

# A General Approach to Integrated Risk Management with Skewed, Fat-tailed Risks<sup>\*</sup>

Joshua V. Rosenberg<sup>†</sup>

Til Schuermann<sup>†</sup>

Federal Reserve Bank of New York  
33 Liberty St.  
New York, NY 10045  
joshua.rosenberg@ny.frb.org, til.schuermann@ny.frb.org

First Draft: July 2004  
This Revision: March 2005  
This Print: March 16, 2005

**Abstract:** Integrated risk management for financial institutions requires an approach for aggregating risk types (market, credit, and operational) whose distributional shapes vary considerably. We construct the joint risk distribution for a typical large, internationally active bank using the method of copulas. This technique allows us to incorporate realistic marginal distributions, both conditional and unconditional, that capture some of the essential empirical features of these risks such as skewness and fat-tails while allowing for a rich dependence structure. We explore the impact of business mix and inter-risk correlations on total risk, whether measured by value-at-risk or expected shortfall. We then compare the copula-based method with several conventional approaches to computing risk.

**Keywords:** Market risk, credit risk, operational risk, risk diversification, copula

**JEL Codes:** G10, G20, G28, C16

---

<sup>\*</sup> We would like to thank Franklin Allen, Kabir Dutta, Darryll Hendricks, Beverly Hirtle, Philippe Jorion, Ugur Koyluoglu, Glenn Meyers, Andrew Patton, Marc Saidenberg, Bill Schwert (the editor), Daniel Straumann, Stuart Turnbull, Jiang Wang, an anonymous referee, participants at the Federal Reserve Conference on Financial Structure and Regulation (New Orleans, October 2004), and seminar participants at the Federal Reserve Bank of New York and Rutgers University, for their helpful comments, and Megan Todd for her able research assistance.

<sup>†</sup> Any views expressed represent those of the authors only and not necessarily those of the Federal Reserve Bank of New York or the Federal Reserve System.

## 1. Introduction

The goal of integrated risk management in a financial institution is to measure and manage risk and capital across a diverse range of activities in the banking, securities and insurance sectors.<sup>1</sup> This requires an approach for aggregating different risk types and hence risk distributions, a problem found in many applications in finance including risk management and portfolio choice.

At the core of all financial institutions, no matter which sector they operate in, lie three risk types: market, credit, and operational risk.<sup>2</sup> The distributional shapes of each risk type vary considerably. For market risk, we typically see portfolio value distributions that are nearly symmetric and often approximated as normal. Credit and especially operational risk generate more skewed distributions because of occasional, extreme losses. These losses might be due to large lending exposures in the case of credit risk, or large catastrophes such as 9/11, in the case of operational risk.<sup>3</sup>

Some risks types, such as market risk, are more easily characterized and measured than others, such as operational risk, but much less is known about the relationship between the risks. We develop an approach to combine marginal distributions in an internally consistent and realistic manner while preserving important properties of the individual risks, like skewness and

---

<sup>1</sup> Virtually all of the large, multinational financial institutions operating around the world today operate in at least two of the three sectors, which make them financial conglomerates by the definition of the Joint Forum (Joint Forum, 2001, p.5). They define a financial conglomerate as “any group of companies under common control whose exclusive or predominant activities consists of providing significant services in at least two different financial sectors (banking, securities, insurance).”

<sup>2</sup> Credit risk is the dominant risk in a commercial bank and received a formal regulatory capital charge with the first Basel Capital Accord (BCBS, 1988). Market risk capital was added with the 1996 market risk amendment (BCBS, 1996). Under the anticipated New Basel Accord (BCBS, 2004), operational risk is added so that each of these risks carries a regulatory capital charge.

<sup>3</sup> See Jorion (2001) for examples of market risk and Crouhy, Galai and Mark (2001) for examples of all three risk types. Kuritzkes, Schuermann and Weiner (2003) present stylized pictures of a very broad range of risk types faced by a financial conglomerate.

fat-tails. At the same time, our methodology can be implemented when there is limited information about inter-risk dependencies, as in the case in which only correlations are available.

In this paper, we directly construct the joint risk distribution for a typical, large, internationally active bank — one engaged in commercial banking and securities (underwriting) activities — using the method of copulas. Market and credit risk distributions are known to be time varying, reflecting the effects of stochastic volatility. With this in mind, we estimate dynamic models for these risks using a multivariate GARCH approach, and then combine the resulting conditional distributions with operational risk which is taken to be unconditional. Indeed, the risk dynamics will naturally carry through to the joint or total risk distribution via the copula.

Our empirical analysis uses a combination of data from regulatory reports, market data and vendor data. Banking holds particular interest at the moment because of the intense policy debate surrounding the New Basel Capital Accord (BCBS, 2004), and the incorporation into regulation of a new risk type (operational) which has substantially different characteristics from market or credit risk.<sup>4</sup> Although our application is to a bank, it could just as easily be extended to an insurer or, indeed, to a financial conglomerate that combines all three (and not just two) sectors under one roof.

Our study is the first to use publicly available, industry-wide data to perform a comprehensive analysis surrounding the set of factors or dimensions that affect total risk. Specifically, we examine the sensitivity of risk estimates to business mix, dependence structure, risk measure, and estimation method. Few papers conduct such sensitivity analysis — those that do focus on just one or two of these dimensions, and the risk densities themselves are not

---

<sup>4</sup> That banks and insurers are actively wrestling with this issue is shown quite clearly in Joint Forum (2003).

estimated from empirical data. Those studies that do model risk based on empirical data do not conduct sensitivity analysis. Moreover, the results of these studies are difficult to generalize, since their data is often based on the experience of a single institution. By using a panel of quarterly data for a set of large banks, we can have more confidence that our results are indeed representative of a typical institution. In addition, by exploring the impact of business mix, we can comfortably span a wide range of business activities and hence bank types.

We find that given a risk type, total risk is more sensitive to differences in business mix or risk weights than to differences in inter-risk correlations. For a policy maker or practitioner this is good news as it is far harder to estimate inter-risk correlations than to assess business mix. We show a complex relationship between volatility and fat-tails in determining total risk: depending on the setting, they either offset or reinforce each other. For example, as operational exposure is increased relative to market and credit exposure, total risk, whether measured using value-at-risk (VaR) or expected shortfall (ES), first declines significantly (operational risk has lower volatility than the other two) but then flattens as the impact of fatter tails offsets the effect of lower volatility. As correlation of market and credit risk with operational risk is increased, both volatility and fat-tails increase. The choice of copula (normal versus Student-t), which determines the level of tail dependence, has a more modest effect on total risk.

We go on to compare the copula-based method with several conventional approaches to computing risk. These are examples of typical industry practice, and in one case, a calculation related to required regulatory capital. Understanding the accuracy of these easy to implement approximations to the more complicated copula-based approach should therefore have broad appeal. A first approximation of adding up the risks (e.g. VaR) of the marginal distributions significantly overestimates total risk or economic capital, which is not surprising as it assumes

perfect inter-risk correlation. Assuming joint normality of the risk factors imposes tails that are thinner than the empirical estimates and significantly underestimates economic capital. A third “hybrid” approach, which combines marginal risks using a formula that would apply to an elliptical distribution, is surprisingly accurate.

The rest of the paper proceeds as follows. In Section 2, we present a brief overview of the related literature. We then discuss risk measurement (VaR, ES) and risk aggregation in Section 3. Most of the discussion is in terms of VaR, simply because it has become so common, but later we go on to conduct robustness checks using ES and find that all results implied by VaR hold. Section 4 provides an overview of copulas, while Section 5 focuses on the marginal risk distributions by risk type and lays out our general approach. In Section 6, we present our analytical results by examining the impact of business mix, correlations and copula type on the joint risk distribution. Section 7 provides some final comments.

## **2. Related literature**

With the rise of risk management as a distinct discipline in banking and finance, the issue of risk aggregation has only recently become an area of study. Using copulas to this end seems like a natural application. Embrechts, McNeil, and Straumann (1999, 2002) were among the first to introduce this toolkit to the finance literature. Li (2000) provides an application to credit risk and credit derivatives. Frey and McNeil (2001) emphasize the importance of tail dependence and, by introducing copulas, generalize dependence beyond correlation.<sup>5</sup> Other applications of

---

<sup>5</sup> Poon, Rockinger and Tawn (2004) use multivariate extreme value theory instead of copulas to model tail dependence. Their technique is data-intensive and requires empirical observations of joint tail events. Other applications of copulas include Rosenberg (2003), who uses a copula related methodology for multivariate contingent claims pricing. Patton (2002) uses copulas to model exchange rate dependence while Fermanian and Scaillet (2003) analyze copula estimation and testing methods.

copulas for portfolio risk measurement include Bouyé (2001), Longin and Solnik (2001), Glasserman, Heidelberger, and Shahabuddin (2002), and Embrechts, Lindskog, and McNeil (2003).

Studies that focus more narrowly on cross-risk type aggregation for a financial institution are less common. Alexander and Pérezier (2003) use a common risk factor approach to characterize the joint distribution of market and credit risk for a bank. Wang (1998) lays out a series of theoretical arguments and approaches for measuring and modeling enterprise-wide risk for an insurer facing a highly diverse set of marginal risk distributions. Ward and Lee (2002) use a normal copula to aggregate a diverse set of risks, some computed analytically (e.g. credit risk is assumed to follow a beta distribution), some by simulation (e.g. mortality risk for life insurance), to arrive at the total distribution for a diversified insurer.

Dimakos and Aas (2004) estimate the joint loss distribution for a Norwegian bank that also owns a life insurance subsidiary. This study, as well as Ward and Lee (2002), approaches the problem of risk aggregation by considering risks pair-wise. Ward and Lee (2002) use pair-wise roll-ups with a Gaussian copula, while Dimakos and Aas (2004) decompose the joint risk distribution into a set of conditional probabilities and impose sufficient conditional independence so that only pair-wise dependence remains. The total risk is then just the sum of the conditional marginals (plus the unconditional credit risk which serves as their anchor). Their simulations indicate that total risk measured using near tails (95% to 99%) is about 10-12% less than the sum of the individual risks. Using the far tail (99.97%), they find that total risk is often overestimated by more than 20% using the additive method. These results suggest that incorporation of diversification effects can be crucial for accurate risk aggregation, particularly in the tails.

Finally, Kuritzkes, Schuermann and Weiner (2003), hereafter KSW, make a simplifying assumption of joint normality, allowing for a closed-form solution, and use a broad set of parameters to arrive at a range of risk aggregation and diversification results for a financial conglomerate. They find somewhat smaller, but still significant differences between total risk and the sum of individual risks. These differences are about 15% across market, credit and operational risk for a bank, 20% to 25% for insurers, and 5% to 15% for a banc-assurance style financial conglomerate.

### **3. Risk measurement and VaR**

#### *3.1. Risk measurement*

Risk is simply the potential for deviation from expected results, particularly adverse deviation. Behind every risky future cash flow, earnings result, or change in value there lies a probability distribution of potential results. The relative magnitude of risk could be defined by a measure of spread or dispersion in that distribution such as the standard deviation or variance. However, variance is not necessarily sufficient for capturing risk — two distributions with dramatically different shapes and differing amounts of downside risk can have the same variance.

Measures such as skewness and kurtosis can be used to quantify the risk that is not adequately described by variance alone. Another approach is to examine the percentiles of the distribution to provide answers like “99% of the time interest rates will move less than X% in one day and 1% of the time the move will be greater.” This effectively summarizes Value-at-Risk (VaR), a concept we take up in more detail below.

### 3.2. Value-at-Risk

Value-at-risk (VaR) has become a standard for measuring and assessing risk in financial institutions. VaR is broadly defined as a quantile of the distribution of returns (or losses) of the portfolio in question. Some find it convenient to consider the negative of this quantile so that higher values of VaR correspond to higher levels of risk. We take VaR simply to be equal to this quantile so that larger (negative) values correspond to higher levels of risk.

More formally, let  $Y_t$  be the portfolio value at time  $t$ , and define the  $k$ -period ahead portfolio return as  $r_{t+k} = \ln(Y_{t+k}) - \ln(Y_t)$ . We denote the  $(1-\alpha)\%$  VaR estimate at time  $t$  for a  $k$ -period ahead return as  $\text{VaR}_{t+k}(\alpha)$ , so that

$$\Pr(r_{t+k} < \text{VaR}_{t+k}(\alpha)) = \alpha. \quad (3.1)$$

Much as the concept of a sufficient statistic provides a compact representation of the characteristics of the data, so VaR is hoped to give us a similarly compact sufficient risk measure. Christoffersen and Diebold (2000) and Berkowitz (2001) argue that rather than focus on just one number such as VaR, risk managers and, implicitly, regulators should focus on the whole density function of returns, perhaps using techniques such as those laid out in Diebold, Gunther and Tay (1998) and Berkowitz (2001). Nonetheless, interest in a simpler summary measure continues.

Artzner, Delbaen, Eber, and Heath (1997, 1999) lay out a set of criteria necessary for what they call a “coherent” measure of risk. They include homogeneity (larger positions bring greater risk), monotonicity (if a portfolio has systematically lower returns than another for all states of the world, its risk must be greater), sub-additivity (the risk of the sum cannot be greater than the sum of the risks) and the risk-free condition (as the proportion of the portfolio invested in the risk-free asset increases, portfolio risk should decline). Importantly, unless the underlying



risk factors come from an elliptical distribution (the normal belongs to this family), VaR does not satisfy the sub-additivity condition. Thus, a firm could concentrate all of its tail risks in one exposure in such a way that the risk borne by that exposure appears just beyond the overall portfolio VaR threshold (Embrechts, McNeil and Straumann, 1999, 2002).

A related statistic, expected shortfall (ES), also sometimes referred to as conditional VaR, is a coherent risk measure that estimates the mean of the beyond-VaR tail region. Specifically, using Eq. (3.1), ES at time  $t$  over horizon  $k$  at confidence level  $\alpha$ , is defined as

$$ES_{t+k}(\alpha) = E[r_{t+k} | r_{t+k} \leq \text{VaR}_{t+k}(\alpha)]. \quad (3.2)$$

Similar to our reporting convention for VaR, we take ES to be a tail expectation of losses, so that ES is also negative. Our empirical results in Section 6.2 are largely discussed using VaR, but we go on to conduct robustness checks using ES as well.

### 3.3. VaR for portfolios

One of the original approaches for portfolio risk measurement is Markowitz's (1959) mean-variance analysis. Consider a simple case of three assets (or more broadly, three cash-flow generating processes) with returns  $r_x$ ,  $r_y$  and  $r_z$  and weights  $w_x$ ,  $w_y$  and  $w_z$  such that these weights sum to one. The portfolio return is simply  $r_p = w_x r_x + w_y r_y + w_z r_z$  and the variance is

$$\sigma_p^2 = w_x^2 \sigma_x^2 + w_y^2 \sigma_y^2 + w_z^2 \sigma_z^2 + 2w_x w_y \sigma_{x,y} + 2w_x w_z \sigma_{x,z} + 2w_y w_z \sigma_{y,z}, \quad (3.3)$$

where  $\sigma_i^2$  is the variance of the  $i^{\text{th}}$  return and  $\sigma_{i,j}$  is the covariance between return  $i$  and  $j$ .

Then, assuming for simplicity that the portfolio expected return is zero, the portfolio VaR can be written as (see, for instance, Bradley and Taqqu, 2003),

$$\text{VaR}_p(\alpha) = \sigma_p F_p^{-1}(\alpha). \quad (3.4)$$

where  $F_p^{-1}(\alpha)$  is the  $\alpha^{\text{th}}$  quantile of the standardized portfolio return ( $r_p/\sigma_p$ ).  $F_p^{-1}(\alpha)$  is the standardized quantile function, which measures risk in terms of the number of standard deviations from zero. Note that  $F_p^{-1}(\alpha)$  is a negative number (e.g. -3.1 for 0.1<sup>st</sup> percentile of the standard normal distribution), and  $\sigma_p$  is positive; thus,  $\text{VaR}_p(\alpha)$  is also negative.

By substituting the portfolio volatility Eq. (3.3) into the VaR Eq. (3.4), we have

$$\begin{aligned} \text{VaR}_p(\alpha) &= \sqrt{w_x^2\sigma_x^2 + w_y^2\sigma_y^2 + w_z^2\sigma_z^2 + 2w_xw_y\sigma_{x,y} + 2w_xw_z\sigma_{x,z} + 2w_yw_z\sigma_{y,z}} [F_p^{-1}(\alpha)] \\ &= \sqrt{w_x^2\sigma_x^2 [F_p^{-1}(\alpha)]^2 + w_y^2\sigma_y^2 [F_p^{-1}(\alpha)]^2 + w_z^2\sigma_z^2 [F_p^{-1}(\alpha)]^2} \text{sign}[F_p^{-1}(\alpha)] \\ &\quad + \sqrt{+2w_xw_y\sigma_{x,y} [F_p^{-1}(\alpha)]^2 + \dots} \end{aligned} \quad (3.5)$$

Eq. (3.5) tells us that the VaR of the portfolio can be written in terms of the second moments of the marginal returns and the inverse distribution function of the standardized portfolio returns.

