. . . About Higher Moments

Mark P. Kritzman

In financial analysis, a return distribution is commonly described by its expected return and standard deviation. For example, the S&P 500 Index might have an expected return of 10 percent and a standard deviation of 15 percent. By assuming that the returns of the S&P 500 Index conform to a particular distribution, such as a normal distribution, we can infer the entire distribution of returns from the expected return and standard deviation.[1]

The expected value of a distribution is referred to as the first moment of the distribution and is measured by the arithmetic mean of the returns. The variance, which equals the standard deviation squared, is called the second central moment or the second moment about the mean. It measures the dispersion of the observations around the mean. The first central moment is not the mean itself, but rather zero, because central moments are measured relative to the mean.

The normal distribution is symmetric around the mean; hence, the median (the middle value of the distribution) and the mode (the most common value of the distribution) are both equal to the mean. Moreover, the normal distribution has a standard degree of peakedness. These properties of the normal distribution explain why just the mean and standard deviation are sufficient to estimate the entire distribution.

Although investment returns usually are assumed to be approximately normally distributed, this assumption is less likely to hold for very short horizons, such as one day, and for long horizons, such as several years. Moreover, certain assets and investment strategies have properties that produce nonnormal distributions over any horizon. Thus, in some cases, to estimate a return distribution, one must go beyond the first moment and the second central moment to the third central moment, which is called skewness, or the fourth central moment, which is called kurtosis.

Mark P. Kritzman, CFA, is a partner of Windham Capital Management.

Skewness

Skewness, which is illustrated in Figure 1, refers to the asymmetry of a distribution. A distribution that is positively skewed has a long tail on the right side of the distribution and its mean is typically greater than its median, which in turn, is greater than its mode. Because the mean exceeds the median, most of the returns are below the mean, but they are of smaller magnitude than the few returns that are above the mean.

In contrast, a distribution that is negatively skewed, which is shown in Figure 2, has a long tail on the left side of the distribution, indicating that the few outcomes that are below the mean are of greater magnitude than the larger number of outcomes above the mean. Hence, the mean is typically lower than the median, which is lower than the mode.

Skewness is computed as the average of the cubed deviations from the mean and is usually measured by the ratio of this value to the standard deviation cubed; that is,

