

Risk Control and Mitigation

The purpose of having accurate risk measurement tools is to benefit traders and investors. In this chapter, we tie together the modeling tools we have developed and see how to move from risk measurement to risk management. Ultimately, the benefits of risk measurement come from putting investors in a better position to make tradeoffs between risk and return, and between different aspects of risk. Some investors may be more tolerant of volatility, but quite sensitive to the risk of large losses. If adequately compensated, other investors may prefer exposure to large losses, but remain averse to volatility, a risk profile characterizing some option portfolios and senior securitized credit products. Metaphorically, one can imagine investors choosing distributions of returns they prefer over distributions that are inferior from their point of view. It is only a metaphor, in view of the problems in maintaining a distributional hypothesis on returns that we've encountered in earlier chapters. Among the risk management objectives market participants might have are:

Reduce volatility of portfolio returns, including any hedges.

Diversification. Risk measurement tools can be used to ensure the portfolio does not contain any undesired concentrations and to identify exposures to extreme economic scenarios.

Left-tail truncation. Identifying extreme-loss scenarios and the positions that contribute to them can guide investors to reducing the extreme losses with only a small sacrifice of—or even an increase in—expected return.

Groping towards optimum. Risk management tools can help investors quantify risk-reward trade-offs and identify trades that improve their return distributions.

Sizing of trades in accordance with the risk takers' goals. As we have seen in Chapter 2 and will explore further here, the size of a position is not linearly related to its impact on the risk of a portfolio.

Selection of risk factors in accordance with the risk takers' goals. Hedging can reduce or eliminate undesired exposures you don't want, such as foreign-exchange risk in a global equity portfolio.

The improvement in investment results that can be achieved through risk management is sometimes called the "risk management alpha."

Risk measurement can help in several ways that are related to the allocation of *risk capital* or to hedging. The concept of risk capital is central to the *risk budgeting* approach to asset management, in which a desired risk level is chosen for a fund or portfolio, and allocated to a set of asset managers, or to the risk factors or positions to which the portfolio is exposed.

13.1 DEFINING RISK CAPITAL

The term "capital" has a number of meanings in business and economics. In economic theory, it refers to intermediate goods, natural resources that have been combined with labor and are capable of producing consumption goods or other intermediate goods. The consumption goods have utility, so capital goods have value "stored" within them. In accounting theory, capital can refer to all of the resources marshalled by a firm, including both liabilities and owners' equity, or only the latter. Owners' equity is the share of a firm's assets that belongs to the firm's owners. It is defined as a residual, once the liabilities to outside parties are accounted for.

In finance generally, "capital" sometimes refers to assets that generate income by being invested and placed at risk. In this sense, it is closer to the economists' definition. In other contexts, "capital" refers to a stock of assets that buffer a business against bankruptcy, a sense that is closer to the accountants' definition of equity. For large complex financial intermediaries, it can be difficult to distinguish between pure liabilities and owners' equity or the buffer stock, as some hybrid securities, such as preferred shares, have characteristics both of debt (fixed obligations to outsiders) and equity (residual interests of owners).

For hedge funds and other pooled investment vehicles such as mutual funds and ETFs, the notion of equity capital is unambiguous: It is the net asset value (NAV) of the funds placed with the fund's management by investors. The NAV is the amount available to be withdrawn in cash by investors, at the most recent valuation of the positions.¹

The notion of risk capital has elements of all these definitions. We define it as an asset reserve earmarked to cover the largest acceptable loss with a given confidence level: that is, a quantile of the portfolio's or firm's

¹Investors actually withdrawing funds will have to do so within the agreed restrictions placed on withdrawals by the hedge fund, and net of fees.

loss distribution. This definition is closely related to the definition of VaR as the largest loss possible at a given confidence level. Similarly to VaR, other things being equal, imposing a higher confidence level implies that a higher loss level can be tolerated. However, the notion of risk capital, like that of VaR, is not tied to any specific distributional hypothesis, or to a view that VaR can be reliably measured.

In financial firms, equity capital plays a particularly important role, because their balance sheets are part of the production process of intermediation and payment services. A financial firm may lose its ability to raise debt or be seized by regulators long before it becomes bankrupt, or its equity value reaches zero, so a purely bankruptcy-oriented definition of risk capital is not adequate. Most financial firms have far higher leverage than the typical nonfinancial firm, so capital must be carefully managed.

Risk capital is a tool for translating loss measures into equity terms. Acquiring assets puts capital at risk, all the more so when acquired by using borrowed funds. Measuring the risk to equity capital is thus a different approach to measuring risk that is directly related to profitability and the ability to continue the business.

The risk capital approach doesn't necessary take an accounting measure of equity as a starting point. It can rely on a "shadow" measure of equity attributed to a portfolio or other activity, rather than on a balance sheet quantity. It can thereby be applied to entities that don't have an individual corporate existence. The risk capital of a portfolio or an activity is the amount of capital needed to absorb a given loss. It is defined so that, even in an extreme event, the firm, strategy, or account will not lose 100 percent of its capital. A risk capital framework thus requires us to set a loss tolerance, a time horizon for the loss to occur, and a probability or confidence interval with which the loss will not be exceeded over the time horizon. A higher confidence level then requires that we hold more risk capital, so we can be more certain that it will be greater than potential losses.

"Acceptable" can mean very different things to different types of investors or intermediaries. For a nonfinancial corporation, it may mean the largest loss that falls short of triggering insolvency. For a bank or other financial intermediary, it may mean a loss that falls well short of triggering bankruptcy, seizure by bank supervisors, large-scale depositor runs, or investor redemptions. An unacceptable loss may thus be one with implications for liquidity as well as solvency (see Chapter 12). For a pension fund, it might mean the largest loss that still leaves the fund overwhelmingly likely to be able to meet its pension obligations in the future. However, all of these firms have in common that there is some level of loss that would mean the end of their business or would be in some other respect catastrophic.

The first task of risk capital measurement is thus to define that level of loss and to construct a portfolio for which the likelihood of its occurrence

is less than the confidence level. Every investor faces the task, in some form, of quantifying the threshold at which loss becomes catastrophic, and having a portfolio that has at least enough “safe” investments or is constructed in such a way as to avoid it.

Risk capital can also be defined for subportfolios and individual positions as their largest acceptable loss, as well as at the level of a firm or an entire portfolio. One of the issues we discuss in this chapter is how a total pool of risk capital can be allocated to individual investments. It hinges on how we measure the impact on the total risk capital of a single investment or subportfolio. In other words, how much of the total risk capital does it “use up”? To answer this, we need to understand the diversification impact of adding to or subtracting from the position or subportfolio.

This leads to the second key task of risk capital measurement: how the parts relate to the whole. How can we set risk capital in such a way that the risk capital allocations to individual positions, subportfolios, or activities are sensible, but also add up to the total risk capital of the firm? These are large questions, and we endeavor in this chapter to answer them in a small way, within the framework of the standard market risk model developed in Chapters 2 through 5.

13.2 RISK CONTRIBUTIONS

We begin with the second task, finding how much of the total risk in a portfolio is contributed by one position. For example, if a portfolio manager is told to reduce his risk, he needs to know how much risk, if any, he can take off by unwinding a given amount of a specific position. In some cases, risk at the portfolio level could actually be increased by the position reduction. Similarly, an asset manager may want to know how and how much the risk of the portfolio will change if he changes the allocation.

We’ll define *risk contributions*, starting with the simplest case, a long-only asset manager who allocates assets to two types of security, stocks and bonds. There are a number of ways to define a risk contribution, depending on which standard-model risk measure we use, for example, the portfolio variance, volatility, or VaR. It also depends on how we vary the exposure: Do we increase it by a little bit, or do we remove it entirely from the portfolio?

Two useful formal concepts here are:

- **Incremental VaR**, the contribution of an entire position to the total VaR: How much does the portfolio risk change if I add or unwind an entire position?
- **Marginal VaR**, the derivative of the VaR with respect to the size of a position; what happens to risk if I increase or reduce a position by a small amount?

13.2.1 Risk Contributions in a Long-Only Portfolio

In a two-asset allocation problem, assuming asset-class returns are jointly normally distributed, the variance of portfolio returns as a percent of initial market value is given by

$$\sigma_p^2 = \omega_1^2 \sigma_1^2 + \omega_2^2 \sigma_2^2 + 2\omega_1 \omega_2 \sigma_1 \sigma_2 \rho \quad (13.1)$$

The portfolio shares invested in stocks ($m = 1$) and bonds ($m = 2$) are denoted ω_m , with $\omega_1 + \omega_2 = 1$. Stock and bond returns have volatilities σ_1 and σ_2 , and a correlation ρ , and the portfolio volatility is σ_p . We assume that only long positions are taken, so $\omega_m \geq 0$, $m = 1, 2$. This expression should be familiar from our discussion of diversification in Chapter 2 and of portfolio VaR in Chapter 5. The τ -period VaR at a confidence level of α is given by

$$\text{VaR}(\alpha, \tau) = -z_* \sqrt{\tau} \sigma_p$$

where z_* is the ordinate of $N(0, 1)$ corresponding to the selected confidence level α . This expression employs the arithmetic approximation of the delta-normal approach, as in Equation (3.3) or (5.1). For the rest of this section, we simplify the notation by setting the time horizon to one year, so the square-root-of-time term drops out.

