Shortfall: a tail of two parts

***Richard Martin** and **Dirk Tasche** show that the expected shortfall, when used in the conditional independence framework, has an elegant decomposition into systematic (risk-factor-driven) and unsystematic parts. The theory is compared and contrasted with the well-known, and analogous, decomposition for variance*

One of the challenges in trading and risk management of portfolios and portfolio derivatives is understanding where the risk is coming from – a difficulty because there are many underlyings. The initial objective of credit portfolio modelling 10 or so years ago was simply to construct a distribution of loss or profit and loss, thereby only measuring risk at portfolio level. More recently, it has become much easier for market participants to hedge or securitise risk, and various financial instruments are available for the purpose, such as credit default swaps (CDSs), index derivatives and options, collateralised debt obligation (CDO) tranches and securitisation vehicles. The requirement of CDO managers and credit or loan portfolio managers to deal with credit default risk has led to the development of two things: first, correlation models, and second, analytical techniques that can handle asymmetrical risk. With correlation modelling, the subject has by and large converged on the conditional independence (CI) framework, allowing concrete statements to be made about default rate volatilities and the potential losses; this is now used both in risk management and in pricing. As regards portfolio analytics, the necessity for risk measures other than the normal mean-variance has been understood, and it is well established that expected shortfall (ESF) is a good tool, as it is sensitive to tail risk, it is a coherent risk measure (Artzner *et al*, 1999, and Acerbi & Tasche, 2002), and it provides a close link to the notion of tranche payouts and hence to the CDO and securitisation world.

To bring together the ideas of CI, ESF and 'where is our risk coming from?', we discuss in this article the application of ESF to CI models and show how to give a precise quantitative measure of how much risk in the portfolio (model) comes from systematic risk and how much is residual. An immediate and useful consequence of this is that it shows how much risk can be reduced by diversification, or how much can be hedged with portfolio instruments and how much requires bespoke ones.

The above-mentioned ideas have underpinned the Basel II regulatory capital framework (Basel Committee on Banking Supervision, 2006, and Gordy, 2003) too, and although that is not the subject of this article it is worth a brief mention. If we can assume that the world is driven by a univariate risk factor A say, and that the A-conditional expected loss of the portfolio, $\mathbf{E}[Y \mid A]$, is a monotonic function of this factor, and also that the portfolio is 'infinitely fine-grained', then the portfolio loss is a one-to-one transformation of the risk factor: $Y = \mathbf{E}[Y \mid A]$. The Basel II framework chooses a simple prescription, thereby allowing the quantiles of $\mathbf{E}[Y \mid A]$ to be easily computable. Now if we wish to incorporate the effects of unsystematic risk (finite portfolio; large or largish exposures) we can model the loss as $Y = X + U$, with $X = \mathbf{E}[Y \mid A]$ (the loss of the putative 'infinitely granular portfolio', which need not exist in reality) and U denoting an independent Gaussian residual of variance σ^2. The difference between the upper P-quantiles of X and Y is given by the granularity adjustment (GA) formula (Martin & Wilde, 2002, and references therein):

$$VAR_P[Y] \sim VAR_P[X] - \frac{1}{2f(x)}\frac{d}{dx}\sigma^2(x)f(x)\Big|_{x=VAR_P[X]} \qquad (1)$$

where f is the density of X. Incidentally the shortfall-GA[1] is:

$$S_P^+[Y] \sim S_P^+[X] + \frac{1}{2P}\sigma^2(x)f(x)\Big|_{x=VAR_P[X]}$$

which superficially appears nicer than the value-at-risk result, and indeed is analytically superior because the correction is always positive; this is one of several things that we shall be talking about later.[2] Notice that for the purposes of the GA formulas above we do not need A to be univariate; however, if it is not, then it is generally a lot more difficult to calculate the VAR of $\mathbf{E}[Y \mid A]$. From now on, we will not make any assumptions about the distribution of A, or about the conditional distribution of Y on A; we thereby keep everything general.

Continuing from the above ideas, we start by writing:

$$Y = \mu_{Y|A} + \left(Y - \mu_{Y|A}\right)$$

with $\mu_{Y|A} = \mathbf{E}[Y \mid A]$ denoting the conditional mean of Y given A (so it is a random variable). It is then natural to consider the following expression:

$$\mathbf{E}\left[Y \middle| Y > y\right] = \mathbf{E}\left[\mu_{Y|A} \middle| Y > y\right] + \text{remainder}$$

thereby splitting the ESF into two parts. In general, we will not know the distribution of $\mu_{Y|A}$ in 'closed form'. In practice this is not an issue because when calculations are done one is in effect coming up with a large number of scenarios for A and computing the conditional mean $\mathbf{E}[Y \mid A]$ in each; the distribution of $\mu_{Y|A}$ is then approximated by the empirical distribution of the

