Preliminary Draft: Do not Quote

AN EMPIRICAL APPLICATION OF USING EXTREME VALUE THEORY TO MEASURE STOCK MARKET RETURNS

Therese Pactwa*

Arun J Prakash**

St. John's University, New York

Florida International University, Miami

Address for Correspondence
Arun J Prakash
Department of Finance
Florida International University
Miami, Fl 33193
(305) 386- 9042

prakasha@fiu.edu

AN EMPIRICAL APPLICATION OF USING EXTREME VALUE THEORY TO MEASURE STOCK MARKET RETURNS

Abstract

Extreme stock price movements are of great concern to both investors and the entire economy. For investors, a single negative return, or a combination of several smaller returns, can possible wipe out so much capital that the firm or portfolio becomes illiquid or insolvent. If enough investors experience this loss, it could shock the entire economy. An example of such a case is the stock market crash of 1987. Furthermore, there has been a lot of recent interest regarding the increasing volatility of stock prices.

This study presents an analysis of extreme stock price movements. The data utilized was the daily returns for the Standard and Poor's 500 Index from January 3, 1978 to July 30, 2004. Research questions were analyzed using the statistical models provided by extreme value theory. One of the difficulties in examining stock price data is that there is no consensus regarding the correct shape of the distribution function generating the data. An advantage with extreme value theory is that no detailed knowledge of this distribution function is required to apply the asymptotic theory. We focus on the tail of the distribution.

Extreme value theory allows us to estimate a tail index, which we use to derive bounds on the returns for very low probabilities on an excess. Such information is useful in evaluating the volatility of stock prices. There are three possible limit laws for the maximum: Gumbel (thick-tailed), Fréchet (thin-tailed) or Weibull (no tail). Results indicated that extreme returns during the time period studied follow a Fréchet distribution. Thus, this study finds that extreme value analysis is a valuable tool for examining stock price movements and can be more efficient than the usual variance in measuring risk.

AN EMPIRICAL APPLICATION OF USING EXTREME VALUE THEORY TO MEASURE STOCK MARKET RETURNS

In finance, one often works with near continuous-time observed (so-called high density) data. At the same time, marginal distributions are heavy-tailed and return data exhibits clustering of extremes and long-range dependence. Due to these complications there is no universally accepted nor indeed easy model that explains all of these phenomena.

In this paper, we are attempting to estimate the extreme values, or the maximum one-day shifts, of stock market returns. With extreme value theory, we are focusing on the probability distribution of the maxima, rather than of the individual observations. Utilizing this technique, we have the advantage of making as few assumptions as possible about the individual returns.

In the first section, we provide a review of the literature on extreme value theory with applications to finance. Next, we describe the data utilized and its source. We also present an outline of the statistical testing we perform, using both descriptive statistics and extreme value theory. Finally, we conclude with some limitations of the study. These are mainly due to computing limitations because of the size of the data.

I. Introduction

Extreme value theory is a classical topic in probability theory and mathematical statistics. It is the foundation of all statistical techniques for identifying outlying observations. Financial time series data consist of daily or weekly reported speculative prices of assets such as stocks, foreign currencies, or commodities such as coffee, corn, oil, livestock, sugar, etc. Firms perform risk management to guard against the risk of loss due to the fall in prices of financial assets held or issued by the company. What are of importance here are the magnitudes of the changes in prices, rather than the average variations. A single, extremely negative return, or a combination of several smaller negative returns, can possibly wipe out so much capital that the firm or portfolio becomes illiquid or insolvent. Trading limits need to be set. These limits are a function of the probability of a single negative return so large that capital is endangered.

Daily or weekly returns, measured by the relative differences of consecutive prices or differences of log-prices, are the appropriate quantities which must be investigated. The special feature of this type of return series is the alternation between periods of tranquility and volatility.

Most importantly, there is empirical evidence that distributions of returns can possess heavy (fat) tails so that a careful analysis of returns is required. Fat-tailed distributions exhibit more probability mass in the tails than distributions such as the standard normal distribution. This means that extremely high and low realizations will occur more frequently than under the hypothesis of normality.

Studies of financial data are faced with the paradoxical situation that extreme risks are, by definition, rare; whereas significant statistical results can only be achieved if a sufficient number of these events can be analyzed. Unfortunately, the number of data in real cases is relatively limited. One possible solution is to use high frequency data for financial assets. However, a risk manager usually likes to evaluate risk over longer time periods. Thus, it is important to be able to compute the behavior of the tail index under time aggregation. Low frequency risk must be extrapolated from high-frequency data. Extreme value theory places emphasis on the tail behavior as the basis for the analysis. In other words, one is interested in the distribution of the largest order statistic, say M_n, which gives the likelihood of an extreme realization. For this reason, we focus on the tail of the distribution.

This approach is similar in theory to McCulloch (1981), who looked at the interest rate risk of commercial banks. McCulloch presumes a specific distribution for this risk, from which bankruptcy probabilities are calculated. In an earlier contribution, Roy (1952) proposed the more robust safety-first criterion, which does not rely on a specific distribution, but instead uses the Tchebysheff bound.[1] One can improve considerably on this bound by exploiting the limit law of the distribution, which is the approach taken by extreme value theory and in this paper.[2]

Extremal events in finance have the advantage that they are mostly quantifiable in units of money (as opposed to other extremal occurrences, such as floods and earthquakes, which might cost the loss of lives). Such events include the stock market crashes of 1929 and 1987. Incidents in the derivatives markets include the collapse of Barings Bank and the losses of Metallgesellschaft, Proctor and Gamble, Kashima Oil, Orange County (California), and

---

[1] Tchebysheff's theorem applies to any distribution and states that for $k \geq 1$, at least ($1 - \frac{1}{k^2}$) of a set of n-measurements will lie within k standard deviations of their mean.

[2] An extension of this problem is the so-called Value-at-Risk problem. One studies the amount of capital $C(q, T, I)$ which can be invested such that at time T the possible loss exceeds a certain amount I with probability q. See Jorion (1997) for more on Value-at-Risk.

Sumitomo Bank. However, most such events also have a non-quantifiable component. These events can cause a larger shock to the underlying financial system. This is sometimes called contagion risk, or the risk that a securities market decline in one country will spread to another, causing serious market losses, or that the failure of one large financial institution will lead to the failure of others. Thus, it is important to study these circumstances.

In order to determine how extremal events occur, we need to use appropriate mathematical methods in order to explain events that occur with relatively small probability but have a significant influence on the behavior of the whole model. We need to find which distributions and stochastic processes typically describe extremal events. We first try to estimate a crash probability on outcomes worse than a critical loss level. We extrapolate the empirical distribution of the maximum order statistic, in combination with some conditions, to construct a list of probability-quantile combinations. In order to construct the probabilities or quantiles, one relies on the limit law for the maximum. Since we do not know exactly what distribution generates the data, the limit law provides an approximation.

II. Review of Applications of Extreme Value Theory to Finance

Within the context of finance, there has been some work analyzing specific data utilizing statistical methodology on extremes. For example, Rothchild and Stiglitz (1970) use the weight of the tails of two random variables to define risk (rather than variance). McCulloch (1978) studies discontinuities of the price process associated with large falls and rises of the stock market. Parkinson (1980) utilizes extreme values (the high and low prices) of common stocks to provide an estimate for the diffusion constant (or the variance of the rate of return) characterizing a random walk. He shows that this method is far superior to the traditional estimate using closing prices only. Parkinson (1977) also shows how these extreme values can be used to value put options.

The application of extreme value theory to finance has recently received a lot more attention in the literature. Hols and de Vries (1991), de Vries (1994), and Danielsson and de Vries (1997) and Müller, Dacorogna, and Pictet (1998) look at high frequency foreign exchange data. Müller et al. also examine the interbank money market cash interest rates for four major currencies. Longin (1996) analyzes over a century of U.S. stock market data (1885-1990).

Loretan and Phillips (1994) use the extremal index to test the covariance stationarity of both the stock market return and exchange rate return series.

In other types of applications, Lucas and Klaassen (1998) consider the effect of extreme returns on the solution of the asset allocation problem. Kluppelberg and Mikosch (1997) and Mikosch and Nagaev (1997) show how extreme value theory can be used to model large insurance claims. Phoa (1999) examines the use of probability *vis á vis* those recorded for the individual stocks evidences the effects of portfolio diversification whereby stock-specific shocks are mitigated. On the other hand, this probability is on the low side given that since 1899 there have been four days on which the Dow dropped by more then 10 percent.