$$S = \frac{(1/n)\sum_{i=1}^{n} (R_i - \mu)^3}{\sigma^3}, \quad (1)$$

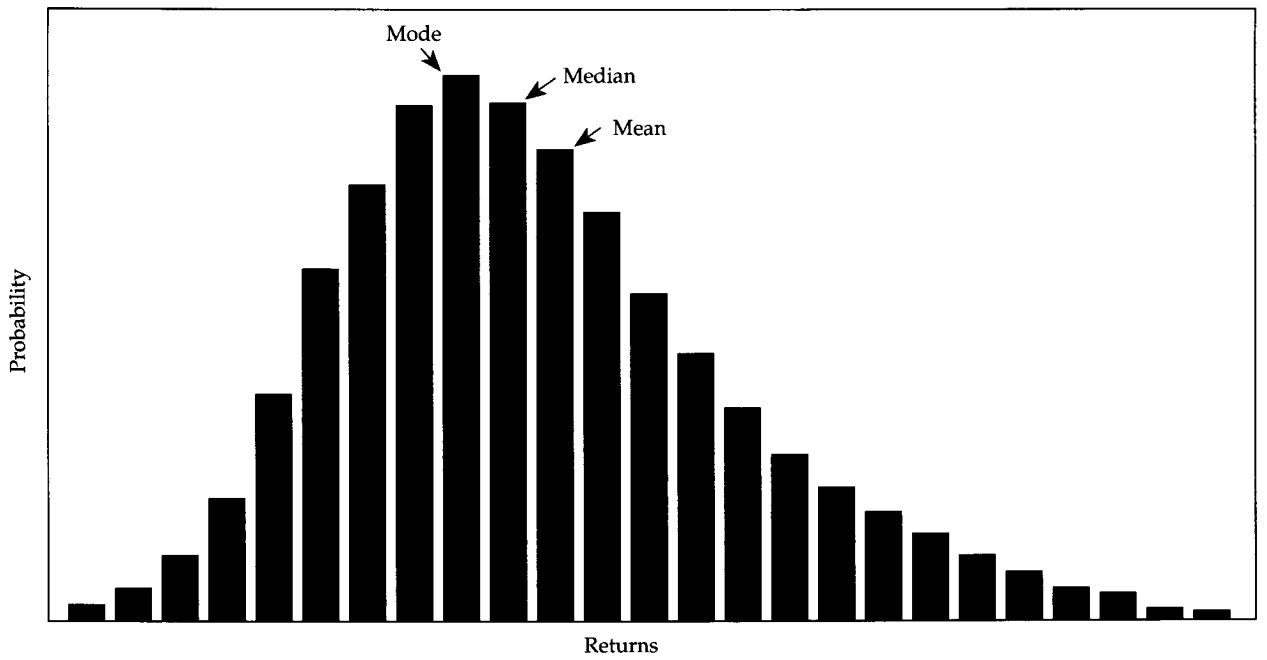
where

S = measure of skewness
n = number of returns
R_i = ith return
μ = arithmetic mean of returns
σ = standard deviation of returns

The distribution of long-horizon returns that arise from compounding independent shorter horizon returns is typically skewed to the right. Suppose, for example, that we select 100 annual returns from an underlying normal distribution of returns that has a mean of 10 percent and a standard deviation of 15 percent, and suppose that we repeat this selection ten times. Hence, we generate ten samples, each consisting of 100 returns. Now, suppose we compound the returns

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

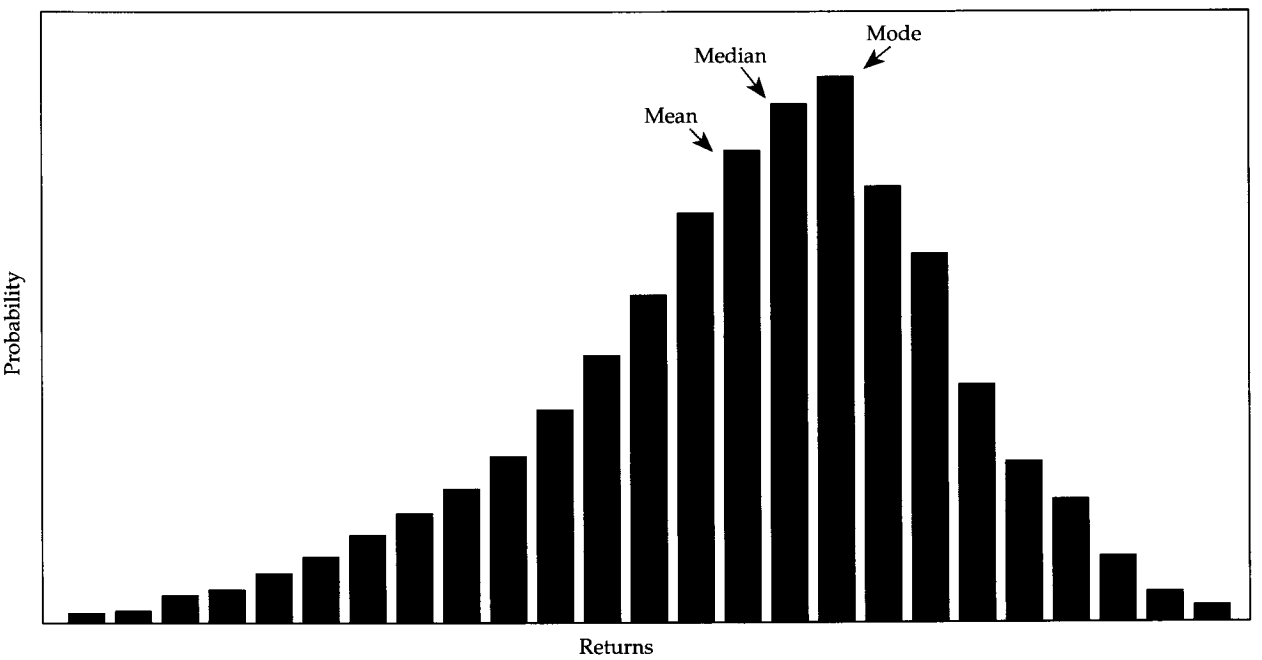
Figure 1. Positively Skewed Distribution



from the ten samples so that we end up with a distribution of 100 cumulative ten-year returns. Table 1 shows the results of such an experiment. The first column shows the average values of the ten samples of annual returns, and the second column shows the values associated with the cumulative ten-year returns.