We can think of the quantile  $F_p^{-1}(\alpha)$  as a providing a scaling factor for each volatility.

Now, suppose the quantiles of the individual standardized returns are the same as for the portfolio returns, i.e.  $F_p^{-1}(\alpha) = F_x^{-1}(\alpha) = F_y^{-1}(\alpha) = F_z^{-1}(\alpha)$ . For example, this is true for the multivariate normal distribution. Given this equality, and under the assumption that  $\text{sign}[F_p^{-1}(\alpha)] = -1$  meaning we are looking at losses, we write portfolio VaR, which we denote

H-VaR (H for hybrid), as

$$\begin{aligned} \text{H-VaR}_p(\alpha) &= -\sqrt{w_x^2\sigma_x^2 [F_x^{-1}(\alpha)]^2 + w_y^2\sigma_y^2 [F_y^{-1}(\alpha)]^2 + w_z^2\sigma_z^2 [F_z^{-1}(\alpha)]^2} \\ &\quad + \sqrt{+2w_xw_y\sigma_{x,y} [F_x^{-1}(\alpha)][F_y^{-1}(\alpha)] + \dots} \\ &= -\sqrt{w_x^2 [\text{VaR}_x(\alpha)]^2 + w_y^2 [\text{VaR}_y(\alpha)]^2 + w_z^2 [\text{VaR}_z(\alpha)]^2} \\ &\quad + \sqrt{+2w_xw_y\rho_{x,y} [(\text{VaR}_x(\alpha))(\text{VaR}_y(\alpha))] + \dots} \end{aligned} \quad (3.6)$$

Eq. (3.6) says that portfolio VaR can be computed using the same formula as portfolio volatility, where each volatility is replaced by the corresponding VaR. If we calculate H-VaR when the

marginals come from different density families, then some volatilities could be overweighted and others underweighted relative to the actual VaR. The net effect will depend on the relationship between the marginal quantiles, the volatilities, and the portfolio quantile. Notice, however, that H-VaR does allow the tail shape of the marginals to affect the portfolio VaR estimate.

Eq. (3.6) simplifies considerably when the individual returns are uncorrelated, i.e.  $\sigma_{i,j} = 0 \forall i, j$ . Then,

$$\text{VaR}_p(\alpha) = -\sqrt{w_x^2 [\text{VaR}_x(\alpha)]^2 + w_y^2 [\text{VaR}_y(\alpha)]^2 + w_z^2 [\text{VaR}_z(\alpha)]^2}.$$

When the risks are perfectly correlated ( $\rho_{i,j} = 1 \forall i, j$ ), Eq. (3.6) becomes (for positive weights):

$$\text{Add-VaR}_p(\alpha) = w_x \text{VaR}_x(\alpha) + w_y \text{VaR}_y(\alpha) + w_z \text{VaR}_z(\alpha). \quad (3.7)$$

We will refer to this as additive VaR or Add-VaR. When risk correlations are less than one, we would expect Add-VaR to overestimate risk. Like H-VaR, Add-VaR allows the tail shape of the marginals to affect the portfolio VaR estimate.<sup>6</sup>

Another special case of Eq. (3.5) is obtained by assuming that the risk distribution is multivariate normal, which makes it also a special case of Eq. (3.6). Then, each of the marginals is normal, and Normal VaR (N-VaR) has standardized quantiles given by the inverse standard normal distribution function ( $\Phi^{-1}(\alpha)$ ),

$$\text{N-VaR}_p(\alpha) = -\sqrt{w_x^2 \sigma_x^2 \Phi^{-1}(\alpha)^2 + w_y^2 \sigma_y^2 \Phi^{-1}(\alpha)^2 + w_z^2 \sigma_z^2 \Phi^{-1}(\alpha)^2 + 2w_x w_y \sigma_{x,y} \Phi^{-1}(\alpha)^2 + \dots}. \quad (3.8)$$

It is clear that N-VaR will be accurate only when the joint risk distribution is multivariate normal. N-VaR is most likely to fail when one or more marginals exhibit significant negative

---

<sup>6</sup> Note that Add-VaR assumes an elliptical underlying distribution and perfect correlation. If the actual distributions are not members of this family, then Add-VaR might not provide an upper bound for VaR. This would be an example of the failure of sub-additivity.

skewness or excess kurtosis. In that case, the normal quantile underestimates the actual marginal quantile.

Clearly, the closeness of each of the VaR approximations (hybrid, additive, and normal) to the actual VaR will depend on the validity of their underlying assumptions. H-VaR has the least restrictive assumptions, since it permits the correlations and marginal quantiles to depend on the data. Add-VaR restricts the correlations to be unity, but does not restrict the quantiles. N-VaR forces the quantiles to come from a normal distribution but allows the correlations to be estimated from the data.

Each of these approximations relies on the assumption that the quantiles of the portfolio are the same as the quantiles of the marginals. When  $F_p^{-1}(\alpha) \neq F_x^{-1}(\alpha) \neq F_y^{-1}(\alpha) \neq F_z^{-1}(\alpha)$ , we are left with the problem of computing Eq. (3.5) when  $F_p^{-1}$  is unknown. To obtain the correct portfolio VaR, we therefore need to obtain the joint return distribution of the portfolio. Copulas allow us to solve this problem by combining the specified marginal distributions with a dependence function to create this joint distribution. The joint distribution can then be used to calculate the quantiles of the portfolio return distribution, since the portfolio returns are weighted averages of the individual returns.

#### **4. Copulas**

Copulas provide important theoretical insights and practical applications in multivariate modeling; an excellent reference is Nelsen (1999). The essential idea of the copula approach is that a joint distribution can be factored into the marginals and a dependence function called a copula. The term copula is based on the notion of “coupling;” the copula couples the marginal distributions together to form a joint distribution. The dependence relationship is entirely

determined by the copula, while scaling and shape (e.g., mean, standard deviation, skewness, and kurtosis) are entirely determined by the marginals.

Using a copula, marginal risks that are initially estimated separately can then be combined in a joint risk distribution that preserves the original characteristics of the marginals. This is sometimes referred to as obtaining a joint density with pre-determined marginals. For example, when individual risk distributions are estimated using heterogeneous dynamic models that cannot be easily combined into a single dynamic model (e.g. different explanatory variables, different variable frequencies, or different model types), the marginals can be thought of as pre-determined. The copula approach is then naturally applied to combine these time-varying marginal risk distributions to obtain a time-varying joint risk distribution.

In addition, there are cases in which marginal risks are not estimated using time-series data (e.g., implied density estimation, survey data, or combination of frequency and severity data). There is then no direct way to create a multivariate dynamic model that incorporates all of the risk types. Once again, the copula method can incorporate these marginal risks into a joint risk distribution.

The copula approach is also useful when “off-the-shelf” multivariate densities inadequately characterize the joint risk distribution. In the risk management context, the multivariate normal distribution is known to poorly fit the skewed, fat-tailed properties of market, credit, and operational risk. Using a copula, parametric or non-parametric marginals with quite different tail shapes can be combined into a joint risk distribution. Joint risk distributions created using copulas can also span a range of dependence types beyond correlation such as tail dependence.

#### 4.1. Representation of a joint distribution using a copula

One of the fundamental results concerning copulas is known as Sklar's Theorem. It states that any joint distribution can be written in terms of a copula and marginal distribution functions. This representation shows that it is possible to separately specify each variable's marginal distribution and the dependence relationship that links these marginals into a joint distribution. The standard representation for the joint distribution  $F_{xy}$  is:

$$F_{x,y}(x, y) = C(F_x(x), F_y(y)), \quad (4.1)$$

where  $C(u, v)$  is the copula, and  $F_x$  as well as  $F_y$  are the cumulative marginal distribution functions.

Alternatively there is the density representation:

$$f_{x,y}(x, y) = f_x(x)f_y(y)c(F_x(x), F_y(y)), \quad (4.2)$$

where  $c(u, v) = \left( \frac{\partial^2}{\partial u \partial v} \right) C(u, v)$  is the copula density, and  $f_x$  as well as  $f_y$  are the marginal density functions.

In Eq. (4.1), the marginal information is embedded in  $F_x(x)$  and  $F_y(y)$ , and the dependence information is captured by  $C(u, v)$ . The copula itself is a joint distribution with uniform marginals, so it maps points on the unit square ( $u, v \in [0, 1] \times [0, 1]$ ) to values between zero and one. The copula relates the quantiles of the two distributions rather than the original variables, so that the copula for two random variables is unaffected by a monotonically increasing transformation of the variables. For example, the copula of the joint distribution of  $X$  and  $Y$  is the same as the copula of the joint distribution of  $\ln(X)$  and  $\exp(Y)$ .

#### 4.2. Constructing a joint distribution using a copula

The copula for any multivariate distribution function can be obtained using *the method of inversion*. This technique factors out the effects of the marginal distributions on the dependence relation by substituting the arguments of the original joint distribution with the marginal quantile functions. Using this method, the copula  $C(u, v)$  of the joint distribution  $F_{x,y}$  is:

$$C(u, v) = F_{x,y} \left( F_x^{-1}(u), F_y^{-1}(v) \right), \quad (4.3)$$

where  $F_x^{-1}$  and  $F_y^{-1}$  are the marginal quantile functions, and  $u$  as well as  $v$  are probabilities.

For example, we can determine the normal copula using the method of inversion. If  $\Phi(x, y; \rho)$  is a bivariate standard normal cumulative distribution function and  $\Phi^{-1}(u)$  and  $\Phi^{-1}(v)$  are standard normal quantile functions, then the normal copula is simply  $\Phi(\Phi^{-1}(u), \Phi^{-1}(v); \rho)$ .

The bivariate normal copula has a single parameter, the correlation coefficient  $\rho$ .

A joint distribution with given marginals and a given copula can be created by plugging-in the marginal distributions into the copula function. To obtain the joint distribution  $F_{a,b}$  using the copula  $C(u, v)$  and the marginals  $F_a$  and  $F_b$ , we have

$$F_{a,b}(a, b) = C(F_a(a), F_b(b)). \quad (4.4)$$

So, to create a bivariate distribution with given marginals and a *normal* copula, one simply inserts the marginal distributions into the normal copula as  $\Phi(F_x^{-1}(u), F_y^{-1}(v); \rho)$ . It is worth noting that the correlation between the risks is not necessarily equal to  $\rho$ . The non-normality of the marginal densities drives a wedge between the  $\rho$  and the Pearson correlation coefficient. In addition, it is possible that some correlations cannot be attained for certain choices of marginals.

### 4.3. Simulation using copula-based multivariate densities

Techniques for simulating realizations from the multivariate normal distribution are well-known. Probably, the most common approach is to draw uncorrelated univariate normal vectors with desired means and standard deviations. These are transformed into correlated draws from a multivariate normal distribution using a Cholesky factorization of the correlation matrix.

A related approach can be used to simulate realizations from a multivariate distribution written in terms of a copula and marginal distributions. For clarity, we illustrate how to simulate random variables  $A$  and  $B$  from a joint distribution generated using a bivariate normal copula and given marginals ( $F_x$  and  $F_y$ ). Notice that  $F_x$  and  $F_y$  can be defined either parametrically or non-parametrically.

First, two uncorrelated standard normal vectors are generated and transformed to correlated standard normal vectors using the lower triangular Cholesky factor.<sup>7</sup> Then, we have  $X, Y \sim \Phi(x, y; \rho)$ .  $X$  and  $Y$  are transformed into realizations from a normal copula by applying the normal cumulative distribution function to each vector [ $U = \Phi(X)$ ,  $V = \Phi(Y)$ ] so that  $U, V \sim \Phi(\Phi^{-1}(u), \Phi^{-1}(v); \rho)$ . Finally,  $A$  and  $B$  are obtained by applying the inverse cumulative distribution function for each marginal to  $U$  and  $V$ , [ $A = F_x^{-1}(U)$ ,  $B = F_y^{-1}(V)$ ].

---

<sup>7</sup> If the two uncorrelated standard normal random variables are  $Z_1$  and  $Z_2$ , then  $X = Z_1$  and  $Y = \rho Z_1 + \sqrt{(1 - \rho^2)} Z_2$ . This method is easily extended to simulate Student-t variates with  $\nu$  degrees of freedom. The Student-t realization is obtained by multiplying each normal pair  $(X_i, Y_i)$  by  $\sqrt{\nu/s_i}$  where  $s_i$  is drawn from a chi-square distribution with  $\nu$  degrees of freedom. Then, the inverse cumulative distribution function for a Student-t( $\nu$ ) is applied instead of the inverse normal to create a Student-t copula realization.



## 5. Marginal risk distributions: market, credit, and operational

In this section, we consider the problem of modeling the marginal risk distributions for a typical large, internationally active bank. This is the kind of bank regulators have in mind when designing capital regulation (BCBS, 2004). As noted earlier, we will focus on three risk-types: market, credit, and operational. Market risk measures the risk of adverse movements in market factors, such as asset prices, foreign exchange rates or interest rates. The risk of loss resulting from failure of obligors to honor their payments is called credit risk, while operational risk can be defined as (BCBS, 2004, §644) “the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events.” This list is hardly exhaustive; for instance, mortality and morbidity risk is typically incurred by life insurers but rarely by banks. But any financial institution is subject to these three risks and hence the focus of our study.

While all banks are subject to minimum capital standards due to credit risk, only banks with a *significant* market risk exposure are required to calculate a risk-based capital ratio that takes into account market risk in addition to credit risk. U.S. regulators deem market risk exposure to be significant if the gross sum of trading assets and liabilities on the bank’s balance sheet exceeds 10 percent of total assets or \$1 billion (USGAO, 1998, p. 121). As reported in Hirtle (2003; Table 1), at the end of 2001 there were 19 bank holding companies (BHCs) that were subject to market risk capital standards. This set of institutions will form the sample of banks for our analysis. As a benchmark, we consider the risk and business activity profile of the median bank from that group for which we have complete data; there are 17 of those.

### 5.1. The bank sample

We have quarterly data for 17 BHCs from 1994Q1 to 2002Q4, obtained from Y-9C regulatory reports, expressed on a pro-forma basis going back in time to account for mergers.

See Table 1 for details. For example, if a bank in 2002 is the result of a merger in 2000, pre-2000 data is merged on a pro-forma basis. We use this data to arrive at typical size and business mix characteristics. Thus for the last year of the sample period the average (median) BHC has \$277bn (\$178bn) in assets, 48% (53%) of which was devoted to lending. Even though only these banks were required to report market risk, the market risk share of the minimum capital requirement was rather small on average (median): 0.18% (0.14%).<sup>8</sup> The average (median) return on assets was 1.07% (1.19%). Because the size distribution of banks is quite skewed – there are a few very large banks – we will use medians to characterize the typical bank.

## *5.2. General approach*

In developing a modeling framework for the joint risk distribution of a bank, we face a number of challenges. First, there is the range of business mix. The “typical” bank we have in mind is engaged in a range of business activities, some of which are more intensive in market risk (like trading), credit risk (like lending), or operational risk (like custody).

A second important consideration is the units in which the risk type is measured and reported. To arrive at an aggregate risk distribution, we need a common currency of risk. Market risk is typically based on the return distribution of the end-of-day positions in the trading book. Similarly, credit risk can be described by a loss or return distribution, typically in terms of percentage exposure. In both cases we need to know the exposure at risk, e.g. trading or lending assets, to compute dollars at risk.

---

<sup>8</sup> This small amount of market risk regulatory capital is not necessarily inconsistent with industry benchmarks of economic capital arising from internal models; those are closer to 20%. Only position risk, i.e. risk arising from trading activities, is measured for purposes of assigning regulatory capital. Interest rate risk due to ALM is typically included in internal models.

Finally, operational risk is not yet part of the regulatory framework – but will be under the New Basel Accord – so typically banks do not currently report operational risk. Data for this risk type is only recently becoming available, and we will rely on external sources to characterize, and calibrate, our operational risk distribution. The measurement unit is in dollars. We convert the operational risk losses into a “return” by normalizing with respect to total assets, since all assets and activities of the bank are in some way subject to this risk type.

The overall or total risk distribution will simply be a weighted combination of the individual risks, where the weights, which add up to one, are determined by the risk-specific exposures. We determine exposures and risk weights as follows. Market risk exposure is equal to trading assets, credit risk exposure is equal to lending assets, and operational risk exposure is equal to total assets. We call the sum of these three exposures the “total book,” and each risk weight is simply the ratio of the exposure to the total book. Based on BHC characteristics at the end of 2002, the risk weights for our benchmark bank are 3.1% market, 29.1% credit, and 67.8% operational. To be sure, these weights could be somewhat arbitrary because they only include on-balance sheet items.<sup>9</sup> Although off-balance sheet items can be larger, this is the usual treatment. We do not consider these weights as benchmarks per se but rather as a useful starting point from which to conduct our business mix experiments. After all, we are interested in understanding the range of total risk outcomes as that business mix changes.

