

Risk Management for Alternative Investments

Prepared for the CAIA Supplementary Level II Book

Philippe Jorion*

June 18, 2012

**Philippe Jorion is a Professor at the Paul Merage School of Business, University of California at Irvine and Managing Director at Pacific Alternative Asset Management Company (PAAMCO). E-mail: pjorion@uci.edu*

© 2012 Philippe Jorion

This chapter provides an overview of financial risk management for alternative investments (AI). These investment products differ from traditional investments, such as stocks, bonds, and cash, and include hedge funds, commodities, real estate, and private equity. Alternatives are often viewed as having relatively low correlations with traditional asset classes, which should provide diversification benefits to the investor's portfolio. They have generally provided good returns with limited volatility or correlation relative to traditional investments. As a result, they are becoming increasingly important in investor portfolios.

Alternatives pose special problems for risk management. They cover a broad range of investment styles. At one end are hedge funds or commodity trading advisors (CTAs) that trade actively, generally using liquid instruments. At the other extreme are funds, such as private equity, where positions are kept for years in assets that are not marked to market.

In each case, risk management is a challenge. Funds that trade actively require a position-based risk management system to monitor and manage their rapidly changing risks. The opposite problem exists for funds that invest in illiquid assets. Illiquidity implies that prices do not change often, which makes it difficult to assess valuation properly, let alone risk.

AI managers generally take views on markets and securities. This process should add value to the investment for a number of reasons. First, AI managers have much wider investment opportunities and are less regulated than managers in traditional asset classes. They can short securities, leverage their portfolio, use derivatives, and generally invest across a broader pool of assets. They can set performance fees. They can impose lockup and minimum redemption notice periods. They do not have to disclose their holdings publicly. Second, AI managers have a stronger financial motivation to perform because of the compensation structure of the industry. Managers receive not only a fixed annual management fee ranging from 1% to 2% of assets under management (AUM), but also an incentive fee that typically represents 20% of annual profits. This helps align the managers' incentives with investors. The prospect of such riches undoubtedly attracts many of the best minds in the business. In the hedge fund industry, Agarwal, Daniel, and Naik [2009] found that greater managerial discretion and managerial incentives are associated with superior performance.

The very features that generate superior performance, however, should cause serious concerns to investors. AI managers can be secretive about their strategy and positions. They have more latitude in setting their net asset value (NAV) than regulated entities, resulting in the possibility of fraud or undue risk exposures that could go undetected and lead to blowups. In particular, incentive fees may tempt the manager to increase risks.

To some extent risks can be mitigated if portfolio managers have invested a substantial fraction of their wealth in the fund itself. For leveraged funds, risk can also be monitored by lenders, such as prime brokers for hedge funds. The prime broker, however, is mainly concerned about losses it could incur if the hedge fund defaults, not necessarily about losses to investors. Therefore, risk monitoring by the prime broker may not be sufficient, which is the reason risk management is particularly important for the alternatives industry. Risk management for alternatives is more difficult, however, than for traditional asset classes.

The purpose of this chapter is to provide an overview of risk management techniques for the alternatives industry. The emphasis is on *market risk*, which is the risk of losses due to movements in financial market prices or volatilities. In investment portfolios, this also includes *credit risk* as changes in perceived default probabilities or actual defaults are incorporated into market prices. *Liquidity risk*, which is the risk of losses due to the need to liquidate positions to meet funding requirements, is also discussed. Investments in alternatives involve operational and business risks as well, but these risks are not considered here.

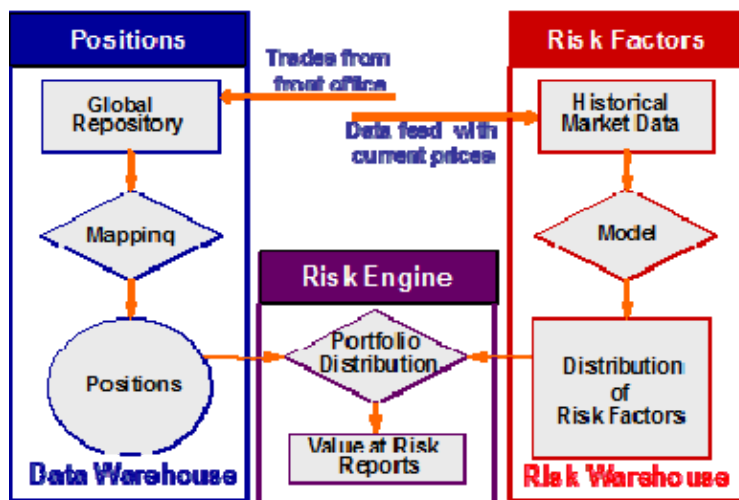
This chapter addresses the following topics: (1) the general design of risk measurement systems, which are constructed from positions, risk factors, and a risk engine, with a comparison of the pros and cons of position-based and returns-based risk measures; (2) how the process of mapping position on risk factors reveals exposures in the portfolio; (3) a review of conventional risk measures, such as leverage and concentration; (4) how to summarize the distribution of a single position or top-level portfolio distribution, comparing various aggregate measures of downside risk, such as standard deviation and value at risk (VAR); (5) an overview of the different approaches to VAR models, including the delta-normal approach, historical simulation, and Monte Carlo simulation; (6) how risk systems can be easily extended

to stress tests and used to manage risk by drilling down into its components; (7) risk measurement problems, such as biases in measures of volatility and correlations with other asset classes, that are created by illiquidity; (8) the limitations of traditional risk measurement systems; and (9) problems posed by the lack of transparency for some alternative investments and proposed solutions to this problem.

RISK MEASUREMENT SYSTEMS

Ideally, market risk should be measured using a position-based risk measurement system as described in Exhibit 1. This involves several steps. First, the risk manager must collect all the current *positions* in the portfolio and map them on the market risk factors via factor *exposures*. Second, the risk manager must construct the statistical distribution of *risk factors* from market data. Third, the risk manager must use the *risk engine* to derive the distribution of profits and losses on the portfolio. This can be summarized by several measures, such as the worst loss at a specified confidence level that is called value at risk (VAR).

Exhibit 1 Components of a Risk Measurement System



The key feature of this system is that it is *position based*. Traditionally, risk has been measured from *returns-based information* (i.e., from the time series of historical returns on the

portfolio). On the one hand, a returns-based risk system is easy and cheap to implement. On the other hand, returns-based measures suffer from severe drawbacks. They offer no information for new instruments and markets. They are completely ineffective for emerging managers or funds that have short track records. Such managers, however, account for a large fraction of the alternatives universe. Returns-based measures do not capture, or rather, are very slow at identifying, style drift. They may not reveal hidden risks.

As an example of this important issue, Lo [2001] considered a hypothetical fund, called Capital Decimation Partners, which appears to perform very well. Based on historical returns, the fund has a high *Sharpe ratio*, defined as the ratio of excess average return to volatility. It turns out, however, that the fund follows a very simple trading strategy, which is to sell out-of-the-money put options on the S&P index. As long as the options are not exercised, the portfolio generates positive and steady returns, which reflect the option premium. On rare occasions, however, the fund could suffer extreme losses. In this case, the returns-based volatility is totally misleading. More generally, returns-based risk measures give little insight into the real risk drivers of portfolio strategy.

Most of these drawbacks are addressed by position-based risk measures. They can be applied to new instruments, markets, and managers. These use the most current position information, which should reveal style drift or hidden risks. For example, Jorion [2007] showed that the risk of Capital Decimation Partners can be captured and controlled effectively by position-based risk systems. In addition, position-based systems can be used for forward-looking stress tests.

Position-based risk systems, however, can be challenging to implement and have drawbacks that risk managers must understand. First, they require more resources and are expensive to implement. A large bank could have several million positions, in which case aggregation at the top level is a major technology challenge. Second, position-based risk measures assume that the portfolio is frozen over the time horizon considered. Taking one month as an example, these risk measures combine the fixed portfolio positions at the beginning of the month with risk factor returns over the month, thus ignoring any active

trading that would take place in practice. To some extent, this problem can be mitigated by more frequent risk measurement. Finally, position-based systems are susceptible to errors and approximations in data and models. They require modeling all positions from the ground up, repricing instruments as a function of movements in the risk factors. The modeling of some instruments can be complex, leading to *model risk*. Even so, position-based risk measures are vastly more informative than returns-based risk measures. This explains why modern risk management systems are built from position-level information.

CONVENTIONAL RISK MEASURES

This section discusses conventional risk measures grouped into factor exposure measures and portfolio exposure measures.

Factor Exposure Measures

Exposures are a major component of position-based risk measurement systems. Their advantage is that they do not consider the range of potential movements in the risk factors and thus do not require assumptions about statistical distributions. This is also a drawback, however, because exposure measures are factor specific and do not aggregate across different types of factors. There is no way, for instance, to combine the duration of the bonds in a portfolio with the beta of the stocks in the portfolio to generate an overall risk measure. Nevertheless, exposures are intuitive to understand and are widely used in risk management and reporting.

Exposures are related to the mapping procedure for positions in Exhibit 1. Mapping is the process of replacing positions by dollar exposures on the risk factors. Consider, for example, a position in a default-free fixed-coupon bond, such as a U.S. Treasury bond. The most important risk factor for this bond is the movement in risk-free yields. Initially, assume that the yield curve is flat and moves in a parallel fashion. For each position, the exposure to this risk factor can be represented by *modified duration* D^* . This is constructed from information about the bond's cash flows and the sequencing of payments. The relative change in the market value of

the position P can be explained by the following combination of this duration and the movement in the risk factor Δy :

$$\frac{\Delta P}{P} = -D^* \Delta y \quad (1)$$

This first-order linear approximation can also be rewritten in terms of dollar duration (D^*P). In the mapping process, the position in the bond can be replaced by its dollar duration,

$$\Delta P = -(D^*P) \Delta y \quad (2)$$

If all N bonds in the portfolio are exposed to the same risk factor, then duration can be aggregated at the top level of the portfolio using the market weights of all positions w_i ,

$$D_p^* = \sum_{i=1}^N w_i D_i^* \quad (3)$$

The same principle applies to other measures of exposure, which are listed in Exhibit 2. These exposures are particularly important to monitor for major market risk factors, such as movements in the general level of equities, movements in risk-free interest rates, and movements in credit spreads. As Equation 3 indicates, exposures are additive across the entire portfolio. As a result, they do not diversify away as the number of positions increases in the portfolio.

