Alternatives to the Standard Market Risk Model

In Chapters 2 through 5, we got pretty far using the standard model of jointly normally distributed asset or risk factor returns. It treats the main-body risks of a portfolio fairly accurately. We now need to take account of the fact that the model is not perfectly accurate. In particular, very large-magnitude returns occur much more frequently than the standard return model predicts, leading to far greater tail risk than the standard risk model acknowledges. The entire distribution of returns, not just the expected return and return volatility, is important to investors.

In this chapter, we look at the behavior of asset prices and alternative models to the joint normal model that might better explain return behavior. We will also see how market prices, especially of options, reflect these alternatives. In Chapter 13, we discuss stress tests, an approach to risk measurement that takes account of the prevalence of extreme returns. Tools such as VaR can help measure the risk of losses that, while large and unpleasant, will be a recurrent cost of doing business. The models described in this chapter and stress testing attempt to measure risks that are life-threatening to a financial firm.

10.1   REAL-WORLD ASSET PRICE BEHAVIOR

We start by comparing the standard model, in which asset prices or risk factors are lognormally distributed, to actual market price behavior, as evidenced by the time-series behavior of their returns. We then provide a few statistical measures of deviations from the normal and a visual tool that summarizes how far asset prices are from normal.

Deviations from the normal model can be summarized under three headings:

Kurtosis (or *leptokurtosis*, literally, "fat tails") is the phenomenon that large returns occur more frequently than is consistent with a normal distribution. The coefficient of kurtosis is the fourth standardized moment of a distribution and provides a statistical measure of the frequency of large positive or negative asset returns. The kurtosis of the normal distribution is precisely 3, so the *kurtosis excess* is defined as the kurtosis coefficient minus 3.

High kurtosis means that returns far above or below the mean occur relatively often, regardless of sign. Since that implies that fewer returns are in the center of the distribution, kurtotic distributions are "peakier" than nonkurtotic ones. If there is enough displacement of probability mass out to the tails, the distribution may exhibit multiple modes or peaks.

The upper panel of Figure 10.1 compares a kurtotic distribution to a normal with same mean and standard deviation.

Skewness. Large moves in asset prices are not necessarily symmetrically distributed; rather, large moves in one direction occur more frequently than in the other. The *skewness coefficient* of a distribution is its third standardized moment and provides a statistical measure of the tendency of large returns, when they occur, to have a particular sign. The normal distribution is symmetrical, that is, its coefficient of skewness is exactly zero. So a high positive or negative skewness coefficient in a return time series is inconsistent with the assumption that returns are normal.

The mean and median of a symmetrical distribution are equal. For a skewed distribution, they are not. The mean is lower than the median for a negatively skewed distribution, that is, one skewed to the left, or having a "fat" left tail. Like kurtosis, skewness can manifest itself in multiple modes in the distribution. The mean is higher than the median for a distribution that is skewed to the right, that is, has unusually large positive returns more often than negative ones.

The lower panel of Figure 10.1 compares a distribution with negative skewness to a normal with same mean and standard deviation. The mean of the skewed distribution is below its median.

Time variation. Asset return distributions are not identical over time. Return volatilities in particular vary, as we noted in introducing volatility estimators in Chapter 3. The variation in volatility behavior is only partly captured by the EWMA/RiskMetrics volatility estimator we described there. In particular, EWMA does not capture "regime changes" and other dramatic and lasting changes in behavior.
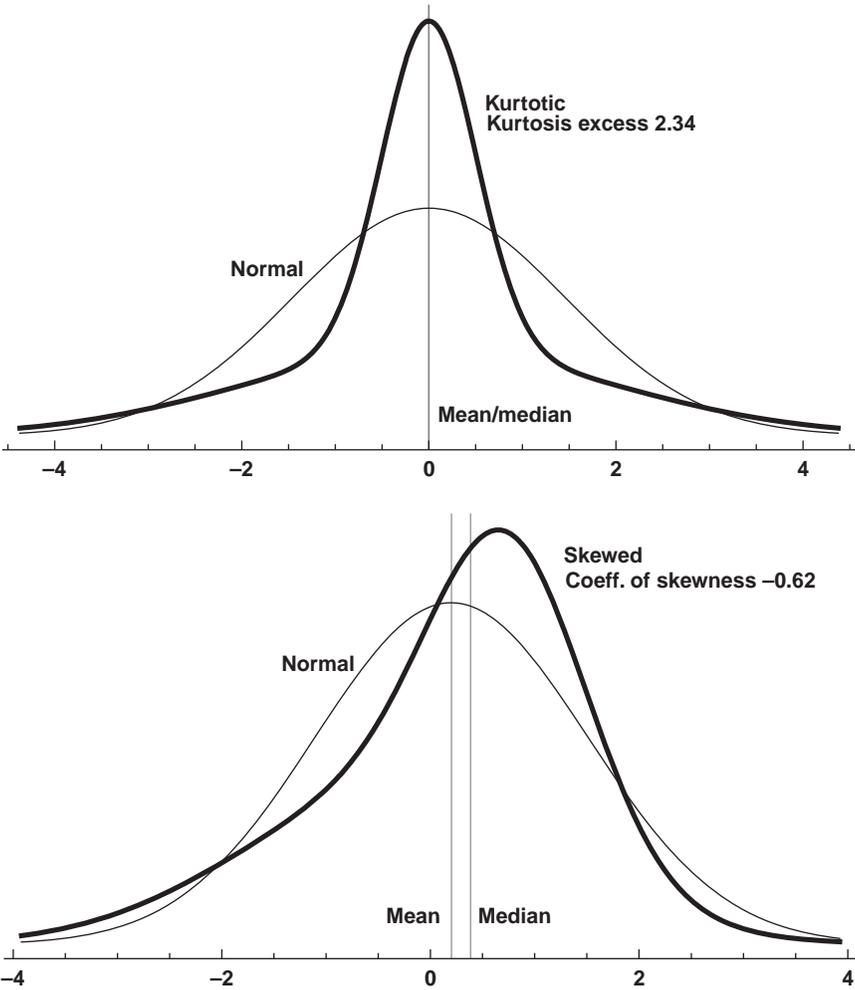
FIGURE 10.1 Normal and Non-Normal Distributions
Upper panel: Kurtotic and normal distributions, both zero mean and with identical variances. The kurtotic distribution is a mixture of two normals, with distributions $N(0, 2.0)$ and $N(0, 0.5)$, each with a probability of 50 percent of being realized. *Lower panel:* Skewed and normal distributions with identical means and variances. The kurtotic distribution is a mixture of two normals, with distributions $N(-0.35, 1.5)$ and $N(0.75, 0.75)$, each with a probability of 50 percent of being realized.

These stylized facts are interrelated. For example, a negative correlation between returns and volatility has been noted; large negative returns for stocks and stock indexes, for example, are more reliably followed by an increase in return volatility than positive returns. This phenomenon is called the *leverage effect*, since it is thought to be related to the use of borrowed funds in establishing positions. We discuss leverage and its impact on financial markets in detail in Chapters 12 and 14.

The anomalies we have just cited are closely related to the implied volatility smile and other departures from the Black-Scholes model predictions of option price patterns described in Chapter 5. The option skew is related to return skewness and to the leverage effect. The volatility smile is related to kurtosis. The term structure of implied volatility is related to the time-variation of volatility. But all the departures from normality in historical return behavior jointly influence the so-called implied volatility biases. Later in this chapter we discuss how to extract the information in option prices about future return behavior more precisely.

It would be easy to assemble a large menagerie of historical asset price return plots that evidence highly non-normal behavior. Let's focus on just three assets that illustrate departures from, as well as very approximate adherence to, the normal model. We'll start with a detailed long-run view of the S&P 500, an important asset in its own right, since it has a very large market value, represents the core asset class of equity securities, and is often used to represent the market portfolio, the universe of all risky assets. It also illustrates many typical features of asset returns. We'll also look at two currency pairs, the exchange rates of the dollar against the euro and the Turkish lira.

Figure 10.2 displays daily returns on the S&P index for the past 83 years.[1] The solid lines are the 99 percent confidence interval for the one-day return, using the EWMA approach of Chapter 3 to update the estimates for each date. Extreme outliers (also called "upcrossings" or "exceedances") lying outside the 99.8 and 99.98 percent confidence intervals, are marked by x's and o's.[2]

---

[1]The S&P 500 index was introduced in 1957. Prior to 1957, data are for the various forerunner indexes published by Standard & Poor's.

[2]The term "exceedance" is typically used in the context of extreme moves in financial returns. "Exceedance" is not found in the online version of the *Oxford English Dictionary*. The term "excession" is typically used in the context of VaR testing, as in the next chapter, to describe returns greater than the VaR shock or losses greater than the VaR. The *Oxford English Dictionary* defines it as "[a] going out or forth," and helpfully provides a usage example from the seventeenth century.
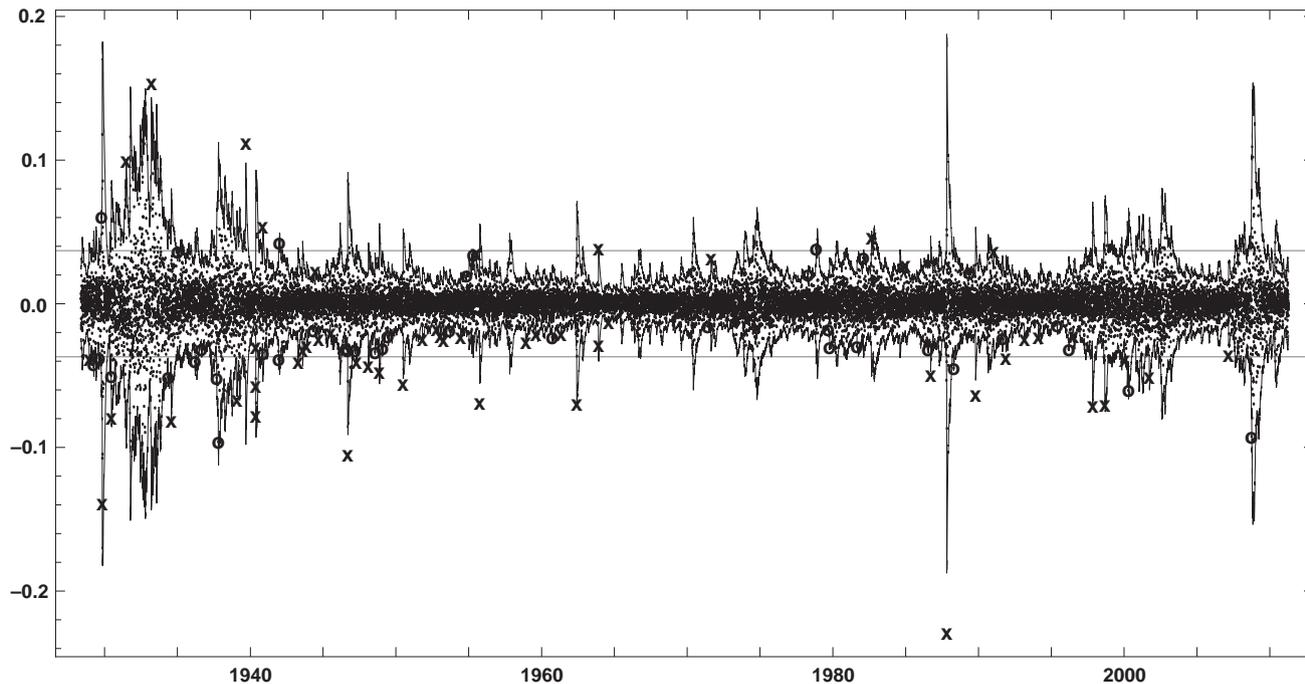
FIGURE 10.2   S&P 500 Daily Returns 1928–2011

The return data cover the interval from Jan. 3, 1928 to Apr. 14, 2011 and are represented by tiny dots. The thin solid plots represent the ±3.09 standard deviation or 9.8 percent forecast confidence interval, based on the EWMA method using 90 days of data and with the standard decay factor of 0.94. Next-day returns outside the 99.98 percent forecast confidence interval are marked by o's. Next-day returns outside the 99.98 percent forecast confidence interval are marked by x's. The horizontal grid lines mark the 99.98 percent forecast confidence interval for returns using the unconditional daily standard deviation of the entire historical sample.

Data source: Bloomberg Financial L.P.

The distribution displays kurtosis—the number of outliers is greater than one would expect if S&P returns were time-varying, but conditionally normally distributed. The excess kurtosis is quite substantial at 19.2. If the distribution were in fact conditionally normal, the "surprise" at observing "too many" exceedances becomes much greater, the higher the confidence level. Suppose daily returns are a normal random variable with a constant volatility. At the 95 percent confidence level, one should expect about one exceedance per month. But the S&P 500 time series exhibits about 30 percent "too many." At the 99.999 percent confidence level, one should expect to see an exceedance only once in a few centuries, but there are 211 times the expected number for the 80 years we observe.

