Article from:

Risks and Rewards Newsletter

October 1998 – Issue No. 31

October's Market Demons: The '87 Stock Market Crash and Likelihood of a Recurrence

by Vinod Chandrashekaran

Editor's Note: *The following article originally appeared in the Winter 1998 issue of The BARRA Newsletter, Horizon, and is reprinted here with permission.*

Monday, October 27, 1997, witnessed a drop of 554 points in the Dow Jones Industrial Average, the largest one-day point drop in the history of the market. Dramatic as it was, this drop was only the twelfth-largest fall in percentage terms. The largest percentage drop in the history of the stock market occurred on Monday, October 19, 1987, when the S&P 500 Index declined by 20.5%. The recent sharp movements witnessed in global markets raise an important question: What is the likelihood of market crashes? This article seeks to provide an answer to this and related questions by focusing mainly on the crash of October 1987. Specifically, we shall seek answers to three questions:

1. Given the history of market returns, was the crash of '87 unusual?
2. How do conditional variance models (such as GARCH) behave around periods of extreme moves in the market?
3. What is the impact of the crash on backtesting and performance evaluation?

I. Was the Crash of '87 Unusual?

The average daily return and the standard deviation of the daily return on the S&P 500 Index over the last two decades have been about 0.066% and 0.96%, respectively. On October 19, 1987, the index had a return of −20.5%, which is approximately a 20-sigma event. If we make the simplifying assumption that daily returns follow a lognormal distribution, then the probability of observing a 20-sigma event is approximately equal to 2.75×10^{-89}. Based on this analysis, we would conclude that the crash of '87 was a rare and unusual event.

Figure 1: Distribution function of the minimum outcome from N draws from a normal distribution

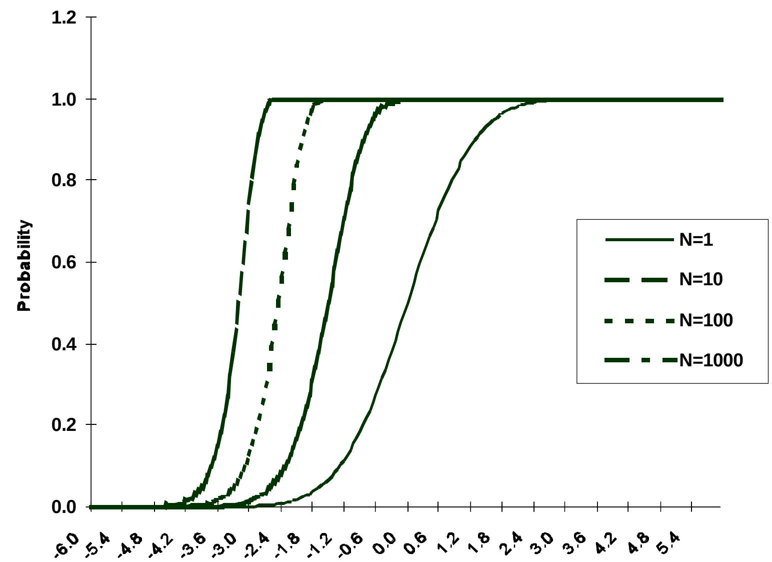

Effects of Repeated Draws from One Distribution

We can learn a bit more about the likelihood of a crash by taking a slightly different perspective. The return of −20.5% does not represent a single draw from a lognormal distribution. The history of publicly available daily returns on the U.S. stock market goes back over 100 years, and the random return on October 19, 1987, represents but one of the over 25,000 daily returns that have been observed over the last century. A more appropriate question to ask is: Given that we have observed 100 years of returns, what is the probability that one of the observed returns is −20.5%? Since the return on the S&P 500 Index on October 19, 1987, is the lowest on record, we can ask this question slightly differently as well: Given that we have observed 100 years of returns, what is the probability that the minimum daily return we will observe is −20.5%?

To see how much difference this perspective makes, let us consider a simple example. Let $X1, \ldots, X_N$ denote N independent draws from a normal distribution with mean zero and standard deviation 1. Define a new random variable Y as follows: $Y = \min(X_1, \ldots, X_N)$. Figure 1 graphs the cumulative distribution function of Y for $N = 1, 10, 100, 1000$. As we would intuitively expect, the figure shows that the distribution of Y shifts to the left as the value of N increases. Table 1 lists the probability that Y is less than −2 (a 2-sigma

continued on page 4, column 1

TABLE 1
Probability That the Minimum Outcome from *N* Draws Is a 2-Sigma or a 3-Sigma Event

N	Prob ($Y < -2$)	Prob ($Y < -3$)
1	0.023	0.001
10	0.214	0.013
100	0.911	0.126
1000	1.000	0.741

October's Market Demons
continued from page 3

event) and the probability that Y is less than -3 (a 3-sigma event).