[1] *The ESF S^+ or S^- is defined by $E[Y \mid Y > y]$ or $E[Y \mid Y < y]$ where Y is the portfolio loss (or value) and y is the VAR at the chosen tail probability. In the case of Y not having a continuous distribution, the definition of ESF should be modified (see Acerbi & Tasche, 2002). The choice of signs >,< reflects the fact that Y can denote portfolio loss (use > sign) or value (use < sign). From now on, we shall stick with the convention that Y denotes loss, to avoid confusion later on, but clearly the theory works equally well with either*

[2] *See also Gordy (2004)*

generated sample. We identify the first term as the contribution of systematic risk to the portfolio ESF (or the 'systematic part' for short). Incidentally, this arises as a natural consequence of the Fourier integral representation of shortfall, used in the saddlepoint approximation (Martin, 2006). The plan for this article is:

- We explore the relationship between this and the well-known analysis of variance formula:

$$\mathbf{V}[Y] = \mathbf{V}\left[\mu_{Y|A}\right] + \mathbf{E}\left[\sigma^2_{Y|A}\right] \tag{2}$$

in which the terms on the right-hand side tell us how much risk comes from the variation of the risk factor(s) ($\mathbf{V}[\mu_{Y|A}]$) and how much comes from residual risk ($\mathbf{E}[\sigma^2_{Y|A}]$). Hence we identify the 'remainder' term in the above equation as the contribution to unsystematic risk. For a multivariate normal portfolio model, and for elliptical distributions (such as Student-*t*), we show that the decompositions are essentially identical.

- We link this to the homogeneity theory of risk measures, and pursue the analysis a little further to show that it does not work properly for VAR.
- We illustrate by way of an example that the ESF and mean-variance decompositions are essentially different for typical credit risk models, which give skewed distributions.
- Finally, to extend things further we consider a notion of 'CI-monotonicity', by which we mean that, in a conditional independence framework, if the systematic or unsystematic risk is increased then the measured risk should increase too. For this not to be so would probably result in some bizarre portfolio optimisations, with an uncorrelated asset being given more exposure 'to reduce risk'. Note that we are not being critical of diversification here: it is legitimate to reduce portfolio risk by increasing allocation to an uncorrelated asset while reducing exposure to others, but the headline risk figure should not decrease if that allocation is increased without a reduction in the others. The VAR does not obey this (and indeed it is not coherent in the Artzner *et al* sense), but ESF does (and that is coherent).

Theory

Systematic decomposition of ESF. We start with our basic definition of the systematic part:

$$\mathbf{E}\left[\mu_{Y|A} \middle| Y > y\right] \tag{3}$$

This can also be written as the mean plus the covariance of the conditional expected loss and the tail probability:

$$\mu_Y + \frac{1}{P^+}\mathbf{V}\left[\mu_{Y|A}, \mathbf{P}\left[Y > y \middle| A\right]\right] \tag{4}$$

with $P^+ = \mathbf{P}[Y > y]$ and $\mu_Y = \mathbf{E}[Y]$. (This is apparent when one expands the covariance as the expectation of the product minus the product of the expectations; the second term cancels the μ_Y.) There is a link with mean-variance theory: the 'mean plus some number of standard deviations' risk measure is:

$$\mu_Y + \eta\sigma_Y = \mu_Y + \frac{\eta}{\sigma_Y}\mathbf{V}\left[\mu_{Y|A}\right] + \frac{\eta}{\sigma_Y}\mathbf{E}\left[\sigma^2_{Y|A}\right] \tag{5}$$

so the second term is the covariance of $\mu_{Y|A}$ with itself ($\times\ \eta/\sigma_Y$). Hence the expressions for systematic risk contribution have in common that they are the mean plus the covariance of the conditional expected loss with something reasonably natural (the conditional tail probability in (4), the conditional expected loss in (5)). As we shall see later, in the case when the joint distribution of asset returns in the portfolio is multivariate normal, the decompositions are in fact identical. But before we do that we want to explore the unsystematic contribution in more detail.

An interpretation using homogeneity. Another derivation of the systematic/unsystematic contributions to ESF exploits only the fact that ESF is a one-homogeneous risk measure. For a risk measure R that is positively homogeneous of degree d (that is, $R[\theta X] = \theta^d R[X]$ for $\theta > 0$), consider the function:

$$G(u,v) = R\left(u\mu_{Y|A} + v\left(Y - \mu_{Y|A}\right)\right)$$

By Euler's lemma:

$$R[Y] = G(1,1) = \frac{1}{d}\left(u\frac{\partial G}{\partial u} + v\frac{\partial G}{\partial v}\right)_{u=v=1} \tag{6}$$