Our *ex ante* weights are quite different from the *ex post* risk contribution from business lines discussed in other studies.<sup>10</sup> For example, KSW report that the average *ex post*

---

<sup>9</sup> This was kindly pointed out by a referee.

<sup>10</sup> The difference arises primarily because the *ex post* contribution depends on both the exposure and the riskiness of the business line, while the *ex ante* weight is only a function of the exposure. For instance, all business activities are exposed to some operational risk (high *ex ante* weight), but this need not result in a large attribution to operational risk capital (*ex post* contribution) if operational activities have relatively low risk. A useful analogy is the example

decomposition based on bank internal models attributes about 20% of economic capital to market risk, 55% to credit risk, and 25% to operational risk. When we perform an ex post risk decomposition using our ex ante benchmark risk weights (Section 6.2), we obtain generally similar results.

For market and credit risk, we model the return due to each risk-type as a function of observable risk factors. This approach allows us to separately measure factor sensitivities and model risk factor dynamics. We can then estimate marginal risk distributions by combining the time-varying risk factor distributions with the factor sensitivities. For operational risk, it is far from clear which observable risk factors might drive operational risk.<sup>11</sup> Thus, we work with the unconditional distribution, the details of which are provided below in Section 5.6.

We proceed by letting bank  $i$  returns from risk type  $j = \{\text{market, credit}\}$  at time  $t$  be driven by a set of observable risk factors  $\mathbf{x}_{j,t}$  as in

$$r_{i,j,t} = \alpha_{i,j} + \boldsymbol{\beta}'_{i,j} \mathbf{x}_{j,t} + \varepsilon_{i,j,t}. \quad (5.1)$$

In Eq. (5.1), bank returns attributable to risk type  $j$  are said to be captured by the set of observable risk factors collected in  $\mathbf{x}_{j,t}$ . We estimate factor sensitivities using panel data, so these are perhaps best interpreted as time-averaged if actual sensitivities are time-varying. The risk factors are themselves subject to complex dynamics which are specified and estimated separately.

Naturally, our chosen factors ( $\mathbf{x}_{j,t}$ ) will not necessarily capture all of the relevant risks. But as long as the omitted factors which appear in  $\varepsilon_{i,j,t}$  are not correlated with observable factors,

---

of a portfolio invested in a low risk and a high risk asset, e.g. U.S. treasuries and the S&P500. One could place a majority of the portfolio in the former (high ex ante weight) but still have the overall portfolio risk be driven largely by the latter (high ex post risk contribution).

<sup>11</sup> An example of this approach for market and credit (but not operational) risk is Alexander and P ezier (2003). See also Jorion (2002) for the importance of market exposures of trading revenues on reported market risk.

we suffer only loss of efficiency but no bias in estimating the bank-specific set of factor sensitivities  $\beta'_{i,j}$ . To be sure, some residual market and credit risk might be captured by our operational risk measure, but a significant amount of the unexplained variation is likely to be bank specific. We analyze residual risk in Section 5.4.

An advantage of proceeding along the lines of Eq. (5.1) is that the time series available for bank data,  $r_{i,j,t}$ , is much shorter and at a lower frequency (quarterly from 1994-2002) than what is available for the risk factors  $\mathbf{x}_{j,t}$  (daily from 1974-2002 for the market risk factors and monthly from 1988-2002 for the credit risk factors). Using the factor approach, we can construct the marginal risk distributions from a much richer history of risk factor movements than what we have seen at a quarterly frequency in the last nine years.

The income from market risk related activities is captured in the regulatory reports as “Trading Revenue.” While fluctuations from this income source include the outcome of position taking, arguably some of this revenue comes from customer-related activities such as fees and commissions. Regulatory reporting, however, is focused on sensitivity of the trading book to risk factors such as equity returns, interest rates and foreign exchange rates. We therefore direct our analysis by regressing trading returns, defined as trading revenue over trading assets, on returns and volatilities of risk factors from a broad set of asset classes in the capital markets: equities, currencies, and interest rates.

Most credit risk arises from lending, which generates fees and interest income. The former is not separately reported, so we will proxy total credit risk related income by net interest income less a charge for incurred losses as proxied by provisions. To define a credit return, we divide this net credit income by the lending assets. In thinking about modeling the variation in this return, credit spreads come immediately to mind as they should yield the appropriate

filtering over a wide set of factors, such as default risk, that we would expect to drive the profits and losses of a bank's loan portfolio.

Having defined the dependent variable in Eq. (5.1) for market and credit risk, we now go on to specify the factor dynamics.

### 5.3. Specification of risk factor dynamics

It is well known that asset returns exhibit stochastic volatility, which can create skewness and fat-tailedness in the resulting risk distributions. To capture this phenomenon, we estimate the joint distribution of the factor risks within a multivariate GARCH framework. In particular, we estimate the asymmetric BEKK model proposed by Kroner and Ng (1998) that allows for different volatility impacts of positive and negative shocks. Asymmetric volatility effects can be important to capture, for example, the negative skewness in equity returns or the positive skewness in changes in credit spreads. This model can be written as

$$\begin{aligned} \mathbf{r}_t &= \boldsymbol{\mu} + \boldsymbol{\alpha}\mathbf{r}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t | \Xi_{t-1} \sim N(\mathbf{0}, \mathbf{H}_t) \\ \mathbf{H}_t &= \boldsymbol{\Omega} + \mathbf{A}'\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}'\mathbf{A} + \mathbf{B}'\mathbf{H}_{t-1}\mathbf{B} + \mathbf{G}'\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}_{t-1}'\mathbf{G}. \end{aligned} \quad (5.2)$$

The first equation models the expected return dynamics of the  $n \times 1$  vector of risk factors ( $\mathbf{r}_t$ ), with a vector of constants ( $\boldsymbol{\mu}$ ) and a first-order autoregressive component with parameter ( $\boldsymbol{\alpha}$ ). The second equation models the volatility dynamics of the factor innovations  $\boldsymbol{\varepsilon}_t$ , conditional on the information set  $\Xi_t$ , where  $\mathbf{H}_t$  is the  $n \times n$  conditional covariance matrix of the  $n$  factors.

Conditional variance depends on a constant parameter matrix ( $\boldsymbol{\Omega}$ ), scaled lagged squared innovations ( $\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}'$ ), the lagged covariance matrix ( $\mathbf{H}_{t-1}$ ), and a matrix of asymmetric terms ( $\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}_{t-1}'$ ) where  $\eta_{it} = \text{Max}[0, -\varepsilon_{it}]$  or  $\text{Max}[0, \varepsilon_{it}]$ .  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{G}$  are  $n \times n$  lower triangular parameter matrices that provide the weights on the lagged squared innovations, lagged covariances, and asymmetric terms. Estimation is performed using maximum likelihood under the assumption of

a Gaussian error density. This procedure has a quasi-maximum likelihood interpretation (Bollerslev and Wooldrige, 1992) and will produce consistent parameter estimates under other error distributions.

One of the well-known difficulties associated with multivariate GARCH model estimation is the large number of parameters to be estimated. Therefore, we make use of a more parsimonious version of the asymmetric BEKK model in which the **A**, **B**, and **G** matrices are constrained to be diagonal. To further reduce the number of parameters required for the market risk model, we implement variance targeting (Engle and Mezrich, 1996). This procedure sets the long-run covariance from the model equal to the sample covariance matrix by imposing the restriction that  $\mathbf{\Omega} = (\mathbf{I} - \mathbf{A} - \mathbf{B})\mathbf{S}$ , where **S** is the sample covariance matrix.

Our market risk model is based on three daily risk factors from 1/1/1974 to 12/31/2002: equity returns, currency returns, and interest rate changes.<sup>12</sup> For credit risk, we would like broad coverage of the credit spectrum, given that not all bond credit ratings are equally liquid. Thus, we use monthly changes in double-A (AA) and triple-B (BBB) credit spreads, for which we have data coverage from December 1988 to the end of 2002.<sup>13</sup>

We find evidence of stochastic volatility in the market and credit risk factors. As shown in Table 2, squared market factors are predictable as evidenced by Ljung-Box statistics significant at the 1% level. Squared credit factors are also predictable with Ljung-Box test p-values of 3% and 10%. We test each factor for an asymmetric volatility effect using a univariate asymmetric GARCH specification. We find that the term  $\text{Max}[0, -\epsilon_{it}]$  is significant for equities,

---

<sup>12</sup> The equity return is the log of the S&P500 return from Datastream. The interest rate measure is the log-difference of the 10-year Constant Maturity Treasury rate from the Board of Governors H.15 release. The currency return is the log difference of the trade-weighted currency index from the Board of Governors G.5 release.

and  $\text{Max}[0, \epsilon_{it}]$  is significant for both credit spreads. These terms are included in estimation of the final model.

Overall, the estimated multivariate GARCH model fits the data well. When we apply Ljung-Box tests to squared standardized model residuals, we find insignificant p-values at the 10% level for all but one risk factor. In Fig. 1, we graph the conditional volatility and conditional correlations of the risk factors. We see that factor volatility is quite variable over the sample period. There are notable volatility spikes for equities in October 1997 (Black Monday, first panel) and for credit spreads in September 1998 (Russian ruble devaluation and Long Term Capital Management recapitalization, third panel). Conditional correlations for equities versus interest rates range from -0.66 to 0.64 (second panel) while conditional correlations for double-A (AA) versus triple-B (BBB) spreads range from 0.59 to 0.99 (fourth panel).

#### *5.4. Market and credit risk*

With the risk factor dynamics in hand, we go on to estimate market and credit factor sensitivities (5.1) over the sample period 1994Q1 to 2002Q4. We test both the factors and the factor volatilities as explanatory variables in our market and credit factor sensitivity regressions. The volatilities used in the regression are quarterly conditional volatilities from the estimated multivariate GARCH models. In the credit regressions, we find that the credit spread volatilities are not statistically significant additions to the credit spreads, so they are omitted from the final model.

Our aim is to characterize the factor sensitivities of a typical BHC. Although sensitivities will vary across banks and over time, we do expect to find certain key commonalities. For

---

<sup>13</sup> We use monthly changes in AA and BBB credit spreads from Bloomberg. The spreads are constructed using the yield-to-maturity from the Merrill Lynch AA and BBB-rated corporate bond indices subtracted from duration-



example, on average, trading returns are likely to be positively related with the equity index, and credit returns are likely to be negatively related to credit spreads. For certain factors, it could be the case that some banks have large positive sensitivities and others have large negative sensitivities. In that case, our average estimate would be close to zero which would not be representative of any of the banks.<sup>14</sup>

We used median asset shares to characterize a typical large BHC for our benchmark risk weights. However, to obtain representative factor sensitivities, the median approach is not as appealing. Using the factor sensitivities of the median size bank would make inefficient use of the 612 bank-quarters we have in our pooled sample. Moreover, using the median sensitivity from each of the 17 bank-specific regressions would not account for correlations between the parameters. Therefore, this method might misestimate the characteristics of the typical bank.

In light of these issues, we thought it most reasonable to follow a pooled OLS regression approach. In Table 3, we present estimates from this regression including bank fixed effects. We also present  $\bar{R}^2$  for the pooled regressions with fixed effects alone and in combination with the factors, as well as the average  $\bar{R}^2$  from the 17 bank-level regressions.

For the bank-level regressions, the market risk factors explain on average about 10% of the variation of returns (as measured by average  $\bar{R}^2$ ) attributable to market risk related activities. In the pooled regression, the addition of the risk factors increases the  $\bar{R}^2$  from 58.3% to 59.7%. We find that the signs of the pooled regression coefficients are broadly in line with our expectations. Specifically, trading returns are positively associated with equity returns as might be expected when equity trading positions are generally net long. Trading returns are also

---

matched Treasury yields from the Federal Reserve Board H.15 release.

positively related to equity volatility. This result could be due to a positive correlation between commission-based trading revenue, trading volume, and volatility. Trading returns are negatively associated with currency volatility. This would be consistent with bank exposure to volatility risk from writing foreign exchange options.

The interest rate and interest rate volatility factors are not statistically significant in the pooled regression, but the signs of the coefficients are as expected. A decline in short-term interest rates has a positive effect on returns, which would be consistent with slight positive duration. And, interest rate volatility is negatively related to trading returns, which could reflect exposure from writing interest rate options or holding long positions in assets with negative convexity such as mortgages.

Turning now to the credit risk regression, we see in Table 4 that the two credit risk factors explain on average about 5% of the variation in credit risk earnings (as measured by average  $\bar{R}^2$ ) across the 17 bank-level regressions. Adding the risk factors to the pooled regression with fixed effects increases the  $\bar{R}^2$  from 56.2% to 56.8%. As spreads on high credit quality (double-A) firms widen, returns attributable to credit activities increase. The opposite is true for spread changes for lower quality (triple-B) firms. This seems reasonable as bank lending portfolios are relatively more concentrated toward the middle and lower end of the credit spectrum. Thus, when the credit quality of those firms deteriorates, returns attributable to bank lending activities should decline.

One reason why the lending portfolio of banks is skewed in this way is due to differential borrower access to capital markets as an alternative source of funding. Large corporates, which tend to have better credit ratings, are able to go to the capital markets instead of banks for

---

<sup>14</sup> We conduct several additional tests and do not find evidence for bimodality of the factor sensitivities from the

funding. This alternative is less available to smaller, lower credit quality firms. As the cost of funding increases, i.e. the spread on high credit quality bonds widens, bank loans become more attractive as a capital source.

In addition, changes in credit spreads need not just reflect credit risk. Huang and Huang (2004) and others have shown that the proportion of the credit spread attributable to default (i.e. to pure credit risk) is relatively low (20-30%) for high quality bonds but higher (60-80%) for speculative grade bonds. As a result, it is not surprising that a widening of double-A spreads actually increases credit related returns for banks.

We are now in a position to generate the market and credit risk distributions by combining the estimated factor sensitivities with simulated annual realizations of the factors based on the estimated multivariate asymmetric GARCH model. In our application, starting values for volatilities are set to their unconditional levels to reflect average market conditions. However, by setting the initial volatility levels to their current values, our model could also be used to estimate time-varying risk densities that depend on current market conditions.

For the market risk factors, which are modeled at a daily frequency, we first simulate a path of 252 daily log-returns and variances. Since the credit risk factors are estimated at a monthly frequency, we simulate a path of 12 log-returns and variances. The sum of the 252 daily or 12 monthly log-returns is equal to a single annual return, and the sum of the daily or monthly variances is equal to a single annual variance. This approach matches the market and credit risk horizons to the same annual horizon as operational risk. We repeat this procedure 200,000 times to create the simulated joint distribution of the annual factors. Each of the

---

bank level regressions.

200,000 simulated annual market or credit returns is just a linear combination of the simulated annual factor realizations multiplied by the factor sensitivities.

### *5.5. Market and credit residual risk*

From the relatively small size of average  $\bar{R}^2$  in the factor sensitivity regressions, it is clear that much of the variation in accounting returns attributable to either market or credit risk related activities is not captured by our set of dynamic risk factors. This should hardly be surprising as bank specific risk often plays an important role in total bank risk. For example, a single large obligor default can have a significant bottom line impact for an individual bank that is not common across the entire industry. Moreover, operating risk and business risk are absorbed into our estimated residual.<sup>15</sup> Operating risk includes operational risk as defined by the New Basel Accord, while business risk includes variation in earnings due to factors such as fluctuations in demand, technology shocks, and changes in the competitive landscape.

In Table 5, we present summaries of the regression residuals for the market and credit risk regressions. We report cross-sectional (across 17 regressions) averages and medians of moments two through four of the residuals. The simulated residuals are discussed below in Section 5.7. The distribution of residuals is asymmetric and fat-tailed for both risk types, though the skew is positive for market and negative for credit risk.

It is also of interest to analyze the impact of these residuals on the total risk distribution. To conduct such an exercise, we need to obtain a representative residual distribution for market

---

<sup>15</sup> Kuritzkes (2002), for instance, contrasts “event” risk, which conforms more closely to the notion of operational risk used here, with “business” risk, which is thought to capture (p. 3) “residual non-financial earnings volatility not attributable to internal or external events.” He finds that event plus business risk make up about 2.5% of total risk weighted assets (or about 30% of the BIS total capital requirement) for these institutions, of which about 58% is residual “business” risk. If we take these figures at face value, it would imply that we are missing about 15-18% of “total” risk.

and credit risk which can then be incorporated into the overall analysis. To avoid creating excess kurtosis by combining densities with different volatilities, we first standardize each of the residuals by the respective bank-risk type residual volatility, i.e. the root mean square error of each bank-risk type regression. We then pool the standardized residuals across all banks for each risk type and proceed to scale them by the median RMSE of each risk type. In this way, the higher moments of the residuals are preserved but at a scale which is representative of this group of large banks. This gives us a sample of 612 residuals for each of the two risk types that can be used to define each non-parametric residual risk distribution.