Although the aggregate shocks covered by the sample period, like the oil shock of 1973, did not lead to excessive daily plunges, they did lead to sustained declines. To see whether extremal analysis does predict these sustained declines in the market indices, Jansen and de Vries (1991) use returns over longer time spans than a day. Probabilities were estimated by using monthly yields in excess of a return over the time span of one year. They determined that the probability on average within any year a monthly market drop exceeds 20 percent or 30 percent is 0.16 and 0.06, respectively, i.e., occurs once in about every 6 or 15 years, respectively. Hence, the probability of a sustained bearish market is quite high. While the events of October 19th may be exceptional, a drop of similar proportions over a somewhat longer period is not unlikely.

In a more recent study, Jones and Wilson (2000) looked at the extraordinary performance of the stock market during the 1990s. For the ten-year period January 1, 1989 through December 31, 1998, the stock market, as measured by the S&P 500, experienced a 20.8 percent annual compound rate of return. Furthermore, over the four-year period 1995 through 1998, the market enjoyed a compound rate of return of approximately 31 percent. For that same four-year period ending in 1998, the only higher geometric mean for a four-year period occurred in 1936, at 32.01 percent -- and this performance followed the greatest single loss in market history (-45.16 percent in 1931) and the third greatest loss in market history in 1930 (-27.41 percent). Based on the assumption of a lognormal stock price distribution, Jones and Wilson calculated probabilities associated with various target rates of return for the period 1920 through 1998. For a four-year investment horizon, the probability of a 30% or more compound rate of return is very small but positive (three to four percent). They also show that the chances for a "Dow 36,000" anytime soon is virtually zero.[3]

III. Description of the Data and the Statistical Testing

As discussed in the previous section, there is little disagreement about the qualitative properties of stock returns. Typically, daily returns are a stationary series that is strongly leptokurtic and possibly exhibits some low order serial dependence. Thus, the data meet the criteria for application of the theorems and the estimators proposed in this paper. The stock market return series utilized is the daily closing prices for the Standard and Poor's Composite stock market index (the "S&P 500") from January 3, 1978 to July 30, 2004 (n = 6,710). The quotes were downloaded from Yahoo.com. This data set is comparable to the data set used by Loretan and Phillips (1994), who analyzed the daily returns from that same index from July 1962 to December 1987 (n = 1,848). The data for the S&P 500 dates back daily until March 1957 and were recorded weekly beginning in January 1918. However, our data set was only available online from January 3, 1978 to the present. As a result, rather than replicating work already performed, we focused on the more recent period.

Research questions were analyzed using the traditional descriptive statistics and the statistical models provided by extreme value theory. For the statistical testing using extreme value analysis, data was entered via the integrated, Pascal-like, programming language XPL. The statistical software utilized for the extreme value analysis was Xtremes[4]. Such data can be visualized by the means of nonparametric tools, such as kernel densities. One may also try to fit a parametric density, e.g., a Gumbel density, with location and scale parameters μ and σ, to the sample density.

A trend is a long-term change of a series of data. Many observations recorded at specific times exhibit a dependence on time. This dependence may be caused by inflation, seasonal effects, etc. There are statistical procedures that measure and remove a trend or seasonal component in a series of data. After having removed this trend, one obtains residuals which may

---

[3]Glassman and Hassett (1999) have written both an article and a book proclaiming their hypothesis that stocks are significantly undervalued because the risk premium on equities has declined. They believed the market in the 1990s was in the midst of a one-time-only rise to the neighborhood of 36,000 for the Dow Jones Industrial Average.

[4] Version 2.1, by the Xtremes Group, Siegen, Germany, included in Reiss and Thomas (1997).

be dealt with. For stock prices, one must first eliminate the positive trend, due to growth of the economy, to get to the martingale model.

Before we can analyze any speculative return series, we must first deal with the fact that prices are computed at equally-spaced moments in time, yet business time does not coincide with physical time. For example, stock markets are closed on weekends and holidays. This poses a problem for the definition of returns at high frequencies. Reiss and Thomas (1997) suggest the following ad hoc procedures to deal with the "weekend effect."

First, Monday returns can be omitted. In this case, we omit the days for which prices are not recorded, including the consecutive day. Thus, after a weekend, the Monday returns are also omitted. Second, Monday returns can be distributed. The return registered after a gap (e.g., the return recorded on Monday) is equally distributed over the relevant days (e.g., if *r* is the return on Monday, then *r*/3 is taken as the return on Saturday, Sunday and Monday).

C. Limitations of the Study

One of the major limitations in this study was the amount of data. Rather than the traditional limitation that the data points are too few, the data points in this case are too many. As a result, the data file for the original sample from 1/3/78 to 5/31/01 was too large for the software package utilized to examine the extreme returns.[5] The test sample had to be reduced so that the software could handle it. This could possibly result in a sample bias error.

IV. The Results

This section of the paper is divided into two parts. We first provide the results from the descriptive statistics employed. Next, we present the results from applying the procedures provided by extreme value analysis.

A. Descriptive Statistics

In this section, we examine the descriptive statistics for our data. We look at frequency histograms and time series properties. We also evaluate the numerical procedures and discuss the volatility of returns. Finally, we compare a few methods for dealing with gaps in returns.

Table 1 below shows frequencies and relative frequencies of extreme daily stock price movements (both gains and losses) for the time period of January 3, 1978 through July 30, 2004. We can see that as we increase the cut-off level as to what we define as extreme returns, the

[5] The test size of the XPL-editor is limited to 64 Kbytes. Data sets are entered using a text editor.

number of observations declines substantially. Plus, there were some years that had no extreme returns at all (e.g., 1992, 1995, and 2004, at least until the end of July). We can also see that the level of incidence has been higher in the more recent years. This information is also presented as a histogram in Figure 10.

Table 1 - Frequencies of Extreme Daily Stock Price Returns (1/3/78 – 7/30/04)

Return		**2%**		**3%**		**4%**
Year	**Frequency**	**Relative Frequency**	**Frequency**	**Relative Frequency**	**Frequency**	**Relative Frequency**
1978	4	0.012	1	0.011	0	0.000
1979	3	0.009	0	0.000	0	0.000
1980	12	0.035	2	0.022	0	0.000
1981	7	0.021	0	0.000	0	0.000
1982	17	0.050	6	0.065	1	0.032
1983	4	0.012	0	0.000	0	0.000
1984	7	0.021	0	0.000	0	0.000
1985	1	0.003	0	0.000	0	0.000
1986	9	0.027	2	0.022	1	0.032
1987	40	0.118	9	0.098	7	0.226
1988	16	0.047	4	0.043	2	0.065
1989	4	0.012	1	0.011	1	0.032
1990	13	0.038	2	0.022	0	0.000
1991	8	0.024	2	0.022	0	0.000
1992	0	0.000	0	0.000	0	0.000
1993	1	0.003	0	0.000	0	0.000
1994	2	0.006	0	0.000	0	0.000
1995	0	0.000	0	0.000	0	0.000
1996	3	0.009	1	0.011	0	0.000
1997	15	0.044	3	0.033	2	0.065
1998	23	0.068	9	0.098	3	0.097
1999	22	0.065	1	0.011	0	0.000
2000	36	0.106	10	0.109	3	0.097
2001	25	0.074	8	0.087	4	0.129
2002	52	0.153	27	0.293	7	0.226
2003	15	0.044	4	0.043	0	0.000
2004	0	0.000	0	0.000	0	0.000
Total	**339**	**1.000**	**92**	**1.000**	**31**	**1.000**

Fig. 10 - Frequencies of Extreme Returns

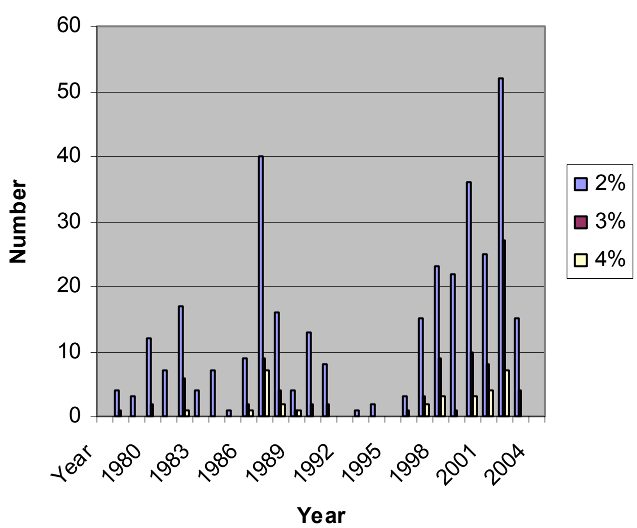

Table 2 reveals the exact dates of extreme returns in excess of three percent (both gains and losses). We can see evidence of the volatility clustering observed by Mandlebrot (1963a and b). Periods of quiescence and turbulence tend to cluster together. In addition, the number of positive (31) and negative (27) returns were fairly equal during the time period studied. These dates will be helpful when identifying the extreme returns in the time series plots shown below in Figures 11 through 20.

