values for the dependent variable. Each Granger Causality Test in this paper will examine the daily, average weekly, and average monthly lags, as described in Equation (6), of each time series that is used as an explanatory variable.

The process of testing for Granger Causality can be further illustrated as follows. Assume there are three time-series being examined: X_t, Y_t, and Z_t. Further assume that two lag levels 1 and 2 for each series are being employed, yielding the following independent variables: X_{t-1}, X_{t-2}, Y_{t-1}, Y_{t-2}, Z_{t-1}, and Z_{t-2}. The following regressions would then be run:

$$X_t = \beta_{X1}X_{t-1} + \beta_{XX}X_{t-2} + \beta_{Y1}Y_{t-1} + \beta_{Y2}Y_{t-2} + \beta_{Z1}Z_{t-1} + \beta_{Z2}Z_{t-2} + \varepsilon_{X,t}$$

$$Y_t = \beta_{X1}X_{t-1} + \beta_{X2}X_{t-2} + \beta_{Y1}Y_{t-1} + \beta_{Y2}Y_{t-2} + \beta_{Z1}Z_{t-1} + \beta_{Z2}Z_{t-2} + \varepsilon_{Y,t}$$

$$Z_t = \beta_{X1}X_{t-1} + \beta_{X2}X_{t-2} + \beta_{Y1}Y_{t-1} + \beta_{Y2}Y_{t-2} + \beta_{Z1}Z_{t-1} + \beta_{Z2}Z_{t-2} + \varepsilon_{Z,t} .$$

Next, suppose we want to test whether or not the series Y_t Granger-causes the series X_t. We would proceed by conducting an F-Test on β_{Y1} and β_{Y2} in the first regression to see whether they are jointly significant. If they are, then we conclude that the past values of the series Y_t contain unique predictive content for the series X_t, above and beyond the predictive content for X_t in X_t itself and Z_t. In other words, Y_t Granger-causes X_t. Now suppose we wanted to see if the series Z_t Granger-causes the series Y_t. If an F-Test on β_{Z1} and β_{Z2} in the second regression indicates they are not jointly significant, then we conclude that the past values of the series Z_t do not contain useful predictive content for Y_t and that Z_t does not Granger-cause Y_t. Thus, we can test time series for predictive content by jointly testing their coefficients in regressions of the type described above.

This concludes the discussion of the methods to be applied in this paper. Now we proceed to examine the results of applying factor analysis and conducting HAR-RV regressions.

6. FACTOR ANALYSIS RESULTS

The results of factor analysis enable the examination of contemporaneous relationships between the realized variances of the stocks in the technology sector. All the techniques described in Section 4 are applied and analyzed in this section, starting with a through explanation of the intuition behind the expectation that the stocks' volatility should be a function of an industry/stock-specific (idiosyncratic) effect and a market (systematic) effect.

6.1 Motivation

The expectation is that there will be two dominating common factors and that the volatility of the ten largest technology stocks would be largely explained by (1) an industry effect, and (2) a market effect. The motivation for this comes from the Market Model, which Sharpe (181) describes as follows:

$$r_{i,t} = \alpha_i + \beta_i r_{m,t} + \varepsilon_{i,t} \qquad (10),$$

where $r_{i,t}$ is the return on stock i, α_i is the intercept, β_i is the sensitivity of the stock's return to the return on an average market portfolio, $r_{m,t}$ is the return on the average market portfolio, and $\varepsilon_{i,t}$ is the error term. The Market Model indicates that the return on a stock is made up of two components – the return on the market and stock-specific idiosyncratic return, represented by the error term as the unexplained part of the total return on the stock. Thus, the question arises whether or not a stock's volatility would also be dependent on both the market and stock-specific factors, such as the industry to which the stock belongs, since we would expect returns and volatility in a particular industry to be correlated. If Equation (10) describes returns accurately, then the variance of a stock i's return would be given by

$$\sigma^2_{r_{i,t}} = \beta_i^2 \sigma^2_{m,t} + \sigma^2_{\varepsilon_{i,t}} \qquad (11);$$

in addition, the covariance of the variances of two stocks i and j would be given by

$$Cov(\sigma^2_{r_{i,t}}, \sigma^2_{r_{j,t}}) = \beta_i^2 \beta_j^2 Var(\sigma^2_{m,t}) \qquad (12),$$

assuming that $\sigma^2_{\varepsilon i,t}$ is independent of $\sigma^2_{\varepsilon j,t}$. Should they not be independent, then the above covariance would be as shown in Equation (13) below:

$$Cov(\sigma^2_{r_{i,t}}, \sigma^2_{r_{j,t}}) = \beta_i^2 \beta_j^2 Var(\sigma^2_{m,t}) + \beta_j^2 Cov(\sigma^2_{m,t}, \sigma^2_{\varepsilon_{i,t}}) + \beta_i^2 Cov(\sigma^2_{m,t}, \sigma^2_{\varepsilon_{j,t}}) + Cov(\sigma^2_{\varepsilon_{i,t}}, \sigma^2_{\varepsilon_{j,t}}) \qquad (13),$$

As has been mentioned before, Table 1 indicates that the beta of every stock being examined is non-zero. Thus, based upon Equations (11), (12), and (13) the variance and covariance of the stocks' returns should have strong influence from the general market and within the industry.

