20

The "At-Risk" Metrics and Measures

Christopher L. Culp
The University of Chicago Booth School of Business

The risk-management process is the process by which an organisation tries to ensure that the risks to which it is exposed are the risks to which it thinks it is and is willing to be exposed. A typical risk-management process begins with the identification of key risks and the articulation by senior managers and directors of the firm's tolerances to those risks. The risk-management process also involves a framework for rectifying unintended deviations between actual and desired risk exposures. That process includes the measurement, reporting, monitoring and control of risks, as well as supervision, oversight, audit and (if necessary) retuning the process (Culp and Mackay 1994; Culp and Planchat 2000; Culp 2001b).

In the risk-measurement part of the risk-management process, risk managers utilise both risk metrics and risk measures. A risk metric is some conceptual way to describe or summarise perceived risk in a single number, whereas a risk measure is an expression or estimate of what the risk manager thinks the risk actually is at the time of measurement (Holton 2003, 2004). A wide range of risk metrics are available for summarising different aspects of financial and insurance risks,[1] and a comparably diverse set of approaches and methods exist for the estimation of risk measures corresponding to those risk metrics. The appropriateness and usefulness of various risk metrics often depend on the type of risk being measured (eg, market, credit, operational or liquidity risk) and the strategic risk-management objective of the institution using the risk metric.

The "At-Risks" are a popular set of risk metrics and measures that enable risk managers to quantify and summarise the risk that an asset, liability, portfolio, project or enterprise may experience an unexpected decline in value or cashflows as market prices change. The original risk metric in the At-Risk family was capital-at-risk (CaR), which summarises (with a certain degree of confidence) the risk that adverse market price movements can precipitate losses in a business unit (or whole enterprise) over a given time period.[2] Early uses of CaR typically involved relatively long time horizons and high confidence levels. For example, a risk manager might select a 99% confidence level for the risk of losses over the next year (with credits applied for potential future changes in the portfolio[3]), in which case the resulting CaR estimate is the annual loss that is not expected to be exceeded more than 1% of the time. So, if the 99% one-year CaR is US$10 million, the risk manager expects with 99% confidence that the business line's loss over the next year will be less than US$10 million. CaR was initially used by banks to ascertain how much capital was needed to support the risk-taking activities of a business unit and to facilitate comparisons of the required capital to support the amount at risk with the expected return on that capital. Indeed, CaR is still used by some institutions as the basis for allocating risk capital and engaging in risk-adjusted performance measurement.

Shortly after the introduction of CaR, some market participants also began to embrace the value-at-risk (VaR) metric. Unlike CaR, VaR originally focused on summarising the risk that a fixed portfolio (ie, with no changes presumed to the current portfolio's composition) might decline in market value over a relatively short time horizon (eg, one day) with a particular level of confidence.[4] Many risk managers, senior executives and regulators embraced the use of VaR for market risk measurement in the early-to-mid-1990s,[5] and its use has proliferated since then. VaR is appealing as a risk metric in part because it is easy to interpret, provides a summary measure of potential extreme losses that can easily be digested by senior managers and directors, and facilitates generally consistent comparisons of risk measures across different financial instruments, business activities and (sometimes) financial institutions.

This reference chapter provides a brief overview to the At-Risk family of risk metrics and measures.[6] We begin with an introduction to the history and basic principles of estimation for traditional VaR. We next consider several important and popular variations on the VaR metric intended to help firms better quantify extreme potential losses in the "tails" of risk distributions. The third section discusses At-Risk measures that are conceptually based on the VaR risk metric but that are used by firms concerned less with market value risk and more with the risks of declines in cashflows or earnings. We then briefly explore modern uses of the CaR metric, both as an economic concept and from a regulatory perspective. A final section concludes and comments on the potential shortcomings of the At-Risks and how those deficiencies can be at least partially addressed through the operation of a robust risk-management process. Although the At-Risks are not risk-measurement panaceas, when they are properly interpreted and appropriately integrated into a well-developed risk-management process they can play an indispensable role in helping risk managers assess whether or not current risk profiles are commensurate with the stated risk targets of their institutions.

VALUE-AT-RISK

History and context

VaR first emerged in the late 1980s and early 1990s as a risk-management tool in the banking community for measuring the market risk of derivatives portfolios. The original goal of VaR was to systematise the measurement of an active trading firm's risk exposures across portfolios of different underlying asset classes (eg, foreign exchange, commodities, interest rates) and financial instruments (eg, futures, swaps, options and securities) for a fixed portfolio over a short time period. VaR made it possible for dealers to estimate enterprise-wide risk and portfolio-level risks that could be compared across trading areas as a means of monitoring and limiting their consolidated financial risks in a consistent and integrated manner.[7]

VaR received its first significant public endorsement in July 1993, when the Global Derivatives Study Group of the Group of Thirty[8] urged derivatives dealers to "use a consistent measure to calculate daily the market risk of their derivatives positions and compare it to market risk limits". The Study Group further recommended that "[m]arket risk is best measured as 'value at risk' using probability analysis based upon a common confidence interval (eg, two standard deviations) and time horizon (eg, a one-day exposure)" (G30 1993).

The Group of Thirty recommendation highlights several specific features of VaR that account for its widespread usage. One such feature of VaR is that it is a consistent metric for describing financial risk. By expressing risk using a possible dollar loss over a specific time horizon (with a specified degree of confidence), VaR facilitates direct comparisons of risk across different business lines and distinct financial products and, sometimes, across financial institutions.

VaR was not, of course, the only risk metric available at that time. Other popular such risk indicia included volatility, the option Greeks (ie, delta, gamma, vega, theta and rho), the dollar value of a one-point price move (DV01) and the present value of a one-point price move (PV01).[9] Apart from differences arising from magnitudes and units of measure, volatility, the Greeks and the DV01/PV01 risk metrics were (and are) useful ways of examining the sensitivities of the prices of various financial instruments to underlying sources of market risk (like asset prices and volatilities) (Culp 2004). Unlike VaR, however, static risk metrics like volatility and the Greeks are not usually "forward-looking" measures of risk.[10] By contrast, VaR is probability-based and forward-looking. With whatever degree of confidence a firm wants to specify, VaR enables the firm to associate a specific loss with that level of confidence in evaluating future potential changes in portfolio value.

Another feature of VaR is its reliance on a common and specific time horizon called the "risk horizon". The risk horizon is chosen by the firm engaging in the VaR calculation and is often based on the amount of time a firm believes it would take to hedge or liquidate a losing position. A one-day risk horizon at, say, the 95% confidence level tells the firm, strictly speaking, that it can expect to lose no more than, say, US$X on the next day with 95% confidence. Firms often go on to assume that the 95% confidence level means they stand to lose in excess of US$X on no more than five days out of the next 100, an inference that holds when certain assumptions are made about the underlying probability distribution.[11]

Another important assumption in almost all VaR calculations is that the portfolio is fixed – ie, the positions in the portfolio do not change over the risk horizon. This assumption of no turnover was not a major issue when VaR first arrived on the scene at derivatives dealers because they were focused on one- or two-day risk horizons and thus found VaR both easy to implement and relatively realistic. But, when it comes to generalising VaR to a longer time horizon (see, eg, the subsequent discussion of the European Union's Solvency II risk-based capital regulatory regime for (re)insurance companies), the assumption of a fixed portfolio can be more problematic.

Estimation

To estimate the VaR of a given portfolio, a firm must generate a probability distribution of possible changes in the value of that portfolio (ie, the "risk distribution") over a specific risk horizon. The VaR of the portfolio is the dollar loss corresponding to some predefined probability level – usually 5%, 1% or less – as defined by the left-hand tail of the distribution.

More formally, denote the current value of some portfolio as v_t. The percentage change in the value of this portfolio from time t through time $t+\tau$ (ie, over the next τ periods) can be defined as follows:[12]

$$R_{t+\tau} = \frac{v_{t+\tau} - v_t}{v_t}$$

The VaR for this portfolio at a given time t then is the $\omega_{t,\tau}$ that solves the following equation:

$$Prob_t[v_t R_{t+\tau} \leq -\omega_{t,\tau}] = \alpha \qquad (20.1)$$

where 1–α% is the confidence level, and τ is the risk horizon. The estimated VaR $\omega_{t,\tau}$ thus is the maximum loss expected to occur (with 1–α% confidence) over the next τ periods. Notice that, in Equation 20.1, the inequality is compared to $-\omega_{t,\tau}$ but we define the VaR as $\omega_{t,\tau}$. We do this because VaR is traditionally expressed as a positive number – in other words, because VaR represents a potential loss, multiplying a negative by a negative yields a positive.

