The Case for
Fully Integrated Models of Economic Capital

19th February 2009, DRAFT VERSION 5

Alexander J. McNeil

Maxwell Institute for the Mathematical Sciences, Edinburgh EH14 4AS, UK
`email: a.j.mcneil@hw.ac.uk`

Gavin Kretzschmar

University of Edinburgh

Axel Kirchner

Barrie & Hibbert Limited

Abstract

Economic capital models are potentially powerful tools for enterprise risk management (ERM), and for the supervisory review process (Pillar 2) of the Basel II and Solvency II regulatory capital frameworks. We argue that, to fulfill this potential, economic capital models need to be fully integrated and to go beyond the more modular approaches that dominate Pillar 1 methodology. In a modular approach capital is determined at business-unit or risk category level (e.g. market, credit and liquidity risk separately) and aggregated ex post by simple summation or correlation-adjusted summation; in a fully integrated approach aggregation occurs implicitly by relating all risks to a common set of fundamental risk drivers.

We explain how calibrated economic scenario generation lies at the heart of a fully integrated approach to modelling the risks on the asset side of a firm's balance sheet and discuss how stochastic scenario generation gives the ideal framework for exploring the diversification benefits that different units or asset classes bring to an enterprise. We explain how this approach allows us to understand the sources of tail risk and gives us a platform for integrated stress testing, sensitivity analysis, and the allocation of capital to business units for risk-adjusted performance comparisons.

Key words: risk management, economic capital, enterprise risk management, Basel II, Solvency II, stochastic models, stress testing

1 Introduction

Economic capital is the capital required by a bank/insurer to limit the probability of insolvency to a given level over a given horizon; see, for example, Burns (2004). As its name suggests, it is an attempt to measure risk in terms of economic realities, rather than in terms of the simplifying rules that dominate Pillar 1 methodology for regulatory capital. Indeed we can interpret Basel II and Solvency II requirements in the Pillar 2 (supervisory review) area as requiring economic capital models. Such a model should afford the opportunity to investigate the true risk sensitivities of an enterprise and engage in a meaningful dialogue with the regulator on capital adequacy. Proponents of economic capital models have more ambitious aims; they see economic capital as a potential enterprise-wide language for discussing risk, comparing returns on risk capital and improving the efficiency of capital usage. In other words, economic capital models can serve as the main platform for the development of a fully-fledged enterprise risk management (ERM) system in a financial institution.

The Basel II accord on Banking Supervision. (2004) articulates a desire for comprehensive enterprise-wide risk assessment as part of the supervisory review. In paragraph 748 it states that "supervisors should assess the degree to which internal targets and processes incorporate the *full range* of material risks faced by the bank" (our italics). It also requires that "supervisors should consider the results of sensitivity analyses and stress tests conducted by the institution and how these results relate to capital plans." Basel II is, however, noticeably silent on the issue of methodology for Pillar II and has been content to adopt a non-prescriptive approach that has left model choice to the banks themselves.

A 2007 study of 17 banks and 16 insurers (IFRI Foundation and CRO Forum (2007)) provides a snapshot of the current state of development of economic capital models across the industry. It finds that "there is significantly increased experience in using Economic Capital across the whole financial services sector (e.g. for banking, frameworks have been in place an average of over 6 years and for insurance 4) and firms now feel broadly comfortable with the accuracy of outputs (75%+ for both insurance and banking)." It goes on to note significant differences in some of the fundamental methodology used to compute economic capital, particularly with regard to treating diversification effects. Banks continue to favour approaches based on variance-covariance or correlation matrices whereas some insurers are pursuing simulation-based approaches. However the study finds that of the surveyed firms "20% of partic-

⋆ The authors thank John Morrison for commenting on a preliminary version of the paper.

ipants use a simple summation approach (which might be considered a Pillar 1 as opposed to a Pillar 2 approach) and of those quantifying diversification benefits, over three quarters use an analytical variance/covariance approach".

However recent stress events in capital markets have served to demonstrate that correlation-based approaches to economic capital fall well short of capturing (and more importantly understanding) risk and asset valuation dependencies. In this paper it will be argued that:

(1) the *modular* approach, based on computation of capital within different business lines and asset classes and aggregation using correlation matrices, is a conceptually flawed procedure;
(2) a *fully-integrated* approach based on an enterprise-wide stochastic model is preferable, but requires a credible economic scenario generation engine;
(3) a fully-integrated solution has the potential to deliver the framework for sensitivity analysis, stress testing, risk-adjusted performance comparison and enterprise steering, which is the ultimate goal of the economic capital project, and which would align with the Pillar II requirements of Basel II and Solvency II.

The apparent advantage of the modular approach is that it offers a simple spreadsheet-based solution to capital calculation in an institution that has decentralized risk modelling to the individual business units. But the price of this simplicity may be risk capital numbers that bear no relationship at all to the stochastic drivers of risk and to their real-world relationships. While modular approaches currently appear to predominate in industry, there is a growing number of examples of risk modelling systems that are moving in the direction of full integration; an example we refer to in this paper is the Barrie & Hibbert economic scenario generator (ESG) (Morrison et al. (2008)).

The paper is structured as follows. In Section 2 we describe a heavily stylized stochastic model to serve as a framework for the computation of economic capital. In Section 3 we look in more detail at the issue of diversification across business lines and asset classes and the drawbacks of modular (asset silo by asset silo) approaches to calculating capital that are not based on fully integrated models.

We discuss the possible architecture of a fully integrated stochastic model in Section 4 while Section 5 describes some of the activities and applications that become possible in a fully integrated model. We devote particular attention to discussions of fully integrated stress testing and fair capital allocation.

2 A Stylized Model of an Enterprise

Economic capital does not yet enjoy an accepted and stable definition; see, for example, Frankland et al. (2008). In this paper we adopt one of the more precise interpretations of the concept, namely that economic capital is the capital that ensures the continued technical solvency of an enterprise over a given time horizon with a given (and extremely conservative) *probability*. It seems clear that this view requires a *stochastic modelling approach*, and we set out such a model in this section. Note that the argument below is heuristic rather than realistic and is intended to apply to a generic enterprise, whether bank, insurer or other.

Suppose that at time t the asset portfolio of the enterprise has value V_t and the liabilities on the balance sheet have fair value B_t. Assume that $V_t > B_t$ so that the enterprise is initially solvent with initial equity $E_t = V_t - B_t$. A risk manager assessing capital adequacy at time t is concerned with fluctuations in these values over the time interval $[t, t+1]$ (typically one year). At the end of the period assets and liabilities will have values V_{t+1} and B_{t+1}.

Throughout the period the enterprise will generate income on its asset portfolio but will also have interest payments to make on liabilities. It may also earn income from providing services (for example a bank will earn income from fees and commissions) but will incur operating costs. Let us suppose the net income from all these sources is I_{t+1}. To ensure solvency at time $t+1$ with a (high) probability α the risk manager requires that

$$P(V_{t+1} - B_{t+1} + I_{t+1} > 0) = \alpha$$

or equivalently that

$$P(E_t + (V_{t+1} - V_t) - (B_{t+1} - B_t) + I_{t+1} > 0) = \alpha.$$

Let us introduce the notation $\Delta_{t+1} = V_{t+1} - V_t$ for asset value growth and the notation $L_{t+1} = -\Delta_{t+1}$ for losses in asset value. In terms of the former we require that

$$P(\Delta_{t+1} + E_t > (B_{t+1} - B_t) - I_{t+1}) = \alpha$$

or in terms of the latter that

$$P(L_{t+1} < E_t + I_{t+1} - (B_{t+1} - B_t)) = \alpha$$

Suppose the enterprise runs a replicating portfolio for liabilities and can be "almost sure" that

$$E(\Delta_{t+1}) \geq (B_{t+1} - B_t) - I_{t+1}$$

so that expected asset value growth more than covers the growth in the value of liabilities plus any shortfall in income. The enterprise would then be conservatively capitalized if

$$P(\Delta_{t+1} - E(\Delta_{t+1}) + E_t > 0) = \alpha$$

or equvalently if

$$P(L_{t+1} - E(L_{t+1}) < E_t) = \alpha.$$

This can be achieved if the initial equity of the company is set as the α-quantile of the distribution of $L_{t+1} - E(L_{t+1})$, which introduce the Value-at-Risk or VaR idea.

