Stochastic Volatility Modeling of the NASDAQ 100 Futures Index

AN INVESTIGATION INTO POSSIBLE STRUCTURAL CHANGES IN NASDAQ VOLATILITY

Motivation: Everyone wants equity liquidity and fewer and fewer agents want to provide it

"Recently, however, the NASDAQ 100's volatility has been increasing as it swats away 10% drops like so many mosquitoes. This behavior is far more consistent with that of physical commodities like soybeans or natural gas, where commercial buyers seek protection from higher prices. Which buyers are panicking for protection? While hard data are not available for NASDAQ 100 futures, we can surmise that the likely suspects are hedge funds that have shorted the stratospheric index, mutual funds that need to buy in order to maintain their competitive rank, and market makers who must take opposite sides of long option trades."

> Howard Simon, "The Energizer NASDAQ", TheStreet.com, February 9, 2000

"Volatility is one of the few things people will agree on in today's market. Though some will say tech stocks will soon climb anew, others think the baton is being passed to the Old Economy stocks for a spell and still others think the whole shebang is heading lower, more chop is a constant in everybody's outlook. 'For the next few months I think we'll see volatility,' said Robert Dickey, managing director of technical analysis at Dain Rauscher Wessels in Minneapolis. 'It's a trader's delight and a real frustration for investors.'"

> Justin Lahart, "Traders Sort Through an Ordinary, Hugely Volatile Session", TheStreet.com, April 10, 2000

It seems like everyone is interested in trading stocks these days. In early April, when the NASDAQ had one of its largest single-day point drops, it was remarkable to me how many people at school were involved in trading (and judging by their faces, margin trading). And most of these people are price-takers. Why is everyone so interested in the equity markets? The spectacular bull run of the 1990s, including most notably the fantastic performance of the NASDAQ in 1999 in which the index yielded in excess of 70%, much of it accumulated in the last quarter, helps explain this phenomenon. Add to that demographic pressure from a maturing baby boom generation, low yields on government fixed income securities and a booming economy and there is a tremendous demand for liquidity.

On the supply side, the traditional providers of liquidity are pulling out of the market-making business (or, at least, they are scaling back their presence). The first critical factor in the withdrawal of liquidity in the capital markets is the failure of large-scale liquidity providers such as Long Term Capital Management. This hedge fund, hedge funds that emulated its style, and investment and commercial bank proprietary trading desks that piggybacked on its trades all were traditional sellers of implied volatility. With the losses of 1998, these sellers of liquidity found it difficult to raise or maintain their capital. Bank management instructed proprietary trading desks to curtail, or, in some cases, eliminate their risk profile. Investment banks merged with commercial banks, as took place in the alliance between Salomon Brothers and Citibank. Other investment banks listed their shares in the public markets and adopted a more fee-reliant business plan, one that de-emphasized the risky profit stream of proprietary trading. This reduction in equity derivatives liquidity was compounded by extraordinary losses in equity derivatives revealed at Union Bank of Switzerland during the implementation of its merger with Swiss Bank. There is also the question of the still-to-be-determined effect that the electronic exchanges (ECNs) like Island and Archipelago on market liquidity.

Question: Did NASDAQ volatility structurally change in August 1998?

This paper discusses a time series analysis of the NASDAQ 100 futures contract (quoted on the Chicago Mercantile Exchange). I picked the NASDAQ because its volatility (both implied and actual) has far exceeded the volatility of the S&P 500 and because NASDAQ volatility has behaved over the past two years more like what one would expect from a commodity than from

Chand Sooran **May 9, 2000**
MIT Sloan School of Management

an equity index. Equity index volatility typically tends to spike on downturns and cool off on rallies, perhaps because of the use of such techniques as portfolio insurance and covered call writing. This typical behavior might also be attributable to risk aversion on the part of individual investors who do not want to short-sell stock, choosing to participate only by buying stock or by closing out existing long positions.

The heavy use of margins by NASDAQ day-traders and speculators may explain some of the smoothing of volatility asymmetry here. Individual investors who buy stock on margin may be more likely to short-sell stock since buying stock and short-selling stock are symmetric in terms of their margin requirements.

Commodity volatility resembles this high-margin NASDAQ model in that players in the commodity markets tend to use futures contracts that require margin from all investors. Commodity volatility is high and does not display a persistent asymmetry related to the direction of the market. In commodity markets, there are both natural sellers (commodity producers) and natural buyers (commodity end users).

I have taken my data from Datastream: symbol NASA100. The data I am using is from the inception of the contract's tracking on January 2, 1983 to February 29, 2000.

I will perform a time series analysis on the NASDAQ 100 futures index and then I will use the residuals from the best time series analysis to model the volatility of the index.

Descriptive Statistics: The data suggests a unit-root process

The first thing that stands out from looking at the chart of the NASDAQ 100's performance is that the series seemed fairly stable until late 1994, after the Federal Reserve's tightening sequence was over and the current bull market began. The series then appeared to have another kind of performance until early 1999 after which it began to gap higher.

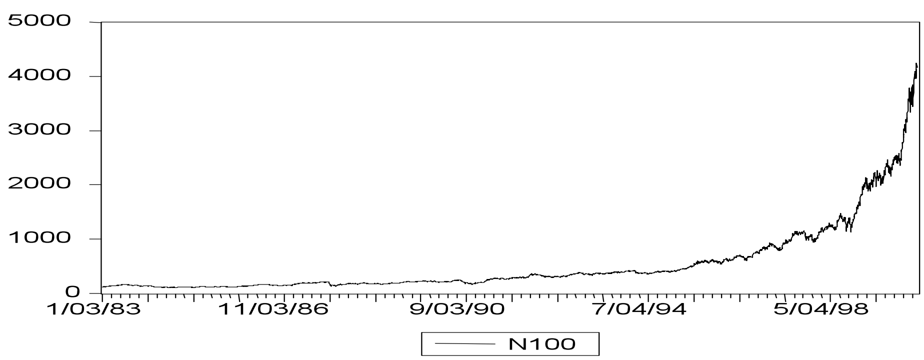

Descriptive Statistics	
Number of Observations	4477
Mean	529.86
Median	278.1723
Maximum	4256.997
Minimum	98.77614
Standard Deviation	656.8254
Skewness	2.707093
Kurtosis	11.1644

The correlogram displays a strong first-order partial autocorrelation (a coefficient of 0.996 and a first-order Q-stat of 4445.8). Notably, the serial correlation does not die down very quickly. The subsequent partial autocorrelations are all negligible. The unit-root nature of the NASDAQ is confirmed by the Augmented Dickey Fuller test statistic for the series. At 10.86631 (with 9 lags and no intercept or trend), the test statistic far exceeds the critical values (-1.9394 at a 5% confidence interval) below which we would reject the null hypothesis that the series is a unit root.

Modeling the NASDAQ 100 as an AR(1) process is consistent with current views of modeling the process for the S&P 500 Index.

Performing the ADF unit root test for the truncated period from March 2, 1999 to February 29, 2000, instead of over the whole period starting on January 1, 1983, we obtain a test statistic of 3.469989 (vs. a 5% critical value of –1.9408 below which we would reject the null hypothesis of a unit root). The correlogram for the truncated sample also shows that autocorrelation dies down much faster than the autocorrelation observed for the whole sample. The first-order Q-stat is 254.48 and the first-order partial autocorrelation is 0.982.