Table 2 - Dates of Returns in Excess of Three Percent (1978 - 2001)

Date	High	Low	Close	Return	Date	High	Low	Close	Return
5/30/00	1421.79	1377.8	1421.29	0.031026	10/19/87	245.28	224.83	224.83	-
0.204677									
9/2/97	927.58	899.47	927.57	0.031241	10/26/87	247.72	227.25	227.66	-0.08283
4/17/00	1399.01	1346.68	1399.01	0.031642	10/27/97	941.64	876.73	876.97	-0.06841
8/27/90	323.11	311.51	321.43	0.031845	8/31/98	1033.47	957.42	957.43	-0.06788
11/30/82	138.97	133.44	138.53	0.032342	1/8/88	261.07	242.94	243.39	-0.06772
10/6/82	126.33	121.82	125.97	0.032795	10/13/89	355.53	332.81	333.64	-0.06117
10/13/00	1373.32	1326.56	1373	0.032851	4/14/00	1440.51	1339.4	1356.1	-0.05883
10/28/99	1341.27	1296.71	1341.26	0.034356	12/21/00	1305.6	1261.16	1264.7	-0.05682
5/31/88	262.15	253.41	262.15	0.03449	10/16/87	298.91	281.52	282.69	-0.05163
4/25/00	1477.60	1429.86	1477.6	0.034495	9/11/86	247.05	234.66	235.17	-0.04809
10/19/00	1389.81	1342.13	1388.52	0.035166	4/14/88	271.57	259.37	259.75	-0.04345
8/20/82	113.49	109.33	113.02	0.035361	3/12/01	1233.42	1176.78	1180.2	-0.04318
9/23/98	1066.09	1029.63	1066.09	0.035421	11/30/87	240.33	225.77	230.3	-0.04173

1/4/88	256.43	247.08	255.94	0.035859	10/25/82	136.97	132.53	133.32	-0.03969
4/22/80	104.02	100.8	103.42	0.036273	10/22/87	258.37	242.99	248.25	-0.03917
1/17/91	327.96	316.25	327.96	0.037323	8/27/98	1084.19	1037.61	1042.5	-0.03843
9/1/98	1000.71	939.98	994.24	0.038447	1/4/00	1455.22	1397.41	1399.3	-0.03753
12/5/00	1376.56	1324.97	1376.24	0.038672	11/15/91	397.15	382.62	382.62	-0.03656
11/3/82	143.50	137.38	142.86	0.039057	8/4/98	1119.73	1071.82	1072.2	-0.03622
4/18/01	1248.42	1191.81	1238.16	0.039135	12/3/87	233.89	225.21	225.21	-0.03526
11/1/78	97.41	94.13	96.85	0.039833	4/3/01	1145.87	1100.19	1106.5	-0.03439
10/15/98	1053.09	1000.12	1047.49	0.041698	3/8/96	653.65	627.63	633.51	-0.03081
4/5/01	1151.47	1103.25	1151.44	0.04368	7/7/86	251.8	243.63	244.05	-0.0307
8/17/82	109.33	104.32	109.03	0.04756	9/30/98	1049.02	1015.73	1017.1	-0.03048
3/16/00	1458.09	1392.15	1457.7	0.048962	8/6/90	344.86	333.27	334.42	-0.03027
10/29/87	246.69	233.38	244.77	0.049299	10/1/98	1017.01	981.29	986.38	-0.03017
1/3/01	1347.76	1274.81	1347.56	0.050099	3/17/80	105.22	101.82	102.25	-0.03007
9/8/98	1023.46	973.89	1023.46	0.050888					
10/28/97	923.09	855.27	921.86	0.051188					
10/20/87	245.61	216.46	236.83	0.053374					
10/21/87	259.26	238.8	258.37	0.090951					

2. Time Series Behavior of Daily prices and returns over the sample period

Next, we examined time series properties of our sample data, entire as well as several subsamples. Upon examining these time series plots, several trends were evident. Up until the 2000 decade, the S&P 500 showed a positive, non-linear trend. This became even more evident in the 1990s. Although prices have increased over time, the rate of return has been much higher in recent years. Looking at the graph of the entire sample, we saw an increasing trend since the late 1970s, then a decline beginning in the year 2000.[6]

We also examined the graphs of the daily returns for the same time periods, respectively. We can observe a clear change in volatility possibly triggered by extreme returns. This is especially evident in the most previous decade of the 1990s and in the more recent data from the new decade. We can visualize the volatility by examining the peaks and troughs of the return data. In addition, we can once again look at the scale differences between the sample periods. These figures show greater volatility in the short run. Comparing the two most recent complete decades, we can seen more extreme spikes in the 1990s than in the 1980s.

[6] Plots can be obtained from the author upon request.

The decade of the 1990s deserved to be singled out because of the incredible performance of the stock market during that time period. The performance of the 1990s in general, and the four-year period 1995 through 1998 in particular, is remarkable. These years saw extraordinary returns not seen since the depressed conditions of the 1930s. Furthermore, the large returns of the 1990s did not follow the spectacular losses as was the case in the 1930s.

3. Numerical Procedures

Next in our analysis, we examine the numerical procedures for describing data. Table 3 presents the results for the entire sample of daily returns from January 1978 through May 2001, and also for the subsample from January 2000 through May 2001. The numbers in Table 3 are fairly typical for daily asset return data and indicate considerable deviation from normality.

Table 3 - Descriptive Statistics

	1/3/78-5/31/01	**1/3/00-5/31/01**
Mean	0.00048964	-0.0003341
Standard Error	0.00013068	0.00076959
Median	0.00047662	-0.00024158
Standard Deviation	0.01005185	0.01454096
Variance	0.00010104	0.00021144
Kurtosis	34.17601344	1.73369641
Skewness	-1.55294109	0.02959499
Range	0.29562782	0.10892456
Minimum	-0.2046765	-0.05882598
Maximum	0.09095132	0.05009858
Sum	2.89721108	-0.11927398
Count	5917	357
Confidence Level (95%)	0.00025612	0.00150837

Looking at the measures of central tendency, we can see a negative mean return for the more recent subsample. This confirms the downward trend of stock prices. For symmetric data, the mean is always equal to the median. For the time period January 1978 through May 2001, we see that the mean is greater that the median. This indicates that the distribution is right-skewed (compared to the left tail, the right tail is elongated). For the shorter time horizon, the opposite occurs. The mean is less than the median, and left-skewness is indicated. Examining this skewness issue further, we take the absolute value of the skewness measure. The higher the absolute value, the greater the skewness. We find skewness in both periods studied. However, the skew is much greater over the longer time horizon.

The kurtosis statistic reflects the peakedness of the center compared to that of the normal distribution. A distribution function with a kurtosis larger than 3.0, which is the kurtosis of a normal distribution function, is fat-tailed or leptokurtic. Our data showed a severe kurtosis in the data covering the longer time horizon.

The measures of variation are then examined. The range is almost three times as great for the longer time horizon, which is not surprising, given that there are more observations over different economic cycles. Since we only use the high and low values, the range is very sensitive to extreme values in the sample. Thus, it is considered to be a weak measure. So, we look at the variance and the standard deviation, which incorporate all of the values in the data set. We see that the variance is higher for the shorter time horizon. Based on the above analysis, we are justified in trying to fit extreme value theory to the data over the time period studied.