The outliers are not evenly divided between positive and negative returns, as would be expected for a symmetric conditional return distribution. Rather, the distribution of S&P returns is skewed, with a coefficient of skewness of −0.48. Large negative returns are predominant, and this disparity also increases with the confidence level, up to a very high confidence level of 99.99 percent. The latter phenomenon reflects the sharp positive returns that often occur after the S&P index has had sustained and sharp negative returns.

Table 10.1 summarizes the results of this exercise, which can be thought of as a test of the accuracy of VaR of the type to be discussed in Chapter 11. The lower limit of the confidence interval plotted in Figure 10.1 can be thought of as the one day VaR at the specified confidence level, expressed as a return, of a long unleveraged S&P 500 position.

Apart from the high-frequency variations in volatility captured by the EWMA confidence-interval limits, the S&P 500 also displays low-frequency changes in volatility that can persist for years. Two periods of extremely high volatility commence in 1929, when the Great Depression began, and in 1937, when a severe unanticipated relapse occurred. They are separated by

TABLE 10.1   Extreme returns in the S&P 500 Index 1928–2011

Confidence level	0.95	0.99	0.9999	0.99999
Number of exceedances	1331	453	77	44
No. negative exceedances	738	288	62	35
No. positive exceedances	593	165	15	9
Ratio negative/positive	1.2445	1.7455	4.1333	3.8889
Rel. frequency of exceedances	0.0639	0.0217	0.0037	0.0021
Expected no. exceedances	1014.5	208.3	2.1	0.2
Actual/expected no. exceedances	1.3	2.2	37.0	211.2

Daily S&P 500 index returns falling outside a confidence interval with the stated confidence level, based on a prior-day EWMA estimate of volatility. Return data are those of Figure 10.1. EWMA estimates use 90 days of data and a decay factor of 0.94.

two years of relatively low volatility. Similarly, volatility was much higher beginning in the late 1990s through 2003 than during the subsequent three years. This unusually low level of volatility, as we see in Chapter 14, was an important feature of the run-up to the subprime crisis.

Figure 10.3 displays similar time-series return plots for the euro and dollar-Turkish lira exchange rates over an approximately 15-year period. These assets display a number of similarities and contrasts to the S&P 500. The euro return history, in the upper panel, is among the return distributions of widely-traded assets that are closest to the normal. While the volatility varies over time, the variations are not dramatic, apart from the increase in volatility coinciding with the subprime crisis. The exceedances from the conditional 99 percent confidence interval are not particularly frequent, nor are exceedances in one direction more frequent than in the other. So the kurtosis and skewness appear low under visual inspection.

The Turkish lira return plot presents a sharp contrast. Its most noteworthy feature is one extremely sharp move, a decline in the lira against the dollar of about 30 percent occurring on February 22, 2001. The lira depreciation was occasioned by Turkey's abandonment of its exchange-rate peg against the dollar. Such currency crashes are not unusual, and are discussed further in Chapter 14. The lira episode is a good illustration of *regime switching*, a shift from one set of statistical parameters governing asset return behavior to another. In the Turkish lira case, the switch is one of both the official exchange-rate regime administered by Turkey as well as a shift in return behavior.

Regime-switching models have been applied using a range of techniques and for a range of assets. One approach, for example, is to model the asset return distribution as a mixture of two normal distributions, as are the distributions displayed in Figure 10.1. One normal distribution describes the behavioral regime currently in place, while the second, perhaps higher-volatility, distribution comes into effect with a probability equal to the mixing parameter.

In foreign exchange, the problem of capturing large moves is called the *peso problem*, since it was first identified in the context of Mexican peso exchange rate behavior in the late 1970s and early 1980s. Forward exchange rates persistently predicted greater peso depreciation against the U.S. dollar than actually took place. This bias could be explained by the occasional drastic depreciations that took place when Mexico carried out a devaluation.

Apart from extreme time variation in return volatility, the Turkish lira also illustrates a high degree of kurtosis and skewness. The statistics describing these phenomena are dominated by the lira's return behavior around exchange-rate regime switches.

However, both examples show that the normal return model with volatility forecast using predominantly recent information is not a bad first
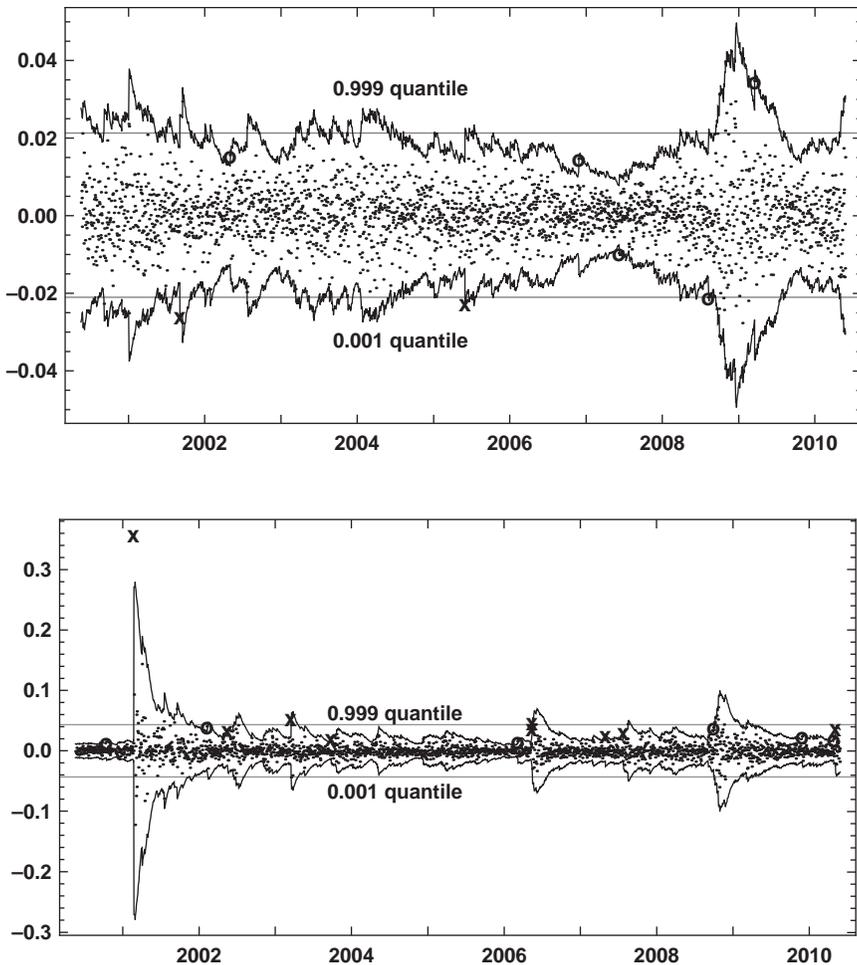
FIGURE 10.3 Statistical Properties of Exchange Rates

Each panel shows the daily return history from Jan. 2, 2000 to June 16, 2010, with the 99.8% confidence interval computed two ways, one using the volatility over the entire observation interval (horizontal grid lines), and the other using the time-varying EWMA volatility estimate. Next-day returns outside the 99.98 percent forecast confidence interval are marked by o's. Next-day returns outside the 99.998 percent forecast confidence interval are marked by x's.

Upper panel: EUR-USD: U.S. dollar price of the foreign currency, so a positive return is a dollar depreciation.

Lower panel: USD-TRY: Foreign currency price of the U.S. dollar, so a positive return is a lira depreciation.

Source: Bloomberg Financial L.P.

approximation to short-term return behavior, as long as it is not used to draw inferences about extreme tail behavior. Only the one wildly extreme lira observation is very far outside the forecast interval. The standard model describes the typical variations in asset values quite well. The shortcoming of the model is its failure to forecast returns that are extremely large and extremely rare.

The anomalies we have just observed are also evident for other assets. Table 10.2 displays the statistical properties of returns on a number of assets in various asset classes. These are based on the unconditional distributions of the asset returns, with volatilities represented by the root means squared (RMS). We use the RMS rather than the standard deviation in the spirit of the zero-mean assumption of Chapter 3. However, mean daily returns are very small for most assets most of the time, so the difference is small.

The VIX, an implied volatility index, has the highest return volatility. Among cash asset classes, equity indexes have the highest volatilities, as seen in both the ranges of returns and the root means squared. The Turkish lira range is very high relative to its volatility, partly due to the influence of the two large devaluation outliers.

Most of the asset returns display at least some distinctly non-normal characteristics. Normality can also be tested formally, for example, via the *Jarque-Bera* test. All the asset returns series of Table 10.2 have very high Jarque-Bera statistics (not displayed) that lead to rejection of the normality hypothesis even for very high confidence levels. Distributional hypotheses generally can be tested via the *Kolmogorov-Smirnov goodness-of-fit test*. The latter test is based on the idea that if a set of observations are generated by a specific distribution, then the largest outlier from that distribution is unlikely to exceed a specified amount.

Kurtosis appears to be universal; even the euro displays mild kurtosis. The Turkish lira has exceptionally high kurtosis. As we see in Chapter 14, currencies are among the few asset prices routinely subjected to price controls. Consequently, they display some of the largest-magnitude returns when authorities are obliged to lift these controls.

Skewness is characteristic of most of these assets. Equity markets exhibit mild downward skewness, while commodities and currencies differ one from the other. Fixed-income futures are skewed toward lower rates. The direction of skewness might well be different in a different observation interval or for a different sample of assets. However, the prevalence of skewness is well-attested.

There are several useful analytical and graphical tools for comparing the distribution of historical returns with a benchmark or model probability distribution. The *kernel estimator* is a technique for estimating the probability density function from a sample of the data it generates. It can

TABLE 10.2   Statistical Properties of Selected Asset Returns

	EUR	JPY	TRY	SPX	IBOV	CL1	GOLDS	CRB	TY1	ED1	VIX
Mean return (bps)	0.04	−0.98	13.49	2.34	33.35	2.50	2.35	1.02	0.41	0.16	0.77
Median return (bps)	0.00	0.00	0.00	5.65	24.23	4.60	0.00	1.05	1.38	0.00	−33.46
Minimum return (%)	−3.38	−6.95	−14.07	−9.47	−39.30	−40.05	−7.24	−6.01	−11.87	−0.81	−35.06
Maximum return (%)	3.47	5.50	35.69	10.96	34.21	22.80	10.24	5.93	3.54	0.93	49.60
Root mean square (% p.a.)	10.60	11.70	22.00	19.04	52.71	41.26	16.50	7.15	7.25	0.98	96.47
Skewness coefficient	0.00	−0.37	5.92	−0.18	0.12	−0.69	0.06	−0.36	−3.87	0.14	0.67
Kurtosis excess	1.78	5.08	137.02	8.38	19.31	16.75	8.90	16.58	96.14	48.71	3.79

Key to columns:

EUR Curncy: Euro spot
JPY Curncy: Japanese yen spot
TRY Curncy New Turkish lira spot
SPX Index: S&P 500 index
IBOV Index: Brazil Bovespa index

CL1 Comdty:generic 1st oil futures
GOLDS Index: Gold spot $/oz.
CRB Index: Commodity Research Bureau/Reuters index
TY1 Comdty: generic 1st 10-year note futures
ED1 Comdty: generic 1st eurodollar futures
VIX Index: CBOE SPX Volatility Index

Data from Jan. 2, 1900, to June 17, 2010.
Data source: Bloomberg Financial L.P.

be thought of as a method of constructing a histogram of the data, but with useful properties such as smoothness and continuity that a histogram lacks. A kernel estimator of a return time series $r_t, t = i, \ldots, T$ can be expressed as

$$\hat{f}(r) = \frac{1}{Th} \sum_{t}^{T} K\left(\frac{r - r_t}{h}\right)$$

To implement the kernel estimator, we need to specify K, the *kernel function*, and h, the *bandwidth*, which controls the smoothness of the kernel estimate. In typical applications, the standard normal density $\phi(\cdot)$ is employed as a kernel function, together with the so-called *optimal bandwidth*

$$h = \frac{1.06\hat{\sigma}}{(T-1)^{\frac{1}{5}}}$$

where $\hat{\sigma}$ is the sample standard deviation of the return time series. Specified this way, the kernel estimator can be viewed as creating a "histogram" in which each observation has a bar to itself. But instead of a bar, the estimator places a small normal density over each observation. Using very tight normal distributions this way would, in aggregate, produce a jagged distribution for the sample as a whole, potentially with many modes. Normals with very high standard deviation would produce a smooth distribution for the whole sample, but would obscure individual data points. The optimal bandwidth takes the sample standard deviation into account in finding a middle ground between these extremes.