The *marginal variance* is found by differentiating Equation (13.1) with respect to the ω_m :

$$\frac{\partial \sigma_p^2}{\partial \omega_m} = 2\omega_m \sigma_m^2 + 2\omega_n \sigma_1 \sigma_2 \rho \quad m \neq n; m, n = 1, 2$$

We can find the change in portfolio volatility by applying the chain rule of calculus. With f and g proper functions of a (possibly) vector-valued argument x , the chain rule states

$$\frac{d}{dx} f[g(x)] = f'(g(x))g'(x)$$

Setting

$$g(x) = \sigma_p^2$$

a function of $\omega = (\omega_1, \omega_2)$, and

$$f(g) = \sqrt{g} \quad \Rightarrow \quad f'(g) = \frac{1}{2} \frac{1}{\sqrt{g}}$$

In other words, the volatility is treated as a function of the variance in computing the derivative, to get the volatility contribution defined as the

marginal volatility, the contribution of position m to the standard deviation of portfolio returns, measured in dollars:

$$\begin{aligned}\frac{\partial \sigma_p}{\partial \omega_m} &= \frac{1}{2} \frac{1}{\sqrt{\sigma_p^2}} \frac{\partial \sigma_p^2}{\partial \omega_i} \\ &= \frac{1}{\sigma_p} (\omega_m \sigma_m^2 + \omega_n \sigma_1 \sigma_2 \rho) \quad m \neq n; m, n = 1, 2\end{aligned}$$

So far, we have defined risk contribution as the marginal impact on portfolio variance or volatility. We can however, also define a marginal contribution to VaR by simply taking a different quantile than the one-sigma quantile. The *marginal VaR* is:

$$\frac{\partial \text{VaR}(\alpha, \tau)}{\partial \omega_m} = -z_* \frac{1}{\sigma_p} (\omega_m \sigma_m^2 + \omega_n \sigma_1 \sigma_2 \rho) \quad m \neq n; m, n = 1, 2$$

The marginal VaR or volatility is the rate at which an increase in allocation to position m increases VaR or volatility. The analogues to the risk contribution metrics presented here can all be computed for expected short-fall as well as for variance, volatility, and VaR. These quantities can also be computed in a simulation as opposed to an analytic/algebraic framework.

The marginal risk contributions have a very useful property, which we will exploit in our discussion of risk capital. As we saw in Chapter 11, VaR is a *homogeneous function* of the investment amounts in the portfolio. As it happens, the portfolio variance and volatility are homogeneous functions of the allocations $\{\omega_2, \omega_1\}$, too.²

Specifically, the portfolio variance is homogeneous of degree 2, and the volatility and VaR are linearly homogeneous; if all the allocations double, that is, the total investment rises from x to $2x$, and the fractions allocated to stocks and bonds are not altered, the variance will quadruple, while the dollar volatility and the VaR will double. The marginal variance, volatility, and VaR therefore have the convenient *Euler property*:

- The sum of the marginal variances of the bond and stock allocations, each multiplied by its allocation, is equal to twice the portfolio variance:

$$\begin{aligned}\sum_{1,2} \omega_m \frac{\partial \sigma_p^2}{\partial \omega_m} &= 2 \left(\sum_{1,2} \omega_m^2 \sigma_m^2 + 2\omega_m \omega_n \sigma_1 \sigma_2 \rho \right) \\ &= 2\sigma_p^2\end{aligned}$$

²Appendix A.6 provides formal statements and proofs of these properties.

- The sum of the marginal volatilities of the bond and stock allocations, each multiplied by its allocation, is equal to the portfolio volatility:

$$\begin{aligned} \sum_{1,2} \omega_m \frac{\partial \sigma_p}{\partial \omega_m} &= \frac{1}{\sigma_p} \sum_{1,2} \omega_m (\omega_m \sigma_m^2 + \omega_n \sigma_1 \sigma_2 \rho) \\ &= \frac{1}{\sigma_p} (\omega_1^2 \sigma_1^2 + \omega_2^2 \sigma_2^2 + 2\omega_1 \omega_2 \sigma_1 \sigma_2 \rho) \\ &= \sigma_p \end{aligned}$$

These allocation-weighted risk contributions are sometimes called the *component* variances, volatilities, and VaR. The component VaR or volatility turns the marginal rate per additional dollar or percent allocated into an amount in dollars.

- The sum of the marginal VaRs of the bond and stock allocations is equal to the portfolio VaR, $z_* \sigma_p$. This is obvious, since the volatility contributions are just the marginal VaRs for $\alpha = 0.8413$, the confidence level at which $z_* = -1.0$.

A final property of the risk contributions can be seen by dividing each component risk contribution by the total risk to get the shares of each risk contribution to the total risk, as measured by variance, volatility, or VaR. The sum of these for the variance is 2, while the sum for the volatility and VaR is 1. We can see this by simply dividing the expressions above through by the respective total risk measure:

$$\begin{aligned} \frac{1}{\sigma_p^2} \sum_{1,2} \omega_m \frac{\partial \sigma_p^2}{\partial \omega_m} &= 2 && \text{(variance)} \\ \frac{1}{\sigma_p} \sum_{1,2} \omega_m \frac{\partial \sigma_p}{\partial \omega_m} &= 1 && \text{(volatility)} \\ \frac{1}{\text{VaR}(\alpha, \tau)} \sum_{1,2} \omega_m \frac{\partial \text{VaR}(\alpha, \tau)}{\partial \omega_m} &= 1 && \text{(VaR)} \end{aligned}$$

In economics, these shares would be called *elasticities*, since they are the product of a mathematical first derivative times the ratio of the function argument with respect to which it being differentiated. The volatility and

VaR elasticities are identical. The fact that they sum to 100 percent is quite useful.³

The following example shows how these measures of risk contribution are calculated.

Example 13.1 (Risk Contributions in a Long-Only Portfolio) Let the stock and bond return volatilities be $\sigma = (\sigma_1, \sigma_2) = (0.10, 0.18)$, and let the correlation between the two asset classes be $\rho = 0.40$. The return covariance matrix, as in Chapter 5, is then

$$\text{diag} \begin{pmatrix} 0.10 \\ 0.18 \end{pmatrix} \begin{pmatrix} 1 & 0.40 \\ 0.40 & 1 \end{pmatrix} \text{diag} \begin{pmatrix} 0.10 \\ 0.18 \end{pmatrix} = \begin{pmatrix} 0.0100 & 0.0072 \\ 0.0072 & 0.0324 \end{pmatrix}$$

with variances and covariances expressed at an annual rate as decimals. Thus the covariance between stock and bond returns is 72 bps per annum.

Suppose the portfolio has a market value of \$300, of which \$100 is invested in bonds and \$200 in stocks. The portfolio variance is then \$1,684 and the volatility is \$41.04.

Next, let's compute the VaR and marginal VaRs of the portfolio. The annual VaR at a 99 percent confidence level is \$95.47. The risk contributions are

	Bonds	Stocks	Total
<i>Variance</i>			
Portfolio variance			1,684.00
Marginal variance	4.88	14.40	
Component variance	488.00	2,880.00	3,368.00
Variance elasticities	0.29	1.71	2.00
<i>Volatility</i>			
Portfolio volatility			41.04
Marginal volatility	0.06	0.18	
Component volatility	5.95	35.09	41.04
Volatility elasticities	0.14	0.86	1.00
<i>VaR</i>			
Portfolio VaR			95.47
Marginal VaR	0.14	0.41	
Component VaR	13.83	81.63	95.47
Variance elasticities	0.14	0.86	1.00

³We run into some terminological confusion here. The RiskMetrics documentation reverses the definitions, so that incremental VaR is the continuous and marginal VaR, the discrete amount concept. The use of "marginal" to describe economic concepts that have differential calculus definitions, such as marginal utility and marginal product, is ingrained in economics, though, and preferable. RiskMetrics labels the *component*, not marginal VaR, as "incremental VaR."

Note that although bonds are $\frac{1}{3}$ of the portfolio by market value, at the margin they contribute only 14.5 percent of the volatility.

The volatility contributions of the two asset classes are very sensitive to all of the parameters—the volatilities, correlation, and relative sizes of the stock and bond allocations. Let’s look at these sensitivities a bit more closely. Figure 13.1 illustrates a few key points, which were foreshadowed in Chapter 2’s discussion of diversification. Each panel shows how risk contributions change with correlation for a given allocation to stocks and bonds.

- When the correlation between the two asset classes are relatively high, the total portfolio volatility is higher, and each asset class tends to be more of a risk amplifier than a risk mitigator. When the correlation is low or negative, the total portfolio volatility is lower, and the risk contributions of the two asset classes diverge more and more. In the baseline Example 13.1, the bond allocation had a positive risk contribution. But as seen in the upper panel of Figure 13.1, if the correlation is low enough, bonds can become a “diversifier,” that is, have a negative risk contribution, without any change in allocation or in the stock or bond return volatilities. This is related to the result we saw in Figure 2.10 of Chapter 2: Correlations don’t have to be negative for a diversification benefit to be present.
- The relative allocations matter a lot. The lower panel of Figure 13.1 shows how the risk contributions vary with correlation if we shift the allocation towards bonds. With a larger bond allocation, bonds never become a diversifier. At a very low correlation near -1 , and with the lower allocation, stocks “overcome” their high volatility and become a diversifier.
- Size and volatility work together. Notice, in the lower panel of Figure 13.1, that for correlations that are not very low, say above -0.4 , the bond and stock allocations have about the same volatility contribution. The large bond allocation, with its lower volatility, has the same impact on portfolio volatility as the smaller stock allocation with its high volatility. In fact, there is a trade-off curve, shown in Figure 13.2, along which we can increase volatility and decrease allocation without changing the risk contribution of the asset class.

13.2.2 Risk Contributions Using Delta Equivalents

We next want to generalize in two ways, by extending the analysis to more than one asset, and by introducing negative allocations, that is short positions with respect to a risk factor. Let’s start with a simple extension




FIGURE 13.1 Risk Contributions in a Long-Only Strategy
 Volatility contributions of the stock and bond allocations as correlation from -1 (complete diversification) to $+1$ (no diversification). The values on the y-axis are the component volatilities for stocks and bonds and the portfolio volatility. The stock and bond volatilities are $(\sigma_1, \sigma_2) = (0.10, 0.18)$. The market value of the portfolio is \$300.

Upper panel: Allocations \$100 to bonds and \$200 to stocks.

Lower panel: Allocations \$200 to bonds and \$100 to stocks.