### *5.6. Operational risk*

For our empirical work, we adopt the definition of operational risk set forth by the Basel Committee. Namely, these are losses due to failure of internal processes, people, systems and external events. For our analysis of operational risk, we use results from de Fournouville, DeJesus-Rueff, Jordan, and Rosengren (2003), hereafter FJJR. They argue that the vendor-provided operational loss databases suffer from a reporting bias that is correlated with size. Small losses are under-reported while large losses are “too big to hide.” Specifically, if we denote  $X$  as a vector of operational losses, then for some reporting threshold  $\tau$ , where  $\tau = \$1$  million in FJJR,  $x = \log(X-\tau)$  is modeled as an exponential density,  $f(x) = \frac{1}{b} e^{-\frac{x}{b}}$ . FJJR incorporate a size-bias correction and estimate the exponential parameter  $b$  to be 0.64 using the OpRisk Analytics database and 0.66 using the OpVantage database. We take the average of these to obtain an exponential parameter of 0.65.

FJJR also report that large, internationally active banks typically experience between 50 and 80 losses exceeding \$1 million per year, the minimum loss value in the datasets. We take the mid-point of their range to obtain a daily probability of an operational loss event of  $65/365$ .<sup>16</sup>

We construct the annual operational loss distribution for our typical bank as follows. To create a single annual loss realization, we first draw 365 Bernoulli trials with  $p = 65/365$ . This reflects the possibility of an operational loss each day. On each day that an operational loss occurs, we draw the log of the dollar loss amount from an exponential distribution with parameter of 0.65. We sum the exponentials of the daily log-losses to create the annual dollar loss. The final operational risk distribution is generated using 200,000 loss years.<sup>17</sup>

### *5.7. Marginal risk distributions*

While each marginal risk is simulated from a parametric model, the introduction of nonlinearities such as GARCH and associated temporal aggregation result in annual risk distributions that might be quite difficult to fit parametrically. This motivates our more robust nonparametric approach for inverse cumulative distribution function estimation. These functions are needed as inputs to the copula to generate the total risk distribution.

For each risk type, we define the empirical inverse cumulative distribution function by sorting the returns and associating with each return its rank in the dataset (divided by 200,000). We then fit a cubic spline to this function. The spline estimate is our inverse cumulative distribution for each risk type.

---

<sup>16</sup> To be sure, we cannot directly verify that these “internationally active banks” represent the median bank in our sample. But at a minimum, this group will likely have regulatory capital assessed for credit and market risk (under the existing capital rules) and thus should overlap substantially with our sample of 17 banks. Note also that FJJR do not find the tail thickness parameter  $b$  to vary with bank size.

<sup>17</sup> If  $\log(X)$  is exponentially distributed, then  $X$  has a Pareto distribution (e.g. Johnson, Kotz, and Balakrishnan, 1994, p. 576). Using the exponential parameter of 0.65, only the first moment of the loss distribution exists. In our

The characteristics of the simulated marginal risk distributions for 200,000 simulated bank-years are presented numerically in Table 6 (columns two through four) and visually in Fig. 2. In terms of assets at risk, market risk has highest volatility ( $\sigma_{mk} = 0.58\%$ ) and thinnest tails ( $\kappa_{mk} = 3.7$ ), while operational risk has lowest volatility ( $\sigma_{op} = 0.04\%$ ) and fattest tails ( $\kappa_{op} = 35.3$ ). Credit risk is in-between with  $\sigma_{cr} = 0.19\%$  and  $\kappa_{cr} = 16.1$ . The market risk distribution is nearly symmetric while the credit risk distribution is moderately left-skewed at  $-1.3$  and operational risk is more significantly skewed at  $-4.5$ .

The operational risk distribution is noticeably different from the other two. Volatility is relatively low, but both skewness and especially kurtosis are quite high. So operational exposure typically has a small impact on risk, but there are occasional extreme losses. This finding is consistent with rare, but enormous operational losses observed at financial institutions such as Allied Irish, Daiwa, and Barings (Jorion, 2001).

At this stage, it is instructive to examine the moments of the simulated market and credit residual risk distributions that are given in Table 5 in the rows labeled “Simulated.” By construction, the volatility of those residuals is the same as the empirical median. As expected, the higher moments of the simulated residuals are similar to those of the actual residuals. Market risk residuals have the positive skew and fat-tails of their systematic counterpart, and credit risk residuals are similarly fat-tailed but negatively skewed. In the risk aggregation analysis (Section 6.3.3), we use these simulated residuals when including residual risk into the total risk distribution.

---

simulations, we set a log-loss greater than 1,000 standard deviations equal to a loss of 1,000 standard deviations. This guarantees the existence of all moments without affecting the tail shape in the region of interest.

## 6. Aggregating the risks

To aggregate the three risk types, we need to know their marginal distributions, the relative weight of each distribution, assign inter-risk correlations, and then specify a copula. We consider the normal copula to be a useful benchmark, since the multivariate normal distribution is commonly used in risk management, an obvious example being RiskMetrics (JP Morgan, 1995). A reasonable first generalization beyond the multivariate normal is to combine more realistic marginals into a joint distribution using a normal copula. The normal copula is also commonly seen in the literature. See, for instance Li (2000) and Ward and Lee (2002).

More recently, the Student-t copula has become prominent (see, for instance, Frey and McNeil, 2001, and Glasserman, Heidelberger, and Shahabuddin, 2002) because it can capture “tail dependence,” which is controlled by the degrees of freedom parameter. In contrast, a normal copula has tail independence. We report risk aggregation results for Student-t copulas with 5 and 10 degrees of freedom in Section 6.3.

### *6.1. Marginal risk-type weights and cross-risk type correlations*

In Section 5.2, we determined benchmark ex ante weights for our large, internationally active bank to be 3.1% for market risk, 29.1% for credit risk, and 67.8% for operational risk. Changing these weights is meant to reflect different business mixes. For example, most regional banks focus their business activities on lending and have little trading activity. As a result they are likely to have relatively more credit risk and less market risk than the base case would suggest. In addition, they could have less exposure to traditional sources of operational risk such as processing or custody related activities. A “processing bank,” conversely, would likely have more operational risk than the average bank. Finally, one would expect a trading-intensive bank to have relatively more market risk.



KSW, in their risk aggregation exercise, examine a broad range of correlations gleaned from several academic and industry studies. They summarize these correlations in their Table 4, which we reproduce here and use in our own analysis as Table 7. The first set is taken from Dimakos and Aas (2004). The second set comes from Ward and Lee (2002) who have separate correlations for market and ALM risk; the table entry is the higher of the two correlations. The third set is cited in Joint Forum (2001, p. 25), and is based on interview findings and reflect the values used by Dutch and other financial conglomerates in their internal capital models. For our benchmark institution, we take a midpoint of those values. The resulting benchmark correlation between market and credit is 50%, and the benchmark correlation of each of those two risk types with operational risk is 20%.<sup>18</sup>

## *6.2. Results*

The total risk distribution using the benchmark weightings and correlations with a normal copula is displayed graphically in Fig. 3 and numerically in the last column of Table 6. The volatility is 0.08% of total book (defined as the sum of trading assets, lending assets and total assets) or 0.11% of total assets. Of course, this distribution inherits features from the marginals. The negative skewness present in the credit and operational marginals is carried over into the overall skewness (-1.1). The kurtosis of total risk at 9.6 is between market at 3.7 and credit at 35.3.

We can now use the characteristics of the benchmark total risk distribution to measure ex post contributions of each risk type. These can be compared with the ex ante business mix

---

<sup>18</sup> When the marginal risks are normal, the correlations for risks combined using a normal copula are equal to the copula parameter values. Otherwise, there will be some difference between the actual correlation and the 'rho values' that are used for the copula. For our analysis across business mix with fixed correlations, the market versus

weights used to estimate VaR.<sup>19</sup> We obtain ex post risk contributions at the benchmark weights of 8.5% for market, 53.1% for credit, and 38.4% for operational risk. As expected, these values differ significantly from the ex ante business mix weights, which were 3.1% for market, 29.1% for credit, and 67.8% for operational risk. However, our ex post contributions are generally comparable to industry economic capital benchmarks; recall these are 20%, 55% and 25% for market, credit and operational risk respectively.

KSW note that the bulk of the economic capital attributed to market risk using bank internal models arises from asset-liability management (ALM), which is not subject to capital regulation and not incorporated in our weights. This difference helps to explain our smaller market risk attribution. Operational risk is difficult to measure precisely, and KSW report large differences in operational capital share allocations across industry studies. Therefore, it is not surprising that our estimate differs somewhat from the industry benchmark.

We start our discussion by focusing on the copula-based results across the two sets of experiments: varying business mix (Fig. 4 and Table 8) and inter-risk correlations (Fig. 5 and Table 9). In the figures, the line labeled Copula uses a normal copula with the benchmark inter-risk correlations, and the next three are approximations that are discussed in detail in Section 6.2.3. Following the legend, the first is the total risk using a normal copula. The second is the hybrid approach, which is based on an elliptical density formula for portfolio VaR and defined in Eq. (3.6). The third is “Additive,” the VaR obtained by simply adding up the VaRs of the

---

credit correlation at  $\rho = 0.50$  is equal to 0.47 and the market/credit versus operational correlation at  $\rho = 0.20$  is 0.16. See also Table 10.

<sup>19</sup> We choose a relatively straightforward approach to risk attribution by applying the H-VaR formula given in Eq. (3.6). This approach takes into account the business mix weights, the marginal VaRs or levels of economic capital, and correlations. More involved methods could be applied, but that would go beyond the scope of our paper; see for instance, Gouriéroux, Laurent and Scaillet (2000) and Koyluoglu and Stoker (2002) for discussion on this topic.

marginals. Finally, the line labeled “Normal” is the VaR under joint normality implying, of course, that the marginal distributions are also normal.

### *6.2.1. Impact of business mix*

In Fig. 4, we see the impact of a change in business mix on the 99.9% VaR (0.1<sup>st</sup> percentile) for total risk distribution at the benchmark correlations. This percentile is commonly used in industry and conforms to the tolerance level of the New Basel Accord. Table 8 reports weights for the three risk types, standardized quantiles, volatilities, Copula-VaRs, and the approximation error for N-VaR and Add-VaR relative to Copula-VaR. Some intermediate values (rows) are omitted for space reasons. The standardized quantile (VaR/volatility) provides a measure of left tail thickness. A normal distribution has a standardized 0.1% quantile of -3.1, so larger negative values indicate a relatively fat left tail.

The top panel in Fig. 4 displays the effects of varying market versus credit weight while holding the operational weight constant at its benchmark level of 67.8%. Thus, the leftmost set of points corresponds to 0% market and 32.2% credit risk, the next is 3.2% market and 29.0% credit, and at the other extreme we have 32.2% market and 0% credit. Naturally, the extreme points are rather unrealistic, but in-between we can view these weights as reasonably reflecting different business mixes from, for example, little to more intensive trading activity.

Initially, with 0% market and 32.2% credit weight, the volatility of the total risk distribution is 0.071% and the 0.1% standardized quantile is -6.15 standard deviations (Table 8, Panel A, first row). At first, VaR declines as market exposure increases. The initial decrease in left tail thickness (market has thinner tails than credit) outpaces the volatility increase (market has higher volatility than credit). At a market risk share larger than 30%, corresponding to a market weight of around 9%, the volatility increase dominates giving the VaR curve a parabolic

shape. VaR ranges from low of -0.434% to a high of -0.630% with the lowest risk achieved near the benchmark exposure weights, corresponding to the italicized row in Table 8, Panel A.

The bottom panel of Fig. 4 and Panel B of Table 8 show the impact of a *change* in operational risk share, while holding the relative proportion of the market and credit exposure constant. In the bottom panel, 0% on the horizontal axis represents the benchmark weighting, 3.1% market, 29.1% credit and 67.8% operational. A shift of +10% from market and credit towards operational would result in weights, in the same order, of 2.8%, 26.2% and 71.0%, and so on (see also Panel B in Table 8). Since it is conceptually not reasonable to assign a weight of zero to operational risk, we adopt this approach based on deviations from the benchmark weighting.

As weights are shifted from market and credit risk, there is an increase in tail thickness (Table 8, Panel B, where the italicized row corresponds to the benchmark). The left tail gets thicker as operational exposure rises. The standardized quantile at baseline weights is  $-5.71$  standard deviations, and it increases to  $-9.09$  standard deviations for a bank with all operational risk. At the same time, as the operational weight increases, the volatility of total risk declines (from 0.105% to 0.041%). The net effect is that the VaR curve first declines as volatility decreases more quickly than tail-thickness increases. At operational weights greater than 50% above the benchmark, the curve flattens reflecting the offsetting effects of decreased volatility and fatter tails.

If we decrease the weight on operational risk and shift it proportionately to market and credit risk, i.e. we move to the left in Fig. 4, overall volatility and risk (as measured by 99.9% Copula-VaR) increase. The standardized quantile at first declines (in absolute value) from  $-5.71$  to  $-5.66$ , only to rise again to about the benchmark level past a 20% shift in the weights from

operational to market and credit risk. The much larger volatility of both market and credit risk dominate the reduction in kurtosis from the operational risk distribution. To be sure, credit risk is itself also quite leptokurtic ( $\kappa_{cr} = 16.1$ ).

### 6.2.2. *Impact of inter-risk correlation*

In Fig. 5, we explore the impact of inter-risk correlation on 99.9% VaR (the 0.1<sup>st</sup> percentile) of the total risk distribution at the benchmark risk weights. Its companion table (Table 9) shows standardized quantiles, volatilities, Copula-VaRs and the approximation error for N-VaR and Add-VaR relative to Copula-VaR. We italicize the row corresponding to the benchmark correlations. The top panels of Fig. 5 and Table 9 display the impact of changing correlation between market and credit risk while correlation of both risk types with operational risk remains constant at the benchmark level of 20%. The bottom panels examine the impact of changing correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%.<sup>20</sup>

As expected, higher levels of market and credit risk correlation lead to higher levels of risk as measured by VaR (Fig. 4, top panel, and Table 9, Panel A).<sup>21</sup> VaR increases from -0.402% at a correlation of zero to -0.438% for a correlation of 0.9. We see that higher market and credit correlation is associated with higher volatility and a slightly thinner left tail. Overall, however, the impact of changing the inter-risk correlation between market and credit risk has a rather small impact on total risk.

---

<sup>20</sup> Of course, one can not construct a correlation matrix from arbitrarily assigned correlations. The Cholesky factorization used for simulation requires the correlation matrix to be invertible. Thus, for some experiments we will not be able to set  $\rho_{i,j}$  for any risk type  $i$  or  $j$  to be arbitrarily high.

<sup>21</sup> Embrechts, McNeil and Straumann (2002) give examples of increased  $\rho$  leading to lower VaR.

Looking now to the bottom panel in Fig. 5 and Panel B in Table 9, VaR increases as we raise the correlation between operational risk and market/credit risk. From initial level of -0.401% at zero correlation, VaR increases to -0.582% at a correlation of 0.8. In this case, higher correlation results in higher volatility and a fatter left tail. Since both of these effects are in the same direction, total risk is more sensitive to the level of operational risk correlation with other risk types than the level of market versus credit correlation.

Deviation from normality in the marginal distributions drives a wedge between  $\rho$  and the correlation coefficient. This is illustrated in Table 10, where we display the relation between the two from simulated returns using a normal copula. The first column corresponds to the top panel of Fig. 5 while the second column corresponds to the bottom panel.

While the market and credit correlation (column 1) is always quite close to the specified  $\rho$  (within 0.03), the operational versus market/credit correlation is somewhat lower (column 2). As  $\rho$  varies from 0.10 to 0.80, market and credit correlation moves from 0.09 to 0.77 but operational correlation moves from 0.08 to 0.67. Even though operational risk correlation has a smaller range for these experiments, VaR is more sensitive to changes in operational rho than market versus credit rho. In our calculation of VaR approximations (including N-VaR), we use the empirical correlation matrix, rather than the  $\rho$  matrix, so that comparison across VaR estimates are based on equal correlations.

### *6.2.3. Approximations to Copula-VaR*

In comparing the three VaR approximations to the copula method (Fig. 4 to Fig. 7, Table 8 and Table 9), we see several patterns. First, Add-VaR always provides the largest VaR estimate, and N-VaR always provides the smallest estimate. Both are significantly biased relative to Copula-VaR. Add-VaR overestimates risk, since it fixes the correlation matrix at

unity, when in fact the empirical correlations are much lower. N-VaR underestimates risk since it uses the lowest standardized quantiles for the marginals, i.e.  $\Phi^{-1}(0.001) = -3.09$  versus actual quantiles of -3.12 (market), -6.29 (credit), and -9.09 (operational). The hybrid approach (H-VaR) tracks Copula-VaR well, but it is also upwardly biased and thus conservative.

In Table 8 and Table 9, we contrast the approximate VaRs with Copula-VaR by computing the approximation error, which is  $(\text{approximation VaR} - \text{Copula-VaR}) / \text{Copula-VaR}$ . In the business mix experiment (Fig. 4 and Table 8), Add-VaR is monotonically decreasing as we shift to a lower volatility risk-type. If we assume the joint distribution is normal (N-VaR), then total risk increases modestly (Fig. 4, top panel, and Table 8, Panel A). Indeed, N-VaR initially underestimates VaR by about 50%, but comes within 5% as the share of market vs. credit risk goes to 100%. Note that, of the three risk types, the market risk distribution is closest to normal. H-VaR is within 14% (no market risk exposure) to 7% (no credit exposure) and has the same parabolic shape as the copula measurement. Add-VaR overestimates Copula-VaR by between 33% (no credit exposure) and 60% (market and credit risk weights of 12.88% and 19.31%).