I calculated the ADF test statistic for the series and for the first difference of the series over five possible samples to test the stationarity of the NASDAQ 100 futures series. The results are summarized in the following table. For each sample size, we cannot reject the null hypothesis of a unit root for the series but we can reject the null hypothesis of a unit root for the first difference of the series.

Therefore, it makes sense to model the NASDAQ 100 time series as an AR(1) process.

ADF Test Statistics	**Series (5% CI)**	**First Difference of Series (5% CI)**
01/03/1983 02/29/2000	10.86631 (-1.9394)	-21.88034 (-1.9394)
01/03/1983 12/30/1994	1.593615 (-1.9394)	-17.92010 (-1.9394)
01/02/1995 07/31/1998	2.677869 (-1.9397)	-9.588391 (-1.9397)
08/31/1998 03/01/1999	1.184498 (-1.9419)	-3.445649 (-1.9419)
03/02/1999 02/29/2000	3.469989 (-1.9408)	-5.377652 (-1.9408)

NASDAQ 100 Returns: Returns are more interesting than levels

For derivatives pricing, we are most interested in the volatility of the returns to a financial price, or in this case a financial index.[1] To calculate the continuously compounded return, I computed the natural logarithm of the current period's price divided by the previous period's price. These are plotted below. Notice how the returns exhibit spiked volatility around the 1987 crash, after which they appear to settle down and how the volatility of returns picks up in a widening band post 1994.

[1] John C. Hull, Options, Futures and Other Derivative Securities, 2nd Edition, 1993, Prentice-Hall, NJ.

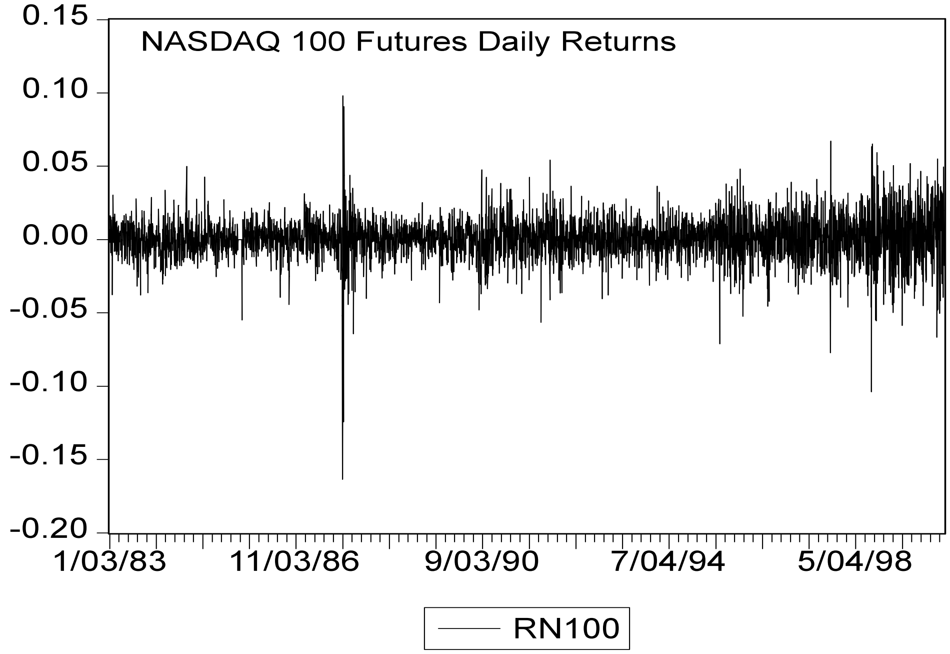

An ADF Unit Root Test on these returns over the entire sample period shows that we can reject the null hypothesis of a unit root at the 1% critical value with the test statistic computed as –21.13945 and the 1% critical value equal to –2.5662. From the correlogram of the returns series, the Q-Statistics from the first partial autocorrelation onwards have a p-value of 0.000, so that we can reject the null hypothesis that there is no autocorrelation up to order 36 (in the case of the test that I ran).

AR Modeling: AR(1) is the most appropriate model using Information Criteria to judge

Since we could not reject the null hypothesis that there was no autocorrelation in the NASDAQ 100 Futures returns, I modeled up to AR(6) (using the same trimmed sample for each) and compared the results based upon the Schwartz Information Criterion. That is, I chose the model with the lowest Schwartz Information Criterion value. These are reported in the table below.

	Schwartz Information Criterion Value
AR(1)	-5.687853
AR(2)	-5.686180
AR(3)	-5.684629
AR(4)	-5.682748
AR(5)	-5.682363
AR(6)	-5.680943

Using the Schwartz Information Criterion, the AR(1) process over the same sample is superior to AR(2) through AR(6). This is consistent with the thinking about the modeling of the returns on the Standard and Poor's Index. This is intuitively palatable since we can see from the data that the NASDAQ 100 futures index is fairly stable for most of its life.

Dependent Variable: RN100
Method: Least Squares
Date: 04/18/00 Time: 23:06
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.000750	0.000210	3.563827	0.0004
RN100(-1)	0.063874	0.014921	4.280745	0.0000

R-squared	0.004080	Mean dependent var	0.000801
Adjusted R-squared	0.003857	S.D. dependent var	0.014083
S.E. of regression	0.014056	Akaike info criterion	-5.691155
Sum squared resid	0.883676	Schwarz criterion	-5.688292
Log likelihood	12735.96	F-statistic	18.32478
Durbin-Watson stat	1.998134	Prob(F-statistic)	0.000019

Note that the AR(1) specification of the returns over the full sample has a Durbin-Watson statistic of 1.998134, very close to 2 suggesting no further autocorrelation.

However, looking at a graph of the residuals from the AR(1) estimate, this time estimated over the full sample period, less the first observation, we see a pattern of residuals that suggests heteroskedasticity, i.e. a process, the variance of which depends on time.

From the correlogram of the residuals and the Q-Statistics reported in this correlogram, we cannot reject the null hypothesis that there is no autocorrelation in the residuals. The Q-Statistic for the first-order partial autocorrelation is 0.0039. Beyond the 21st lag, the Q-Statistic has a p-value less than 0.05 (i.e. the 5% confidence interval).

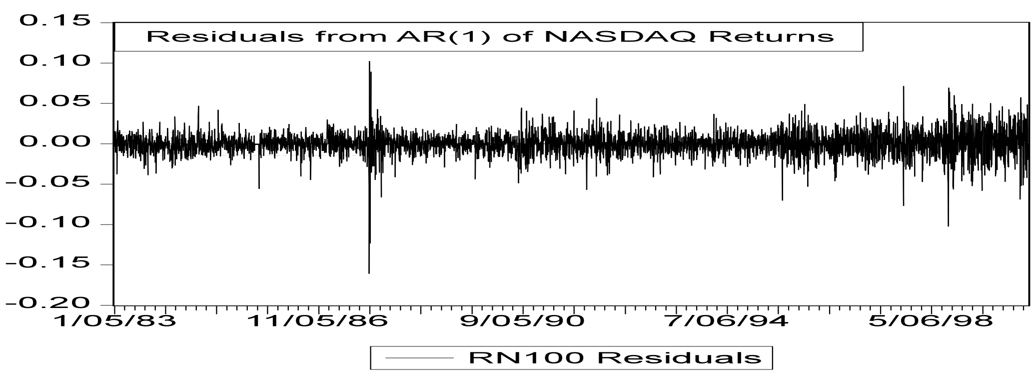

Heteroskedasticity: Various tests confirm that NASDAQ volatility is not constant

The first check that I examined was a correlogram of the squared residuals generated by the AR(1) time series specification in EViews.