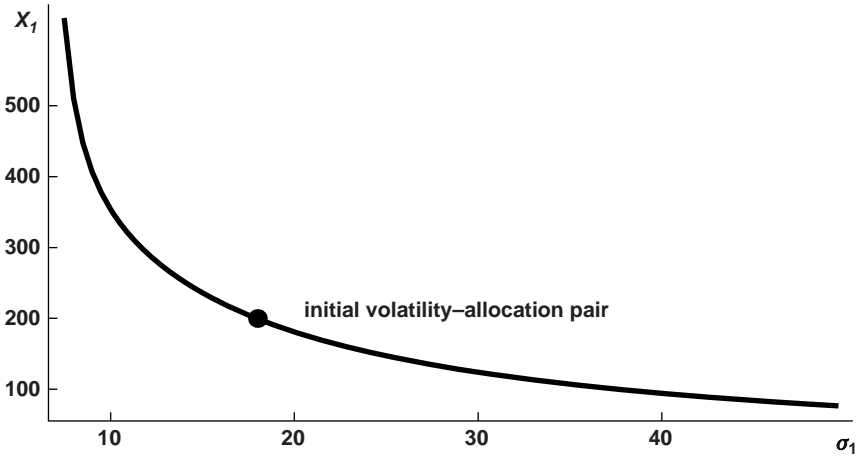


FIGURE 13.2 Allocation, Volatility, and Constant Risk Contribution
 The plot shows how much the stock allocation needs to increase as the volatility decreases in order to keep the risk contribution constant. Note that as the volatility becomes quite low, the allocation must increase at an increasing rate per volatility point. Note also that, at a stock allocation greater than 300, we must have a short position in bonds.

of the previous long-only example. Imagine a portfolio consisting entirely of stocks, but long some stocks and short others. Hedge fund managers following an *equity market-neutral* or *long-short strategy*, rather than traditional asset managers, might have such a portfolio. We’ll elaborate with an example.

Example 13.2 (Volatility Contribution in a Long-Short Strategy) We now consider an allocation to two different stock portfolios. The long stock portfolio has a market value d_l , and the long stock portfolio has a market value d_s , with $(d_l, d_s) = (500, -500)$, measured in millions of dollars. The two subportfolios have identical return volatilities $\sigma = 0.20$, and the correlation between the two subportfolios returns is high, $\rho = 0.75$. We would expect such a high correlation, unless the stocks have exceptionally low betas, because most of the risk in both subportfolios is systematic stock market risk.

The return covariance matrix is then

$$\text{diag} \begin{pmatrix} 0.20 \\ 0.20 \end{pmatrix} \begin{pmatrix} 1 & 0.75 \\ 0.75 & 1 \end{pmatrix} \text{diag} \begin{pmatrix} 0.20 \\ 0.20 \end{pmatrix} = \begin{pmatrix} 0.04 & 0.03 \\ 0.03 & 0.04 \end{pmatrix}$$

with variances and covariances expressed at an annual rate as decimals.

The risk contributions are

	Longs	Shorts	Total
<i>Variance</i>			
Portfolio variance			5,000.00
Marginal variance	10.00	-10.00	
Component variance	5,000.00	5,000.00	10,000.00
Variance elasticities	1.00	1.00	2.00
<i>Volatility</i>			
Portfolio volatility			70.71
Marginal volatility	0.07	-0.07	
Component volatility	35.36	35.36	70.71
Volatility elasticities	0.50	0.50	1.00
<i>VaR</i>			
Portfolio VaR			164.50
Marginal VaR	0.16	-0.16	
Component VaR	82.25	82.25	164.50
Variance elasticities	0.50	0.50	1.00

In this example, the short stock portfolio has a negative marginal contribution, equal in magnitude to the positive contribution of the long portfolio. It is a diversifier, so increasing the short portfolio by a small amount would, at the margin, offset the risk impact of an increase in the long portfolio. But both the longs and the shorts make an equal and positive contribution to risk, as seen in the component contributions. The negative marginal contribution is multiplied by a negative portfolio allocation to get a positive contribution for the short portfolio. That guarantees that growing the entire portfolio, while keeping equal long and short shares, would grow the risk proportionally.

We can generalize this to a larger portfolio with many positive and negative delta equivalents. From here on, let's focus exclusively on the marginal VaRs. As we know by now, the marginal volatilities are closely related to them, while the marginal variances are, in the end, less interesting because their elasticities add up to 200 rather than 100 percent of the portfolio variance.

To express the marginal VaR formally for a portfolio, we need the following rule from matrix algebra. For any proper matrix Σ and conformable vector \mathbf{d} ,

$$\frac{\partial}{\partial \mathbf{d}} \mathbf{d}' \Sigma \mathbf{d} = 2 \Sigma \mathbf{d} = 2 \begin{pmatrix} \sum_m^M d_m \sigma_{1m} \\ \sum_m^M d_m \sigma_{2m} \\ \vdots \\ \sum_m^M d_m \sigma_{Nm} \end{pmatrix}$$

Using the chain rule, we get

$$\frac{\partial}{\partial \mathbf{d}} \sqrt{\mathbf{d}' \Sigma \mathbf{d}} = \frac{\Sigma \mathbf{d}}{\sqrt{\mathbf{d}' \Sigma \mathbf{d}}}$$

We can now use these facts from matrix algebra to define the important marginal concepts and derive some of their properties. The derivative of the variance with respect to position m is defined as

$$\frac{\partial}{\partial d_m} \mathbf{d}' \Sigma \mathbf{d}$$

that is, the increase in variance resulting from an increase in the delta equivalent of position m by \$1.

The marginal VaR is

$$\begin{aligned} \text{MVaR}_{mt}(\alpha, \tau) &\equiv d_m \frac{\partial}{\partial d_m} \text{VaR}_t(\alpha, \tau) \\ &= -z_* \sqrt{\tau} d_m \frac{\partial}{\partial d_m} \sqrt{\mathbf{d}' \Sigma \mathbf{d}} \\ &= -z_* \sqrt{\tau} d_m \frac{\Sigma \mathbf{d}}{\sqrt{\mathbf{d}' \Sigma \mathbf{d}}} \end{aligned}$$

The vector of marginal VaRs is then

$$\text{MVaR}_t(\alpha, \tau) = -z_* \sqrt{\tau} \frac{\mathbf{d}' \Sigma \mathbf{d}}{\sqrt{\mathbf{d}' \Sigma \mathbf{d}}}$$

We can now confirm the Euler property—the sum of the marginal VaRs is equal to the VaR—for general portfolios:

$$\begin{aligned}
 \sum_m \text{MVaR}_{mt}(\alpha, \tau) &= -z_* \sqrt{\tau} \sum_m d_m \frac{\Sigma \mathbf{d}}{\sqrt{\mathbf{d}' \Sigma \mathbf{d}}} \\
 &= \frac{z_*^2 \tau \sum_m d_m \Sigma \delta}{\text{VaR}_t(\alpha, \tau)} \\
 &= \frac{z_*^2 \tau \mathbf{d}' \Sigma \mathbf{d}}{\text{VaR}_t(\alpha, \tau)} \\
 &= \frac{\text{VaR}_t(\alpha, \tau)^2}{\text{VaR}_t(\alpha, \tau)} \\
 &= \text{VaR}_t(\alpha, \tau)
 \end{aligned}$$

The *incremental VaR* is computed as the difference between the VaR with and without a particular position or subportfolio:

$$\text{IVaR}_{mt}(\alpha, \tau) \equiv \text{VaR}_t(\alpha, \tau)(\text{entire portfolio}) - \text{VaR}_t(\alpha, \tau)(\text{partial portfolio})$$

Both measures have in common that a position that reduces risk by either diversifying or hedging the portfolio will have a low or even negative marginal and/or incremental VaR.

13.2.3 Risk Capital Measurement for Quantitative Strategies

For some investments and some investment companies, especially hedge funds, it can be difficult to define the size or “amount” of the position, or the level of the activity, in a way that accommodates the VaR contribution calculations we have just outlined. We first describe a few examples of strategies for which the delta equivalents of the positions in the portfolio at a point in time don’t capture the exposure sizes appropriately. We can, however, use risk capital calculations to provide an alternative way to measure the level of activity of these strategies and capture the size at which they are being run.

Examples of such strategies include statistical arbitrage, equity and credit correlation trading, and gamma trading and other systematic option strategies. Such strategies are typically *market-neutral*, meaning there are both long and short positions, intended to eliminate net exposure to the stock market or credit spreads, or that the portfolio has zero net market value, that is, the market value of the long positions is close to that of the short positions. Each position in such strategies may have a

well-defined delta, but returns on the strategy as a whole have a very different distribution—in particular, the potential losses may be far greater—than a delta-based VaR would predict.

These strategies are generally leveraged; the financing is provided by counterparties or by an exchange clearing house. The strategies often in addition have embedded leverage, and involve short positions. The cash investment in such strategies is generally used to put up margin requirements and may have little relation to the volatility of returns relative to the cash invested. If the cash requirements are low, a prudent portfolio manager may hold a cash reserve, as a risk management tool, in order to increase his equity in the strategy. Risk capital calculations can be a useful tool to determine the appropriate level of such additional cash buffers. The additional cash buffer is then set so that the total equity against the strategy is adequate risk capital.

We have briefly encountered some of these strategies in Chapter 1, in describing the growth of hedge funds and other large capital pools, in Chapter 10, in describing correlation trading, and in Chapter 12, describing the range of techniques for increasing leverage in trading strategies. To help understand the example we'll provide in a moment, let's further describe some of these quantitative strategies:

Statistical arbitrage attempts to profit from transitory differences between actual prices and the fair market or forecast value according to a model. The word “arbitrage” is of course a misnomer, since market prices may not converge to their model values quickly or at all. There are two general orientations for such models, which are typically applied in equity and futures markets. They may rely on price forecasts from a fundamental-data model extrapolation from time-series, or technical trading models, to arrive at short-term asset price forecasts. These strategies rely on or “use” market liquidity, that is, the ability to buy and sell without materially affecting prices.

Or they may attempt to detect and exploit transitory imbalances in order flow, acting as liquidity providers at a very granular level of the price discovery process. Such models, called *high-frequency* or *algorithmic trading*, are also most frequently applied to assets traded on organized exchanges, such as equities and futures, where there is strong, but fluctuating, two-way order flow in most stocks. High-frequency trading uses computing power to rapidly detect these imbalances and enter orders on the “other side.” If not filled, the orders are rapidly withdrawn. While generally acting as liquidity providers, the strategies have come under scrutiny following the May 6, 2010 “flash crash,” in which stock prices

experienced extremely high intraday volatility. Some observers held high frequency traders responsible for the volatility by placing large automated sell orders or by withdrawing entirely from trading.

The trading frequency of such strategies varies very widely, from fractions of a second to several weeks, as measured by the average length of time a dollar's worth of stock resides in the portfolio. The portfolios are typically adjusted so that the net market exposure is close to zero. For this reason, the cash investment can be relatively small. For equity portfolios, the cash requirements will be greater than for exchange-traded derivatives strategies.