Given that the annual returns were drawn

Figure 2. Negatively Skewed Distribution



Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Table 1. The Effect of Compounding on Skewness

	Average of Ten Samples of Normally Distributed Annual Returns	Cumulative Ten-Year Returns
Mean	0.1066	1.7580
Median	0.1105	1.3211
Standard deviation	0.1496	1.3073
Skewness	−0.1163	1.3634

from a normal distribution, the average mean and the average median are not very far apart and the average skewness of the distribution is close to zero, although slightly negative. The distribution of the cumulative ten-year returns is significantly positively skewed, however. Moreover, the mean of the cumulative ten-year returns is significantly greater than the median, despite the fact that the average median of the annual samples exceeds their average mean.

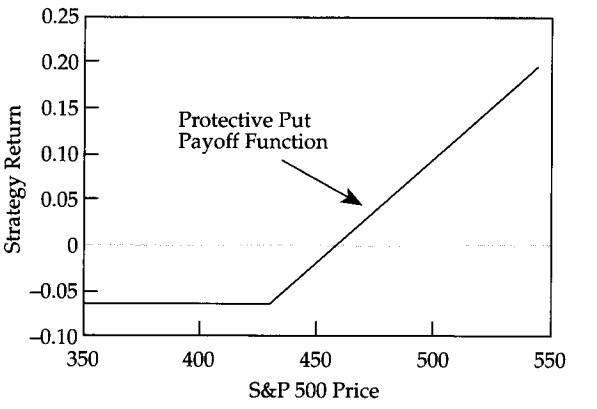
The process of compounding introduces skewness because compounded favorable returns have a greater impact than compounded unfavorable returns of equal magnitude. For example, two consecutive 10 percent returns increase an asset's value by 21 percent, whereas two consecutive −10 percent returns decrease an asset's value by only 19 percent.

The skewness that arises from compounding independent returns causes the returns to be log-normally distributed, which implies that the logarithms of the quantities 1 plus the compounded returns are normally distributed.[2] This relationship is convenient because, by transforming returns into the logarithms of their wealth relatives, we can ignore the skewness of the distribution of the underlying returns and use the normal distribution to estimate the probability of experiencing various logarithmic results. These results, of course, must be transformed back into their original units.[3]

Now, consider another process that results in a skewed return distribution. Suppose we purchase a put option to protect an investment in an S&P 500 Index fund that is currently valued at $450.00. Furthermore, suppose that the strike price of the option is $425.00, that it expires in one year, and that it costs $10.00. Figure 3 shows the returns of such a strategy contingent upon various values for the S&P Index.

If the S&P 500 Index has a 10 percent expected return (including the reinvestment of dividends) and a 15 percent standard deviation, and if its annual returns are approximately normally distributed, we can assign probabilities that the Index will equal or fall below various values. Moreover, because we can map these values precisely onto returns for the protective put strategy, we can generate the Index's probability distribution. Table 2 demonstrates this transformation.

Figure 3. Contingent Returns of Protective Put Strategy



Column 1 of Table 2 shows possible prices for the S&P 500 Index at the end of the one-year horizon. Column 2 shows the corresponding probability of falling below the prices in Column 1. Column 3 shows the values of a protective put strategy contingent on the S&P prices in Column 1. These values are based on an initial value of $460.00 for the protective put strategy—a $450.00 investment in the S&P Index and a $10.00 investment in the put option. The protective put strategy returns shown in Column 4 are derived by dividing the contingent values of the protective put strategy by 460 and subtracting 1. Column 5 shows the frequency distribution of the protective put strategy returns. For example, there is a 16.79 percent chance that the strategy will return precisely −6.52 percent because this return would obtain for any S&P value equal to or less than $430.00. Also, because the probability of achieving a return between −6.52 percent and −4.35 percent equals 3.97 percent, the chance of experiencing a return below −4.35 percent is 20.76 percent (16.79 + 3.97).