ARCH or some sort of heteroskedastic pattern in the residuals is suggested by the non-zero partial autocorrelations with statistically significant Q-Statistics with lags out to roughly lag 6 or 7. The correlogram of the squared residuals from the returns AR(1) process is printed in the Appendix.

Next, I considered the histogram of the residuals, specifically looking at the Jarque-Bera statistic. This test statistic gives an assessment of the normality of the distribution of the series in question, here the residuals from the AR(1) specification of the returns of the NASDAQ.

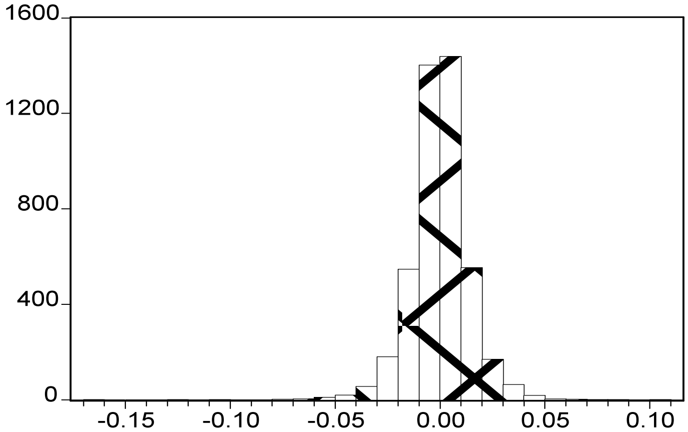

Series: Residuals	
Sample 1/05/1983 2/29/2000	
Observations 4475	
Mean	5.75E-19
Median	6.02E-05
Maximum	0.102221
Minimum	-0.160835
Std. Dev.	0.014054
Skewness	-0.585491
Kurtosis	11.58554
Jarque-Bera	13999.80
Probability	0.000000

Under the null hypothesis that the distribution is a normal one, the Jarque-Bera probability gives the probability that the Jarque-Bera statistic exceeds the observed value in the case of the null hypothesis. Here, because the probability is 0, we reject the null hypothesis of normally distributed residuals. The distribution of the residuals appears to exhibit a remarkable degree of kurtosis (with a value of 11.58554 compared to a value of 3 for the normal distribution).

The following table summarizes the results of different Breusch-Godfrey Serial Correlation LM Tests on the residuals of the AR(1) specification applied to the NASDAQ returns, with test lags up to lag 12. The null hypothesis for this test is that there is no serial correlation in the residuals up to lag order p. We cannot reject the null hypothesis up to an order of around 6, using roughly a 5-7% degree of confidence interval. The p-value for lag order 7 clearly rejects the null hypothesis of no serial correlation in the residuals up to order 7 at the 5% confidence interval.

P	B-G Serial Correlation LM Test Statistic	P-Value
1	0.933690	0.333906
2	1.434488	0.488096
3	2.442247	0.485819
4	2.498476	0.644909
5	9.575452	0.088198
6	10.86165	0.092748
7	15.80402	0.026968
8	15.81079	0.045170
9	16.46096	0.057857
10	17.05233	0.073214
11	18.57261	0.069210
12	19.88711	0.069252

Under the Engle ARCH LM test, the null hypothesis is that there is no autocorrelated conditional heteroskedasticity in the residuals up to order q. The following table summarizes the results of ARCH LM tests up to order 12.

Q	Engle ARCH LM Test Statistic	P-Value
1	284.3722	0.0000
2	506.9169	0.0000
3	659.1000	0.0000
4	662.4057	0.0000
5	796.6272	0.0000
6	796.4770	0.0000
7	811.5320	0.0000

8	824.3319	0.0000
9	825.8176	0.0000
10	835.6991	0.0000
11	836.7724	0.0000
12	842.8438	0.0000

The ARCH LM affirms our intuition that some sort of persistence characterizes the variance of the residuals. Options traders in financial markets often remark that when the price action is volatile, it tends to be volatile for a while and when trading is calm, it tends to be so for a period of time as well. Volatility in financial markets appears to trend.

ARCH(1): My first attempt to explain the serially correlated residuals is straightforward

I have chosen the benchmark for the modeling of the heteroskedastic residuals from the AR(1) specification of the NASDAQ 100 returns to be the ARCH(1) specification. This is a natural starting point and I anticipate that it will not be suitable because of the persistence in the residuals observed with the heteroskedasticity tests above. However, it provides a useful point of reference for our subsequent analysis.

The objective is to model the error process in such a way that the leftover is white noise. I can check the effectiveness of any ARCH-type specification by using the same tests that I used to identify the heteroskedasticity of the residuals from the AR(1) specification of the NASDAQ returns in the first case.

Dependent Variable: RN100
Method: ML – ARCH
Date: 04/18/00 Time: 23:21
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 18 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000789	0.000188	4.203338	0.0000
RN100(-1)	0.087570	0.013605	6.436376	0.0000
		Variance Equation		
C	0.000140	2.53E-06	55.32047	0.0000
ARCH(1)	0.283636	0.014012	20.24174	0.0000

R-squared	0.003502	Mean dependent var	0.000801
Adjusted R-squared	0.002833	S.D. dependent var	0.014083
S.E. of regression	0.014063	Akaike info criterion	-5.794778
Sum squared resid	0.884190	Schwarz criterion	-5.789052
Log likelihood	12969.82	F-statistic	5.236868
Durbin-Watson stat	2.044721	Prob(F-statistic)	0.001316

In the ARCH(1) specification, looking at the Q-Statistics from the correlogram of the residuals, I cannot reject the null hypothesis that there is no serial correlation in the residuals until approximately order 22. See the Appendix for the ARCH(1) residuals correlogram.

However, the Q-Statistics from the correlogram of the squared residuals lead me to reject the null hypothesis that the residuals are not serially correlated from the first partial autocorrelation. See the Appendix, as well.

Looking at the histogram of the residuals, the Jarque-Bera test statistic rejects the null hypothesis that the distribution of the residuals is a normal one.

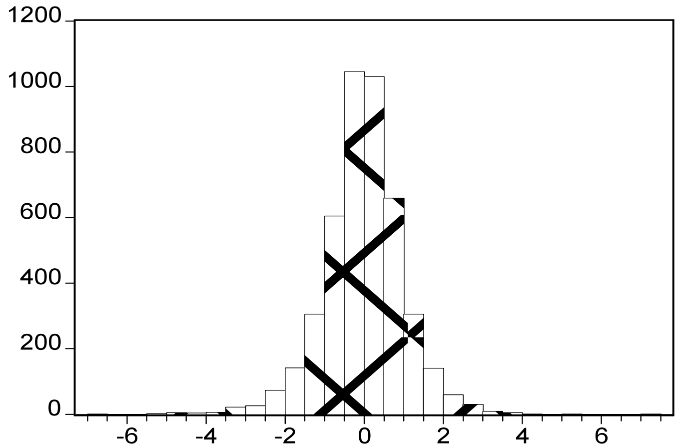

Series: Standardized Residuals	
Sample 1/05/1983 2/29/2000	
Observations 4475	
Mean	-0.009825
Median	0.000352
Maximum	7.003742
Minimum	-6.621500
Std. Dev.	1.000083
Skewness	-0.269218
Kurtosis	6.052298
Jarque-Bera	1791.200
Probability	0.000000

Finally, we can reject the null hypothesis of homoskedastic residuals using the ARCH LM Test of order 2, for which we obtain a test statistic of 58.99215 and a p-value of 0.0000.