Gamma trading attempts to exploit the tendency of option implied volatilities to exceed realized volatility over the life of the option. The trader sells an option, typically at-the-money with a short time to maturity, and delta hedges. The option hedge is adjusted frequently, at least daily. If the underlying asset price moves significantly less than predicted by the implied volatility, the time decay on the short option position will exceed the trading costs, and the effect of the asset price move on the option value will be offset nearly exactly by that on the asset position itself. Such strategies can be carried out for a variety of asset types, such as currencies, equities, money markets, and fixed-income securities, and in a number of ways, such as exchange-traded futures and options, OTC options, or variance swaps. The cash requirements of the strategy can be quite low.

Convertible bond trading is a long version of the gamma trading strategy. It can be profitable when the options embedded in the bonds are cheap relative to exchange-traded and OTC options on the underlying stock, and realized volatility is high. Convertible bond trading is also typically levered, but not extremely so.

Even if the risk contributions of other strategies the hedge fund pursues would otherwise be straightforward, the presence of one or more of these quantitative strategies will make it hard to compute them, since all the risk contributions depend on all of the allocations in the portfolio. This makes sizing any of the strategies more difficult.

Typically, however, portfolio managers work with some measure of the size or level at which the strategy is being carried out. There are a few such rule-of-thumb approaches to measuring the size of the strategy. The first is to measure the activity level of the strategy by the amount of cash used, all of which is generally devoted to margin requirements with exchanges, intermediaries, or counterparties. As noted, since some of these strategies can be levered up quite a bit, and also may have considerable embedded leverage, this is potentially quite misleading. A strategy may appear "small" measured

by the cash it employs, but very large measured by its potential losses, implying that an appropriate risk capital is higher than the rule of thumb states.

The second frequently encountered rule-of-thumb is to use some measure of long market exposure. For example, the size at which a credit correlation strategy is run may be stated as the notional amount or market value on each side, or as a delta or spread⁰¹ on each side, neither of which is closely related to potential loss.

Especially for strategies that can be highly leveraged, the size of the strategy may be more accurately measured by the risk capital the trader imputes to the strategy. In effect, this approach sets an equity measure for the strategy that is distinct from the out-of-pocket cost of implementation. The trader implicitly sets aside a cash reserve in addition to the required margin in order to run the strategy at a prudent level, given the amount of capital he has. This is a means of reducing the probability of “bankrupting” the strategy, or of forcing it to be unwound immediately. The trader can meet margin calls out of the cash reserve rather than being forced out of positions immediately at a loss. Another strategy might have cash funding requirements that are high relative to the risk of the strategy. It can then be viewed as freeing up risk capital that is then available for other strategies.

Once measures of each quantitative strategy’s size are established, they can be treated as deltas in a portfolio of assets with correlated returns. The return volatility and covariances with respect to other strategies can be computed using the dollar returns and the risk capital measure. An important assumption in this approach is that the level of the fund’s own activity does not affect returns in the market via adverse price impact or other market liquidity effects. For example, if the fund doubles the level at which it operates the strategy, we assume it doesn’t lower the rate of return by causing the arbitrage opportunity to disappear or haircuts to increase. This is consistent with the assumption typically underpinning these strategies, that market and funding liquidity are adequate to support them. The subprime crisis revealed how unreliable this assumption can be.

In this approach, just scaling the strategy up or down won’t affect risk measurement. To see this, imagine doubling the measured risk capital of one strategy arbitrarily. This means that the return volatility and all of the covariances with that strategy must be halved, since the dollar returns haven’t changed. This has no effect on either the portfolio volatility or the marginal risk contributions.

In the rest of this section, we set out an example of this approach, showing how risk capital computations can be used to appropriately size and assess capital charges for quantitative strategies. We imagine a hedge fund operating two quantitative strategies, which we’ll call gamma trading and




FIGURE 13.3 Simulated Hedge Fund Strategy Returns

The plot shows the P&L, in millions of dollars, for the two strategies. The gamma trading dollar returns are marked by small circles, and the stat arb returns by x's.

statistical arbitrage for concreteness. The fund is able to accurately estimate the strategies' expected returns, and return variances and covariances. It may do so by actually running the strategies at a specified size and observing daily P&L over some period, or it may have backtested the models on historical data, so the P&L time series is based on "paper trading." An example of what such results might look like for one year is displayed in Figure 13.3. For concreteness, we assume that the strategies' sizes are initially measured by the amount of cash they require, one of the sizing rules-of-thumb mentioned above, so returns are measured as each day's dollar P&L divided by the cash employed in the strategy. For the gamma strategy, the margin is \$50, and for stat arb, \$200.

Based on these results of the backtest, the fund estimates the returns and volatilities of the two strategies as

	Gamma ($m = 1$)	Stat arb ($m = 2$)
Mean annual excess return (μ_m , %)	15.0	10.0
Annual volatility (σ_m , %)	75.0	35.0

The correlation is estimated as 0.40.

We also assume that the hedge fund operates no other strategies, and that there are no investor redemptions and subscriptions, so the fund NAV is affected only by P&L. The gross dollar return to the fund's investors equals the sum of the dollar returns or P&L of the two strategies, and the rate of

return is equal to the dollar return divided by the NAV of the fund. We ignore the “level-changing” effect of the fund’s P&L, which reduces returns when returns have been positive, and vice versa.

Next, suppose that the fund manager wishes to increase the risk capital supporting each strategy by holding a cash reserve against it, over and above the cash required for financing. We look at this first from the point of view of the individual strategies, as though each were run in a dedicated single-strategy fund of its own, and then from the point of view of the fund as a whole. Essentially, the manager is reducing the VaR or the volatility of each strategy to a target level by setting aside extra cash. Framed in terms of volatility, the portfolio manager might want to reduce the annual return volatility to say, 25 percent. To reduce the probability of exhausting equity and “bankrupting” the strategy. The additional cash is used as additional equity, not as a liquidity reserve.

Denote the *share* of the additional cash in the total risk capital of strategy m by χ_m , $m = 1, 2$. For a target volatility of, say, 0.25, χ_m is

$$\chi_m = 1 - \frac{0.25}{\sigma_m} \quad m = 1, 2$$

If, for example, the target volatility is half the strategy’s actual return volatility, the risk capital doubles, since the appropriate additional cash buffer is then half the total risk capital. If the target volatility happens to equal the strategy’s volatility, then no additional risk capital is set aside. In our example, for each dollar of required margin, setting aside additional equity capital, held in a liquid, risk-free form, of \$2.00 for the gamma trading strategy and \$0.40 for statistical arbitrage, will reduce the volatility of each strategy to 25 percent. The dollar amount of additional cash required to achieve this suppression of volatility is $\frac{\chi_m}{1-\chi_m}$ times the required margin. Returns are reduced by the same proportion as volatility, to $\mu_m(1 - \chi_m)$. (We assume the return on the additional cash is zero.) In our example, we have:

	Gamma	Stat arb
χ_m	0.667	0.286
Additional risk capital (\$)	100	80
Total risk capital (\$)	150	280
Return on risk capital ($\mu_m(1 - \chi_m)$, %)	5.00	7.14
Volatility on risk capital ($\sigma_m(1 - \chi_m)$, %)	25.0	25.0

Returns are diluted in the same proportion as the volatility.

So far, we have viewed each strategy in isolation. That is, these results are applied to each strategy as if it were a standalone fund. Next, we use our definition of component VaR (or volatility) and VaR elasticities to attribute

the fund's NAV to the two strategies, that is, to determine the amount of risk capital each strategy is using. Imagine investors have placed \$500 with the fund; that is, it starts with an NAV of \$500. Initially, the fund manager operates the two strategies using no additional cash buffer, and allocates 20 percent of the fund to gamma and 80 percent to stat arb. Applying the algebra developed earlier in this chapter, the risk contributions are

	Gamma	Stat arb	Fund
<i>Variance</i> (\$, ann.)			
Portfolio variance			33,625.00
Marginal variance	196.50	119.00	
Component variance	19,650.00	47,600.00	67,250.00
Variance elasticities	58.44	141.56	200.00
<i>Volatility</i> (%)			
Portfolio volatility (% ann.)			36.67
Marginal volatility	53.58	32.45	
Component volatility	10.72	25.96	36.67
Volatility elasticities	29.22	70.78	100.00
<i>VaR</i>			
Portfolio VaR (\$, $\tau = \frac{1}{12}$, $\alpha = 0.99$)			123.14
Marginal VaR	0.36	0.22	
Component VaR	35.98	87.16	123.14
VaR elasticities	29.22	70.78	100.00

Although the gamma strategy uses only $\frac{1}{5}$ of the cash invested, its risk contribution is nearly 30 percent.

We can use risk contributions to make a risk-based allocation of the fund's resources. Suppose the fund manager wants to allocate 25 percent of its risk capital to the gamma strategy and 75 percent to statistical arbitrage. Further, it has a target volatility of 25 percent per annum for the fund as a whole, which is exceeded by its return volatility of 36.67 with the current allocation of 20 percent gamma, 80 percent stat arb, and no cash buffer. The fund needs to determine the level at which to operate each strategy and the total additional cash reserve needed to achieve its target volatility and risk capital allocation. It can then attribute the additional cash as risk capital to the two strategies. By solving this problem, the fund manager links the risk tolerance of the fund, expressed through its target volatility, and the strategy size or allocation decision.

To solve this problem, we first find the "cash allocation" shares: ω_1 for the gamma trading strategy, ω_2 for stat arb, and χ the cash buffer, expressed as fractions of the fund's initial \$500 NAV. The ω_m are the shares

of the fund’s initial cash resources used to fund the strategies and are equal to the cash margin required for each, divided by the initial NAV. These shares plus the cash buffer share must sum to unity. Once the cash buffer χ required to achieve the target volatility is determined, it is allocated to the two strategies consistently with the total risk capital allocation. This second step is carried out within the fund; it is an internal reckoning, not a transaction with the market.