In this paper we will concentrate on the asset side of the balance sheet and consider capital to be determined by analysis of the loss distribution of L_{t+1}. We note, however, that there is no reason why full stochastic modelling of both assets and liabilities, or equivalently of net asset value, cannot be undertaken.

The VaR concept, whereby capital is calculated as a quantile of the (mean-corrected) loss distribution, has now been widely adopted, but VaR is only one possible measure of risk that can be applied to a loss distribution; the risk measure variously known as expected shortfall, conditional VaR or tail VaR is an attractive alternative with better theoretical properties Tasche (2002).

As indicated above, it is usual in economic capital calculation to subtract *expected loss* $E(L_{t+1})$ from the measure of risk; see also Burns (2004). If we used the VaR risk measure at level α (say 99.97%) the economic capital would be determined to be

$$\mathrm{EC} = \mathrm{VaR}_\alpha(L_{t+1}) - E(L_{t+1}) = \mathrm{VaR}_\alpha\left(L_{t+1} - E(L_{t+1})\right),$$

where VaR_α denotes the α-quantile operator on the distribution of a random variable. This capital would be interpreted as being sufficient to reduce the probability of insolvency to $1-\alpha$ (say 3 in 10000) over the time period $[t, t+1]$.

2.1 The Challenge of Valuation

Valuing the portfolio in the present (V_t) and in the future (V_{t+1}) is a significant challenge. The ideal method is *market-consistent* (or fair-value) valuation which is an attempt to determine what we would get for the assets if we attempted to dispose of them on the market.

Consider first the market-consistent valuation of the portfolio at time t (now). For a large part of the portfolio it ought to be possible to simply look up currently quoted prices. This approach is known by practitioners as *marking-to-market*, but its feasibility is contingent on access to prices for a liquid market in the assets in question. While it is certainly true that, for example, equities, bonds, foreign currency and exchange-traded derivatives can be marked to market, the situation is more difficult for such assets as over-the-counter (OTC) derivatives, books of mortgage loans and the various classes of structured (or securitized) debt instruments, such as MBSs (mortgage-backed securities) and CDOs (collaterized debt obligations).

If no market prices are available for certain assets or if trading is very thin then certain instruments can be valued by *marking-to-model*. This is the practice of assigning a value to a position based on a pricing model that makes reference to other underlying instruments whose values may be observed in the market. As such it can be viewed as a kind of interpolation procedure. But the quality of a mark-to-model valuation is questionable in a period of financial market stress; many of the valuation models prove to be fair weather models that do not inspire high levels of confidence in a crisis.

Events of recent times have shown the difficulties of market-consistent valuation in challenging market conditions. The credit write-downs have resulted from a continual process of trying to determine a market-consistent value for illiquid securitized debt assets in a climate of great uncertainty concerning the value of the underlying debt. Inevitably many of these assets have been marked-to-model and the dramatic revisions in values as model parameters have changed have contributed to the popular opinion that valuation has a large element of “mark-to-myth”. See Crouhy et al. (2007) for further discussion of the problems of valuation during the subprime credit crisis.

Note that our sketch of the different levels of market-consistent valuation mirrors proposed developments in accounting practice that are currently taking place. The international accounting standards boards (IASB) are soliciting comments on proposed amendments to IFRS 7, the standard that relates to international fair-value accounting. The aim is to improve the quality of financial disclosures by using a three-level fair value hierarchy with (roughly speaking) mark-to-market valuation at level 1, mark-to-model with fully observed market inputs at level 2, and mark-to-model with unobserved inputs at level 3.

Valuing assets at time t (now) is difficult, but valuing them at time $t+1$ is an order of magnitude more challenging. Since market prices in the future are unknown, they must be forecast, but it would be an impossible task to forecast prices for all assets. In practice we project forward a future market in certain fundamental underlying *risk factors* that affect all asset valuations; examples of such risk factors are as representative equity prices (perhaps concentrating on indexes), interest rates, currencies and yields on bonds.

While certain assets (such as share or bond holdings) may be marked directly to this hypothetical market, other positions (such as derivatives) will generally be marked to a model that takes the forecasts of these underlyings as inputs.

The parameters in these valuation formulas (for example, the volatility parameter in the Black Scholes option pricing formula) also have to be forecast and are viewed as elements in the set of risk factors.

A further issue complicates the determination of V_{t+1} and that is the fact that the asset mix in the portfolio may conceivably change over the period $[t, t+1]$ as a result of active portfolio management. It will be necessary to specify rebalancing strategies and rules.

2.2 The Role of Economic Scenario Generation

Projecting forward the underlying risk factors for purposes of valuation at $t+1$ is the role that can be filled by an economic scenario generator. Suppose we denote the risk factors at time t by the vector $\mathbf{Z}_t = (Z_{t1}, \ldots, Z_{td})$. We set up a multivariate stochastic process $\mathbf{Z} = (\mathbf{Z}_s, s \geq t)$ which projects the values of these risk factors into the future and gives us snapshots $\mathbf{Z}_s$ of the economy at future times $s \geq t$. The value of the portfolio at future times V_s can be considered as a random variable of the form

$$V_s = f_s(\mathbf{Z}_s, s) \tag{1}$$

where f_s is a function that we will refer to as the *portfolio mapping at time* s. It contains information about the portfolio composition at time s and incorporates the valuation formulas that are used to value the more complex (derivative) assets with respect to the underlying risk factors $\mathbf{Z}_s$. Note that, in general, it depends not only on the value of the risk factors at time s, but also on the time s itself; this is because the value of a derivative position with maturity/expiry T typically depends on the remaining time to maturity $T-s$.

Note also that there is a time subscript on the mapping function f_s to allow for the possibility of dynamic rebalancing which could change the entire composition of the mapping over time.

An economic scenario generator generally takes a Monte Carlo (simulation) approach and generates a series of realisations/paths $(\mathbf{Z}_s(\omega_i), s \geq t)$ for $i = 1, \ldots, m$. Each ω_i is in effect the label for a particular economic scenario and $(\mathbf{Z}_s(\omega_i), s \geq t)$ is the manifestation of that scenario in terms of a path for the fundamental risk factors.

2.3 Conditional versus unconditional

There are subtle questions concerning the stochastic model that is used to generate $\mathbf{Z} = (\mathbf{Z}_s, s > t)$. The most natural form of projection is conditional on past realized values of the process $(\mathbf{Z}_s, s \leq t)$. This effectively conditions the projection on the current and recent economic climate, whether benign or stressed. There is some debate whether this is always desirable or appropriate. In the case of a bank an argument could be made for using a conditional calibration for trading book assets but it may be less appropriate for assets in the banking book that are typically held to maturity. Moreover, in a benign epoch a conditional projection may fail to anticipate stress events over the horizon, whereas in volatile periods economic capital calculations based on conditional projections will share that volatility and actions taken on the basis of these figures may be pro-cyclical and serve to amplify volatility.

An alternative approach is to calibrate the stochastic process $\mathbf{Z} = (\mathbf{Z}_s, s > t)$ based on longer-term information including historical tail events and to generate an unconditional realization that does not overweight the most recent realized values.

In Pillar 1 regulatory methodology a mixing of conditional and unconditional approaches tends to take place. Basel II market risk measurement practice varies: it may be based on the *variance-covariance* (or delta-normal) method which usually involves time-series-based forecasts of risk factor volatilities and correlations and is thus conditional; or it may be based on *historical simulation* which is an unconditional method involving resampling of historical data in the absence of a dynamic model; see McNeil et al. (2005) for more detail. The Basel II formula used in the internal-ratings-based (IRB) approach to credit risk can be implemented in both a conditional or unconditional manner, depending on whether the internal assessments of PDs (probabilities of default) are considered to be *point-in-time* or *through-the-cycle*.