4. The Volatility of Returns

In Figures 10 and 11, squared daily returns for the series July 2, 1962 through December 31, 1987 and January 3, 2000 and May 31, 2001, respectively, are visualized in scatter plots. These illustrations show that there are periods of tranquility and volatility of the return series. This shows that while the return process is a fair game, risk is spread unevenly. Hence, the risk of an investment is somewhat clustered through time and somewhat predictable. We say that returns are conditionally heteroscedastic (the variance is varying conditional on the past). These charts illustrate the Martingale property of return time series data. We can also see evidence of greater volatility over the shorter time horizon.

5. Gaps in Returns

Next, we look at the effects of gaps in the data, or the weekend effect. To accomplish this, we transform *prices p(t)* into returns *r(t)* = log *p(t)* - log *p(t* - 1), whereby gaps in returns are dealt with in the following manner: 1) omit the days for which prices are not recorded including the consecutive day (e.g., omit Monday); 2) equally distribute the return registered after a gap over the relevant days (e.g., if r is the return on Monday, then r/3 is taken as the return on Saturday, Sunday and Monday); and 3) do not change the original data set.

The returns under these three scenarios are displayed in scatterplots in Figures 12 through 17. Figures 11 through 14 examine the returns from the period July 2, 1962 through December 31, 1987. When we look closely at the graphs, we can see that there are extraordinarily high and low returns to the stock market index which are not displayed after returns are distributed (especially when we look at the low returns in October 1987). We also see that returns are homogeneously scattered around zero. The serial autocorrelation for the gap omitted returns are practically equal to zero, which supports the martingale hypothesis. In particular, the expectation of returns conditioned on the past is equal to zero.

Figures 15 through 17 reveal the same type of analysis but on the more recent time period of January 3, 2000 through May 31, 2001. The same properties described above are even more evident with less data points. Again, we see more extreme returns when Monday returns are unchanged. Omitting these returns from the data will result in an underestimation of risk, especially that related to negative returns, which is the risk of most concern to investors.

Figure 10 – Squared Daily Returns (7/62-12/87)

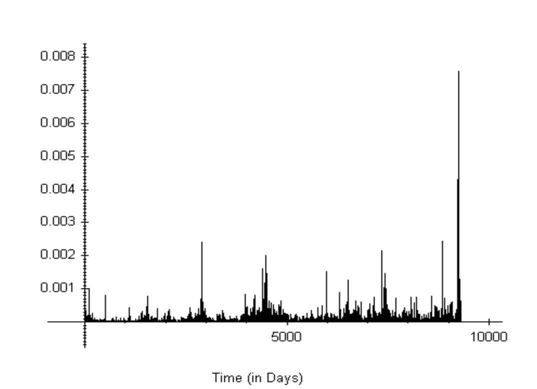

Figure 11 –Squared Daily Returns (1/00-5/01)

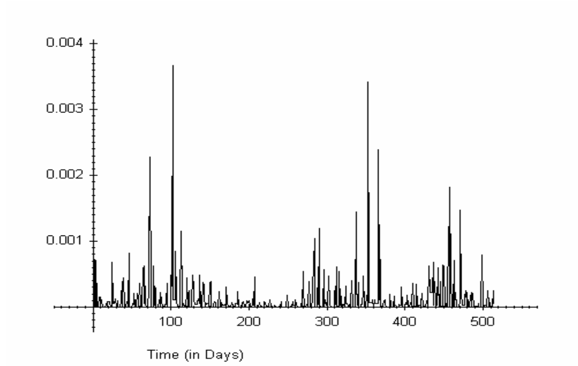

Figure 12 - Returns with Mondays Omitted (7/2/62 – 12/31/87)

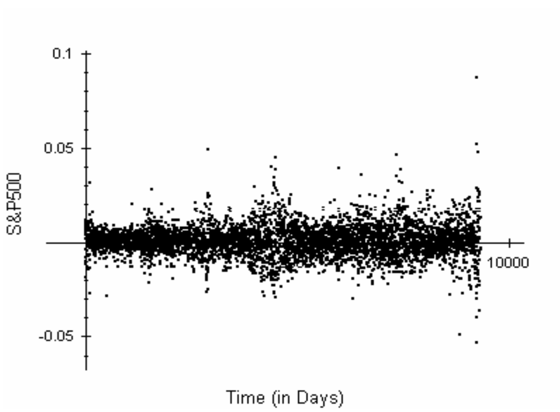

Figure 13 - Returns with Mondays Distributed (7/2/62 – 12/31/87)

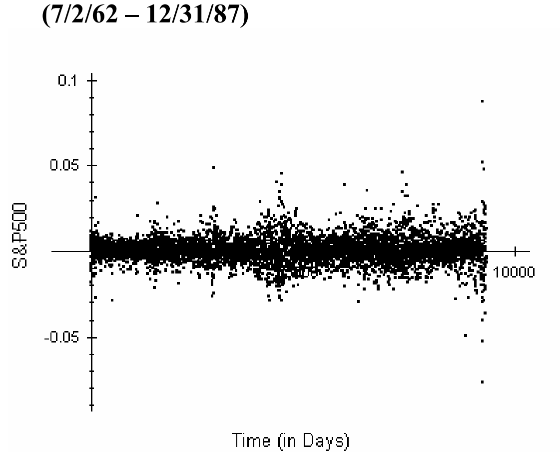

Figure 14 - Returns with Mondays Unchanged (7/2/62 – 12/31/87)

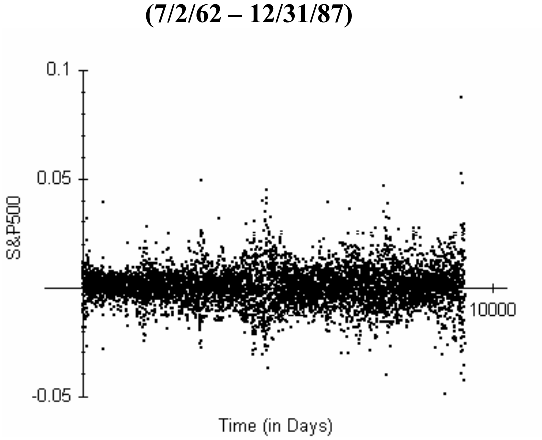

Figure 15 - Return with Mondays Omitted (1/3/00 – 5/31/01)

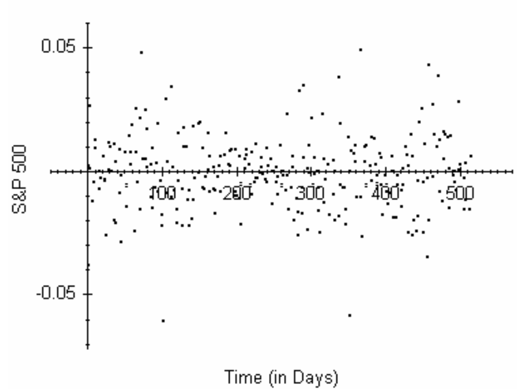

Figure 16 - Returns with Mondays Distributed (1/3/00 – 5/31/01)

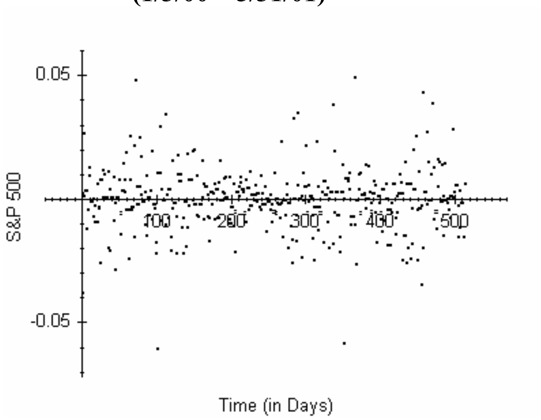

Figure 17 - Return with Mondays Unchanged (1/3/00 – 5/31/01)

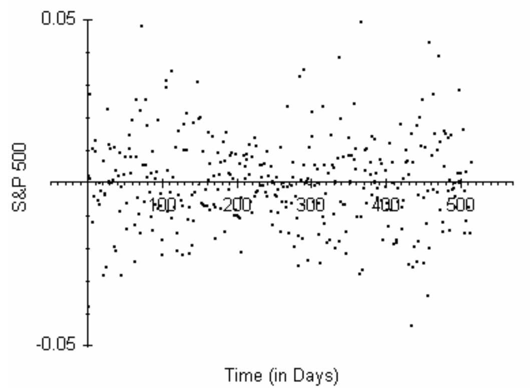

B. Extreme Value Analysis

1. Non-parametric Techniques

Non-parametric techniques include visual methods, as well as estimators and goodness of fit tests. The visual tools for examining data with extreme values include sample distribution functions, quantile functions, histograms, and kernel densities. Quantile functions and kernel densities are the most useful. Sample distribution functions are utilized for smaller sample sizes and histograms are for data in group form. Sample quantile functions are the diagonals of sample distribution functions. Kernel densities depend upon the bandwidth b and the number of the kernel j.