The desired risk capital allocation can be most simply expressed using the volatility elasticities

$$\omega_m \frac{\omega_m \sigma_m^2 + \omega_n \sigma_1 \sigma_2 \rho}{\omega' \Sigma \omega} \quad m, n = 1, 2, m \neq n$$

where $\omega' = (\omega_1, \omega_2)$. The fund manager must solve the following equations for the ω_m and χ :

$$\begin{aligned} \omega_1 + \omega_2 + \chi &= 1 \\ \omega_1 \frac{\omega_1 \sigma_1^2 + \omega_2 \sigma_1 \sigma_2 \rho}{\omega' \Sigma \omega} &= 0.25 \\ \sigma_p &= \sqrt{\omega' \Sigma \omega} = 0.25 \end{aligned} \tag{13.2}$$

The elasticities and the portfolio volatility are nonlinear functions of the ω_m and of the volatility and correlation parameters, so a numerical search procedure must be used to obtain the solution.⁴

The first condition in the optimization problem (13.2) is the “budget constraint,” and limits the manager to the funds placed by investors. The second condition expresses the fund manager’s decision to set the risk capital allocation of the gamma strategy to 25 percent. No additional condition is required to impose the risk capital allocation of stat arb to 75 percent, because the volatility elasticities, like the cash allocation shares, sum to unity. The third and last condition sets the fund’s return volatility to 25 percent, its overall risk limit.

The results are

	Gamma	Stat arb	Cash buffer
Allocation of cash (%)	12.18	56.86	30.96
Required margin (\$)	60.92	284.29	154.79
Attribution of risk capital (\$)	125.00	375.00	
Attribution of cash buffer (\$)	64.08	90.71	

⁴The results in the examples were computed using *Mathematica*’s FindRoot algorithm.

Although the fund maintains a cash reserve of nearly \$155, a bit more than 30 percent of the funds placed by investors, it is “fully invested” from a risk standpoint, since it is at its desired volatility ceiling. Of this additional cash, \$64.08 is attributed to the gamma and \$90.71 to the stat arb strategy. In contrast to the tripling of the capital attributed to the gamma strategy as a standalone strategy, as part of this portfolio, the gamma strategy’s risk capital is merely doubled. The reason is that its marginal volatility is already reduced by the diversification benefit it enjoys as a relatively small allocation within the fund. The required margin is the capital actually invested in the strategy; additional cash is attributed from the reserve to fill out the risk capital. Expected fund returns are 37.57 percent.

This approach also provides a basis on which to assess traders a risk capital charge. It would be based on the risk capital allocations of \$125 and \$375, rather than on the amounts of cash actually used in the strategies.

In this example, both strategies have risk capital that is higher than the required margin. If we relax the target volatility constraint to 35 percent, setting $\sigma_p = 0.35$ in the numerical search problem that solves the set of equations (13.2), we get a contrasting result. The results are now

	Gamma	Stat arb	Cash buffer
Allocation of cash (%)	17.06	79.60	3.34
Required margin (\$)	85.29	398.00	16.71
Attribution of risk capital (\$)	125.00	375.00	
Attribution of cash buffer (\$)	39.71	−23.00	

Expected returns, of course, are higher at 52.59 percent. The additional cash reserve is much lower, \$16.71. The interesting feature in this example is that the additional risk capital of \$39.71 is funded mainly by the excess capital in the stat arb strategy. Stat arb is less highly levered in the sense that a given cash outlay provides a less volatile return, so it frees up risk capital for gamma when a higher level of risk is desired. It would be assessed a capital charge corresponding to less cash than it actually uses.

Note also that the risk capital allocation works here because of the low required cash margin of the gamma strategy. If the gamma strategy used cash equal to its required risk capital, it couldn’t be ramped up to attain a higher risk/volatility target. If the market doesn’t provide enough “slack” in its margin requirements, or if the manager sets the fund’s risk tolerance high enough, he fund manager can’t use the discrepancies between market and model risk assessments to increase leverage.

This approach can be extended to any number of strategies to determine the level at which to operate each strategy consistently with the desired risk allocation. The examples focused on quantitative strategies; the disparity between the cash investment made in quantitative strategies and their actual risks tends to be large, making a risk capital approach to allocation particularly apt. But it can be applied as readily to conventional strategies, and provides an example of how risk analysis can be integrated into the investment process.

The approach sketched here has important limitations. It is based on portfolio and marginal volatility. But as we have seen, volatility as a measure of risk is inadequate if returns don't follow the standard model. But our example is suggestive of the ways in which quantitative risk data can aid investors in achieving a risk and return profile that is closer to their objectives.

In the quantitative strategy example above, and in the risk budgeting approach generally, the fund manager may assess a capital charge for the risk capital allocated. This is sometimes called a *transfer pricing* problem, as it shares characteristics with the problem of pricing transfers of goods within a nonfinancial firm that don't pass through a market pricing process. Among the reasons for risk capital charges is to create incentives for portfolio managers to identify risk-minimizing investment approaches, and as a part of compensation mechanisms in which portfolio managers are rewarded for excess returns only, rather than for the absolute level of returns.

Capital charges are generally based on an estimate of the cost of capital. There are two bases for such cost calculations:

1. The cost of equity capital is the minimum expected return on equity that will induce investors to put their own capital behind an investment. That is, the cost of capital is the hurdle rate required for the fund to retain capital. Because equity is the riskiest position in the capital structure, expected returns on equity generally must be high to induce investment.
2. The cost of debt financing is the blended rate paid for different types of debt. It depends on interest rates on long- and short-term, and secured and unsecured debt. A firm or investor can make choices among these sources of funding to arrive at a funding blend that works best for it. The cost of debt is lower than that of equity, so a blended cost of financing will be lower than an equity hurdle rate. Capital charges based on borrowing costs can be viewed as a "cost-recovery" approach.

13.3 STRESS TESTING

Stress testing is an approach to risk measurement that attempts to grapple with the difficulties of modeling returns statistically. It posits the size of

shocks directly, or with a “light” modeling apparatus. It is also sometimes called *scenario analysis*, since it asks the question, “what happens to our firm (or portfolio) if the following asset price changes take place.” In spite—or because—of its light modeling infrastructure, stress testing has become a very widespread approach to risk measurement. While only indirectly related to internal risk management by banks, an important example were the supervisory stress testing procedures carried out in the United States in early 2009 and in Europe in mid-2010. The U.S. exercise, repeated in 2011 and to be conducted regularly thereafter, helped authorities determine which banks would be obliged to raise additional equity capital or prevented from paying out dividends.

In a stress test, we compute the losses resulting from sharp adverse price moves, or from an adverse macroeconomic scenario. The challenge of stress testing is to design scenarios that are extreme enough to constitute a fair test of severe potential loss, but are still plausible and could realistically take place. One can distinguish two schools of thought on stress testing by the extent to which macroeconomic design, as opposed to statistical analysis of risk-factor behavior, enters into scenario design.

Another purpose of stress testing is to examine scenarios in which market prices behave in a way unlikely in the classical normal-return model. To a large extent, increasing reliance on stress testing is a response to the many drawbacks of the market and credit risk modeling approaches we have been studying in Chapters 2 through 9. We have reviewed some critiques of these models in Chapters 10 and 11. One possible response to the drawbacks of the standard statistical models is to develop better statistical models or techniques, and much valuable discussion has explored the potential advantages of, say, extreme value theory (EVT) techniques, or the use of implied volatility data. However, these discussions have not led to a definitive and widely shared conclusion about a superior alternative to the standard model. Another potential response is agnosticism about the capacity of risk managers to make useful quantitative statements about tail events and their impact on the firm’s or investor’s returns. But pure agnosticism does not yield practical tools that firms and investors can use to manage risk day to day. Stress testing is a third type of response. Rather than relying entirely on the results of a model to obtain a “likely worst case,” stress testing skirts the modeling issues by placing the scenario itself, rather than the model, in the forefront.

Stress testing has also become more important over time because it is easier to communicate stress test results than model results to most audiences. Most investors and financial firm executives do not have specialized training in statistical modeling. It is remarkable that VaR, a fairly sophisticated risk measure involving quantiles and distributional hypotheses, as well as a good understanding of its limitations, have extended as far into

mainstream financial discussion as they have. Even if better statistical models could provide the solution to the limitations of the normal distribution, it is unlikely that the necessary fluency in the statistical issues would keep up. To be useful, and to having an impact on decision making, risk measures have to be understandable. It is therefore appropriate to develop risk measures that can be understood and acted on without extensive quantitative training.

Stress testing need not stand in opposition to statistical modeling. As noted, some approaches to stress testing build on the results of return models. Depending on one's view of the standard models, stress testing can be considered a supplement or a substitute for the joint normally distribution and other non-normal model-based approaches. We may use historical data, or human judgment, or both, to create likely worst-case scenarios.

13.3.1 An Example of Stress Testing

Consider the sample portfolio we studied in Chapter 5. The second column of Table 13.1 displays the volatilities of the risk factors, computed using the EWMA approach, as in Chapter 5. The parametric estimates of the annualized marginal volatility for each of the six risk factors are displayed

TABLE 13.1 Volatility and Shocks

Factor	Vol	MVol	VaR Shocks		Stress Test Shocks		Price Units
			% Shock	Δ Price	% Shock	Δ Price	
EUR	5.70	2.16	0.42	-5.42	-5.0	-64.32	USD ticks
JPY	6.44	1.30	0.25	-21.51	-5.0	-425.86	$10^6 \times \Delta$ USD price of yen)
SPX	7.80	-0.64	-0.13	1.73	-10.0	-137.69	points
GT10	14.77	-2.88	-0.56	0.20	-10.0	3.49	bond points
XU100	20.18	17.38	3.38	-1,340.61	-15.0	-5,944.08	points
TRL	12.36	9.62	1.87	-12.92	-7.5	-51.76	$10^9 \times \Delta$ (USD price of TRL)

“Vol” represents the volatility of each risk factor at an annual rate. For GT10, the vol is the yield vol. “MVol” represents the marginal volatility contribution as defined above, again, as a yield vol for GT10. The VaR shock is 3.09 times the MVol, expressed at a daily rate. The stress shock is the risk factor return stipulated by the scenario. Both the VaR and stress shocks are therefore expressed as a percent change in the risk factor. The volatilities and shocks are all expressed in percent. The “ Δ price” corresponding to the VaR and stress scenarios states the price units, as detailed in the last column.

in the third column. Recall that the marginal volatility, like the marginal VaR, takes into account not only the return volatilities of the risk factors, but also their return correlations and the portfolio weights.