Empirical/historical VaR

The estimation of VaR using historical data is known as empirical or historical VaR (HVaR). If we are estimating HVaR using discrete distributions (eg, histograms of historical data), Equation 20.1 can be expressed more simply as:

$$\omega_{t,\tau} = -v_t Q_{\alpha,t+\tau} \tag{20.2}$$

where $Q_{\alpha,t+\tau}$ denotes the α^{th} percentile of the distribution of $R_{t+\tau}$. In other words, α% of the observations in the distribution of $R_{t+\tau}$ are below the VaR estimate $\omega_{t,\tau}$.

For example, Figure 20.1 shows the values for the S&P500 stock index (SPX) from January 2009 through April 2012. The black bars in Figure 20.2 define the empirical frequency distribution of the daily percentage changes in the value of the SPX over that time period.[13]

Figure 20.1 S&P500, 02/01/09–30/04/12

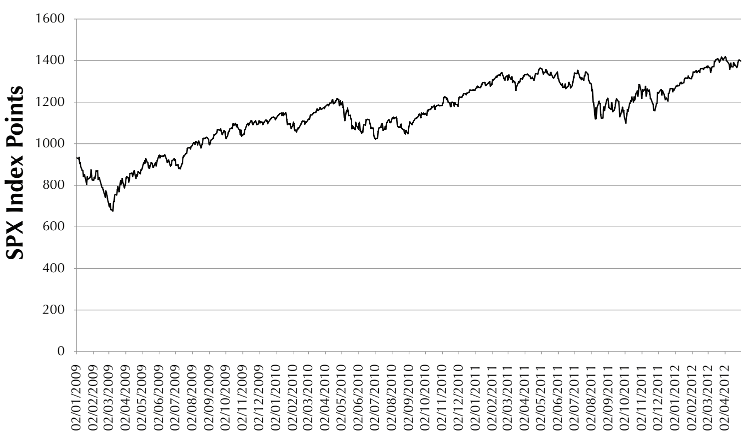

Source: Bloomberg

Suppose that a risk manager is calculating an estimate of the one-day VaR (at the 95% confidence level) on April 30, 2012, when the SPX closed at 1,397.91 and when the portfolio was valued at, say, US$1,397.91. If a risk manager assumes that this historical return distribution is a good representation of the possible percentage change in the SPX over the next day (ie, from April 30 to May 1), then the one-day VaR can be estimated from the 5th percentile return in Figure 20.2, which is −2.341% and is denoted on Figure 20.2 as a black diamond. The 95% confidence one-day VaR is then

$$\omega_{t,1} = -v_t Q_{\alpha,t+1} = (-\$1{,}397.91)(-2.341\%) = \$32.72$$

So, the risk manager expects with 95% confidence that the portfolio will not decline by more than US$32.72 between April 30, 2012, and May 1, 2012.

Figure 20.2 Distributions of daily percentage changes in the SPX, 02/01/09–30/04/12

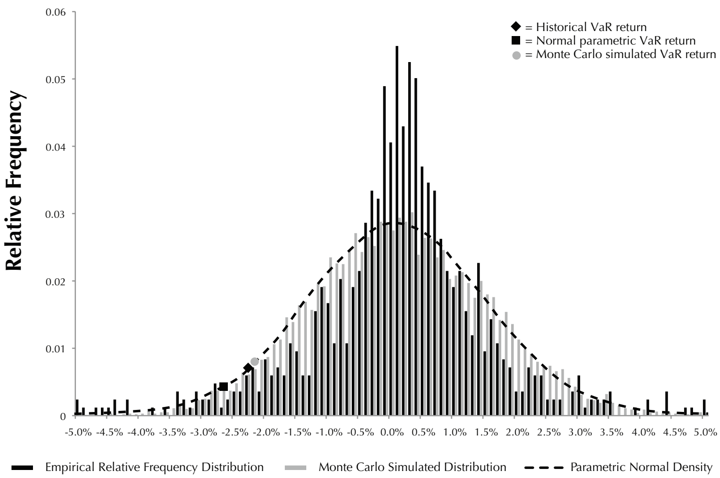

Parametric VaR

Instead of using historical data to generate the risk distribution, we can also describe the probability distribution of $R_{t+\tau}$ using some continuous probability density function. If that distribution at time t is denoted $f_t(R_{t+\tau})$, then the VaR estimate is the value $\omega_{t,\tau}$ that leaves α% of the probability distribution in the left-hand tail of the distribution. Expressed using only returns (ie, normalising the value of the portfolio to unity), the VaR is now the $\omega_{t,\tau}$ that solves

$$\int_{-\infty}^{-\omega_{t,\tau}} f_t(R_{t+\tau})dR_{t+\tau} = \alpha \qquad 20.3$$

Although Equation 20.3 can in principle be solved for any distribution $f_t(R_{t+\tau})$ by numerical integration, a more common approach is to estimate a parametric VaR (PVaR) by assuming a probability distribution for $R_{t+\tau}$ that can be described by a parameter vector $\boldsymbol{\theta}$ of sufficient statistics[14] that fully characterise the distribution.

For example, the ubiquitous normal distribution has sufficient statistics of location and scale (ie, mean μ and variance σ^2) that fully characterise the distribution – ie, $\boldsymbol{\theta}=\mu,\sigma^2$ are sufficient statistics for the normal distribution. So, if we assume that $R_{t+\tau}\sim N(\mu_t,\sigma_t^2)$, where μ_t is the time t conditional mean of the return distribution and σ_t^2 as the time t conditional variance, then we can express the one-day VaR $\omega_{t,1}$ as

$$\omega_{t,1} = -v_t(\mu_t - z_\alpha \sigma_t) \qquad 20.4$$

where z_α is the critical value of the standard normal distribution that leaves $\alpha\%$ of the probability of the distribution in the tail. If we further assume the presumed-normal return distribution is both independently and identically distributed and is stable over time (ie, $f_t(R)=f(R)\forall t \text{ where } R\sim NID(\mu,\sigma^2)$), then we can also calculate a multi-period extension of Equation 20.4 for users with risk horizons longer than one day. Specifically, the τ-day VaR is

$$\omega_{t,\tau} = -v_t(\mu\tau - z_\alpha \sigma\sqrt{\tau}) \qquad 20.5$$

The assumption of normally distributed changes or percentage changes in value became prevalent beginning in 1994, when JP Morgan introduced its RiskMetrics framework for VaR measurement. (JP Morgan-Reuters 1996) Based on a parametric normal estimation methodology, the RiskMetrics PVaR approach was popularised by the decision of JP Morgan to distribute certain volatility and correlation data that greatly simplified VaR estimation for smaller firms which did not necessarily collect or have easy access to the required data inputs for a sufficiently wide range of products.[15] Partly if not largely as a result of the widespread availability of the RiskMetrics data and the relative ease of implementing the approach, the use of normal PVaR boomed in the 1990s.[16]

Returning to Figure 20.2, we see that the black dashed line shaped like a bell curve represents the probability density function for daily percentage changes in the SPX if a risk manager assumes that those changes are drawn from a normal distribution. For comparability to the empirical histogram, we set μ=0.0006 and σ=0.0140, which correspond to the mean and standard deviation (respectively) of daily returns on the SPX over the sample period. In that case, the one-day VaR (at the 95% confidence level) can be estimated using the April 30, 2012, portfolio value of US$1,397.91 as

$$\omega_{t,1} = -v_t(\mu - z_\alpha \sigma) = -US\$1{,}397.91(0.0006 - 1.96 \cdot 0.0140) = US\$37.52$$

The SPX return that corresponds to the VaR estimate (ie, (0.0006−1.96·0.0140)) is −2.684%, which is indicated on Figure 20.2 with a black square.

That the PVaR based on the assumption of normally distributed changes in the SPX is not equal to the HVaR in the sample period is not surprising given the data. As Figure 20.2 illustrates, the empirical distribution is classically "leptokurtic" – ie, it has a more peaked centre and fatter tails than the normal distribution. In this particular case, the distribution is much heavier at the centre than the normal distribution, which pulls the HVaR upwards *vis-à-vis* the normal PVaR.