This point can be demonstrated by considering a portfolio of N stocks, where returns are driven by a general equity index R_M plus residual effects ε , which, as a first approximation, are assumed independent across stocks,

$$R_i = \beta_i R_M + \varepsilon_i \quad (4)$$

The return of a stock portfolio can be written as

$$R_p = \sum_{i=1}^N w_i R_i = \sum_{i=1}^N w_i \beta_i R_M + \sum_{i=1}^N w_i \varepsilon_i = \beta_p R_M + \sum_{i=1}^N w_i \varepsilon_i \quad (5)$$

where β_p is the portfolio beta. As a result, the variance can be decomposed into two terms,¹

$$V(R_p) = \beta_p^2 V(R_M) + \sum_{i=1}^N w_i^2 V(\varepsilon_i) \quad (6)$$

Exhibit 2 Measures of Exposure

Risk Factor	Exposure
Movements in equity index price	Beta
Movements in the risk-free rate	Duration
Quadratic move in rates	Convexity
Movements in credit spreads	Spread duration
Movements in the risk factor	Delta
Quadratic move in risk factor	Gamma
Implied volatility	Vega
Default	Jump to recovery

As the portfolio becomes more diversified, the second term becomes smaller.² In contrast, the first term depends on the average portfolio beta and the variance of the market factor only. Because the average beta does not depend on the number of positions, it is not a diversifiable exposure. This is why institutional investors, who typically have large direct allocations to equities, should also monitor the beta exposure of their alternative investments to be aware of their total exposure to equities.

While useful, these measures of exposure have limitations. Linear exposures do not account for large movements in the risk factors. Quadratic measures improve the approximation but only to some extent. In addition, exposures do not aggregate across risk factors, which is why statistical risk measures are also needed.

Portfolio Exposure Measures

Conventional portfolio exposure measures provide very simple indicators of total risk. The most common family of measures is based on leverage. Consider, for instance, a stock-only hedge fund with the balance sheet described in Exhibit 3. The fund starts with \$100 in equity, borrows \$20 from the broker, and purchases \$120 in some stocks. The fund then borrows and short sells \$80 worth of other stocks.

Exhibit 3 Hypothetical Hedge Fund Balance Sheet

Assets		Liabilities	
\$120	long stock	\$80	short stock
\$80	cash lent to stock owner	\$20	loan
		\$100	equity

Define now V_A , V_L , V_S , and V_E as the market value of total assets, long stock positions, short stock positions, and equity, respectively (in absolute values). For a regular corporation, balance sheet leverage is conventionally measured by V_A / V_E . For investment funds, cash assets and liabilities are ignored. The usual measures of leverage are

- Long leverage, or V_L / V_E
- Short leverage, or V_S / V_E
- Gross leverage, or $(V_L + V_S) / V_E$
- Net leverage, or $(V_L - V_S) / V_E$

Each of these measures has a different use and interpretation. Generally, higher leverage indicates higher risk. Long leverage, for instance, is the inverse of the drop in the value of the long positions that would wipe out the equity, assuming other positions are not changed. In this case, long leverage is $\$120 / \$100 = 1.2$. Hence, if the long positions were to fall by $1 / 1.2 = 83.33\%$, the portfolio would lose $\$120 \times 83.33\% = \100 , which would wipe out the equity of $\$100$. Similarly, short leverage is $\$80 / \$100 = 0.8$, meaning that if the short positions went up by $1 / 0.8 = 125\%$, the equity would be wiped out. In this case, the portfolio would lose $\$80 \times 125\% = \100 .

An even worse scenario considers the gross leverage, which is $(120 + 80) / 100 = 2.0$ in this case. Disaster would happen if the longs were to go down by 50% and the shorts up by 50%. Of course, it is highly unlikely that both the long and short positions would go in the worst possible direction at the same time. Net leverage, which is $(120 - 80) / 100 = 0.4$ is more meaningful for this reason. It means that the equity would be wiped out if both longs and shorts went

down by 1/0.4, or 250%. The loss in this case would be $\$120 \times 250\% - \$80 \times 250\% = \$100$.

The advantage of these measures is that they can be constructed from portfolio listing information. The disadvantage, however, is that they are based on simplistic assumptions, which is that all positions among assets and/or liabilities move by the same amount. This may be acceptable for all-equity portfolios, but certainly less so for fixed-income products. For the latter, market values can be adjusted to 10-year equivalents. In addition, these leverage measures do not consider off-balance sheet items or the quality of financing.

Other measures of risk involve classifying the market value of the portfolio into different categories: asset class, industry concentration, region, issuer market capitalization, issuer style (e.g., value or growth), debt credit rating, debt duration, and so on. These are simple measures of diversification. Measures of concentration can be also reported, such as the list of positions with the largest long and short market values.

STATISTICAL RISK MEASURES: SINGLE INVESTMENT OR PORTFOLIO

This section illustrates how to compute measures of market risk for a single investment or at the top level of an investment portfolio. Consider for example a hedge fund trader with a position in a foreign currency, say, \$4 billion short the yen against the dollar. How can we describe the potential loss on this position over the next day?

This example is particularly appropriate because the risk factor, the yen/dollar exchange rate, is priced in a liquid market for which there is a long history that spans quiet and turbulent times. Hence, historical data should be a good guide from which to build the statistical distribution of future risks.

Building the Distribution

To answer this question, we use 10 years of historical daily data on the yen/dollar rate (1999–2008) and simulate a daily return. The simulated daily return in dollars is then

$$R_t(\$) = -Q_0(\$) [S_t - S_{t-1}]/S_{t-1}$$

where Q_0 is the current dollar value of the position and S_t is the spot rate in dollar per yen. For instance, the exchange rates on December 31, 1998, and January 4, 1999, are 112.80 and 111.65 yen/dollar, respectively. As the usual convention in this market is to quote the exchange rate in yen, we need to invert it to measure dollar values. The simulated return is then $R_2(\$) = -\$4 \text{ billion} (1/111.65 - 1/112.80)/(1/112.80) = -\$4 \text{ billion} \times 1.03\% = -\41.2 million . Repeating this operation over the entire sample, or $N = 2,539$ trading days, creates a time series of fictitious returns, which is plotted in Exhibit 4.

This approach is position-based because it uses the most current position, which is Q_0 . In contrast, a returns-based approach would use the history of profits and losses (P&L) for the trader. This is largely irrelevant, however, if the trader changes the portfolio substantially.

The statistical distribution of P&L can be summarized by a *histogram*, which compiles the number of observations within ranges, as shown in Exhibit 5. For example, there are five cases of a loss worse than $-\$120$ million, none between $-\$120$ million and $-\$115$ million, and so on. This entire distribution should be of interest to the risk manager. Generally, this can be described by the *probability density function*, or pdf, $f(x)$.

Exhibit 4 Time Series of Simulated Daily Returns on Portfolio (\$ Millions)





Exhibit 5 Distribution of Simulated Daily Returns on Portfolio (\$ Millions)



Summarizing the Distribution

Single summary statistics usefully describe the distribution of profits and losses. Define x_i as the value of an observation, and N as the number of observations. The *mean* μ is the first moment, or expectation, of X ,

$$E(X) = \frac{1}{N} \sum_{i=1}^N x_i \tag{7}$$

In this case, the mean is $-\$0.43$ million. As we shall see, this is small compared to typical risk measures.

The dispersion can be assessed by the *standard deviation* (SD), usually defined as σ . This is constructed from the variance, or second moment, as

$$SD(X) = \sqrt{\frac{1}{N-1} \sum_{i=1}^N [x_i - E(X)]^2} \tag{8}$$

In this example, the standard deviation, or volatility of returns, is \$26.6 million. Distributions with greater volatility are more risky. This measure, however, is symmetric and treats

equally both the positive and negative observations of like size. Another measure that focuses on the downside risk is the semi-standard deviation. Define N_L as the number of points below zero. The risk measure is

$$SD_L(X) = \sqrt{\frac{1}{N_L} \sum_{i=1}^N [\text{Min}(x_i, 0)]^2} \quad (9)$$

In this case, the measure is \$28.1 million, slightly greater than the standard deviation. This suggests that the distribution has a longer tail on the downside than on the upside.

Symmetry can be summarized by the *skewness* coefficient, which is the scaled third moment S . This is the expectation of the deviation from the mean to the third power,

$$S = \frac{E[(x - \mu)^3]}{\sigma^3} \quad (10)$$

Negative skewness indicates a long left tail, or the possibility of larger losses than gains. In our example, the skewness is -0.45 , which is slightly negative. Generally, a skewness coefficient below -1 should be a source for concern.

The size of the tails can be assessed by the *excess kurtosis* coefficient K , which is the scaled fourth moment in excess of 3, or

$$K = \frac{E[(x - \mu)^4]}{\sigma^4} - 3 \quad (11)$$

An excess kurtosis greater than 0 indicates that the distribution has fatter tails than a normal distribution and, hence, may generate more extreme values. In our example, the excess kurtosis is 3.34, which reveals much fatter tails than in a normal distribution. This is indeed typical of most financial series. Generally, an excess kurtosis coefficient above 2 should be a source for concern.

Another measure of downside risk is the lower *quantile*, which is the cutoff value q that correspond to a prespecified confidence level c ,

$$P(X \geq q) \geq c \quad (12)$$

Note that this is defined in terms of the cumulative probability to the right of q . The cumulative probability to its left is $1 - c$.

The quantile is usually transformed into a positive number that represents a loss, expressed in dollars or whichever currency is used. This is also known as value at risk (VAR), or the worst loss, such that there is a low, prespecified probability that the actual loss will be larger, $\text{VAR} = -q$. For example, suppose that we pick a 95% confidence level. We first compute the number of observations required in the left tail from $(1 - c) N = 5\% \times 2,543 = 126.95$. We then sort observations from the lowest return to the highest. Starting from the bottom, the observations ranked 126 and 127 are $-\$42.41$ and $-\$42.40$, respectively, with frequencies of 4.963% and 5.002%, respectively. Hence, $q = -\$42.41$, and VAR is $\$42.41$ million.³ The risk manager can then give the following economic interpretation to this number: Under normal market conditions, the most the portfolio can lose over one day is about $\$42$ million at the 95% confidence level.

VAR has become widely used as a statistical measure of portfolio risk. Notably, it is used by the Basel Committee [1996] as the basis for the market risk charge for commercial banks.⁴ This is the amount of capital that the bank must keep on its books as a buffer against trading losses. The advantage of VAR is that it takes into account the shape of the distribution function. Negative skewness or high kurtosis will be reflected in VAR.

VAR reporting is also required for investment funds in the European Union. These funds, known as UCITS (Undertaking for Collective Investments in Transferable Securities,) include hedge fund-type structures with derivatives. The UCITS directive requires “sophisticated” funds to measure market risk using a 99%, 1-month VAR. For these funds, VAR cannot exceed 20% of NAV.

A disadvantage of VAR, however, is that it sheds no light on the size of losses once VAR is exceeded. A complementary risk measure is the *conditional VAR* (CVAR), which is the average of losses beyond VAR. Using the ranked observations, we have M losses up to VAR.