The VaR shock is the percentage change in the risk factor in the VaR scenario and is related to the marginal volatility by

$$\text{VaR shock} = -z_* \times \text{marginal volatility}$$

where z_* is the ordinate of $N(0, 1)$ corresponding to the selected confidence level α . In a portfolio context, the VaR shock depends not only on the the asset return volatility, but also on the correlation to other returns and thus on the composition of the portfolio. It is displayed for each risk factor in the next two columns of Table 13.1. The change in the value of the risk factor—the risk factor return—is the VaR shock times the current level of the risk factor. To facilitate comparison to the stress scenario, we set a high confidence level $\alpha = 0.999$, so $z_* = 3.09$. That corresponds to a loss size, which, within the model, could be expected to occur on about one trading day in four years.

The stress shocks are displayed in the next column. While the VaR shocks are outputs from the statistical model used to calculate the volatilities, the stress shocks are “designed,” that is, chosen to represent a highly adverse scenario. To facilitate comparison with the VaR analysis, they are to be understood as one-day returns to the risk factors. The stress shocks are set to represent a pattern frequently observed in the financial market crises of recent decades, as we discuss in Chapter 14: the U.S. dollar appreciates sharply against other currencies, while U.S. Treasury bond yields and all equity markets drop sharply. Corresponding to the decline in yields, U.S. Treasury prices rise sharply; this aspect of the scenario design is consistent with the “flight to quality” or “risk off” behavior typically seen in financial markets under stress. The same flight to quality pattern would tend to drive the dollar sharply higher, but without knowing more about the specific background of a crisis, it would be hard to predict whether the euro would appreciate against the yen, or vice versa, or neither. The drop in equity prices is large in all markets, but is even worse for emerging than for developed-country markets.

The results for the positions and the portfolio are displayed in the table below. For consistency with the VaR estimates, stress losses (gains) are represented as positive (negative) U.S. dollar amounts. There are a number of notable differences from the VaR estimates.

- The P&L estimates are much larger, of course, since the stress scenario shocks are much larger than the VaR shocks. In particular, the Turkish

stock market position becomes an even more dominant driver of the aggregate risk of the portfolio.

- The positions that would be expected to act as risk mitigants, such as the short S&P 500 position, have a much larger effect. In the VaR estimates, it is a minor risk contributor. In the stress scenario, it helps a great deal to limit portfolio losses.
- The two major-currency positions offset one another completely rather than partially. Of course, this is an artifact of the stress scenario, which assigns identical returns to the two currencies.
- The U.S. Treasury position, which in “normal” markets is a small risk contributor, becomes a risk mitigant in the stress scenario.

	MVaR	Stress Loss
long EUR	4,214	50,000
short JPY	-2,526	-50,000
short SPX	1,254	-100,000
long GT10	2,003	-34,937
long XU100	52,553	213,750
Portfolio	57,498	78,813

The portfolio stress loss is significantly worse than that estimated by the VaR. Partly, that is due to the fact that the VaR analysis is based on a historical observation window in which volatilities were relatively low. But it also reflects the fact that the stress scenario has been designed to take fat tails into account. The scenario is intended to reflect a risk event that occurs with a probability of about 0.1 percent, and is therefore appropriately a much larger loss than the normal return quantile in the VaR estimate.

But the stress scenario has placed not only the overall loss, but also the portfolio construction in a different light. Analyzed only with VaR tools, this looked like a fairly low-volatility portfolio, with the Turkish stock index as the one “high-octane” component, the tail wagging the portfolio dog. The stress test, if it embodies a reasonable scenario, shows that the other portfolio elements are more dynamic than the VaR analysis alone revealed, based as it is on a low-volatility sample period and a thin-tailed return model. Finally, the diversification characteristics of the portfolio look very different. Overall, because the the long U.S. Treasury and short S&P index positions have much larger gains in the stress scenario than in the VaR analysis, the entire portfolio looks better hedged for a crisis.

13.3.2 Types of Stress Tests

We have alluded to the different approaches to constructing scenarios, with one focusing more on finding empirically reasonable shocks to risk factors, while another looks more to the “story” being told by the scenario. This reflects the different purposes of stress tests. On the one hand, stress tests are designed for ensuring a firm’s capital adequacy and survival through a crisis. Therefore, both risk appetite and economic analysis have important roles to play in determining stress scenarios. On the other hand, they are designed to take account of alternative return distribution and volatility models to the normal.

Stress tests must be formulated with a specific discrete time horizon, that is, the return scenario is posited to take place, say, over one day or one month. Unless the stress test horizon is extremely short, we need specific assumptions about what trading of positions and hedging, and at what prices, will take place within the horizon of the stress test. Some trading assumptions, such as ongoing delta hedging of option positions, are perhaps reasonable accommodations to dynamic strategies. Permitting other positions to change so as to protect against loss within a stress scenario raises the possibility that the stress test results will understate losses, since the point of a stress test is to explore potential losses given the current portfolio or book of business. More importantly, the deterioration of market and funding liquidity and impairment of market functioning is also likely to limit the amount of trading that can be done in stress conditions. Stress scenarios should therefore take current positions for the most part as unalterable.

The stress test in our example had a horizon of one day, and is consistent with a no-trading assumption. The stress test result is then just a mark-to-market P&L of the portfolio. A longer horizon may be more useful for many firms. Typically a span of one month or one calendar quarter is needed for even a relatively contained financial crisis to play out. But a longer horizon than one quarter is unrealistic if a no-trading assumption is imposed.

Stress tests have been classified into several types. *Historical stress tests* apply shocks that have actually been realized in a historical crisis to the portfolio. Table 1.1 of Chapter 1 provides some examples of episodes that could be incorporated into a historical stress test. Several ambiguities need to be resolved in carrying out a historical stress test. For example, the worst loss for a particular risk factor, say the dollar-Mexican peso exchange rate, might have been realized one week into a crisis involving Mexican markets, while the worst loss for the Mexican stock market might not have been realized until several weeks later. The stress loss for a portfolio will be greater if the worst losses are bundled together as an instantaneous mark-to-market P&L than if a particular historical interval is chosen.

Historical stress tests are important and valuable as a point of reference, but history never repeats itself. An alternative approach is based on possible future events though it may be informed by history. Most stress tests in practice are of this type. But if history is no longer a rigid guide, how do we design the scenarios? One important principle is to identify portfolio vulnerabilities and see to it that they are properly stressed. VaR analysis can be a useful complement to stress testing, as it has a capacity to identify subtle vulnerabilities in a complex portfolio that are not obvious when looking at line-item positions. Discussion with traders is also important in identifying potential stresses.

Other approaches to stress testing are less dependent on judgment and rely more on algorithms for evaluating the portfolio in different scenarios. Such approaches capture both the interaction of multiple risk factors in generating losses, such as occurs in option positions, as well as the susceptibility of the specific portfolio to particular combinations of factor returns. In the *factor-push approach*, many combinations of risk factor returns are tested, and the stress loss is taken as the largest portfolio loss that occurs. The highest and lowest returns for each risk factor are set at some reasonably large potential range, resulting in a grid of shocks. The stress test result is the largest portfolio loss resulting from this grid of stresses. An example is the Chicago Mercantile Exchange's (CME) Standard Portfolio Analysis of Risk (SPAN) system, used in establishing net margin requirements for futures positions. Such approaches can have many dimensions, since there are potentially many combinations of risk factors to search over, and the portfolio must be revalued in each one. A class of tools for limiting the range of risk factor combinations using distributional assumptions is called the *maximum loss approach*.

An example of the potential usefulness of factor-push approaches is the Amaranth hedge fund collapse of September 2006. Amaranth had put on large positions in the calendar spread between natural gas futures with different expiries. The capital at risk from this one set of positions was comparable in magnitude to the entire NAV of the fund. A large bet on a tightening of these spreads to historical norms went awry as spreads widened instead. Historical data on natural gas calendar spreads and their volatility would have underestimated the potential loss to the trade. A stress test based on a wide enough range of possible values of the calendar spreads involved might have identified the potential loss. It is, of course, unknown what degree of awareness the Amaranth fund managers possessed of the potential loss, and thus, whether the risk had been consciously taken on.

In many cases, we are interested in shocking only some risk factors in a portfolio. This raises the issue of how to treat the remaining risk factors. One approach is to use their last-observed values, that is, set their returns to zero,

in computing portfolio losses in the stress scenario. This may lead to unrealistic scenarios. Another approach, sometimes called *predictive stress testing*, is to use their conditional expected values. The expected values are conditioned on the values taken by the stressed risk factors, using the estimated correlation matrix of all the risk factors. This in turn raises a final issue in the treatment of the risk factor correlations; correlations between many risk factor returns are higher during periods of financial stress. We discuss the behavior of correlation in financial crises in more detail in Chapter 14.

In spite of its difficulties, stress testing has taken on great importance relative to VaR in recent years. At one point, stress testing was discussed as a complement to VaR and in the context of VaR modeling, but the emphasis has now shifted, and the results of VaR analysis are now apt to be reported as a stress scenario among others. Stress testing should be carried out regularly, as a regular part of the risk reporting cycle. The scenario design should be varied as portfolio concentrations and market conditions evolve.

13.4 SIZING POSITIONS

Determining the appropriate size of positions is one of the major decisions investors and traders must make, alongside choosing which trade ideas to adopt and to discard, and determining hedging policy. They need to avoid excessive position concentration and achieve diversification. In investment management, this is the allocation decision.

Risk capital calculations can be helpful in determining position size. Identifying large risk contributions provides a more reliable guide to concentration than notional size. In this section, we compare a number of tools used to identify concentrations and guide the search for diversification.

13.4.1 Diversification

We have discussed diversification in a number of contexts. Here, we provide some quantitative measures for market and credit risk.

A common method for measuring the degree of diversification from a market risk standpoint is to compare the VaR of a portfolio to the sum of the VaRs of the individual positions. Let $\mathbf{x} = (x_1, \dots, x_M)'$ denote a portfolio. The *diversification benefit* is defined as difference between the sum of the single-position VaRs and the portfolio VaR:

$$\sum_{m=1}^M \text{VaR}_t(\alpha, \tau)(x_m) - \text{VaR}_t(\alpha, \tau)(\mathbf{x})$$

This quantity is expected to be non-negative. But there are cases, uncommon but not pathological, in which the diversification effect is negative. VaR consequently violates the axiom of subadditivity, as discussed in Chapter 11). In almost all cases, however, the diversification benefit will be positive and make sense.