Table 1 clearly shows that events that may be viewed as very unlikely to occur become much more likely to occur when we take into account the fact that we are making repeated draws from the same distribution. For example, the likelihood of a 3-sigma event when we make a single draw is 0.1%. In contrast, if we sample 1000 times, the likelihood that the minimum draw is less than -3 sigma is 74.1%. An inspection of the numbers in the table reveals another interesting fact: For small values of N (e.g., $N=1$, 10), the probabilities within each column increase linearly with N. For example, the probability that Y is less than -2 when $N=10$ (0.214) is approximately ten times the probability that Y is less than -2 when $N=1$ (0.023). It can be shown that this approximate linear relationship holds for small probability events (such as 2-sigma events under a lognormal distribution) and small values of N.

The Effect of Increasing Observation Horizons

We now evaluate the likelihood of a crash, using this slightly different perspective. Assume that daily returns on the S&P 500 Index are drawn from a lognormal distribution with mean and standard deviation equal to the sample mean (0.066%) and sample standard deviation (0.96%) observed over the last two decades [1]. Using these assumptions, we can construct the theoretical cumulative probability distribution function for the minimum daily return observed over horizons ranging from one day to 100 years (see Figure 2). Figure 2 shows that tail events become much more likely as we increase the observation horizon. For example, the likelihood that the minimum negative daily return is -4% or lower is 0.0012% over a given day but increases to 26.27% over a 100-year horizon. However, in spite of the increase in likelihood of tail events due to an increase in the number of observations, it is clear from the figure that a minimum return of -20.5% is still virtually impossible to explain using data on daily returns.

It is clear from the preceding discussion that increasing the observation horizon will increase the likelihood of tail events. Since the lognormal distribution assigns positive probabilities to returns in the range (-100% ∞) there surely must be an observation horizon over which a minimum daily return of -20.5% is likely. But this line of inquiry is not very satisfying. For example, we would not find it very comforting to know that a minimum daily return of -20.5% is very likely to happen over a million years! Instead of increasing the observation horizon, we investigate two other avenues of research:

- How does the analysis above change if we increase the return horizon from daily to monthly?
- How would a change in the distributional assumption affect our conclusions?

Figure 2: Distribution function of the minimum daily return over horizon of length T (assuming lognormal distribution)

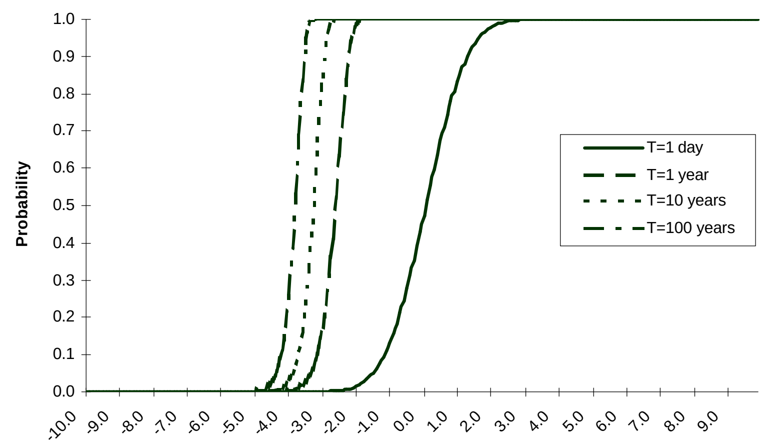

Figure 3: Distribution function of the minimum monthly return over horizon of length T (assuming lognomal distribution)

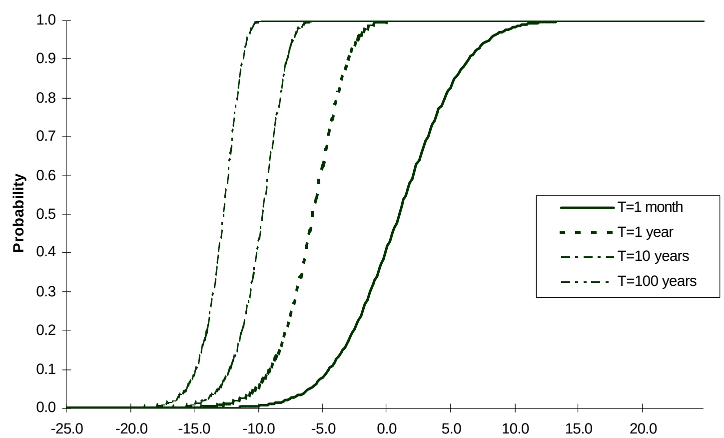

The Effect of Changing Return Horizons

First, we examine the effect of a change in the return horizon on our conclusions. The monthly mean return on the S&P 500 Index is about 1%, with a

continued on page 5, column 1

October's Market Demons
continued from page 4

standard deviation of about 4%. Using these statistics, Figure 3 graphs the cumulative probability distribution function for the minimum monthly return over observation horizons ranging from one month to 100 years. The probability that the minimum monthly return over a 100-year observation period is less than −21.5% (which was the return on the S&P 500 Index over the month of October 1987) is approximately 0.0067%. These numbers suggest that even when we look at monthly returns, the market crash represents a very unlikely event.