The literal definition of economic capital in terms of limiting the probability of insolvency over a defined time horizon seems to come down on the side of conditional calibration as being appropriate, but a less literal interpretation allows us more space for interpretation. There might be an argument for computing the capital as a maximum of the numbers derived from a conditional and an unconditional calibration; this would have the effect of smoothing out the numbers in more benign periods and elevating them only in crises, so that the procyclicalty problem is lessened (although still present in crises).

3 Dealing with Diversification

The pooling of risks across portfolios, business lines, organizations achieves diversification and intuitively this benefit depends on the degree of dependence between the pooled risks. It is generally accepted that aggregate economic capital should reflect the diversification benefit and therefore be less than the sum of the economic capital requirements that each line would require on a stand-alone basis.

An enterprise-wide stochastic risk model, which models the mutual dependence of asset value changes across an enterprise on a number of risk factor processes, permits us to evaluate overall economic capital for the enterprise and compare this with the sum of stand-alone capital requirements. It also permits us to compare the aggregate capital obtained in this way with the aggregate capital obtained from what we term *a modular approach*.

A company without an enterprise-wide stochastic model of economic capital is forced to compute its capital requirements from the ground up, sub-portfolio by sub-portfolio, taking account of assumed diversification in a final aggregation step which generally involves assumptions about correlations between sub-portfolios. This modular (or silo) approach features prominently in the Solvency II regulatory framework for insurers.

Let us suppose that we divide our institution by business line or asset class into d sub-portfolios. For each sub-portfolio $j = 1, \ldots, d$ we have to consider possible losses

$$L_{j,t+1} = -\Delta_{j,t+1} = -(V_{j,t+1} - V_{j,t}),$$

which aggregate by simple summation to give the overall value change of the enterprise

$$L_{t+1} = -(V_{t+1} - V_t) = -\left(\sum_{j=1}^{d} V_{j,t+1} - \sum_{j=1}^{d} V_{j,t}\right) = \sum_{j=1}^{d} L_{j,t+1}. \tag{2}$$

In what follows we suppress the time subscripts $t+1$ and concentrate on cross-sectional behaviour for a fixed time period.

3.1 Capital by asset class or risk factor?

Imagine an enterprise as a matrix: the rows are the various asset classes (or business lines and units) which sum to give the overall asset position of the firm as in (2); the columns are the various risk factors (interest rate risk, credit risk, equity market risk, etc.) which affect valuations of assets.

In this paper we take the view that the only natural way to disaggregate the enterprise is by asset class (by row); it is much more artificial to attempt to disaggregate the enterprise by risk factor (by column). We can remove entire classes of assets from the enterprise by selling them and it is quite natural to think of the capital required to hold that asset class as a stand-alone entity. In contrast it is much less natural to think of the capital required for, say, credit risk alone. Credit risk cannot be simply sold out of the enterprise. At best appropriate hedging stategies with derivatives can be adopted to mitigate credit risk. The fundamental distinction is that a risk factor has a pervasive and generally non-linear influence across the assets of an enterprise and cannot easily be isolated, whereas asset classes sum up as in (2) and can generally be decoupled from each other.

Unfortunately Basel II and Solvency II have fostered a style of analysis that is based on disaggregation by risk factors. This largely reflects the development of the regulation where attention has focussed in turn on market risk, credit risk, operational risk, interest-rate risk in the banking book, liquidity risk etc.. Considering the first two of these, it has become standard practice to compute capital for market risk and credit risk separately and to aggregate these by summation, or correlation-adjusted summation to account for imperfect dependence between the two. However, there is no satisfactory theory to support aggregation of this kind. As pointed out in Breuer et al. (2008) the justification runs along the lines of dividing a bank's assets into banking and trading book assets and equating credit risk to the risk of the banking book and market risk to the risk of the trading book. Such assumptions are untenable since many positions depend on both risk factors. Moreover, Breuer et al. show that there can be complex interaction between the risk factors so that the capital required to cover a position may exceed the sum of the capital required to cover its market risk element and the capital required to cover its credit risk element. Alessandri and Drehmann (2008) discuss similar conceptual issues with respect to the integration of credit and interest-rate risk, although they find empirical evidence to suggest that the problem of adverse interaction of these risk types is less likely.

It is worth noting that the use of stochastic scenario generation as the basis of risk modelling does facilitate the common practice of "switching off" certain risk factors and observing the effects on asset positions of the stochastic behaviour of a single risk factor. In terms of the framework set out in (1) the switching off of risk factors corresponds to fixing certain component of $\mathbf{Z}_s$ at deterministic values. The exercise is artificial since risk factors cannot be turned off in reality and there are dangers in neglecting the risks arising from the interaction of risk factors as Breuer et al. (2008) point out. Despite lacking a clear theoretical justfication, it could be argued that the practice gives indications of the relative impacts of risk factors and, consequently, clues about the kind of hedging behaviour that might produce the greatest mitigation of risk.

3.2 Diversification and correlation

Attempting to express diversification across asset classes in terms of correlation is a problematic area. Much of the lay-person's intuition about correlation is based on a multivariate normal world view and there are subtle pitfalls to be aware of when we step outside this idealized world; see Embrechts et al. (1999) and Embrechts et al. (2002). For example, a UK Solvency II publication (Treasury and FSA (2006)) makes the seemingly reasonable claim that "diversification benefits can be assessed by correlations between different risk categories. A correlation of +100% means that two variables will fall and rise in lock-step; any correlation below this indicates the potential for diversification benefits."

The last sentence is not true of Pearson's product-moment correlation, i.e. the standard linear correlation of common usage, although it would be true of a rank correlation, such as Spearman's rho or Kendall's tau.

To explain the fallacy we will give a definition to the informal concept of "lock-step" used above by equating it to comonotonicity. This means all risks are increasing functions of a common underlying risk: $(L_1, \ldots, L_d) = (v_1(Z), \ldots, v_d(Z))$. Such risks would be considered undiversifiable since a high realization of Z will lead to large losses across the enterprise with no possibility that any $L_i = v_i(Z)$ can be offset by another $L_j = v_j(Z)$.

While it is certainly true that a correlation equal to one, implies that two risks are comonotonic, it is not necessarily true that two comonotonic risks always have a correlation of one. We can consider models where individual risks move in lock-step (are undiversifiable), but have an arbitrarily small correlation. For two given distributions, the *attainable correlations* form a sub-interval of $[-1, 1]$. The upper bound corresponds to comonotonicity, and the lower to countermonotonicity (negative lock-step), but these need not be 1 and -1. An example is given in McNeil et al. (2005) showing how these bounds vary for two lognormal risks.

3.3 The Modular Approach

In a modular approach to capital adequacy individual risks at sub-unit level are transformed into capital charges $\mathrm{EC}_1, \ldots, \mathrm{EC}_d$. These are then combined to calculate the overall economic capital EC.

This method is described in CEIOPS-06 (2006) (page 71) as a *bottom-up* approach, although we will avoid the confusing terminology of bottom-up and *top-down* and contrast the modular approach with the *fully integrated* approach. By the latter we mean an approach where overall economic capital is calculated from a single model of the enterprise that describes the mutual dependence of different risks on common risk drivers as described in Section 2.

In the modular approach the combination of individual capital charges is usually based on a calculation of the following kind:

$$\mathrm{EC} = \sqrt{\sum_{i=1}^{d}\sum_{j=1}^{d} \rho_{ij}\mathrm{EC}_i\mathrm{EC}_j} \tag{3}$$

where the ρ_{ij} are the "correlations" between the risks. See, for example CEIOPS-06 (2006), page 98. It may be reasonably asked, what principle underlies this method of aggregation?