The full data set from July 2, 1962 through December 31, 1987 is represented by a kernel density in Figure 18. The automatic bandwidth selection by means of cross-validation leads to the choice of $b = 0.0031167$. The kernel density selected was the Epanechnikov kernel $\{0.75(1-z^2)\}$, which is the optimal kernel under a certain condition. We do not want to choose a bandwidth too large, or over smoothing of the data will occur. In Figure 18, the peaked curve has the $b = 0.0031167$ and the smoother curve underneath has a $b = 0.01$. We can see that the curve is far from the shape of a normal curve. One of the benefits of visualizing data with kernel densities is that they help us in identifying subpopulations in data. In this case, we do not see any evidence of a subpopulation. We can also check to see if the densities are bimodal or platykurtic (thin tails).

We also provide an analysis of the lower and upper extremes (the tails). Figure 20 shows the kernel densities for the upper tail. The peaked curve represents the auto selection bandwidth of 0.0012142. The smoother curve has a $b = 0.01$. Figure 21 reveals the kernel densities for the lower tail. The peaked curve represents the auto selection bandwidth of 0.0013582, while the smoother curve has the larger $b = 0.01$.

Figures 22 through 23 represent the kernel densities from the more recent period of January 3, 2000 through May 31, 2001. Figure 23 represents the entire data set. The automatic bandwidth selection provided a $b = 0.0052645$. We were consistent with the choice of the larger $b = 0.01$. Comparing these densities with those of Figure 18, we can still see a peakedness in the curve. However, there is an indication of a sub-population in the lower tail. Again, the larger the b, the smoother the curve.

When we look at the upper tail in Figure 22 and the lower tail in Figure 23, we see greater evidence of skewness and several sub-populations in the tails. The curve is not as smooth. This is also evidenced in the kernel density for the lower tail. The automatic bandwidths in these cases were 0.0023400 and 0.0026030 for the upper and lower tails respectively. These are compared to the larger $b = 0.01$.

In Figure 24, we show the sample quantile function for data covering the period June 2, 1962 through December 31, 1987. We can see the similarity to the Gumbel quantile function in Figure 4. Both the Fréchet and the Weibull distributions can be deduced from the Gumbel by means of transformation:

$T(x) = \log(x)$ for the Fréchet, and

$T(x) = -1/x$ for the Weibull.

With the Fréchet, the left end point is always zero and all potential measurements are positive.

In Figure 25, we can visualize the sample quantile function for the upper tail of the distribution. It closely resembles a Fréchet quantile function was given in Figure 5. In Figure 26, we look at the sample quantile function for the lower tail. It also resembles a Fréchet quantile function.

Figures 27 through 28 reveal the sample quantile functions for the more recent data. We see evidence of exactly the same shapes of the curves as found with the older data. The curves are just not as smooth.

Figure 18 –Kernel Densities (7/2/62 -12/31/87)

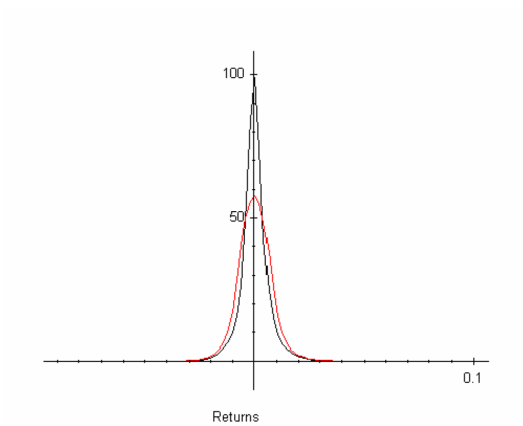

Figure 19 –Kernel Densities for the Upper Tail (7/2/62 -12/31/87)

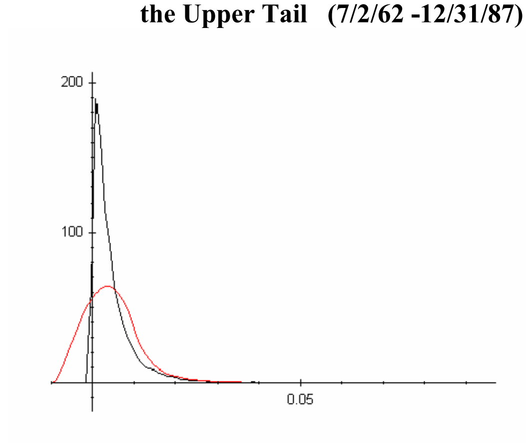

Figure 20 –Kernel Densities for the Lower Tail (7/2/62 -12/31/87)

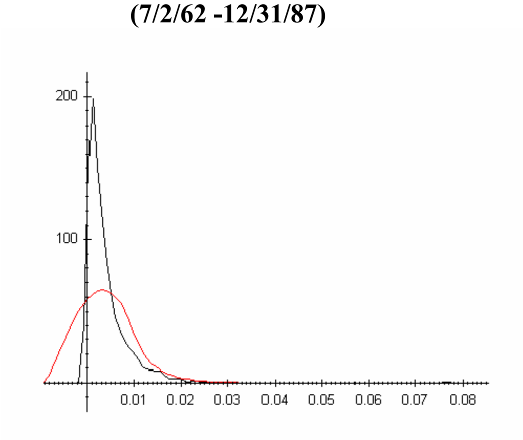

Figure 21 –Kernel Densities (1/31/00 -5/31/01)

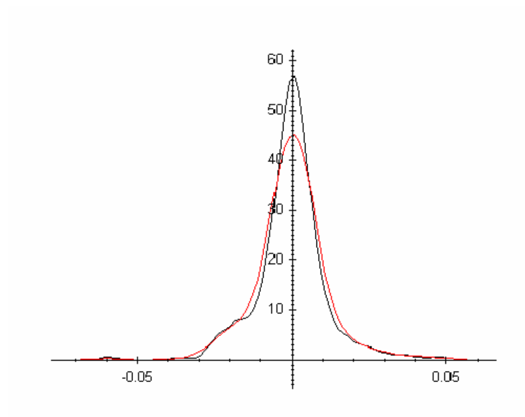

Figure 22 –Kernel Densities for the Upper Tail (1/31/00 -5/31/01)

Figure 23 –Kernel Densities for the Lower Tail (1/31/00 -5/31/01)

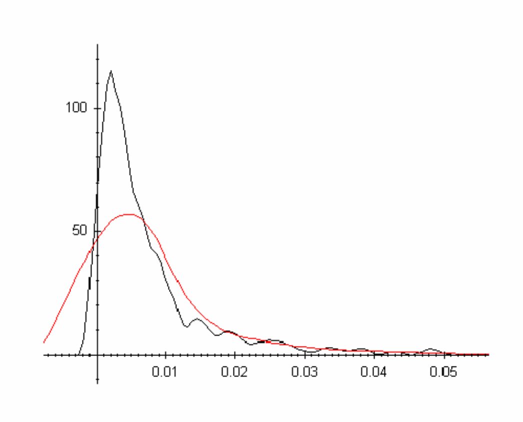

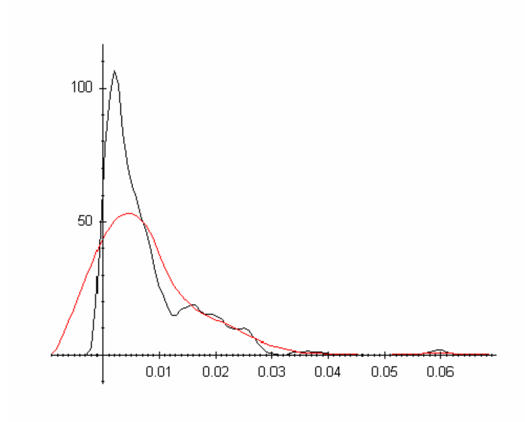

Figure 24 – Sample Quantile Function (7/2/62 -12/31/87)