6.2 Correlation among the stocks' Realized Variances

In anticipation of the factor analysis, we also examine how well the realized variances of the each of the stocks, as well as the market, correlate with each other. These results are summarized in Table 2.[1] It is evident that all the stocks' realized variations have high correlations with each other and are in the range `0.7 - 0.8`. The market's realized variation, on the other hand, only has a medium correlation with all the stocks, around `0.5` in each case. The results of adding Google's realized variation into the dataset are in Table 3. It is clear that all the correlations reduce noticeably. The original stocks' Realized Volatilities still have medium correlation with each other; however, Google's realized variation has exceptionally low correlation with every other stock and the market.

6.3 Standard Factor Analysis

We now look at the first factor analysis, shown in Table 4. The upper part of the table lists all the discovered common factors and their eigenvalues. Factor 1 (eigenvalue = `7.72630`) is very significant. Factor 2 (eigenvalue = `0.15161`) may be interpreted as a weak, but not entirely insignificant factor. The remaining factors do not have adequate eigenvalues and will be ignored hereafter.

Conclusions from the factor analysis may now be drawn, keeping in mind the caveat that interpreting a factor analysis requires sifting through different rotations for

[1] As the variable names in Table 1 suggest, these correlations (as well as the subsequent factor analyses) are on the logs of the Daily realized variances. The rationale for taking logs prior to any analysis here has been explained in Section 3.

interpretable results, which is why its second purpose was stated to be to provide firm hypotheses to be tested using other methods. The lower part of Table 4 contains factor loadings and uniquenesses. All the stocks' variances correlate very strongly with Factor 1 and also exhibit high communality. This appears to confirm part of our hypothesis that a dominant industry effect would explain a significant portion of the stocks' volatility. That the market's volatility has relatively low correlation with this factor and high uniqueness provides further credence to this interpretation. Factor 2 has very low factor loadings throughout, even with the market. This suggests that a market effect, if any, is not discernible here and as such, the market does not appear to play a significant role in explaining technology stocks' volatility. Again, the market's high uniqueness suggests that its volatility has little in common with that of the stocks. Figure 1 (a plot of the factor loadings of the two most significant factors) shows that the market appears to be quite distinct from the stocks, all of which are clustered together.

Next, the Stata output is forcibly restricted to only two factors by giving the command that only factors with eigenvalue of `0.1` or greater be retained. Table 5 indicates that the two retained factors have identical loadings as the two factors in Table 4. However, it is important to note that the uniqueness for each of the variables has increased slightly, which indicates that removing the so-called insignificant factors does indeed reduce the measured communality.

6.4 Factor Rotations

We now examine the results of Factor Rotations in this case. Table 6 shows the correlations among the factors produced by an oblique rotation on the factor analyses in Tables 4 and 5. Practically every correlation is really weak. The only medium correlation is between the original Factors 4 and 5, which are the two most insignificant factors. This can be attributed to sampling errors and indicates that this is not a robust measure at all. We conclude that an oblique rotation will not be helpful in enhancing our previous analysis, since such a rotation admits meaningful interpretations only when the factors produced have medium to strong correlation.

An orthogonal rotation yields uncorrelated factors by construction, with these factors either having a very strong or very weak correlation with all the variables by design. At first glance, this would appear to be ideal for identifying distinct industry and market effects. The results of an orthogonal rotation on the original Factor Analysis of Table 4 are in Table 7. An orthogonal rotation would in all likelihood be helpful if there were multiple strong common factors; however, in this case this method does not provide meaningful results – the correlations (i.e., factor loadings) tend to be in the 0.3 – 0.6 range, i.e., most of the stocks have medium correlation with factors. This is probably because this technique places a high premium on the original Factors 2 and 3, which were previously deemed to be insignificant.

However, upon examining closely the new factor loadings on the factors F1 (Table 5) and F2 (Table 5) in Table 7, which result from an orthogonal rotation on the Factor Analysis of Table 5 (where Stata was forced to restrict output to only two factors), we can perhaps see a really basic Diversified/Business Technology effect. The loadings on F1 for Cisco, Dell, IBM, Microsoft, and Oracle are actually quite high and certainly seem to be in a higher range than the loadings for the other stocks. These 5 stocks are for companies that are major players in both Personal and Business Computing ("Diversified"), or are really strong in Business and Server Computing. Thus, we may conclude that there appears to be some evidence of a Diversified/Business Technology effect here, although with the caveat that EMC appears to be a notable exception.