Proposition 3.1 *The aggregation rule (3) may be justified when the following two assumptions hold.*

(1) The risks $(L_1, \dots, L_d)$ have a joint elliptical *distribution with $\rho(L_i, L_j) = \rho_{ij}$.*
(2) Economic capital for a risk X is computed using to a law-invariant, positive-homogeneous *and* translation-preserving *risk measure ϱ according to $\mathrm{EC} = \varrho(X) - E(X)$.*

Examples of elliptical distributions are the multivariate normal and the heavier-tailed multivariate Student t distribution. Law-invariant risk measures depend only on the distribution function of X. Positive-homogeneous, translation preserving risk measures satisfy the requirements that $\varrho(\lambda X) = \lambda\varrho(X)$ and $\varrho(X + c) = \varrho(X) + c$ for constants $\lambda > 0$ and $c \in \mathbb{R}$. Examples are Value-at-Risk (VaR) and expected shortfall.

A full proof of the proposition above can be based on Theorem 6.8 and Proposition 6.13 in Embrechts et al. (2002). We give a demonstration in the case where the risk measure is $\varrho = \mathrm{VaR}_\alpha$ for $\alpha > 0.5$ and the risks have a joint normal distribution. If $L = L_1 + \cdots + L_d$, then

$$\frac{L - E(L)}{\mathrm{sd}(L)} \sim N(0,1) \quad \text{and} \quad \frac{L_i - E(L_i)}{\mathrm{sd}(L_i)} \sim N(0,1)$$

and it follows from the positive-homogeneity and translation-preservation properties of VaR that $\mathrm{VaR}_\alpha(L) - E(L) = \lambda_\alpha\,\mathrm{sd}(L)$ and $\mathrm{VaR}_\alpha(L_i) - E(L_i) = \lambda_\alpha\,\mathrm{sd}(L_i)$ where λ_α is the α-quantile of standard normal.

We use the fact that (irrespective of the distribution of $(L_1, \dots, L_d)$) the standard deviation aggregates according to

$$\mathrm{sd}(L) = \sqrt{\sum_{i=1}^{d}\sum_{j=1}^{d} \rho_{ij}\,\mathrm{sd}(L_i)\,\mathrm{sd}(L_j)} \tag{4}$$

to infer that

$$\mathrm{VaR}_\alpha(L) - E(L) = \sqrt{\sum_{i=1}^{d}\sum_{j=1}^{d} \rho_{ij}\,(\mathrm{VaR}_\alpha(L_i) - E(L_i))\,(\mathrm{VaR}_\alpha(L_j) - E(L_j))},$$

which justifies (3).

There are a number of issues with the modular approach embodied in (3). First, it is only underpinned by theoretical principles in a very specific and unrealistic model of the risk universe. The potential losses in each of the subunits are unlikely to follow an elliptical distribution, which would preclude all possibilities of skewness and impose a common symmetric form on each of the L_i.

Second, it is dependent on the widely misunderstood concept of correlation. How exactly are the correlations to be determined? The kinds of risks where we have reliable empirical experience of typical correlation values are in the minority (for example, equity market risks at shorter time horizons). In the absence of empirical experience we are reliant on "expert opinions" but there are a number of consistency requirements that the experts must respect. As we have observed, every ρ_{ij} must be compatible with the distribution of L_i and L_j. Moreover, the matrix (ρ_{ij}) must be positive semi-definite. It is quite easy to specify nonsensical correlation matrices.

More fundamentally, a correlation only provides a scalar summary of the dependence between two risks and does not attempt to explain the origins of that dependence. The idea that stress events may lead to higher levels of "extremal dependence" or "tail dependence" between risks has been absorbed by risk practitioners but the modular approach does not offer a natural way of modelling this phenomenon.

In CEIOPS-06 (2006) (page 142) the advice is given that "when selecting correlation coefficients, allowance should be made for tail correlation. To allow for this, the correlations used should be higher than simple analysis of relevant data would indicate." But clearly this is an unsatisfactory approach that does not address the fundamental question of where the tail correlation comes from; in contrast a fully integrated model does provide a framework for attempts to model the sources of tail correlation.

In recent years the problems surrounding the use of correlation as an integration tool have become more widely known, although this has had limited effect on risk management practice. There has, however, been some interest in the use of copulas for modelling dependence of risks across business lines and asset classes.

In contrast to correlation, a copula does offer a complete description of the dependence between risks, rather than a scalar summary thereof; this description takes the form of a probability distribution. Suppose that the individual risks $L_i, i = 1, \ldots, d$ are modelled using univariate models $F_i(x) = P(L_i \leq x)$. Then by choosing a suitable copula C we can build up a complete multivariate distributional model for $(L_1, \ldots, L_d)$ of the form

$$F(x_1, \ldots, x_d) = P(L_1 \leq x_1, \ldots, L_d \leq x_d) = C(F_1(x_1), \ldots, F_d(x_d)).$$

The tricky consistency requirements of working with correlation matrices are avoided, since any copula C can be used to build a valid multivariate model, and we have the possibility of choosing copulas that are known to yield higher levels of tail dependence. To compute aggregate capital in a copula-based model, Monte Carlo simulation will generally be necessary, since simple analytical formulas like (3) are seldom available.

In insurance, aggregation of loss distributions with Gauss copulas has a long history and a method described by Iman and Conover (1982) is widely used; see CAS (2006) for some discussion of this and related methods. However there has also been interest in using alternative copulas, like the t copula, to obtain, higher levels of tail dependence for a prescribed correlation level; see also Daul et al. (2003).

But the Achilles heel of the modular method - the calibration issue - is not avoided by the use of copulas. In the absence of multivariate loss data, which would allow us to estimate C empirically, we are again reliant on the opinions of experts. This time the experts must decide on an appropriate copula family and then determine the parameters of the copula. For the Gauss copula (and for the t copula) this is usually done by fixing a matrix of rank correlations and inferring copula parameters from these. Asking experts to fix matrices of *rank* correlation parameters in the absence of data may be asking too much.

The general problem of the modular approach remains that the specification of dependence is exogenous to the individual risk models. It involves an ex post, phenomenological description of the effects of dependence rather than an explicit structural attempt to explain the causes of dependence.

Losses in sub-portfolios depend on value changes ($L_{j,t+1} = -\Delta_{j,t+1} = -(V_{j,t+1} - V_{j,t})$) and future valuations are driven by fundamental risk factors $\mathbf{Z}_{t+1}^{(j)}$ according to $V_{j,t+1} = f_{j,t+1}(\mathbf{Z}_{t+1}^{(j)}, t+1)$. Many of these risk factors, for example those describing the structure of the yield curve or the average performance of equity markets, are common to many sub-portfolios of assets.

This is the origin of dependence in a fully integrated model: correlation arises from the mutual dependence of future values across an enterprise on a set of common risk drivers. Fully integrated models are common factor models. The risk factors $\mathbf{Z}_{t+1}^{(j)}$ that enter into the future valuation of sub-portfolio j contain a subset in common with the risk factors $\mathbf{Z}_{t+1}^{(k)}$ that enter into the future valuation of sub-portfolio k. These common factors are the drivers of dependence between $V_{j,t+1}$ and $V_{k,t+1}$ and consequently between $L_{j,t+1}$ and $L_{k,t+1}$. The dependence arises endogenously through the specification of the model.

While, conceptually, this is a much more satisfactory way of modelling dependence, the performance of a fully integrated model is crucially dependent on the "correct" identification of common risk drivers, the "correct" specification of the mapping functions $f_{j,t+1}$ and their "correct" calibration. Of course no models are ever correct, but the common factor models we build must be as plausible as we can make them and particular effort must go into the modelling of the common factors and the linking of these factors to risks at sub-portfolio level. When we first set up a fully integrated model, sensitivity analyses are necessary to help us understand the critical components of the model that have the largest impact on model outputs. They can show us where to focus resources in gathering empirical evidence to support the model and to refine its detailed calibration.