Figure 25 – Sample Quantile Function for the Upper Tail (7/2/62 -12/31/87)

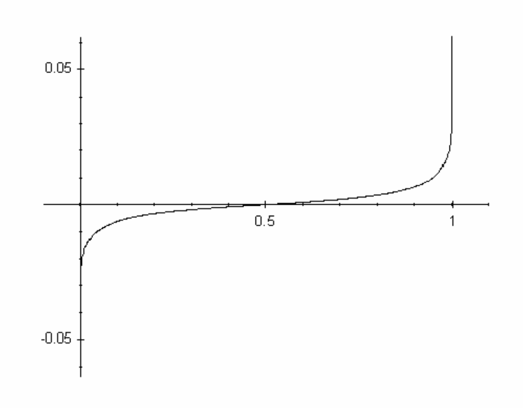

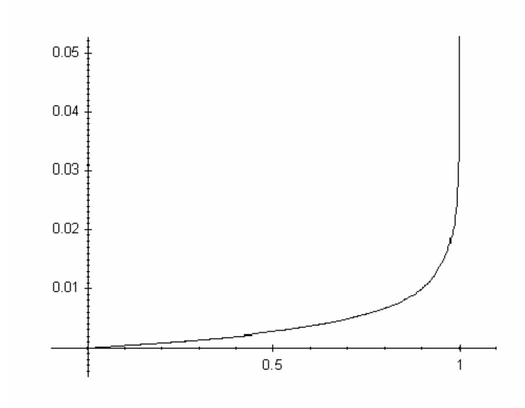

Figure 26 – Sample Quantile Function for the Lower Tail (7/2/62 -12/31/87)

Figure 27 – Sample Quantile Function (1/3/00 - 5/31/01)

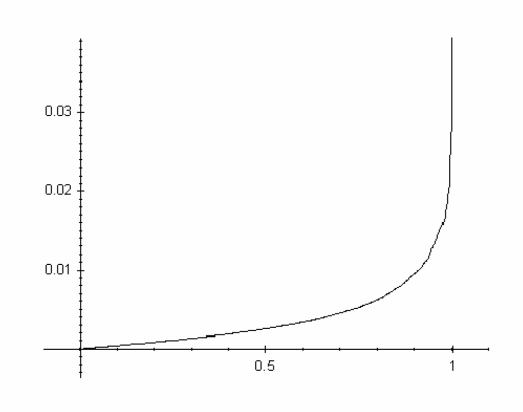

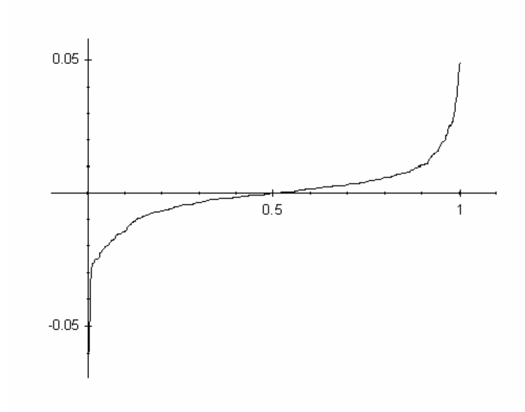

Figure 28 – Sample Quantile Function for the UpperTail (1/3/00 - 5/31/01)

Figure 29 – Sample Quantile Function for the Lower Tail (1/3/00 - 5/31/01)

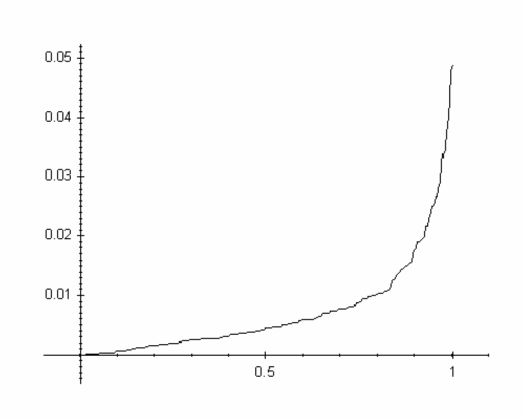

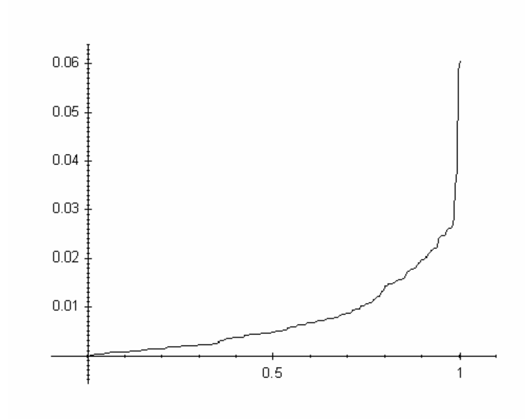

Estimators utilized include the maximum likelihood estimators (MLE), where the μ_n and σ_n of the location and scale parameters must be evaluated numerically, and the moment estimators, where the estimators of μ_n and σ_n are deduced from the sample mean and variance. In addition, we also used the Hill (1975) estimator. Table 4 shows the results of the estimates of the tail indices of stock market returns.

Table 4. - Estimates of Tail Indices

(7/2/62-12/31/87)		(1/3/00 - 5/31/01)	
Entire Data Set:		**Entire Data Set:**	
Hill	2.02024	Hill	3.52981
Moment	6.82101	Moment	3.31285
MLE	4.48998	MLE	3.22537
Upper Tail:		**Upper Tail:**	
Hill	2.44652	Hill	2.45084
Moment	9.3395	Moment	-12.4556
MLE	7.31303	MLE	-3.23043
Lower Tail:		**Lower Tail:**	
Hill	2.74657	Hill	1.63348
Moment	5.3695	Moment	18.4415
MLE	4.66684	MLE	-7.01677

Loretan and Phillips (1994) and Reiss and Thomas (1997) report that one must take heavy upper and lower tails with a tail index around 3.00 into account. Loretan and Phillips look at daily returns for the S&P 500 from July 1962 through December 1987 and monthly returns for a broad index of stocks from January 1834 to December 1987. They find that the point estimates are almost all less than 4.00; they range from 2.5 to 3.2 for the monthly stock market return series and from 3.1 to 3.8 for the daily stock market return series.

Table 5 provides the sample functional parameters for all of the data sets examined. We see evidence of higher skewness when only looking at the tails of the distribution. Comparing the two time periods for the entire data set, the earlier time period has a lower minimum and a higher maximum. The more recent period has a higher variance and a lower (negative) mean.

Comparing the upper tails, the earlier time period has a lower minimum and a lower maximum. The more recent period has a higher mean and a higher variance. Looking at the lower tails, we found a lower minimum and a higher maximum in the earlier time period. We found a higher mean and a higher variance in the lower tails for the more recent period. Thus, in all cases the variance was higher in the shorter, and more recent, time horizon. In addition, the minimums were all lower for the earlier time period.

Table 5. Sample Functional Parameters

(7/2/62-12/31/87)		**(1/3/00 - 5/31/01)**	
Entire Data Set:		**Entire Data Set:**	
Sample Size	9313	Sample Size	514
Minimum	-0.076332	Minimum	-0.060627
Maximum	0.0870888	Maximum	0.04884
Median	2.459E-16	Median	8.36E-05
Mean	0.0015966	Mean	-0.000285
Mean Deviation	0.0064557	Mean Deviation	-0.001392
Variance	4.17E-05	Variance	0.0001298
Skewness	0.0815873	Skewness	-0.091055
Upper Tail:		**Upper Tail:**	
Sample Size	4678	Sample Size	260
Minimum	9.107E-18	Minimum	3.322E-16
Maximum	0.0870888	Maximum	0.048884
Median	0.0027427	Median	0.004465
Mean	0.0043078	Mean	0.0071678
Mean Deviation	0.0050946	Mean Deviation	0.0084004
Variance	2.6E-05	Variance	7.06E-05
Skewness	3.49313	Skewness	2.3698
Lower Tail:		**Lower Tail:**	

Sample Size	4635	Sample Size	254
Minimum	2.168E-18	Minimum	9.498E-17
Maximum	0.0763324	Maximum	0.0606272
Median	0.0025229	Median	0.0048973
Mean	0.0040269	Mean	0.0079136
Mean Deviation	0.004759	Mean Deviation	0.0086786
Variance	2.26E-05	Variance	0.0001
Skewness	4.28911	Skewness	2.41929

2. Parametric Techniques

We can also visualize data by parametric distribution functions or densities. Thereby, one can visually control the validity of a parametric model. One should choose the parametric distribution by minimizing the maximization from the curves. We look at the Q-Q plots to see if there is a deviation from a straight line. If there is a strong deviation, then either the shape parameter is inaccurate or the model selection is wrong.