Figure 4 plots these probabilities as a function of the protective put strategy returns. It reveals that overlaying a protective put option on an asset with a symmetric distribution truncates the left tail

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Table 2. Probability Distribution of Put Strategy Returns

(1) S&P Price	(2) Probability of Falling Below S&P Price	(3) Value of PP Strategy	(4) PP Strategy Return	(5) Probability between Successive Returns
350	1.59%	430	−6.52%	0.00%
360	2.28	430	−6.52	0.00
370	3.20	430	−6.52	0.00
380	4.42	430	−6.52	0.00
390	5.99	430	−6.52	0.00
400	7.97	430	−6.52	0.00
410	10.40	430	−6.52	0.00
420	13.32	430	−6.52	0.00
430	16.79	430	−6.52	16.79
440	20.76	440	−4.35	3.97
450	25.25	450	−2.17	4.49
460	30.20	460	0.00	4.95
470	35.54	470	2.17	5.34
480	41.22	480	4.35	5.68
490	47.05	490	6.52	5.83
500	52.95	500	8.70	5.90
510	58.78	510	10.87	5.83
520	64.46	520	13.04	5.68
530	69.80	530	15.22	5.34
540	74.75	540	17.39	4.95
550	79.24	550	19.57	4.49
560	83.21	560	21.74	3.97
570	86.68	570	23.91	3.47
580	89.60	580	26.09	2.92
590	92.03	590	28.26	2.43
600	94.01	600	30.43	1.98
610	95.58	610	32.61	1.57
620	96.80	620	34.78	1.22
630	97.72	630	36.96	0.92
640	98.41	640	39.13	0.69

of the distribution and thereby imparts positive skewness. Most option strategies result in a skewed return distribution. Writing covered calls, for example, prevents the writer of the options from participating in profits generated from increases above the exercise price. Thus, the right side of the distribution is truncated at the exercise price, which results in negative skewness.

Many assets include embedded options. For example, convertible bonds give the owner a call option on the stock of the firm. Similarly, callable bonds grant the issuer of the bonds a call option on its debt. Clearly, these assets have skewed returns. Diversified portfolios that include assets with embedded options may also exhibit significant skewness.

Skewness also arises from dynamic trading strategies such as portfolio insurance. This strategy, which reduces exposure to a risky asset as the asset's price falls and increases exposure to a risky asset as it rises, produces a positively skewed return distribution.

Kurtosis

The fourth central moment is called kurtosis, which refers to the peakedness of a distribution. It equals the average of the deviations from the mean raised to the fourth power and is measured as a ratio of this quantity to the standard deviation raised to the fourth power. In equation form,

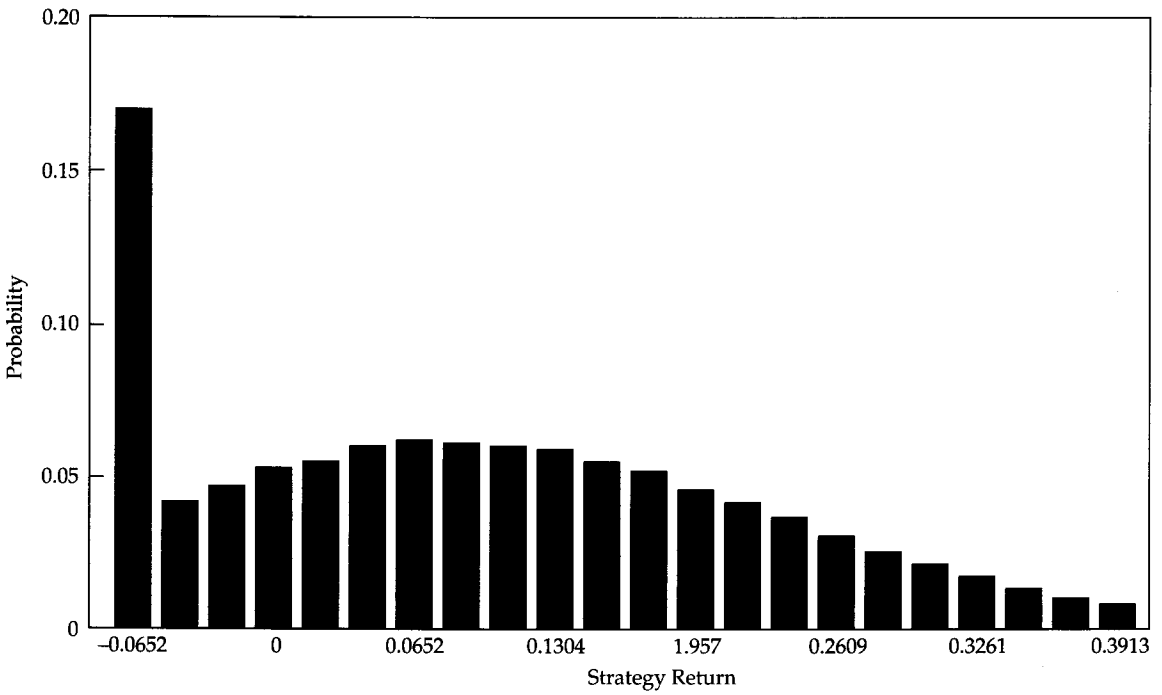
$$K = \frac{(1/n)\sum_{i=1}^{n}(R_i - \mu)^4}{\sigma^4}, \quad (2)$$