The Effect of Changing Distribution Assumptions

Next we examine the effect of a change in the distributional assumption on our results. It is well-known that the unconditional distribution of stock returns is characterized by the presence of fat tails. A direct implication of this is that tail events are more likely than the lognormal distribution would predict. This line of inquiry has a long history. Fama [2] concluded that stock returns appeared to be drawn from a member of the stable Paretian family of distributions with infinite variance. The normal distribution belongs to the stable Paretian class and is the only member of this class with finite variance. Subsequent researchers have shown that if the time series of market returns is drawn from normal distributions with time-varying variances, then the unconditional distribution of market returns would have fat tails.

One popular alternative to the lognormal assumption is to assume that the unconditional distribution of stock returns is log-*t*. The *log-t* distribution arises when stock returns for each period are lognormally distributed, with each period's variance being drawn from an inverted gamma distribution. If a random variable *U* has a log-*t* distribution with *v* degrees of freedom, then log $(U) \equiv t_v$, where t_v follows a *t*-distribution with *v* degrees of freedom. The expected value of log (*U*) is zero, and the variance of log (*U*) is equal to:

$$\frac{v}{v-2} .$$

Let: $$\sigma_v = \left(\frac{v}{v-2}\right)^{1/2} .$$

Figure 4: Distribution function of the minimum daily return over horizon of length T (assuming t-distribution with 5 df)

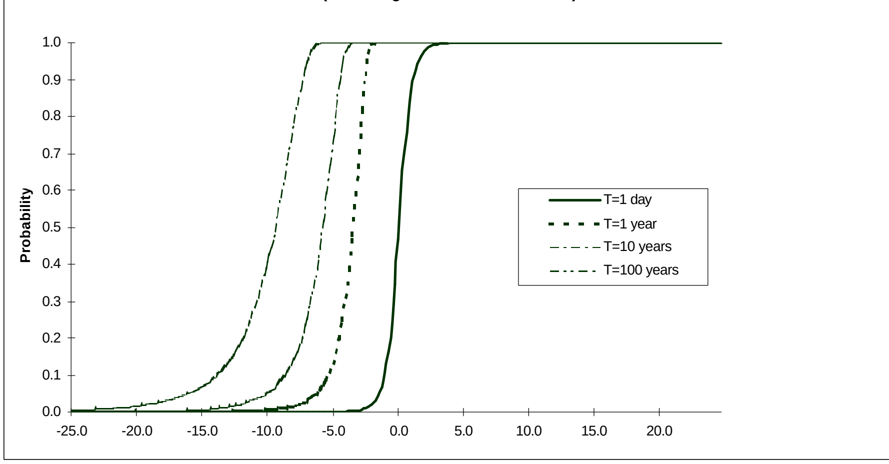

Figure 5: Distribution function of the minimum daily return over horizon of length T (assuming t-distribution with 3 df)

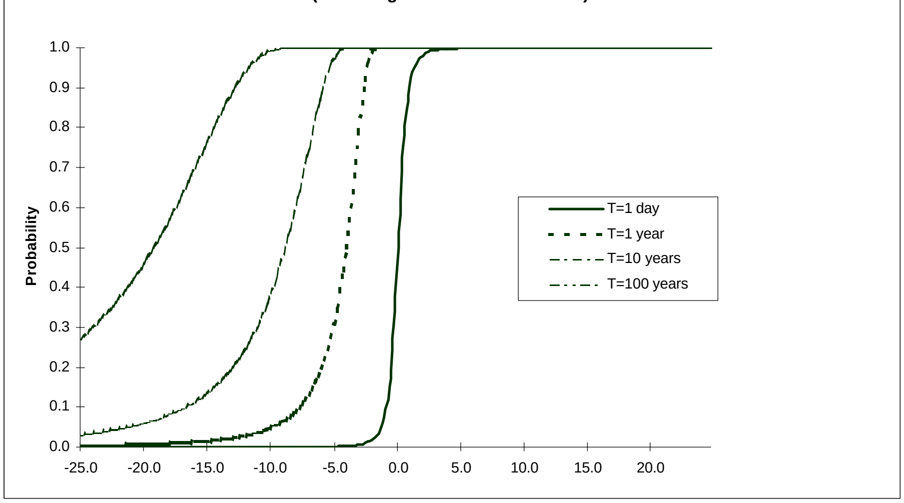

Let *r* denote the log of 1 plus the rate of return on the market. The mean and standard deviation of *r* are denoted by μ and σ respectively.