What must be stated clearly is that fully integrated models demand a much higher level of expertise and understanding in those who construct and operate them. Moreover, fully integrated models, being based on economic scenario generation, will generally be used in a Monte Carlo simulation mode, as we observed in Section 2.2. This may impose a considerable computational burden and require long run times to obtain sufficiently rich results.

4 Architecture of a Fully Integrated Model

This paper is not about the specific design of any particular model, but for the purposes of debate we sketch out a possible architecture. The extent of the model, and the risks that are covered, will depend on the nature of the firm under consideration; see Sweeting (2007) for a comparative survey of the risks that affect banks, insurance companies and pension funds. Our treatment is oriented towards the risks that face the assets of a typical bank. Figure 1 provides a diagrammatic overview of the components of a possible model. Footnotes in this section give an idea of the choices that have been made in the construction of a particular example of an economic scenario generator, namely the Barrie & Hibbert ESG ; further details may be found in Morrison et al. (2008).

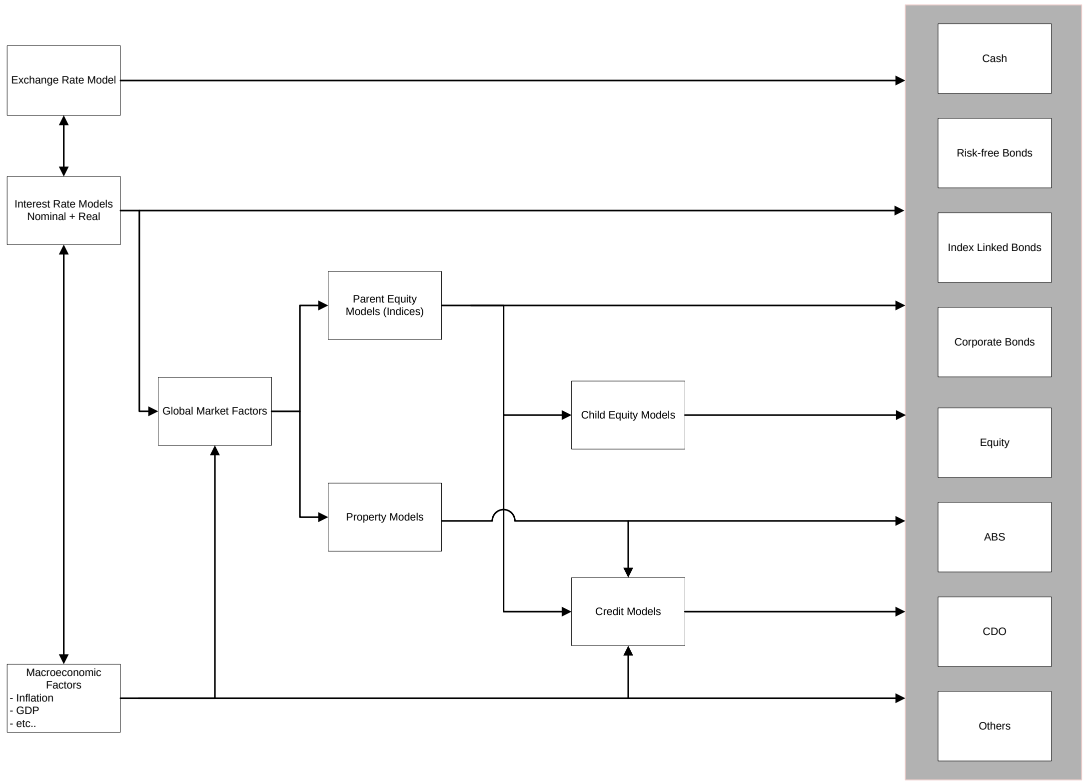

Fig. 1. Overview of the architecture of a possible fully integrated model of economic capital. In IT parlance this corresponds to a "use case" diagram showing how scenarios from interlinked stochastic risk factor models feed through to value changes at asset level.

Interest rate models. Fundamental to the architecture is a model for interest rates that will permit the valuation of future cash flows at any point in time. The literature on interest rate models is vast and the menu of options extensive, including classical models of the short rate, Heath-Jarrow-Morton style models of forward rates and LIBOR market models; see, for example, Brigo and Mercurio (2006).

The model chosen should be able to capture the dynamics of the entire yield curve and permit an easy passage from the dynamics under the so-called "real-world probability measure" to the dynamics under a "risk neutral pricing measure" which excludes the possibility of arbitrage in bond prices.

The macroeconomy. Interest rates are an essential feature of the macroeconomy, but other macroeconomic variables, such as GDP and inflation will also feature in a typical economic scenario generator. They may serve, for example, as important predictors of the credit cycle so that they will have an influence on the capital required for credit-risky assets. Inflation is also a necessary risk factor for valuing index-linked investments at future time points.

In Figure 1 we show an interaction between interest rates and other macroeconomic variables, and there are various ways this interaction can be modelled. An interesting modern framework for this integration is the so-called macro-finance framework where macroeconomic factors and latent factors are incorporated in an affine multifactor model of the dynamics of interest rates; see, for example, Spencer (2008). [1]

Exchange rate models. Projections of exchange rates are of course essential for valuing assets denominated in currencies other than the base or domestic rate. A conventional approach is based on a purchasing power parity (PPP) argument and assumes that real exchange rates are equal to, or oscillate around, a long-run equilibrium level over time, while nominal exchange rates depart from real exchange rates according to differences in inflation in the two countries. Recall that the real exchange rate is given by

$$Y_t = X_t^{\frac{\mathrm{d}}{\mathrm{f}}} \frac{P_t^{\mathrm{f}}}{P_t^{\mathrm{d}}}$$

where P_t^{f} and P_t^{d} are price levels in the foreign and domestic market, and $X_t^{\frac{\mathrm{d}}{\mathrm{f}}}$ is the nominal exchange rate expressed as the cost in domestic currency of a unit of foreign currency.

By linking exchange rates to price levels and thus inflation, the foreign exchange rate model is dependent on other features of the macroeconomy that are coupled to inflation, such as interest rates. [2]

[1] The Barrie & Hibbert ESG adopts a standard no-arbitrage approach to term structure modelling with latent factors. Macroeconomic variables do not feature explicitly in the model. Inflation is constructed as the differential between real and nominal exchange rates. For real interest rates, which may become negative, a 2-factor Vasicek model is assumed, while for nominal rates the positivity requirement leads to a choice between a 2-factor Black-Karasinski model or LIBOR market model.

[2] The ESG allows short term deviations in the real exchange rate by modelling Y_t by

Market factors. A major feature of any asset model is the methodology for modelling the performance of equity markets, including the dependencies between equity markets in various economies (UK, US, Europe, Asia, etc.) and in various industrial sectors within economies. The scale of the modelling problem means that factor modelling approaches are invariably adopted.

There are a number of approaches to factor modelling of stock market returns, although a division into three broad categories is useful Connor (1995). *Macroeconomic* factor models relate stock returns to time series of macroeconomic driving factors, such as interest rates and GDP changes. *Fundamental* factor models define driving factors associated with stock attributes, such as industry category, company size and company value; they then construct these factors as portfolio returns for portfolios of stocks sharing the attributes. *Statistical* factor models attempt to discover a number of latent factors that are explanatory of returns in a statistical sense; there are a number of techniques for extracting the factors, such as principal component analysis (PCA) and standard multivariate factor analysis.

It is generally found that fundamental and statistical factor models outperform macroeconomic factor model; Connor (1995) finds that fundamental factor models have a slight edge over statistical factor models. However, fundamental factor models tend to be appropriate for building detailed models of returns for individual stocks rather than explaining more aggregate returns for whole economies, or industry sectors within economies.

Equity returns. The level of detail of a model of equity returns may vary, but it is useful to be able to model both average returns within economies and industrial sectors (index or *parent* equity returns) as well as returns of specific shares within those economies and sectors (*child* equity returns). Index returns may feed into models of credit risk, for example, while child equity returns are used to model large holdings in individual stocks.