For credit risk, a common measure of diversification is the *Herfindahl index*. It is equal to the sum of the squares of the share of each credit in the portfolio. If there is only one credit, the index is equal to unity. If there are n credits, each with an equal share, the index is equal to n^{-1} .

Another diversification measure for credit risk is the *diversity score*. In Chapter 8, we saw that dividing a credit portfolio with a fixed par value into smaller pieces with uncorrelated defaults, but the same default rate, progressively reduces the fraction of the portfolio that defaults, and thus the credit VaR, up to a limit determined by the uniform default rate. The diversity score provides a comparison between such a granular portfolio and a congruent one with correlated defaults. The portfolio of correlated credits will have a high diversity score if its credit risk is closer to that of the granular, uncorrelated one and a low diversity score if its credit risk is closer to that of a single large credit. As noted earlier, because credit events are generally low-probability, so that credit returns are more fat-tailed than market returns, more granularity is required to reduce idiosyncratic risk to a desired level than for an equity portfolio.

13.4.2 Optimization and Implied Views

Reverse optimization or *implied views* is another tool, emerging from the risk capital framework, to help portfolio managers decide whether to invest more or less in particular assets. Recall from Chapter 2 that in an efficient portfolio, the expected excess return of each of the M assets is equal to its beta to the portfolio, multiplied by the excess return of the portfolio. This condition was expressed in Equation (2.3), reproduced here:

$$\mu_m - r_f = \beta_m(\mu_p - r_f) \quad m = 1, \dots, M$$

Each of the β_m is the ratio of the covariance of asset- i excess return with the portfolio excess return to the portfolio excess return variance:

$$\beta_m = \frac{\sigma_{mp}}{\sigma_p^2} \quad m = 1, \dots, M$$

Putting these two expressions together, we conclude that, in an efficient portfolio,

$$\frac{\mu_m - r_f}{\mu_n - r_f} = \frac{\sigma_{mp}}{\sigma_{np}} = \frac{\sum_k^M \delta_k \sigma_{mk}}{\sum_k^M \delta_k \sigma_{nk}} \quad m, n = 1, \dots, M$$

for any pair of investments m and n in the portfolio. The second equality puts the implied view in the context of the delta-normal approach to portfolio risk measurement of Chapter 5.

This is a remarkably far-reaching result. An efficient portfolio is not necessarily optimal. Efficiency, as noted in Chapter 2, is a condition of “minimal rationality” in portfolio construction. A reallocation of the available capital that makes a nonefficient portfolio into an efficient one is not a trade-off; it is a zero-cost improvement.

The result states that, if the portfolio is efficient, if the asset returns are multivariate normal, and if we are confident that our estimates of the variances and covariances of the asset returns are accurate, then both of the following statements hold:

1. If we have estimated or know the excess return on even one investment in the portfolio, the portfolio itself implicitly reveals all the other expected excess returns.
2. Even if we have no estimates of excess return, the portfolio itself reveals the *ratios* of expected excess return of any pair of investments.

An asset manager may not have formulated an explicit excess return estimate for any of the assets in the portfolio. But he will certainly claim to have constructed an efficient portfolio. Suppose the marginal VaR of asset m is twice that of asset n . He can then step back and consider whether he really expects the return of asset m to be double that of asset n .

We can relate these implied views to our measures of risk contribution:

$$\frac{\mu_m - r_f}{\mu_n - r_f} = \frac{\sigma_{mp}}{\sigma_{np}} = \frac{\text{MVaR}_{mt}(\alpha, \tau)}{\text{MVaR}_{nt}(\alpha, \tau)} \quad m, n = 1, \dots, M$$

The ratio of the marginal VaRs are equal to the ratios of the expected excess returns.

13.5 RISK REPORTING

We have presented a range of risk statistics so far: risk factor sensitivities, VaR and related concepts, stress test results, and risk capital measures. In order to be of any use, they must be presented in the form of reports to portfolio managers and other decision makers, and interpreted. Chapter 11 briefly summarized the elements of a risk measurement system, which would typically include a reporting layer. Let's use the example of the portfolio delta-normal VaR computed earlier to see how we can use these statistics to better understand the portfolio.

We start with a VaR report on the portfolio by position. The first two columns of data display the marginal and incremental VaRs of the factors. The complement VaR is the VaR of the portfolio after the position has been removed or hedged completely. The standalone VaR, finally, is the VaR of the position viewed in isolation.

Position	MVaR	IVaR	Complement VaR	Standalone VaR
long EUR	3,172	2,435	40,849	8,349
short JPY	-1,902	-2,838	46,122	9,438
short SPX	944	-563	43,847	11,435
long GT10	1,508	825	42,460	7,734
long XU100	39,562	31,925	11,360	41,779
Total	43,285			78,735

The first noteworthy feature of this portfolio is that it has considerable diversification in it. This can be seen in the fact that most of the marginal VaRs are considerably smaller than the standalone VaRs. Viewed differently, the sum of the standalone VaRs at \$78,735 is nearly double the portfolio VaR at \$43,285. Nonetheless, most of the risk in the portfolio appears to be coming from the XU100 position even though the market values of all the positions are equal. Its marginal VaR is about 90 percent of the total, and is almost as great as its standalone VaR. The VaR report immediately reveals that there is a concentrated source of risk in the portfolio.

We can also report the risk of the portfolio by risk factor. We add a column showing the annualized volatility of each risk factor, measured using the exponentially weighted moving average (EWMA) algorithm and expressed in percent. The volatility of GT10, the 10-year Treasury note, has

been converted from a yield to a price volatility so that it can be compared to the other factor volatilities.

Risk Factor	σ_n	MVaR	IVaR	Complement VaR	Standalone VaR
EUR	5.70	3,172	2,435	40,849	8,349
JPY	6.44	-1,902	-2,838	46,122	9,438
SPX	7.80	944	-563	43,847	11,435
GT10	5.28	1,508	825	42,460	7,734
XU100	20.18	25,468	19,967	23,317	29,578
TRL	12.36	14,094	11,958	31,326	18,108
Total		43,285			84,642

For the first four positions, since they are each a linear function of the risk factor in the same order in the list, the VaR results are identical to those reported by position. Note that the sum of standalone risk factor VaRs is somewhat higher, since we have broken out two of the risk factors. For the Turkish stock index (XU100) position, we can gain some intuition into how the high marginal risk is generated. The volatilities of the local currency XU100 and of the Turkish lira return are the two highest in the book. Moreover, as we saw in Chapter 5, these two returns have a fairly high positive correlation of 0.51, since sharp local-currency stock market declines tend to coincide with depreciation of the local currency against the dollar. The two risk factors do not have any large negative correlations with the other long risk factors or positive ones with the short risk factors that might offset their risk contributions. We can also see that about 40 percent of the risk of the XU100 position comes from the currency exposure it generates, and could be eliminated by hedging that risk.

To understand the portfolio better, we can group the positions into strategies, based on a trade thesis common to several positions, or on the use of a position to hedge an unwanted risk in other positions in the group. (In this example, as it happens, no positions are there as a deliberate hedge, though the short S&P index position will act as one in a stress scenario.) We can break our portfolio into three subportfolios, with VaR statistics as reported here:

Strategy	Standalone VaR	MVaR	IVaR	Complement VaR
Euro bullish/yen bearish	6,374	1,270	809	42,476
U.S. economy bearish	13,190	2,452	444	42,840
Turkish stock market	41,779	39,562	29,349	13,935

Euro bullish/yen bearish strategy (positions 1 and 2), both short the dollar against the other two major currencies. This strategy has the smallest standalone VaR. It is also the strategy with the smallest impact on the rest of the portfolio.

Some characteristics of the positions increase risk: EUR and JPY returns are positively correlated with the returns of the Istanbul Stock Exchange index (ISE 100), a long risk factor, and negatively with the 10-year U.S. government yield, a short risk factor. On the other hand, both EUR and JPY returns have relatively low volatility. Moreover, since the portfolio is long EUR risk, it contributes risk, while the short JPY position subtracts from the risk of the portfolio, as seen from its negative marginal and incremental VaR. The returns of the two positions are highly correlated with one another, and their correlations with the other risk factors in the portfolio are similar, so their risks are more or less offsetting. Together, therefore, this dollar-neutral bet on the euro-yen exchange rate has only a small net risk, viewed in isolation, as seen from the small standalone VaR, and it has only a small overall impact on the rest of the portfolio, as seen from the small incremental VaR.

U.S. economy bearish strategy (positions 3 and 4), which will both perform well if there is a material deterioration in U.S. growth prospects. The trade will also perform well if there is an increase in risk aversion, or, in the extreme, a “flight to quality” due to fear of a financial crisis. In that sense, the entire trade could be used as a portfolio hedge for a wide range of portfolios, such as the typical institutional portfolio, that are generally long risk assets.

This strategy also has a small impact on the risk of the entire portfolio (low incremental VaR). However, its standalone VaR is fairly high, since the risk factors underlying the two positions have an estimated correlation close to zero, limiting the diversification benefit. While one would expect a strong positive correlation between yields and equity indexes during a stress event, estimated correlations are often low. As we see in the next chapter, equity and bond return correlations are susceptible to rapid change during financial crises.

Turkish stock market strategy (position 5). This position’s impact on the portfolio is overwhelming, primarily because the volatilities of both of its risk factors are high, and because the risk factors are positively correlated: A rise in the Turkish stock market is associated with a weakening of the lira against the dollar.

Marginal and incremental VaR, when used together with stress test results and with nonquantitative, judgmental data, can help us understand a set of positions as a portfolio, and to trace through the ways in which sets of positions within the portfolio contribute to, or mitigate, portfolio risk.

13.6 HEDGING AND BASIS RISK

There are two ways to mitigate risk: reducing positions and hedging. Risk management encompasses ensuring that exposures to all risk factors are the ones desired by the risk taker, ascertaining that hedges are effective, and seeing that risk exposures are sized in accordance with the risk taker's goals. We have discussed position sizing in the context of diversification and risk capital. In this section, we focus on problems and issues with hedging in the context of trading and investment risk. Hedging is an issue that affects not only traders, but every market participant, since all are exposed to a range of risks, and all must decide which to bear and which, if possible and cost effective, to mitigate.