In our study, we assume that:

$$\sigma_v \left(\frac{r-\mu}{\sigma}\right) .$$

follows a *t*-distribution with *v* degrees of freedom. We will present results for the cases $v=5$ and $v=3$. A point worth noting about *t*-distributions is that all even moments of orders equal to or greater than the v^{th} moment are infinite. So, for example, when $v=5$, even moments of order 6 and above are infinite; and when $v=3$, even moments of order 4 and above are infinite. (Note that in the latter case the distribution has infinite kurtosis.)

Figures 4 and 5 display the cumulative probability distribution function for the minimum daily return for observation horizons ranging from one day to 100 years for $v=5$ and $v=3$, respectively. An examination of the graphs reveals that, as anticipated, tail events

continued on page 6, column 1

October's Market Demons
continued from page 5

are now much more likely than they were under the lognormal distribution. When the t-distribution has 5 degrees of freedom, the probability that the minimum daily return observed over 100 years is less than −20.5% is about 1.52%—still a relatively unlikely event. For a t-distribution with 4 degrees of freedom (not graphed), the crash probability is 8.93%, still a low number. In contrast, when the t-distribution has 3 degrees of freedom, the same probability jumps to 43.41%—nearly even odds of a crash over 100 years!

Our answer to the question "Was the crash of '87 unusual?" is thus somewhat tentative. When the crash is viewed not in isolation but as the worst outcome of a number of draws from the same lognormal distribution, then its likelihood increases—but not to a level that makes it very likely to happen. On the other hand, when we consider fat-tailed distributions (e.g., the log-t distribution), we see that a market crash becomes more likely to occur.

We have confined ourselves to a study of the distribution of the minimum daily return over various horizons, using a variety of assumptions regarding the distribution of daily stock returns. In principle, this analysis can be extended to a study of the likelihood of the worst K returns ($K=1,2,\ldots$) over the past 100 years. For example, the second smallest daily return on the S&P 500 Index over the past 100 years was −12.3%, on October 28, 1929. Using the distribution function of order statistics, we can extend the above analysis to study the likelihood that the two smallest returns on the S&P 500 are −20.5% and −12.3%

Alternative Return Distributions Compared

Given our results, a natural question to ask is: What is the true unconditional distribution of stock returns? Table 2 sheds some light on this question by tabulating selected theoretical and empirically observed percentile points for daily returns on the S&P 500 Index. The theoretical distributions have been calibrated to have the same mean and standard deviation as the sample mean (0.066%) and sample standard deviation (0.96%).

An inspection of Table 2 reveals that the empirically observed mid-range percentile points (e.g., the 25th and 75th percentiles) are closer to the theoretical values for the two t-distributions, while the extreme percentiles (e.g., the first and 99th percentiles) are closer to those of the lognormal distribution. The values in Table 2 do not offer clear evidence on the appropriate distributional form for index returns. It would be useful to look at higher order moments to get some more clues. Table 3 presents the skewness and excess kurtosis coefficients for the observed time series of returns and the values implied by the theoretical distributions considered above.

Table 3 shows the dramatic effect of the crash on the sample skewness and excess kurtosis coefficients. When the crash is included, it is clear that it is

continued on page 7, column 1

TABLE 2
Theoretical and Empirical Percentile Points for the Distribution of Daily Returns on the S&P 500 index (Returns Expressed in %)

Percentile	Empirical	Lognormal	Log-*t* with 5 Degrees of Freedom	Log-*t* with 3 Degrees of Freedom
1	−2.29	−2.17	−2.44	−2.45
5	−1.33	−1.51	−1.43	−1.24
10	−0.92	−1.17	−1.03	−0.84
25	−0.38	−0.58	−0.48	−0.36
50	0.07	0.07	0.07	0.07
75	0.54	0.71	0.61	0.49
90	1.07	1.30	1.16	0.97
95	1.49	1.65	1.57	1.37
99	2.35	2.30	2.57	2.58

TABLE 3
Theoretical and Empirically Observed Skewness and Excess Kurtosis Coefficients for Daily Returns on the S&P 500 Index

Statistics	Empirical (Including 10/19/87)	Empirical (Excluding 10/19/87)	Lognormal	Log-*t* with 5 Degrees of Freedom	Log-*t* with 3 Degrees of Freedom
Skewness	−3.30	−0.24	0	0	0
Excess Kurtosis	79.67	7.53	0	6	∞

TABLE 4
Kolmogorov-Smirnov (KS) and Kuiper (KP) Test Statistics for Hypotheses Regarding the Distribution of Daily Stock Returns

Distribution	KS Test Statistic	*P*-Value of KS Test Statistic	KP Test Statistic	*P*-Value of KP Test Statistic
Lognormal	0.083	0.0001	0.153	0.0001
Log-*t* w/5 Degrees of Freedom	0.049	0.0001	0.086	0.0001
Log-*t* w/3 Degrees of Freedom	0.018	0.1306	0.028	0.0310

October Market Demons
continued from page 6

difficult to reconcile the sample higher order moments with the theoretical moments of any single distribution considered above. On the other hand, when the crash is excluded, the log-*t* distribution with 5 degrees of freedom appears to have predicted moments that match the empirically observed moments closely. However, since the crash did occur, it is debatable whether it should be dropped from the analysis simply because it represents an inconvenient data point!