In the bottom panel of Fig. 4 and Panel B of Table 8, we see that as the business mix shifts entirely towards operational risk, copula, H-VaR and Add-VaR converge. Copula-VaR and H-VaR pick up the tradeoff between lower volatility and heavier tails, so they both flatten out. N-VaR shrinks steadily as it is driven entirely by volatility (which is lowest for operational risk) and ignores the very heavy tails inherited from the operational risk distribution.

Moving on to the correlation experiments (Fig. 5 and Table 9), we find that H-VaR tracks Copula-VaR fairly well, although H-VaR is always more conservative. The maximum deviation of H-VaR from Copula-VaR is 17% across all correlation experiments with an average

approximation error of 14% for market/credit and 9% for operational correlations. The average approximation error for N-VaR is -45% when varying market versus credit correlations and -48% when varying the operational correlation. This reflects the overly thin left tail of the normal distribution compared to the actual risk distributions. N-VaR also changes linearly with rho, while Copula-VaR and H-VaR reflect non-linear interactions of the tails with the correlation.

For the correlation experiments, we also compute a “diversification benefit,” which is the percentage decrease in VaR using the copula approach versus Add-VaR:  $(\text{Add-VaR} - \text{Copula-VaR}) / \text{Add-VaR}$ . This is of interest because it provides a benchmark for the degree of conservatism that comes with the additive VaR rule. This conservatism might be thought of as providing a buffer against parameter and model uncertainty. Given the benchmark business mix weights, we find diversification benefits of 33% to 39% when varying market versus credit correlations and 11% to 39% when varying the operational correlation. These are generally higher than the 15% diversification benefit reported by KSW, although KSW use the 99% VaR level and assume joint normality. Our results are closer to the 20+% reported by Dimakos and Aas (2004), who also use a normal copula and look at the far tail (99.97% in their case).

Taken together these results imply that the interaction between measures of correlation and the higher moments of the marginals is quite subtle and complicated. Add-VaR systematically overestimates total risk while N-VaR underestimates total risk. H-VaR provides a fairly close approximation.



### 6.3. Robustness tests

#### 6.3.1. Risk measure

Since VaR is not a coherent risk measure, we repeat the set of experiments using the expected shortfall (ES) measure as defined in Eq. (3.2). The general conclusions are no different from our previous findings using VaR. By way of illustration, we present in Fig. 6 two experiments focusing on the risk sensitivity to operational exposure and operational correlation.

The top panel of Fig. 6 shows the impact on 99.9% ES of a *change* in operational risk weight, while holding the relative proportion of the market and credit weight constant. This is analogous to the VaR results displayed the bottom panel in Fig. 4. Similarly, the correlation experiment in the bottom panel of Fig. 6 examines the impact on 99.9% ES of changes in correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%. This figure is analogous to the bottom panel of Fig. 5.

While expected shortfall is always greater than VaR (as guaranteed by the definition), their sensitivities to changes in business mix and correlation are very similar. For example, in the bottom panel of Fig. 6, copula ES follows the same pattern as copula VaR (bottom panel of Fig. 5), but is consistently about 5% to 10% larger. Remarkably, even for expected shortfall, the hybrid ES comes quite close to the copula ES (on average within 13%, nearly the same as for VaR), albeit always with a positive bias.

#### 6.3.2. Copula choice

To measure VaR sensitivity to the choice of copula, we repeat the experiments forming joint distributions using Student-t copulas to allow for tail dependence, which is determined by

the degrees of freedom parameter ( $\nu$ ). Choices for strong tail dependence in the literature range from  $\nu$  of 3 to 7 (Glasserman, Heidelberger and Shahabuddin, 2002) or 3 to 8 (Demarta and McNeil, 2004), so we take a mid-point of  $\nu = 5$ . As a case between strong tail dependence ( $\nu = 5$ ) and tail independence (normal copula,  $\nu = \infty$ ), we use  $\nu = 10$ . See also Embrechts, McNeil and Straumann (2002).

The results are presented in Fig. 7, where the top panel shows the business mix experiment with operational risk, and the bottom panel shows the inter-risk correlation experiment, analogous to the previous figure on expected shortfall. Two features stand out in both experiments. First, there is relatively little difference in risk across copula choice, although as expected, measured risk is higher when allowing for tail dependence. In the top panel where we vary business mix, the average difference of VaR using the  $t(10)$  ( $t(5)$ ) instead of the Gaussian copula is 4% (9%). The maximum deviation is 9% (17%) when operational exposure is 30% higher than the benchmark. Looking at the correlation experiment in the bottom panel, the average differences are very similar to the previous experiment at 4% (10%). The maximum deviation of 5% (14%) occurs at  $\rho = 0.1$ , which is slightly below the benchmark correlation of 0.2.

Second, in both experiments, risk implied by the hybrid method is very close to the  $t(5)$  copula. The average deviation in both cases is just 1% with a maximum of 2% for the business mix and 3% for the correlation experiment. The positive bias of the hybrid method relative to the Gaussian copula observed earlier is on the order of incorporating some tail dependence, which from a practical perspective is perhaps not an unreasonable degree of conservatism.<sup>22</sup>

---

<sup>22</sup> Although we report only results for these two experiments, they are very similar across all experiments, whether risk is measured by VaR or ES.

### 6.3.3. *Residual risk*

As discussed in Section 5.4, the regressions for market and credit risk are unlikely to capture all relevant risk for each risk type. Some of the “residual” in both cases will be operational risk as defined here, as well as business and other types of risk. Nonetheless, it can be instructive to examine the impact of this “residual” risk on the total risk distribution by adding the simulated residuals for market and credit risk, as described in Section 5.5, to their respective risk types. Since operational risk is not modeled using the factor approach, we have no “residual” here.

In keeping with the set of experiments presented in Section 6.3, the top panel in Fig. 8 shows the business mix experiment with operational risk. This can be compared to the bottom panel in Fig. 4. The bottom panel of Fig. 8 shows the inter-risk correlation experiment, which can be compared to the bottom panel in Fig. 5.

These new results are broadly in line with the experiments without residual risk. Consider first copula-VaR; we notice immediately by looking at the scale of the vertical axis that the total amount of risk has nearly doubled. This reflects the large share of unexplained variance in our factor regressions. However, the basic features of the residual risk distributions, as characterized by their moments, are quite similar to the systematic risk distributions (see previous discussion in Section 5.7 as well as Table 5). Thus, we find that while the scaling has changed, the basic shapes of the VaR curves are similar with or without residual risk.

The results of the business mix experiment shown in the top panel of Fig. 8 are comparable to the no residual risk experiment in the bottom panel of Fig. 4, albeit with a steeper slope. At 100% operational risk weight, total risk is unchanged from the earlier experiment. When residual volatility is added to the high volatility risk types (market and credit), total risk is

higher when these types have positive weight. Note that the flattening of the total risk curve happens later when residual risk is included. In this case, a greater operational weight is required for the fat left tail of operational risk to offset the new higher volatilities of market and credit risk.

These basic observations carry over to the correlation experiment shown in the bottom panel of Fig. 8. Copula-VaR is substantially higher now that residual risk is included for market and credit risk. However, copula-VaR is now much less sensitive to changes in operational correlation with the other risk types so the VaR curves are flatter. As this correlation increases from 0.0 to 0.8, the proportional VaR increase with no residual risk is 45% (bottom panel of Fig. 5) compared to 7% with residual risk. Interestingly, the hybrid approximation seems to pick up more of the fat-tails in operational risk as  $\rho$  increases. Previously H-VaR converged to copula-VaR, whereas now it remains below (i.e. risk is higher) across the spectrum of  $\rho$  values.

#### *6.4. Sensitivity analysis*

Finally, to assess the relative importance of variation in the determinants of total risk across all of the business mix and correlation experiments, we calculate the range of VaRs across each experiment type. Then, we take the absolute difference between the largest and smallest VaR (within the experiment) and divide by the average VaR. The result is the proportional range of VaRs due to variation of the chosen factor (Table 11). A wide range of VaRs indicates high sensitivity to the factor being varied. As an example, for changes in operational risk correlation, we have the largest VaR of  $-0.582\%$  at a correlation of 0.8 and the smallest VaR of  $-0.401\%$  at a correlation of zero. Using the average VaR of  $-0.482\%$ , the relative sensitivity is 45%. To measure sensitivity to approximation choice, we average the approximation errors across business mix and correlation experiments. For risk sensitivity to copula choice, we take the

average proportional VaR difference using a Student-t(5) copula versus a normal copula across the business mix and correlation experiments reported in Fig. 7. For these last two measures, we first average within each experiment and then across experiments.

The greatest risk sensitivities are to operational risk exposure, approximation choice (for Add-VaR and N-VaR), and operational risk correlation. These are followed by market versus credit risk exposure, copula choice, and market versus credit correlation. From this ordering, we see that assumptions about operational exposures and correlations are much more important for accurate risk estimates than assumptions about relative market and credit exposures or correlations. Risk measurements are also quite sensitive to approximation choice. Adding up the marginal risks or using joint normal risk distribution result in approximation errors of 45% or -42%. Finally, the choice of copula seems to have a relatively modest impact with the next to lowest risk sensitivity of 10%.

## **7. Final Comments**

How would one aggregate market, credit, and operational risk distributions to arrive at a total risk distribution for a financial institution? One solution is to simply add up the risk for each type. Among other things, this assumes no diversification benefits that might accrue across risk types due to non-perfect correlation. Our analysis shows that the additive approach overestimates risk by more than 40%. A typical alternative, assuming joint normality of the risks, underestimates risk by a similar amount.

In this paper, we implement an alternative method to combine these risks using copulas. Our approach forms a joint distribution from specified marginals in an internally consistent and realistic manner while preserving important properties about the individual risks. We find that

interaction at the portfolio level between measures of correlation and the higher moments of the marginals is quite subtle and complicated.

Because the operational risk distribution has much heavier tails than a normal distribution, we find that risk is especially sensitive to the chosen level of operational exposure and correlation with market and credit risk. Our results suggest that approximations that might work well for combining market and credit risk tend to fail in combining the three risks because of the operational risk's unique characteristics. This could have important implications for risk managers and regulators since not only is operational risk difficult to measure, but it also requires special care in aggregation.

We also identify a hybrid approximation that tracks the copula method quite well, especially if the risks are thought to exhibit tail dependence. The hybrid approach has the advantage of being relatively easy to compute as only information about the marginals and the correlation matrix are needed.

What are some implications of our findings for the capital allocation problem faced by the financial institution? We take as given business mix weights, although, of course, these are the result of an internal process that balances risk and expected return. We abstract from this decision process by setting expected returns to zero. To be sure, the choice of business mix depends not only on risk but also on expected profits, lest one conclude that all banks would optimally choose to lower their risk by reducing lending and trading activities. When the risk measure used is sensitive to tail events, business lines with low volatility and high skewness could be less attractive, require more capital, and require higher expected return.

## References

- Alexander, C., Pézier, J., 2003. On the Aggregation of Firm-Wide Market and Credit Risks. ISMA Centre Discussion Papers in Finance 2003-13.
- Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1999. Coherent Measures of Risk. *Mathematical Finance* 9, 203-228.
- Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1997. Thinking coherently. *Risk* 10, 68–71.
- Basel Committee on Banking Supervision, 1988. Internal Convergence of Capital Measurement and Capital Standards. <http://www.bis.org/publ/bcbs04A.pdf>, July.
- Basel Committee on Banking Supervision, 1996. Amendment to the Capital Accord to Incorporate Market Risks. <http://www.bis.org/publ/bcbs24a.htm>, January.
- Basel Committee on Banking Supervision, 2004. International Convergence of Capital Measurement and Capital Standards: A Revised Framework. <http://www.bis.org/publ/bcbs107.htm>, June.
- Berkowitz, J., 2001. Testing the Accuracy of Density Forecasts, applications to Risk Management. *Journal of Business and Economic Statistics* 19, 465-474.
- Bollerslev, T., Wooldridge, J.M., 1992. Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances. *Economics Letters* 12, 143-172.
- Bouyé, E., 2001. Multivariate Extremes at Work for Portfolio Risk Measurement. FERC working paper, WP01-10, (September); available at [http://users.wbs.ac.uk/cms\\_attachment\\_handler.cfm?f=mev2.pdf&t=mev2.pdf](http://users.wbs.ac.uk/cms_attachment_handler.cfm?f=mev2.pdf&t=mev2.pdf).
- Bradley, B.O., Taqqu, M.S., 2003. Financial Risk and Heavy Tails. In: S.T. Rachev (Ed.), *Handbook of Heavy-Tailed Distributions in Finance*. Elsevier, Rotterdam, pp. 35-103.
- Christoffersen, P., Diebold F.X., 2000. How Relevant is Volatility Forecasting for Financial Risk Management? *Review of Economics and Statistics* 82, 12-22.
- Crouhy, M., Galai, D., Mark, R., 2001. *Risk Management*. McGraw Hill, New York.
- De Fountnouvelle, P., DeJesus-Rueff, V., Jordan, J., Rosengren, E., 2003. Using Loss Data to Quantify Operational Risk. Federal Reserve Bank of Boston Working Paper.
- Demarta, S., McNeil, A.J., 2004. The t Copula and Related Copulas. forthcoming, *International Statistical Review*.
- Diebold, F.X., Gunther, T.A., Tay, A.S., 1998. Evaluating Density Forecasts. *International Economic Review* 39, 863-883.
- Dimakos, X.K., Aas, K., 2004. Integrated Risk Modeling. *Statistical Modelling* 4, 265-277.
- Embrechts, P., Lindskog, F., McNeil, A.J., 2003. Modeling Dependence with Copulas and Applications to Risk Management. In S. Rachev (Ed.), *Handbook of Heavy Tailed Distributions in Finance*. Elsevier, Rotterdam, pp. 329-384.
- Embrechts, P., McNeil, A.J., Straumann, D., 1999, Correlation: Pitfalls and Alternatives. *Risk* 12, 69-71.

- Embrechts, P., McNeil, A.J., Straumann, D., 2002. Correlation and Dependence in Risk Management: Properties and Pitfalls. In: M.A.H. Dempster (Ed.) Risk Management: Value at Risk and Beyond. Cambridge University Press, Cambridge, UK, pp. 176-223.
- Engle, R.F., Mezrich, J., 1996. GARCH for Groups. Risk 9, 36-40.
- Fermanian, J.-D., Scaillet, O., 2003. Nonparametric Estimation of Copulas for Time Series. Journal of Risk 5, 25-54.
- Frey, R., McNeil, A.J., 2001. Modelling Dependent Defaults. Unpublished working paper, ETH Zurich.
- Glasserman, P., Heidelberger, P., Shahabuddin, P., 2002. Portfolio Value-at-Risk with Heavy-Tailed Risk Factors. Mathematical Finance 12, 239-270.
- Gourieroux, C., Laurent, J.P., Scaillet, O., 2000. Sensitivity Analysis of Values at Risk. Journal of Empirical Finance 7, 225-245.
- Hirtle, B., 2003. What Market Risk Capital Reporting Tells Us About Bank Risk. Federal Reserve Bank of New York Economic Policy Review 9, 37-54.
- Huang, J.-Z., and Huang, M., 2004. How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk? Unpublished working paper, Graduate School of Business, Stanford University.
- Johnson, N.L., Kotz, S., Balakrishnan, N., 1994. Discrete Multivariate Distributions. Wiley, New York.
- Joint Forum, 2001. Risk Management Practices and Regulatory Capital, Cross-Sectoral Comparison. November; available at <http://www.bis.org/publ/joint03.pdf>.
- Joint Forum. 2003. Trends in Risk Integration and Aggregation. August; available at <http://www.bis.org/publ/joint07.pdf>.
- Jorion, P., 2001. Value at Risk: The Benchmark for Managing Financial Risk 2nd Edition, McGraw Hill, New York, :.
- Jorion, P., 2002. How Informative are Value-at-Risk Disclosures? Accounting Review 77, 911-31.
- J.P. Morgan, 1995, RiskMetrics Technical Manual, New York.
- Koyluoglu, H.U., Stoker, J., 2002. Honour your Contribution. Risk 15, 90-94.
- Kroner, K.F., Ng, V.K., 1998. Modeling Asymmetric Comovements of Asset Returns. Review of Financial Studies 11, 817-844.
- Kuritzkes, A., Schuermann, T., Weiner, S.M., 2003. Risk Measurement, Risk Management and Capital Adequacy of Financial Conglomerates. In Herring, R. and R. Litan (Eds.), Brookings-Wharton Papers in Financial Services 2003, pp. 141-194.
- Kuritzkes, A., 2002. Operational Risk Capital: A Problem of Definition. Journal of Risk Finance, Fall, 1-10.
- Li, D.X., 2000. On Default Correlation: A Copula Function Approach. Journal of Fixed Income, March, 43 - 54.