As a consequence, one is led to consider $d^{-1}(\partial G/\partial u)_{u=v=1}$ and $d^{-1}(\partial G/\partial v)_{u=v=1}$ as, respectively, contributions to systematic and unsystematic risk. Applied to the ESF, these give ($d = 1$):

$$\left(\frac{\partial G}{\partial u}\right)_{u=v=1} = \mathbf{E}\left[\mu_{Y|A} \middle| Y > y\right], \quad \left(\frac{\partial G}{\partial v}\right)_{u=v=1} = \mathbf{E}\left[Y - \mu_{Y|A} \middle| Y > y\right]$$

in agreement with (3).[3] Moreover, in the mean-variance case we have (R = variance, $d = 2$):

$$\frac{1}{2}\left(\frac{\partial G}{\partial u}\right)_{u=v=1} = \mathbf{V}\left[\mu_{Y|A}\right], \quad \frac{1}{2}\left(\frac{\partial G}{\partial v}\right)_{u=v=1} = \mathbf{E}\left[\sigma^2_{Y|A}\right]$$

as desired. In general, the ratio (systematic/total) can be thought of as a generalised R-squared (coefficient of regression), for it tells us what proportion of the total risk of Y is explained by $\mu_{Y|A}$, the best prediction in the least-squares sense of Y given A.

The unsystematic contribution is positive. The following argument shows how the ESF of the 'infintely granular portfolio' and the ESF of the real portfolio are related. We have:

$$\begin{aligned}\mathbf{E}\left[\mu_{Y|A} \middle| Y > y\right] &\le \mathbf{E}\left[\mu_{Y|A} \middle| \mu_{Y|A} > y^*\right] \\ &= \mathbf{E}\left[Y \middle| \mu_{Y|A} > y^*\right] \le \mathbf{E}\left[Y \middle| Y > y\right]\end{aligned} \tag{7}$$

Here, y is the VAR of the portfolio and y^* is the VAR of $\mu_{Y|A}$ at the same tail probability. In (7), the first expression is the contribution of the systematic risk to the whole; the second expression is the ESF of the systematic part considered in isolation; and the fourth is the ESF of the whole portfolio. Notice that the equivalent expression for mean-variance is:

$$\sigma_Y^{-1}\mathbf{V}\left[\mu_{Y|A}\right] \le \mathbf{V}\left[\mu_{Y|A}\right]^{1/2} \le \sigma_Y \tag{8}$$

with σ_Y the standard deviation of Y. Although the mean-variance result is obvious from (2), the ESF result (7) is not, so we prove it now.

The cleanest way of dealing with both inequalities at the same time is to use this result: if V, W are any two random variables and v, w are chosen such that $\mathbf{P}[V > v] = \mathbf{P}[W > w]$, then:

$$\mathbf{E}\left[V \middle| W > w\right] \le \mathbf{E}\left[V \middle| V > v\right]$$

For the proof, note that if B denotes any random variable bounded to [0, 1] then:

[3] *See Tasche (1999), Cor. 5.7, for the result on derivatives of ESF*

1 Adding unsystematic risk can cause VAR to reduce

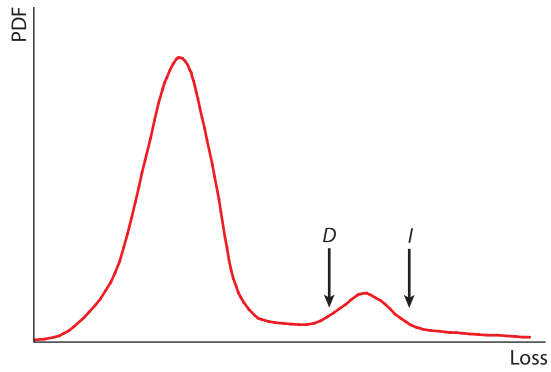

The graph shows the PDF of the 'infinitely granular portfolio'. For quantiles near *I*, the VAR will increase when a small amount of uncorrelated risk is added, but for those near *D*, it will decrease. This anomaly cannot occur with ESF

2 Loss distributions for the test portfolio with different correlation assumptions

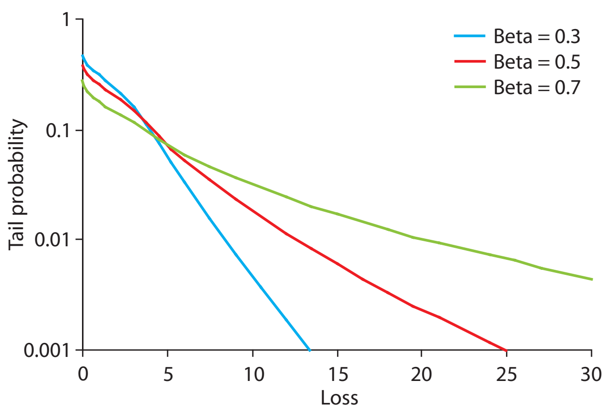

$$\mathbf{E}\left[(V-v)\left(B-\mathbf{1}\left[V>v\right]\right)\right]\le 0$$

for if $V > v$ then $\mathbf{1}[V > v] = 1$ and $B - \mathbf{1}[V > v] \le 0$, while if $V \le v$ then $\mathbf{1}[V > v] = 0$ and $B - \mathbf{1}[V > v] \ge 0$, so either way the product is ≤ 0. Then set $B = \mathbf{1}[W > w]$ and rearrange. The central equality follows by consideration of $\mathbf{E}[Y\mathbf{1}[\mu_{Y|A} > y^*]$: first condition on A, turning Y into $\mu_{Y|A}$, and then integrate A out.