6.5 The results of including Google (2004 – 2008)

Finally, we study the effects of adding Google into the dataset (please see Tables 8 and 9). As has been emphasized previously, data for Google are only available from 2004 onwards, since that is when the company went public. The market's communality is almost exactly the same as its previous communality (see Table 4). Some of the stocks from before also have higher uniqueness and lower communality. However, the most important result to note is that Google has extremely high uniqueness (`0.8878`) and very low correlations with all of the common factors, ranging from `0.09 – 0.20`. The common factors explain practically none of Google's realized volatility and this suggests that there is not much to pursue by way of trying to establish communality between Google's realized variation and that of other technology stocks and the market. This view is further reinforced by the fact that the factor loadings for all the stocks other than Google exhibit almost no change from Table 8 to Table 9. This suggests that including Google in the analysis has no significant impact on the preexisting correlation structure within the technology sector and that Google's variance is fundamentally different from that of technology companies. A possible explanation for Google's separation from both the technology industry and the market, presented visually in Figure 2, is that Google's revenues derive largely from advertisements placed on searches conducted through its search engine, unlike the other technology firms, all of which generate revenue primarily through providing hardware and software.

7. REALIZED VARIATION PREDICTIVE CONTENT REGRESSIONS

This section examines the results of applying the regressions described in Section 5 to examine the dynamic relationships between the realized variances of the stocks in the technology sector. Since stock price data for Google are available only from the time of its Initial Public Offering in 2004, first regressions are conducted on all the other stocks for the period 1997 – 2008. After analyzing these regressions, the results of including Google in regressions on all the stocks over the period 2004 – 2008 are examined.

7.1 Results from 1997-2008 (excluding Google)

We examine the results of regressing every stock's daily realized variation, as well as that of the market, against the lagged values of all ten stocks and the market, utilizing the HAR-RV model as described in Equation (6). We then conduct Granger Causality Tests using heteroskedasticity-robust standard errors as explained in Section 5.

Table 10 compares the explanatory power of only the stocks found to be significant in each case through the Granger Causality Tests with that of all the stocks and the market. We see that the overwhelming majority of the explanatory power (94.40% or more, in some cases nearly 100%) is contained in the series deemed significant by Granger Causality Tests.

The results of all the Granger Causality Tests are summarized in Table 11. In this table, reading across a row shows all the stocks the realized variations of which Granger-cause a particular stock's realized variation. Reading down a column provides a quick view of all the stocks that have realized variations Granger-caused by a particular stock's realized variation.

As can be seen by the significance all across the diagonal, the past volatility of every series is useful in predicting that series' own volatility. There are several other interesting observations to be made here. Hewlett-Packard, which went through a lot of turmoil in the shape of a chaotic merger with Compaq that went to a shareholder vote, as well as extensive restructuring, is the only stock for which the market has significant predictive power. Microsoft, which makes software to go with Dell and Intel's products, is very significantly explained by them. EMC Corporation and IBM are major rivals in the server business and all the stocks that are significant for the former are also significant for the latter. Most technology companies extensively use Intel-made processors and are also significantly explained by Intel. It should be noted that for the majority of the time over which the data was obtained, IBM was in both the personal computer as well as office server business, and we see that it is significantly explained by companies that relate strongly to both these lines of business. Xerox provides relatively unique products and this is borne out by the fact that its volatility only has significant predictive content for Texas Instruments, another technology company that provides services different from most of the companies in the sector. In general, it appears that the volatility of a company in a particular line of business within the technology sector tends to have significant predictive content for the volatility of other companies in that specific focus within the overall technology sector. Furthermore, the market's realized variation does not have significant predictive content for any technology stock's realized variation, other than Hewlett-Packard, which is the one stock that experienced great tumult over this period. This suggests that the market's realized variation's apparent significance in explaining HP's realized variation could be attributed to too much noise crowding out the real signal sent out by HP's realized variation.

Overall, the results of the Granger Causality Tests make tremendous intuitive sense.

7.2 Results from 2004-2008 (including Google)

The final step in the regression analysis is to examine how much predictive content Google's realized variation contains for the other stocks and the market, and vice versa. Table 12 indicates that only 36.6% of Google's realized variation is explained by the regression, which goes along with the previously established disconnect between

Google and the rest of the industry and the market. Google's realized variation contains significant predictive content only for Microsoft and Dell. It must be mentioned that this is not due to an overall decrease in explanatory power in the regression – the adjusted R^2 in most cases is comparable to the corresponding R^2 in Table 10. Table 13 indicates that neither the market, nor any technology stock (other than Google itself) contains significant predictive content for Google. Thus, we may conclude that the results from the regression corroborate those obtained through factor analysis – that its unique revenue model and limited time as a publicly held company make the nature of its volatility very different from that of other technology companies and the market.