Volatility Modeling: I will use different models of volatility to test for structural change

In the second part of this paper, I will introduce more sophisticated ARCH-based models to address the heteroskedasticity of the residuals obtained from the AR(1) specification of the returns of the NASDAQ 100 futures index.

Specifically, I intend to use the following models to test the hypothesis that NASDAQ volatility structurally changed in August 1998:

1. Autocorrelated Conditional Heteroskedasticity (ARCH)
2. Generalized Autocorrelated Conditional Heteroskedasticity (GARCH)
3. Threshold Autocorrelated Conditional Heteroskedasticity (TARCH)

In each case, I will:

1. Apply the error model to the AR(1) specification of the NASDAQ returns
2. Determine the "best" version of each error model and
3. Test for a structural change in the volatility of returns by inserting a dummy variable in the error model.

I will determine the "best" version of each error model by using likelihood ratio tests, testing down from a big structure to a smaller structure, i.e. from a more generalized specification to a more restricted one.

This dummy variable CRISIS will be equal to zero before August 1, 1998 and it will take the value 1 after August 1, 1998.

ARCH

The first specification of the error model that I examined was the Autocorrelated Conditional Heteroskedasticity specification, applied to the AR(1) model of the NASDAQ 100 returns.

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2$$

I am modeling the variance as having some persistence, in this case persistence determined by the first order lag of the squared residual from the AR(1) specification. I can extend this to order ARCH(p) of the underlying AR(1) specification by adding p lagged squared residual terms in the conditional variance specification.

I did this for ARCH(1) to ARCH(9), using the same sample, and I compared the value of each specification by using the log-likelihood function values for each specification. These are as follows:

ARCH(p)	Loglikelihood Function
1	12969.816280
2	13084.863506
3	13135.158182
4	13177.822363
5	13210.297626
6	13231.699582
7	13234.877406
8	13238.175605
9	13246.828138

I can use these Loglikelihood function values to obtain test statistics with which we can compare the efficacy of the different specifications of the variance model. These test statistics are distributed as Chi-Squared with the number of degrees of freedom equal to the number of restrictions. Comparing an ARCH(2) to a restricted ARCH(1) specification means that the test statistic would be distributed as Chi-Squared with one degree of freedom, for example.

ARCH(1) is a restricted version of ARCH(9) in which the coefficients on the lagged squared residual terms for lags 2 through 9 are restricted to be equal to zero. I used a likelihood ratio test to compare the different ARCH models. I found that I rejected every ARCH(.) specification other than the ARCH(9) specification, i.e. the least restricted one. See the Appendix for a table listing the different likelihood ratio tests I performed and a description of the serial correlation properties of both the residuals and the squared residuals in these cases. The ARCH(9) specification has no serial correlation out to the 23rd lag in the residuals and no serial correlation in the squared residuals beyond the 36th lag.

The estimation output for the ARCH(9) specification, using an AR(1) process for the NASDAQ 100 returns series is as follows:

Dependent Variable: RN100
Method: ML – ARCH
Date: 05/07/00 Time: 13:57
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 22 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000775	0.000170	4.563279	0.0000
RN100(-1)	0.111788	0.016739	6.678488	0.0000
	Variance Equation			
C	6.25E-05	2.00E-06	31.21588	0.0000
ARCH(1)	0.128716	0.009609	13.39556	0.0000
ARCH(2)	0.115777	0.014234	8.133612	0.0000
ARCH(3)	0.063849	0.013870	4.603359	0.0000
ARCH(4)	0.107363	0.014021	7.657208	0.0000
ARCH(5)	0.084439	0.016072	5.253953	0.0000
ARCH(6)	0.070065	0.013872	5.050792	0.0000
ARCH(7)	0.013883	0.012923	1.074243	0.2827
ARCH(8)	0.033181	0.013674	2.426618	0.0152
ARCH(9)	0.060400	0.013175	4.584573	0.0000

R-squared	0.001764	Mean dependent var	0.000801
Adjusted R-squared	-0.000697	S.D. dependent var	0.014083
S.E. of regression	0.014088	Akaike info criterion	-5.915007
Sum squared resid	0.885731	Schwarz criterion	-5.897828
Log likelihood	13246.83	F-statistic	0.716901
Durbin-Watson stat	2.092052	Prob(F-statistic)	0.723399

The only ARCH lag with an insignificant t-statistic is the 7th lag. The p-value for the F-Statistic for the null hypothesis that all of the coefficients are zero is 0.723399. Therefore, we cannot reject the null hypothesis that all of the coefficients are jointly equal to zero.

I then tested the possibility of a structural change in the volatility of the NASDAQ 100 return by including a dummy variable, called CRISIS, that is equal to zero before August 1, 1998 and equal to 1 after August 1, 1998 in the ARCH(9)-AR(1) specification. This dummy variable is included in the variance specification so that the variance set up looks as follows:

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1\varepsilon_{t-1}^2 + \alpha_2\varepsilon_{t-2}^2 + \alpha_3\varepsilon_{t-3}^2 + \alpha_4\varepsilon_{t-4}^2 + \alpha_5\varepsilon_{t-5}^2 + \alpha_6\varepsilon_{t-6}^2 + \alpha_7\varepsilon_{t-7}^2 + \alpha_8\varepsilon_{t-8}^2 + \alpha_9\varepsilon_{t-9}^2 + \beta CRISIS_t$$

Dependent Variable: RN100
Method: ML – ARCH
Date: 05/07/00 Time: 14:03
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 25 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000762	0.000173	4.406647	0.0000
RN100(-1)	0.110399	0.016893	6.535248	0.0000
	Variance Equation			
C	6.89E-05	2.25E-06	30.62097	0.0000
ARCH(1)	0.112082	0.008242	13.59960	0.0000
ARCH(2)	0.099804	0.014416	6.923287	0.0000
ARCH(3)	0.054902	0.013455	4.080254	0.0000
ARCH(4)	0.073653	0.013902	5.297987	0.0000
ARCH(5)	0.079447	0.016031	4.955955	0.0000
ARCH(6)	0.057731	0.013213	4.369134	0.0000

ARCH(7)	0.008124	0.011879	0.683860	0.4941
ARCH(8)	0.027610	0.013294	2.076918	0.0378
ARCH(9)	0.040971	0.012688	3.229229	0.0012
CRISIS	0.000194	3.20E-05	6.060045	0.0000

R-squared	0.001903	Mean dependent var	0.000801
Adjusted R-squared	-0.000781	S.D. dependent var	0.014083
S.E. of regression	0.014088	Akaike info criterion	-5.926869
Sum squared resid	0.885608	Schwarz criterion	-5.908259
Log likelihood	13274.37	F-statistic	0.709055
Durbin-Watson stat	2.089366	Prob(F-statistic)	0.744108

The crisis variable has a t-statistic equal to 6.060045 with a corresponding p-value of 0.0000. The F-Statistic for the joint null hypothesis that all of the slope coefficients is equal to zero is still such that we cannot reject the null hypothesis. Notice also that including the crisis variable increases the p-value for the t-statistic of the ARCH(7) coefficient, meaning that it is increasingly less likely that we can reject the null hypothesis that the coefficient α_7 is zero.