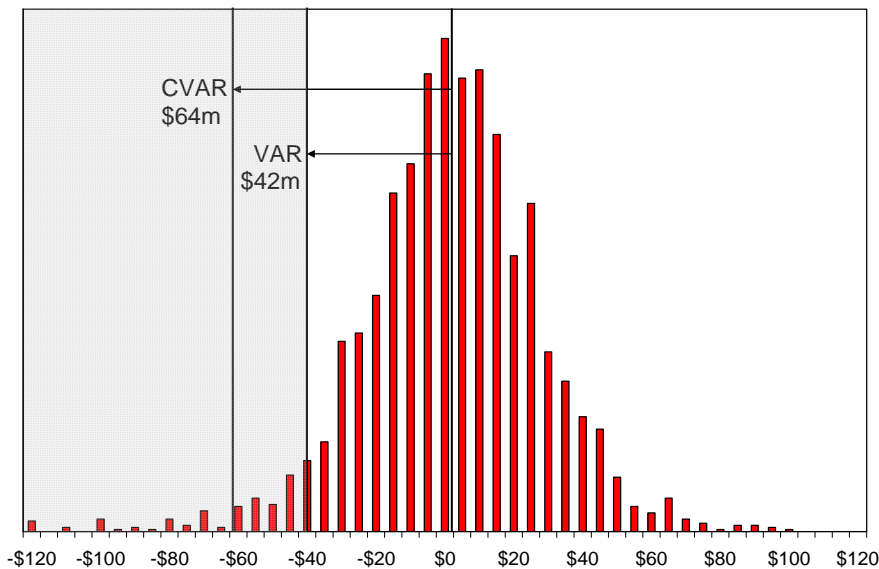
The CVAR is then

$$\text{CVAR} = \frac{1}{M} \sum_{i=1}^M (-x_i) \tag{13}$$

Exhibit 6 displays the VAR and CVAR risk measures for this sample. Here, CVAR is \$63.6 million. By construction, this must be greater than VAR. Generally, the two numbers are similar in terms of order of magnitude. In this case, the CVAR is 50% greater than the VAR of \$42 million. A portfolio could contain short positions in out-of-the-money options that could lose a lot of money if exercised. If this were the case, CVAR could be several times VAR. This raises a red flag that the portfolio is exposed to extreme risks.

Finally, it should be noted that even CVAR does not characterize the absolute worst loss. This is basically impossible to ascertain if movements in risk factors are unbounded.

Exhibit 6 Risk Measures for the Empirical Distribution (\$ Millions)



Parametric vs. Non-Parametric Approaches

The risk manager, however, may decide that the distribution of returns could be well described by a *parametric* distribution, such as the normal distribution. This considerably simplifies the analysis because the distribution is then characterized solely by two parameters, its mean μ and standard deviation σ . The quantile around the mean becomes a multiple of σ , using a multiplier α that depends on the confidence level. For example, if Z has a standard normal distribution and $c = 95\%$, we know from statistical tables that $P(Z \geq -1.645) = 95\%$,⁵ so that $\alpha = 1.645$. Hence, VAR can be defined as

$$\text{VAR} = \alpha\sigma \quad (14)$$

where σ is measured in dollar terms. This considers risk in terms of the deviation from the mean of the distribution on the target date. Another approach is to define risk in terms of changes from the initial portfolio value, in which case the formula for VAR should adjust for the mean, $\text{VAR} = \alpha\sigma - \mu$. However, it is common to ignore the mean for two reasons. First, when the estimation interval is small (i.e., daily), μ , is generally small, in which case it would be sensible to set it to zero.

Second, estimates of μ , are less accurate over short horizons, which implies that the estimated value of μ , is typically not statistically different from zero.

If σ is measured in terms of rates of return, it should be multiplied by the current value of the portfolio W , so that $\text{VAR} = \alpha\sigma W$. In our example, $\text{VAR} = 1.645 \times 0.664\% \times \$4 \text{ billion} = 1.645 \times \$26.6 = \$43.7 \text{ million}$.⁶ Note that it is close to the empirical, non-parametric VAR of \$42.4 million. At higher confidence levels, however, these two numbers start to diverge from each other because actual distributions have fatter tails than the normal.

Exhibit 7 displays the fitted normal distribution. Note that, relative to Exhibit 5, the tails are much thinner. This confirms the previous observation that the empirical kurtosis of the data is greater than that of a normal distribution.

Exhibit 7 Risk Measures for the Normal Distribution (\$ Millions)

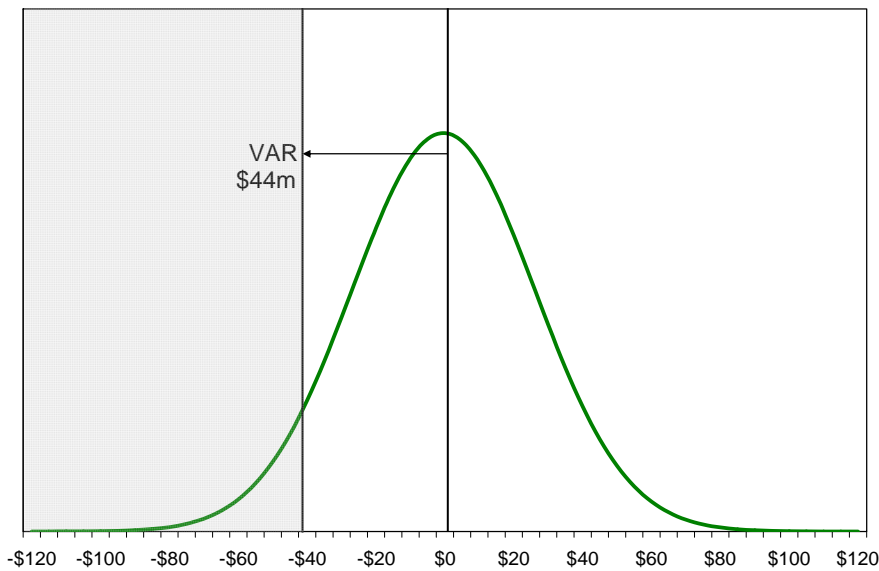


Exhibit 8 reports the quantiles, α , of a standardized normal distribution. For a confidence level of 95%, for example, the multiplier α is 1.645. Exhibit 8 also reports the multiplier corresponding to the conditional VAR. For a confidence level of 95%, this is 2.063. By construction, this number must be greater than α . The two numbers, however, are similar in magnitude. Here, CVAR is 25% greater than VAR.

Exhibit 8 Lower Quantiles of the Standardized Normal Distribution

	Confidence level (%)						
	99.99	99.9	99.0	97.5	95.0	90.0	50.0
Quantile ($-\alpha$)	-3.719	-3.090	-2.326	-1.960	-1.645	-1.282	-0.000
$E(Z Z<-\alpha)$	-3.957	-3.367	-2.665	-2.338	-2.063	-1.755	-0.798

If the risk manager believes the distribution of the variable under consideration is substantially different from a normal distribution (e.g., has fatter tails), the manager could use the multiplier for another parametric distribution, such as the Student t .⁷ In this case, the multiplier a

will be higher. More generally, the first three or four moments can be used to adjust the normal quantile using the *Cornish–Fisher expansion*. The Cornish-Fisher expansion is a method that allows us to estimate quantiles of an arbitrary distribution from its moments. We illustrate this method with the first three moments, up to the skewness S . The Cornish-Fisher expansion is

$$VAR = \alpha' \sigma \quad (15)$$

Here α' is related to the original α according to the following relationship:

$$\alpha' = \alpha + \frac{1}{6}(\alpha^2 - 1)S \quad (16)$$

As an example, with $S = -0.5$, the coefficient at the 95% level of confidence is increased from 1.645 to 1.787. More negative skewness indeed means that the distribution is more risky. With a normal distribution, $S = 0$ and α remains 1.645 as expected.⁸

The parametric approach must be more efficient than a non-parametric approach because it makes a strong assumption about the shape of the distribution (provided the assumption is correct). In contrast, a non-parametric approach makes no such hypothesis—other than assuming that the past is representative of the future.

The increase in the VAR precision can be traced to the fact that the computation of the standard deviation uses all the data points in the sample and, as a result, is estimated rather precisely. In contrast, the quantile only uses the values of the two numbers around the cutoff point. As a result, the sample quantile is much less precisely estimated, or has substantial *estimation error*. In other words, another data sample could yield a totally different number, especially if the confidence level is high. When VAR is estimated from the standard deviation, it is much less susceptible to variations in the data. Therefore, the parametric method is more efficient. This reflects the general principle in statistics that putting more structure on a model will give more precise results, provided that the assumptions are valid.

Choice of Horizon and Confidence Level

To measure risk, we need to define the horizon and, for VAR-type measures, the confidence level. Consider first the choice of the horizon. For trading portfolios, this is

typically short term, such as one day. For investment portfolios, the horizon is longer, typically from one month to one year.

Longer horizons increase risk measures. This can be shown in the case where returns are identically and independently distributed across subperiods. Consider, for example, the return on the short yen position previously discussed, but over two consecutive days. If we measure returns in logarithmic form, the two-day return is $R_{12} = \ln(S_2/S_0) = \ln(S_1/S_0) + \ln(S_2/S_1) = R_1 + R_2$ and the variance is

$$\sigma_{12}^2 = \sigma_1^2 + \sigma_2^2 + 2 \text{Cov}(R_1, R_2) \quad (17)$$

If returns are independent from one day to the next, the covariance term is zero. If distributions are identical, we have $\sigma_2^2 = \sigma_1^2$ and the 2-day variance reduces to $\sigma_{12}^2 = 2\sigma_1^2$. This shows that the variance increases linearly with time and thus the volatility increases with the square root of time. More generally, defining T as the number of days, we have

$$\sigma_T = \sigma_1 \sqrt{T} \quad (18)$$

The same adjustment applies to VAR when daily returns have normal distributions, because a linear combination of jointly normal variables is itself normal. As a result, both sides of Equation 18 can be multiplied by α , which gives the square root of time rule,

$$\text{VAR}_T = \text{VAR}_1 \sqrt{T} \quad (19)$$

For instance, in our hedge fund case, the daily VAR was \$43.7 million, assuming a normal density; extrapolating to one month, or 21 trading days, gives $\$43.7 \times 21 = \200.1 million. This assumes that daily returns are uncorrelated. In the case of the yen/dollar rate from before, this is indeed verified because the first-order correlation, or *autocorrelation*, coefficient is -0.034 only, with a standard error of 0.020 . The t -statistic is small, at $t = -0.034/0.020 = -1.7$, indicating the absence of statistical significance. This suggests that returns for one day are not useful to forecast returns the next day. Later, we will see that this assumption of zero autocorrelation does not hold well for less-liquid investments.

The choice of the horizon depends on the use of VAR. If the goal is to provide an accurate

measure of downside risk, the horizon should be relatively short, ideally less than the average period for major portfolio rebalancing. In contrast, if the VAR number is being used to decide how much capital to set aside to avoid bankruptcy, then a long horizon is advisable. This is because institutions need to have enough time for corrective action as problems start to develop. The Basel rules require a 10-day horizon for market risk and an annual horizon for credit and operational risk.