Simulation-based VaR

Figure 20.2 also shows a third approach to VaR estimation, which combines explicit distributional assumptions with historical parameter estimates. Known as simulation-based VaR, the approach involves the Monte Carlo simulation of a large number of potential one-day changes in the SPX based on an assumed stochastic process that governs the probabilistic evolution of the portfolio value (Picoult, 1998).

Specifically, suppose that changes in the index evolve over time according to a standard geometric Brownian motion process $dv = \mu v dt + v \sigma dZ$, where dZ is a Brownian motion (ie, $dZ \approx \varepsilon\sqrt{dt}$, where $\varepsilon \sim N(0,1)$). Using Itô's lemma, log changes in the index can be expressed as

$$dlnv = (\mu - \frac{1}{2}\sigma^2)dt + \sigma dZ$$

We then run 10,000 day-ahead simulations of this process for dt=1. To estimate the parameters of the process, we use historical daily log changes in the index from January 2009 through April 2012 from which we estimate μ=0.0001 and σ=0.0140. The grey bars in Figure 20.2 define the frequency distribution for these 10,000 simulated one-day percentage changes in the SPX. The 5th percentile of those simulated returns is −2.24%, denoted on Figure 20.2 as a grey circle. That implies a 95% confidence one-day VaR of $31.32:

$$\omega_{t,1} = -v_t Q_{\alpha,t+1} = (-\$1{,}397.91)(-2.24\%) = \$31.32$$

Choice of methodology and inputs and backtesting the model

The above three general VaR estimation approaches – ie, historical, parametric and simulation-based – are all based to some extent on the subjective judgement of the risk manager and the use of historical data. None of these methods is inherently more reliable than any other to measure the risk of uncertain future outcomes. In practice, variations across specific measurement methodologies, moreover, are considerably more diverse and complex than the simple illustrative examples discussed above.

VaR estimates are also sensitive to the input data used. If historical values are used to estimate parameters for PVaR estimates or Monte Carlo models or to generate whole empirical distributions, a significant choice facing a modeller is the relevant historical period (which, as we explain later, may also be affected by regulatory considerations). The risk manager faces the usual tradeoff endemic to empirical financial economics – ie, the use of more recent data may be better to capture short-term market conditions, whereas the use of a longer time series may better reflect long-term risks and generate a more conservative estimate of the tails of the return distribution. Many institutions use both – eg, define a daily volatility parameter as the higher of the volatility in the last 90 days and the volatility the week after Lehman Brothers failed. Prior to the credit crisis, an institution might have instead looked at the higher of the last 90 days and the volatility in September 1998 during which Long-Term Capital Management sustained large losses and precipitated a brief disruption in interbank funding markets.

When estimating PVaR, particular attention must be paid to how the volatility parameter is measured. Indeed, all three of the basic VaR estimation methods outlined above will experience problems if a constant volatility input is assumed. Significant empirical research has documented that virtually all asset markets exhibit time-varying volatility. As a result, "violations" of a VaR model (ie, situations in which the actual loss exceeds the VaR estimate over the relevant risk horizon) will tend to cluster during periods of increasing volatility. If a constant volatility is assumed, losses will tend to exceed the estimated VaR more often when volatility rises and less often when volatility falls.

Risk managers incorporate time-varying volatility into VaR calculations using a variety of different methods. For example, the aforementioned RiskMetrics approach to volatility estimation uses an exponentially weighted moving average volatility measure to give relatively higher weight to more recent data without throwing out older data (JP Morgan-Reuters 1996). An unconditional volatility estimate or a moving-average volatility, by contrast, would give no weight to observations outside the estimation window and would give equal weight to observations within the sample period. Other more complex methods of historical volatility estimation are also sometimes used, such as estimating volatility using one of the family of autoregressive conditional heteroscedasticity (ARCH) models.[17] Finally, many PVaR estimates utilise implied volatility from traded option prices as a more forward-looking and market-based measure of volatility.[18]

The "quality" of a VaR model is usually assessed based on the out-of-sample performance of the model using historical data. In other words, VaR models are "backtested" to determine how they would have performed on a given past date if the VaR was computed using information available only up to that past date. The results are then compared to actual changes in portfolio values to estimate various measures of risk coverage. Examples of typical VaR backtesting methods can be found in Kupiec (1995), Hendricks (1996), Christoffersen (1998), Diebold, Gunther, and Tay (1998), Lopez (1999), Berkowitz (2001), Campbell (2007), Van Roekel (2008), and Brown (2011).[19] In fact, many institutions rely on more than one VaR approach and then either evaluate all of the results or use backtesting to fine-tune their choice of model.

Absolute versus benchmark-relative VaR

VaR can be used to measure the absolute risk of a portfolio or the risk of a target portfolio relative to some other benchmark portfolio. For an organisation using VaR to help administer a system of absolute risk limits *vis-à-vis* predefined risk targets, absolute VaR is the most sensible choice. But for organisations whose risk-management objectives are defined relative to other organisations or portfolios, benchmark-relative VaR is usually more relevant.

For example, a well-capitalised, well-funded and highly liquid asset manager may be far more interested in managing its risk of underperforming a benchmark than its absolute exposure to losses. In such a situation, the same basic risk-measurement methods discussed above still apply, but now the definition of the underlying loss distribution will be a relative return distribution. Specifically, define

$$r^e_{t+\tau} = R^T_{t+\tau} - R^B_{t+\tau}$$

where $R^T_{t+\tau}$ is the return on the target portfolio from t to $t+\tau$ and $R^B_{t+\tau}$ is the return on the benchmark portfolio over the same period. The distribution of $r^e_{t+\tau}$ then serves as the basis for benchmark-relative VaR estimates.

One popular way to measure benchmark-relative VaR is to rely on the performance measure known as "tracking error", which is defined as the volatility of the difference in a return on some target investment portfolio T *vis-à-vis* the benchmark portfolio B to which portfolio T's performance is compared and assessed. Tracking error is defined for some performance evaluation horizon over the last k periods as

$$\sigma^{T,B}_{t,t-k} = \sqrt{\frac{1}{k-1}\sum_{j=0}^{k-1}\left(r^e_{t+\tau-j} - \mu^{T,B}\right)^2} \tag{20.6}$$

If we assume that $r^e_{t+\tau} \sim N(\mu^{T,B}, \sigma^{T,B})\ \forall t$, then we can express the τ-period Tracking Error VaR ("TE-VaR")[20] at the confidence level $(1-\alpha)$ as

$$\omega_{t,\tau} = -v_t(\mu^{T,B}\tau - z_\alpha \sigma^{T,B}\sqrt{\tau}) \tag{20.7}$$

VaR and market liquidity risk

As noted, VaR is principally a market risk metric. Closely related to market risk, however, is market liquidity risk (also known as asset liquidity risk) (G30 1993; Culp 2001b; Brunnermeier and Pedersen 2009). Market liquidity risk is the risk that a position cannot be hedged or liquidated (at all or at reasonable bid–ask spreads) during periods of significant market illiquidity – a situation that was all too often observed in late 2007 and 2008 during the spread of the global credit crisis.

Market liquidity risks can be difficult to model in a VaR framework, especially for risk horizons longer than a few days. Some institutions try to approximate this risk by assuming that an extreme realisation of the bid–ask spread is a suitable market liquidity risk proxy. A risk metric that is useful for such risk managers is liquidity-adjusted VaR (LVaR) (Jorion 2007). Specifically, suppose changes in the bid–ask spread at time t for some asset j (denoted $S_{t,j}$) are presumed normally distributed. The liquidity-shocked bid–ask spread can be defined for some confidence level $(1-\gamma)$ as

$$\frac{1}{2}V_{t,j}(\mu_{t,j,S}+z_{\gamma}\sigma_{t,j,S})$$

where $\mu_{t,j,S}$ and $\sigma_{t,j,S}$ are the mean and volatility of the spread, and where $V_{t,j}$ is the value of the asset. Notice that the spread is multiplied by one-half to reflect the fact that a liquidation would constitute only a sale and not a round turn, whereas the actual spread measures the cost of supplying immediacy for a round trip.[21] One-day parametric LVaR is then defined as[22]

$$LVaR_{t,1}=-V_{t,j}(\mu_{t,j}-z_{\alpha}\sigma_{t,j})+\frac{1}{2}V_{t,j}(\mu_{t,j,S}+z_{\gamma}\sigma_{t,j,S}) \qquad 20.8$$

If a firm prefers not to adopt the parametric normal approach to VaR measurement, spread risk as a measure of asset liquidation costs can be incorporated into the other approaches, as well. In a Monte Carlo LVaR simulation, for example, the spread can be presumed to evolve according to a stochastic process of its own that is correlated with the evolution of the asset price or return. Bivariate Monte Carlo then can be utilised to estimate LVaR. Similarly, historical spreads can be used to reflect liquidity costs in either historical or historical-simulation LVaR estimates.