8. CONCLUSION

This paper employs two major analysis techniques to see how the volatilities of the major technology stocks and that of the market relate to each other – both in terms of measuring their communality, as well as for seeing how much predictive content they have for each other. The market, which was found to have practically no communality with the technology stocks by conducting factor analysis, is significant only for explaining Hewlett-Packard's volatility, the one stock that experienced great tumult over this period, which suggests that there is too much noise clouding the true signal of HP's volatility. Otherwise, it would not make much sense for the market's volatility to have predictive content for only one technology stock's volatility.

Ultimately, we see that utilizing the information contained in high frequency data provides useful insight into the nature of equity volatility within the technology sector. Employing techniques that take advantage of this information, we determine that some long held beliefs on the nature of equity volatility are quite robust, while others require examination in greater depth – it was established that, as suggested by the Market Model, a stock's volatility has a major idiosyncratic component. On the other hand, it appears that the Market Model's implication of a systematic component to equity volatility can be called into question – at best, our results are indicative of a rather weak and insignificant systematic component. The industry effect completely overshadows any impact the market might have on technology stock volatility. To the author's knowledge, these results are unprecedented. The significance of these results is that analysts focusing on the technology sector can focus primarily on trends within the industry and the specific stock they are examining, without being too concerned with external events in the market that have little apparent influence on the technology sector.

It would not be prudent, however, to dismiss the market's role in influencing equity volatility. As can be seen from the current financial situation, the technology sector is relatively insulated from the general market, compared to other sectors, such as the banking sector. Examining this in more depth by applying one or both of the techniques utilized in this paper, as well as other methods, to determine the extent of the interdependency of the volatility of stocks within other industries, as well as the nature of their association with the market's volatility, would constitute interesting and useful future research.

9. TABLES

Table 1: Betas of stocks analyzed (taken from Yahoo! Finance)

Stock	CSCO	DELL	EMC	HPQ	IBM	INTC	MSFT	ORCL	TXN	XRX	GOOG
Beta	1.37	1.43	1.29	0.99	1.05	1.21	1.07	1.10	1.03	1.27	1.64

Table 2: Correlations between the natural logs of each stock's realized variation 1997-2008 (excluding Google)

ln	**CSCO**	**DELL**	**EMC**	**HPQ**	**IBM**	**INTC**	**MSFT**	**ORCL**	**TXN**	**XRX**	**SP**
CSCO	1.0000										
DELL	0.7365	1.0000									
EMC	0.7814	0.6563	1.0000								
HPQ	0.7677	0.6778	0.7366	1.0000							
IBM	0.7665	0.7293	0.7057	0.7553	1.0000						
INTC	0.7555	0.7320	0.7044	0.6962	0.6827	1.0000					
MSFT	0.8208	0.7474	0.7271	0.7511	0.8030	0.7334	1.0000				
ORCL	0.8295	0.7527	0.7527	0.7580	0.7707	0.7443	0.8121	1.0000			
TXN	0.7984	0.6751	0.7718	0.7713	0.7336	0.7335	0.7484	0.7918	1.0000		
XRX	0.7046	0.6384	0.6871	0.7127	0.6831	0.7308	0.6836	0.6914	0.7114	1.0000	
SP	0.5068	0.4750	0.5505	0.5245	0.4958	0.4995	0.4956	0.5416	0.5298	0.4930	1.0000

Table 3: Correlations between the natural logs of each stock's realized variation 2004-2008 (including Google)

ln	**CSCO**	**DELL**	**EMC**	**HPQ**	**IBM**	**INTC**	**MSFT**	**ORCL**	**TXN**	**XRX**	**GOOG**	**SP**
CSCO	1.0000											
DELL	0.3607	1.0000										
EMC	0.5599	0.3505	1.0000									
HPQ	0.4893	0.2661	0.4616	1.0000								
IBM	0.5405	0.3910	0.5157	0.4967	1.0000							
INTC	0.3655	0.3706	0.3986	0.2638	0.2781	1.0000						
MSFT	0.5899	0.4208	0.4703	0.4377	0.6232	0.2687	1.0000					
ORCL	0.5648	0.3352	0.5174	0.4321	0.4408	0.3795	0.4562	1.0000				
TXN	0.5755	0.2069	0.4473	0.4489	0.4079	0.3459	0.3579	0.4622	1.0000			
XRX	0.2674	0.1802	0.3354	0.3286	0.3731	0.3223	0.2377	0.2951	0.2684	1.0000		
GOOG	**0.1151**	**0.1524**	**0.2188**	**0.0990**	**0.1823**	**0.1426**	**0.0914**	**0.1355**	**0.0842**	**0.1575**	**1.0000**	
SP	0.4169	0.0304	0.3153	0.3616	0.2396	0.1893	0.2297	0.3169	0.4774	0.2362	0.0551	1.0000