- Longin, F. Solnik, B., 2001. Extreme Correlation of International Equity Markets. *Journal of Finance* 56, 649-676.
- Markowitz, H., 1959. *Efficient Diversification of Investments*, John Wiley, New York.
- Nelsen, R.B., 1999. *An Introduction to Copulas*. Springer-Verlag, New York.
- Patton, A., 2002. Modeling Time-Varying Exchange Rate Dependence Using the Conditional Copula. UCSD Discussion Paper 01-09.
- Poon, S.-H., Rockinger, M., Tawn, J., 2004. Extreme Value Dependence in Financial Markets: Diagnostics, Models and Financial Implications. *Review of Financial Studies* 17, 581-610.
- Rosenberg, J., 2003. Nonparametric Pricing of Multivariate Contingent Claims. *Journal of Derivatives* 10, 9-26.
- United States General Accounting Office (USGAO), 1998. Risk-Based Capital: Regulatory and Industry Approaches to Capital and Risk. Report to the Chairman, Committee on Banking, Housing, and Urban Affairs, U.S. Senate and the Chairman, Committee on Banking and Financial Services, House of Representatives, GAO/GGD-98-153.
- Wang, S.S., 1998. Aggregation of Correlated Risk Portfolios: Models and Algorithms. *Proceedings of the Casualty Actuarial Society* 85, 848-939.
- Ward, L.S., Lee, D.H., 2002. Practical Application of the Risk-Adjusted Return on Capital Framework. CAS Forum Summer 2002, Dynamic Financial Analysis Discussion Papers; available at <http://www.casact.org/pubs/forum/02sforum/02sftoc.htm>.

Table 1: Summary table for 17 bank holding companies (BHCs) in the sample, millions US dollars, year-end 2002. Stocks (e.g. assets) are year-end, flows (e.g. net income) are the sum of four quarters. Means and medians are taken over the four quarters. The data source is Y-9C BHC regulatory reports.

<b>Bank Holding Company</b>	<b>Total Assets</b>	<b>Lending Assets</b>	<b>Trading Assets</b>	<b>Percent Lending Assets</b>	<b>Percent other Assets</b>	<b>Net Income</b>	<b>Net Interest Income</b>	<b>Provisions</b>	<b>Trading Revenue</b>	<b>Minimum Market Risk Capital Requirement</b>	<b>Market Risk Share of Capital Requirement</b>	<b>Return on Assets</b>
ABN Amro N. America	139,605	61,343	3,106	43.9%	53.8%	434	2,349	145	147	121	0.16%	0.31%
Bank of America	660,458	356,451	95,829	54.0%	31.5%	9,249	20,210	1,165	778	2,313	0.44%	1.40%
Bank One	277,383	148,125	11,000	53.4%	42.6%	3,295	8,569	628	224	140	0.06%	1.19%
Citigroup	1,097,190	477,323	155,208	43.5%	42.3%	15,276	37,882	2,690	4,513	2,446	0.35%	1.39%
FleetBoston Financial	190,589	119,427	4,486	62.7%	35.0%	1,188	6,369	813	584	251	0.14%	0.62%
HSBC N. America	113,371	59,649	14,592	52.6%	34.5%	860	2,813	42	96	172	0.24%	0.76%
J.P. Morgan Chase	758,800	216,253	248,301	28.5%	38.8%	1,663	11,526	921	2,594	2,663	0.58%	0.22%
Keycorp	84,710	62,457	2,561	73.7%	23.2%	979	2,762	147	65	15	0.02%	1.16%
Mellon Financial	36,306	8,420	1,911	23.2%	71.5%	682	697	6	183	60	0.23%	1.88%
PNC	66,410	37,089	926	55.8%	42.8%	1,184	2,198	65	95	16	0.03%	1.78%
State Street	85,794	5,128	3,435	6.0%	90.0%	1,015	979	1	259	27	0.10%	1.18%
Suntrust Banks	117,323	80,916	1,718	69.0%	29.6%	1,332	3,244	97	120	20	0.02%	1.14%
Deutsche Bank (Taunus)	224,278	16,843	64,270	7.5%	63.8%	-2,109	1,029	97	-165	243	0.00%	-0.94%
Bank of New York	77,576	32,015	7,309	41.3%	49.3%	902	1,739	93	234	43	0.06%	1.16%
U.S. Bancorp	180,027	120,410	898	66.9%	32.6%	3,289	6,840	349	227	47	0.03%	1.83%
Wachovia	341,839	171,801	33,155	50.3%	40.0%	3,579	10,021	308	24	505	0.19%	1.05%
Wells Fargo	349,259	254,453	10,167	72.9%	24.2%	5,435	14,858	438	321	374	0.13%	1.56%
Mean	277,241	127,009	37,591	48%	44%	2,839	7,887	456	606	559	0.18%	1.07%
Median	177,941	76,343	8,131	53%	39%	1,188	3,244	195	224	168	0.14%	1.19%

Table 2: We estimate a diagonalized asymmetric multivariate GARCH model (Kroner and Ng, 1998) for market and credit risk factors using the method of maximum likelihood. For the market risk factor model, equation 1 is for equity returns, equation 2 is for currency returns, and equation 3 is for log-differenced interest rates. This model is specified using variance targeting, so standard errors are not computed for the parameters of the omega matrix, which are functions of sample covariances and volatility model parameters. For the credit risk factor model, equation 1 is for AA credit spreads and equation 2 is for BBB credit spreads. Data definitions are given in Section 5.3. Note that the volatility parameters  $\omega$ ,  $a$ ,  $b$ ,  $g$  correspond to the square root of the corresponding parameters in a univariate GARCH model. Ljung-Box p-values are reported for 20 lags in the market model and 3 lags in the credit model for squared returns and standardized squared returns ( $\varepsilon_t/\sqrt{h_t}$ ) using estimated conditional volatilities.

$$\mathbf{r}_t = \boldsymbol{\mu} + \mathbf{A}\mathbf{r}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t | \Xi_{t-1} \sim N(\mathbf{0}, \mathbf{H}_t)$$

$$\mathbf{H}_t = \boldsymbol{\Omega} + \mathbf{A}' \boldsymbol{\varepsilon}_{t-1} \boldsymbol{\varepsilon}_{t-1}' \mathbf{A} + \mathbf{B}' \mathbf{H}_{t-1} \mathbf{B} + \mathbf{G}' \boldsymbol{\eta}_{t-1} \boldsymbol{\eta}_{t-1}' \mathbf{G}.$$

**Estimated Asymmetric Multivariate GARCH Risk Factor Models**

	<u>Market risk factor model</u>			<u>Credit risk factor model</u>		
	Parameter estimate	Standard error	P-Value	Parameter estimate	Standard error	P-Value
$\mu_1$	2.51E-04	1.07E-04	0.0188	0.5300	1.0357	0.6089
$\mu_2$	2.52E-05	3.80E-05	0.5067	1.6946	1.3182	0.1986
$\mu_3$	1.73E-05	8.87E-05	0.8456			
$\alpha_1$	0.0819	0.0130	<0.001	0.2420	0.0634	<0.001
$\alpha_2$	0.0495	0.0118	<0.001	0.3086	0.0948	0.0011
$\alpha_3$	0.0970	0.0119	<0.001			
$\omega_{11}$	1.12E-06			2.7851	0.5576	<0.001
$\omega_{22}$	3.89E-08			1.9908	0.4310	<0.001
$\omega_{33}$	4.34E-07					
$\omega_{21}$	2.20E-08			-0.0069	34.1960	0.9998
$\omega_{31}$	-3.00E-07					
$\omega_{32}$	1.37E-08					
$a_1$	0.1830	0.0075	<0.001	0.1563	0.1603	0.3295
$a_2$	0.1760	0.0038	<0.001	0.0787	0.1922	0.6823
$a_3$	0.2217	0.0050	<0.001			
$b_1$	0.9635	0.0013	<0.001	0.8630	0.0449	<0.001
$b_2$	0.9831	0.0007	<0.001	0.9117	0.0246	<0.001
$b_3$	0.9725	0.0013	<0.001			
$g_1$	0.2367	0.0047	<0.001	0.6818	0.1187	<0.001
$g_2$				0.7266	0.0857	<0.001
Number of observations (log-likelihood)		7199 (78621)		169 (-1141)		
Eq. 1: Ljung-Box p-value (squared returns)		<0.001		0.030		
Eq. 2: Ljung-Box p-value (squared returns)		<0.001		0.101		
Eq. 3: Ljung-Box p-value (squared returns)		<0.001				
Eq. 1: Ljung-Box p-value (squared standardized residuals)		0.953		0.842		
Eq. 2: Ljung-Box p-value (squared standardized residuals)		0.048		0.949		
Eq. 3: Ljung-Box p-value (squared standardized residuals)		0.532				

Table 3: Estimates of market risk regressions using pooled OLS with bank fixed effects. Dependent variable is as annualized trading revenue divided by trading assets. Data is from Y-9C bank holding company regulatory reports over the period 1994Q1-2002Q4 using the bank holding company sample discussed in Section 5.1. Equity return is the annualized log-return of the S&P500 from Datastream.  $\Delta$  interest rate is the change in the 10-year Constant Maturity Treasury (CMT) rate from the Federal Reserve Board H.15 release. Currency return is the annualized log difference of the trade-weighted currency index from the Federal Reserve Board G.5 release. Equity, interest rate, and currency volatilities are the log of daily variances from the trivariate GARCH model (Table 2) summed over the quarter and then scaled to annual terms.

<b>Variable</b>	<b>Parameter Estimate (<math>\hat{\beta}</math>)</b>	<b>Standard Error</b>	<b>P-value</b>
Equity return	0.0204	0.0058	0.0005
Equity volatility	0.0122	0.0033	0.0002
$\Delta$ interest rate	-0.0056	0.0042	0.1896
Interest rate volatility	-0.0030	0.0052	0.5591
Currency return	-0.0048	0.0148	0.7442
Currency volatility	-0.0113	0.0052	0.0302
$\bar{R}^2$ fixed effects only	58.3%		
$\bar{R}^2$ fixed effects and factors	59.7%		
Average $\bar{R}^2$ from bank-level regressions	10.2%		
Number of observations	612		

Table 4: Estimates of credit risk regressions using pooled OLS with bank fixed effects. Dependent variable is annualized net interest income less provisions. Data is from Y-9C bank holding company regulatory reports over the period 1994Q1-2002Q4 using the bank holding company sample discussed in Section 5.2. Explanatory variables are change in the AA and BBB spreads. We use monthly data on the yield-to-maturity and Macaulay duration from the Merrill Lynch AA and BBB-rated corporate bond indices obtained from Bloomberg. We compute spreads by subtracting from each yield-to-maturity a weighted average of the 5, 7, and 10 year Constant Maturity Treasury (CMT) rates. The weights are based on the difference between the duration of the relevant index and the duration of a par 5, 7, or 10 year Treasury paying the relevant CMT rate.

<b>Variable</b>	<b>Parameter Estimate (<math>\hat{\beta}</math>)</b>	<b>Standard Error</b>	<b>P-value</b>
$\Delta$ AA spread	$5.80 \times 10^{-5}$	$1.59 \times 10^{-5}$	0.0003
$\Delta$ BBB spread	$-3.38 \times 10^{-5}$	$9.36 \times 10^{-6}$	0.0003
$\bar{R}^2$ fixed effects only	56.2%		
$\bar{R}^2$ with fixed effects and factors	56.8%		
Average $\bar{R}^2$ from bank-level regressions	4.8%		
Number of observations	612		

Table 5: Characteristics of moments of residuals across 17 bank specific regressions. These bank specific regressions are estimated using the same specification as the pooled OLS models. Details of market risk regression are given in Table 3, while details of credit risk regression are given in Table 4. In the top panel, the first row represents the cross-sectional average of the standard deviation, skewness and kurtosis of the 17 market risk regression residuals. The bottom panel reports the same statistics for the credit risk regression residuals. Because these distributions are skewed, the median is different from the mean. The rows labeled “Simulated” denote the moments of the simulated residuals; see Section 5.4.

<b>Market Risk Residuals</b>			
	Std. Dev.	Skewness	Kurtosis
Mean	2.27%	0.5	8.4
Median	1.18%	0.3	7.5
Simulated	1.18%	0.5	6.4
<b>Credit Risk Residuals</b>			
Mean	1.25%	-0.6	7.3
Median	1.04%	-0.8	5.5
Simulated	1.04%	-0.6	6.0

Table 6: Characteristics of simulated marginal risk distributions for 200,000 bank-years using methods described in Section 5.7. All values are annualized, and the mean is zero by construction. The total risk distribution is generated by simulating the weighted average of market (3.1%), credit (29.1%), and operational (67.8%) returns from a joint distribution defined using a normal copula with the benchmark correlations (market versus credit = 0.5, market/credit versus operational = 0.2). Assets at risk for market risk is equal to trading assets, for credit risk is equal to lending assets, and for operational risk is equal to total assets. Total book is the the sum of these three assets at risk.

	<b>Market Risk</b>	<b>Credit Risk</b>	<b>Operational Risk</b>	<b>Total Risk</b>
$\sigma$ (assets at risk)	0.58%	0.19%	0.04%	0.11%
$\sigma$ (total book)	0.02%	0.05%	0.03%	0.08%
Skewness	0.2	-1.3	-4.5	-1.1
Kurtosis	3.7	16.1	35.3	9.6
0.1 percentile	-1.81%	-1.20%	-0.37%	-0.43%
0.5 percentile	-1.47%	-0.72%	-0.27%	-0.29%
1 <sup>st</sup> percentile	-1.32%	-0.57%	-0.17%	-0.24%

Table 7: Inter-risk correlations (Table 4 from Kuritzkes, Schuermann and Weiner (KSW), 2003). The entries correspond to examples of inter-risk correlations from three studies (academic and industry), as summarized in KSW, that are closely related to the risk types analyzed in this paper. The first set is taken from Dimakos and Aas (2004). The second set comes from Ward and Lee (2002) who have separate correlations for market and asset-liability management risk; the table entry is the higher of the two correlations. The third set is cited in Joint Forum (2001, p. 25), and is based on interview findings and reflect the values used by Dutch and other financial conglomerates in their internal capital models.

	<b>Credit</b>	<b>Market/Asset- Liability Management</b>	<b>Operational and Other</b>
<b>Credit</b>	100%	30%	20%
		30%	44%
		80%	40%
<b>Market/Asset- Liability Management</b>	30%	100%	20%
	30%		13%
	80%		40%
<b>Operational and Other</b>	20%	20%	
	44%	13%	100%
	40%	40%	



Table 8: Standardized quantile and volatility for 99.9% VaR (0.1% tail) of the total risk distribution as business mix changes. Business mix and risk weights are defined in Section 5.2. In Panel A, each row represents market weight versus credit weight while holding the operational weight constant. The italicized row is closest to the benchmark risk weights: 3.1% for market risk, 29.1% for credit risk and 67.8% for operational risk. Some intermediate values (rows) are omitted for space reasons. In Panel B, each row represents the *change* in operational weight, while holding the relative weights of market and credit constant. The italicized row represents the benchmark risk weights. The standardized quantiles for the market, credit and operational risk distributions are -3.12, -6.29 and -9.09 respectively. The marginal volatilities in same order are 0.58%, 0.19% and 0.04%. Add-VaR (Eq. (3.7)) adds the VaRs of the marginal risk distributions, N-VaR (Eq. (3.8)) assumes joint normality, and H-VaR (Eq. (3.6)) combines marginal risks a formula that would apply to an elliptical distribution. Approximation error is (approximation VaR – Copula VaR) / Copula VaR. See also Fig. 4.