where

K = measure of kurtosis
n = number of returns
R_i = ith return
μ = arithmetic mean of returns
σ = standard deviation of returns

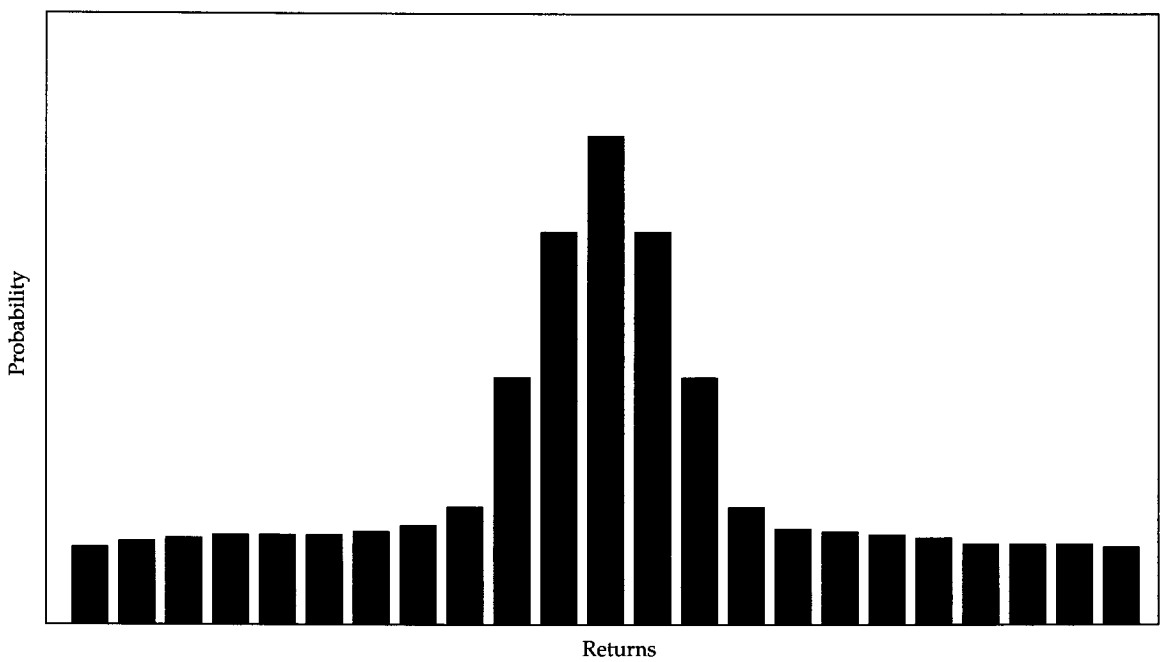
Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Figure 4. Probability Distribution of Protective Put Strategy



A normal distribution has a kurtosis value equal to 3. A distribution that has wide tails and a tall narrow peak is called leptokurtic; its kurtosis will exceed 3. Compared with a normal distribution, a larger fraction of the returns are at the extremes rather than slightly above or below the mean of the distribution. This distribution is shown in Figure 5.