In this part of the analysis, we only looked at the tails of the distributions. Figures 30 through 4 show the Q-Q plots for the data from the period July 1962 through December 1987. Figure 30 presents the Hill estimator, while Figure 31 shows the MLE estimator and Figure 32 shows the moment estimator. All of these are for the upper tails. All three estimators provide a fairly linear Q-Q plot. Figures 34 through 36 examine the same estimators for the lower tails. Again, we see fairly linear Q-Q plots and a strong similarity between the MLE and the moment estimators.

Figure 36 through 41 provide the Q-Q plots for the more recent data. The lines are not as linear as we saw earlier with the longer time horizon. It is interesting to note that the curves in Figures 37 and 38 are in a negative quadrant when using both the MLE and the moment estimators. In every case of using the Hill estimator, was equal to zero. However, in this case, when looking at the upper tail of the more recent data, the MLE and moment estimators provided and that was far from zero. In addition, the tail index estimates were negative. We can still see evidence of a linear trend though. From these graphs, we can conclude that extreme value theory provides a good estimate for the tails of stock price returns.

Figure 30 – Q-Q Plot of Upper Tail Using Hill Estimator (7/62 – 12/87)

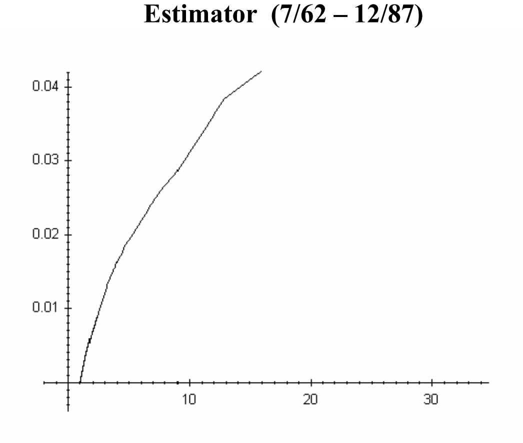

Figure 31 – Q-Q Plot of Upper Tail Using MLE Estimator (7/62 – 12/87)

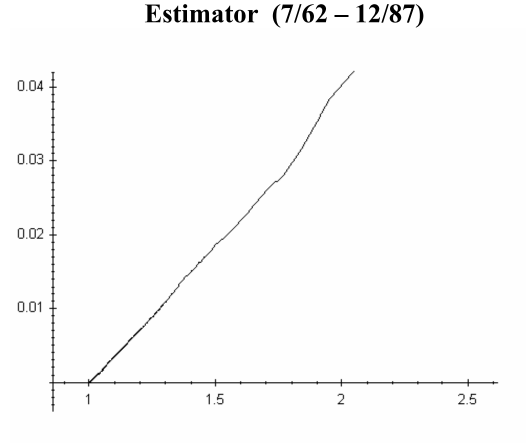

Figure 32 – Q-Q Plot of Upper Tail Using Movement Estimator (7/62 – 12/87)

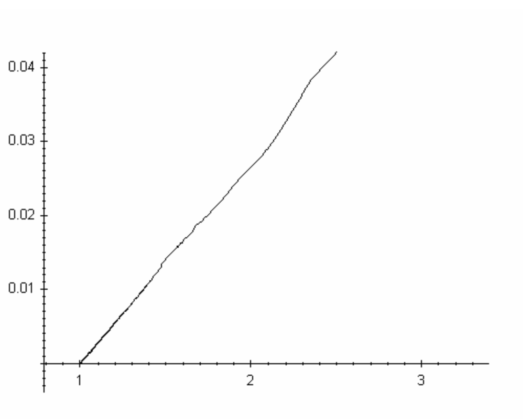

Figure 33 – Q-Q Plot of Lower Tail Using Hill Estimator (7/62 – 12/87)

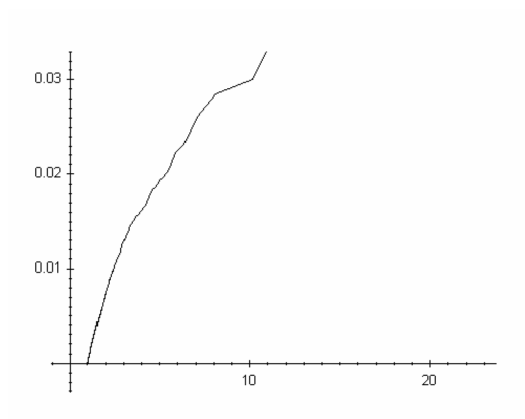

Figure 34 – Q-Q Plot of Lower Tail Using MLE Estimator (7/62 – 12/87)

Figure 35 – Q-Q Plot of Lower Tail Using Movement Estimator (7/62 – 12/87)

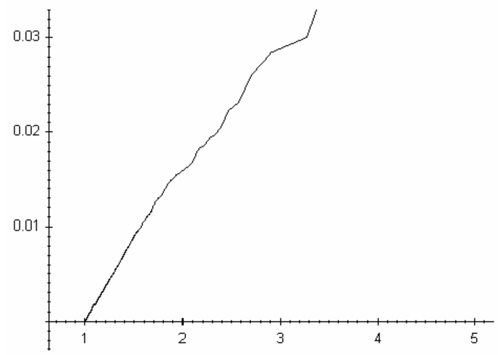

Figure 36 – Q-Q Plot of Upper Tail Using Hill Estimator (1/00 – 5/01)

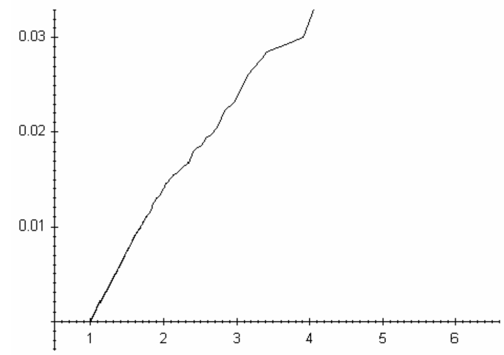

Figure 37 – Q-Q Plot of Upper Tail Using MLE Estimator (7/62 – 12/87)

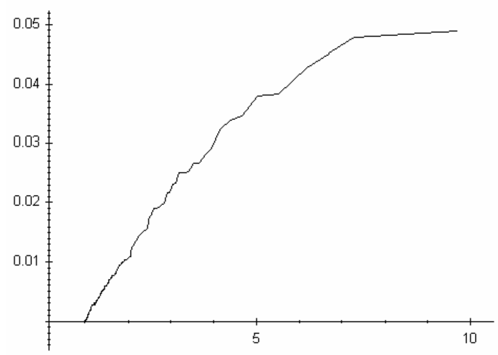

Figure 38 – Q-Q Plot of Upper Tail Using Movement Estimator (1/00 – 5/01)

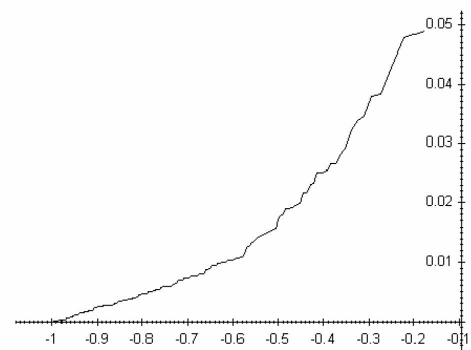

Figure 39 – Q-Q Plot of Lower Tail Using Hill Estimator (1/00 – 5/01)