### Business Mix

<b>Panel A: Market versus Credit Risk</b>									
Share of Market vs. Credit	Market Weight	Credit Weight	Operat. Weight	99.9% VaR	Volatility	Standardized Quantile	Approximation Error		
							N-VaR to Copula	H-VaR to Copula	Add-VaR to Copula
0%	0.00%	32.19%	67.81%	-0.439%	0.071%	-6.15	-50%	13%	46%
<i>10%</i>	<i>3.22%</i>	<i>28.97%</i>	<i>67.81%</i>	<i>-0.434%</i>	<i>0.076%</i>	<i>-5.70</i>	-46%	13%	52%
20%	6.44%	25.75%	67.81%	-0.434%	0.084%	-5.15	-40%	13%	56%
40%	12.88%	19.31%	67.81%	-0.449%	0.107%	-4.21	-27%	14%	60%
60%	19.31%	12.88%	67.81%	-0.481%	0.133%	-3.61	-14%	14%	57%
80%	25.75%	6.44%	67.81%	-0.549%	0.163%	-3.37	-8%	10%	45%
90%	28.97%	3.22%	67.81%	-0.591%	0.178%	-3.32	-7%	8%	38%
100%	32.19%	0.00%	67.81%	-0.630%	0.194%	-3.26	-5%	7%	33%

<b>Panel B: Operational versus Market and Credit Risk</b>									
Share Shifted from Market and Credit to Operational	Market Weight	Credit Weight	Operat. Weight	99.9% VaR	Volatility	Standardized Quantile	Approximation Error		
							N-VaR to Copula	H-VaR to Copula	Add-VaR to Copula
-50%	4.65%	43.64%	51.71%	-0.599%	0.105%	-5.72	-46%	5%	34%
-30%	4.03%	37.82%	58.15%	-0.530%	0.093%	-5.70	-46%	8%	40%
-20%	3.72%	34.91%	61.37%	-0.497%	0.087%	-5.70	-46%	9%	44%
-10%	3.41%	32.00%	64.59%	-0.461%	0.082%	-5.66	-45%	12%	49%
<i>0%</i>	<i>3.10%</i>	<i>29.09%</i>	<i>67.81%</i>	<i>-0.434%</i>	<i>0.076%</i>	<i>-5.71</i>	-46%	13%	52%
10%	2.79%	26.18%	71.03%	-0.408%	0.071%	-5.77	-46%	14%	54%
20%	2.48%	23.27%	74.25%	-0.387%	0.065%	-5.91	-48%	15%	55%
30%	2.17%	20.36%	77.47%	-0.366%	0.061%	-6.05	-49%	16%	56%
50%	1.55%	14.55%	83.90%	-0.355%	0.052%	-6.86	-55%	10%	45%
80%	0.62%	5.82%	93.56%	-0.362%	0.043%	-8.47	-64%	2%	19%
100%	0.00%	0.00%	100.00%	-0.373%	0.041%	-9.09	-66%	0%	0%

Table 9: Standardized quantile and volatility for 99.9% VaR (0.1% tail) of the total risk distribution as inter-risk correlation changes. The first set are the results of changing correlation between market and credit risk while the correlation of both risk types to operational risk remains constant at benchmark level of 20%. The second set describes the impact of changing correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%. Note that the maximum correlation possible for the operational risk experiment (bottom panel) is 0.8; see also Footnote 20 and Table 10 below. The italicized row represents the benchmark correlations. Both sets of results: benchmark weights are 3.1% for market risk, 29.1% for credit risk and 67.8% for operational risk. Add-VaR (Eq. (3.7)) adds the VaRs of the marginal risk distributions, N-VaR (Eq. (3.8)) assumes joint normality, and H-VaR (Eq. (3.6)) combines marginal risks a formula that would apply to an elliptical distribution. Approximation error is (approximation VaR – Copula VaR) / Copula VaR. Diversification benefit in the last column is (Add-VaR – Copula VaR)/Add-VaR. Note that the standardized quantile for the market, credit and operational risk distributions are -3.12, -6.29 and -9.09 respectively. The marginal volatilities in the same order are 0.58%, 0.19% and 0.04%. See also Fig. 5.

### Correlation

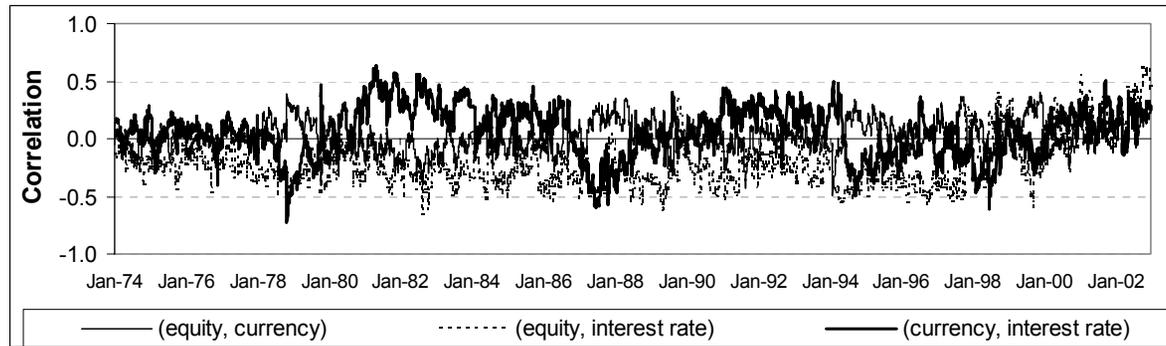
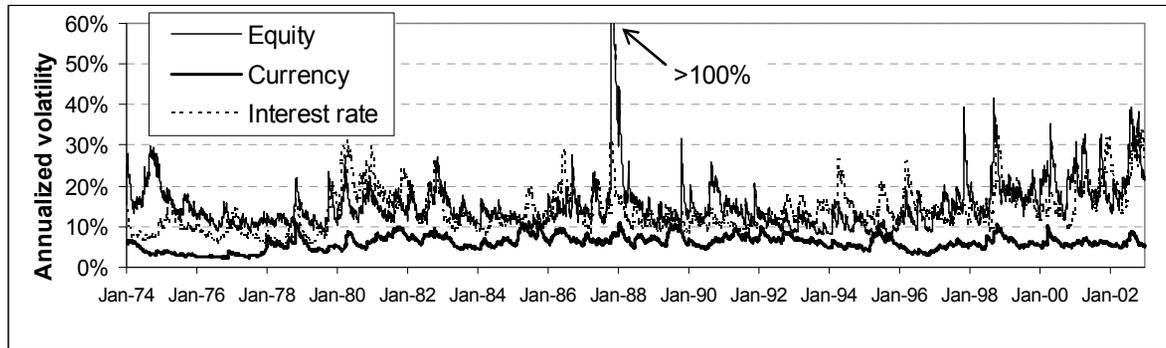
<b>Panel A: Market versus Credit</b>							
Rho	99.9% VaR	Volatility	Standard ized Quantile	<u>Approximation Error</u>			
				N-VaR to Copula	H-VaR to Copula	Add-VaR to Copula	Diversification Benefit (Copula)
0	-0.402%	0.069%	-5.79	-47%	17%	63%	39%
0.1	-0.406%	0.071%	-5.74	-46%	17%	62%	38%
0.2	-0.414%	0.072%	-5.74	-46%	14%	58%	37%
0.3	-0.426%	0.073%	-5.80	-47%	13%	54%	35%
0.4	-0.426%	0.075%	-5.70	-46%	14%	54%	35%
<i>0.5</i>	<i>-0.434%</i>	<i>0.076%</i>	<i>-5.71</i>	<i>-46%</i>	<i>13%</i>	<i>52%</i>	<i>34%</i>
0.6	-0.433%	0.077%	-5.60	-45%	15%	53%	35%
0.7	-0.437%	0.079%	-5.56	-44%	14%	50%	34%
0.8	-0.434%	0.080%	-5.44	-43%	14%	50%	33%
0.9	-0.438%	0.081%	-5.40	-43%	15%	50%	33%
<b>Panel B: Operational versus Market and Credit</b>							
0	-0.401%	0.072%	-5.59	-45%	14%	65%	39%
0.1	-0.418%	0.074%	-5.66	-45%	13%	58%	37%
<i>0.2</i>	<i>-0.434%</i>	<i>0.076%</i>	<i>-5.71</i>	<i>-46%</i>	<i>13%</i>	<i>52%</i>	<i>34%</i>
0.3	-0.452%	0.078%	-5.78	-47%	12%	45%	31%
0.4	-0.473%	0.080%	-5.90	-48%	10%	39%	28%
0.5	-0.499%	0.082%	-6.05	-49%	8%	32%	24%
0.6	-0.522%	0.085%	-6.17	-50%	6%	26%	20%
0.7	-0.554%	0.087%	-6.38	-52%	3%	19%	16%
0.8	-0.582%	0.089%	-6.54	-53%	1%	13%	11%

Table 10: Relation between rho and correlation from a multivariate distribution created using simulated returns (see Table 6) and a normal copula with benchmark weights (3.1% for market risk, 29.1% for credit risk and 67.8% for operational risk). The second column corresponds to the top panel in Fig. 5 where the correlation between market and credit is allowed to change while each of their correlations with operational risk remains fixed at the benchmark level ( $\rho = 20\%$ ). The third column corresponds to the bottom panel of Fig. 5 where the correlation between market and credit stays fixed at the benchmark level ( $\rho = 50\%$ ) as it changes with respect to operational risk and the other two risk types.

<b>rho</b>	<b>correlation (market, credit)</b>	<b>correlation (operational, {market, credit})</b>
0.1	0.09	0.08
0.2	0.19	0.16
0.3	0.28	0.24
0.4	0.38	0.32
0.5	0.47	0.40
0.6	0.57	0.49
0.7	0.67	0.58
0.8	0.77	0.67

Table 11: Risk sensitivities. Range of 99.9% VaRs (0.1% tail) across different factors and experiments in descending order. Average risk sensitivity is  $|\text{Maximum VaR} - \text{Minimum VaR}| / \text{Average VaR}$ , using the maximum or minimum VaR entries across the respective business mix (Table 8) and correlation (Table 9) experiments. Risk sensitivities for exposure experiments are from Table 8 and for correlation experiments from Table 9. We calculate approximation error as  $(\text{approximation VaR} - \text{Copula VaR}) / \text{Copula VaR}$ . We report averages of these values under “Approximation choice” in the table. Add-VaR (Eq. (3.7)) adds the VaRs of the marginal risk distributions, N-VaR (Eq. (3.8)) assumes joint normality, and H-VaR (Eq. (3.6)) combines marginal risks a formula that would apply to an elliptical distribution. Copula choice is the average proportional difference for VaR using the Student-t(5) copula versus the normal copula versus across the business mix and correlation experiments reported in Fig. 7.

<b>Factor</b>	<b>Average Risk Sensitivity</b>
Operational risk exposure	58%
Approximation choice	
Add-VaR	45%
N-VaR	-42%
H-VaR	11%
Operational risk correlation	45%
Market versus credit risk exposure	40%
Copula choice	10%
Market versus credit risk correlation	8%



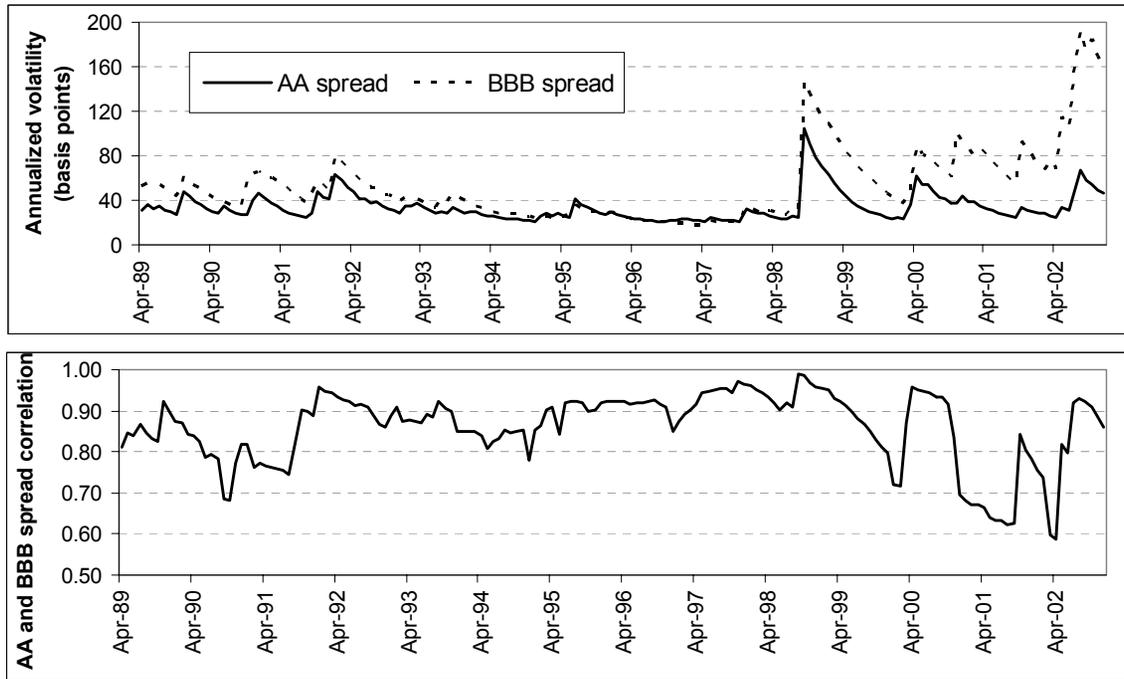


Fig. 1: Conditional volatilities and correlations from estimated multivariate GARCH model (Table 2). Top two panels plot market risk factor volatilities and correlations over the period 1974-2002. Bottom two panels plot credit risk factor volatilities and correlations over the period 1989-2002. Data definitions are given in Section 5.3.

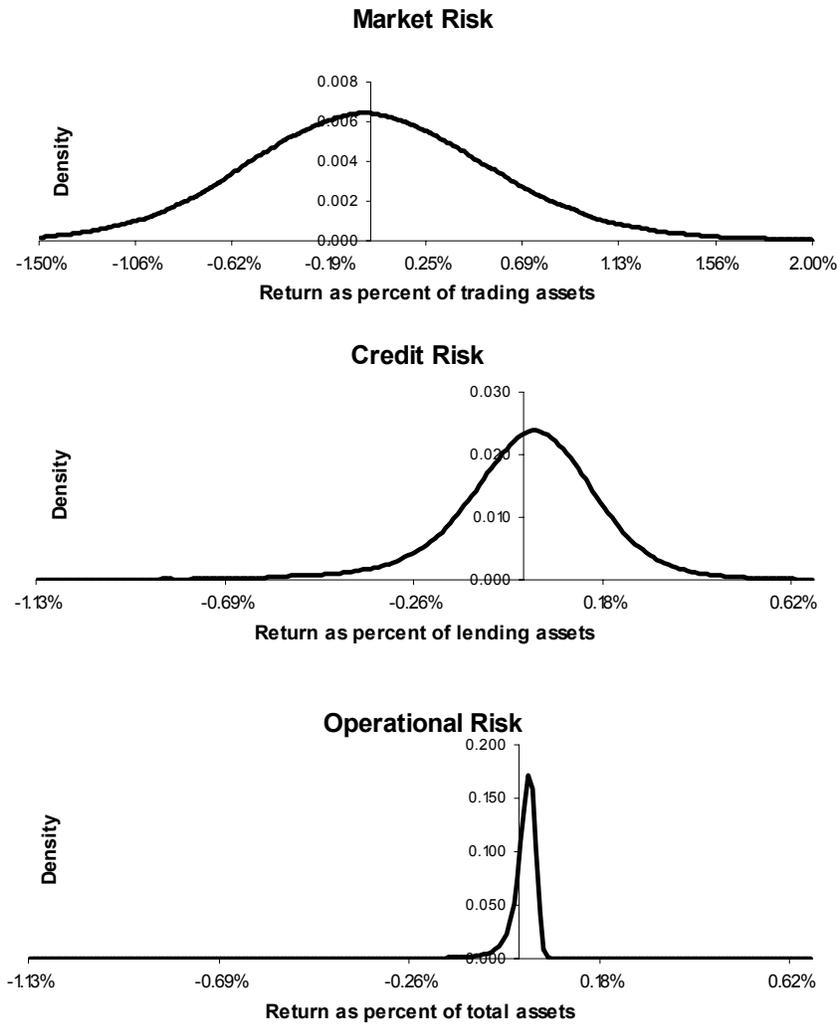


Fig. 2: Market, credit and operational risk distributions estimated using simulated factor returns and estimated factor exposures based on methods described in Section 5.7 and summarized in Table 6. The three panels display kernel density plot (Epanechnikov kernel, Silverman's optimal bandwidth) of simulated marginal risk distributions for 200,000 bank-years.

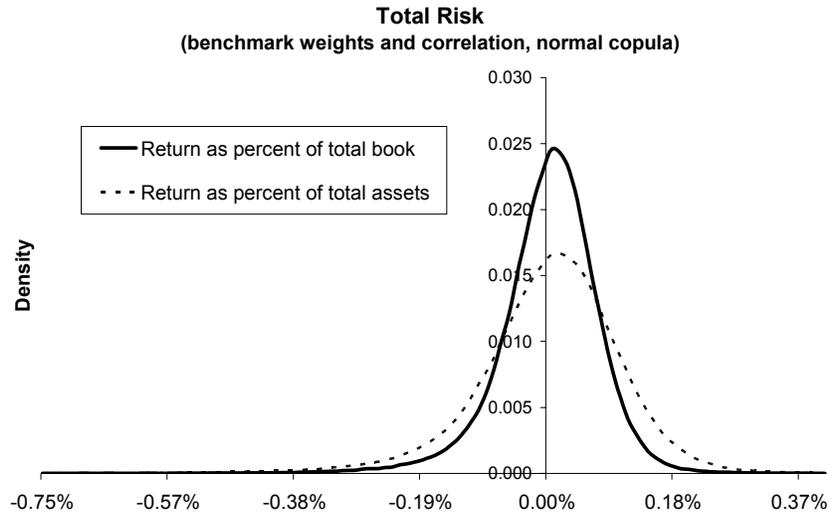


Fig. 3: Total risk distribution estimated using market, credit and operational risks (see Table 6 and Fig. 2). These are combined using a normal copula, benchmark weights (3.1% for market risk, 29.1% for credit risk and 67.8% for operational risk) and benchmark correlations (50% for market and credit, and 20% of those two risk types with operational risk). The figure displays a kernel density plot (Epanechnikov kernel, Silverman’s optimal bandwidth) using 200,000 simulated bank years. Total book is the the sum of trading assets, lending assets, and total assets.