Next, we turn to the choice of the confidence level. The higher the confidence level is, the greater the VAR measure. Assuming a normal distribution, we can use the quantiles in Exhibit 8 to adjust the 95% VAR to a different confidence level. For example, the 99% VAR would be \$43.7 times (2.326/1.645), or \$61.8 million. From the empirical distribution, the non-parametric VAR is \$75.3 million. In this case, the normal-based VAR understates the empirical VAR.

As with the horizon, the choice of the confidence level depends on the use of VAR. If the goal is to provide a general measure of downside risk, the confidence level should not be too high, typically 95% or 99% as required by the Basel Committee. Here, what really matters is consistency of the VAR confidence level across trading desks or across time. In contrast, if the VAR number is being used to decide how much capital to set aside to avoid bankruptcy, then a high confidence level is advisable to keep the fund safe. The Basel rules require a 99.9% confidence level for credit and operational risk.

Institutions now routinely report measures of *economic capital*, which is the amount of capital an institution would voluntarily set aside to support its business activities. This is typically estimated as a VAR measure derived from the distribution of total profits and losses at a very high confidence level such as 99.97% over a year. This approach is fraught with problems, however. The first one is that the institution must take into account all of its risks and measure their distribution properly. The second is that very high confidence levels make it very difficult to estimate VAR measures precisely, because there are few, if any, observations in the left tail.⁹

Long-Term Capital Management (LTCM) is an example of a fund that blew up because it

did not have enough capital.¹⁰ At the beginning of 1998, LTCM thought that \$4.7 billion of capital was more than sufficient to absorb a worst-case situation. By August of that year, the fund had lost \$2.4 billion. It was unable to raise additional funds and did not materially change its risk profile. By September 23, 1998, the fund had lost another \$2 billion, forcing a Fed-orchestrated bailout. The portfolio managers had badly underestimated how much they could lose.

Backtesting

No risk measurement system would be complete without a process for backtesting. This involves systematic comparisons of the actual returns with the risk forecasts. With a well-calibrated system, the number of losses worse than VAR, also called exceptions, should correspond closely to the confidence level. For example, backtests of a 1-day VAR at the 99% level of confidence over a period of one year should yield, on average, two to three exceptions per year (more precisely, 1% times 252 trading days in a year, or about 2.5 observations). Too many exceptions should cause the risk manager to reexamine the models.

To implement backtests, the risk manager needs to construct a decision rule. As an example, the Basel Committee setup a simple system for verifying the risk numbers reported by banks, with a “green” zone for up to 4 exceptions, a “yellow” zone for 5 to 9, and a “red” zone for 10 and above. In other words, the model fails the backtest if we observe more than four exceptions over the last year. There is no perfect rule, however, in the presence of uncertainty. When the model is correctly specified, the probability of observing 5 or more exceptions is 10.8%, which is the Type I error rate. This reflects bad luck, perhaps unusually volatile markets.

More generally, a simple decision rule can be constructed as follows. Define x as the observed number of exceptions over the last T observations. If the VAR confidence level is $c = 1 - p$, we should expect to see pT exceptions on average. Then compute the statistic

$$z = \frac{x - pT}{\sqrt{p(1-p)T}} \quad (20)$$

This is approximately distributed as a standard normal variable. Hence, if z is too large (e.g., above 2), the model failed, with a Type I error rate of about 5%. The onus is then on the risk manager to understand why this has happened and how to improve the model.

Modeling Changes in Volatility

As we saw earlier, the estimate of volatility is a critical input for calculating VAR and other risk measures. Volatility can change over time, however, and needs to be monitored. Suppose we observe N daily observations on the rate of return, r , of an asset and we wish to forecast the variance over the next day, t . The conventional method for computing the variance is, from Equation 8,

$$\sigma_t^2 = \frac{1}{N-1} \sum_{i=1}^N (r_{t-i} - \mu)^2 \quad (21)$$

In this expression, all observations have the same weight. If the financial environment changes, however, it is more appropriate to assign relatively higher weights to the most recent observations. A popular approach to such weighting scheme is the *exponentially weighted moving average* (EWMA) model, where the variance forecast is

$$\sigma_t^2 = (1-\lambda) \sum_{i=1}^N \lambda^{i-1} (r_{t-i} - \mu)^2 + \lambda^N \sigma_{t-N}^2 \quad (22)$$

Here, λ must be assigned a value between 0 and 1. This *decay factor* determines the pattern of weights, which decreases as the observation gets older. If N is large enough, the last term, will be negligible. The EWMA volatility is typically written in the recursive form

$$\sigma_t^2 = (1-\lambda)(r_{t-1} - \mu)^2 + \lambda\sigma_{t-1}^2 \quad (23)$$

Hence, the variance forecast is a weighted average of the recent innovation squared and of the previous day's variance.

As an example, suppose $\lambda = 0.94$ and that the latest volatility forecast is 1%. Assume that we observe a change in price, away from the mean, of 3%. The new volatility forecast is then

$$\sigma_t^2 = (1 - 0.94) \times 0.03^2 + 0.94 \times 0.01^2 = 0.0001480$$

$$\sigma_t = \sqrt{0.0001480} = 1.22\%$$

This shows that a shock of a size greater than the current volatility of 1% pushes up the volatility forecast from 1% to 1.22%. The extent of this effect depends on the decay factor, λ . A lower value assigns more weight to recent observations.¹¹

The EWMA model is a special case of the class of volatility models known as *generalized autoregressive conditional heteroskedastic* (GARCH). In the GARCH (1,1) model, the day t forecast includes one lag of the innovation and one lag of the variance

$$\sigma_t^2 = \omega + \alpha \times (r_{t-1} - \mu)^2 + \beta \times \sigma_{t-1}^2 \quad (24)$$

In the case of the EWMA model, $\omega = 0$, $\alpha = 1 - \lambda$, $\beta = \lambda$. In the GARCH model, the constant can be interpreted as a long-run variance $\bar{\sigma}^2$ times $(1 - \alpha - \beta)$. In addition, the GARCH (1,1) model does not force α and β to sum to one, which generates more realistic dynamics in the variance forecast. The GARCH forecast is a weighted average of the long-run variance, of the squared innovation, and of the previous variance.

To illustrate, suppose a GARCH(1,1) is estimated with the following parameters:

$$\sigma_t^2 = 0.000006 + 0.05 \times (r_{t-1} - \mu)^2 + 0.90 \times \sigma_{t-1}^2$$

As in the previous example, suppose that the current standard deviation is 1% and the current excess return is 3%. The estimated volatility for day t will be

$$\sigma_t^2 = 0.000006 + 0.05 \times 0.03^2 + 0.90 \times 0.01^2 = 0.0001410$$

$$\sigma_t = \sqrt{0.0001410} = 1.19\%$$

Further, these coefficients imply a long-run volatility of


$$\bar{\sigma}^2 = \frac{\omega}{(1 - \alpha - \beta)} = \frac{0.000006}{(1 - 0.05 - 0.90)} = 0.000120$$

$$\bar{\sigma} = \sqrt{0.000120} = 1.10\%$$

To illustrate, Exhibit 9 displays the GARCH volatility forecast for the S&P 500 Stock Index. The average volatility is 1.2% daily, which translates into approximately 19% annually.

There are wide fluctuations, however, around this average, which are tracked by this GARCH model. In particular, volatility spiked up to more than 5% (or 80% annually) after the Lehman bankruptcy in September 2008. Other periods, in particular 2004 to 2006, were particularly quiet, with volatility below 1%. In summary, these models do adapt to changing financial environments and allow more responsive measures of risk.

Exhibit 9 Daily Volatility Forecast (GARCH) for the S&P 500 Index



RISK MEASUREMENT METHODS

In this section, the most common risk measurement methods are described. These include VAR approaches, risk decomposition, and stress tests.

VAR Approaches

Three major methods are used for computing VAR across large portfolios. The methods can be generally classified into *linear methods* and *full valuation methods*. Linear methods

replace the positions by their linear exposures on risk factors (e.g., bonds by their dollar duration and options by their delta). Full valuation methods, in contrast, revalue all the instruments for the new values of the risk factors. Such methods are more complex and take longer to run but are generally more accurate.

The first method is called *delta-normal*, or *variance-covariance*. This involves, first, a linear mapping of the positions onto the risk factors, resulting in a vector of dollar exposures, x . This is a linear valuation method. Next, the risk manager computes the covariance matrix of the risk factors, Σ , typically from historical data, which can be constructed to place more weight on more recent data, as in the EWMA approach. The variance of the portfolio is then computed from $\sigma_p^2 = x' \Sigma x$. Assuming, for instance, a normal distribution gives

$$VAR = \alpha \sigma_p = \alpha \sqrt{x' \Sigma x} \quad (25)$$

Take for instance a bond portfolio with a value of \$1 million and duration of 10 years. The dollar exposure x to movements in yield is then the dollar duration, as in Equation 2. Suppose that the volatility of monthly changes in 10-year Treasury yields has been $\sigma=0.31\%$. The 95% monthly VAR is then $\alpha \times (V \times D) \times \sigma = 1.645 \times \$1\text{MM} \times 10 \times 0.31\% = \$50,995$.

This method is very simple and quick to implement. Unfortunately, it is inappropriate if the portfolio has nonlinear instruments such as options, or if its distribution is strongly non-normal. If it is symmetric, however, a simple solution is to use the multiplier α from a distribution with fatter tails.

The second method is called *historical simulation*. This is a full valuation method that simulates movements in the risk factors from their recent history. The current portfolio value is P_t , which is a function of N current risk factors at time t , $P_t = P[f_{1,t}, f_{2,t}, \dots, f_{N,t}]$. We sample the changes in factor movements from the historical distribution, without replacement. The first change $k=1$ comes from yesterday's movements $j = t-1$, the second from the day before, and so on,

$$\Delta f^k = \{\Delta f_{1,j}, \Delta f_{2,j}, \dots, \Delta f_{N,j}\} \quad (26)$$

Next, we construct hypothetical factor values, starting from the current ones. For factor i , the construction is $f_i^k = f_{i,t} + \Delta f_{i,j}$, which is used to reprice the portfolio, $P^k = P[f_1^k, f_2^k, \dots, f_N^k]$. Suppose, for instance, that the current yield is now at 1.89%. Last month, it moved from 2.08 to 1.89, which is a change of -0.19%. We then apply this change to the current value, giving $1.89 - 0.19 = 1.70\%$ and reprice the portfolio using this hypothetical yield. This creates a hypothetical gain of \$19,009. Repeating this operation using the entire historical window, we can then sort the portfolio values to build the distribution of returns and report VAR as the sample quantile.