In addition to intuitiveness, the kernel estimator has the virtue of being quite easy to implement. Figure 10.4 displays an important example, the VIX implied volatility index, that is pertinent to the discussion of option risk measurement in Chapter 5. Together with the sample moments displayed above, this example provides the useful insight that the VIX does not differ more than many cash assets from the normal distribution. It has significant, but not extreme kurtosis, and a noticeable positive skew; that is, large changes in implied volatility tend to be increases. This provides some support for using implied volatility returns as a risk factor in an option portfolio VaR estimate based on lognormal returns.

Another useful graphical tool for comparing historical returns to a benchmark distribution, such as the normal, is the *QQ plot*, short for "quantile quantile" plot, in which the quantiles of the historical and benchmark distributions are plotted against one another.

To generate a QQ plot, we need two things. First, we need the *order statistics* of the historical return series; that is, we order the returns in
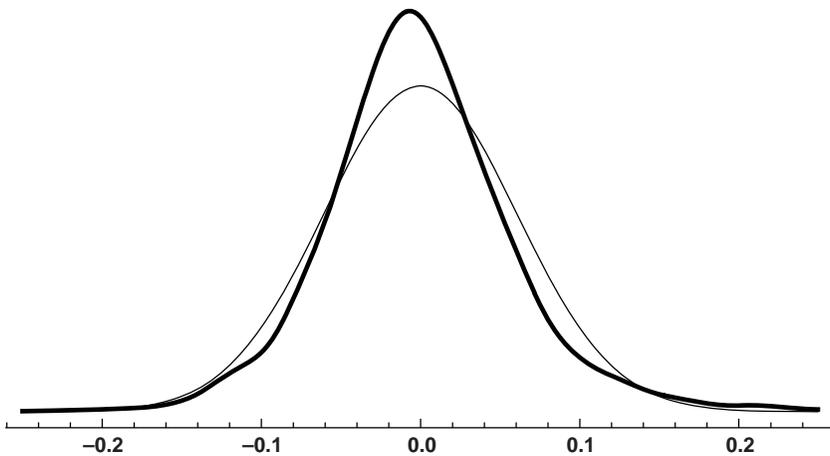
FIGURE 10.4   Kernel Estimate of the Distribution of VIX Returns
The thick plot is the kernel estimate of logarithmic changes in the VIX index,
January 2, 1900, to June 17, 2010, using an optimal bandwidth. The thin plot is a
normal distribution with mean and variance equal to the sample mean and
variance of the data.
Data source: Bloomberg Financial L.P.

ascending order by size rather than by date. Second, we need a benchmark
probability distribution, that is, the distribution against which we will com-
pare the historical return distribution at hand. Typically, we use the normal
distribution as a benchmark for historical asset returns, but we could also
use another specific alternative distribution to the normal.

Suppose we have $T + 1$ data points in our risk factor time series and
therefore T return data points. The fraction $\frac{i}{T}$ is the relative place of the
i-th order statistic in the return series. Based solely on the data sample, we
would say that the probability of the return having a realized value less than
or equal to the i-th order statistic is $\frac{i}{T}$. The likelihood of a return less than
or equal to the smallest (most negative) return (the first-order statistic) is
estimated at $\frac{1}{T}$, while that of a return less than or equal to the largest return
is estimated at 1.

The abscissa of the i-th point on a QQ plot is thus the i-th order
statistic. The ordinate of the i-th point is the $\frac{i}{T}$-quantile of the benchmark
distribution. For example, if $T = 1,000$, about four years of daily data, the
10th smallest return should be close to the first percentile of the benchmark
distribution. If the benchmark distribution is, say, $N(0, \sigma)$, with σ equal to
the historical daily volatility of the time series, then the pair consisting of
the 10th order statistic of the return series and the 0.01 quantile of $N(0, \sigma)$
constitute one point on the QQ plot. If the historical returns are, in fact,

drawn from the benchmark distribution, the two values are likely to be close together. The QQ-plot will then lie close to a 45-degree line through the origin.

What if the two distributions are dissimilar? Let's answer using the normal distribution, the usual benchmark for financial return behavior. The plot may take on a variety of shapes, depending on what kinds of deviations from the benchmark the historical returns display:

- Kurtosis manifests itself in an S-shaped plot. The extreme values occur more frequently than one would expect in a normal distribution, so the largest-magnitude negative historical returns will be smaller (more negative) than normal returns with the same expected frequency. The largest positive historical returns will be bigger than normal returns with the same expected frequency.

    The QQ plot will consequently lie *above* the 45°-line left of the origin, and *below* the 45°-line right of the origin. The QQ plot has to cross the 45°-line somewhere in between, hence the S shape.
- Suppose positive skewness, that is, to the right tail, is present in a historical return distribution, so large positive returns occur more frequently than negative ones. Then the largest historical returns will exceed normal returns with the same expected frequency. This manifests itself in a QQ plot that is below the 45° line to the right of the origin, but not as far above it to the left. The graph will be asymmetrical.
- If the historical return distribution evidences both skewness and kurtosis, the shape depends on which influence is more powerful. Kurtosis will impose a symmetrical S shape, while skewness will drag one arm further above or below the 45° line than the other.
- If the mean of the historical returns is higher than that of the benchmark, the QQ plot will be shifted, and lie to the left or right of the 45° line, rather than intersecting it at the origin. If the variance of the historical returns is higher than that of the benchmark, the QQ plot will be flatter and closer to a straight line.

There is a small problem in constructing a QQ plot using the tabulated quantiles of the benchmark distribution. The relative frequency of the T-th order statistic is 1, but the normal quantile of 1 is infinite. We therefore use the order statistics of a set of T random numbers generated by the benchmark distribution instead of the theoretical quantiles. This introduces some simulation noise (less, the more observations we have), but permits us to plot even the extreme returns in the historical data and is.

We illustrate these patterns in Figure 10.5 with two QQ plots of the currency returns displayed in Figure 10.3. The two plots contrast sharply in some ways, and are similar in others. The QQ plot for the EUR-USD
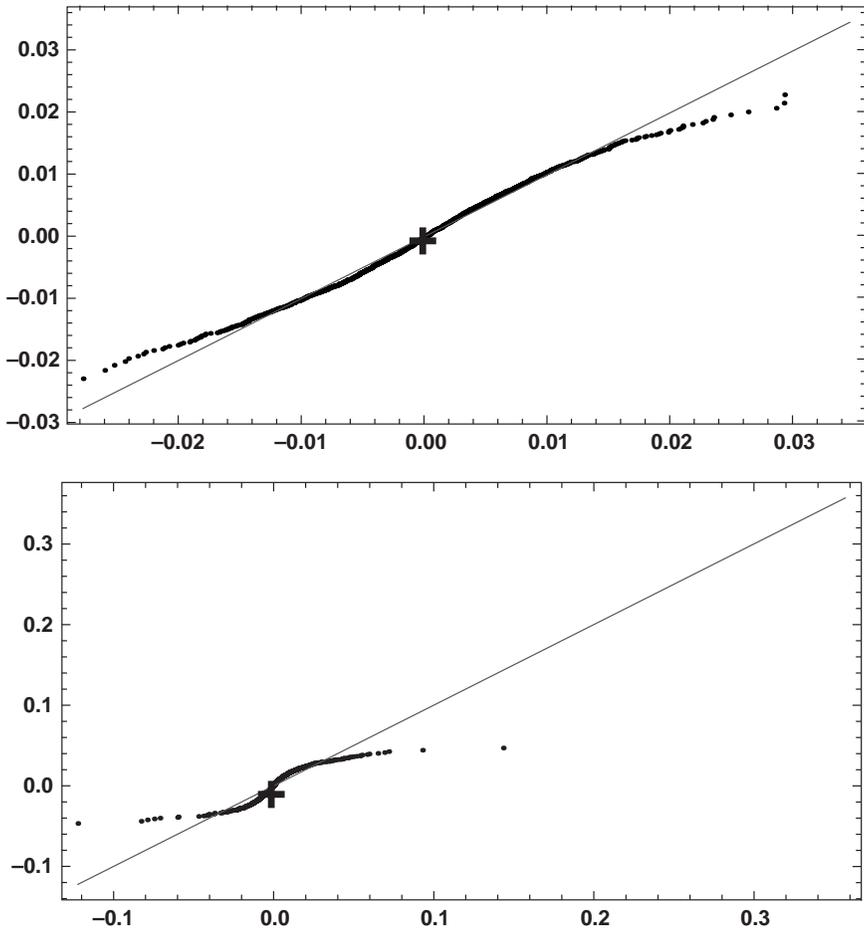
FIGURE 10.5 QQ Plot of USD Exchange Rates against the Euro and Turkish Lira Quantiles of daily returns, Jan. 2, 2000 to June 16, 2010, plotted against simulations from a normal distribution with mean zero and standard deviation equal to the historical volatility at a daily rate. The cross marks the origin.
Upper panel: EUR-USD exchange rate (USD per EUR). Standard deviation equal to 0.61 percent. Positive returns correspond to dollar depreciation.
Lower panel: USD-TRY exchange rate (TRY per USD). Standard deviation equal to 1.17 percent. Positive returns correspond to lira depreciation.
Source: Bloomberg Financial L.P.

exchange rate appears to be close to normally distributed. It lies reasonably close to the $45°$ line, though there appear to be some extreme negative and positive returns compared to the normal. This is consistent with the modest kurtosis excess we saw in the statistical summary. The USD-TRY plot in the lower panel shows very high kurtosis and considerable skewness. For example, the point corresponding to the largest USD-TRY return in the sample is further below the $45°$ line than the smallest return is above it. Both plots show only small deviations from the zero-mean return assumption. Overall, the appearance of both QQ plots is consistent with the statistical analysis presented earlier and summarized in Table 10.2.

10.2 ALTERNATIVE MODELING APPROACHES

A great variety of alternatives to the standard model have been put forward to better account for and forecast asset return behavior. In this section, we provide a few examples that are suggestive of the range of approaches: a specific alternative hypothesis to the stochastic process followed by asset prices, and a set of models focusing on forecasts of extreme returns.

10.2.1 Jump-Diffusion Models

Alternative models of the stochastic process that asset returns follow may more fully capture their behavior than the standard model. One such alternative, the *jump-diffusion model*, builds on the standard model of geometric Brownian motion. The jump-diffusion combines geometric Brownian motion with a second process, in which the asset price makes discontinuous moves at random times. This *jump process* is similar to the model we used to study defaults in Chapter 7. The major difference is that in default modeling, only one default event can occur; the default time is modeled as the first—and one and only—arrival time of a Poisson-distributed event, in that case default. A jump-diffusion process permits more than one jump to occur.

   The stochastic process thus imagines the asset price following a diffusion punctuated by large moves at random, Poisson-distributed times. To help compare it to geometric Brownian motion, we can write the stochastic differential equation (SDE) of an asset price S_t following a jump-diffusion as

$$dS_t = \left(\mu + \frac{1}{2}\sigma^2 - \lambda \mathbf{E}\left[k_t\right]\right) S_t dt + \sigma S_t dW_t + k_t S_t dq_t$$

   This differs from geometric Brownian motion, defined by Equation (2.1), by the addition of the jump term $k_t S_t dq_t$, and the corresponding
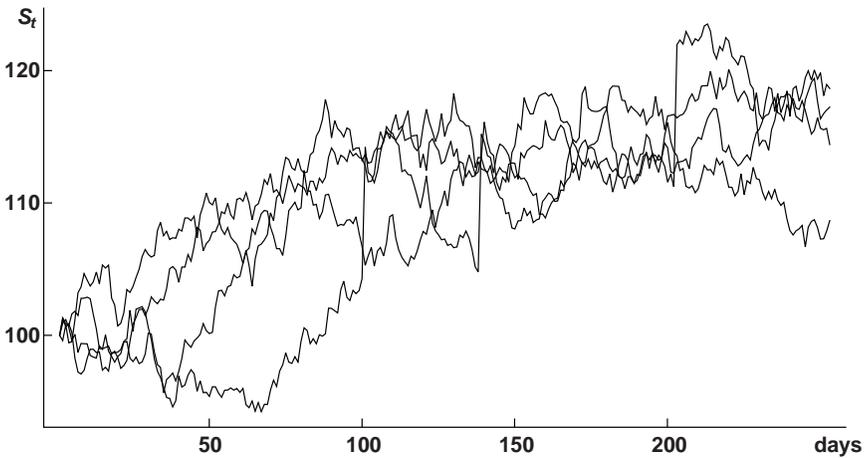
**FIGURE 10.6** Jump-Diffusion Process: Asset Price Level
Four simulations of the price level over time S_t of an asset following a
jump-diffusion process with $\mu = 0$ and $\sigma = 0.12$ (12 percent) at an annual rate.
Jumps are Poisson distributed with frequency parameter $\lambda = \frac{1}{252}$ and a
nonstochastic jump size $k = 0.10$ (10 percent). The initial price of the asset is
$S_0 = 100$.

adjustment of the mean increment to the asset price by the expected value
of the jump size $\lambda \mathbf{E}[k_t] S_t$. The size of the jump may be modeled as random
or deterministic. The (possibly random) jump size at time t is given by k_t,
measured as a percent of the current asset price S_t. The jump probability
is driven by a parameter λ, which plays a similar role to the hazard rate of
Chapter 7; dq_t is an increment to a Poisson process with

$$dq_t = \begin{Bmatrix} 1 \\ 0 \end{Bmatrix} \quad \text{with probability} \quad \begin{Bmatrix} \lambda dt \\ 1 - \lambda dt \end{Bmatrix}$$

Figure 10.6 displays four realizations of a jump-diffusion in which the
jumps are modeled as deterministic 10 percent increases in the asset price.
Figure 10.7 displays a simulation of a time series of returns from the same
jump-diffusion process. Figures 10.6 and 10.7 correspond to Figures 2.5 and
2.6 illustrating the behavior of a pure diffusion or random walk process.