From (7) we can see immediately that the contribution of the unsystematic part of the ESF, $\mathbf{E}[Y - \mu_{Y|A} \mid Y > y]$, is necessarily positive. The same is true for variance: $\sigma_Y - \sigma_Y^{-1}\,\mathbf{V}[\mu_{Y|A}] > 0$. These inequalities can be thought of as consequences of the convexity of the risk measures. Indeed, the property that risk contributions cannot exceed the corresponding stand-alone risk (that is, equations (7) and (8)) follows from the subadditivity of the risk measure, if it is positively homogeneous of degree one (see Tasche 2002, Proposition 2.5). If we try and construct the analogue of (7) for VAR, it fails. From (6), the contribution of systematic risk to VAR is $\mathbf{E}[\mu_{Y|A} \mid Y = y]$ (the proof follows arguments given in Gouriéroux, Laurent & Scaillet, 2000). But now:

$$\begin{aligned}\mathbf{E}\left[\mu_{Y|A}\middle|Y=y\right] &\nless \mathbf{E}\left[\mu_{Y|A}\middle|\mu_{Y|A}=y^*\right]\equiv y^*\\ &=\mathbf{E}\left[Y\middle|\mu_{Y|A}=y^*\right]\nless \mathbf{E}\left[Y\middle|Y=y\right]\equiv y\end{aligned} \tag{9}$$

(The symbol '$a \nless b$' means, somewhat non-standardly, 'a might not be less than b', that is, that the assertion $a < b$ may be false.) Again y and y^* denote the VAR of Y and of $\mu_{Y|A}$ at the same upper tail probability. We have found only artificial examples in which the first inequality is violated, but the second one is easier to demonstrate: the GA formula (1) shows that if the density of $\mathbf{E}[Y \mid A]$ is increasing in the part of the tail at which the VAR falls (see figure 1) then the effect of adding small amounts of unsystematic risk is to reduce the VAR. By contrast, we commented in the introduction to this article that the sensitivity of ESF to small amounts of Gaussian unsystematic risk is always positive, and we have shown here that it is positive for any amount of unsystematic risk. Therefore, the anomaly never occurs.

Examples with different portfolio models

■ **Multivariate normal model and elliptic model.** We drew a comparison between the mean-variance and ESF frameworks earlier (equations (4) and (5)). Pursuing this line a bit further, one might ask whether the ESF and mean-variance decompositions give the same result in any particular case. We start with the multivariate normal portfolio and risk factor, on which the ESF and mean-variance measures are equivalent. Let the risk factor be $A \sim N(0, \Sigma)$ and write:

$$Y = \mu + k'A + U$$

where k is the vector of the factor weights and $U \sim N(0, \sigma_U^2)$ is independent of A and represents the unsystematic risk. Then:

$$\begin{aligned}\mu_{Y|A} &= \mu + k'A\\ \sigma^2_{Y|A} &= \sigma^2_U\\ f_{Y|A}(y) &= \frac{1}{\sigma_U}\phi\left(\frac{y-\mu-k'A}{\sigma_U}\right)\\ \mathbf{P}\left[Y>y\middle|A\right] &= \Phi\left(\frac{\mu+k'A-y}{\sigma_U}\right)\end{aligned}$$

and so the systematic contribution to ESF is:

$$\begin{aligned}&\frac{1}{P^+}\mathbf{E}\left[(\mu+k'A)\Phi\left(\frac{\mu+k'A-y}{\sigma_U}\right)\right]\\ &=\frac{1}{P^+}\left[\mu\Phi\left(\frac{\mu-y}{\sigma_Y}\right)+\frac{k'\Sigma k}{\sigma_Y}\phi\left(\frac{\mu-y}{\sigma_Y}\right)\right]\end{aligned}$$

(as the integration over A can be done in closed form) while the unsystematic contribution is:

$$\frac{1}{P^+}\mathbf{E}\left[\sigma_U^2 f_{Y|A}\right]=\frac{1}{P^+}\frac{\sigma_U^2}{\sigma_Y}\phi\left(\frac{\mu-y}{\sigma_Y}\right)$$

where $\sigma_Y^2 = k'\Sigma k + \sigma_U^2$ is the unconditional variance of Y. Note that $P^+ = \Phi((\mu - y)/\sigma_Y)$ so the first term in the systematic part is just μ, as expected.