This can be achieved in a hierarchy of factor models: at the top level a factor model of parent equity returns can attempt to explain dependencies across economies and sectors in terms of a smaller number of drivers of the world economy; at the lower level a factor model of child equity returns can attempt to explain dependencies of stock returns within an economy/sector conditional on parent equity returns and possibly other factors. [3]

a mean-reverting process but otherwise follows the PPP approach. Since inflation is modelled as the differential of real and nominal rates in any economy, the evolution of exchange rates is thus linked to the evolution of interest rates.

[3] The ESG adopts this two-level approach. The top-level factor model of parent equity returns is a statistical factor model, where the factors are constructed by PCA. The lower-level model of child equity returns is in effect a one-factor (macroeco-

Property. Property can be modelled as an equity-like asset whose return is driven by the same kind of factors that effect worldwide equity returns, whether observable macroeconomic factors or abstract statistical factors. As well as the obvious importance of property as an asset class in its own right, the role of property valuations in the credit cycle (as evidenced by the current credit crisis) make property an important component of a complete scenario generation model for economic capital purposes. [4]

Credit risk. Value changes of credit risky assets are driven by real defaults as well as perceptions of the credit quality of non-defaulted debt, as reflected both in the ratings assigned by rating agencies and the prices of debt-related instruments such as corporate bonds and credit default swaps (CDSs).

A model of credit risk for economic capital purposes must combine a model for defaults and rating transitions with a model for prices of (or, equivalently, yield spreads on) credit-risky instruments. It needs to be calibrated to these two sources of information: empirical frequencies of default and migration events on the one hand and market prices on the other. Since the former give information about real world probabilities of default and the latter give information about market-implied (or risk-neutral) probabilities of default, we require a common modelling framework that allows easy passage between the real-world and risk-neutral probability measures. [5]

While quantification of default risk, credit quality and credit risk premia at the level of individual obligors is important, it is arguably more important to model the dependencies between obligors. Economic capital models are used to measure the risk of entire asset classes and the tail of the portfolio credit loss distribution for an entire class of credit-risky assets is mainly driven by the level of dependence between obligors (Frey and McNeil (2003)).

The dominant methodology for modelling portfolio credit risk is the Vasicek extension (Vasicek (1997)) of Merton's structural model of credit risk for an

---

nomic style) CAPM model where the parent equity return is the factor. All equity returns are modelled as excess returns over the risk-free return in the respective economies.

[4] The ESG models property returns in terms of the same set of statistical factors that drive worldwide equity returns.

[5] This, at least, has been the ideal of credit risk modelling. But credit spreads also carry a premium for liquidity in addition to a premium for pure credit risk and this serves to distort the relationship between market-implied and real-world probabilities of default. The distortion is greatest in periods of stress, such as the current credit crunch, when the appetite for credit risk is depressed and markets are illiquid. There have been a number of papers that attempt to disentangle liquidity and credit risk in spread data including Webber (2007) and Houweling et al. (2005).

individual obligor (Merton (1974)). Merton's idea that the default risk of a firm is driven by stochastic fluctuations in its asset value with respect to deterministic future liabilities is extended to a portfolio model by considering dependent asset value processes with a factor model structure. In Vasicek's approach a one-factor, Gaussian model of portfolio credit risk is derived from an assumption of correlated Brownian motions for log asset value processes. This has become an industry standard and influenced the development of the IRB approach to Basel II, Pillar 1 risk measurement as well as the valuation of structured credit, where a reduced form version of the model known as the one-factor Gauss copula model is used for CDOs.

The idea that is relevant for integrated economic capital is that of using standard linear factor models to explain the dependence between variables which represent the "distance-to-default" or "ability-to-pay" of obligors. These variables would represent the difference between asset values and liabilities in a Merton interpretation but tend to be handled more abstractly. Under Gaussian assumptions for the factor model, credit events for different obligors are conditionally independent given the factors. In principal, any number of factors can be used that are believed to affect firm valuations; they can be factors describing the macroeconomy or factors describing the performance of equity markets, which are both readily available in the economic capital model that we have sketched out. [6]

5 The Added Value of Fully Integrated Models

We have argued that fully integrated models are the conceptually superior approach to modelling dependence and diversification across an enterprise. In practical terms they also provide the framework for a number of activities that are not possible, or are more difficult, in a modular approach.

5.1 Mapping the Risk Landscape

As Figure 1 and Section 4 indicate, fully integrated models can be viewed as maps of the complex risk landscape as it effects a financial enterprise. A model permits a discussion of risk to take a more tangible form, subject to the caveat

[6] In the ESG credit dependence is modelled within economies (or possibly sectors within economies) by assuming a one-factor model where the factor is the return on the parent equity index for that economy (or sector). At enterprise-wide level, for a company exposed to multiple economies, the model can be viewed as multifactor. Portfolios of credit risky assets for different economies are dependent because the parent equity indexes are dependent.

that we are clear about the levels of uncertainty (or risk) to which the entire exercise is subject.

Model risk. Model risk is the risk of using a misspecified model to measure financial risk. In as much as all models are "wrong", model risk is always present to some degree when we use a model to draw inferences about the real world. We should distinguish between model risk at the level of the overall economic scenario generation model and model risk at the lower level of asset valuation models, treated below. Model risk at scenario generation level is "working with the wrong map", a map which is not just slightly off in its positioning of features but one which bears no useful relationship to the true landscape. The only possible approach to this risk is a process of continual refinement of the model where we "backtest" what we observe with what we would have predicted and make appropriate adjustments. This is a process without any real end where occasional seismic shifts may prompt more substantial revisions.

Parameter/calibration risk. If the scenario generation model is "correct" in its articulation of the key mechanisms of the economy but has been incorrectly calibrated to current conditions, we talk of parameter risk. It is really another form of model risk, or working with the wrong model, but it is useful to distinguish parameter risk from "wrong kind of model" risk because these relate to two distinct strands of statistical or econometric inference: parameter inference and goodness-of-fit testing respectively.

It is important not to forget parameter risk when it comes to using the model for risk quantification. Unfortunately it has become standard practice to quote risk measures like VaR as point estimates, without error bounds. But a fully integrated model does give the framework to construct errors by generating scenarios over a probability distribution of plausible parameter values.

Valuation model risk. Valuation models for many asset classes are subject to model risk (both "wrong kind of model" and parameter risk) but, if we are prepared to accept the model of the economy as plausible, the scenario generation framework gives us a test laboratory for examining how different valuation model assumptions affect pricing, and for assessing the impact of pricing differences on measures of risk and capital.

Monte Carlo error. Risk and capital numbers derived from the scenario generation model are Monte Carlo estimates of the true model quantities. For example, a 99% VaR estimate derived as the 10th worst of 1000 scenarios is an estimate of the true underlying VaR implied by the model, which cannot in general be obtained by analytical means. Accuracy of the estimate is dependent on the number of Monte Carlo replications. Where run times place limits on the number of replications, there are a number of methods of improving the efficiency of Monte Carlo by reducing the variance of the Monte Carlo estimator. These include antithetic variates and importance sampling (Robert and Casella (1999)), although the latter method is not easy to implement in a complex economic scenario generation setting. A certain amount of Monte Carlo error will be present in model outputs.

5.2 Fully Integrated Models for Stress Testing

We have argued that stochastically generated *scenarios* are the key to integrated risk models. In this section we explore the role of *stress scenarios* in fully integrated risk models; in other words we consider the relationship between fully integrated risk models and stress testing of assets.

There is a tendency to see stress testing as a distinct activity from standard risk quantification, particularly in market risk as a supplement to "VaR analysis" Kupiec (1998). However, we subscribe to the view expressed in Berkowitz (2004) that stress testing can and should be brought under the umbrella of the stochastic risk model and that stress testing that is conducted outside the stochastic risk model is of limited use.