We now turn to more formal tests of the distribution of daily stock returns—namely, the Kolmogorov- Smirnov test and the Kuiper test. Table 4 presents the test statistics and the associated significance levels.

The test statistics in Table 4 strongly reject the null hypotheses that daily returns arise from a lognormal distribution or from a log-*t* distribution with 5 degrees of freedom. The null hypothesis of log-*t* with 3 degrees of freedom fails to be rejected by both tests at the 1% level but is rejected by the Kuiper test at the 5% level. The results of the formal tests are thus consistent with our earlier findings and strongly suggest that the unconditional distribution of daily returns is fat-tailed with very large (possibly infinite) higher-order moments.

TABLE 5
Daily GARCH-Forecast-Standardized Residuals for the Days Surrounding the October 1987 Crash

Date	Return on S&P 500 Index (%)	GARCH-Predicted Standard Deviation (%)	GARCH-Forecast-Standardized Residual
10/13/87	1.66	1.06	1.51
10/14/87	-2.95	1.08	-2.78
10/15/87	-2.34	1.20	-2.01
10/16/87	-5.16	1.25	-4.17
10/19/87	-20.47	1.55	-13.25
10/20/87	5.33	4.01	1.31
10/21/87	9.10	4.04	2.24
10/22/87	-3.92	4.28	-0.93
10/23/87	-0.01	4.25	-0.02
10/26/87	-8.28	4.15	-2.01

TABLE 6
Monthly GARCH-Forecast-Standardized Residuals for the Months Surrounding the Crash in October 1987

Month	Return on S&P 500 Index (%)	GARCH-Predicted Standard Deviation (%)	GARCH-Forecast-Standardized Residual
September 1987	-2.20	4.36	-0.75
October 1987	-21.52	4.34	-5.20
November 1987	-8.16	7.41	-1.25
December 1987	7.35	6.93	0.91

II. GARCH Forecasts Around Periods of Extreme Market Movements

In Part I of this article we studied the unconditional distribution of stock returns over the last 100 years with special focus on the likelihood of a market crash. Our conclusion was that a market crash has nearly even odds of occurring over a period of 100 years if the unconditional distribution of daily stock returns arises from a fat-tailed distribution with very large (possibly infinite) higher-order moments. It has been widely documented that such an unconditional distribution is consistent with each period's returns being conditionally lognormally distributed with time-varying conditional variances. In this section, we focus on a particular parameterization of the conditional variance structure—namely, the GARCH (1,1) model—and study the behavior of this model around periods of extreme market movements.

GARCH Applied to October 1987

To perform this study, we estimated separate GARCH(1,1) models using daily and monthly returns on the S&P 500 Index. The daily model was estimated using returns over the period March 1980 through September 1987 (1,906 days), and the monthly model used data from January 1973 through September 1987 (177 months). Table 5 presents the GARCH-forecast-standardized residuals and other numbers of interest for the days surrounding the crash in October 1987. As the estimation period for the models excluded October 1987, our reported results are out-of-sample.

Table 5 documents a number of interesting facts:

- In response to sharp market moves in the days immediately preceding the crash, the GARCH forecast of the standard deviation for October 19, 1987, was about 50% higher than it had been about a week before the crash.
- The crash return constitutes a 13-sigma event relative to the GARCH forecast volatility for October 19, 1987, in contrast to the 20-sigma characterization of the crash using unconditional moments of the distribution of daily returns.
- After the crash, GARCH forecast volatility rises to a level of over 4% per day, which causes many of the sharp post-crash market movements to be classified as "normal" events that are plausible even if daily returns are conditionally lognormally distributed.

A look at the time series of GARCH forecasts shows that the predicted volatility continues to be very high for several weeks after the crash. For example, the daily GARCH forecast as of the end of December 1987 (using data through December 1987 to estimate

continued on page 8, column 1

October Market Demons
continued from page 7

the GARCH parameters) was 1.60%—a number that is about 50% higher than pre-crash forecasts. This is a manifestation of the well-known high degree of persistence in daily GARCH forecasts. Table 6 shows GARCH-standardized residuals using the monthly GARCH model. The monthly GARCH forecasts also rise sharply following the month of the crash and continue to remain high for a few months after the crash.

The fundamental intuition built into GARCH models is the notion of volatility clustering—i.e., periods of high volatility are likely to be followed by more periods of high volatility. If historical volatility is low, then GARCH models will continue to forecast low volatility. Although "outliers" are not ruled out even when using GARCH forecasts, the distinguishing feature of an accurate GARCH model is that these outliers would be randomly distributed in time, in contrast to the forecasts of naïve models where outliers would appear clustered together. Judged by this metric, Tables 5 and 6 provide anecdotal evidence that, although the crash itself appears as an outlier, GARCH models are at least partially successful in explaining the sharp movements around the period of the market crash.