With smaller, more frequent jumps that are hard to discern visually,
the jump-diffusion model can mimic a wide variety of asset-price behaviors.
The model can generate return time series that exhibit kurtosis as well as
skewness. The jumps per se generate kurtosis; if the jumps are modeled so
that they are not zero on average, they generate skewness.
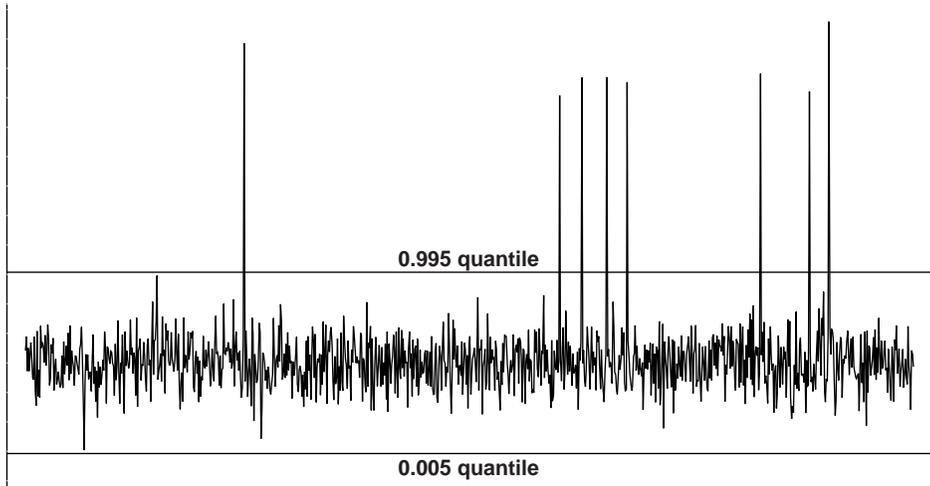
FIGURE 10.7   Jump-Diffusion: Daily Returns
Simulation of 1,000 sequential steps of a jump-diffusion process with $\mu = 0$ and $\sigma = 0.12$ (12 percent) at an annual rate. The y-axis displays returns in percent. Jumps are Poisson distributed with frequency parameter $\lambda = \frac{1}{252}$ and a nonstochastic jump size $k = 0.10$ (10 percent). The horizontal grid lines mark the 99 percent confidence interval.

With larger, less frequent jumps, the jump-diffusion model can mimic assets that are subject to sudden, large moves. One example is currency crises, which we have illustrated here with the Turkish lira and discuss further in Chapter 14. Another category of asset prices subject to sudden, drastic breaks in behavior, and well-described by the jump-diffusion model, are certain equities. For example, pharmaceutical and biotechnology companies are often highly dependent on one product or patent. Figure 10.8 displays an example, the stock price of Elan Corporation, a pharmaceutical company focused on diseases of the brain. Elan is highly dependent on a very small number of drugs for future revenue. On five occasions in the past decade, the firm has lost 50 percent or more of its market value in a single day, due to failure to obtain regulatory approval for a drug, a drastic swing in reported earnings, or reports of side effects from one of its key products. Its kurtosis excess is over 122, rivaling that of the Turkish lira exchange rate, and its skewness coefficient is an extremely large-magnitude –7.7.

10.2.2   Extreme Value Theory

A branch of statistics called *extreme value theory* (EVT) provides a somewhat different approach. Rather than looking for a parametric family of
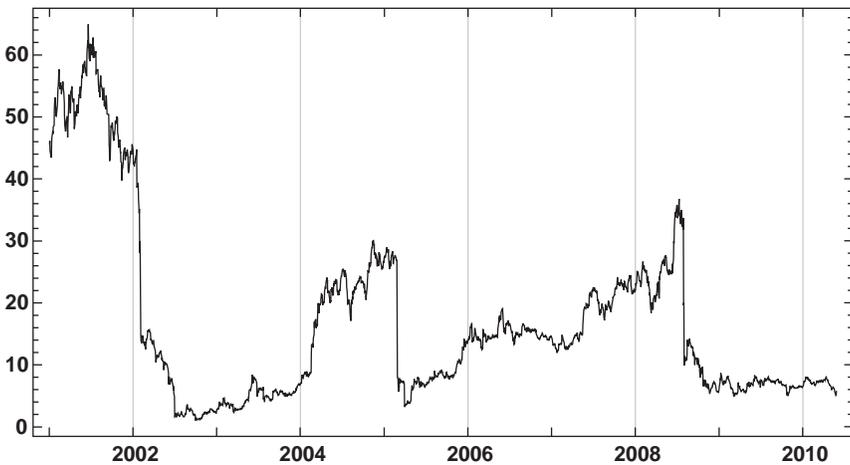
FIGURE 10.8 Elan Corporation Stock Price
Price of Elan Corporation PLC common shares, Jan. 2, 2001 to May 28, 2010.
Source: Bloomberg Financial L.P.

distributions or stochastic processes that can better explain observed returns than does the normal, it looks for techniques that can summarize and, hopefully, forecast extreme returns from a wide range of distributions. EVT concerns itself with the distribution of tail events, rather than with the distribution of main-body returns. Corresponding to this more data-focused approach, the user of these tools often has to make judgements about what constitutes extreme returns in a time series.

We'll illustrate these concepts using the long time series of S&P 500 returns underlying Figure 10.2. We start by presenting, in the table below, and in Figure 10.9, the long-term characteristics of the return series, as we did for the past two decades in Table 10.2. The returns display considerable kurtosis, and some negative skewness.

Long-term statistical properties of daily S&P 500 returns

Number of observations (*NOBS*)	20,921
Mean return (bps)	2.06
Median return (bps)	4.63
Minimum return (%)	−22.90
Maximum return (%)	15.37
Root mean square (%, annual)	18.96
Skewness coefficient	−0.46
Kurtosis excess	19.20
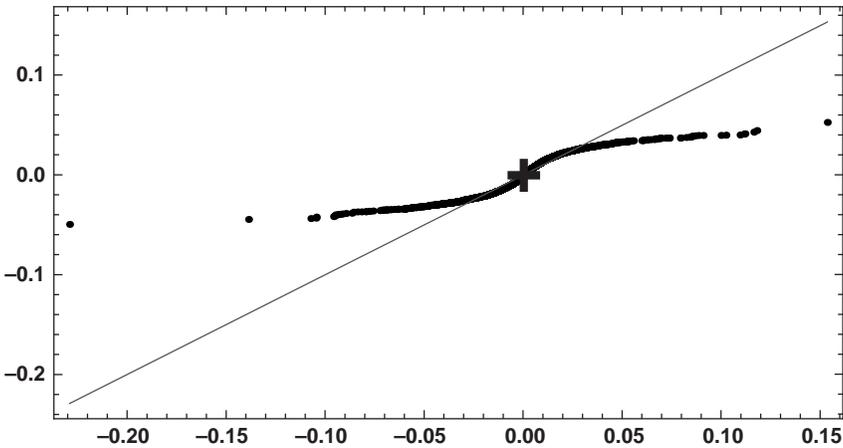
FIGURE 10.9   QQ Plot of the S&P 500
Quantiles of daily returns of the S&P 500 index, 1929–2011, plotted against
simulations from a normal distribution with mean zero and a standard deviation
equal to 1.2 percent, the daily historical volatility of the S&P 500 series over the
full sample. The data are those plotted in Figure 10.2.
Data source: Bloomberg Financial L.P.

In finance applications, EVT focuses on two ways of defining extreme
returns:

1. The maximum or minimum of a return stream
2. The size and probability of outliers, that is, returns larger or smaller
   than some large-magnitude threshold

A set of statistical tools has been developed to characterize these types
of extreme events. The starting point for both is a set of independently and
identically distributed (i.i.d.) returns $r_1, \ldots, r_t$. Nothing else is said at the
outset about their distribution function, which we denote $F(r_t)$. The largest
return (or, if we multiply the entire series by -1, the largest loss) is denoted
m_t. The order statistics of the data set are denoted $r_{1,t}, \ldots, r_{t,t}$.

For the fat-tailed distributions typically encountered in finance, the dis-
tribution of the normalized maximum return $\frac{m_t - a_t}{b_t}$ converges to the *Frechèt
distribution*:

$$\mathbf{P}\left[\frac{m_t - a_t}{b_t} \leq x\right] \to \exp\left[-\left(\frac{x}{\alpha}\right)^{-\alpha}\right] \quad \text{for} \quad t \to \infty \quad x > 0, \alpha > 0$$

where a_t and b_t are sequences of normalizing constants, playing the same
role as the mean and standard deviation in standardizing a normal random

variable. The Frechèt is an example of an *extreme value distribution*. (An analogous asymptotic distribution exists for thin-tailed distributions.) Under a few additional technical conditions, outliers in a Frechèt-distributed return series follow a *power law*, that is, if r_t is a large return, then

$$\mathbf{P}[r_t \geq x] = 1 - F(r_t) = -x^{-\alpha} \quad x, \alpha > 0$$

We call α the *tail index*. The tail index of the normal distribution is $\alpha = 2$. For fat-tailed financial security returns, we expect to see a tail index in excess of 2.

It's worth reflecting for a moment on how remarkable these statements are. We've said nothing about the distribution of the r_t, beyond their being independent and identically distributed (i.i.d.) and having a fat-tailed distribution. But we've been able to state distribution functions for the largest return in the series and for large returns generally, just as the central limit theorem tells us that the distributions of suitably normalized means of i.i.d. random variables converge to the normal distribution. Still, the power law is not a magic bullet; the i.i.d. assumption, for example, is not valid for returns that display stochastic volatility and other time-varying return properties.

A simple estimator for the tail index is called *Hill's estimator*. We set a standard of extreme return, somewhat arbitrarily, by identifying them with the k largest returns in the data set. The estimate is then

$$\hat{\alpha} = \left[\frac{1}{k}\sum_{i=1}^{k}\log(r_{i,t}) - \log(r_{k,t})\right]^{-1}$$

A problem with this simple estimator is that it can vary quite a bit with the threshold for defining extreme returns. If the threshold, which we'll denote u, is low, then k is high, and the estimate may include many observations that are from the main body rather than the tail of the distribution. That will bias $\hat{\alpha}$ downward. If the threshold is high and k is a small number, there may be too few observations to accurately estimate α. We can see this tendency in the estimates in the next example.

Example 10.1 (Hill's Estimator)   Each row of the next table shows the Hill estimate $\hat{\alpha}$ applied to S&P 500 returns that are below the threshold in the first column. The negative return threshold is converted to a positive one by multiplying the entire return time series by -1. The second column displays the number of observations k that are included in the estimate for each return threshold u, the third column displays their relative frequency in the time series, and the fourth column displays the Hill estimate.

But although $\hat{\alpha}$ varies widely as k changes, there is some consolation in the example, too. The estimates appear to be converging toward a value of about 4.5, consistent with the pronounced, but not extreme degree of kurtosis visible in Figure 10.9 and reflected in the coefficient of kurtosis excess.

u	k	$k \times NOBS^{-1}$ (%)	$\hat{\alpha}$
-0.020	751	3.590	2.334
-0.030	329	1.573	2.937
-0.050	79	0.378	3.447
-0.075	22	0.105	4.474
-0.100	5	0.024	4.559

Next, we give an example of how EVT can be used to provide not only a characterization of the return distribution, but also risk measures. We needed to select a return threshold that defined extreme returns in order to estimate the tail index. Similarly, we need to specify a threshold in order to estimate the probability that we will observe a return that is even larger than the threshold by a given amount. The amount by which a particular return exceeds a predefined standard of extremeness is called an *exceedance* or *peak over threshold* (so the approach is sometimes fetchingly referred to as POT modeling).