To compare this with the mean-variance framework, let the mean-variance risk measure be mean plus η standard deviations, that is, $\mu_Y + \eta\sigma_Y$. Then the systematic and unsystematic parts are respectively:

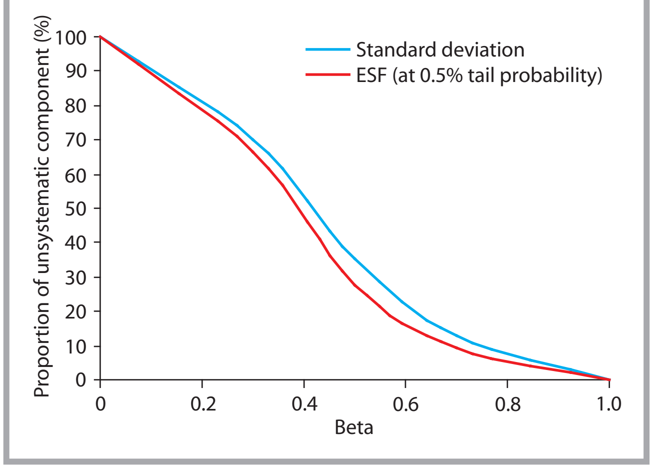

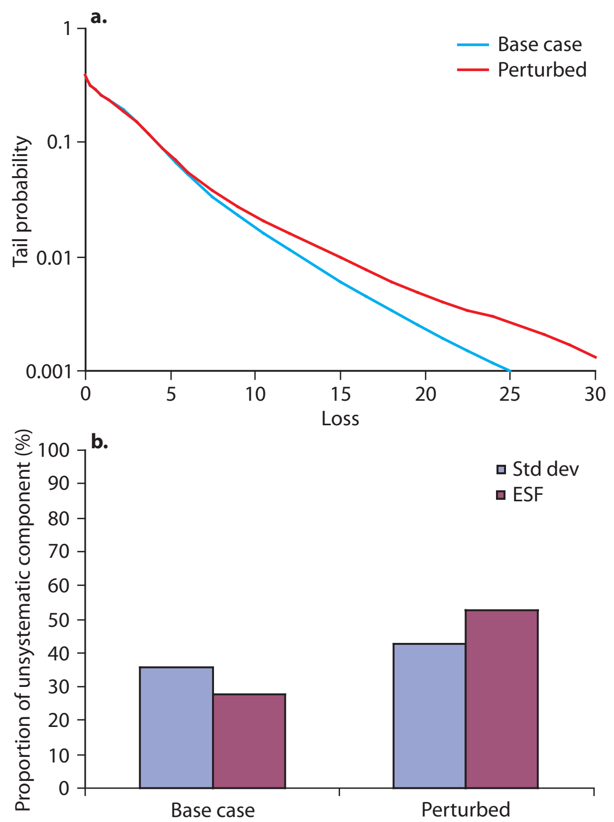

The 'base case' is figure 2 with β = 0.5. In (a) the overall risk is seen to increase. In (b) the proportion of unsystematic risk has increased and the increase is substantially more for ESF than it is for standard deviation

$$\mu + \eta \frac{k'\Sigma k}{\sigma_Y} \quad \text{and} \quad \eta \frac{\sigma_U^2}{\sigma_Y}$$

Setting $\eta = \phi(\Phi^{-1}(P^+))/P^+$, a definition that depends only on the choice of tail probability rather than on the portfolio in question, makes the ESF and mean-variance decompositions agree. The systematic part is proportional to the square of the correlations (that is, 'k squared') and the unsystematic part is inversely proportional to the portfolio size (not inversely proportional to its square root).

As η is independent of the portfolio mean and variance, the result extends automatically to elliptical distributions of finite variance, as the elliptical model is obtained from the normal model by making the variance random and then integrating it out (for example, for Student-*t* it is reciprocal-gamma distributed).

■ **Loan portfolio.** This example, by contrast, shows that ESF and mean-variance do not always agree. We take as a test example a hypothetical portfolio of 50 default-risky loans. For simplicity we assume that the recovery rates are zero, the exposures lie in the range one to five and the mean default probabilities (roughly) in the range 0.2% to 2%. The correlation is modelled by a one-factor Gaussian copula, by which we mean that the risk factor is $V \sim N(0, 1)$ and the conditional default probability is given by the usual formula:

$$p_{i|V} = \Phi\left(\frac{\Phi^{-1}(\bar{p}_i) - \beta_i V}{\sqrt{1-\beta_i^2}}\right)$$

As we are going to vary β, for simplicity we give each asset the same β, though of course this is not necessary.

In figure 2, we show the loss distributions for three different values of β – 0.3, 0.5 and 0.7 – which in the context of a real portfolio correspond to low, moderate and high degrees of correlation. As expected, they become fatter-tailed as the correlation is raised. More important are the systematic/unsystematic decompositions (figure 3). They are shown for standard deviation and for ESF with an upper tail probability of 0.5% (to make the two more directly comparable, the expected loss is deducted from the systematic contribution to ESF). For low correlation, the risk is mostly unsystematic, whereas for high correlation it is mostly systematic, as expected. In this respect the two risk measures give pretty much the same results.