Can We Predict an Abrupt Market Transition?

Since GARCH models use a weighted average of historical realized volatility to predict future volatility, in times of a transition from a low-volatility regime to a high-volatility regime the first few sharp movements may appear as outliers that are unanticipated by the GARCH model. An interesting question we might ask is: Are there other techniques that might be used to predict extreme market movements? This question is of clear interest in the current regime since popular debate in the weeks leading up to October 27, 1997, centered on comparisons with October 1987 and on the likelihood of another market crash.

One obvious answer is to look at S&P 500 Index option-implied volatility forecasts. Figure 6 shows the time-series evolution of the S&P 500 Index level over the past one year and the S&P 500 Index (SPX) option-implied volatility at the beginning of each month from July 1996 through July 1997. The annualized average implied volatility using near-term

Figure 6: S&P500 level and SPX-option-implied volatility over the last one year

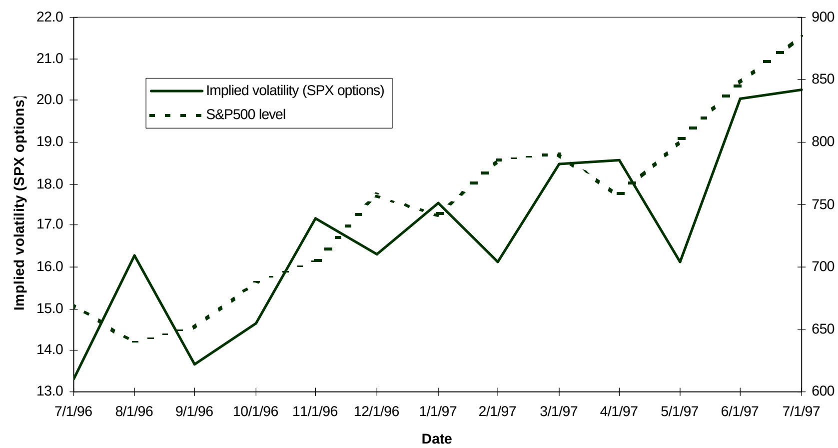

(less than one month to maturity), near-the-money options has risen from approximately 13.29% on July 1, 1996, to about 20.25% on July 1, 1997. Over the same time period, the S&P 500 Index has risen from 670 to 885. Somewhat surprisingly, over a number of months (e.g., May and June 1997) increases in the S&P 500 Index have been accompanied by increases in option-implied volatility, an observation which is at odds with the "leverage effect" (i.e., the usually negative relationship between price movements and volatility).

One explanation for Figure 6 is that options market participants expected the S&P 500 Index to have higher short-term volatility in the coming weeks and months. In contrast to the high implied volatility forecasts, the conditional variance prediction of GARCH models ranges from approximately 14.10% as of July 1, 1996, to approximately 15.55% as of July 1, 1997. Since the sharp movements that were anticipated by options market participants were realized in October 1997, we would expect that GARCH forecasts would have also risen subsequent to the first few sharp movements in the market.

Our study of option-implied volatility over the past one year suggests that we can incorporate "forward-looking" information in volatility forecasts by combining option-implied volatility with GARCH forecasts. For example, we could estimate a GARCH model using option-implied volatility as one of the variables in the conditional variance equation. Studies by Day and Lewis [3] and Lamoureux and Lastrapes [4] suggest that these two sources of information are complementary.

In summary, our study of GARCH-standardized residuals around the period of the crash of October 1987 shows that, while the crash itself was an outlier, most of the market volatility subsequent to the crash can be fully accounted for using GARCH forecasts. GARCH models use the presence or absence of outliers to predict subsequent increases or decreases in volatility. Hence, while outliers may exist even when using GARCH forecasts, these outliers are likely to be randomly dispersed through time. In the current regime, we saw that option-implied volatility as of July 1, 1997, appeared to be much higher than GARCH forecasts. One explanation for this finding is that options market participants expected to see higher volatility in the coming weeks/months for the S&P 500 Index. Since the expected increase in volatility has been realized, we would expect that GARCH forecasts will also respond.

III. Impact of the Crash on Backtesting and Performance Evaluation

In the previous sections we have studied issues relating to the likelihood of a crash (unconditional study) and the behavior of GARCH forecasts of the S&P 500 Index volatility around the period of the crash (conditional study). In this section, we provide some thoughts on

continued on page 9, column 1

October Market Demons
continued from page 8

the influence of the crash on backtesting investment strategies and on performance evaluation.