The above liquidity-risk-adjusted measures of market risk are static because they assume no change in the composition of the portfolio over the VaR horizon. Consequently, the market impact is ignored, or, equivalently, is presumed not to affect the spread in a predictable manner. As noted, however, the time horizon of the liquidation also affects the risk of the asset sale. A slower liquidation implies better pricing and less market impact, but leaves the firm exposed to market risk for a longer period of time.

TAIL-BASED EXTENSIONS TO VAR

One drawback of VaR is that it does not provide information about the size of the potential losses in the tail of the risk distribution. If $\omega_{t,\tau}$ is the estimated VaR corresponding to a (1–α)% confidence level over τ periods, that means that α% of the potential losses are below $\omega_{t,\tau}$. If α=5%, τ=1 day, and $\omega_{t,\tau}$= US$1 million, for example, then 5% of potential daily losses are expected to exceed US$1 million. But those potential losses in excess of US$1 million could as easily include losses that are either one dollar above or that are significantly greater than US$1 million. As a result, some organisations adopt alternative risk metrics and measures either to augment or replace traditional VaR in an effort to try and quantify with greater precision the potential losses within the tail of the underlying risk distribution. We discuss two examples of such tail-based extensions to VaR in the sections below.

Expected shortfall/CVaR/TVaR

The most common tail-based risk metric used in financial and insurance applications is known (equivalently) as expected shortfall, conditional VaR (CVaR), or tail VaR (TVaR). TVaR summarises the conditional expected loss on a portfolio below its estimated VaR. TVaR as a risk metric thus relies on the estimated VaR measure. TVaR can be viewed as the average of all potential losses that are worse than the estimated VaR. For a continuous risk distribution $f_t(R_{t+\tau})$,

$$TVaR = -\frac{v_t}{\alpha} \int_{-\infty}^{-\omega_{t,\tau}} R_{t+\tau} f_t(R_{t+\tau}) dR_{t+\tau} \tag{20.9}$$

For example, if the 95% VaR over the next month is US$1 million, the corresponding TVaR would be the average of all potential portfolio losses when losses exceed US$1 million over the next month (collectively accounting for the remaining 5% of potential losses). For example, a VaR of US$1 million might correspond to a TVaR of

US$1.5 million. So, in situations where the loss exceeds US$1 million, the expected loss in excess of US$1 million is US$1.5 million. The TVaR is, however, a conditional expected loss, and actual losses could, of course, be much larger.

Extreme-value VaR

Instead of choosing a different risk metric such as TVaR to summarise the risks in the left-hand tail of a risk distribution, some prefer instead to rely on statistical techniques for measuring extreme values and drawing related probabilistic inferences using parametric approaches. VaR estimates based on extreme-value theory (EVT) are sometimes known as extreme-value VaR (EV-VaR) measures. Embrechts, Klüppelberg and Mikosch (2003), Dowd (2005), and McNeil, Frey and Embrechts (2005) offer excellent reviews of the most prevalent EV-VaR estimation methods, and Coles (2001) and de Haan and Ferreira (2006) provide good introductions to EVT more generally.[23] We discuss two popular EV-VaR estimation approaches below.

EV-VaR and the generalised extreme-value distribution

Suppose the cumulative distribution of one-period returns $F(r)=Prob[R\leq r]$ is unknown but that we can sample n observations from $F(r)$, denoted $R^{(1)},\ldots,R^{(n)}$.[24] We can then define the "maximum" as

$$M_n=max\{abs[R^{(1)},\ldots,R^{(n)}]\}$$

In other words, we treat the largest or smallest return in the sample of n observations as the maximum. Then we can define the distribution of the maximum return M_n asymptotically as follows:[25]

$$\lim_{n\to\infty} M_n = H_{\xi,\mu,\sigma}(r)=\begin{cases} exp\left\{-\left(1+\xi\left(\dfrac{R-\mu}{\sigma}\right)\right)^{-1/\xi}\right\} & if \quad \xi\neq 0 \\ exp\left\{-exp\left[\left(\dfrac{R-\mu}{\sigma}\right)\right]\right\} & if \quad \xi=0 \end{cases} \tag{20.10}$$

The distribution of $H_{\xi,\mu,\sigma}$ as described in Equation 20.10 is known as the generalised extreme-value (GEV) distribution. The parameters μ and σ are measures of the location and scale of the distribution of

M_n, and the parameter ξ is known as the "tail index" and describes the thickness and/or shape of the loss tail. When ξ=0, the GEV distribution is equivalent to the Gumbel distribution, which has tails that are similar to the tails of the normal or lognormal distribution. When ξ>0, the GEV distribution is equivalent to the Fréchet distribution, which exhibits fat tails *vis-à-vis* the normal distribution and thus is often considered a more realistic distribution of financial asset returns.[26] Normalised maximums drawn from the Student's t-distribution and the inverse gamma distribution also converge to the Fréchet distribution.

Note that the distribution $H_{\xi,\mu,\sigma}(r)$ is not a distribution of returns but is rather a distribution of the extreme losses (expressed as returns) from the distribution of returns. To obtain the corresponding VaR measure, we thus wish to find

$$Prob[M_n \leq M_n^*] = H_{\xi,\mu,\sigma} = (Prob[r \leq M_n^*])^n = (1-\alpha)^n$$

where α is the desired VaR confidence level and where M_n^* is some threshold extreme loss (Dowd 2005).

If we use the Gumbel distribution (ie, ξ=0), the one-period VaR at the 95% confidence level[27] can be computed follows:

$$\omega_{t,1,n} = -v_t[\mu - \sigma(ln(-nln(1-\alpha)))]$$

The VaR estimate is now subscripted with n to emphasise that the estimate (and parameters, μ, σ, and ξ) are specific to the extreme values drawn from a sample size of n.

For example, let n=100 and assume ξ=0.25. Using our previous parameter estimates of μ=0.0006 and σ=0.0140 for the S&P500, the 95th percentile EV-VaR (Gumbel) is

$$\begin{aligned} \omega_{t,1,100} &= -\$1,397.91[0.0006 - 0.0140(ln(-100ln(0.95)))] \\ &= -\$1,397.91 \cdot -2.23\% = \$31.17 \end{aligned}$$

Recall that the 95% VaR estimates discussed earlier using the historical, parametric and Monte Carlo simulation approaches yielded VaR measures of US$32.72, US$37.52 and US$31.32, respectively. So, the Gumbel EV-VaR is quite close to the Monte Carlo VaR estimate and within sight of the other two traditional VaR approaches.

Now suppose we instead use the Fréchet distribution, in which case EV-VaR is

$$\omega_{t,1,n} = -vt[\mu - \frac{\sigma}{\xi}(1-(-nln(1-\alpha))^{-\xi})] \quad \xi > 0$$

Again using our earlier parameter estimates (and now set $\xi=0.25$), the EV-VaR (Fréchet) is

$$\begin{aligned} \omega_{t,1,100} &= -\$1{,}397.91[0.0006 - \frac{0.0140}{0.25}(1-(-100ln0.95)) - 0.25 \\ &= -\$1{,}397.91 \cdot -1.8189\% = \$25.43 \end{aligned}$$

The Fréchet EV-VaR estimate of US$25.43 thus is distinctly smaller than the Gumbel EV-VaR estimate for the same set of parameter values, and is smaller than all of the traditional VaR measures reported earlier.

GEV EV-VaR estimates, however, are highly sensitive to the estimate of the tail index parameter. For example, if we assume instead that $\xi=-0.25$ (ie, the Weibull distribution), we compute an estimated EV-VaR of $40.36, which is now larger than any of the three traditional VaR estimates and the Gumbel and Fréchet EV-VaR estimates.

Indeed, one criticism of the parametric EV-VaR approach is precisely that it depends heavily on the value of the tail index parameter, which is usually estimated from historical data. But, if we are interested in extreme values, the historical data may be lacking those extreme values in sufficient quantities to yield a robust estimate for the tail index parameter. The original premise of VaR, after all, was that potential losses could be measured over short time horizons and fixed portfolios up to a high degree of confidence, but could not necessarily be measured confidently in the extreme tails of loss distributions.