Doing a Wald coefficient test for the Crisis coefficient also yielded a p-value of zero for the null hypothesis that the Crisis coefficient was equal to zero.

Wald Test:
Equation: ARCH9AR1CRISIS

Null Hypothesis:	C(13) = 0		
F-statistic	36.72414	Probability	0.000000
Chi-square	36.72414	Probability	0.000000

I also did a likelihood ratio test in which I compared the restricted ARCH(9) specification without the CRISIS dummy variable to the ARCH(9) specification that included the CRISIS dummy variable. I obtained a test statistic of 55.0835936. The critical value under a Chi-Squared distribution with one degree of freedom is 3.8414553. Therefore, we can reject the null hypothesis that the coefficient on the CRISIS dummy variable is zero.

The coefficient for the crisis variable is very small. This means that the mean for the variance term shifts higher by a smaller, but statistically significant, amount in the post-crisis period. We cannot infer anything about the differences in symmetry before and after the crisis pivotal date from this simple ARCH specification of the variance.

The inclusion of the dummy variable seems to have little effect on the coefficients in the AR(1) specification. Plotting the difference between the conditional variance from the ARCH(9)-AR(1) specification with the crisis dummy variable in the variance specification and the conditional variance from the ARCH(9)-AR(1) specification without the crisis variable confirms this description. Before the crisis pivotal date, the crisis variance is below the non-crisis variance, by a small amount. After the crisis pivotal date, the crisis variance is above the non-crisis variance uniformly.

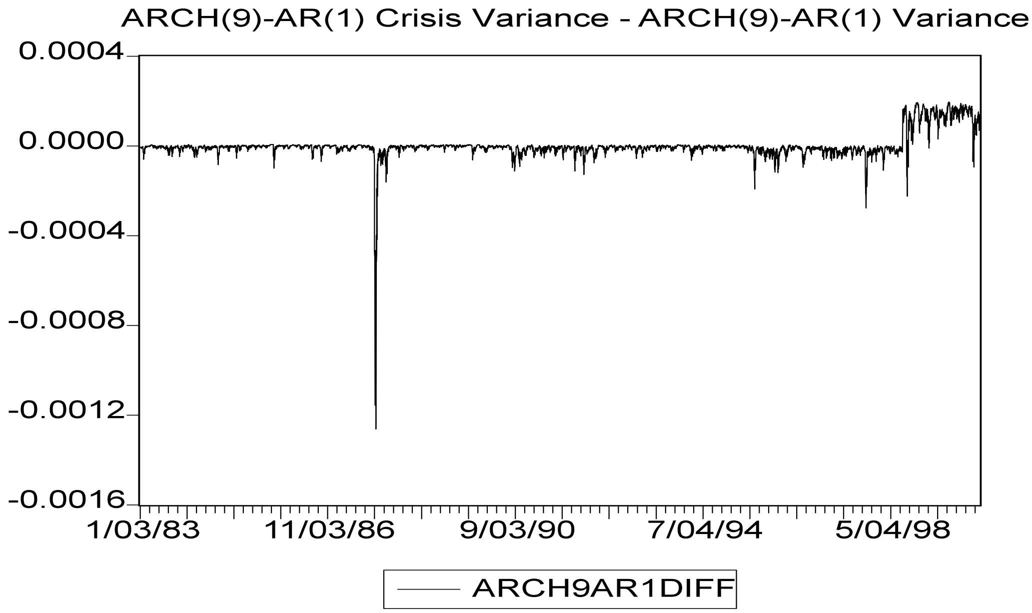

GARCH

The GARCH specification includes in the variance model lagged terms from the forecast variance, as well as lagged terms for the squared residual (as in the ARCH specification). The GARCH(1,1) specification is set up as follows:

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \delta\sigma_{t-1}^2$$

In general a GARCH(p,q) model has p lagged squared residual terms and q lagged forecast variance terms. Note that the ARCH(p) specification is a restricted form of the GARCH(p,q) specification of the variance in which the lagged forecast variance terms are ignored.

Including lagged forecast variance terms may make sense for a financial time series if market participants are engaged in the process of updating their forecasts of variance as time progresses. This may turn out to be a more important part of the NASDAQ volatility story than lagged squared residual terms from the AR(1) process, which in and of themselves have a less intuitive meaning in the context of a traded financial series.

The crisis variable can also be included here as a dummy variable in the variance specification, in the same way that we included the crisis variable as a dummy variable in the ARCH specification. For example in the GARCH(1,1) specification,

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \delta\sigma_{t-1}^2 + CRISIS_t$$

I first performed a series of GARCH(p,1) specifications and I obtained the following log-likelihood values from the estimation output.

GARCH()	Log-likelihood Function

GARCH(1,1)	13264.90399
GARCH(2,1)	13270.31474
GARCH(3,1)	13282.22611
GARCH(4,1)	13283.37438
GARCH(5,1)	13276.48047
GARCH(6,1)	13289.40504
GARCH(7,1)	13291.20439
GARCH(8,1)	13281.98015
GARCH(9,1)	13288.47226

Something odd occurs at GARCH(7,1). The Loglikelihood value peaks at GARCH(7,1). Since GARCH(7,1) is a restricted version of GARCH(9,1), the Loglikelihood value for GARCH(9,1) should be higher than the Loglikelihood value for GARCH(7,1), if only marginally so, as we observed in the ARCH likelihood ratio testing above.

Performing a likelihood-ratio test in which we compare GARCH(9,1) and GARCH(1,1), we obtain a test statistic of 47.13653 (with a 5% critical value of 15.507312). We reject the null hypothesis that the restrictions hold true.

However, looking at the estimation output for GARCH(9,1), we can see that most of the ARCH coefficients beyond the 1st lag have t-statistics that do not allow us to reject the null hypothesis that these coefficients are individually equal to zero. The single GARCH coefficient has a p-value of 0.0000 for the null hypothesis that it is zero. This makes sense given the intuition that traders adjust their forecast variances to reflect the evolution of prices, a phenomenon that corresponds to the GARCH terms.

Dependent Variable: RN100
Method: ML – ARCH
Date: 05/07/00 Time: 15:13
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 38 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000732	0.000168	4.351340	0.0000
RN100(-1)	0.118141	0.016752	7.052444	0.0000
		Variance Equation		
C	1.07E-06	2.28E-07	4.680486	0.0000
ARCH(1)	0.127668	0.007874	16.21445	0.0000
ARCH(2)	-0.014586	0.018274	-0.798210	0.4247
ARCH(3)	-0.065523	0.022075	-2.968273	0.0030
ARCH(4)	0.032266	0.024523	1.315735	0.1883
ARCH(5)	-0.026305	0.023160	-1.135780	0.2560
ARCH(6)	0.000256	0.014533	0.017616	0.9859
ARCH(7)	-0.032704	0.021802	-1.500049	0.1336
ARCH(8)	-0.011056	0.018024	-0.613433	0.5396
ARCH(9)	0.011452	0.008823	1.297900	0.1943
GARCH(1)	0.971857	0.004678	207.7708	0.0000

R-squared	0.001132	Mean dependent var	0.000801
Adjusted R-squared	-0.001554	S.D. dependent var	0.014083
S.E. of regression	0.014094	Akaike info criterion	-5.933172
Sum squared resid	0.886292	Schwarz criterion	-5.914562
Log likelihood	13288.47	F-statistic	0.421339

Durbin-Watson stat	2.104433	Prob(F-statistic)	0.955984

Furthermore, the F-Statistic for the null hypothesis that all of the coefficients are jointly equal to zero returns a p-value of 0.955984. We cannot reject this null hypothesis.