The Importance of Time Period Choice

The first point to note is that the time horizon over which backtesting and/or performance evaluation are conducted will determine the extent to which excluding the crash will affect the reported results. Figure 7 plots the ex-post Sharpe ratio on the S&P 500 Index over horizons of 1, 5, 10, and 20 years, including and excluding the crash. As the horizon lengthens, we see that the Sharpe ratio when the crash is included gradually approaches the Sharpe ratio excluding the crash. We should point out that because we need very large sample sizes to estimate mean returns accurately, the two sets of Sharpe ratios are not statistically distinguishable from one another (i.e., they are within two standard errors of each other).

The second point is that excluding the crash can have dramatic implications for the profitability of certain types of strategies. For example, Sheikh [5] demonstrates that a strategy of buying the S&P 500 Index plus writing out-of-the-money puts on the index was a profitable strategy (relative to buying the S&P 500 Index) over periods strictly before and strictly after the crash. The post-crash period that Sheikh studied was August 1988 through February 1995. In contrast, a similar strategy that was put in place starting in September 1987 lagged the cumulative return on the S&P 500 Index, even after over seven years (as of February 1995). In other words, the loss suffered in the month of the crash was more than the gains made by the strategy over the next seven years!

The third point is that it is a good idea to run backtests over historical periods that represent different regimes—e.g., bull and bear markets, periods of low volatility and high volatility, etc. Figure 8 shows the cumulative return on the S&P 500 Index over the 10-year period January 1987 through December 1996. It is evident from the figure that there have not been too many bad months, especially over the past five years. The crash represents a useful observation precisely because it was a particularly bad month. Including this observation in backtests serves as a check on the robustness of proposed investment strategies.

Figure 7: Annualized Sharpe ratio for S&P500 index (over various horizons) including and excluding October 1987

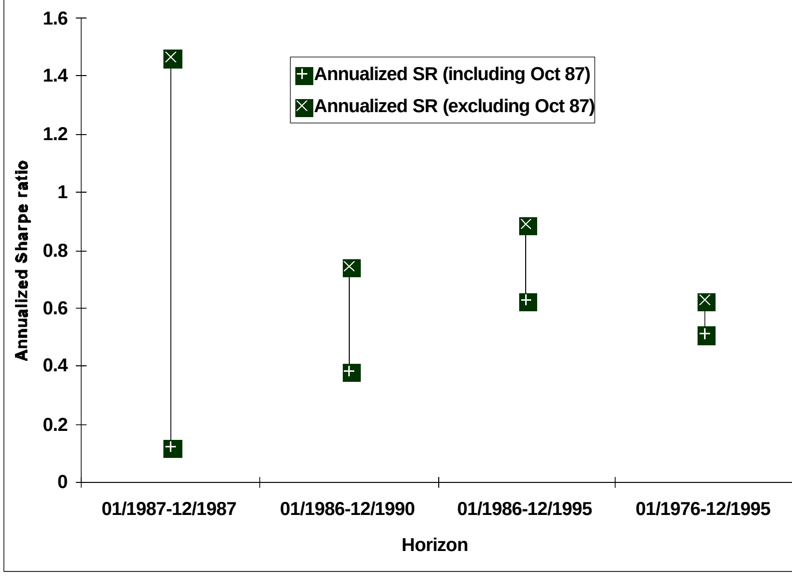

Should the Crash Be Included in Performance Studies?

Finally, we consider performance evaluation in the presence of the crash. As the above analysis of the Sharpe ratio suggests, the total risk/return picture, especially over smaller horizons, differs significantly depending on whether or not the crash is included in the sample. For an active manager who is usually fully invested in equities, including the crash does not bias performance results since the active manager is evaluated based on his or her active risk/return profile (i.e., risk and return net of the market).

Let us consider the more difficult question of an active manager who aims to achieve superior returns by forecasting the returns on the S&P 500 Index (that is, by timing the market). Let $r_{B,t}$ be the excess return on the S&P 500 Index in period t and let μ_B denote the per-period long-run expected excess return on the index. Each period, the market timer has a forecast of the excess return on the index over its long-run average. In symbols, for each period the market timer has a forecast Δf_B of the value of $r_{B,t} - \mu_B$. Let λ_{BT} denote the risk aversion coefficient of the investor for benchmark *timing* and let $\hat{\sigma}^2_{B,t}$ be the investor's forecast of the variance on the index over period t. Then, the optimal active beta position for the investor is given by:

$$\beta^*_{p,A,t} = \frac{\Delta f_{B,t}}{2\lambda_{BT}\hat{\sigma}^2_{B,t}}$$

Grinold and Kahn [6] discuss the appropriate objective function for an active manager and derive the optimal active beta policy stated above. We conducted a simulation study using the actual history of realized market returns over the period January 1987 through December 1996. The market timer is assumed to make monthly forecasts of the index return. Each month, the market timer receives a signal $g_{B,t}$ as follows:

$$g_{B,t} = IC\left(r_{B,t} - m_B\right) + s\sqrt{1-IC^2}\,u_t$$

IC = the information coefficient of the manager
m_B = the average excess return on the index over the 10-year sample period
s = the sample standard deviation of the excess return on the index, and
u_t = a random number drawn from a distribution with zero mean and unit standard deviation.

continued on page 10, column 1

October Market Demons
continued from page 9

$r_{B,t}$ = is set equal to the observed excess return on the S&P 500 Index in month t.