Table 4: Factor Analysis on RVs 1997-2008 (excluding Google)

	Eigenvalue	Difference	Proportion	Cumulative
Factor 1	7.7263	7.57469	1.0076	1.0076
Factor 2	0.15161	0.06516	0.0198	1.0274

Variable (ln)	F1	F2	Uniqueness	Communality
CSCO	0.9026	-0.0467	0.1685	0.8315
DELL	0.8173	-0.1607	0.292	0.708
EMC	0.8467	0.145	0.2563	0.7437
HPQ	0.8573	0.0712	0.2471	0.7529
IBM	0.8570	-0.1231	0.2344	0.7656
INTC	0.8404	0.0224	0.2539	0.7461
MSFT	0.8843	-0.1698	0.1869	0.8131
ORCL	0.8974	-0.0659	0.1818	0.8182
TXN	0.8735	0.1224	0.2164	0.7836
XRX	0.8035	0.1191	0.3204	0.6796
SP	0.5959	0.1397	0.6159	0.3841

Only Factors 1 and 2 are shown as Factors 3-5 had trivial eigenvalues, indicating that they are insignificant

Table 5: Same as Table 4, with new Uniqueness and Communality numbers, by employing mineigen(0.1), which forcibly restricts output to two factors, by enforcing a minimum eigenvalue of 0.1, instead of the default 0

Variable (ln)	F1	F2	Uniqueness	Communality
CSCO	0.9026	-0.0467	0.1832	0.8315
DELL	0.8173	-0.1607	0.3062	0.7080
EMC	0.8467	0.145	0.262	0.7437
HPQ	0.8573	0.0712	0.2599	0.7529
IBM	0.8570	-0.1231	0.2504	0.7656
INTC	0.8404	0.0224	0.2932	0.7461
MSFT	0.8843	-0.1698	0.1892	0.8131
ORCL	0.8974	-0.0659	0.1903	0.8182
TXN	0.8735	0.1224	0.222	0.7836
XRX	0.8035	0.1191	0.3402	0.6796
SP	0.5959	0.1397	0.6255	0.3841

Table 6: Correlations among factors after Oblique rotations on Tables 4 and 5

Factor	F1(4)	F2(4)	F3(4)	F4(4)	F5(4)		F1(5)	F2(5)
F1(4)	1					F1(5)	1	
F2(4)	.1478	1				F2(5)	.01936	1
F3(4)	0.1210	.04684	1					
F4(4)	.09771	0.2418	0.1628	1				
F5(4)	.04321	.03125	-0.1314	-0.4188	1			

F1(4)-F5(4) are the 5 factors and their correlations after an oblique rotation on Table 4. F1(5) and F2(5) are the 2 factors and their correlations after an oblique rotation on Table 5.

Table 7: Factor Loadings after Orthogonal rotation

Variable (ln)	F1 (Table 4)	F2 (Table 4)	F1 (Table 5)	F2 (Table 5)
CSCO	0.6351	0.5473	0.7568	0.4940
DELL	0.6159	0.3768	0.7550	0.3517
EMC	0.4651	0.6574	0.5987	0.6160
HPQ	0.5284	0.6106	0.6507	0.5627
IBM	0.6653	0.4774	0.7650	0.4055
INTC	0.4782	0.4859	0.6658	0.5133
MSFT	0.7089	0.4522	0.8145	0.3838
ORCL	0.6423	0.5397	0.7639	0.4755
TXN	0.4983	0.6512	0.6337	0.6136
XRX	0.4092	0.5519	0.5790	0.5697
SP	0.2910	0.4894	0.3991	0.4639

The two most significant factors after performing Orthogonal Rotations on the Factors in Tables 4 and 5 (Uniqueness and Communality are not shown as they do not change)

Table 8: Factor Analysis on RVs, 2004-2008, excluding Google

	Eigenvalue	Difference	Proportion	Cumulative
Factor 1	4.31447	3.80405	0.9678	0.9678
Factor 2	0.51042	0.24101	0.1145	1.0823
Factor 3	0.26941	0.11068	0.0604	1.1427