Taking β = 0.5, we greatly increase the exposures to three of the entities, from 5, 5, 5 to 25, 12, 15, respectively. This has a greater impact on tail risk than on standard deviation. Further, it is risk that principally arises from single-name exposures. We should therefore expect the unsystematic component to change much more for ESF than for standard deviation – and indeed figure 4b shows that the numerics confirm this.

CI-monotonicity

Here, we move on to the more general theme of what happens when a new risky instrument is added to a portfolio. We do this from the perspective of what happens when we start varying the amounts of systematic and unsystematic risk. To start the ensuing discussion, it is convenient to write down the expression for the ESF where the conditional distribution of $Y \mid A$ is assumed to be normal:

$$\mathbf{E}[Y|Y>y] = \frac{1}{P^+}\mathbf{E}\left[\mu_{Y|A}\Phi\left(\frac{\mu_{Y|A}-y}{\sigma_{Y|A}}\right) + \sigma_{Y|A}\phi\left(\frac{\mu_{Y|A}-y}{\sigma_{Y|A}}\right)\right] \quad (10)$$

(see also Martin, 2006). As the expression obtained in (10) is of the form $(\ldots)\mu_{Y|A} + (\ldots)\sigma_{Y|A}$, it seems reasonable to try to interpret the coefficients as the sensitivities of ESF to $\mu_{Y|A}$ and $\sigma_{Y|A}$. Of course, there is the complication that the coefficients themselves also depend on $\mu_{Y|A}$ and $\sigma_{Y|A}$. Nevertheless, this is an avenue worth exploring. We consider an arbitrary conditional independence model and ask how the risk measure changes when the loss distribution is altered. In performing this perturbation, we treat systematic and unsystematic risk separately.

The model is that, conditionally on the risk factor A, the density of loss is assumed to be:

$$f_{Y|A}(y) = \frac{1}{\sigma_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)$$

for some probability density function ψ that is normalised to have mean zero and variance one. The subscript in ψ_A indicates that the shape of the distribution may also depend on the risk factor. The case that we have just been considering is that the loss is normally distributed conditionally on the risk factor, that is, $\psi_A(x) = \phi(x)$. The interpretation of $\mu_{Y|A}$ is that its variation (as A varies) constitutes the systematic risk, while $\sigma^2_{Y|A}$ is the unsystematic risk.[4] We would like the risk measure R to have the property that it is an increasing function of $\mu_{Y|A}$ and $\sigma^2_{Y|A}$ (which are themselves functions of A). That is:

$$\frac{\partial R[Y]}{\partial \mu_{Y|A}} > 0, \quad \frac{\partial R[Y]}{\partial \sigma_{Y|A}} > 0$$

Note that we are differentiating with respect to the conditional mean and standard deviation. If these conditions are satisfied, we say that R is CI-monotonic. Interestingly, neither the traditional 'mean plus some number of standard deviations' measure, nor VAR, is CI-monotonic.

To show that the risk measure $\mu_Y + \eta\sigma_Y$ is not CI-monotonic, we write it out as:

$$\begin{aligned} R[Y] &= \mu_Y + \eta\sigma_Y \\ &= \mathbf{E}\left[\mu_{Y|A}\right] + \eta\left(\mathbf{E}\left[\left(\mu_{Y|A}-\mu_Y\right)^2\right] + \mathbf{E}\left[\sigma^2_{Y|A}\right]\right)^{1/2} \end{aligned}$$

We now need to bear in mind that $\mu_{Y|A}$ and $\sigma^2_{Y|A}$ are random variables, and hence are functions (defined on the space of possible values of A and taking real values). Perturbing $\mu_{Y|A} \to \mu_{Y|A} + \delta\mu_{Y|A}$, etc, and thereby differentiating $R[Y]$ with respect to these functions, we find the change in R to be:

$$\delta R = \mathbf{E}\left[\left(1 + \eta\frac{\mu_{Y|A}-\mu_Y}{\sigma_Y}\right)\delta\mu_{Y|A} + \frac{\eta\sigma_{Y|A}}{\sigma_Y}\delta\sigma_{Y|A}\right]$$

So the coefficient of $\delta\mu_{Y|A}$ is not necessarily positive: if the conditional mean loss is smaller than the unconditional mean then increasing it causes the variance to reduce. The coefficient of $\delta\sigma_{Y|A}$ is positive, though: the sensitivity to unsystematic risk is 'correct'.