This method is intuitive because losses can be traced to a particular historical episode. It does not assume a normal distribution and instead uses the actual historical distribution. It can also handle options. These properties explain why this is the most widely used method for VAR. On the other hand, the method relies on a short moving window (typically one to four years) to infer the factor distribution. If this window omits some major risks or covers an unusually quiet period, the method will understate risk. To illustrate, let us go back to Exhibit 9. At the end of 2006, a *backward-looking window* based on the last year would have indicated a very low risk level. This understated risk for 2007 and 2008. In response, the Basel Committee [2009a] now requires banks to compute their capital requirements from a combination of the usual VAR measure as well as a *stressed VAR*, which uses factor moves over a continuous 12-month period of significant financial stress.

The third method is called *Monte Carlo simulation* and is very similar to the historical simulation period except that factor movements are sampled from a prespecified distribution,

$$\Delta f^k \approx g(\Delta f; \theta) \quad (27)$$

where g is the joint distribution and θ the parameters. We could run millions of simulated scenarios k , each case revaluing the entire portfolio. VAR is then computed from the distribution of changes in portfolio values. This method is very flexible because it can accommodate many types of stochastic processes. It will, however, take more computational

time and is less intuitive than other methods. Mistakes in the specification of the model are not as easy to identify. Thus, this approach is more powerful, but is subject to *model risk*.

Risk Decomposition

The goal of risk measurement systems should be to provide much more than a single summary measure of risk. They should also help the portfolio manager understand the sources of risk and drill down to the level of subportfolios and even individual positions. *Marginal risk* provides such information, representing the change in risk due to a small increase in one of the allocations. For simplicity, we can focus on risk measures that are based on the standard deviation because these lead to analytical expressions. Define x_i as the size of the dollar position in asset or risk factor i . Using $\text{VAR} = \alpha \sigma_P W$ as the risk measure, the marginal risk of position i in portfolio P , MRISK, is defined as the partial derivative

$$\text{MRISK}_i = \frac{\partial \text{VAR}}{\partial x_i} = \frac{\partial (\alpha \sigma_P W)}{\partial x_i} = \alpha \frac{\text{Cov}(R_i, R_P)}{\sigma_P} = \alpha \beta_{i,P} \sigma_P = \alpha \rho_{i,P} \sigma_i \quad (28)$$

The MRISK of an allocation is given by $(\text{VAR}/W)\beta_{i,P}$. This means that the change in the VAR of a portfolio resulting from a small change in the size of a position is proportional to the *beta* of the position with respect to the portfolio (note that this beta is not the traditional systematic risk exposure, which is the beta to a major stock market index.)

MRISK is a unitless measure because it is constructed as the ratio of a dollar VAR to a dollar change in the position. Here, β is defined from a regression of risk factor i on the portfolio. A large value for β indicates that a small addition to this position will have a relatively large effect on the portfolio risk. Hence, positions with large betas should be cut first because they will lead to the greatest reduction in risk. Alternatively, the positions can be kept in the portfolio if they have comparatively high expected returns. Whichever the choice, the portfolio manager should be fully aware of the risk implications of the positions.

This tool can be expanded to measure the contribution to the portfolio risk, CRISK, which is obtained by multiplying the marginal risk for position i by its weight in the portfolio,

$$\text{CRISK}_i = x_i(\alpha\beta_{i,p}\sigma_p) = x_i \times \text{MRISK}_i \quad (29)$$

Component VAR is measured in dollars, as is VAR. Given the definition of MRISK, we can see that the CRISK of a position is $\text{VAR} \times w_i \times \beta_{i,p}$, where w_i is the weight of position i in the portfolio. We can write x_i in terms of w_i times the dollar value W of the portfolio: $x_i = w_i W$. Because the beta of a portfolio with itself is one, the weighted sum of $w_i \beta_{i,p}$ across the N risk factors is guaranteed to be one. Hence, this proves that the sum of the risk contributions adds up exactly to the total portfolio risk, RISK,

$$\text{RISK} = \alpha\sigma_p W = \alpha\left(\sum_{i=1}^N w_i \beta_{i,p}\right)\sigma_p W = \sum_{i=1}^N x_i(\alpha\beta_{i,p}\sigma_p) = \sum_{i=1}^N \text{CRISK}_i \quad (30)$$

Therefore, we have shown how to decompose RISK into an additive and exhaustive decomposition.

Component VAR provides an additive decomposition of the portfolio VAR. This decomposition is not obvious because it depends on the weight of each risk factor in the portfolio, its volatility, and its correlation to the entire portfolio. Positions that hedge the portfolio risk will have negative component VAR. Positions can be ranked in order of decreasing importance of component VAR. Those at the top, generally above 5% of the total, are called *hot spots*. They should be closely examined by the portfolio manager because they contribute most to the risk of the portfolio. The portfolio manager should make sure that these are not unintended bets, but rather that they are justified by views.

As an example, consider our previous portfolio that was short \$4 billion in yen, to which is added a long position of \$1 billion in euros. The two currencies have a slightly positive correlation of 0.28. Exhibit 10 displays the risk decomposition. Recall that the stand-alone position in the yen had a daily VAR of \$43.7 million at the 95% level of confidence. The combined portfolio now has a total VAR of \$41.9 million, which is lower due to diversification effects. For the yen position, the marginal VAR is the change in portfolio VAR after adding \$1 million to the position. If so, VAR changes from \$41.9392 to \$41.9286, which is a change of -0.0106 . Therefore, the negative marginal VAR entry for the yen indicates that the adding to the position, or bringing it towards zero, should reduce risk.

Exhibit 10 Risk Decomposition of Currency Portfolio (\$ Millions)

	Market Value x_i	Volatility σ_i	Risk $\alpha\sigma_i W$	Marginal Risk MRISK _{<i>i</i>}	Component Risk x_i MRISK _{<i>i</i>}
\$/Yen	-\$4,000	0.66%	\$43.7	-0.0106	\$42.4
\$/EUR	\$1,000	0.63%	\$10.4	-0.0005	-\$0.5
Total	-\$3,000		\$41.9		\$41.9

Next, multiplying this by the position of -\$4 billion gives a component VAR of \$42.4 million for the yen. The component VAR for the euro is negative, reflecting the diversification benefit of adding the second currency to the portfolio.

In this case, the risk decomposition clearly shows that the risk of the total portfolio is driven by the position in the yen. The portfolio manager should have a strong view on the yen to justify the risk taken. In contrast, the position in the euro can be justified simply on risk reduction grounds.

More generally, this analysis can be done in reverse. *Risk budgeting* is the process by which an investor selects a total risk budget for the fund that is then parceled out to various investments and positions. In this case, the focus is on the risk allocation instead of the usual market value allocation.

Stress Tests

As previously mentioned, the statistical distribution of risk factors is typically estimated over a short historical window. This may miss major movements in risk factors that occur infrequently. As a result, VAR measures must be complemented by stress tests. Risk managers typically assess extreme scenarios, such as the stock market crash of 1987, currency devaluations, the credit crisis that started in 2007, and so on.

In fact, the Federal Reserve Bank applied a stress test to large U.S. banks to ascertain

whether they could absorb losses in an adverse economic environment during 2009 and 2010. Scenarios that form the basis of stress tests can be taken from *historical* episodes. Alternatively, *prospective scenarios* are built from scratch, specifically tailored to the fund's portfolio. The Basel Committee [2009b] described how to construct stress scenarios. VAR systems can easily accommodate new scenarios that are handled just like any other period in the historical simulation window.

Scenario analysis is also routinely used to set margin requirements by prime brokers and clearing counterparties, often in combination with VAR measures. As an example, consider a portfolio with many short and long option positions on the same underlying asset. In this case, notional amounts are rather meaningless because some of the positions could be fully or partially offsetting each other. The risk could be much greater, or less, than the net amount initially invested. The counterparty would make sure that the margin requirement is sufficient by building a battery of scenarios with a range of movements in the asset price and its implied volatility. The entire portfolio is repriced in each scenario. The margin is then set as the worst loss across all scenarios. The advantage of scenarios is that they can help to uncover situations that are plausible, yet have no recent historical precedent. Thus, stress tests are absolutely necessary complements to statistical risk measures such as VAR.

ILLIQUIDITY

So far, we have assumed that the balance sheet of the fund was rather liquid. This is generally the case for some categories of alternative investments, such as global macro funds, commodity trading advisors (CTAs), and long-short equity funds. These funds invest in major currencies, large stocks, and Treasury bills and bonds, which are very liquid. Some *over-the-counter* (OTC) instruments, such as corporate debt, are generally less liquid because they trade infrequently, e.g. once a week. At the lower end of the liquidity range are real estate funds, private equity funds, and venture capital funds, where transactions cannot be conducted for years.

Illiquidity and Risk Measures

Risk measures are negatively affected by asset illiquidity, which is the risk of losses due to the market impact of liquidating the positions. Illiquid assets trade infrequently. They have wide bid–ask spreads and large price impact. The *price impact function* describes how far down the price would have to move to sell a specific position.

Instrument liquidity risk creates a major problem for the measurement of risk. After all, risk measures represent potential changes in market prices. If historical prices do not change frequently enough, traditional risk measures cannot be accurate. Worse, they will tend to underestimate the true economics risks.

Consider, for example, a private equity fund that invests in distressed debt (i.e., debt issued by companies in financial distress or in bankruptcy). This debt trades infrequently, perhaps once a month. Typically, these funds report their net asset value at the end of each month. If the bonds are not liquid, it is unlikely that all bonds will have market-clearing prices on the last day of the month. Instead, the valuation could be based on a trade in the middle of the month. This is why the end-of-month price is called *stale*. Unfortunately, this distorts several risk measures.¹²

The first effect is that the reported monthly volatility is biased downward. This occurs because prices are based on trades during the month, which is similar to an averaging process. Movements in monthly averages are less volatile than movements based on end-of-period values. As an example, a moving average of a price with a window of 20 days will be smoother than the most recent price.

The second effect is that monthly changes will display positive autocorrelation meaning that movement in one direction during one month will tend to persist the following month because they are not fully captured by reported prices in the first month.

To illustrate, take an example where a bond value goes from \$100 to \$110 by the end of the next month and then stays at \$110 the second month. We only observe prices mid-month, say \$105 during the first month and \$110 during the second. The true monthly returns are +10% and 0%. Instead, the observed returns are +5% and approximately +5%. This proves

our two points. First, the volatility of the observed returns is less than that of true returns. Second, the two observed returns are highly correlated.

This autocorrelation can be measured using the regression,

$$R_t = \alpha + \rho R_{t-1} + \varepsilon \quad (31)$$

where ρ is the autocorrelation coefficient. It is called first-order because it relates returns to those lagged by one period. In practice, positive values above 0.1 indicate potential illiquidity problems.