The probability of a return in excess of a threshold u is $1 - F(u)$, where $F(\cdot)$, again, is the distribution function of the r_t, about which little has had to be said. The probability of a return r_t that exceeds u by $y_t = r_t - u$ (the exceedance over threshold) is $1 - F(u + y_t)$, and the probability of realizing a return r_t between u and $u + y_t$ is $F(u + y_t) - F(u)$. Therefore, the conditional probability of a return between u and $u + y_t$ given that $r_t > u$ is

$$\mathbf{P}\left[r_t - u | r_t > u\right] = \frac{F(u + y_t) - F(u)}{1 - F(u)}$$

For i.i.d. fat-tailed returns, exceedances over thresholds follow a *generalized Pareto distribution* (GPD), with cumulative distribution function $G(y_t; \alpha, \beta, u)$:

$$\mathbf{P}\left[r_t - u | r_t > u\right] = G(y_t; \alpha, \beta, u) = 1 - \left(1 + \frac{y_t}{\alpha\beta}\right)^{-\alpha} \quad 0 < \alpha < \infty$$

where α is the tail index and β is a normalizing constant associated with the standard deviation. The parameters α and β can be estimated from data via maximum likelihood.

The estimation procedure starts by obtaining the density of the GPD $g(y_t; \alpha, \beta, u)$ by differentiating $G(\cdot)$ w.r.t. y_t:

$$g(y_t; \alpha, \beta, u) = \left[\frac{1}{\beta} \left(1 + \frac{y_t}{\alpha\beta} \right) \right]^{-(1+\alpha)}$$

The log-likelihood function is therefore

$$\sum_i^k \log \left\{ \left[\frac{1}{\beta} \left(1 + \frac{y_{(i,t)}}{\alpha\beta} \right) \right]^{-(1+\alpha)} \right\}$$

where k is the number of exceedances corresponding to u and $y_{(i,t)} = |r_{(i,t)}| - u$, $i = 1, \ldots, k$ is the set of exceedances in the data. To estimate the parameters α and β, we numerically find the values $\hat{\alpha}$ and $\hat{\beta}$ that maximize the log-likelihood function. The estimate of $G(y_t; \alpha, \beta, u)$ is then

$$G(y_t; \hat{\alpha}, \hat{\beta}, u) = 1 - \left(1 + \frac{y_t}{\hat{\alpha}\hat{\beta}} \right)^{-\hat{\alpha}}$$

We can combine this with the definition of the *conditional* probability distribution of exceedances to estimate the tail return distribution. The natural estimate of the probability $1 - F(u)$ of the conditioning event—a return less than or equal to the threshold u—is $\frac{k}{NOBS}$, the frequency of returns less than u in the sample. Our estimate of the conditional probability of an exceedance over u less than or equal to y_t is $G(y_t; \hat{\alpha}, \hat{\beta}, u)$, so the conditional probability of an exceedance greater than y_t is

$$1 - G(y_t; \hat{\alpha}, \hat{\beta}, u) = \left(1 + \frac{y_t}{\hat{\alpha}\hat{\beta}} \right)^{-\hat{\alpha}}$$

The *unconditional* probability of a return r_t in excess of u is therefore

$$\frac{k}{NOBS} \left(1 + \frac{r_t - u}{\hat{\alpha}\hat{\beta}} \right)^{-\hat{\alpha}}$$

We can use this probability distribution to calculate the probability of a very low return, or to find the VaR, that is, a low-return quantile. The next example illustrates.

Example 10.2 (Estimating VaR with EVT) Continuing the S&P 500 example, let's set a threshold of –7.5 percent. The 22 days on which returns were that low or lower are tabulated here, together with the order statistics of the corresponding exceedances $y_{(i,t)}$. The exceedance column of data will be entered into the log-likelihood function.

i	Date	$r_{(i,t)}$	$y_{(i,t)}$
1	24Sep1931	−0.0757	0.0007
2	31May1932	−0.0775	0.0025
3	14May1940	−0.0776	0.0026
4	09Oct2008	−0.0792	0.0042
5	16Jun1930	−0.0794	0.0044
6	26Jul1934	−0.0815	0.0065
7	12Aug1932	−0.0836	0.0086
8	05Oct1932	−0.0855	0.0105
9	26Oct1987	−0.0864	0.0114
10	10Oct1932	−0.0893	0.0143
11	21Jul1933	−0.0911	0.0161
12	29Sep2008	−0.0920	0.0170
13	20Jul1933	−0.0930	0.0180
14	01Dec2008	−0.0935	0.0185
15	15Oct2008	−0.0947	0.0197
16	05Oct1931	−0.0951	0.0201
17	18Oct1937	−0.0956	0.0206
18	03Sep1946	−0.1044	0.0294
19	06Nov1929	−0.1045	0.0295
20	29Oct1929	−0.1071	0.0321
21	28Oct1929	−0.1386	0.0636
22	19Oct1987	−0.2290	0.1540

Our estimated parameters are $\hat{\alpha} = 4.514$, not far from the Hill estimate of the tail index, and $\hat{\beta} = 0.0177$, higher, but not by much, than the full-sample standard deviation of returns. They are obtained numerically as the values that maximize the likelihood function.

We have $k = 22$ and $NOBS = 20{,}921$. Our estimate of the probability of a return less than −7.5 percent is $22 \times 20{,}921^{-1} = 0.001052$. The estimated probability of a return of, say, −10 percent is therefore

$$0.000307 = 0.001052 \left(1 + \frac{0.10 - 0.075}{4.514 \times 0.0177}\right)^{-4.514}$$

or 3.1 basis points. The actual frequency of returns of −10 percent or less in the sample is 2.4 basis points.

The VaR shock at a confidence level of, say, 99.99 percent is the number r that satisfies

$$0.00001 = 0.001052 \left(1 + \frac{r - 0.075}{4.514 \times 0.0177}\right)^{-4.514}$$

or 12.96 percent. In other words, using our estimated extreme value distribution, we would expect to see a one-day decline of about 13 percent or worse roughly once in 40 years. This is not a surprising estimate in view of the two observations of such large outliers in our 80-odd years of daily data.

10.3 THE EVIDENCE ON NON-NORMALITY IN DERIVATIVES PRICES

In Chapter 2, we defined a risk-neutral probability distribution, the probability distribution of the future asset price that is implied by current market prices of assets. The risk-neutral distribution is contrasted with the real-life, "subjective" or "physical" probability distribution. The term "subjective" focuses on the fact that it is the distribution that "the market" or the representative agent believes in when assessing future returns and making investment decisions. The term "physical," particularly in the context of a specific model, focuses on the fact that it is the distribution the model posits is true.

We can use risk-neutral distributions to obtain information about what the market—that is, the consensus expressed in prices—thinks about the distribution of asset returns. Option prices contain a great deal of information about market perceptions of the distribution of future asset prices, and they adjust to take account of the deviations from the standard model we noted at the beginning of this chapter. The information is expressed through the implied volatility smile, but is masked in two ways. First, it is embedded in option prices and needs to be extracted, via techniques we describe shortly. Second, the information is on risk-neutral rather than physical probability distributions, and is therefore blended with information about market preferences concerning risk. In this section, we see how to use option prices to derive risk-neutral distributions.

10.3.1 Option-Based Risk-Neutral Distributions

Chapter 5 introduced the so-called Black-Scholes option biases or anomalies, the important ways in which actually observed option prices differ from the predictions of the model. We also saw examples of one aspect of this phenomenon, the implied volatility smile: the cross-section, at a point in time, of the implied volatilities of European call or put options on the same underlying and with the same maturity, but different exercise prices. Chapter 5 was focused on identifying sources of option risk and applying appropriate measures of that risk. But the option biases and the volatility smile also have great significance for the study of the real-world behavior of asset returns.

The Black-Scholes model is similar to the standard risk-management model of conditionally normal returns. The models differ mainly in that in

the Black-Scholes model does not allow for variation in volatility over time. The option biases are driven by market perceptions and expectations of future returns as well as by deviations of realized return behavior from the standard model. The volatility smile results from kurtosis in returns, and the expectation that it will persist in the future. Option skew expresses an expectation that large-magnitude returns in a particular direction will predominate.

Risk appetites and the desire to hedge against rare events play the key role here. Some option biases, such as the equity market put skew—the tendency for low-strike options to be expensive relative to high-strike options—have been remarkably persistent. The put skew manifested itself through many years in which the actual behavior of equity returns was much closer to the normal model. Only with the subprime crisis did the "crash insurance" it seems to have captured appear fully warranted. The persistent put skew is similar in this respect to the peso problem in currency forward prices, discussed earlier in this chapter.

The techniques of this section build on the asset pricing model of Chapter 2. There, we saw that risk-neutral probabilities are equal to the present values of elementary claims that provide a payoff of \$1 in one specific future state, and 0 otherwise. The value of an elementary claim is related to the state of the world—feast or famine—it is associated with, and the representative agent's, that is, the market's, desire to hedge against low-consumption states.

Chapter 2 is a finite-state setting, with discrete states. In a continuous-state setting, we have a probability density function, rather than a set of discrete probabilities. To understand how to actually estimate the risk-neutral distribution using option prices, we reframe the discussion in terms of a continuous risk-neutral density function $\tilde{\pi}(S_T)$, where S_T is the future asset price. But we return to the finite-state to describe practical procedures for estimating risk-neutral probabilities.

In a number of areas of applied finance, as noted, a market index, typically the S&P 500, is a proxy for the state of the economy, and its future random value is an index of the future state. "States" are modeled as "realizations of future price." Taking that analogy a step further, if we can estimate the risk-neutral density of the S&P 500 at some future date, we have a proxy for the state price density.

Option Prices and Risk-Neutral Distributions   We start by presenting an important result, known as the *Breeden-Litzenberger formula*, about the relationship between prices of European call options and the risk-neutral density. Specifically, the mathematical first derivative of the call option's value with respect to the strike price is closely related to the risk-neutral probability that the future asset price will be no higher than the strike price.

The payoff at maturity to a European call option maturing at time T, with an exercise price X, is $\max(S_T - X, 0)$. The observed market value at

time t of a European call option is therefore the present expected value of that payoff under the risk-neutral distribution.[3] In the Black-Scholes model, in which perfect delta hedging is possible, that expected value is evaluated using the risk-free rate, rather than a discount factor that includes a risk premium. For our purposes, since we are trying to extract a risk-neutral, rather than the subjective probability, it is appropriate to use the risk-free rate even without the Black-Scholes model assumptions. We are not formulating an alternative or independent model to the market of what the asset price ought to be. Rather, we can describe what we are trying to do in two complementary ways:

1. Find a probability distribution that matches up with the market value, thus blending the risk premiums embedded in observable asset prices into probabilities assigned to various outcomes
2. Find the subjective probability distribution that a representative agent would have to have in his head, if he were indifferent to risk and market prices were as we find them

Matching the option price to the risk-neutral present expected value of that payoff gives us

$$c(t, \tau, X) = e^{-r\tau} \tilde{E} \left[\max(S_T - X, 0)\right] = e^{-r\tau} \int_X^\infty (s - X)\tilde{\pi}(s)ds$$

where
$\qquad S_T =$ terminal, or time-T, asset price
$\qquad X =$ exercise price
$\qquad \tau \equiv T - t =$ time to maturity
$\qquad c(t, \tau, X) =$ observed time-t price of an option struck at X and maturing at time T
$\qquad \tilde{E}\left[\cdot\right] =$ an expectation taken under the risk-neutral probability measure
$\qquad \tilde{\pi}(\cdot) =$ risk-neutral probability density of S_T, conditional on S_t
$\qquad r =$ continuously compounded risk-free rate, assumed constant over both time and the term structure of interest rates

---

[3]We specify European options since they have a fixed maturity date on which they can be exercised. American options, which can be exercised at any time prior to maturity, are not associated with one fixed forecast horizon, and consequently don't lend themselves as well to estimating fixed-horizon risk-neutral distributions. Most exchange-traded options and options on futures have American-style exercise.

We're assuming for simplicity that the asset pays no dividend or other cash flow. Taking the integral over the interval $[X, \infty]$ lets us substitute out the $\max(\cdot)$ function in the expression for the expected value of the payoff, the difference between S_T and X, given that S_T is greater.