Variable (ln)	F1	F2	F3	Uniqueness	Communality
CSCO	0.7873	0.0665	-0.1031	0.3369	0.6631
DELL	0.4774	-0.3552	0.1226	0.6235	0.3765
EMC	0.7108	-0.0229	0.0777	0.4881	0.5119
HPQ	0.6451	0.0874	-0.0813	0.5517	0.4483
IBM	0.7152	-0.1991	-0.1613	0.3975	0.6025
INTC	0.5019	-0.0597	0.3444	0.6244	0.3756
MSFT	0.6886	-0.2494	-0.2224	0.413	0.587
ORCL	0.6825	0.0229	0.0763	0.5193	0.4807
TXN	0.6540	0.3067	0.0120	0.4731	0.5269
XRX	0.4468	0.0306	0.1706	0.7070	0.2930
SP	0.4633	0.4132	-0.0445	0.6124	0.3876

Table 9: Factor Analysis on RVs, 2004-2008, including Google

	Eigenvalue	Difference	Proportion	Cumulative
Factor 1	4.35962	3.83497	0.9575	0.9575
Factor 2	0.52465	0.21712	0.1152	1.0727
Factor 3	0.30754	0.13033	0.0675	1.1403

Variable (ln)	F1	F2	F3	Uniqueness	Communality
CSCO	0.7846	-0.0888	-0.1337	0.3373	0.6627
DELL	0.4804	0.3584	0.0635	0.6201	0.3799
EMC	0.7155	0.0368	0.0986	0.4719	0.5281
HPQ	0.6433	-0.0978	-0.0645	0.5530	0.4470
IBM	0.7173	0.1913	-0.1443	0.3933	0.6067
INTC	0.5040	0.0798	0.3065	0.6281	0.3719
MSFT	0.6863	0.2165	-0.2690	0.4094	0.5906
ORCL	0.6820	-0.0261	0.0507	0.5193	0.4807
TXN	0.6512	-0.3145	0.0148	0.4717	0.5283
XRX	0.4503	-0.0060	0.2156	0.7041	0.2959
GOOG	0.2082	0.1214	0.1879	0.8878	0.1122
SP	0.4609	-0.4162	-0.0030	0.6131	0.3869

Table 10: Explanatory power of lagged values of stocks, as determined by Granger Causality Tests

DEPENDENT VARIABLE	**ALL EXPLANATORY VARIABLES AND ADJUSTED R^2**		**SIGNIFICANT STOCKS ONLY AND ADJUSTED R^2**		**% R^2 RETAINED**
CSCO	11 x 3	0.6863	4 x 3	0.6773	98.69%
DELL	11 x 3	0.7159	4 x 3	0.6758	94.40%
EMC	11 x 3	0.6464	4 x 3	0.6344	98.14%
HPQ	11 x 3	0.5433	5 x 3	0.5334	98.18%
IBM	11 x 3	0.5900	7 x 3	0.5875	99.58%
INTC	11 x 3	0.6451	4 x 3	0.6351	98.45%
MSFT	11 x 3	0.6555	4 x 3	0.6489	98.99%
ORCL	11 x 3	0.6680	4 x 3	0.6605	98.88%
TXN	11 x 3	0.6721	3 x 3	0.6432	95.70%
XRX	11 x 3	0.5725	3 x 3	0.5623	98.22%
SP[2]	11 x 3	0.3530	3 x 3	0.3526	99.89%

A comparison of how well each stock's daily realized variance is predicted by all 3 lag levels of all 10 stocks and the market, and how well it is predicted by only the stocks deemed significant by Granger causality tests (see Table 10 for which stocks are significant in which case). The calculation depicted in the second and fourth columns is designed to indicate the number of regressors in that particular regression. For example, the entry for Cisco is 4 x 3. This indicates that 4 of the 11 stocks were found to be significant and since there are 3 lag levels for each stock, there are 4 x 3 = 12 regressors in that particular regression.

[2] SP = S&P 500 (all others are standard ticker symbols, in order from top to bottom: Cisco, Dell, EMC Corporation, Hewlett Packard, IBM, Intel, Microsoft, Oracle, Texas Instruments, and Xerox).