Therefore, I will consider GARCH(1,1) through GARCH(1,9).

Again, I notice that there is something strange about the 7th lag, this time on the GARCH term, when comparing Loglikelihood values.

GARCH()	Log-likelihood Function
GARCH(1,1)	13264.90399
GARCH(1,2)	13266.05432
GARCH(1,3)	13268.35859
GARCH(1,4)	13271.43544
GARCH(1,5)	13272.61838
GARCH(1,6)	13272.07602
GARCH(1,7)	13272.00905
GARCH(1,8)	13276.96115
GARCH(1,9)	13277.11525

This time, I started with the GARCH(1,9) specification and began by comparing it to the GARCH(1,8) specification, using a likelihood ratio test. I found that the GARCH(1,8) was the superior specification of the GARCH(1,q) setups.

Unrestricted Specification	Restricted Specification	Number of Restrictions	Test Statistic	5% Critical Value	5% P-Value	Reject/Not Restrictions
GARCH(1,9)	GARCH(1,1)	8	24.422509	15.50731249	0.001946131	Reject
GARCH(1,9)	GARCH(1,8)	1	0.308191	3.841455338	0.57879239	Not Reject
GARCH(1,8)	GARCH(1,7)	1	9.904211	3.841455338	0.00164901	Reject
GARCH(1,8)	GARCH(1,6)	2	9.770272	5.991476357	0.007558097	Reject
GARCH(1,8)	GARCH(1,5)	3	8.685548	7.814724703	0.033777462	Reject
GARCH(1,8)	GARCH(1,4)	4	11.051420	9.487728465	0.025992167	Reject
GARCH(1,8)	GARCH(1,3)	5	17.205127	11.07048257	0.004126757	Reject
GARCH(1,8)	GARCH(1,2)	6	21.813665	12.59157742	0.001308694	Reject
GARCH(1,8)	GARCH(1,1)	7	24.114318	14.06712726	0.001087812	Reject

See the Appendix for a comparison of the serial correlation properties of the residuals and the squared residuals for the GARCH(1,8) case.

The estimation output from the GARCH(1,8) case is:

Dependent Variable: RN100
Method: ML - ARCH
Date: 05/07/00 Time: 15:50
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 31 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000700	0.000176	3.969811	0.0001

RN100(-1)	0.115488	0.017318	6.668834	0.0000
		Variance Equation		
C	7.10E-06	6.75E-07	10.52936	0.0000
ARCH(1)	0.126142	0.007058	17.87150	0.0000
GARCH(1)	0.465515	0.092597	5.027339	0.0000
GARCH(2)	0.103991	0.135736	0.766128	0.4436
GARCH(3)	0.123381	0.147956	0.833906	0.4043
GARCH(4)	0.064826	0.145589	0.445264	0.6561
GARCH(5)	-0.006150	0.135539	-0.045376	0.9638
GARCH(6)	-0.149670	0.140379	-1.066190	0.2863
GARCH(7)	0.116124	0.142411	0.815414	0.4148
GARCH(8)	0.115570	0.104443	1.106539	0.2685

R-squared	0.001416	Mean dependent var	0.000801
Adjusted R-squared	-0.001046	S.D. dependent var	0.014083
S.E. of regression	0.014090	Akaike info criterion	-5.928474
Sum squared resid	0.886040	Schwarz criterion	-5.911295
Log likelihood	13276.96	F-statistic	0.575134
Durbin-Watson stat	2.099289	Prob(F-statistic)	0.850581

I then tested the possibility of a structural change in the volatility of the NASDAQ 100 return by including in the GARCH(1,8)-AR(1) specification a dummy variable, called CRISIS, that is equal to zero before August 1, 1998 and equal to 1 after August 1, 1998. This dummy variable is included in the variance specification so that the variance set up looks as follows:

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1\varepsilon_{t-1}^2 + \delta_1\sigma_{t-1}^2 + \delta_2\sigma_{t-2}^2 + \delta_3\sigma_{t-3}^2 + \delta_4\sigma_{t-4}^2 + \delta_5\sigma_{t-5}^2 + \delta_6\sigma_{t-6}^2 + \delta_7\sigma_{t-7}^2 + \delta_8\sigma_{t-8}^2 + \beta CRISIS_t$$

The estimation output is as follows:

Dependent Variable: RN100
Method: ML - ARCH
Date: 05/07/00 Time: 16:31
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 34 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000697	0.000177	3.940981	0.0001
RN100(-1)	0.115341	0.017368	6.641105	0.0000
		Variance Equation		
C	1.04E-05	9.61E-07	10.86414	0.0000
ARCH(1)	0.127744	0.006701	19.06390	0.0000
GARCH(1)	0.384745	0.085513	4.499257	0.0000
GARCH(2)	0.194780	0.123501	1.577148	0.1148
GARCH(3)	0.183715	0.130713	1.405488	0.1599
GARCH(4)	0.028704	0.123348	0.232704	0.8160
GARCH(5)	-0.063411	0.115786	-0.547652	0.5839
GARCH(6)	-0.201583	0.115507	-1.745201	0.0809
GARCH(7)	0.118091	0.116614	1.012673	0.3112
GARCH(8)	0.156781	0.094970	1.650848	0.0988
CRISIS	3.26E-05	7.43E-06	4.389337	0.0000

R-squared	0.001430	Mean dependent var	0.000801
Adjusted R-squared	-0.001255	S.D. dependent var	0.014083
S.E. of regression	0.014092	Akaike info criterion	-5.933443
Sum squared resid	0.886027	Schwarz criterion	-5.914832
Log likelihood	13289.08	F-statistic	0.532620
Durbin-Watson stat	2.099001	Prob(F-statistic)	0.894932

The t-statistic for the CRISIS coefficient has a p-value of zero (in large part due to the small standard error for the coefficient estimate) stating that we would reject the null hypothesis that the coefficient on CRISIS is zero.

A Wald test for this coefficient confirms this result.

Wald Test:
Equation: GARCH18AR1CRISIS

Null Hypothesis:	C(13)=0		
F-statistic	19.26628	Probability	0.000012
Chi-square	19.26628	Probability	0.000011

A likelihood ratio test comparing the GARCH(1,8)-AR(1) model without the crisis variable (the restricted specification) to the GARCH(1,8)-AR(1) model with the crisis variable (the unrestricted specification) also confirms this result. The computed test statistic is 24.2376928, which is much greater than the significant value of a Chi-Squared distribution with one degree of freedom, equal to 3.8414553, leading us to reject the null hypothesis that the restriction holds true.

Plotting the difference of the variance from a GARCH(1,8) specification that does not include the crisis variable and the GARCH(1,8) specification that does include the crisis variable, we see the similar pattern to the one observed in the ARCH case. There is a discrete jump in the mean of the variance at the CRISIS pivotal date of August 1, 1998.