EV-VaR and the generalised Pareto distribution

Another popular tail-based extension to traditional VaR is EV-VaR estimated using the generalised Pareto distribution (GPD). That approach begins with the distribution function $F(l)$ for an independent and identically distributed sample of losses l and a specified target

T. In returns space, l can be interpreted as the absolute values of all negative returns, so that negative returns become positive losses and positive returns are discarded. The conditional distribution of losses in excess of T is

$$Prob[l-T \leq L | l>T] = F_T(l) = \frac{F(l+T)-F(T)}{1-F(T)}$$

As the loss threshold T gets large, the conditional distribution $F_T(l)$ converges to the GPD:

$$\lim_{T \to \infty} F_T(l) = H_{\xi,\theta}(l)$$

where $H_{\xi,\theta}(l)$ is the GPD, or

$$H_{\xi,\theta}(l) = \begin{cases} 1 - e^{-l/\theta} & if \quad \xi = 0 \\ 1 - \left(1 + \dfrac{\xi l}{\theta}\right)^{-1/\xi} & if \quad \xi \neq 0 \end{cases}$$

which is defined for $\xi \geq 0$ over $l \geq 0$ and for $\xi < 0$ over $0 \leq l \leq -\frac{\theta}{\xi}$. The two parameters of the GPD, ξ and θ, represent the tail index parameter as discussed above and a scale parameter, respectively. Dowd (2005) shows how to estimate EV-VaR from the above. But he is also quick to note that the choice of the threshold T must be high enough to satisfy the convergence of the conditional loss distribution to the GPD and yet not so high that too few actual data points are available to obtain reliable parameter estimates – not an easy empirical problem to solve.

Below-target risk

A popular alternative to GEV- and GPD-based EV-VaR is to estimate downside risk based on the distribution of losses in excess of a target amount. One simple approach to estimate tail losses is known as below-target risk (BTR). For a chosen risk threshold or target T and a continuous risk distribution $f_t(R_{t+\tau})$ that can be estimated with numerical integration, BTR is defined as

$$BTR = -v_t \int_{-\infty}^{T} (T - R_{t+\tau})^2 f_t(R_{t+\tau})\, dR_{t+\tau} \qquad 20.11$$

BTR expressed in the simple form in Equation 20.11 is also known as below-target variance or downside semi-variance when $T=\mu$ (Culp, Mensink and Tanner, 1997; Culp and Mensink, 1999).

STRATEGIC RISK-MANAGEMENT OBJECTIVES AND THE AT-RISKS

To be of practical use for risk-management purposes, a risk metric must comport with the primary risk that an institution is seeking to manage. VaR presumes a focus by management on the risk of changes in the value of a portfolio, but this need not be (and often is not) a firm's primary strategic risk-management focus. Instead, firms may be more interested in managing cashflows- or earnings-at-risk.

Cashflows-at-risk

A cashflow risk manager uses risk management to reduce cashflow volatility and increase debt capacity.[28] Cashflow-based risk management is distinct from value-based risk management mainly in the sense of timing. A value risk manager is concerned about the value of a portfolio either at a specific point in time (eg, when debt must be retired) or over regular intervals (eg, monthly changes in value). A cashflow risk manager, by contrast, is concerned with cashflows whenever they might occur. The two are related, of course, because value is just the expected discounted present value of future cashflows. But a firm concerned with cashflows is concerned with them over time, not collapsed back to a single point or interval in time using present values.

Cashflow risk managers often rely on a VaR-like risk metric known as cashflows-at-risk (CFaR). Unlike VaR, the risk distribution underlying a CFaR calculation is the per-period change in cashflows or net available liquidity. For example, denote the change in the available funds or liquidity of a firm between measurement date t and future cashflow date τ as $\Delta\ell_{t+\tau}=\ell_{t+\tau}-\ell_t$, where ℓ_t denotes the firm's availability liquid funds and assets at time t. The CFaR is defined in a manner analogous to the VaR in Equation 20.1 as:

$$Prob_t[\Delta\ell_{t+\tau}\leq-\mathcal{L}_{t,\tau}]=\alpha \qquad 20.12$$

where $1-\alpha$ is the confidence level, and $\mathcal{L}_{t,\tau}$ is the estimated CFaR. $\mathcal{L}_{t,\tau}$ thus represents the depletion in funding liquidity that a firm expects to occur with $(1-\alpha)\%$ confidence over the next τ periods.

The basic mathematics and statistics for generating the distribution of $\Delta\ell_{t+\tau}$ in order to estimate the CFaR $\mathcal{L}_{t,\tau}$ are similar to the VaR estimation methods discussed earlier. In addition, CFaR estimates are commonly supplemented with liquidity stress tests (just as VaR estimates are often augmented with market risk stress tests, as well).

Earnings-at-risk

Apart from value and cashflow risk-management objectives, some firms also focus on reducing the variability of earnings as a primary risk-management objective. For such firms, the proper risk metric is generally earnings-at-risk (EaR).

The basic principles for measuring EaR are similar to those discussed above for VaR and CFaR, although the underlying earnings data presents certain specific complications for risk managers. For example, the relative infrequency of earnings data generally requires a longer risk horizon to match accounting disclosure and reporting horizons and the availability of input earnings data, which can limit the value of EaR as a short-term risk-management tool. In addition, EaR is highly sensitive to extant accounting guidance and rules, and it can be challenging to incorporate potential changes in or subjective assessments of compliance by auditors with such rules over long time horizons.

CAPITAL-AT-RISK

As noted in the introduction, the original at-risk metric was capital-at-risk (CaR). Risk capital may be defined most generally as financial capital (usually equity) or cash allocated to a risky portfolio of assets and/or liabilities to absorb potential shortfalls in net asset value (ie, losses) (Merton and Perold 1993). Some firms still utilise CaR today for the same purpose. In addition, most regulated banks and insurance companies are required to measure CaR for the purpose of complying with risk-based minimum capital requirements. These two conceptions of CaR, however, can be quite different.

Economic CaR

Economic CaR is an institution's risk metric for ascertaining the capital required to support a business line or portfolio and can be viewed as a fairly conservative estimate of the "maximum reasonable loss" on that portfolio or business line. Economic CaR thus can be viewed as similar to VaR, albeit not always with the usual VaR assumptions of a fixed portfolio and a short time horizon.

Economic CaR is sometimes used exclusively as a risk-management tool, much like the other At-Risks. In some cases, however, institutions may instead use CaR as the building-block measure of risk in the *ex ante* allocation of risk capital or as the risk metric in *ex post* risk-adjusted performance measurement (ie, a process used by some firms to quantify the performance of business units by comparing realised profits to the risk capital utilised to produce those losses) (eg, Zaik, Walter, Kelling and James 1996).

On an *ex ante* basis, CaR can be used for risk-adjusted capital allocation – ie, the process by which a firm allocates risk capital to business or trading units (Culp 2000). Some *ex ante* CaR allocation schemes are informal attempts by banks and insurance companies to assign risk capital based on a comparison of risk capital requirements to hurdle rates. Risk-adjusted return on capital (RAROC), for example, is a popular measure of return per unit of risk capital and is defined as the expected net economic income of a business line relative to its scaled economic CaR. RAROC is considered particularly appealing by some firms because of its close relation (when properly measured) to economic value-added (EVA) or shareholder value-added (SVA). Specifically, allocating risk capital to a business when its RAROC exceeds the firm's weighted-average cost of capital is synonymous with allocating capital to a business unit when the projected EVA of that business unit is positive (Culp 2006).

Risk capital allocation can also be a more formal process in which financial institutions define a total amount of risk capital and then allocate that capital to various business lines or portfolios in a "risk budget". The basic premise of a risk budget is to assign a target amount of risk to a portfolio commensurate with the strategic objectives and expected return target for that portfolio. Portfolio managers then have discretion as to how to "spend" their budgeted amount of risk on specific trades and strategies to try to achieve their expected return target (Rahl 2000; Pearson 2002; Litterman 2003).

The measurement of CaR for risk capital-allocation purposes gives rise to certain issues that are not necessarily important for institutions measuring the At-Risk of portfolios in isolation but that can be very important for firms using CaR to allocate scarce risk capital. In particular, not all measures of CaR for individual portfolios or business units result in a full allocation or utilisation of the enterprise-wide CaR of the firm (ie, risk capital may remain idle and unallocated depending on which CaR measure is used).