Table 11: Results of Granger Causality Tests on HAR-RV regressions

	CSCO	DELL	EMC	HPQ	IBM	INTC	MSFT	ORCL	TXN	XRX	SP
CSCO	0.0016 **	0 ***	0.2085	0.7840	0.0722	0 ***	0.749	0.0050 **	0.3356	0.2286	0.5129
DELL	0.0014 **	0 ***	0.2775	0.1909	0.9186	0 ***	0.0004 ***	0.3170	0.1678	0.9767	0.3382
EMC	0.238	0.0004 ***	0 ***	0.2911	0.1106	0.2917	0.2338	0 ***	0.0567	0.0313 *	0.3196
HPQ	0.0945	0.0007 ***	0.0778	0.0280 *	0.3548	0 ***	0.5503	0.4819	0.0838	0.0254 *	0.0107 *
IBM	0.4403	0.0024 **	0.0056 *	0.8393	0 ***	0.0004 ***	0.0404 *	0.0050 **	0.5228	0.0078 **	0.1266
INTC	0.0017 **	0.1238	0.4709	0.0006 ***	0.2013	0 ***	0.1599	0.1128	0.2422	0.0373 *	0.7811
MSFT	0.0066 **	0 ***	0.9695	0.2042	0.3800	0 ***	0 ***	0.0842	0.8907	0.1190	0.1152
ORCL	0.4150	0.0023 **	0.1560	0.2365	0 ***	0.0009 ***	0.5620	0 ***	0.8430	0.0876	0.0863
TXN	0.1783	0.0523	0.2712	0.2608	0.6585	0 ***	0.2000	0.8545	0 ***	0.0013 **	0.0996
XRX	0.2322	0.1596	0.589	0.3381	0.8444	0.0130 *	0.1165	0.0039 **	0.3278	0 ***	0.5006
SP	0.6582	0.8839	0.2999	0 ***	0.0392 *	0.5802	0.8537	0.3708	0.8436	0.8767	0 ***
* p < 0.05, ** p < 0.01, *** p < 0.001											

The results of regressing the Daily Realized Variance of each of the stocks and the market against all 3 lag levels of all 10 stocks and the market, from 1997-2008. The first column indicates the stock whose daily RV is the regressand for that particular regression. For example, the third row indicates that when EMC's daily RV is the regressand, the RVs for Dell (DELL), EMC (EMC) itself, Oracle (ORCL), and Xerox (XRX) pass their respective Granger Causality Tests, while the other stocks and the market (denoted by SP) all fail this test. The p-values shown are obtained by conducting a Granger causality test on all 3 lag levels of the stock for that particular regression. Since 3 lag levels were tested for each stock, the degrees of freedom also equal 3 for each test.

Table 12: The predictive content contained in past lagged values of Google's RV for all technology stocks and the market

Dependent Variable	Adjusted R^2 (all 12 x 3 = 36 regressors)	Granger Test on Google
CSCO	0.6139	0.552
DELL	0.6175	0.0009***
EMC	0.6344	0.064
HPQ	0.5331	0.1779
IBM	0.6693	0.0001
INTC	0.3616	0.1133
MSFT	0.6615	0.0362*
ORCL	0.4673	0.0584
TXN	0.5209	0.8825
XRX	0.1230	0.3246
SP	0.1285	0.2702
GOOG	0.3660	0***
* p < 0.05, ** p < 0.01, *** p < 0.001		

Table 13: The predictive content for Google's RV contained in past lagged values of all other technology stocks and the market

Adjusted R^2	0.3660
Explanatory Variable	**Result of Granger Test**
CSCO	0.0657
DELL	0.4591
EMC	0.5408
HPQ	0.4681
IBM	0.7695
INTC	0.0809
MSFT	0.2422
ORCL	0.1908
TXN	0.6566
XRX	0.0573
SP	0.8506
GOOG	0***
* p < 0.05, ** p < 0.01, *** p < 0.001	

10. FIGURES

Figure 1: Factor Loadings[3] on the two factors with the highest eigenvalues in the Factor Analysis of the natural logs of Daily Realized Variations (corresponds to Table 4)

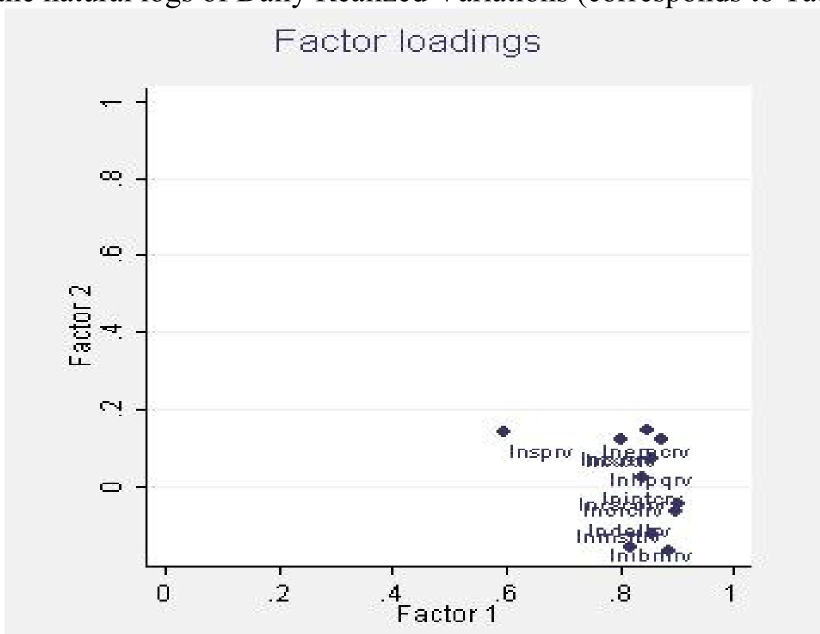

Figure 2: Factor Loadings corresponding to Table 9. Note how distinct Google's volatility is from the market as well as all the other technology stocks.