We define the risk-neutral probabilities by matching them up with the option prices. There is an important assumption behind this, as we saw in Chapter 2: the absence of arbitrage opportunities. The risk-neutral probabilities may be quite far from the physical probabilities market participants actually believe in. But the market prices cannot contain opportunities to make money without risk; otherwise the risk-neutral probabilities would not be well-defined. Another assumption is that there are observable call prices for any strike price X. The sparser the option price data, the more difficult is the empirical estimation process.

Differentiating the no-arbitrage market call price with respect to the exercise price X, we have

$$
\begin{aligned}
\frac{\partial}{\partial X}c(t, \tau, X) &= e^{-r\tau}\frac{\partial}{\partial X}\int_{X}^{\infty}(s - X)\tilde{\pi}(s)ds \\
&= -e^{-r\tau}\int_{X}^{\infty}\tilde{\pi}(s)ds \\
&= -e^{-r\tau}\left[\int_{-\infty}^{\infty}\tilde{\pi}(s)ds - \int_{-\infty}^{X}\tilde{\pi}(s)ds\right] \\
&= e^{-r\tau}\left[\int_{0}^{X}\tilde{\pi}(s)ds - 1\right]
\end{aligned}
$$

In the second line of this derivation, we used Leibniz's Rule to differentiate with respect to an integration limit. In the third line, we recognized that the lower limit of integration can't be less than zero, because asset prices can't be negative. We also split the integral into the difference of two integrals. In the fourth line, we recognized that $\int_{-\infty}^{\infty}\tilde{\pi}(s)ds = 1$, because $\tilde{\pi}(\cdot)$ is a probability density function and, again, that $\int_{-\infty}^{X}\tilde{\pi}(s)ds = \int_{0}^{X}\tilde{\pi}(s)ds$.

This result implies that the risk-neutral cumulative distribution function of the future asset price is equal to one plus the future value of the "exercise price delta" of the market price of a European call:

$$
\tilde{\Pi}(X) \equiv \int_{0}^{X}\tilde{\pi}(s)ds = 1 + e^{r\tau}\frac{\partial}{\partial X}c(t, \tau, X)
$$

Differentiate again to see that the risk-neutral probability density function is the future value of the second derivative of the call price with respect to the exercise price:

$$\tilde{\pi}(X) = e^{r\tau} \frac{\partial^2}{\partial X^2} c(t, \tau, X)$$

Extracting the risk-neutral density from option prices preserves an important property of asset prices, namely, that the expected value of the future asset price under the risk-neutral probability distribution equals the current forward asset price. To see this, consider a call option with an exercise price of zero:

$$c(t, \tau, 0) = e^{-r\tau} \int_0^\infty s\tilde{\pi}(s)ds$$
$$= e^{-r\tau} \tilde{E}[S_T]$$
$$= e^{-r\tau} F_{t,T}$$

where $F_{t,T}$ is the forward price. The last line of this derivation holds by virtue of the definition of a forward price.

Using call prices is a slightly convoluted path to the risk-neutral cumulative distribution function. Put prices are more direct: The first derivative of the price of a European put $p(t, \tau, X)$ with respect to the exercise price is the future value of the risk-free cumulative distribution function itself. The payoff at maturity of a put with an exercise price X, maturing at time T, is $\max(X - S_T, 0)$. The current value of a put is therefore

$$p(t, \tau, X) = e^{-r\tau} \tilde{E}[\max(X - S_T, 0)] = e^{-r\tau} \int_0^X (X - s)\tilde{\pi}(s)ds$$

and its first derivative with respect to X is

$$\frac{\partial}{\partial X} p(t, \tau, X) = e^{-r\tau} \int_0^X \tilde{\pi}(s)ds$$
$$= e^{-r\tau} \tilde{\Pi}(S_T)$$

The second derivative with respect to the exercise price is identical to that of a call.

Pricing Elementary Claims from Options   Now let's return to the finite-state setup for more intuition on these results and to actually estimate risk-neutral distributions. We start by setting out a discretized version of our option-based estimate of the risk-neutral cumulative probability distribution and density functions. The "exercise price delta" and thus the CDF can be approximated by

$$\tilde{\Pi}(X) \approx 1 + e^{r\tau} \frac{1}{\Delta} \left[c \left(t, \tau, X + \frac{\Delta}{2} \right) - c \left(t, \tau, X - \frac{\Delta}{2} \right) \right]$$

As $\Delta \to 0$, this should get very close to the CDF. Similarly, the PDF can be approximated as

$$\tilde{\pi}(X) \approx \frac{1}{\Delta} \left[\tilde{\Pi} \left(X + \frac{\Delta}{2} \right) - \tilde{\Pi} \left(X - \frac{\Delta}{2} \right) \right]$$

$$= \frac{1}{\Delta} \left\{ 1 + e^{r\tau} \frac{1}{\Delta} [c(t, \tau, X + \Delta) - c(t, \tau, X)] \right\}$$

$$- \frac{1}{\Delta} \left\{ 1 + e^{r\tau} \frac{1}{\Delta} [c(t, \tau, X) - c(t, \tau, X - \Delta)] \right\}$$

$$= e^{r\tau} \frac{1}{\Delta^2} [c(t, \tau, X + \Delta) + c(t, \tau, X - \Delta) - 2c(t, \tau, X)]$$

In Chapter 5, we discussed option spreads and combinations, option portfolios that combine puts and calls, such as straddles. That discussion was in the context of vega risk, and we saw how these combinations embed information about the volatility smile and how to take the smile into account in order to accurately measure vega risk. In the present context, we will use the volatility smile to estimate risk-neutral distributions.

For the next bit of analysis, we rely on an option spread called a *butterfly*, which consists of long positions in two calls with different exercise prices, and a short position in a call with an exercise price midway between those of the long calls. Figure 10.10 shows the payoff profile of a butterfly with exercise prices 99, 100, and 101, and centered at 100. It corresponds to $X = 100$ and $\Delta = 1$.

Butterflies can be used to construct claims on an asset that pay off if the realized future asset price falls in a narrow range. A butterfly must be priced as if it were a lottery ticket paying off if that particular range for the future asset prices is realized. If the calls we are scrutinizing are options on a broad index such as the S&P 500, which is often taken as the price of a claim on future consumption, then the prices of butterflies are proxies for elementary
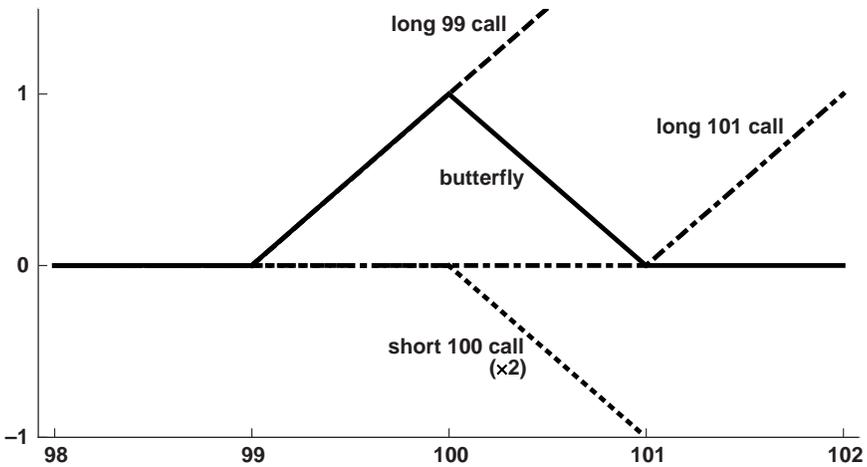
FIGURE 10.10 Constructing a Long Butterfly
The butterfly is constructed by combining long positions of one call option in each of the "flanking" strikes ($X = 99$ and $X = 101$) and a short position in two call options struck at the "center" of the butterfly ($X = 100$).

claims and, as we saw in Chapter 2, for a probability measure on the future asset price.

Suppose the future values of the asset can take on only integer values, that is, the state space is not only countable, but also doesn't have many elements. A long butterfly centered at 100 will pay exactly $1 if the future index level is 100 and zero otherwise. Absence of arbitrage implies that the price of the butterfly must equal the price of an elementary claim that pays $1 if the future price is 100 and zero otherwise. The value of the butterfly is
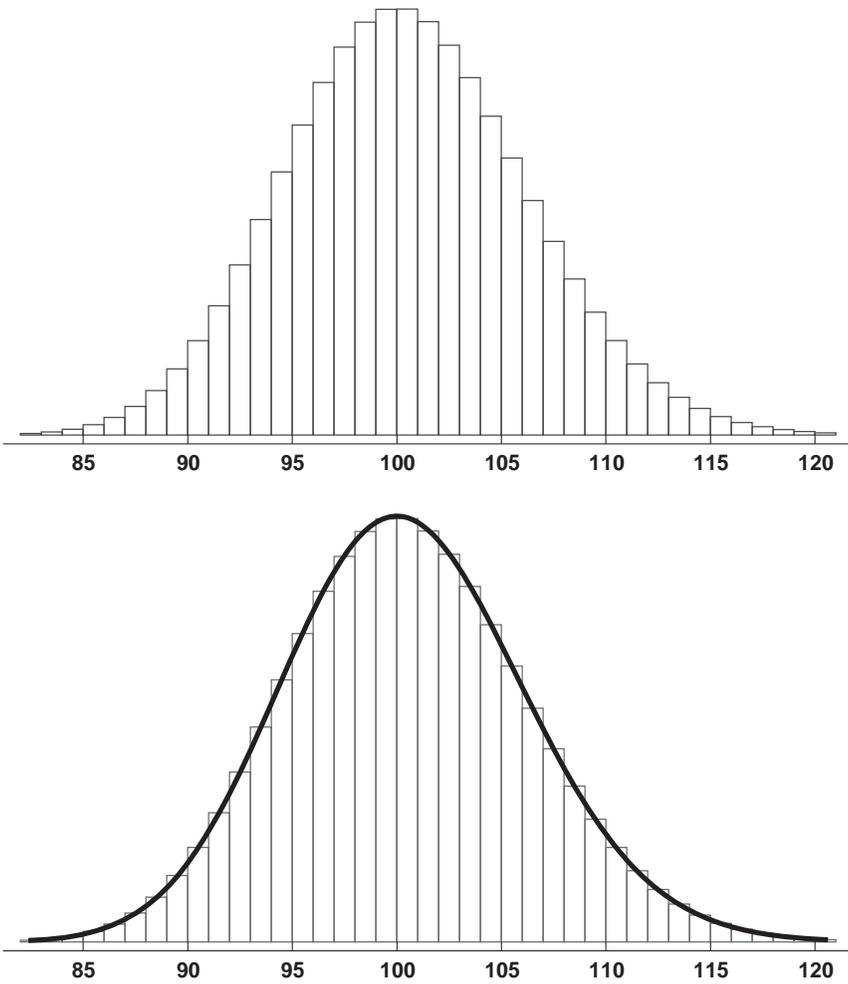
$$c(t, \tau, 99) + c(t, \tau, 101) - 2c(t, \tau, 100)$$

Let's now imagine that we "undiscretize" the state space by letting the future index take on noninteger values, say, at quarter-point intervals. We can still create a butterfly that pays $1 conditional on $S_T = 100$, but its value is now

$$\frac{1}{0.25} [c(t, \tau, 99.75) + c(t, \tau, 100.25) - 2c(t, \tau, 100)]$$

For any level of the terminal price $S_T = X$, and any price interval Δ, we have

$$\frac{1}{\Delta} [c(t, \tau, X - \Delta) + c(t, \tau, X + \Delta) - 2c(t, \tau, X)]$$

FIGURE 10.11   State Prices and the Risk Neutral Density
Upper panel: State prices along the price axis. The state prices are computed by assuming a specific risk neutral distribution for the future value of the index.
Lower panel: State prices and the risk neutral density.

The expression above gives the value of the elementary claim itself. These are equal to the *areas* of the rectangles in the upper panel of Figure 10.11. The *height* of each of these rectangles is approximately equal to the present value of the risk-neutral density of S_T, evaluated at X:

$$\frac{1}{\Delta^2} [c(t, \tau, X - \Delta) + c(t, \tau, X + \Delta) - 2c(t, \tau, X)]$$

Finally, taking $\Delta \rightarrow 0$, we get the expression for a continuous risk-neutral density derived above. This gives us some intuition on the Breeden-Litzenberger formula, and shows us the way to actually estimate a risk-neutral density.