CaR can be measured on an undiversified, diversified, marginal or co-measurement basis (as well as a few others). Undiversified CaR is the CaR of a portfolio without regard to correlations with other portfolios or business units at the firm, whereas diversified CaR takes cross-portfolio correlations into account. Marginal CaR, in turn, is the difference between enterprise-wide CaR with and without a business line included in the calculation (Culp 2001). Undiversified, diversified, and marginal CaR measures, however, do not generally result in a full allocation of risk capital for a firm – ie, the sum of the individual CaR estimates assigned to each business unit is usually less than the enterprise-wide CaR.

Various alternative CaR measures have been proposed to overcome this problem, as discussed, for example, in Venter and Major (2003). One such measure is co-VaR, which is based on a risk metric like TVaR or XTVaR (Kreps 2005). TVaR (ie, expected shortfall) was discussed earlier, and XTVaR is simply the expected shortfall relative to the mean of the distribution (as opposed to the VaR as in TVaR). XTVaR is defined as

$$XTVaR = -\frac{1}{2} v_t \int_{-\infty}^{-\mu_k} R_{t+\tau} f_t(R_{t+\tau})\, dR_{t+\tau} \tag{20.13}$$

where μ_k is the average return for business unit k. If all business units in a firm have the same expected return μ_k, then $XTVaR \approx TVaR - \mu$. The co-TVaR or co-XTVaR of a business unit then is the specific expected shortfall (relative to a target or the mean, respectively) for a given business unit, and the sum of those measures should in principle equal the aggregate TVaR or XTVaR (Venter and Major, 2003).

Regulatory CaR

The At-Risks have been embraced by both banking and insurance regulators in the administration of minimum risk-based capital requirements for banks and insurance companies. Care must be taken, however, not to equate automatically regulatory CaR with economic CaR (as discussed in the previous section). Although the concepts and risk metrics are similar, the risk measures that emerge from an institution's economic CaR measurement process used for capital allocation and/or risk-adjusted performance measurement may be quite different from the regulatory CaR estimated by the same institution for purpose of compliance with regulatory capital requirements.

Banking

The foundational guidance for national banking regulators to assess the capital adequacy of banks is provided by the Basel Committee on Banking Supervision of the Bank for International Settlements (BIS). The original Basel Accord was promulgated in 1988 and specified minimum risk-based capital requirements for internationally active banking institutions. In the 1988 Accord, banks were required to hold capital to absorb potential losses arising from the credit risk of certain assets (BIS 1988). The 1988 Accord was silent on capital requirements for market and operational risks. The risk-based capital charges for credit-sensitive assets, moreover, were fixed charges based solely on the type of obligor giving rise to the credit exposure – eg, no capital was required for extensions of credit to OECD sovereign obligors, whereas credit extended to AAA-rated corporations carried a full 100% risk weight.

In the early 1990s, central bankers began to perceive a need for capital requirements related to the market risk of assets. Although banking regulators did not require banks to use CaR or VaR as market risk metrics, a strong bias began to emerge among regulators by the mid-1990s in favour of that approach. The first explicit use of VaR in regulatory capital requirements occurred with the 1996 market risk amendment to the Accord (BIS 2005).

The BIS began to permit and embrace the use of VaR for estimating capital requirements associated with market risk in January 1996 (BIS 2005[29]). Specifically, the market risk amendments to the Accord specified both a standardised measurement for determining the capital required for market risks and permitted certain banks to use an internal models-based approach to capital adequacy determination for market risks. Although banks were allowed discretion in the exact modelling approach they used under the internal models approach, a minimum standard was articulated in which banks would measure market risk using a 99^{th}-percentile VaR measure with a 10-day risk horizon and at least one year of historical data for parameter estimation (BIS 2005, pp. 40–1).

In 2004, the BIS promulgated the first substantive and complete revision of the 1988 Basel Accord in what is generally known as "Basel II". Basel II retained the internal models approach contained in the 1996 market risk amendments and thus continued the use of VaR by banks for measuring the CaR of their market risk-sensitive assets (BIS 2004).

In response to the financial crisis that began in 2007, the BIS in 2010 proposed a further modification of the Accord known as "Basel III" (BIS 2010, 2011). In its initial proposal, the BIS retained the VaR-based approach of Basel II with some changes. For example, Basel III would require banks to engage in enhanced VaR measurement based on "a continuous 12-month period of significant financial stress" (BIS 2011, p. 3). In addition, Basel III obliges banks to measure the market risk arising from the expected counterparty risk of over-the-counter derivatives. These credit value adjustments (CVAs) to capital requirements are generally computed based on changes in counterparty credit spreads as reflected in banks' VaR models for bonds (or, more commonly, their credit-default-swap-implied synthetic bond price VaR equivalents) (BIS 2011, pp. 31–4).

In May 2012, the BIS issued a new consultative document in which it proposed to replace VaR altogether as the key market risk metric with TVaR, ostensibly in order to capture tail risk better in regulatory CaR measurements (BIS 2012). The proposal was met with mixed reactions from market participants. Some view the shift of Basel towards CaR measurement based on the tails of loss distributions as constructive, whereas others view the change as introducing a significant degree of ambiguity into regulatory capital requirements. For example, Aaron Brown – chief risk officer at AQR Capital and a respected risk-management expert – commented as follows: "VaR is an objective measure that can be determined with calculable error bounds. Expected shortfall will always be an opinion, because it is infinitely sensitive to events with infinitesimal probability. All the evidence in the history of the universe can never tell you expected shortfall” (Carver 2012).

Insurance

Minimum capital requirements for (re)insurance companies are remarkably complex and disparate across international borders. Nevertheless, most countries specify some kind of minimum capital requirements for (re)insurance underwriters.

In the US, the regulation of primary insurance carriers has previously been state-specific (with most states deferring to the New York State Insurance Commission when crafting their own local regulations). Recent post-crisis regulatory reforms will change that, but, as of this writing, minimum capital requirements for underwriters continue to be administered by state insurance regulators. Although state regulators have had significant latitude and discretion in administering minimum capital requirements, the National Association of Insurance Commissioners (NAIC) defines a set of standardised risk-based capital requirements in an effort to promote conformity. The NAIC standards attempt to require insurers to hold an amount of capital deemed adequate to cover most of their major risks. As in the original Basel Accord and the standardised approach in Basel II, fixed risk weights are defined for risky assets, liabilities and premium writings. The size of exposures are adjusted with risk-weighting factors, and the aggregate weighted risk exposure defines an insurer’s “authorised control level”, which is then compared with its total adjusted capital to determine capital adequacy. Forward-looking risk metrics such as VaR and CaR play no role in current US risk-based capital requirements for (re)insurers.

In the European Union, capital requirements for insurance underwriters were previously based on a “solvency margin,” defined broadly as the minimum relation required between surplus (akin to equity) and premiums written and either claims incurred (non-life) or mathematical reserves (life). Beginning in 2014,[30] however, European (re)insurance companies will be subject to the new Solvency II risk-based capital regime. Solvency II defines key risk-based capital requirements for insurance companies known as a minimum capital requirement (MCR) and a solvency capital requirement (SCR). Like Basel II, a (re)insurance company can calculate its SCR using either a standardised approach or internal models. If a carrier chooses the latter, the internal model-based SCR calculation must be based on a 99.5% VaR with a one-year risk horizon. Unlike traditional VaR measures based solely on market risk, moreover, the Solvency II SCR requires carriers to consider all risks, including market, credit, operational and underwriting risks.

DRAWBACKS OF THE AT-RISKS AND CONCLUDING THOUGHTS

Although useful for many types of risks, the At-Risk metrics are not risk-measurement panaceas. VaR came under a particularly strong attack for its inability to capture the market movements that occurred during the global credit crisis that began in 2007.[31] Many such criticisms, however, were statements of the obvious – ie, VaR cannot "predict" crises or reliably describe losses in the tail of a risk distribution. Indeed, the original appeal of the At-Risk metrics was that they tend to be relatively stable and reliable up to a reasonable level of confidence, but the risks in the tails of the risk distributions (that VaR does not purport to measure) are inherently less predictable and measurable.

Even setting aside the somewhat obvious criticism of the At-Risks that they cannot predict tail losses, the At-Risks do nevertheless have several other shortcomings. For example, some of the At-Risks (including VaR) are not "coherent" risk metrics because they are not "subadditive" – ie, the aggregate risk of a portfolio comprising several subportfolios might be greater than the risks of the constituent subportfolios if viewed in isolation.[32] Compared with alternative risk metrics such as TVaR (which is subadditive), VaR thus has some conceptual limitations. Despite this theoretical limitation, it is unclear whether the coherence of a risk measure actually makes much difference for actual risk-management purposes and matters to risk managers.