To show that VAR is not CI-monotonic, we start with the expression for the lower tail probability:

$$P^- = \mathbf{E}\left[\Psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\right]$$

(here $\Psi_A(y) = \int_{-\infty}^{y}\psi_A(x)dx$). Taking differentials with respect to $\mu_{Y|A}$, $\sigma_{Y|A}$ gives:

$$\begin{aligned} \delta P^- = \mathbf{E}\Bigg[&\frac{1}{\sigma_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\delta y - \frac{1}{\sigma_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right) \\ &\times \delta\mu_{Y|A} - \frac{y-\mu_{Y|A}}{\sigma^2_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\delta\sigma_{Y|A}\Bigg] \end{aligned} \quad (11)$$

To find the sensitivity of VAR, that is, to find δy, we equate δP^- to zero. In doing so, we find that the coefficient of $\delta\mu_{Y|A}$ is positive, but that of $\delta\sigma_{Y|A}$ is negative if the conditional mean $\mu_{Y|A}$ exceeds the VAR y. Hence it is possible for an increase in unsystematic risk (conditional variance) to decrease the VAR – as we found in figure 1 and also as the GA formula (1) warned us would happen.

However, ESF is CI-monotonic, as we now show. Denoting the VAR by y, we have the following expression for the ESF:

$$S^+ = \frac{1}{P^+}\mathbf{E}\left[\int_y^\infty \frac{x}{\sigma_{Y|A}}\psi_A\left(\frac{x-\mu_{Y|A}}{\sigma_{Y|A}}\right)dx\right]$$

Taking differentials, integrating by parts and tidying up, we have:

$$\begin{aligned} \delta\left(P^+S^+\right) = &-\mathbf{E}\left[\frac{y}{\sigma_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\right]\delta y \\ &+\mathbf{E}\left[\left[\int_y^\infty \frac{1}{\sigma_{Y|A}}\psi_A\left(\frac{x-\mu_{Y|A}}{\sigma_{Y|A}}\right)dx + \frac{y}{\sigma_{Y|A}}\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\right]\delta\mu_{Y|A}\right] \\ &+\mathbf{E}\Bigg[\Bigg[-\int_y^\infty \frac{\mu_{Y|A}-x}{\sigma^2_{Y|A}}\psi_A\left(\frac{x-\mu_{Y|A}}{\sigma_{Y|A}}\right)dx + \frac{y\left(y-\mu_{Y|A}\right)}{\sigma^2_{Y|A}} \\ &\times\psi_A\left(\frac{y-\mu_{Y|A}}{\sigma_{Y|A}}\right)\Bigg]\delta\sigma_{Y|A}\Bigg] \end{aligned}$$

As with the VAR, we require the tail probability to be unaffected by the perturbation, so the quantile (y) has to vary while $\delta P^- = 0$. Using (11) we eliminate the y-differential by considering $\delta(P^+S) + y\delta P^-$. This, after some tidying up, gives us:

$$\begin{aligned} \delta\left(P^+S^+\right) = &\mathbf{E}\left[\int_y^\infty \frac{1}{\sigma_{Y|A}}\psi_A\left(\frac{x-\mu_{Y|A}}{\sigma_{Y|A}}\right)dx\,\delta\mu_{Y|A}\right] \\ &+\mathbf{E}\left[\int_y^\infty \frac{x-\mu_{Y|A}}{\sigma^2_{Y|A}}\psi_A\left(\frac{x-\mu_{Y|A}}{\sigma_{Y|A}}\right)dx\,\delta\sigma_{Y|A}\right] \end{aligned}$$

The desired CI-monotonicity property is proved once we are happy that the integrals in both square-bracketed terms are positive. The first clearly is because the integrand is positive, but the second requires a little more work. Condition on A: if $y > \mu_{Y|A}$ then the integrand is positive, while if not we write the integral as $\int_{-\infty}^{\infty} - \int_{-\infty}^{y}$, in which the first integral is zero (by normalisation of ψ_A) and the integrand in the second is negative.

The proof of the CI-monotonicity of ESF shows that ESF is not the only coherent risk measure that has this property. Any risk measure that can be written as a linear combination (with positive coefficients that add up to one) of ESFs at different levels P^+

[4] There is no immediate interpretation of higher-order conditional moments in terms of systematic and unsystematic risk. Therefore, the definition of CI-monotonicity involves only the first two moments

will also be CI-monotonic. Theorem 7 in Kusuoka (2001) shows that all law-invariant and comonotone-additive coherent risk measures can be represented this way. As shown in Theorem 3.6 of Tasche (2002), these risk measures are just the spectral measures that were introduced by Acerbi (2002).

Law-invariance of a risk measure means that the value of the risk measure depends on its argument only through the probability distribution of the argument. This property is implicitly required by the definition of CI-monotonicity. Thus, there are no CI-monotonic risk measures that are not law-invariant. Comonotone additivity of a risk measure means that the measure is additive with respect to comonotonic arguments, that is, with respect to random variables that can be represented as increasing functions of one common random factor. Theorem 4 in Kusuoka (2001) implies that, in general, law-invariant coherent risk measures that are not comonotone-additive will not be CI-monotonic.