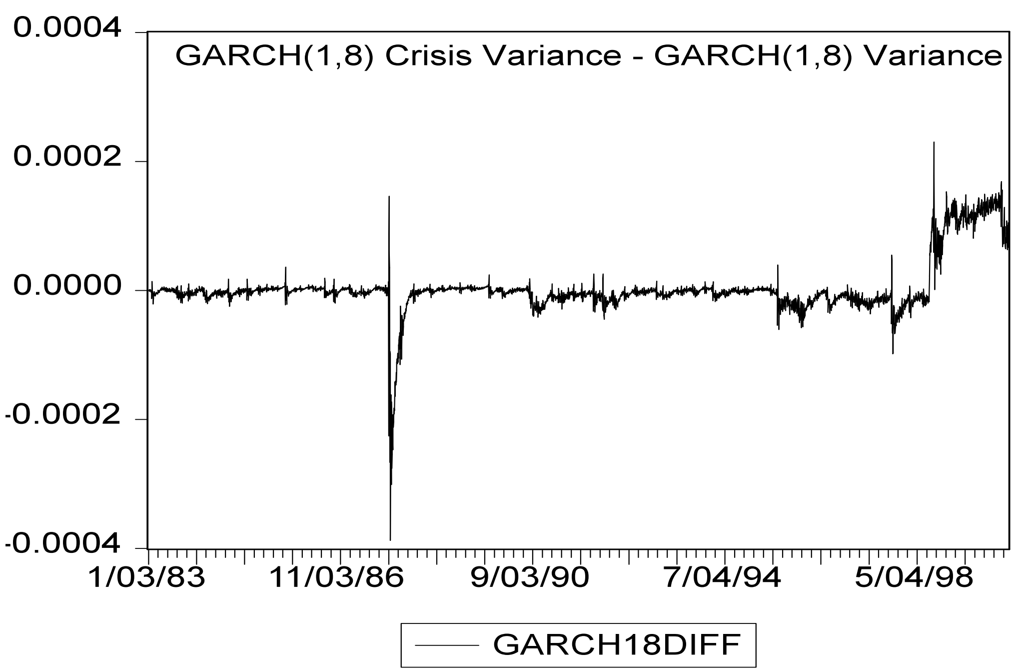

Note that the difference is exaggerated during the stock market crash of 1987 because the GARCH(1,8) specification, before August 1, 1998, without the crisis variable generally overestimates the variance compared to the GARCH(1,8) specification with the crisis variable (i.e. the difference between the two tends to be negative). This effect is exploded during the unusual volatility of the Fall of 1987.

TARCH

The Threshold Autoregressive Conditional Heteroskedasticity Model is a generalization of the standard GARCH model to include in the variance forecast an asymmetric reaction to (in this case) positive and negative lagged residuals in the NASDAQ 100 performance. TARCH(1,1) may be represented as,

$$RN100 = C + \gamma + \gamma_1 RN100_{-1} + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \delta \sigma_{t-1}^2 + \beta \varepsilon_{t-1} d_{t-1}$$

$d_{t-1} = 1$ if $\varepsilon_t < 0$ and $d_{t-1} = 0$ otherwise

Higher-order TARCH specifications would resemble higher-order GARCH setups with the inclusion of the single $\beta\varepsilon_{t-1}d_{t-1}$ term.

If β is equal to zero, then the impact of a negative residual (bad news) and the impact of a positive residual (good news) is asymmetric.

TARCH is compelling here in light of the institutional factors involved in equity trading. Recall that I stated in the description of the problem that equity volatility typically tends to be asymmetric, demonstrating higher volatility in the event of a downturn in the equity market than it does in the case in which returns are positive. This is certainly the case for the “skew” in the equity options markets in which deep out-of-the-money equity calls often trade at an implied volatility discount to the implied volatility of at-the-money options and out-of-the-money puts trade at an implied volatility premium to the at-the-money options. This may be due to investor reticence to hold losing positions, the effect of margin in clearing out losing positions early (and the use of margin predominantly to buy stocks by individual investors), institutional obstacles to short-selling and the use of strategies such as covered call writing that involve selling out-of-the-money calls.

I ran TARCH(1,1) to TARCH(1,9), using the same logic as in the GARCH modeling to restrict the ARCH term p to 1. With a series of likelihood ratio tests, I found that the most appropriate TARCH representation here of the variance from the AR(1) specification of the underlying NASDAQ 100 index was the TARCH(1,6).

Unrestricted Specification	Restricted Specification	Number of Restrictions	Test Statistic	5% Critical Value	5% P-Value	Reject/Not Restrictions
TARCH(1,9)	TARCH(1,1)	8	35.708911	15.50731249	1.98508E-05	Reject
TARCH(1,9)	TARCH(1,8)	1	0.316201	3.841455338	0.573899581	Not Reject
TARCH(1,8)	TARCH(1,7)	1	3.611317	3.841455338	0.057387679	Not Reject
TARCH(1,7)	TARCH(1,6)	1	2.657365	3.841455338	0.103071364	Not Reject
TARCH(1,6)	TARCH(1,5)	1	12.378271	3.841455338	0.000434359	Reject
TARCH(1,6)	TARCH(1,4)	2	8.852922	5.991476357	0.011956732	Reject
TARCH(1,6)	TARCH(1,3)	3	21.739599	7.814724703	7.38965E-05	Reject
TARCH(1,6)	TARCH(1,2)	4	28.437117	9.487728465	1.01703E-05	Reject
TARCH(1,6)	TARCH(1,1)	5	29.124028	11.07048257	2.1924E-05	Reject

In the TARCH(1,6) setup without the inclusion of the CRISIS variable, I obtain the following estimation output:

Dependent Variable: RN100
Method: ML - ARCH
Date: 05/07/00 Time: 17:04
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 23 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000451	0.000176	2.561299	0.0104
RN100(-1)	0.123196	0.017072	7.216442	0.0000
	Variance Equation			
C	6.97E-06	7.52E-07	9.271190	0.0000
ARCH(1)	0.054131	0.009029	5.995224	0.0000
(RESID<0)*ARCH(1)	0.110852	0.012533	8.844792	0.0000
GARCH(1)	0.765817	0.103283	7.414759	0.0000
GARCH(2)	-0.154275	0.112317	-1.373562	0.1696
GARCH(3)	0.139320	0.103040	1.352094	0.1763
GARCH(4)	0.292116	0.106922	2.732036	0.0063
GARCH(5)	-0.570967	0.117804	-4.846743	0.0000
GARCH(6)	0.379015	0.079793	4.749990	0.0000

R-squared	0.000242	Mean dependent var	0.000801
Adjusted R-squared	-0.001997	S.D. dependent var	0.014083
S.E. of regression	0.014097	Akaike info criterion	-5.938032
Sum squared resid	0.887082	Schwarz criterion	-5.922285
Log likelihood	13297.35	F-statistic	0.108138
Durbin-Watson stat	2.113564	Prob(F-statistic)	0.999753

Including the CRISIS variable in the TARCH(1,6) setup, we obtain,

Dependent Variable: RN100
Method: ML - ARCH
Date: 05/07/00 Time: 17:16
Sample(adjusted): 1/05/1983 2/29/2000
Included observations: 4475 after adjusting endpoints
Convergence achieved after 81 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000463	0.000177	2.622442	0.0087
RN100(-1)	0.121199	0.016944	7.153141	0.0000
	Variance Equation			
C	1.05E-05	1.03E-06	10.19026	0.0000
ARCH(1)	0.030622	0.008592	3.563793	0.0004
(RESID<0)*ARCH(1)	0.147968	0.013271	11.14969	0.0000
GARCH(1)	0.843259	0.112962	7.464976	0.0000
GARCH(2)	-0.356090	0.174965	-2.035207	0.0418
GARCH(3)	0.150670	0.183652	0.820407	0.4120
GARCH(4)	0.211940	0.181610	1.167006	0.2432
GARCH(5)	-0.368510	0.141733	-2.600023	0.0093
GARCH(6)	0.340857	0.064351	5.296846	0.0000
CRISIS	3.29E-05	5.95E-06	5.521486	0.0000

R-squared	0.000501	Mean dependent var	0.000801

Adjusted R-squared	-0.001962	S.D. dependent var	0.014083
S.E. of regression	0.014097	Akaike info criterion	-5.948503
Sum squared resid	0.886852	Schwarz criterion	-5.931324
Log likelihood	13321.77	F-statistic	0.203529
Durbin-Watson stat	2.109754	Prob(F-statistic)	0.997457

The t-statistic for the CRISIS coefficient has a p-value of zero, leading us to reject the null hypothesis that the CRISIS coefficient is zero.