We've mentioned that finding enough option prices is one of the challenges in estimating $\tilde{\pi}(S_T)$. The expression above spells out exactly what we need to estimate it at one specific point X: We need prices of three options, with exercise prices X, $X + \Delta$, and $X - \Delta$, all with the same maturity and all observed at the same time. Since there generally aren't that many different options trading on a single underlying asset, we will have to interpolate between the handful of observable option prices.

10.3.2 Risk-Neutral Asset Price Probability Distributions

Implementing this approach is difficult because it requires, in principle, a set of options with exercise prices closely spaced, in increments of Δ, on the asset price axis. In practice, not enough option contracts, with different exercise prices on a given asset with a given maturity, trade simultaneously. The option price data must also be of good quality; since we take second differences of the European call price as a function of strike to approximate the risk-neutral density, even small errors, such as rounding prices to the nearest tick, can lead to large anomalies in the estimated density. Carrying out these techniques therefore requires intensive data preparation to eliminate or correct flawed prices, and techniques for extracting information efficiently from the prices that remain after filtering.

Apart from lack of high-quality data, but also as a result of it, risk-neutral densities extracted from options can display violations of no-arbitrage conditions. Two in particular are worth mentioning:

1. The value of a call or put must be a convex function of the exercise price. Observed option prices interpolated ones can violate this no-arbitrage condition in a minor way over small intervals on the exercise price axis. This leads to computation of negative probability densities over the interval.
2. The mean of the risk-neutral density may not exactly equal the forward or futures price of the underlying asset.

A number of techniques have been developed to estimate risk-neutral densities from the generally sparse available data. We'll give a few examples of the results, based on a simple approach to estimation. It starts with the option data themselves. In order to have a European call price function that

is as smooth as possible, it helps to begin with option-implied volatilities as a function of the Black-Scholes call option delta rather than of exercise price. We'll denote the date-t implied volatility of a European option on a given underlying asset, maturing at time $T = t + \tau$, and with a delta δ, by $\sigma(t, \tau, \delta)$.

In the foreign exchange markets, options actually trade in these terms. In Chapter 5, we described combinations of options called risk reversals and strangles. These combinations can be readily converted into prices of individual options with specified deltas. For example, consider a 25-delta one-month strangle. Its price is quoted as the difference between the average implied vols of the 25-delta put and call, and the at-the-money forward (ATMF) put or call vol $\sigma(t, \tau, 0.50)$:

$$\text{strangle price} = \frac{1}{2}[\sigma(t, \tau, 0.25) + \sigma(t, \tau, 0.75)] - \sigma(t, \tau, 0.50)$$

This is equivalent to the implied vol of a butterfly set at the 25- and 75-delta strikes. The risk reversal quote is the implied vol spread between the two "wing" options:

$$\text{risk reversal price} = \sigma(t, \tau, 0.25) - \sigma(t, \tau, 0.75)$$

Note that strangle and risk reversal are quoted as vol spreads, while the ATMF is a vol level.[4] Using these definitions, the vol levels of options with different deltas can be recovered from the strangle, risk reversal, and ATMF quotes:

$$\sigma(t, \tau, 0.25) = \sigma(t, \tau, 0.50) + \text{strangle price} + \frac{1}{2} \times \text{risk reversal price}$$

$$\sigma(t, \tau, 0.75) = \sigma(t, \tau, 0.50) + \text{strangle price} - \frac{1}{2} \times \text{risk reversal price}$$

We can carry out the same operations for the 10-delta risk reversal and strangle. For most major currency pairs, these prices can all be readily obtained, for a wide range of tenors from overnight to several years.

Once we have a set of implied volatilities for different deltas, we can interpolate between them. There are a number of ways to do this, including the parametric approach of least-squares fitting and the nonparametric

---

[4]Note also as a minor detail that the ATM or ATMF option will have a delta close to but not exactly equal to 0.50.

approaches of applying a spline or an interpolating polynomial. However we go about it, the result is a function $\sigma(t, \tau, \delta)$, defined for any call delta $0 < \delta < 1$.

The next step is to find the exercise price X corresponding to $\sigma(t, \tau, \delta)$ for each delta. The quotation convention for implied volatilities is the Black-Scholes model, even though markets are perfectly aware the model is not accurate; it merely provides a set of units. The delta corresponding to an implied volatility is therefore the Black-Scholes delta, rather than the "true" sensitivity of the option value to changes in the underlying asset value. The Black-Scholes delta is

$$e^{-r^*\tau} \Phi \left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}} \right]$$

where r^* is the continuously compounded dividend, interest rate, or cost of carry of the underlying asset. For any particular value of the delta δ°, we can solve

$$\delta^\circ = e^{-r^*\tau} \Phi \left[\frac{\ln\left(\frac{S_t}{X}\right) + \left(r - r^* + \frac{\sigma(t, \tau, \delta^\circ)^2}{2}\right)\tau}{\sigma(t, \tau, \delta^\circ)\sqrt{\tau}} \right]$$

numerically for X to derive a volatility function $\sigma(t, \tau, X)$. The search algorithm would do so by finding the pair (σ, X) that lies on the interpolated volatility smile $\sigma(t, \tau, \delta)$ at the point δ°, and also returns δ° when substituted into the Black-Scholes delta.

The last step is to calculate the risk-neutral distribution. We substitute the volatility function $\sigma(t, \tau, X)$ into the Black-Scholes formula for the value of a European call option $v(S_t, \tau, X, \sigma, r, q)$ to obtain $v[S_t, \tau, X, \sigma(t, \tau, X), r, q]$. The volatility function $\sigma(t, \tau, X)$ is an estimate of the Black-Scholes implied volatility, that is, the volatility that, for exercise price X, would match the Black-Scholes formula to the market option price. So $v[S_t, \tau, X, \sigma(t, \tau, X), r, q]$ is an estimate of the market price of a call option with any X. In other words, we set

$$c(t, \tau, X) = v[S_t, \tau, X, \sigma(t, \tau, X), r, q]$$

Although we can differentiate $v(S_t, \tau, X, \sigma, r, q)$ algebraically, we generally will have to differentiate $v[S_t, \sigma(t, \tau, X, \tau, X), r, q]$ numerically. The cumulative probability distribution of the asset price is estimated by

$$\tilde{\Pi}(X) \approx 1 + e^{r\tau} \frac{1}{\Delta} \left\{ v \left[S_t, \tau, X + \frac{\Delta}{2}, \sigma \left(t, \tau, X + \frac{\Delta}{2} \right), r, q \right] \right.$$
$$\left. -v \left[S_t, \tau, X - \frac{\Delta}{2}, \sigma \left(t, \tau, X - \frac{\Delta}{2} \right), r, q \right] \right\}$$

and the density by

$$\tilde{\pi}(X) \approx e^{r\tau} \frac{1}{\Delta^2} \{ v[S_t, \tau, X + \Delta, \sigma(t, \tau, X + \Delta), r, q]$$
$$+ v[S_t, \tau, X - \Delta, \sigma(t, \tau, X - \Delta), r, q] - 2v[S_t, \tau, X, \sigma(t, \tau, X), r, q] \}$$

We don't have to believe in the validity or perfect accuracy of the Black-Scholes model at all to use it this way. All we need is our estimates of observed volatilities and the Breeden-Litzenberger relationship.

Example 10.3 Let's look at an example of how a risk-neutral distribution is constructed for the S&P index. The data used are displayed in the following table:

	September 29, 2008	May 27, 2010
Index	1,106.39	1,103.06
Risk-free rate	0.06	0.15
Dividend rate	2.72	1.99
10-δ call vol	44.08	20.03
25-δ call vol	46.29	21.52
40-δ call vol	48.79	23.30
50-δ call vol	50.62	24.67
60-δ call vol	52.65	26.38
75-δ call vol	56.39	29.95
90-δ call vol	61.88	37.20

Index: S&P index closing index level; risk-free rate: one month U.S. Treasury bill rate. All volatilities are of 1-month options, annualized and expressed in percent.

Figure 10.12 displays $\sigma(t, \tau, \delta)$ and $\sigma(t, \tau, X)$, the two versions of the interpolated volatility smile. The curve on the left, $\sigma(t, \tau, \delta)$, interpolates
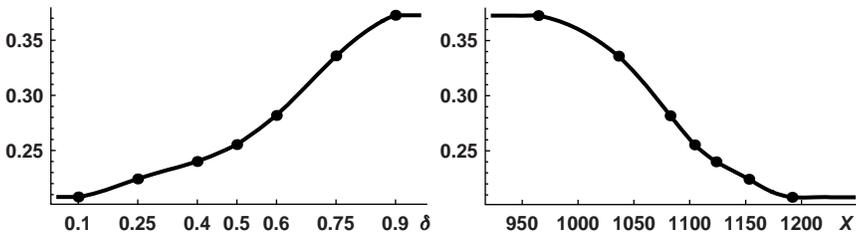
FIGURE 10.12   Fitted Implied Volatility Smile
The left panel displays the volatility smile in delta-vol space, the right panel in
strike-vol space, one month options on the S&P index, May 27, 2010. The points
represent actually observed implied volatilities, expressed as decimals.
*Data source:* Bloomberg Financial L.P.

between the volatility levels derived from the option quotes displayed in the
table above. The curve on the right is $\sigma(t, \tau, X)$, the one required for the next
step in obtaining a risk-neutral distribution, in which we estimate the prob-
abilities and the probability density as first and second differences in the call
option value. Note that the high implied volatilities are those for options
with high call deltas, corresponding to low strikes in index terms and low
put deltas.

The volatility smiles displayed in Figure 10.12 are interpolated by fitting
a cubic spline to seven observed implied volatilities for call deltas equal
to (0.10, 0.25, 0.40, 0.50, 0.60, 0.75, 0.90). The spline is "clamped" at the
endpoints by imposing the condition that the first derivative, that is, the
slope of the volatility smile, is equal to zero at $\sigma(t, \tau, 0.10)$ and $\sigma(t, \tau, 0.90)$,
the endpoint implied volatilities.[5]

The resulting risk-neutral densities for the two dates are displayed in
Figure 10.13. Each plot represents the probability density of the S&P index
one month in the future. The index levels on these two dates happen to
be very close to one another, just above 1,100, but the distributions of the
future exchange rate implied by options are very different.[6] The density on
September 29, 2008, is drawn from a high-volatility regime, just after the
Lehman bankruptcy. It is very dispersed around the mean and has high
skew and kurtosis. On May 27, 2010, in contrast, the distribution has a
much lower variance, in spite of the European debt crisis then currently in

---

[5]For more detail on how to compute cubic splines and clamped cubic splines, see
Stoer and Bulirsch (1993) and Klugman, Panjer, and Willmot (2008).
[6]The technique used is similar to that of Malz (1997). The main difference is to
employ a cubic spline rather than polynomial interpolation through the observed
volatilities to represent the volatility smile.
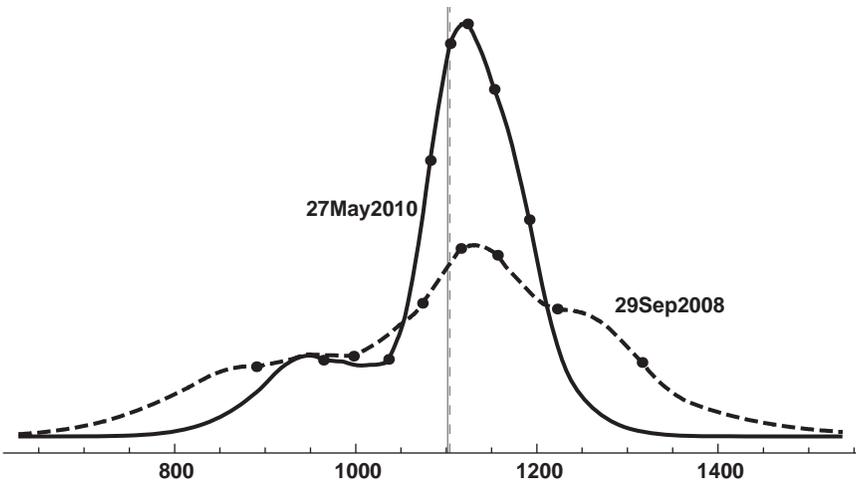
FIGURE 10.13 Estimated Risk-Neutral Densities
One-month risk-neutral density functions for the S&P index, May 27, 2010, and
September 29, 2008. The vertical grid lines represent the forward stock index level
for each date. The points correspond to the exercise prices for which implied
volatilities are directly observed.
Data source: Bloomberg Financial L.P.

full blossom, making the skew even more apparent to the eye. Both show a
pronounced negative skew toward lower index levels.

   One way to compare the two distributions is to compute the risk-neutral
probability of shocks of different sizes. The table below summarizes the risk-
neutral probabilities of arithmetic returns in percent on the index over the
month following the observation date. Thus we see that the risk-neutral
probability that the index would fall by 25 percent over the subsequent
month was over 4 times higher just after the Lehman bankruptcy as in the
spring of 2010.