More importantly, the At-Risks are only as good as the underlying data used to translate generic risk metrics into actual risk measures. No matter how reliable and relevant the At-Risks might be for a firm in estimating the risk of a given portfolio, the resulting risk measure will not be helpful if, for example, certain material risk exposures are omitted from the calculation or if the other aspects of the risk-management process of the firm (eg, risk reporting and control) are not able to address the deficiencies in and/or outputs from At-Risk models.[33] Even if implemented using reasonable methods, At-Risk measures are merely measures of risk and do not guarantee the proper management of those risks. As noted earlier, the risk-management process includes various components, of which risk measurement is only one. Even the best possible risk metric will not help a company if the resulting risk measures are not integrated into an appropriate risk-management process (Culp, Miller and Neves 1998).

The At-Risk estimation methodology can also impact the usefulness of the model. For example, estimating a one-quarter VaR based on an assumption of no changes in the underlying portfolio and no time variation in the parameters of the presumed risk distribution (or no presumed variation in the distribution itself) can lead to misguided risk inferences. The At-Risks are also subject to "model risk" – ie, the risk that the valuation and risk-measurement models (and their input parameters) used to estimate a risk measure are inaccurate, flawed or excessively unrealistic (Derman 1996; Rebonato 2003; Alexander 2009).

Despite the conceptual flaws of the At-Risks, their critics often fail to recognise that one of the most pronounced benefits of using the At-Risks as key risk metrics is to facilitate systematic communications and discussions among senior managers when the At-Risk metrics raise red flags (eg, Brown 2008). A prudent and robust risk-management framework, in fact, should utilise both the relevant At-Risk metric(s) and other basic measures of risk to provide a broad overview of a portfolio's actual risk exposure, of which the At-Risks are just one component. For example, risk managers should not pay so much attention to VaR that they stop evaluating other risk metrics (eg, Sharpe and Sortino ratios and alpha) and other risk-measurement methodologies (eg, scenario analysis and stress testing). But nor should risk managers dismiss the At-Risks because some underlying assumptions required for At-Risk estimation are a bit unrealistic. As with all financial models,[34] what is important with the At-Risks is to understand their uses and their limitations and to interpret the results accordingly in the context of a well-defined and robust risk-management process.

I am grateful to Aaron Brown, Andria van der Merwe, Andrea Neves, David Ross, Pietro Veronesi and Jonathan Williams for helpful comments on earlier drafts. The usual disclaimer applies, and the comments and opinions expressed here are solely my own and do not necessarily reflect the views of those acknowledged, my clients, any institutions with which I am affiliated or the clients and customers of any institutions with which I am affiliated. In compliance with the University of Chicago Booth School of Business's academic disclosure policy, I also hereby acknowledge that I am periodically retained to provide consulting services on the issues discussed in this chapter to financial-services firms, non-financial corporations, asset managers, law firms and other organisations.

1 For a discussion of the distinctions between "financial" and "insurance" risks, see, eg, Culp (2001a, 2006, 2008).

2 For an informative and interesting perspective on the history of the At-Risks, see Brown (2011).

3 For example, if a position in the portfolio was subject to a 1% stop-loss provision, the CaR estimate would reflect only the 1% change and presume the position was liquidated when the stop-loss threshold was reached.

4 In the late 1980s when CaR and VaR first emerged, it was also common for CaR estimates to take into account long-term asset price and volatility trends. VaR, by contrast, was generally measured based on very recent market volatility (and thus did not depend as much on the current level of asset prices).

5 See, eg, FRB (1993) and G30 (1993).

6 This chapter is not intended to provide an exhaustive discussion or comprehensive literature survey of the At-Risks, but rather an introduction to the main concepts and issues.

7 For further discussion, see Smithson (1998), Crouhy, Galai and Mark (2001), Stulz (2003) and Jorion (2007).

8 At the time, the Global Derivatives Study Group was a group representing the over-the-counter derivatives dealer community. The Study Group also received input from academics, regulators, and non-dealer market participants.

9 Not all of these other available risk metrics were "consistent". For example, the DV01 or Greek of an asset is sensitive to the units of measurement, which inhibits direct comparisons of these risk metrics across different asset classes.

10 Implied volatility (ie, volatility implied by observed traded option prices) is a forward-looking risk metric.

11 This interpretation assumes that asset price changes are independently and identically distributed – ie, that price changes are drawn from essentially the same distribution every day.

12 If the portfolio has any interim cashflows between t and $t+\tau$, they would also need to be taken into account if the objective is to measure a "return" rather than just a percentage change.

13 The empirical distribution is a distribution of relative frequencies – ie, the total number of observations in the underlying data in each interval divided by the total number of observations.

14 See, eg, Mood, Graybill, and Boes (1974).

15 In 1996, JP Morgan entered into a partnership with Reuters to distribute the RiskMetrics data. The group was subsequently spun out into a separate company called the RiskMetrics Group, which was later acquired by MSCI.

16 So pervasive was the use of normal PVaR that some wrongly consider PVaR (or even VaR more generally) to be synonymous with the assumption of normality in the risk distribution. But even in the confines of PVaR estimation, a wide range of distributions can be used to approximate the potential changes in value of a portfolio. See, eg, Holton (2003), Dowd (2005) and Jorion (2007).

17 See, eg, Engle (1982), Bollerslev (1986) and Nelson (1990, 1991).
18 See, eg, Giot (2005).
19 Backtesting can also be done using less rigorous rules of thumb that are arguably often just as good as the more complicated approaches.
20 See, eg, Jorion (2007).
21 In reality, a risk manager is more likely to assume "wrong-way" liquidity – ie, that relatively less liquid assets need to be sold in a declining market. In that case, the appropriate presumptive sale price would be the offer price and not the bid–ask midpoint.
22 See Bangia, Diebold, Schuermann and Stroughair (1999a, b); see also Jorion (2007).
23 The discussion in this section relies on these various references heavily; interested readers should refer to the indicated citations for additional discussion and details.
24 The sample likely will be drawn from a historical return distribution.
25 This expression requires that $1 + \frac{\xi(r_{t,\tau} - \mu)}{\sigma} > 0$.
26 Less complicated ways of capturing fat tails in VaR measurements often use the traditional PVaR approach, but, instead of assuming a normal distribution, are based on a fat-tailed distribution like the Student's *t*-distribution.
27 Note that we would not normally estimate EV-VaR at the 95% confidence level. Because we care about extreme values, more typical confidence levels would be 99.5% or 99.9%. We use 95% here for comparability of our examples to the discussion earlier.
28 See, eg, Froot, Scharfstein and Stein (1993).
29 The market risk amendment to the Accord was first released in January 1996 and then subsequently amended in September 1997 and November 2005. The cited 2005 document is the version that supersedes the previous versions.
30 The final implementation date for Solvency II is subject to some debate.
31 See, eg, Einhorn (2008) and Nocera (2009). For a different perspective, see, eg, Brown (2008). A "coherent" risk measure is typically defined as a risk measure that exhibits monotonicity, subadditivity, translation invariance, positive homogeneity, and relevance. See Artzner, Delbaen, Eber and Heath (1999).
32 See, eg, Artzner, Delbaen, Eber and Heath (1999), Dowd (2005), Hull (2006) and Munenzon (2010).
33 For some examples, see Culp, Miller and Neves (1998).
34 For a thoughtful perspective on this topic, see Derman (2011).

REFERENCES

Alexander, C., 2009, *Market Risk Analysis Vol IV: Value-at-Risk Models* (New York: John Wiley & Sons).

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath, 1999, "Coherent Measures of Risk", *Mathematical Finance* 9(3) (July).

Bangia, A., F. Diebold, T. Schuermann, and J. Stroughair, 1999a, "Liquidity on the Outside", *Risk* 12 (June).

Bangia, A., F. Diebold, T. Schuermann, and J. Stroughair, 1999b, "Modeling Liquidity Risk, With Implications for Traditional Market Risk Measurement and Management", Wharton Financial Institutions Center Working Paper 99-06.

Bank for International Settlements (BIS), 1988, "International Convergence of Capital Measurement and Capital Standards", Basel Committee on Banking Supervision, July.