This positive autocorrelation substantially increases the volatility over longer horizons. Consider the example in Equation 17 where we extrapolated the one-period volatility to two periods. Initially, we assumed that movements were uncorrelated across periods, which led to the square-root-of-time rule, i.e., an adjustment of $\sqrt{2}=1.41$. Now assume a non-zero first-order autocorrelation coefficient ρ . The multiple- period variance is now

$$\sigma_{12}^2 = \sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2 = \sigma_1^2 + \sigma_1^2 + 2\rho\sigma_1^2 = \sigma_1^2 2(1 + \rho) \quad (32)$$

With an autocorrelation of $\rho= 0.5$, the adjustment factor to the volatility changes from $\sqrt{2} = 1.41$ to $\sqrt{2(1 + \rho)} = 1.73$. Thus, the risk should be higher by $(1.73-1.41)/1.41 = 22\%$. In general, the variance over N periods can be written as

$$V(\sum_{i=1}^N R_i) = \sigma_1^2 [N + 2(N - 1)\rho + 2(N - 2)\rho^2 + \dots + 2(1)\rho^{N-1}] \quad (33)$$

As a result, the widespread method of annualizing monthly data by multiplying by the square root of 12 understates the annual risk. This can be adjusted using Equation 33 instead.

Alternatively, we can construct an adjusted series,

$$R_t^* = \frac{1}{1 - \rho} R_t - \frac{\rho}{1 - \rho} R_{t-1} \quad (34)$$

When $\rho=0$, this collapses to the usual return $R_t^* = R_t$. A positive value for ρ increases the volatility of the adjusted series R_t^* .¹³ This adjustment method was originally developed to deal with the observed smoothing of real estate prices. Because of high transaction costs and the long time needed to close a real estate transaction, prices do not immediately adjust to

new information.

A third, related effect is that measures of systematic risk will be systematically biased downward. Consider an asset with a monthly return of R_t . If the market I_t goes up during a month, only a fraction of this increase will be reflected in the NAV, leading to a beta measure that is too low. This can be corrected using lags,

$$R_t = \alpha + \beta_0 I_t + \beta_{-1} I_{t-1} + \beta_{-2} I_{t-2} + \varepsilon_t \quad (35)$$

The corrected beta is then the sum of the contemporaneous beta plus betas on lagged values of the index. This beta, called *Dimson beta*, is defined as $\hat{\beta} = \beta_0 + \beta_{-1} + \beta_{-2}$. Here we arbitrarily included two lagged values of the index. In practice, lags would be added up to the point where their coefficient is no longer significant.

More generally, correlations of illiquid assets with other asset classes are biased downward. This is a serious issue when “low correlations” are used as a major argument for investing in new asset classes.

Equations 34 and 35 provide an adjustment to risk measures for short-term returns, typically monthly. Another, simpler, approach is to extend the return interval (e.g., to take quarterly or even annual steps instead of monthly steps). Unfortunately, this approach leads to less-precise risk estimates because the number of independent data points shrinks quickly. For example, 10 years of monthly data yield 120 data points for the monthly volatility and beta estimates. Using annual returns creates 10 data points only.

Even with these statistical adjustments, historical data have limitations. For private equity (PE) funds, valuations are based on unrealized, as well as realized investments, and thus introduce noise and potential biases due to subjective accounting treatment. In this case, position-based information can be useful as well. For private equity, positions in nontraded stocks could be replaced by positions in traded stocks in equivalent industries, countries, and of like size. This mapping process would certainly create better risk measures than those based on investments carried at cost. Ljungqvist and Richardson [2003], for example, estimated the systematic risk of PE funds by identifying the companies held in each fund and assigning

them the beta of publicly traded firms in the same industry. They report an average beta of 1.1. Hence, PE funds that are more leveraged than typical public equities can have high beta. Ljungqvist and Richardson also found that PE funds tend to be concentrated in one or two industries, which creates higher risk.

Illiquidity can have a major effect on the risk-adjusted performance of alternative investments. For instance, performance is often evaluated with the Sharpe ratio (SR), which is the ratio of the average return on the portfolio \bar{R}_p in excess of the risk-free rate over the volatility,

$$SR_p = \frac{\bar{R}_p - R_F}{\sigma_p} \tag{36}$$

Exhibit 11 compares the total returns on indices representing (1) publicly traded U.S. stocks (S & P 500 Index); (2) hedge funds (CSFB Global Index); and (3) private equity (PE) funds (Cambridge Associates Index). The exhibit displays the annualized performance estimated from quarterly data measured over the period from 1994 to 2008. For example, the usual risk measures panel shows that the volatility of PE is 11.7%. This seems lower than the 16.8% risk of U.S. stocks; similarly, the PE beta is only 0.54. These numbers are misleading, however, because the autocorrelation of the PE index is very high, at 0.45.

Exhibit 11 Comparison of Performance of U.S. Stocks, Hedge Funds, and Private Equity, 1994–2008 (from quarterly data, annualized)

Asset Class	Average	Conventional Risk Measures			Auto-correlation	Adjusted Risk Measures		
		Std.Dev.	Beta	Sharpe Ratio		Std.Dev.	Beta	Sharpe Ratio
US Stocks	7.7%	16.8%	1.00	0.24	0.09	16.8%	1.00	0.24
Hedge Funds	8.9%	8.9%	0.35	0.57	0.23	10.6%	0.35	0.48
Private Equity	14.0%	11.7%	0.54	0.88	0.45	16.2%	0.86	0.64

Source: Author’s computations. The “usual risk measures” transform the quarterly standard deviation to an annualized measure by multiplying by the square root of four. The “adjusted risk measures” take autocorrelation into account and accordingly adjust the standard deviation, beta, and Sharpe ratio.

The right panel reports adjusted risk measures using Equations 33 and 35, the latter with three lags. The volatility and beta of the PE index are now markedly higher, at 16.2% and 0.86, respectively. Taking annual steps produces similar results, with estimates of 17.4% and 1.07, respectively. As a result, the Sharpe ratio of which appeared several times higher than that of U.S. stocks, drops from 0.88 to 0.64, a considerable difference.

For the hedge fund index, the corrections are minor. The autocorrelation is small, leading to slightly higher volatility. There are no significant lags on the market. The Sharpe ratio drops from 0.57 to 0.48, which is a smaller change. The risk-adjusted performance is still twice that for U.S. stocks.

Thus it is important to correct for illiquidity effects when evaluating risk-adjusted performance. Conroy and Harris [2007] reached even stronger conclusions. Based on a number of other indices over the period 1989 to 2005, they show that the volatility and beta of private equity are higher than that of U.S. stocks. As a result, they argue that when correctly adjusted for risk, the performance of private equity has been hardly better than that of U.S. equities. Similarly, Jegadeesh et al. [2010] examine the prices of publicly traded funds that invest in unlisted PE funds. This allows them to avoid the selection bias problem that affects all PE databases. They find that unlisted PE funds are expected to earn positive abnormal returns of approximately 0.50% per year only.

Forced Liquidation Risk

Illiquidity causes another type of risk, which cannot be as easily measured as market risk. Funds that are leveraged may face funding requirements that could force them to sell assets in order to raise cash. Thus, *funding liquidity risk*, which arises when the firm cannot meet cash flow or collateral needs, can cause *asset liquidity risk*, which is the risk of losses due to the price impact of large asset sales. Liquidity risk, however, is complex cannot be reduced to simple quantitative rules.

Commercial banks are naturally exposed to this type of risk. On the liability side, they

raise deposits, a form of short-term debt, that are used to invest in long-term assets, such as loans. Even if the bank is solvent (i.e., the value of assets exceed the value of liabilities), it might run into difficulties if depositors demand their money all at one time or, in other words, a “bank run” occurs.

Hedge funds are also exposed to liquidation risk, especially when they have high leverage. Exhibit 12 links sources of liquidation risk to a hedge fund balance sheet. Asset liquidity risk arises on the asset side and is a function of the size of the positions as well as of the price impact of a trade. On the liability side, funding risk arises when the hedge fund cannot rollover funding from its broker or when losses in marked-to-market positions or increases in haircuts lead to cash requirements for additional margin. Liability funding risk is a major source of risk for hedge funds because failing to meet a margin call can cause a lender to seize the collateral for the margin loan, and thus forcing liquidation of the fund. In these situations, the portfolio manager loses control of the investment strategy, which can lead to a *blowup*. Funding risk also arises when the fund faces investor redemptions.

Exhibit 12 Balance Sheet and Sources of Liquidation Risk for a Hedge Fund

Assets	Liabilities
Size of positions	Funding
Price impact	Mark-to-market, haircuts
	Equity
	Investor redemptions

Alternatives managers typically try to manage their liquidation risk by matching the horizon of their assets and liabilities. Funds that invest in highly liquid assets, such as CTAs that deal in exchange-traded futures, can allow daily investor redemptions. But funds that invest in illiquid securities, such as distressed debt, impose long *lockup periods*, meaning that investors cannot redeem their investment for an extended, predetermined period of time. For hedge funds, lockup periods average three months but they can extend up to five years. When

redemptions are allowed, a minimum *notice period* can be required. Funds also often have *gate* that limit the amount of withdrawals each period to a fraction of the equity investment. In extreme cases, funds generally have the ability to impose an outright *suspension* of redemptions.

For private equity funds, whether illiquidity is a problem depends on the capital structure. Some private equity categories have no leverage. An example is a *venture capital* fund that makes equity investments in start-up ventures. Because there is no debt, the asset side is matched with the liability side, which consists of investor equity that may not be redeemed for a long period. Other categories have leverage. The best example consists of *leveraged buyouts* (LBOs), in which public firms go private by repurchasing all outstanding shares. The acquisition is financed by a large proportion of debt, typically from 60% to 80% of the transaction value. This can include *senior debt* and *subordinated debt*, also called *mezzanine debt*. In this case, the risk is not being able to roll over the debt. Often, however, a large fraction of debt consists of short-term bank bridge loans that may have to be repaid after only two years. The short-term nature of these loans can cause liquidity problems. During the market turmoil that started in 2007, bank refinancing did indeed become very difficult. In response, PE firms issued *capital calls* to their investors.¹⁴ Such capital calls help PE firms manage their liquidity risk.¹⁵

LIMITATIONS OF CONVENTIONAL RISK MEASURES

A number of limitations are associated with conventional risk measures. This section discusses the limitations and highlights what to look for.

General Limitations

A good risk manager should be keenly aware of the limitations of conventional risk measures. First, although statistical risk measures, such as VAR, are designed to give a sense of the potential extent of losses, they certainly do not describe the absolute worst loss. The risk manager can increase the confidence level so as to experience fewer exceptions but this will

create other problems. Due to the paucity of data in the tails, the VAR measures are increasingly unreliable at higher confidence levels, even when distributions are stationary.