Shock	May 27, 2010	September 29, 2008
−33.3	0.23	2.03
−25.0	1.48	6.18
−10.0	9.36	23.69
0.0	41.50	46.80
10.0	95.53	74.57
25.0	99.99	96.56
33.3	100.00	99.10

Another way to compare the distributions is through its quantiles. The next table displays the first and fifth percentiles of the risk-neutral distribution on the two dates:

	0.01 Quantile			0.05 Quantile	
	SPX close	SPX level	loss (%)	SPX level	loss (%)
May 27, 2010	1,103.06	852	22.6	917	16.7
September 29, 2008	1,106.39	716	35.1	809	26.7

The September 29, 2008, distribution attributes a probability of 1 percent to a one-month decline in excess of 35.1 percent in the index, to a level of 716. The later distribution states that a smaller decline, of only 22.6 percent or more, has a 1 percent likelihood of occurring.

Risk-neutral distributions are useful in several ways. A straightforward application is to draw inferences about the probabilities being assigned by the market, in its collective wisdom, to the realization of different asset price levels or events. For a broad index, such as the S&P 500, the probability that a low future index level is realized can be interpreted as the probability of an adverse state of the economy.

A fair critique of such inferences is that the probabilities obtained in this way are risk-neutral rather than physical, so treating them as straightforward forecasts is unwarranted. However, they nonetheless contain valuable information. Consider the S&P example. We cannot discern whether the dramatic fall in the risk-neutral probability of a massive decline in the index reflects a change in the market's belief in or fear of a market crash, but we know it must be some combination of the two, and the fact of the change is important to both traders and policy makers. The increased risk-neutral probability can also be interpreted as a large increase in the market's willingness to pay up to hedge against that event. That might indicate that some market participants with sensitivity to mark-to-market losses have long positions they are uncomfortable with.

Another important application of risk-neutral distributions is in estimating risk aversion. If we can obtain a reliable estimate of the physical probability distribution of a representative index such as the S&P 500, the differences from the risk-neutral distribution permit us to draw inferences about risk premiums. The physical distribution might, for example, be based on historical volatility estimates such as those presented in Chapter 3. One such approach, called *implied binomial trees*, has become important in developing tools for pricing exotic derivatives. It extracts not only the risk-neutral distribution of returns at one future point in time, but also the stochastic process followed by returns over time.

10.3.3 Implied Correlations

Risk-neutral probability distributions are one of several types of information about asset return distributions contained in market prices. Chapter 1 used data on nominal and inflation-protected bond yields to obtain risk-neutral estimates of future inflation rates (see Figure 1.14). In Chapters 7 and 8, we saw how to obtain risk-neutral estimates of default probabilities and correlations from credit derivatives prices. In this section, we extract data on return correlations among individual stocks, called the *risk-neutral implied equity correlation*, using the implied volatilities of individual stocks in an index and the implied volatility of the index itself. We describe how equity implied correlation is computed and see how it has behaved in recent years.

An equity index is a weighted sum of the constituent stocks. Its returns can be expressed as:

$$r_{\text{index},t} = \sum_{n}^{N} \omega_{nt} r_{nt}$$

where $r_{\text{index},t}$ represents the time-t index return, and ω_{nt} and r_{nt} the time-t weights and returns on the $n = 1, \ldots, N$ constituent stocks. The index return volatility $\sigma_{\text{index},t}^2$ is related to the N volatilities of individual stock returns by

$$\sigma_{\text{index},t}^2 = \sum_{i} \omega_{nt}^2 \sigma_{nt}^2 + 2 \sum_{n} \sum_{m<n} \omega_{mt} \omega_{nt} \sigma_{mt} \sigma_{nt} \rho_{mn,t}$$

where $\rho_{mn,t}$ is the time-t correlation between returns on stocks m and n. Note that the index volatility cannot exceed the average individual stock volatility; there cannot be negative diversification in the index as a whole.

Let's make a simplifying assumption, that the pairwise correlation $\rho_{mn,t} = \rho_t, \forall m, n$. This is analogous to the assumption that a single copula correlation or a single beta drives the pairwise default correlations in a portfolio credit model. We can then estimate an implied correlation by using the relationship

$$\rho_t = \frac{\sigma_{\text{index},t}^2 - \sum_{i} \omega_{nt}^2 \sigma_{nt}^2}{2 \sum_{i} \sum_{m<n} \omega_{mt} \omega_{nt} \sigma_{mt} \sigma_{nt}}$$

and substituting the index and individual stock implied volatilities for σ_t and $\sigma_{n,t}$.

As with any risk-neutral quantity, the implied correlation will differ from the actual implied correlation by an unobservable risk premium. The

risk premium will be driven by the estimated correlations and the urgency of hedging individual stock exposures. For a given expected correlation, the implied correlation or the risk premium will be higher when market participants are less eager to have undiversified individual equity exposures or are more averse to bearing systematic risk.

To get some additional intuition into this relationship, suppose all the stocks in the index are identical and the index is equally weighted, that is, $\sigma_{nt} = \sigma_t$ and $\omega_{nt} = 1/N, \forall n$. Then

$$\rho_t = \frac{\sigma_{\text{index},t}^2 - NN^{-2}\sigma_t^2}{2N^{-2}\dfrac{N(N-1)}{2}\sigma_t^2}$$

$$= \frac{\sigma_{\text{index},t}^2 - N^{-1}\sigma_t^2}{(N-1)N^{-1}\sigma_t^2}$$

$$\approx \left(\frac{\sigma_{\text{index},t}}{\sigma_t}\right)^2$$

In this simplification, the implied correlation is close to zero (unity) when the index volatility is small (large) relative to the typical individual stock volatility. During financial crises, index as well as single-stock volatilities rise sharply, but index volatility rises faster, as market participants flee systematic risk, driving implied correlation higher.

Just as the option skew can be interpreted as an indicator of the market's perception or fear of large asset returns, implied equity correlation can be interpreted as an indicator of the perceptions or fear of systemic risk. Figure 10.14 shows that implied correlation peaks at times of market stress. At the worst point of the subprime crisis, it drew close to its maximum possible value of 1.

Implied correlation is also a market risk factor. Just as with volatility, traders can take positions on the difference between the current and anticipated implied equity correlation over some future horizon. There are several ways to execute such *dispersion trades*. There can also be hard-to-detect implied correlation risk in any equity option portfolio hedged using index options, even if the trade is not focused on implied correlation. The primary motivation for such trades can include profiting from long gamma, from the difference between implied and realized volatility, or from individual option vega. Some examples of portfolios exposed to implied correlation include:

- Some traders hold portfolios of individual options. These may be arbitrage portfolios taking long positions in options considered overpriced and short positions in options considered underpriced, or as part of
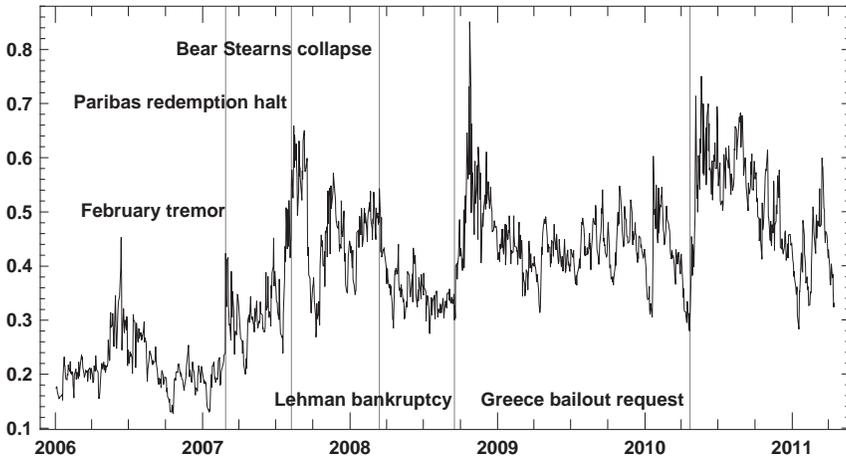
FIGURE 10.14 Risk-Neutral Implied Equity Correlation
Implied correlation of the S&P 500 index, using the largest 193 stocks and
market-capitalization weights and constituents as of Apr. 15, 2011. The included
stocks account for about 80 percent of the total market capitalization of the index.
Data source: Bloomberg Financial L.P.

long-short equity portfolios. The net long or short exposure of such
a portfolio is generally hedged with index options. A portfolio with a
substantial long individual option "overhang," hedged with short in-
dex options, will experience losses if implied correlation rises in a stress
environment.

- *Variance swaps* are OTC derivatives contracts in which one party pays a
  fixed premium and receives the squared returns on a stock or on a basket
  of stocks. They can be used for hedging or to take advantage of the fact
  that implied volatility generally exceeds realized volatility. Arbitrage po-
  sitions in variance swaps on individual stocks are generally hedged with
  variance swaps on equity indexes. When implied correlation increases,
  such portfolios are overhedged. The short position in index variance
  swaps has losses that may exceed the gains on the single-stock variance
  swaps.
- Consider a portfolio of convertible bonds in which the credit exposure
  is hedged with long protection positions in credit default swaps on the
  convertible bond issuers, the risk-free curve risk is hedged with payer
  swaps or short government bond positions, the option risk is delta
  hedged, and the vega risk is hedged via index options. Such a portfolio
  has risks quite similar to a correlation portfolio of individual equity
  options hedged with equity index options.

Another important implied correlation that can be extracted from op-tion prices is that between two currency pairs, for example, the correlation between the exchange rates of the U.S. dollar against the euro and the Japanese yen, or that between pound sterling's exchange rates against the euro and the U.S. dollar. This implied correlation can be extracted using prices of the three distinct European options on the three currency pairs involved. For example, to estimate the correlation between EUR-USD and USD-JPY, we would require options on the two major currency pairs, and, in addition, prices of options on the EUR-JPY exchange rate.

FURTHER READING

Two classic postwar papers studying asset return distributions are Fama (1965) and Mandelbrot (1963). Cont (2001) is a more recent summary of asset return behavior; de Vries (1994) summarizes the statistical properties of foreign exchange rates. Glosten, Jagannathan, and Runkle (1993) discusses the asymmetric impact of positive and negative returns on volatility.

Duffie and Pan (1997); Hull and White (1998); and Glasserman, Heidelberger, and Shahabuddin (2002) discuss fat-tailed return distribu-tions and alternatives to the standard model in the context of VaR model-ing. An adaptation of VaR to encompass extreme returns is applied to the Asian crisis of 1997 in Pownall and Koedijk (1999). See also Rosenberg and Schuermann (2006). The term "peso problem" was introduced in Krasker (1980). See Evans (1996) for a survey of the earlier literature on the peso problem.

Kernel density estimators are discussed in Silverman (1986) and Klugman, Panjer, and Willmot (2008).

The classic discussion of jump-diffusions and their properties is Mer-ton (1976). Empirical applications include Ball and Torous (1985), Jorion (1988), and Akgiray and Booth (1988).

Gabaix (2009) discusses finance applications of power laws. An early empirical application in finance was to currency pegs; see, for example, Nieuwland, Verschoor and Wolff (1994). On Hill's estimator, see Bangia, Diebold, Schuermann, and Stroughair (1999). Embrechts, Klüppelberg, and Mikosch (2004) is the standard work on extreme value theory, and while ad-vanced, has a number of accessible introductory sections. Klugman, Panjer, and Willmot (2008) includes an introduction to extreme value distributions. Introductions to extreme value theory are also provided by Longin (2000) and Neftci (2000).

Surveys of techniques for extracting risk-neutral distributions from op-tion prices include Jackwerth (1999, 2004), Mandler (2003), and Bliss and

Panigirtzoglou (2002). The Breeden-Litzenberger theorem was first stated in Breeden and Litzenberger (1978). Examples of the use of these techniques include Bates (1991), Shimko (1993), and Malz (1996, 1997). Examples of central bank use of these techniques in drawing inferences about market sentiment include Bahra (1996) and Clews, Panigirtzoglou, and Proudman (2000).

Amin and Ng (1997) and Christensen and Prabhala (1998) discuss "second-moment efficiency," the question of whether implied volatility is an accurate forecast of future realized volatility. The relationship between risk-neutral and physical distributions is studied in Aït-Sahalia and Lo (2000), Rosenberg and Engle (2002), and Bliss and Panigirtzoglou (2004). See Rubinstein (1994) and Derman and Kani (1994) on the implied binomial tree approach.

Implied correlations between exchange rates are studied in Campa and Chang (1998) and Lopez and Walter (2000).