Similarly, CI-monotonic risk measures are not coherent in general. Consider, for instance, the risk measure defined as the third power of ESF. This risk measure is CI-monotonic but not one-homogeneous and hence not coherent.

Concluding remarks

We have not said much about what the risk factor A corresponds to in practice, because an interpretation is unnecessary for the development of the theory. In one-factor credit risk models, A corresponds to the 'general state of the economy'. More complex portfolios require multi-factor models, and it is quite common[5] to use industrial sectors to represent the extra factors. In a multi-factor implementation, the theory allows the following distinction to be made in a fairly precise way. Consider two portfolios, both of which have significant exposure in (say) the auto sector: portfolio A's exposure arises from having many auto issuers, but portfolio B's arises from a large exposure to (say) Ford. If we simply aggregated the issuers' risk contributions by sector, we might find that the pictures for A and B were similar, but the technique described in this article would show that B had significantly more unsystematic risk. This would point to the conclusion that the right hedge for portfolio A was auto basket protection, and for portfolio B a single-name CDS on Ford.

A question that naturally arises is whether portfolio risk can be split down by issuer and by 'systematic/unsystematic' at the same time. In the context of the example in figure 4, we are inviting the conclusion that most of the unsystematic risk comes from assets to which the exposures were increased, because those are big single-name exposures that would require single-name hedges. The answer is yes, and this will be presented in forthcoming work. ■

Richard Martin is head of quantitative credit strategy at Credit Suisse in London. Dirk Tasche is risk analyst at the Deutsche Bundesbank in Frankfurt.
The authors would like to thank two anonymous referees for providing useful comments on an earlier draft of this article. The opinions expressed in this article are those of the authors and do not necessarily reflect the views of their institutions.
Email: richard.martin@credit-suisse.com, tasche@ma.tum.de

[5] *For example, in Credit Suisse's PortfolioRisk+, at* www.portfolioriskplus.com

References

Acerbi C, 2002
Spectral measures of risk: a coherent representation of subjective risk aversion
Journal of Banking and Finance 26(7), pages 1,505–1,518

Acerbi C and D Tasche, 2002
On the coherence of expected shortfall
Journal of Banking and Finance 26(7), pages 1,487–1,503

Artzner P, F Delbaen, J-M Eber and D Heath, 1999
Coherent measures of risk
Mathematical Finance 9, pages 203–228

Basel Committee on Banking Supervision, 2006
Basel II: international convergence of capital measurement and capital standards: a revised framework – comprehensive version

Gordy M, 2003
A risk-factor model foundation for ratings-based bank capital rules
Journal of Financial Intermediation 12(3), pages 199–232

Gordy M, 2004
Granularity adjustment in porfolio credit risk management
In Risk Measures for the 21st Century, edited by G Szegö, Wiley

Gouriéroux C, J-P Laurent and O Scaillet, 2000
Sensitivity analysis of values at risk
Journal of Empirical Finance 7(3), pages 225–245

Kusuoka S, 2001
On law invariant risk measures
Advances in Mathematical Economics 3, pages 83–85

Martin R, 2006
The saddlepoint method and portfolio optionalities
Risk December, pages 93–95

Martin R and T Wilde, 2002
Unsystematic credit risk
Risk November, pages 123–128

Tasche D, 1999
Risk contributions and performance measurement
Working paper, Technische Universität Munchen

Tasche D, 2002
Expected shortfall and beyond
Journal of Banking and Finance 26(7), pages 1,519–1,533

Guidelines for the submission of technical articles

Risk welcomes the submission of technical articles on topics relevant to our readership. Core areas include market and credit risk measurement and management, the pricing and hedging of derivatives and/or structured securities, and the theoretical modelling and empirical observation of markets and portfolios. This list is not an exhaustive one.

The most important publication criteria are originality, exclusivity and relevance – we attempt to strike a balance between these. Given that *Risk* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the technical editor and his team, submissions are logged, and checked against the criteria above. Articles that fail to meet the criteria are rejected at this stage.

Articles are then sent to one or more anonymous referees for peer review. Our referees are drawn from the research groups, risk management departments and trading desks of major financial institutions, in addition to academia. Many have already published articles in *Risk*. Depending on the feedback from referees, the technical editor makes a decision to reject or accept the submitted article. His decision is final.

We also welcome the submission of brief communications. These are also peer-reviewed contributions to *Risk* but the process is less formal than for full-length technical articles. Typically, brief communications address an extension or implementation issue arising from a full-length article that, while satisfying our originality, exclusivity and relevance requirements, does not deserve full-length treatment.

Submissions should be sent to the technical team at *technical@incisivemedia.com*. The preferred format is MS Word, although Adobe PDFs are acceptable. The maximum recommended length for articles is 3,500 words, and for brief communications 1,000 words, with some allowance for charts and/or formulas. We expect all articles and communications to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations. Authors should allow four to eight weeks for the refereeing process.