Figure 5. Leptokurtic Distribution



Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

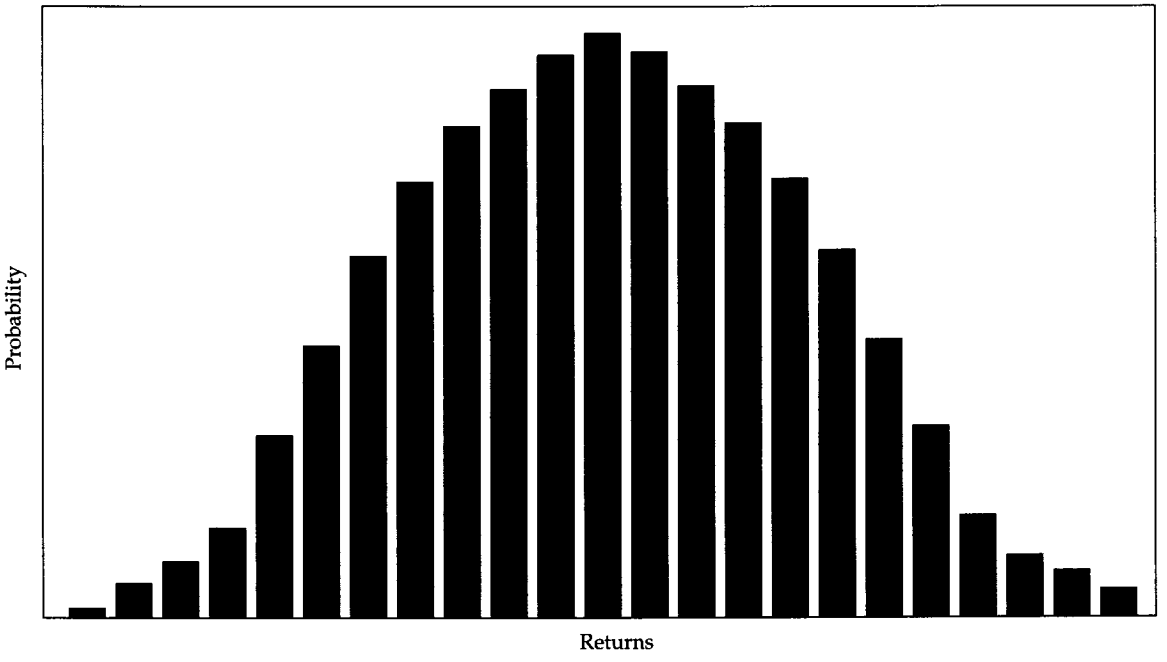
A distribution that has thin tails and a relatively flat middle is called platykurtic. Its kurtosis will be less than 3. Relative to a normal distribution, a larger fraction of the returns are clustered around the mean, as shown in Figure 6.

Return series that are characterized by jumps as opposed to more continuous changes will tend to be leptokurtic. Consider an example in which a country's economic policymakers decide to manage their exchange rate relative to the currency of an important trading partner. Specifically, they control it such that it will not increase or decrease more than 1 percent in any quarter. Suppose, however, that a black market in the currency emerges among traders; hence, the policymakers can observe the true market prices. Although the policymakers believe they can control the short-term volatility of the currency, they recognize that periodically they must reset the currency's value to the market exchange rate. Therefore, they revalue the currency every two years to accord with the black market exchange rate. Although this example may seem conveniently contrived, it is not qualitatively different from many actual exchange rate systems and other regulatory mechanisms such as circuit breakers on commodities and securities exchanges.

Table 3 compares the black market returns with the regulated returns. The market exchange rates are based on random quarterly returns drawn from a normal distribution with a mean of zero percent and a standard deviation of 5 percent. The regulated exchange rates are based on returns that equal the minimum of the positive market return (1 percent) or the maximum of the negative market return (−1 percent), except for every eighth quarter. The eighth-quarter returns are derived by resetting the regulated exchange rate to the black market exchange rate. As this experiment reveals, the regulatory process induces a high level of kurtosis to the returns series: 21.52 versus 3.25 for the market returns. (Remember that a normal distribution has a kurtosis measure of 3.) The intuition behind this result is that the regulatory process dampens moderately deviant returns, forcing them closer to the mean than they would have been otherwise; at the end of every second year, however, the revaluation process produces highly deviant returns, thereby fattening the tails of the distribution.