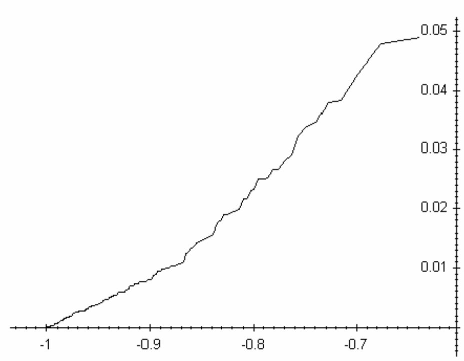

Figure 40 – Q-Q Plot of Lower Tail Using MLE Estimator (1/00 – 5/01)

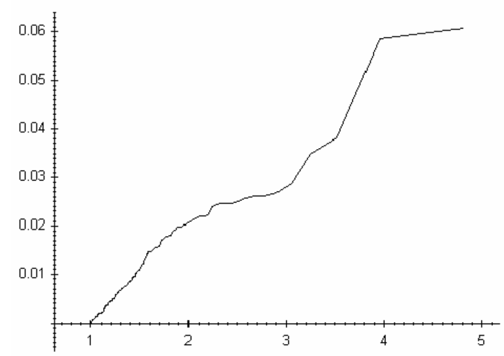

Figure 41 – Q-Q Plot of Lower Tail Using Movement Estimator (1/00 – 5/01)

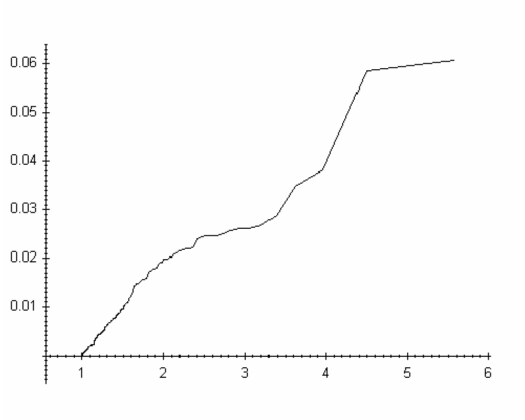

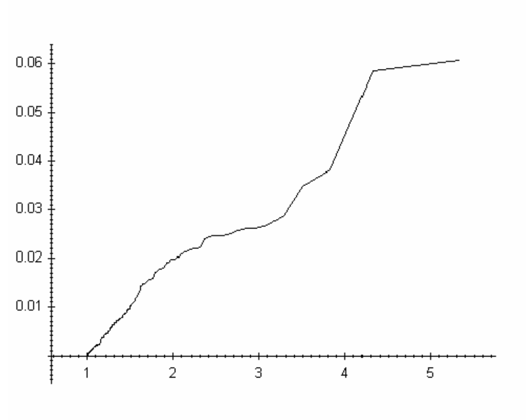

VII. SUMMARY AND RECOMMENDATIONS

Our primary aim of this paper was an attempt to fit extreme value distributions to daily stock price returns. Specifically, we tested the null hypothesis of $H_0 : \alpha < 2$ versus the alternative hypothesis of $H_1 : \alpha \geq 2$. We were able to utilize extreme value analysis and we found that in the majority of the cases studied the α was greater than 2.0. Both the Student's *t* and the ARCH process allow for $\alpha \geq 2$. The leptokurtic stable hypothesis requires α <2.

We know now that the greatest bull market in history is over. For the two most recent decades, we saw a generally positive, non-linear trend in stock market prices. Now, in the new millennium, we are evidencing a downward trend. Will this trend continue? Will the probabilities for extreme returns continue to change? Due to the random nature of stock prices returns, these are questions that can only be answered in the future. However, given the tools provided here, one can study the clustering of variance to make a formulated opinion regarding the future.

This study finds that extreme value analysis is a valuable tool for examining stock price movements and can be more efficient than the usual variance in measuring risk. We find that the investment horizon and the method utilized to estimate the tails play an important role in the determination of the tail index. A shorter time horizon, especially one with high volatility, shows more mass in the tails. Therefore, careful consideration must be given to the time period studied as well as the assumptions of the model used.

Investors trade off return and risk. Portfolio managers need to take the predictability of the risk of an investment into consideration when constructing and monitoring their portfolios. A recommendation for future research would be to utilize the entire data set of daily stock price returns from 1962 through the present. Extreme value analysis provides a better fit when there are more observations. In addition, one could adjust the time series for the linear growth in the economy. In sum, extreme value theory is a valuable tool to analyze stock price returns and it is recommended that a portfolio manager utilize the theory to its fullest extent.

References

Daníelsson, Jón, and Casper G. de Vries. (1997) "Tail Index and Quantile Estimation with Very High Frequency Data. *Journal of Empirical Finance*, 4, 241-257.

Glassman, James K., and Kevin A. Hassett. (1999) "Dow 36,000." *Atlantic Monthly*, September, 37-58.

_______. (1999) Dow 36,000. Times Business, New York.

Hols, Martien C.A.B., and Capser G. de Vries. (1991) "The Limiting Distribution of Extremal Exchange Rate Returns." *Journal of Applied Econometrics*, 6, 287-302.

Jansen, Dennis W., and Casper G. de Vries. (1991) "On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective." *The Review of Economics and Statistics*, 18-24.

Jones, Charles P., and Jack W. Wilson. (2000) "Stock Returns in the 1990s: Implications for the Future." *The Journal of Private Portfolio Management,* Spring, 51-60.

Jorion, Philippe. (1997) *Value at Risk: The New Benchmark for Controlling Market Risk*. Irwin Professional Publishing, Chicago.

Klüppelberg, C., and T. Mikosch. (1997) "Large Deviations of Heavy-Tailed Random Sums with Applications to Insurance and Finance." *J. Appl. Probab*., 34, 293-308.

Longin, François. (1996) "The Asymptotic Distribution of Extreme Stock Market Returns." *Journal of Business,* 63, 383-408.

Loretan, Mico, and Peter C.B. Phillips. (1994) "Testing the Covariance Stationarity of Heavy-Tailed Time Series: An Overview of the Theory with Applications to Several Financial Datasets. *Journal of Empirical Finance*, 1(2), 211-248.

Lucas, André, and Pieter Klaassen. (1998) "Extreme Returns, Downside Risk, and Optimal Asset Allocation." *The Journal of Portfolio Management*, Fall, 71-79.

Mandlebrot, Benoit. (1963a) "New Methods in Statistical Economics". *The Journal of Political Economy,* October, 421-440.

_______. (1963b) "The Variation of Certain Speculative Prices". *Journal of Business*, 36, 394-419.

McCulloch, J.H. (1978) "Continuous Time Processes with Stable Increment." *Journal of Business*, 51, 601-619.

_______. (1981) "Interest Rate Risk and Capital Adequacy for Traditional Banks and Financial Intermediaries." In S.J. Maisel (ed.), *Risk and Capital Adequacy in Commercial Banks*. University of Chicago Press, Chicago, 223-248

Mikosch, T., and A.V. Nagaez. (1998) "Large Deviations of Heavy-Tailed Sums with Applications in Insurance." *Extremes,* 1(1), 81-110.

Müller, Ulrich A., Michael M. Dacorogna, and Olivier V. Pictet. (1998) "Heavy Tails in High-Frequency Financial Data." *In A Practical Guide to Heavy Tails*: *Statistical Techniques and Applications*, Robert J. Adler, Raisa E. Feldman, and Murad S. Taqqu (eds.), Birkhäuser, Boston, 55-77.

Parkinson, Michael. (1977) "Option Pricing: The American Put." *Journal of Business*, 50, 21-36.

_______. (1980) "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business,* 53(1), 61-65.

Phoa, Wesley. (1999) "Estimating Credit Spread Risk Using Extreme Value Theory." *The Journal of Portfolio Management*, Spring, 69-73.

Rothschild, M., and Stiglitz, J.E. (1970) "Increasing Risk: I. A Definition." *Journal of Economic Theory*, 2, 225-243.

Roy, A.D. (1952) "Safety First and the Holding of Assets." *Econometrica,* 431-449

Vries, Casper G. de. (1994) "Stylized Facts of Nominal Exchange Rate Returns," in *The Handbook of International Macroeconomics*, ed. by Frederick Van Der Ploeg, Blackwell Publishing, Cambridge, MA, 348-389.