While definitions of stress testing are hard to find, Berkowitz find that the essence of the practice lies in "choosing scenarios that are costly and rare, and then putting them to the valuation model". But rare scenarios that are chosen outside a stochastic model do not generally come with reliable assessments of likelihood attached and the mere knowledge that they could be very destructive may not be useful in formulating decisions since, as Berkowitz points out, it will always be possible to devise a scenario that causes ruin, but which lacks any plausibility when benchmarked against our understanding of the world to date.

Stress scenarios are best formalized as tail events arising from probability distributions. Consider a stochastic risk model where changes in value on the asset side (Δ_{t+1}) are related to changes in the risk factors ($\mathbf{Z}_t \rightarrow \mathbf{Z}_{t+1}$) by

$$\Delta_{t+1} = V_{t+1} - V_t = f_{t+1}(\mathbf{Z}_{t+1}, t+1) - f_t(\mathbf{Z}_t, t).$$

There are therefore two distributions to consider: the univariate distribution of Δ_{t+1} and the multivariate distribution of $\Delta Z_{t+1} = (\mathbf{Z}_{t+1} - \mathbf{Z}_t)$. It is probably more common to think of stress testing in terms of how tail scenarios for ΔZ_{t+1} play out as "effects" on the asset valuation side, and this is possible in a stochastic scenario generation set-up. However, arguably the opposite approach, of identifying tail events for Δ_{t+1} and tracing their "causes" back to the risk factor side, is more natural and more instructive. We will refer to the two approaches as "factor stress testing" and "asset-response stress testing" respectively.

Asset-response stress testing.

This form of stress testing (to adapt the words of Berkowitz above) consists of putting many randomly generated scenarios to the valuation models and identifying and understanding the scenarios that lead to the costliest effects on asset holdings. In normal operation mode, running, say, several thousands of annual risk factor paths, some of these will result in tail scenarios for Δ_{t+1}. Of 1000 scenarios the 10 that lead to the largest losses in value can be viewed as indicative of century events; the worst scenario of all is a first indication of a millennium event. Thus these scenarios come with assessments of likelihood pre-attached and address the quantification of stress testing in a way that the Basel II recommendations do not.

Having identified such stress scenarios, we should attempt to understand them, to work out the conjunctions of risk factor changes that can lead to them, and to draw inferences concerning possible mitigating risk management actions.

Obtaining sensible quantitative results is predicated on the conviction that, if we have specified the stochastic scenario generation model "correctly" and calibrated it "prudently" to data and expert beliefs, then all of the tail scenarios that should concern us are "in the model" and will emerge with appropriate frequencies over the course of stochastic simulation runs. Thus there should be no need to enrich the model with additional, and more arbitrary stress scenarios. However, the issues of model and parameter risk (as identified in the previous section) must always be considered and it will be necessary to perform sensitivity analyses to see how robust our opinions about the sources of tail risk are to changes in the model and its parameterization.

Factor stress testing.

In this approach it is usual to choose risk factor scenarios from the past which have led to financial distress or to invent future scenarios which are deemed possible. Risk factor scenarios are point scenarios that say, for instance, that interest rates rise by 200 basis points, equity markets fall by 40% and property values fall by a third.

Past scenarios for risk factor changes should appear plausible with respect to the distribution of $\Delta\mathbf{Z}_{t+1}$, or they cast doubt on the validity of its model and/or its parameterization. If events of this magnitude have occurred in the past then events of this magnitude or worse should be assigned a non-negligible likelihood of occurring again. Where the dimensionality of the risk factor vector is moderate, this likelihood can be roughly computed by generating a large number of stochastic scenarios and then counting the proportion for which each component of the risk factor vector is at least as extreme as the corresponding component of the scenario vector.

Fictive scenarios are more problematic: if experts are convinced of their plausibility then they should appear possible within the context of the distribution of $\Delta \mathbf{Z}_{t+1}$ in the same way as historical scenarios. If they are created without strong conviction and are at odds with the model then they represent the kind of valueless, frequency-less scenario that we should really discount, preferring instead to have faith in our model calibrations.

More pragmatically, even if it is against the philosophy of a fully integrated stochastic model, it is possible to incorporate a set of arbitrary point scenarios in the stochastic sampling procedure using an approach along the lines of the general framework for stress testing suggested by Berkowitz (2004). Let $\mathbf{Z}_{t+1}^{(1)}, \ldots, \mathbf{Z}_{t+1}^{(k)}$ denote a set of such point scenarios and assume that they are equally likely extreme events. With probability p draw a random point scenario from the set of k such scenarios and with probability $1-p$ generate a genuine stochastic scenario $\mathbf{Z}_{t+1}$ from the fully integrated scenario generator. This effectively assigns a frequency of k/p to the point scenario so that, for example, if we are performing an analysis of annual asset fluctuations, and we set $p = 0.01$ and $k = 10$, then each scenario is effectively treated as a 1000-year event.

5.3 Capital Allocation

Economic capital calculated for an asset portfolio or enterprise can be broken up into pieces that are attributable to sub-portfolios or business units. This process of capital allocation can be used as the basis of risk-adjusted performance comparison across sub-portfolios and there is now a considerable literature on the theory of fair allocation of capital including Tasche (1999), Denault (2001), Kalkbrener (2002).

The generic principle that is commonly adopted is known as Euler allocation and is applicable to any positive homogeneous risk measure. Suppose that the multivariate distribution of the vector $(L_1, \ldots, L_d)$, representing future losses in d sub-portfolios, is held fixed. We consider hypothetical portfolio losses of the form $L(\boldsymbol{\lambda}) = \sum_{i=1}^{d} \lambda_i L_i$ where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_d)$ is a vector of unconstrained portfolio weights. In reality our portfolio loss will be $L = L(\mathbf{1}) = \sum_{i=1}^{d} L_i$ but by considering $L(\boldsymbol{\lambda})$ for general $\boldsymbol{\lambda}$ we can examine how the risk changes as we hypothetically vary the size of each sub-portfolio.

Let ϱ be a positive homogeneous, translation preserving risk measure (like VaR or expected shortfall) and assume that overall economic capital is determined according to $\mathrm{EC}(\boldsymbol{\lambda}) = \varrho(L(\boldsymbol{\lambda})) - E(L(\boldsymbol{\lambda}))$. Then the Euler capital allocation to the ith risk will be given by

$$\mathrm{AC}_i = \frac{\partial \varrho\left(\mathrm{EC}\left(\boldsymbol{\lambda}\right)\right)}{\partial \lambda_i}|_{\boldsymbol{\lambda}=\mathbf{1}} = \frac{\partial \varrho\left(L\left(\boldsymbol{\lambda}\right)\right)}{\partial \lambda_i}|_{\boldsymbol{\lambda}=\mathbf{1}} - E(L_i) \; ;$$

we will write this as

$$\mathrm{AC}_i = \varrho(L_i \mid L) - E(L_i)$$

for simplicity. Note that the individual contributions add up to give a complete capital allocation: $\sum_{i=1}^d \mathrm{AC}_i = \varrho(L) - E(L) = \mathrm{EC}$.

Using arguments in Tasche (2000) it can be shown that under some technical assumptions on the distribution of $(L_1, \ldots, L_d)$ (fulfilled, for example, by the existence of a joint probability density) the term $\varrho(L_i \mid L)$ takes the following forms in the case of VaR and expected shortfall:

$$\begin{aligned} \mathrm{VaR}_\alpha(L_i \mid L) &= E(L_i \mid L = \mathrm{VaR}_\alpha(L)) \\ \mathrm{ES}_\alpha(L_i \mid L) &= E(L_i \mid L \geq \mathrm{VaR}_\alpha(L)) \,. \end{aligned}$$

The forms of these expressions reveal, at least in theory, how the economic capital contribution may be estimated using the Monte Carlo output from an economic scenario generator. For example, in the case of expected shortfall, we would average the losses in each subportfolio over all scenarios where the total portfolio loss exceeded the Value at risk. In practice, the problem of rare event simulation arises, and long run times may be necessary to get accurate results. But the main point is that the necessary prerequisite for computing allocations is a fully-specified joint model for $(L_1, \ldots, L_d)$ and this is delivered by a fully integrated model but not by a modular model and correlation matrix. (The modular approach with copulas would permit the calculation of allocations.)