Daily returns also tend to be leptokurtic. Table 4 shows the kurtosis of daily and monthly currency returns for the 20-year period beginning January 1, 1974, and ending December 31, 1993.

Figure 6. Platykurtic Distribution



Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Table 3. Jump Process Introduced by Exchange Rate Management

Period	Market Returns	Market Exchange Rates	Regulated Returns	Regulated Exchange Rates
1	0.49%	1.0049	0.49%	1.0049
2	2.04	1.0254	1.00	1.0149
3	2.26	1.0485	1.00	1.0251
4	4.23	1.0928	1.00	1.0353
5	−3.95	1.0497	−1.00	1.0250
6	1.76	1.0682	1.00	1.0352
7	−5.11	1.0135	−1.00	1.0249
8	−3.38	0.9793	−4.45	0.9793
9	−3.67	0.9434	−1.00	0.9695
10	−0.07	0.9427	−0.07	0.9687
11	1.25	0.9545	1.00	0.9784
12	−3.12	0.9246	−1.00	0.9246
13	2.79	0.9504	1.00	0.9504
14	−3.71	0.9151	−1.00	0.9409
15	4.88	0.9597	1.00	0.9503
16	−2.81	0.9327	−1.85	0.9327
17	−1.74	0.9165	−1.00	0.9234
18	3.37	0.9474	1.00	0.9327
19	5.80	1.0023	1.00	0.9420
20	0.17	1.0040	0.17	0.9436
21	−1.80	0.9859	−1.00	0.9342
22	6.58	1.0507	1.00	0.9435
23	−1.83	1.0315	−1.00	0.9341
24	8.28	1.1170	19.58	1.1170
25	1.92	1.1384	1.00	1.1282
26	−4.54	1.0868	−1.00	1.1169
27	−9.33	0.9854	−1.00	1.1057
28	1.13	0.9965	1.00	1.1168
29	0.87	1.0052	0.87	1.1265
30	12.46	1.1304	1.00	1.1378
31	−4.76	1.0766	−1.00	1.1264
32	−4.38	1.0294	−8.61	1.0294
33	−7.57	0.9515	−1.00	1.0191
34	−5.68	0.8974	−1.00	1.0089
35	−0.57	0.8923	−0.57	1.0031
36	−0.92	0.8841	−0.92	0.9939
37	3.80	0.9177	1.00	1.0039
38	−1.12	0.9074	−1.00	0.9938
39	5.72	0.9593	1.00	1.0038
40	−0.14	0.9580	−4.56	0.9580
Kurtosis	3.25		21.52	

Although monthly currency returns are only slightly leptokurtic, daily currency returns are significantly so. This empirical tendency shows up in other returns as well, including stock returns. This result may arise from price jumps that occur in response to the accumulated information that is released during nontrading hours, especially over weekends. As the measurement interval increases, these price jumps cancel out, which explains why monthly returns typically are less leptokurtic than daily returns.

Table 4. Kurtosis of Monthly and Daily Currency Returns, January 1, 1974–December 31, 1993

Currency	Monthly	Daily
British pound	4.20	7.03
German mark	3.23	23.80
French franc	3.67	16.43
Swiss franc	3.40	31.82
Japanese yen	3.46	106.91

Although the assumption that asset returns are normally distributed is convenient, in many situations, it is inappropriate. I have shown that independent returns compounded over long horizons are lognormally distributed. Moreover, op-

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

tion strategies and dynamic trading strategies result in skewed distributions. Finally, conditions that produce price jumps typically lead to leptokurtic return distributions. The benefit of convenience may not always outweigh the cost of imprecision.

FOOTNOTES

1. For a more detailed discussion of this notion, see M. Kritzman, "What Practitioners Need to Know About Uncertainty," *Financial Analysts Journal* (March/April 1991):17–21.
2. We should also expect annual returns or returns of any periodicity to be lognormally distributed. Because short horizon returns are relatively small, however, they do not differ significantly from the logarithms of their wealth relatives. Thus, their true lognormal distribution is well approximated by a normal distribution.
3. For a more in-depth review of this topic, see M. Kritzman, "What Practitioners Need to Know About Future Value," *Financial Analysts Journal* (May/June 1994):12–15.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.