Second, modern risk measures are based on current positions that are assumed fixed over the time horizon. In practice, dynamic trading could increase or decrease risk. Such changes can be identified by backtesting both actual returns and *hypothetical returns*. The latter recreate the holding-period return assuming a frozen portfolio. If the backtest fails for actual returns, but not hypothetical returns, the risk manager can conclude that the model is well calibrated but that actual trading increases the risk profile.

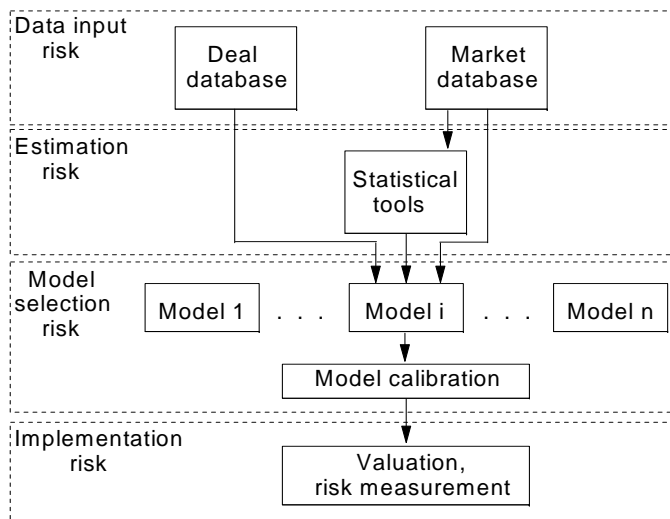
Third, as previously mentioned, all risk systems involve simplifications, obtained by mapping the positions on the selected risk factors. These simplifications could create “holes” in the risk systems. Many hedge funds, for example, take positions in corporate bonds that are hedged by purchasing *credit default swaps* (CDS). Normally, losses on the bonds should be offset by gains on the CDS. If the risk system maps both bonds and CDS on the same curve, the net exposure is zero, so that there appears to be no risk. During 2008, however, the basis between bond and CDS spreads widened sharply, causing mark-to-market losses for many funds. These losses were not anticipated by most risk systems. The design of risk management systems depends on the trading strategy and requires experienced risk managers.

More generally, *model risk* can occur at various stages of the risk management process. Exhibit 13 shows that errors can arise when trades and market data are entered into the system, when risk factors are statistically modeled, during the mapping process, and even during implementation.

Finally, most statistical risk measures assume that the recent past is a good representation of the future. This may not be the case, however, if the recent past has been unusually quiet or if it contains none of the events that are likely to develop in the future. This may not be the case, however, if the recent past has been unusually quiet or if it contains none of the events that are likely to develop in the future. As the Counterparty Risk Management Policy Group [2008] put it, “Risk monitoring and risk management cannot be left to quantitative risk metrics, which by nature are backward looking.” This is why stress tests are required in addition to the

statistical risk measures. The need for stress tests is an issue particularly in a period of rising volatility. Models based on simple moving averages respond slowly to these changes and systematically underestimate future risk. Models such as the EWMA that place more weight on more recent data will respond more quickly to rising volatility.

Exhibit 13 Model Risk



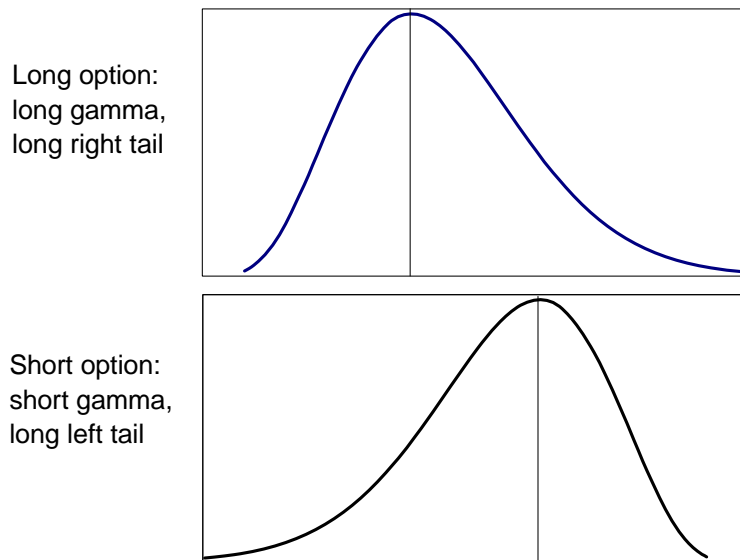
Finally, most statistical risk measures assume that the recent past is a good representation of the future. This may not be the case, however, if the recent past has been unusually quiet or if it contains none of the events that are likely to develop in the future. This may not be the case, however, if the recent past has been unusually quiet or if it contains none of the events that are likely to develop in the future. As the Counterparty Risk Management Policy Group [2008] put it, “Risk monitoring and risk management cannot be left to quantitative risk metrics, which by nature are backward looking.” This is why stress tests are required in addition to the statistical risk measures. The need for stress tests is an issue particularly in a period of rising volatility. Models based on simple moving averages respond slowly to these changes and systematically underestimate future risk. Models such as the EWMA that place more weight on more recent data will respond more quickly to rising volatility.

Even so, some risks are totally outside the scope of most scenarios. In 2008, many risk models failed largely due to “unknown unknowns.” One type of unknown is *regulatory risk*. Two examples of regulatory risk are sudden restrictions on short-sales, which wreaked havoc on hedging strategies in 2008, and structural changes, such as the conversion of investment banks to commercial banks that accelerated the deleveraging of the financial industry. Similarly, it is difficult to account fully for *counterparty risk*. It is not enough to know your counterparty. You need to know your counterparty’s counterparties too. In other words, these risks are network externalities. Understanding the full consequences of Lehman Brothers’ failure would have required information on the entire topology of the financial network. Such contagion effects transform traditional risks into *systemic risk*, which can only be handled by the regulators or the government.

Things to Watch For

Risk managers should thoroughly understand the risk profile of the investment strategy. Some types of investments, such as small stocks or private equity, involve an upfront investment that can be returned several times if successful. This strategy is similar to a *long option* position, in which the upfront payment is the maximum loss. As shown in Exhibit 14, this type of distribution has a long right tail, or positive skewness, which is a desirable feature. Long option positions can only lose the premium paid, but can generate a return many times the amount of the premium.

Such distributions can also be created by dynamic trading. For example, adding to a position after experiencing gains replicates the payoff from a long option position. This is typical of many *trend-following systems*. Similarly, *stop-loss rules* cut positions after losses are incurred. Traditional risk measures, however, assume that the portfolio is fixed and may miss this behavior. In such cases, traditional risk measures will overestimate risk, which is conservative.

Exhibit 14 Distribution of Investment Payoffs

Such distributions can also be created by dynamic trading. For example, adding to a position after experiencing gains replicates the payoff from a long option position. This is typical of many *trend-following systems*. Similarly, *stop-loss rules* cut positions after losses are incurred. Traditional risk measures, however, assume that the portfolio is fixed and may miss this behavior. In such cases, traditional risk measures will overestimate risk, which is conservative.

Much more dangerous are situations in which the portfolio manager holds instruments with embedded *short option* positions or when dynamic trading replicates a short option position. In these cases, the distribution has a long left tail, meaning that the investment is exposed to very large losses.

These positions are, unfortunately, quite common. For instance, any investment in a credit-sensitive bond can be decomposed into a long position in a risk-free bond plus a short position in a credit default swap. The best that can happen is that all the coupons and principal will be paid back in time, in which case the actual return is basically the yield to maturity. So, the upside is limited. In contrast, a default can completely wipe out the investment.

The problem is that portfolio managers tend to like short positions in options because they deliver consistent outperformance as long as the options are not exercised. A good example is that of Bernard Madoff, who reportedly lost \$21 billion of investor money. He attracted so much money because his funds delivered good and steady returns, which turned out to have been fabricated. This fraud is the largest Ponzi scheme in history.

A different example is that of UBS. During 2007, UBS suffered losses of \$ 19 billion on super- senior, triple-A rated, tranches of pools backed by mortgage debt, also called *asset-backed securities* (ABS). Investing in super-senior tranches can be viewed as selling out-of-the-money put options. As long as the real estate market continues to rise, the default rate on sub-prime debt will remain relatively low and the super-senior debt remains safe, experiencing no price volatility. However, if the real estate market corrects, as it did in 2007, the put options will move in-the-money and lead to large losses on the super senior debt.

None of these movements were captured in the previous historical data because of the sustained appreciation in the housing market until 2007 as well as because of the inherent nonlinearity in these securities. Instead of modeling how these structures depended on real estate prices, some banks simply chose to “map” the super-senior debt on AAA-rated corporate bond curves. Such gross simplification assumed that these tranches had no credit risk and it also totally ignored their nonlinearities - an example of a flawed mapping process.

As a result, UBS did not impose internal risk-based capital charges for units of the bank that invested in these asset-backed securities. Because these securities returned a wide spread over LIBOR and because internal capital was only charged LIBOR, this was an arbitrage opportunity. Not unexpectedly, these securities found their way into the CDO warehousing book, into the trading book, into the liquid Treasury book, and into a hedge fund subsidiary. As reported later by UBS [2008], there was no monitoring of net or gross concentrations of positions in this asset class at the firm-wide level. By the start of 2007, the notional exposure had grown to approximately \$100 billion. At the time, UBS had about \$33 billion in tangible equity capital.

In conclusion, while the credit crisis that started in 2007 admittedly led to extreme and totally unexpected movements in risk factors, there were also notable failures in some risk models. As the SEC Senior Supervisor Group [2008] report indicates, financial institutions that did poorly used outdated or inflexible assumptions in their risk models. These examples demonstrate the limitations of risk systems. Risk managers should be aware of potential weaknesses in conventional risk measures and continuously reassess their effectiveness.

TRANSPARENCY

This section addresses the problems of non- transparency in alternative investments and discusses possible solutions.

Problems with Non-Transparency

Managers of alternative investments are generally reluctant to reveal information about their positions, but this lack of transparency has serious disadvantages for investors. Disclosure allows *risk monitoring* of the fund, which is especially useful with active trading by helping to avoid situations in which the portfolio manager unexpectedly increases leverage or changes style. Closer monitoring of the fund can also decrease the probability of fraud and, more generally, blowups.¹⁶

Disclosure is also important for *risk aggregation*. The investor should know how the fund interacts with other assets in the portfolio. Whether the fund has a positive or negative correlation with the rest of the portfolio affects the total portfolio risk.

In 2008, two blue-ribbon private-sector committees established by the President's Working Group (PWG) released separate sets of best practices for hedge fund investments. One report reflected the viewpoint of asset managers; the other report was written by investors. The two reports offer strikingly different perspectives on the need for disclosures and transparency. The investor committee (PWG, 2008a) states, "A key concern for investors is that hedge funds' lack of transparency may lead to unexpected risk exposures. ... Hedge fund managers typically cite commercial reasons for providing little transparency. There are

sometimes legitimate competitive reasons for keeping information confidential, but often there are not.” The term “transparency” is mentioned 16 times in this report, as opposed to “confidential,” which is mentioned only once. In contrast, the term “transparency” is not mentioned even once in the asset manager report (PWG, 2008b), as opposed to the term “confidential” which is mentioned eight times.