A further development, described in Tasche (2006), is the calculation of diversification scores to give a measure of the extent of diversification in the total portfolio of an enterprise. A global diversification index can be calculated as

$$\mathrm{DI} = \frac{\varrho(L) - E(L)}{\sum_{i=1}^d \varrho(L_i) - E(L)} = \frac{\mathrm{EC}}{\sum_{i=1}^d \mathrm{EC}_i}.$$

This is simply the total economic capital for the portfolio divided by the sum of stand-alone economic capital amounts for the sub-portfolios.

A sub-portfolio diversification index can be calculated as

$$\mathrm{DI}_i = \frac{\varrho(L_i \mid L) - E(L_i)}{\varrho(L_i) - E(L_i)} = \frac{\mathrm{AC}_i}{\mathrm{EC}_i}.$$

This shows the reduction in capital that the sub-portfolio enjoys through being part of the enterprise. Where the ratio is small, this is an indication that sub-portfolio i is well-diversified with respect to the rest of the enterprise. If the global diversification is less impressive, it may be possible to gain a global improvement by increasing the size of sub-portfolio i at the expense of other sub-portfolios.

Note that the problem of how to allocate capital to single risk factors which have a non-linear impact on the value of a portfolio has also been considered by Tasche (2007) although the theory is more complicated and its implementation in the context of an economic scenario generator seems less feasible than allocation with respect to asset classes.

6 Conclusion

Current events illustrate dramatically the importance of considering extreme scenarios, and dependencies between risk factors during extreme scenarios, when assessing capital adequacy. Arguably, a satisfactory treatment of dependence and tail dependence is the most serious challenge in quantitative risk management. Isolated extreme movements may disappear in the wash of diversification, but joint extremal behaviour of important risk factors can generate the systematic risk that puts companies at risk (and the systemic risk that puts economies at risk).

We have argued against a modular approach to calculating economic capital. The simple correlation-based aggregation of capital charges for different asset classes or risk factors is flawed from both a theoretical and practical perspective. Most seriously, it provides no understanding of the sources of dependence and no real insight about how we should act to manage the risk.

We have argued in favour of a fully integrated approach to economic capital, which attempts to model directly the dependence of asset classes across an enterprise on key risk drivers, and we have suggested that the complexity of this challenge is best addressed with an economic scenario generation approach based on stochastic models and Monte Carlo simulation. A fully integrated approach can deliver a framework for sensitivity analysis, stress testing, risk-adjusted performance comparison and enterprise steering.

There is much talk at present about enterprise risk management (ERM) as a new holistic discipline for running companies. In this vision economic capital models are used to articulate risk sensitivities, to set risk controls and limits, to compute and allocate capital, to inform decision making and to facilitate the development of a true "risk culture" within an enterprise. It is our belief that only a fully integrated economic capital model can form the foundation on which the ERM project can be built.

References

Alessandri, P. and Drehmann, M. (2008). An economic capital model integrating credit and interest rate risk. Preprint.

Berkowitz, J. (2004). A coherent framework for stress-testing. In Jorion, P., editor, *Innovations in Risk Management: Seminal Papers from the Journal of Risk*. Risk Books.

Breuer, T., Jandacka, M., Rheinberger, K., and Summer, M. (2008). Inter-risk diversification effects of integrated market and credit risk analysis. preprint.

Brigo, D. and Mercurio, F. (2006). *Interest Rate Models - Theory and Practice: With Smile, Inflation and Credit*. Springer Verlag, Berlin, 2nd edition.

Burns, R. (2004). Economic capital and the assessment of capital adequacy. *FDIC Supervisory Insights*, 1(2):5–16.

CAS (2006). *The Report of the Research Working Party on Correlations and Dependencies Among All Risk Sources*. Casualty Actuary Society. CAS Forum.

CEIOPS-06 (2006). *Draft advice to the European Commission in the Framework of the Solvency II project on Pillar I issues - further advice*. CEIOPS (Committee of European Insurance and Occupational Pensions Supervisors).

Connor, G. (1995). The three types of factor models: a comparison of their explanatory power. *Financial Analysts Journal*.

Crouhy, M., Jarrow, R., and Turnbull, S. (2007). The subprime credit crisis of 07. unpublished.

Daul, S., De Giorgi, E., Lindskog, F., and McNeil, A. (2003). The grouped t-copula with an application to credit risk. *Risk*, 16(11):73–76.

Denault, M. (2001). Coherent allocation of risk capital. *Journal of Risk*, 4(1).

Embrechts, P., McNeil, A., and Straumann, D. (1999). Correlation: pitfalls and alternatives. *Risk*, 12(5):93–113.

Embrechts, P., McNeil, A., and Straumann, D. (2002). Correlation and dependency in risk management: properties and pitfalls. In Dempster, M., editor, *Risk Management: Value at Risk and Beyond*, pages 176–223. Cambridge University Press, Cambridge.

Frankland, R., Smith, A., Wilkins, T., E., V., Holtham, A., Biffis, E., Eshun, S., and Dullaway, D. (2008). Modelling extreme market events. Technical report, Institute & Faculty of Actuaries. Report of the benchmarking stochastic models working party.

Frey, R. and McNeil, A. (2003). Dependent defaults in models of portfolio credit risk. *J. Risk*, 6(1):59–92.

Houweling, P., Mentinck, A., and Vorst, T. (2005). Comparing possible proxies of corporate bond liquidity. *Journal of Banking and Finance*, 29:1331–1358.

IFRI Foundation and CRO Forum (2007). Insights from the joint IFRI/CRO Forum survey on Economic Capital practice and applications. KPMG Business Advisory Services.

Iman, R. L. and Conover, W. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics — Simulation and Computation*, 11:311–334.

Kalkbrener, M. (2002). An axiomatic approach to capital allocation. Preprint, Deutsche Bank AG, Frankfurt.

Kupiec, P. (1998). Stress testing in a Value at Risk framework. *Journal of Derivatives*, 6:7–24.

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

Merton, R. (1974). On the pricing of corporate debt: The risk structure of interest rates. *J. Finance*, 29:449–470.

Morrison, S., Kirchner, A., and Kretzschmar, G. (2008). Barrie & Hibbert Economic Scenario Generator - Technical Documentation. Technical report, Barrie & Hibbert Internal Report.

on Banking Supervision., B. C. (June 2004). International Convergence of Capital Measurement and Capital Standards. A Revised Framework. Bank of International Settlements.

Robert, C. and Casella, G. (1999). *Monte Carlo Statistical Methods*. Springer, New York.

Spencer, P. (2008). Stochastic volatility in a macro-finance model of the U.S. term structure of interest rates 1961–2004. *Journal of Money, Credit and Banking*, 40(6):1177–1215.

Sweeting, P. (2007). Modelling and managing risk. Technical report, Institute of Actuaries and Faculty of Actuaries.

Tasche, D. (1999). Risk contributions and performance measurement. Preprint, TU-Munich.

Tasche, D. (2000). Conditional expectation as a quantile derivative. Preprint, TU-Munich.

Tasche, D. (2002). Expected shortfall and beyond. *J. Banking Finance*, 26:1519–1533.

Tasche, D. (2006). measuring sectoral diversification in an asymptotic multifactor framework. *Journal of Credit Risk*, 2(3):33–55.

Tasche, D. (2007). Capital Allocation to Business Units and Sub-Portfolios: the Euler Principle.

Treasury and FSA (2006). *Solvency II: a new framework for prudential regulation of insurance in the EU*. HM Treasury and FSA. Discussion paper.

Vasicek, O. (1997). The loan loss distribution. Preprint, KMV Corporation.

Webber, L. (2007). Decomposing corporate bond spreads. *Bank of England Quarterly Bulletin*.