There is no bright line between decisions on whether and how to hedge and other investment decisions. "Hedging" describes exposures that are not the core of a trade thesis, but are bundled with the securities through which the thesis is expressed. Hedging involves weighing risk against return, but is generally couched more in terms of the cost of hedging. Effective hedging reduces the volatility of the portfolio expressing the thesis, or eliminates an unwanted exposure that is not part of the thesis.

The term *basis risk* is generally used to describe the risk that two very similar, but not quite identical, securities will diverge or converge in price to the detriment of the investor. There is no clear standard of when two securities are similar enough to describe the relative price risk as "basis" rather than "market risk." Basis risk is one of the key risks to which a hedged portfolio is exposed. It can be thought of as the risk that a hedge position fails to fulfill its purpose.

Some important examples are:

The Treasury bond basis is the difference between prices of U.S. Treasury notes and bonds in the cash market and the corresponding futures prices. Cash market prices are typically somewhat higher than futures prices because the seller of a Treasury futures acquires a *delivery option* from the futures buyer. At any point in time, several cash notes or bonds are eligible to be delivered by the futures seller to the buyer to satisfy the seller's delivery obligation. As interest rates change, the identity of the *cheapest-to-deliver* security

that is, the cheapest security among all those eligible, may change, but the futures seller may always discharge his obligation by delivering whatever note or bond is cheapest-to-deliver. The value of this option reduces the value of the futures relative to the cash market.

The bond-CDS basis. The spread over the Libor curve of a corporate bond in the cash market is typically not precisely equal to the premium of a CDS on the class of bonds of that issuer and seniority. Similar spreads exist between CDS indexes, such as the CDX, and indexes of CDS on asset-backed securities, on the one hand, and indexes of the spreads on underlying cash bonds. The difference can be positive or negative, but is typically small, since a large difference invites market participants to place trades that would profit from a reversal.

During the subprime crisis, however, the bond-CDS basis became unprecedentedly wide for many bonds, with CDS spreads much tighter than those of cash bonds. This phenomenon was driven by liquidity. Funding liquidity drove many market participants to attempt to raise cash by selling assets that had been financed in part by borrowing, usually in collateral markets, and could no longer be financed on the same terms as before the crisis, if at all. The preponderance of offers also impaired transactions liquidity. Together, these liquidity-based forces drove cash spreads wider than CDS. A market participant wishing to take advantage of this gap would have had to buy bonds in the cash market and buy CDS protection. The position would have had a positive cash flow, and, as noted in Chapter 11, such “arbitrage” trades are much-prized by traders. Very few market participants were in a position to do so, however, since buying cash bonds was a capital-intensive activity at a time of dire shortage of capital or “balance sheet.”

Figure 13.4 illustrates this phenomenon with daily differences between spreads on Citigroup 10-year senior unsecured bonds and 10-year CDS spreads. The basis was close to zero prior to the subprime crisis, but reached a peak of close to 500 basis points, as Citi unsecured bond prices, like those of many other money-center banks, and liquidity in the financial-issuer bond market, reached their nadir in March of 2009. As can be seen, the basis not only widened, but also became very volatile.

Other important examples of basis risk arise in structured credit trading and became important drivers of large losses during the early phases of the subprime crisis. Traders sought to hedge the credit and market risk in investment-grade residential mortgage-based security (RMBS) by going long

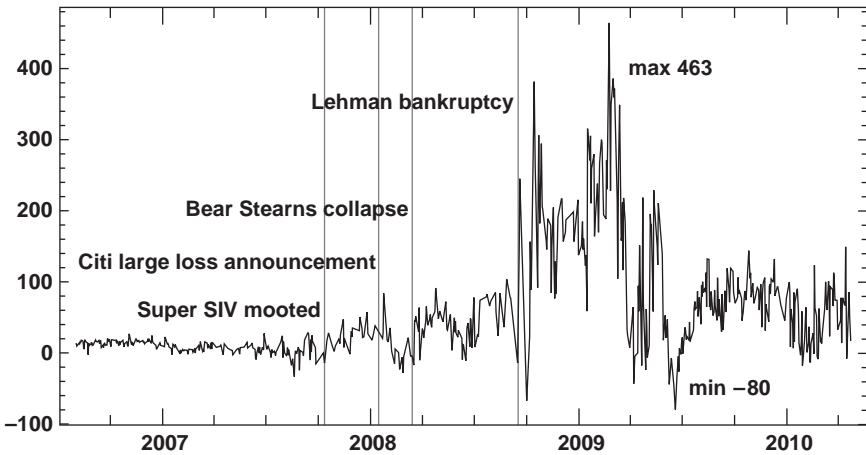


FIGURE 13.4 Citigroup CDS Basis 2007–2010

Difference between the spread over Libor (z -spread) of Citigroup Inc. bonds and the premium on Citigroup 10-year CDS on senior unsecured debt (bond spread minus CDS premium). The bond spread is blended from spreads on two senior unsecured issues: the 4.7% maturing May 29, 2015 (CUSIP 172967CY5) and the 5.85% maturing August 2, 2016 (CUSIP 172967DQ1).

Source: Bloomberg Financial L.P.

the ABX index, described in Chapter 11. But the ABX, while it captured the severe price declines in almost all RMBS over longer periods, was only very loosely tied to the performance of any particular portfolio of RMBS over any shorter time frame. As seen in Figure 11.4, the ABX indexes experienced several short-lived rallies in 2007 and 2008. Losses by some investors in lower-rated and lower-quality subprime RMBS were reportedly exacerbated by losses on ABX hedges.

In some instances, traders hedged positions in mortgage and non-mortgage structured credit with positions in the IG corporate credit indexes rather than the ABX. The corporate spread products have the virtue of relative liquidity, making it less costly to adjust hedges frequently as positions change. But as can be seen in Figure 14.14 and the table following it, such a hedge would have been disastrous during the subprime crisis. Investment-grade and high-yield corporate spreads rose by a factor of about 8 or 9 between mid-2007 and the end of 2009. Spreads on the investment-grade structured products that were to be hedged widened by a factor of about 100, in some cases considerably more.

A basis can open up almost anywhere. One might use a hedging instrument only to find that it diverges in some surprising way from the risk factor

one had hoped to offset. Because of issues that can loosely be described as “basis risk,” most exposures cannot be precisely hedged. There is a spectrum of hedging accuracy, ranging from simple hedges that are relatively easy to accurately gauge and vary over time, to positions that are close to unhedgable.

Some exposures can be hedged precisely enough that they can be treated as routine. Currency hedges on foreign exchange-denominated fixed-income securities, for example, are relatively easy to measure. Consider a long government bond position denominated in foreign currency. Assume the bond has no credit risk. The currency risk of the position can be hedged by selling forward the foreign-exchange proceeds of coupon and redemption payments, thus locking in the current forward foreign-exchange rates. This hedge could be combined with the initial purchase of foreign currency needed to buy the bond in a currency swap, probably executed with the dealer through which the bond is purchased. The investor now faces only the desired exposure to interest rates.

Even this currency hedge is not perfect if the security will not be held to maturity, since there is then price risk in the bond. Even if foreign risk-free rates fall, so the trade turns out well, the domestic-currency proceeds of the bond are uncertain, and the return on the investment is therefore also uncertain, particularly if the position is highly leveraged. The trader also has counterparty risk exposure through the currency swap hedge. For some currencies, there may also be liquidity risk.

If there is significant price risk in the security, the foreign-exchange risk on the profits can be material. Examples are equity or credit-risky bond positions. This risk is sometimes called *quanto risk*, after a type of exotic option introduced in the late 1980s. Quanto risk occurs when correlation risk is embedded in the price risk of a single position.

Many trade ideas require risk-free rate hedges, because the desired exposure is a credit spread, but the trade is executed through cash securities that pay a coupon and incorporate a risk-free rate as well as spread component. For example, a long investor in a U.S. fixed-coupon corporate bond may sell Treasury bonds or pay fixed in an interest-rate swap. The net cash flow then consists only of the spread at the time the trade is entered into. Such hedges can also generally be put on with high accuracy. However, even in this simple example, there is a potential for basis risk. Corporate spreads are generally set by the market relative to swap rather than Treasury rates, particularly for lower credit quality bonds. But the hedging vehicle chosen may in any event prove to be the wrong one, if the swap spread, the spread between government bond yields and swap rates, changes materially. Swap spread volatility has been extremely high at times, as can be seen in Figure 14.15.

Another issue with this type of hedge is that credit-risky securities with low spreads that trade on price can hit a “ceiling” at or slightly above par. Even if interest rates fall, the security may not appreciate in price, introducing negative convexity into its price behavior. This widens its spread, but if the clientele for the security is narrow, and the demand curve is flat at above-par prices, the widening may not generate much additional demand. One solution to this problem is to use swaptions or options on Treasury or eurodollar futures to hedge, thus matching more closely the security’s convexity profile.

The most difficult types of risk to hedge are distressed and bankrupt bonds, and other securities with binary return profiles. In some cases, such as merger arbitrage, there is a natural hedge for a binary event. In general, however, hedges with continuous price behavior will not perform well.

FURTHER READING

The algebra of risk contributions is laid out in Sharpe (2002). Sharpe (2002) and Froot and Stein (1998) provide introductions to the risk capital and risk budgeting approach to investment. A general discussion of the concept of risk capital is provided by Merton and Perold (1993) and Perold (2005). See Litterman (1996) on implied views.

For an overview of stress testing, see Schachter (2010). Kupiec (1995a) and Mina and Xiao (2001) discuss issues around stress tests of subportfolios. Kim and Finger (2000) discusses stress testing techniques that take account of “regime changes” in correlation. Committee on the Global Financial System (2005) describes large banks’ stress testing practices.

Breuer, Jandačka, Rheinberger, and Summer (2009) discusses systematic approaches to finding appropriate stress scenarios. See Amato and Remolona (2005) on the difficulty of achieving diversification in credit portfolios.

The supervisory stress tests conducted in the United States in 2009 are described in Hirtle, Schuermann, and Stiroh (2009), those in Europe in 2010 in Committee of European Banking Supervisors (2010). Alfaro and Drehmann’s (2009) discussion of the difficulty of designing stress tests focuses on supervisory stress tests, but applies in large part to internal risk management.

Fung and Hsieh (1997, 2001) discuss dynamic strategies. Hedge fund risk management is discussed by Lo (2001), Asness, Liew, and Krail (2001), and Lhabitant (2001). Burghardt and Belton (2005) discuss the Treasury basis.