Greater disclosure is resisted on the grounds that it would reveal proprietary information, leading to the possibility that a third-party might trade against the fund. This threat, however, comes from the broker-dealer community, generally not from investors. If this is an issue, confidentiality agreements should prevent leakages of sensitive information. AI managers generally prefer to release such information, whether directly or through affiliates, to investors who have no trading operations and who would not be able to profit from these data. Recipients of position-level information should have internal controls to prevent the dissemination and inappropriate use of this information.

Another argument that is sometimes advanced is the lack of investor sophistication. In other words, disclosing positions would give too much information to investors who might not be able to use it. This is a “paternalistic” view, however. Many investors do have the capabilities to use the information and should have the choice to do so.

A final and more subtle argument is that requiring transparency creates a *selection bias* in managers. It is sometimes asserted that the very best hedge fund managers have all the assets they need and do not need to offer transparency. If so, requiring transparency reduces the pool of hedge fund managers to those with lower performance. Aggarwal and Jorion (2012), however, report that this argument has no empirical support.

Solutions for Transparency

These arguments can be addressed with a number of solutions, in particular for hedge funds. The first consists of external risk measurement services. These firms receive the individual positions of funds, after signing the proper nondisclosure agreements, and provide aggregate risk measures to investors. This solution partially solves the problems of risk

aggregation and managers' widespread reluctance to disclose detailed information about their positions. However, risk service providers have little incentive to model risk as accurately as possible because they do not have a stake in the portfolio performance; for example, they rarely perform backtesting.

Another solution, which is still fairly rare, is to invest through a fund of funds that has position-level information. A fund of funds with no related trading operations is more likely to earn the trust of hedge fund managers. Also, large funds of funds should have the capabilities to process this information, because building risk systems is a complex undertaking that benefits from economies of scale. As a result, such funds of funds can perform the risk monitoring and measurement function for the investor. This position-level information can also be used to provide independent checks on the valuation of assets in the portfolio and to improve the portfolio construction process, thereby justifying the added fee for the fund of funds.

CONCLUSION

The alternative investments industry has thrived because of its good performance, which is explained by a combination of investment flexibility and strong financial incentives for fund managers. These features, however, should also cause concerns because they may lead fund managers to take on too much risk. Indeed, hedge fund failures, or blowups, seem to occur on a regular basis. Risk should be managed at the level of the fund, by the portfolio manager, and at the investor level, either directly or indirectly through risk aggregation services or funds of funds.

Relative to the traditional asset management industry, however, risk management is a special challenge for alternatives. Alternative products run the entire gamut of investment styles. At one end are CTAs with frenetic trading activity. At the other extreme are private equity funds, where investments are not traded, hard to value, and locked for years.

In each case, risk measures ideally should be based on position-level information because returns-based risk measures have severe drawbacks. First, the length of the time series may

not be long enough for meaningful risk analysis. Second, risk measures based on older returns may no longer be relevant. This is particularly problematic given the wide investment latitude some managers have and how quickly they can trade in and out of positions. Amaranth Advisors LLC, for instance, started as a convertible bond trading fund and then morphed into a predominantly highly leveraged natural gas trading fund. Such change would be very difficult to identify from returns data.

Illiquid assets pose different problems. Stale prices create biases in risk measures, causing volatility and systematic risk to be understated. This has implications for the role of these assets in portfolio allocation and for risk-adjusted performance measures.

Overall, this chapter has described several approaches to manage risk. Risk managers should use exposure measures, statistical risk measures, and stress tests. As we have seen, the design of an effective risk system requires a thorough understanding of the underlying trading strategies. It requires simplifications that recognize the tradeoff between speed and accuracy. Overall, risk management for alternatives is still as much an art as a science. Using common sense is important when interpreting risk numbers.

ENDNOTES

¹Note that there is no covariance term between the market and residual effects because these are independent by virtue of the regression framework; in addition, there are no covariance terms between residual effects because these are assumed to be independent of each other.

²This can be proved in the simple case where all weights are the same $w=1/N$ and all the residual variances are equal. As N increases, the second term then becomes $\sum_{i=1}^N w_i^2 V(\varepsilon_i) = \sum_{i=1}^N (1/N)^2 V(\varepsilon) = N(1/N)^2 V(\varepsilon) = (1/N)V(\varepsilon)$, which goes to zero as the number of stocks N increases. So, residual risk is diversifiable, unlike market exposure.

³Sometimes, this number is expressed in terms of deviations from the mean. In this example, the mean is close to zero and, as a result, the VAR number would not change much.

⁴The Basel Committee on Banking Supervision consists of central bankers from the Group of Ten (G-10) countries. It prescribes minimum standards to regulate internationally active commercial banks.

⁵Note that 1.645 is the standard normal deviate for a one-tailed probability of 95%. For a normal distribution, the deviate for a two-tailed probability of 95% is 1.96. This is the case because 2.5% of the distribution is below -1.96 and 2.5% is above $+1.96$. So, the usual association of α around 2 for a 95% confidence level corresponds to a two-tailed test.

⁶Note that the number 0.664% is the previously reported standard deviation of returns of \$26.6 million expressed as a percentage of the portfolio value of \$4 billion.

⁷The Student t is a symmetric probability distribution where the thickness of tails depends on a parameter N called degrees of freedom. As N tends to infinity, the distribution tends to the normal pdf. As N decreases, the distribution has increasingly fatter tails. Thus the parameter can be chosen to fit the empirical data.

⁸The expanded form of the Cornish–Fisher formula calculates VAR using kurtosis as well.

⁹Rebonato [2007] provided a lucid criticism of economic capital measures, which he calls

“science fiction” numbers. For instance, assessing an empirical VAR measure at the 99.97% level of confidence would require 3 observations in the left tail out of 10,000 annual observations.

¹⁰Jorion [2000] provided a risk management perspective of LTCM.

¹¹RiskMetrics uses the EWMA model with $\lambda = 0.94$ to estimate the daily volatility of various instruments.

¹² See also Getmansky et al. [2004].

¹³ Peterson and Grier [2006] explain how to adjust returns series that are artificially smooth for the purpose of computing covariance matrices, which are essential inputs into asset allocation.

¹⁴A capital call occurs when a PE manager, usually the general partner, requests than an investor in the fund (a limited partner) provides additional capital. When entering a new PE investment, a limited partner typically injects initial funding and also agrees to provide additional capital over time, up to a maximum amount.

¹⁵The other side of the coin, however, is that investors may be forced to invest additional money precisely at the same time as turmoil in financial markets is creating losses on the rest of their portfolio. Siegel [2008] argued that these liquidity considerations are important when evaluating allocations across asset classes.

¹⁶ Christory, Daul, and Giraud [2006] examine the characteristics of hedge funds that blowup. Over the 1994 to 2004, they report an average probability of default of 0.30% per annum. Most of the blowups observed are attributed to operational problems such as fraud, which can be minimized through a due diligence process and continuous monitoring.

REFERENCES

- Agarwal, Vikas, Naveen Daniel, and Narayan Naik. 2009. “Role of Managerial Incentives and Discretion in Hedge Fund Performance.” *Journal of Finance* 64, 2221--2256.
- Aggarwal, Raj and Philippe Jorion. 2012. “Is There a Cost to Transparency?” *Financial Analysts Journal* 68, 108--123.
- Basel Committee on Banking Supervision. 2009a. *Revisions to the Basel II Market Risk Framework*. BIS, Basel.
- Basel Committee on Banking Supervision. 2009b. *Principles for Sound Stress Testing Practices and Supervision*. BIS, Basel.
- Basel Committee on Banking Supervision. 1996. *Amendment to the Basel Capital Accord to Incorporate Market Risk*. BIS, Basel.
- Christory, Corentin, Stéphane Daul, and Jean-René Giraud. 2006. “Quantification of Hedge Fund Default Risk.” *Journal of Alternative Investments* 9, 71--86.
- Conroy, Robert and Robert Harris. 2007. “How Good are Private Equity Returns?” *Journal of Applied Corporate Finance* 19, 96--108.
- Counterparty Risk Management Policy Group. 2008. *Containing Systemic Risk: The Road to Reform*. CRMPG, NY.
- Dimson, Elroy. 1979. “Risk Measurement When Shares are Subject to Infrequent Trading.” *Journal of Financial Economics* 7, 197--226.
- Getmansky, Mila, Andrew Lo, and Igor Makarov. 2004. “An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns.” *Journal of Financial Economics* 74, 529--609.
- Jegadeesh, Narasimhan, Roman Kräussl, and Joshua Pollet. 2010. “Risk and Expected Returns of Private Equity Investments: Evidence Based on Market Prices.” *Working Paper*, Emory University.
- Jorion, Philippe. 2007. “Risk management for hedge funds with position information.” *Journal of Portfolio Management* 34, 127--134.

- Jorion, Philippe. 2006. *Value at Risk: The New Benchmark to Manage Financial Risk*. McGrawHill, NY.
- Jorion, Philippe. 2000. “Risk Management Lessons from Long-Term Capital Management.” *European Financial Management* 6, 277—300.
- Ljungqvist, Alexander and Matt Richardson. 2003. “The Cash Flow, Return and Risk Characteristics of Private Equity.” *Working Paper*, NYU.
- Lo, Andrew. 2001. “Risk Management for Hedge Funds: Introduction and Overview.” *Financial Analysts Journal* 57, 16--33.
- Morgan, J.P. 1995. *Riskmetrics Technical Manual*. J.P. Morgan, New York.
- Peterson, Steven and John Grier. 2006. “Covariance Misspecification in Asset Allocation.” *Financial Analysts Journal* 62, 76--85.
- President’s Working Group on Financial Markets. 2008a. *Report of the Investors’ Committee: Principles and Best Practices for Hedge Fund Investors*. PWGFM, Washington.
- President’s Working Group on Financial Markets. 2008b. *Report of the Asset Managers’ Committee: Best Practices for the Hedge Fund Industry*. PWGFM, Washington.
- Rebonato, Riccardo. 2007. *Plight of the Fortune Tellers*. Princeton University Press, Princeton.
- Senior Supervisors Group. 2008. *Observations on Risk Management Practices during the Recent Market Turbulence*. SEC, Washington.
- Siegel, Laurence. 2008. “Alternatives and Liquidity: Will Spending and Capital Calls Eat Your ‘Modern’ Portfolio?” *Journal of Portfolio Management* 35, 103--114.
- UBS. 2008. *Shareholder Report on UBS's Write-Downs*. UBS, Zurich.