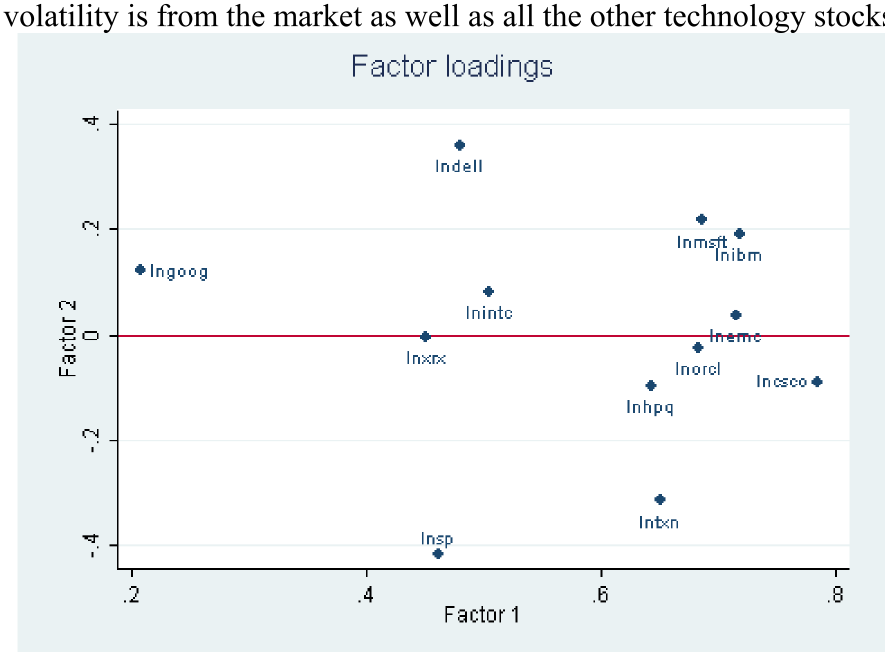

[3] Recall that a factor loading is the correlation between the factor and the variable in question.

11. REFERENCES

Abdi, Herve. "Factor Rotations in Factor Analyses." 2003. UT Dallas. 15 Oct. 2008 <http://www.utdallas.edu/~herve/abdi-rotations-pretty.pdf>.

Andersen, T.G. and T. Bollerslev (1998). "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885-905.

Andersen, T. G., Bollerslev T., Diebold F. X., and P. Labys (2003). "Modeling and forecasting realized volatility." *Econometrica* 71, 579-625.

Bollerslev, T. (2006). "Financial Market Volatility: From ARCH and GARCH to High-Frequency Data and Realized Volatility." Zeuthen Lectures, University of Copenhagen.

Corsi, F. (2003). "A simple long memory model of realized volatility." Unpublished manuscript, University of Southern Switzerland.

Garson, David. Quantitative Research in Public Administration. NCSU. 29 Apr. 2008 <http://faculty.chass.ncsu.edu/garson/pa765/index.htm>.

Gorsuch, R. L. 1974. Factor Analysis. Philadelphia: W. B. Saunders Company.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37, 424-438.

Hull, John C. Risk Management and Financial Institutions. Upper Saddle River: Prentice Hall, 2006.

Kim, J. O. and C. W. Mueller. 1978a. *Introduction to factor analysis. What it is and how to do it.* Sage University Paper Series on Quantitative Applications in the Social Sciences, 07-013. Thousand Oaks, CA: Sage.

Levy, H. and Post, T. Investments. Essex, England: Pearson Education Limited, 2005. 498-501.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., and von Weizsäcker, J.E. (1997), "Volatilities of Different Time Resolutions - Analyzing the Dynamics of Market Components." *Journal of Empirical Finance*, 4, 213-239.

Multivariate Statistics Reference Manual. College Station: StataCorp LP, 2005.

Sharpe, William C., Gordon J. Alexander, and Jeffrey W. Bailey. Investments. Upper Saddle River: Prentice Hall, 1999.

Spearman, C. 1904. General intelligence objectively determined and measured. *American Journal of Psychology* 15: 201-293.

Stock, James H., and Mark W. Watson. Introduction to Econometrics. New York: Addison-Wesley, 2006.