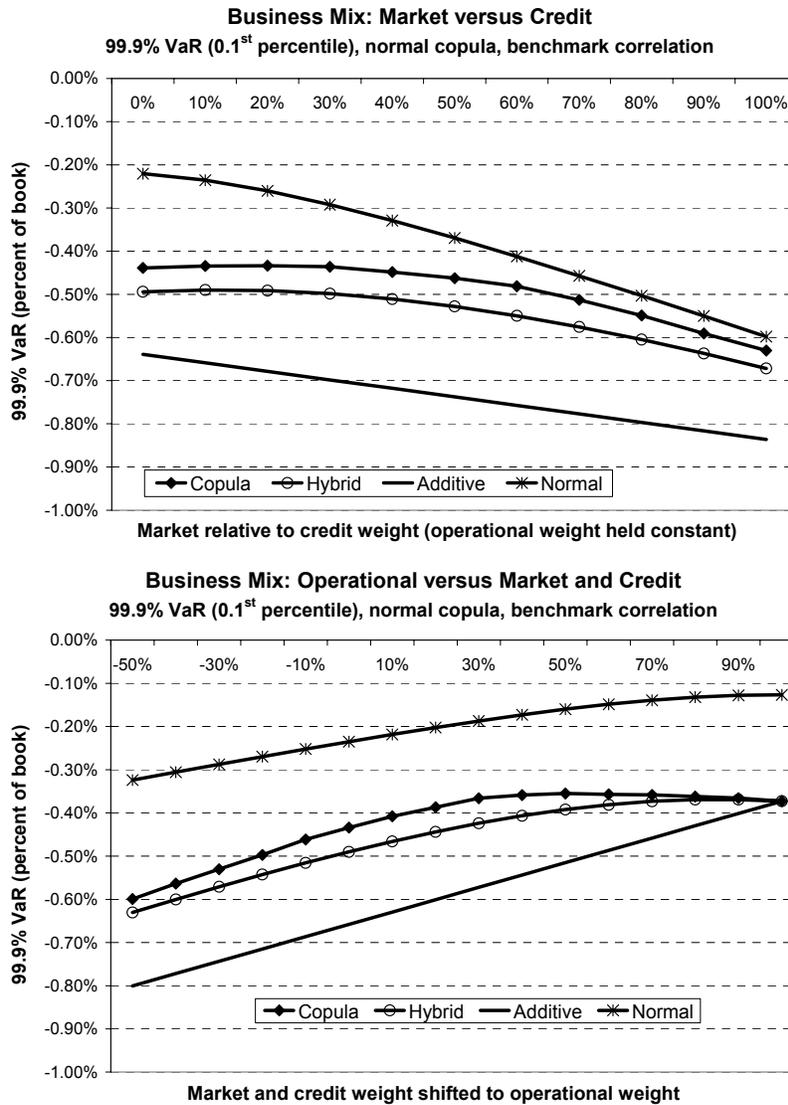


Fig. 4: Impact of business mix; 99.9% VaR (0.1% tail) of the total risk distribution. Additional detail given in Table 8. Top panel: Impact of market weight relative to credit weight while holding operational weight constant. Bottom panel: Impact of shift in market and credit weight to operational weight while holding relative market and credit weights constant. Both panels: normal copula with the benchmark inter-risk correlations of 50% for market and credit, and 20% of those two risk types with operational risk. In the legend, Copula is VaR computed using the copula-based total risk distribution, Hybrid (Eq. (3.6)) is an approximate VaR approach that combines marginal risks using a formula that would apply to an elliptical distribution, Additive adds up the marginal VaRs (Eq. (3.7)), and Normal (Eq. (3.8)) is VaR computed assuming joint normality.

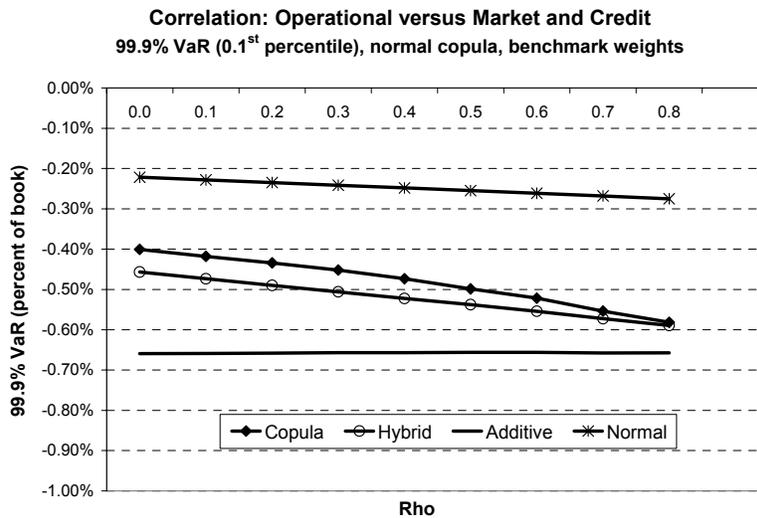
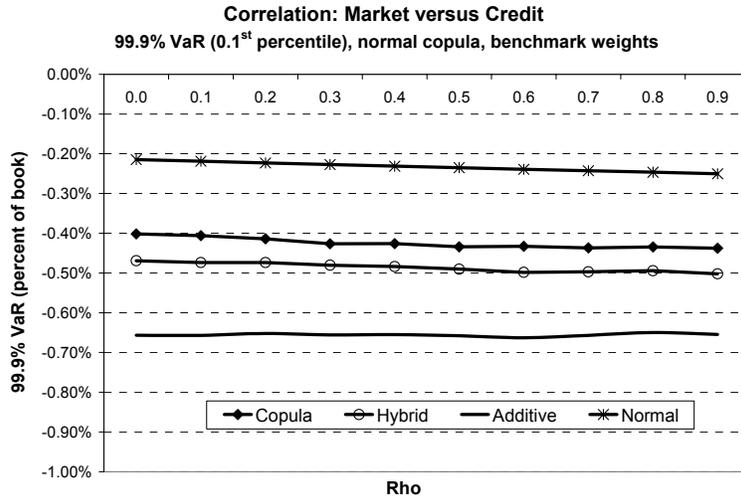


Fig. 5: Impact of inter-risk correlation ( $\rho$ ); 99.9% VaR (0.1% tail) of the total risk distribution. Additional detail given in Table 9. Top panel: Impact of changing correlation between market and credit risk while correlation of both risk types to operational risk remains constant at benchmark level of 20%. Bottom panel: impact of changing correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%. Both panels: benchmark weights are 3.1% for market risk, 29.1% for credit risk and 67.8% for operational risk. In the legend, Copula is VaR computed using the copula-based total risk distribution, Hybrid (Eq. (3.6)) is an approximate VaR approach that combines marginal risks using a formula that would apply to an elliptical distribution, Additive adds up the marginal VaRs (Eq. (3.7)), and Normal (Eq. (3.8)) is VaR computed assuming joint normality.

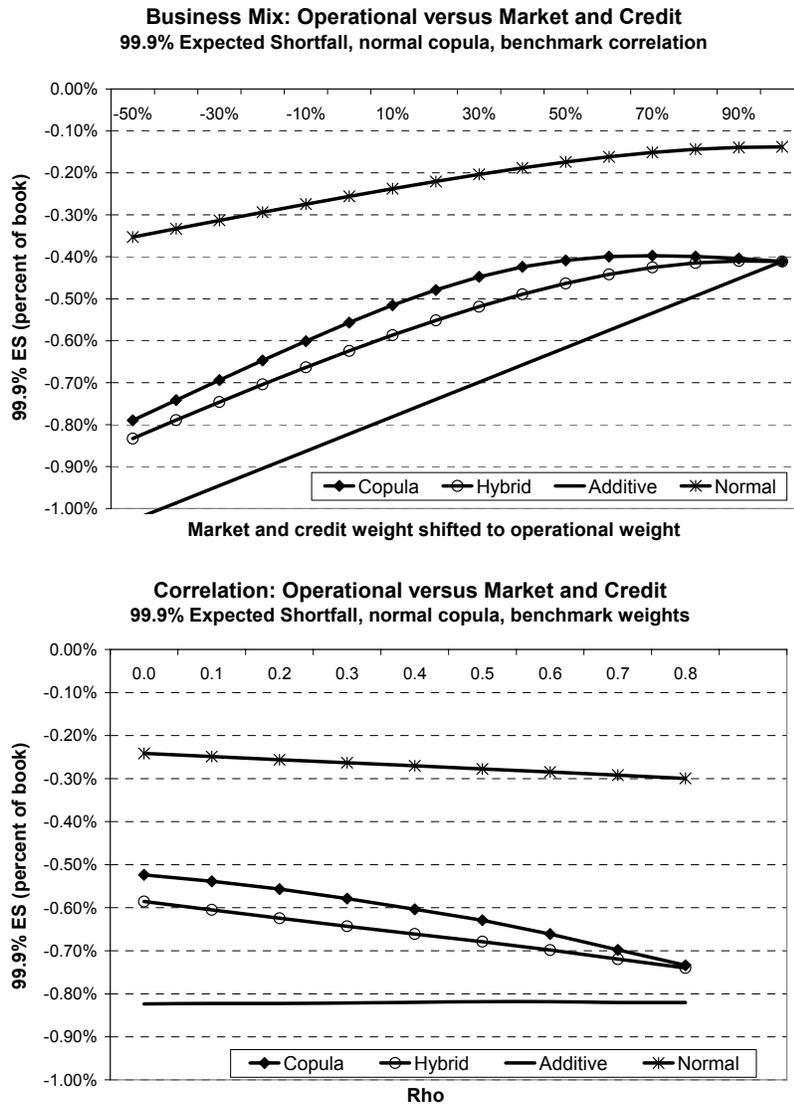


Fig. 6: Impact of business mix and inter-risk correlation on the 99.9% expected shortfall (ES) of the total risk distribution. Top panel: Impact of shift in market and credit weight to operational weight while holding relative market and credit weights constant; analogous to bottom panel of Fig. 4. Bottom panel: impact of changing correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%; analogous to bottom panel of Fig. 5. In the legend, Copula is VaR computed using the copula-based total risk distribution, Hybrid (Eq. (3.6)) is an approximate VaR approach that combines marginal risks using a formula that would apply to an elliptical distribution, Additive adds up the marginal VaRs (Eq. (3.7)), and Normal (Eq. (3.8)) is VaR computed assuming joint normality.

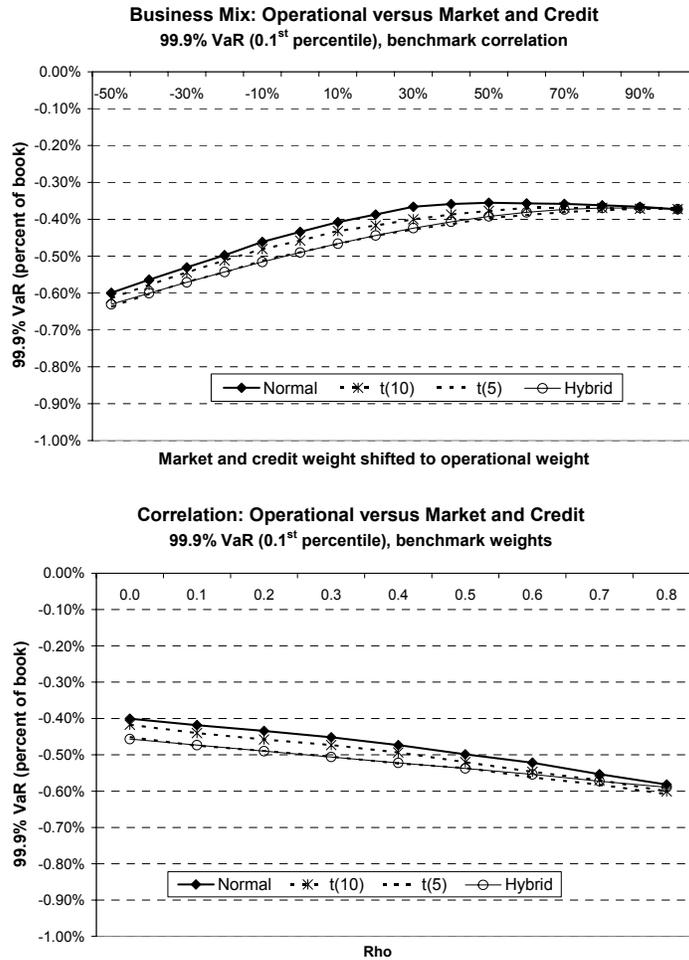


Fig. 7: Sensitivity to choice of copula. The 99.9% VaR (0.1% tail) is computed for joint distributions created using three different copulas: Normal is computed using the normal copula, t(10) and t(5) are computed using the Student-t( $\nu$ ) copula for  $\nu = 5, 10$ . Hybrid (Eq. (3.6)) is an approximate VaR approach that combines marginal risks using a formula that would apply to an elliptical distribution. Impact of business mix (top panel) and inter-risk correlation (bottom panel) on the 99.9% VaR (0.1% tail) of the total risk distribution (analogous to Fig. 6 above).

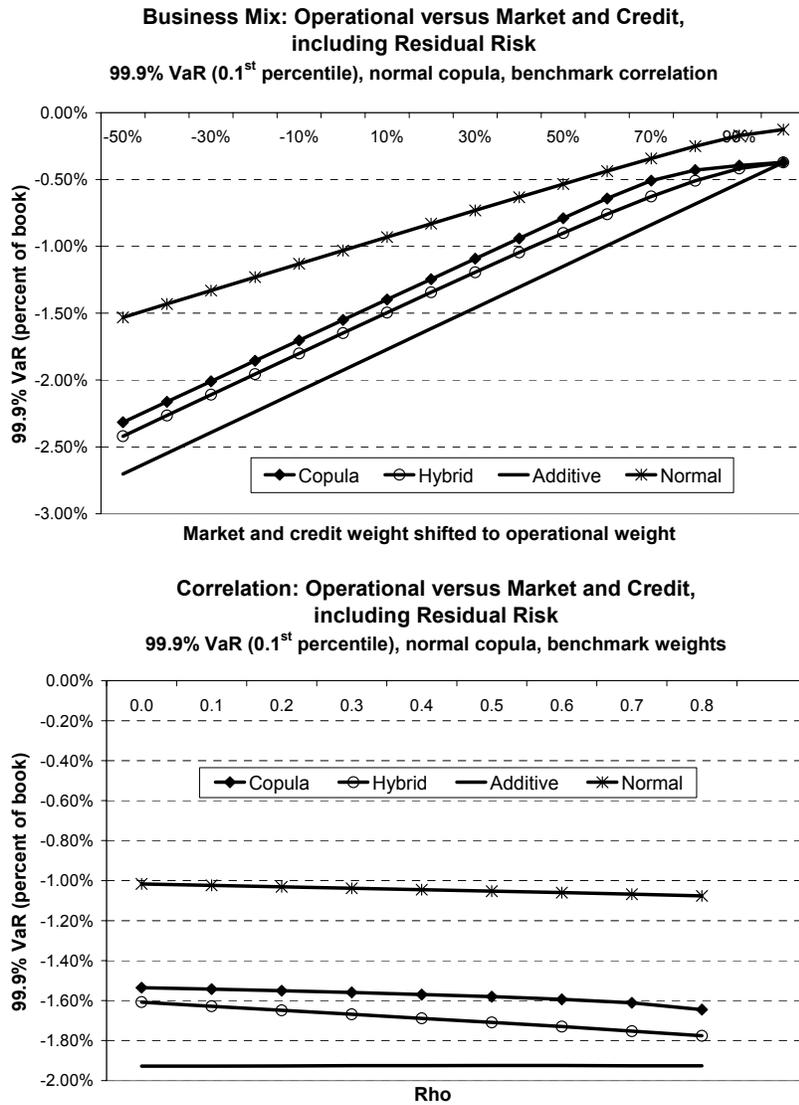


Fig. 8: Impact of residual risk on the 99.9% VaR (0.1% tail) of the total risk distribution. Top panel: Impact of shift in market and credit weight to operational weight while holding relative market and credit weights constant; analogous to bottom panel of Fig. 4. Bottom panel: impact of changing correlation between operational and the other two risk types while keeping the correlation between market and credit constant at benchmark level of 50%; analogous to bottom panel of Fig. 5. In the legend, Copula is VaR computed using the copula-based total risk distribution, Hybrid (Eq. (3.6)) is an approximate VaR approach that combines marginal risks using a formula that would apply to an elliptical distribution, Additive adds up the marginal VaRs (Eq. (3.7)), and Normal (Eq. (3.8)) is VaR computed assuming joint normality.