Given this signal, the manager constructs an optimal forecast of the excess return on the index as follows:

$$\Delta f_{B,t} = E\left(r_{B,t} - \mu_B \middle| g_{B,t}\right) = (IC) g_{B,t}$$

For simplicity, we assume that the investor's forecast of the variance of the index return in month t, $\hat{\sigma}^2_{B,t}$ is equal to s^2 for all months. For a given sequence of signal realizations, we can derive the corresponding time series of active beta positions. Using the actual history of market returns, we can then compute the ex-post information ratio for the investor. We considered three different *IC levels; 0.05, 0.10,* and *0.15*. For each *IC*, we ran 100 simulations of the entire 10-year history from January 1987 through December 1996. Table 7 reports the average ex-post information ratios across these simulations.

Table 7 shows that there are no significant differences between the two columns of information ratios. In other words, including the crash does not appear to make a difference for the performance evaluation of a market timer.

Summary

In this article, we have presented some perspectives on the crash of October 1987. We found that the likelihood of a market crash increases dramatically if the unconditional distribution of stock returns is fat-tailed with very large (possibly infinite) higher-order moments. Our study of GARCH forecasts showed that, with the exception of the crash itself, these forecasts were at least partially successful in capturing sharp movements around the period of the crash. We found that option-implied volatility has increased dramatically over the past one year, suggesting that the market expects higher volatility in the weeks/months ahead. Finally, we offered some thoughts on the impact of the crash on backtesting and performance evaluation. We showed via a simulation study that including the month of the crash does not have a significant effect on the ex-post information ratios of a market timer.

Figure 8: Cumulative return on the S&P500 index since January 1st, 1987

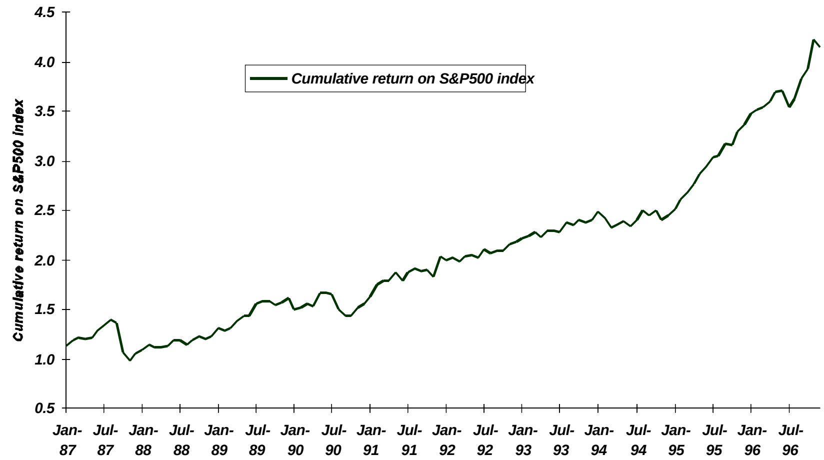

End Notes

1. Throughout this section, we assume that the time series of daily returns *iid* (independent and identically distributed) draws from the specified unconditional distribution. In Part II of this article we will explore the behavior of conditional variances around the period of the crash.

2. Fama, Eugene F., "The Behavior of Stock Market Prices," *Journal of Business*, no. 38, 1995, pp.34–105.
3. Day, Theodore E. and Craig M. Lewis, "; Stock Market Volatility and the Information Content of Stock Index Options," *Journal of Econometrics*, v.52, 1992, pp.267–287.
4. Lamoureux, Christopher and William Lastrapes, "Forecasting Stock Return Variance: Towards an Understanding of Stochastic Implied Volatility," *The Review of Financial Studies*, no. 6, 1993, pp.293–326.
5. Sheikh, Aamir, "Portfolio Construction with Derivatives," BARRA's 19th Annual Equity Research Seminar, Pebble Beach, CA, June 11–14, 1995.
6. Grinold, Richard C. and Ronald N. Kahn, Active Portfolio Management, Probus Publishing, Chicago, IL, 1995.

Vinod Chandrashekaran is a Senior Consultant in the Research Group at BARRA Inc. in Berkeley, California.

TABLE 7
Ex-post Information Ratios (IR) for Market Timing
Over the Period January 1987 through December 1996
(Average over 100 Simulations)

Information Coefficient	Average Ex-post IR (including October 1987)	Average Ex-post IR (excluding October 1987)
0.05	0.191	0.195
0.10	0.396	0.377
0.15	0.467	0.448