Bank for International Settlements (BIS), 2005, "Amendment to the Capital Accord to Incorporate Market Risks", Basel Committee on Banking Supervision, November.

Bank for International Settlements (BIS), 2010, "Basel III: International Framework for Liquidity Risk Measurement, Standards and Monitoring", Basel Committee on Banking Supervision, December.

Bank for International Settlements (BIS), 2011, "Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems", Basel Committee on Banking Supervision, June.

Bank for International Settlements (BIS), 2012, "Fundamental Review of the Trading Book", Basel Committee on Banking Supervision, May.

Berkowitz, J., 2001, "Testing Density Forecasts, With Applications to Risk Management", *Journal of Business and Economic Statistics* 19(4).

Board of Governors of the Federal Reserve System (cited in text as "FRB"), 1993, SR Letter 93-69.

Brown, A., 2008, "Counterpoint: Capital Inadequacy", *Risk Professional* (June/July).

Brown, A., 2011, *Red-Blooded Risk* (New York: John Wiley & Sons).

Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics* 31.

Brunnermeier, M. K., and L. H. Pedersen, 2009, "Market Liquidity and Funding Liquidity", *Review of Financial Studies* 22(6).

Campbell, S. D., 2007, "A Review of Backtesting and Backtesting Procedures", *Journal of Risk* 9(2) (Winter).

Carver, L., 2012, "Basel Committee Proposes Scrapping VaR, *Central Banking* (May 4).

Christoffersen, P. F., 1998, "Evaluating Interval Forecasts", *International Economic Review* 39(4) (November).

Coles, S., 2001, *An Introduction to Statistical Modeling of Extreme Values* (Berlin: Springer).

Crouhy, M., D. Galai, and R. Mark, 2001, *Risk Management* (New York: McGraw-Hill).

Culp, C. L., 2000, "*Ex Ante* vs. *Ex Post* RAROC", *Derivatives Quarterly* 7(1) (Fall).

Culp, C. L., 2001a, "The Revolution in Corporate Risk Management: A Decade of Innovations in Process and Products", *Journal of Applied Corporate Finance* 14(4) (Winter).

Culp, C. L., 2001b, *The Risk Management Process* (New York: John Wiley & Sons).

Culp, C. L., 2004, *Risk Transfer: Derivatives in Theory and Practice* (New York: John Wiley & Sons).

Culp, C. L., 2006, *Structured Finance and Insurance* (New York: John Wiley & Sons).

Culp, C. L., 2008, "Risk and Risk Management", in F. J. Fabozzi (ed.), *Handbook of Finance, Vol. III: Valuation, Financial Modeling, and Quantitative Tools* (New York: John Wiley & Sons).

Culp, C. L., and R. J. Mackay, 1994, "Managing Derivatives Risk: A Strategic Guide", in *1995 Handbook of Business Strategy* (New York: Faulkner & Gray).

Culp, C. L., and R. Mensink, 1999, "Measuring Risk for Asset Allocation, Performance Evaluation, and Risk Control: Different Problems, Different Solutions", *Journal of Performance Measurement* (Fall).

Culp, C. L., and P. Planchat, 2000, "New Risk Culture: An Opportunity for Business Growth and Innovation", *Derivatives Quarterly* 6(4) (Summer).

Culp, C. L., M. H. Miller and A. M. P. Neves, 1998, "Value at Risk: Uses and Abuses", *Journal of Applied Corporate Finance* 10(4) (Winter).

Culp, C. L., K. Tanner and R. Mensink, 1997, "Risk, Returns and Retirement", *Risk* 10(10) (October).

De Haan, L., and A. Ferreira, 2006, *Extreme Value Theory: An Introduction* (Berlin: Springer).

De la Pena, V. H., R. Rivera and J. Ruiz-Mata, 2007, "Quality Control of Risk Measures: Backtesting VAR Models", *Journal of Risk* 9(2) (Winter).

Derman, E., 1996, "Model Risk", Goldman Sachs Quantitative Strategies Research Notes, April.

Derman, E., 2011, *Models Behaving Badly* (New York: Free Press).

Diebold, F. X., T. A. Gunther and A. S. Tay, 1998, "Evaluating Density Forecasts", *International Economic Review* 39(4).

Dowd, K., 2005, *Measuring Market Risk*, 2nd edn (New York: John Wiley & Sons).

Einhorn, D., 2008, "Private Profits and Socialized Risks", *Risk Professional* (June/July).

Embrechts, P., C. Klüppelberg, and T. Mikosch, 2003, *Modelling Extremal Events for Insurance and Finance* (Berlin: Springer).

Engle, R. F., 1982, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica* 50(4) (July).

Froot, K., D. Scharfstein and J. Stein, 1993, "Risk Management: Coordinating Corporate Investment and Financing Policies", *Journal of Finance* 48.

Giot, P., 2005, "Implied Volatility Indexes and Daily Value at Risk Models", *Journal of Derivatives* 12(4) (Summer).

Global Derivatives Study Group (G30), 1993, "Derivatives: Practices and Principles", Group of Thirty.

Hendricks, D., 1996, "Evaluating Value-at-Risk Models Using Historical Data", *FRBNY Economic Policy Review* (April).

Holton, G. A., 2003, *Value-at-Risk* (New York: Academic Press).

Holton, G. A., 2004, "Defining Risk", *Financial Analysts Journal* 60(6).

Hull, J., 2006, "VaR vs. Expected Shortfall", *Risk* (December).

Jorion, P., 2007, *Value at Risk*, 3rd edn (New York: McGraw-Hill).

JP Morgan-Reuters, 1996, "RiskMetrics – Technical Document", 4th edn, December 17.

Kreps, R., 2005, "Riskiness Leverage Models", *Proceedings of the Casualty Actuarial Society Casualty Actuarial Society – Arlington, Virginia* 92.

Kupiec, P. H., 1995, "Techniques for Verifying the Accuracy of Risk Measurement Models", *Journal of Derivatives* 3(2) (Winter).

Litterman, R., 2003, *Modern Investment Management: An Equilibrium Approach* (New York: John Wiley & Sons).

Lopez, J. A., 1999, "Methods for Evaluating Value-at-Risk Estimates", *FRBSF Economic Review* (2).

McNeil, A. J., R. Frey and P. Embrechts, 2005, *Quantitative Risk Management: Concepts, Techniques, and Tools* (Princeton: Princeton University Press).

Merton, R. C., and A. Perold, 1993, "Theory of Risk Capital in Financial Firms", *Journal of Applied Corporate Finance* 6(3) (Fall).

Mood, A. M., F. A. Graybill and D. C. Boes, 1974, *Introduction to the Theory of Statistics* (New York: McGraw-Hill).

Munenzon, M., 2010, "Risk Measurement from Theory to Practice: Is Your Risk Metric Coherent and Empirically Justified?", SSRN Working Paper (May 12).

Nelson, D. B., 1990, "ARCH Models as Diffusion Approximations", *Journal of Econometrics* 45.

Nelson, D. B., 1991, "Conditional Heteroskedasticity in Asset Returns: A New Approach", *Econometrica* 59(2) (March).

Nocera, J., 2009, "Risk Mismanagement", *New York Times* (January 4).

Pearson, N. D., 2002, *Risk Budgeting* (New York: John Wiley & Sons).

Picoult, E., 1998, "Calculating Value-at-Risk with Monte Carlo Simulation", in B.

Dupire (ed.), *Monte Carlo: Methodologies and Applications for Pricing and Risk Management* (London: Risk Books).

Rahl, L. (ed.), 2000, *Risk Budgeting: A New Approach to Investing* (London: Risk Books).

Rebonato, R., 2003, "Theory and Practice of Model Risk Management", in P. Field (ed.), *Modern Risk Management: A History* (London: Risk Books).

Smithson, C. W., 1998, *Managing Financial Risk* (New York: McGraw-Hill).

Stulz, R., 2003, *Risk Management and Derivatives* (Mason, OH: South-Western).

Van Roekel, G. J., 2008, "Extended Analysis of Back Testing Framework for Value at Risk". Master Thesis at University of Twente, August

Venter, G., and J. A. Major, 2003, "Allocating Capital by Risk Measures: A Systematic Survey", *Guy Carpenter Views* (August).

Zaik, E., J. Walter, G. Kelling and C. James, 1996, "RAROC at Bank of America: From Theory to Practice", *Journal of Applied Corporate Finance* 9(2) (Summer).