Similarly, the Wald test rejects the null hypothesis that the CRISIS coefficient is zero.

Wald Test:
Equation: TARCH16CRISIS

Null Hypothesis:	C(12) = 0		
F-statistic	30.48681	Probability	0.000000
Chi-square	30.48681	Probability	0.000000

A likelihood ratio test comparing the general TARCH(1,6) setup (i.e. including the CRISIS variable) to a restricted TARCH(1,6) specification (without the CRISIS variable) yields a test statistic of 48.8539848 which is much greater than the critical value for a Chi-Squared Distribution with one degree of freedom of 3.8414553. We reject the null hypothesis that the coefficient on the CRISIS variable is zero.

Plotting the variance obtained from the TARCH(1,6) specification including the CRISIS variable, we see that (with the exception of the Fall of 1987), variance appears to be quite stable until the crisis of August 1998.

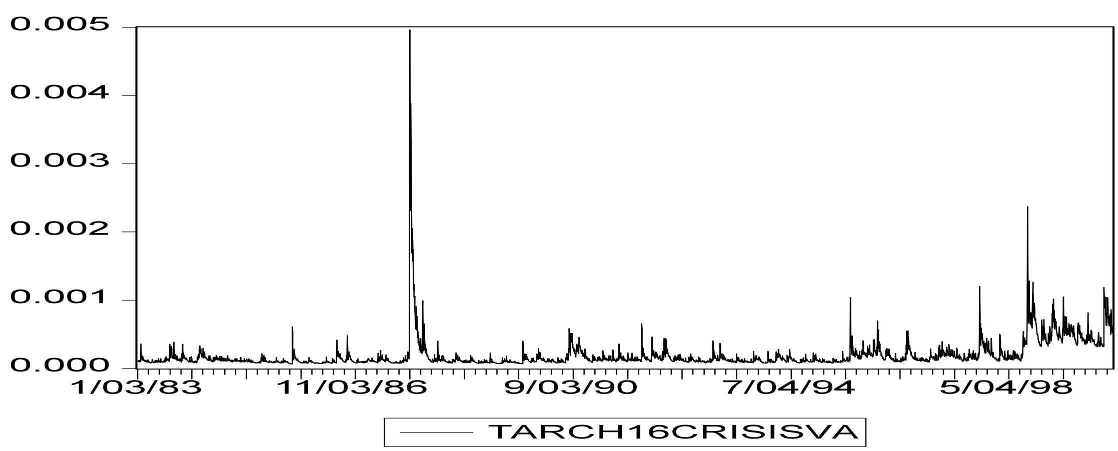

Finally, looking at the difference between the TARCH(1,6) variance with and without the inclusion of the CRISIS variable,

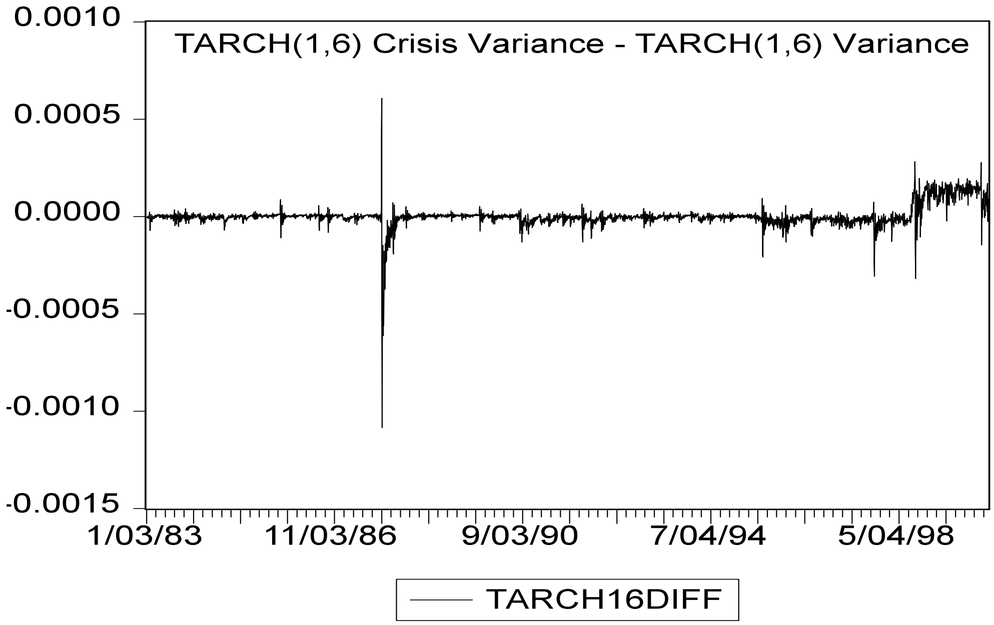

We observe the discrete change in volatility in August 1998 attributable to the Russian default crisis.

Conclusion: NASDAQ volatility changed structurally with the Crisis of August 1998

The first person to suggest to me that NASDAQ volatility (and equity volatility, generally) had undergone a structural change with the crisis of the Fall of 1998 and the ensuing developments in the equity markets was the head of equity derivatives for Lehman Brothers, the bank for which I will be working this summer. It was critically important for him (and everyone working for him) to understand this because, if true, this conclusion has far-reaching implications.

I have shown using a variety of specifications of NASDAQ volatility, including Autogregressive Conditional Heteroskedasticity, Generalized Autogregressive Conditional Heteroskedasticity and Threshold Autoregressive Conditional Heteroskedasticity, that NASDAQ volatility changed structurally in August 1998.

Risk management is predicated on forecasts of variance and correlation, developed from historical data applying these kinds of models. The conclusion of this paper is that the assumptions underlying the models embedded in current risk management technologies applied to the day-to-day operations of most dealer trading desks, hedge fund trading desks and corporate treasuries needs to be re-evaluated in order to ensure that it is still appropriate.

Traders who make markets ("market makers"), while already cognizant of the growing difficulty in obtaining liquidity, can quantify the additional risk of being involved in market-making roles and can justify potentially larger bid-offer spreads in derivative products. Often, market-makers inherit positions, as a consequence of the services they provide. This analysis suggests that carrying such residual positions is riskier than it might appear, based upon long-term historical data. This study suggests a different strategy, a more nimble one to be sure, is appropriate for institutions that seek to maintain a presence as market-makers.

Corporate treasures and investors must be prepared to shorten their funding and investment horizons and to be opportunistic when developing strategies for their involvement in the equity markets. Enhanced volatility is making day-to-day trading more important, in order to preserve capital as well as to take advantage of market opportunities.