

Federal Reserve Bank of New York
Staff Reports

A General Approach to Integrated Risk Management
with Skewed, Fat-Tailed Risks

Joshua V. Rosenberg
Til Schuermann

Staff Report no. 185
May 2004

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

A General Approach to Integrated Risk Management with Skewed, Fat-Tailed Risks

Joshua V. Rosenberg and Til Schuermann

Federal Reserve Bank of New York Staff Reports, no. 185

May 2004

JEL classification: G10, G20, G28, C16

Abstract

The goal of integrated risk management in a financial institution is to measure and manage risk and capital across a range of diverse business activities. This requires an approach for aggregating risk types (market, credit, and operational) whose distributional shapes vary considerably. In this paper, we use the method of copulas to construct the joint risk distribution for a typical large, internationally active bank. This technique allows us to incorporate realistic marginal distributions 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 find that given a risk type, total risk is more sensitive to differences in business mix or risk weights than it is to differences in inter-risk correlations. A complex relationship between volatility and fat tails exists in determining the total risk: whether they offset or reinforce each other will depend on the setting. The choice of copula (normal versus student-t), which determines the level of tail dependence, has a more modest effect on risk. We then compare the copula-based method with several conventional approaches to computing risk, each of which may be thought of as an approximation. One easily implemented approximation, which uses empirical correlations and quantile estimates, tracks the copula approach surprisingly well. In contrast, the additive approximation, which assumes no diversification benefit, typically overestimates risk by about 30 to 40 percent.

Key words: market risk, credit risk, operational risk, risk diversification, copula

The authors thank Franklin Allen, Ugur Koyluoglu, Andrew Patton, Marc Saidenberg, Daniel Straumann, and participants in seminars at the Federal Reserve Bank of New York and Rutgers University for their helpful comments, and Megan Todd for her able research assistance. Rosenberg: Research and Market Analysis Group, Federal Reserve Bank of New York (e-mail: joshua.rosenberg@ny.frb.org). Schuermann: Research and Market Analysis Group, Federal Reserve Bank of New York (e-mail: til.schuermann@ny.frb.org). The views expressed in the paper are those of the authors and do not necessarily reflect the position 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.² 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.³ 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 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.⁴

Some risks are better characterized and measured (e.g. market risk) than others (e.g. 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 fat-tails. At the

² Virtually all of the large, multinational financial institutions operating around the world today operate in at least two of the three sectors, which makes 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).”

³ 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). With the anticipation of the New Basel Accord (BCBS, 2001) operational risk is added so that each of these risks will carry a regulatory capital charge.

⁴ 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.

Federal Reserve Bank of New York
Staff Reports

A General Approach to Integrated Risk Management
with Skewed, Fat-Tailed Risk

Joshua V. Rosenberg
Til Schuermann

Staff Report no. 185
May 2004

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

A General Approach to Integrated Risk Management with Skewed, Fat-Tailed Risk

Joshua V. Rosenberg and Til Schuermann

Federal Reserve Bank of New York Staff Reports, no. 185

May 2004

JEL classification: G10, G20, G28, C16

Abstract

The goal of integrated risk management in a financial institution is to measure and manage risk and capital across a range of diverse business activities. This requires an approach for aggregating risk types (market, credit, and operational) whose distributional shapes vary considerably. In this paper, we use the method of copulas to construct the joint risk distribution for a typical large, internationally active bank. This technique allows us to incorporate realistic marginal distributions 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 find that given a risk type, total risk is more sensitive to differences in business mix or risk weights than it is to differences in inter-risk correlations. A complex relationship between volatility and fat tails exists in determining the total risk: whether they offset or reinforce each other will depend on the setting. The choice of copula (normal versus student-t), which determines the level of tail dependence, has a more modest effect on risk. We then compare the copula-based method with several conventional approaches to computing risk, each of which may be thought of as an approximation. One easily implemented approximation, which uses empirical correlations and quantile estimates, tracks the copula approach surprisingly well. In contrast, the additive approximation, which assumes no diversification benefit, typically overestimates risk by about 30-40 percent

Key words: market risk, credit risk, operational risk, risk diversification, copula

The authors thank Franklin Allen, Ugur Koyluoglu, Andrew Patton, Marc Saidenberg, Daniel Straumann, and participants in seminars at the Federal Reserve Bank of New York and Rutgers University for their helpful comments, and Megan Todd for her able research assistance. Rosenberg: Research and Market Analysis Group, Federal Reserve Bank of New York (e-mail: joshua.rosenberg@ny.frb.org). Schuermann: Research and Market Analysis Group, Federal Reserve Bank of New York (e-mail: til.schuermann@ny.frb.org). The views expressed in the paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

same time, our methodology can be implemented when there is limited information about inter-risk dependencies, as in the case where 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. We estimate market, credit, and operational risk distributions using a combination of data from regulatory reports, market data and vendor data. 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. Banking also holds particular interest at the moment because of the intense policy debate surrounding the New Basel Capital Accord (BCBS, 2001), and the incorporation into regulation of a new risk type (operational) which has substantially different characteristics than market or credit risk.⁵

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. An overview of the relevant literature can be found in Section 2. Specifically, we examine the sensitivity of risk estimates to business mix, dependence structure, risk measure and estimation method. Few papers conduct such sensitivity analysis, and those that do focus on just one or two dimensions, and the risk densities themselves are not estimated from empirical data. The studies that model the risks based on empirical data do not conduct any sensitivity analysis. Moreover, the results of these studies are difficult to generalize, since their data is 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

⁵ That banks and insurers are actively wrestling with this issue is shown quite clearly in Joint Forum (2003).

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 document a complex relationship between volatility and fat-tails in determining the 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. In contrast, 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 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 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. In Section 3 we discuss risk measurement (VaR, ES) and risk aggregation. 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 bank management, 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.⁶ Further analytical results in this vein are presented in Schönbucher (2002).⁷ Other application of copulas for portfolio risk measurement include Bouyé (2001), Bouyé et. al. (2001), Longin and Solnick

⁶ 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.

⁷ 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 Durrleman, Nikeghbali, and Roncalli (2000) and Fermanian and Scaillet (2003) analyze copula estimation and testing methods.

(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. Wang (1999, 2002) 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 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 (2002) 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 pairwise. Ward and Lee (2002) use pairwise roll-ups with a Gaussian copula, while Dimakos and Aas (2002) decompose the joint risk distribution into a set of conditional probabilities and impose sufficient conditional independence so that only pairwise 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. 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.

More accurate risk aggregation methods may also be used to attribute the risk contributions back to their marginal sources. This is, of course, crucial to capital allocation across business lines. Although clearly an important issue, it is beyond the scope of this paper. Koyluoglu and Stoker (2002) show clearly just how complex this attribution problem can be, and their work suggests possible avenues for future application of our technique.

Industry studies such as cited in Kuritzkes (2002) and KSW report ex-post risk contributions from individual risk types. KSW report that, on average, about 20% of economic capital, as computed by banks' own internal models, can be attributed to market risk, about 55% to credit risk and the remainder, about 25%, to operational risk. Note that this remainder includes business risk, which is outside the scope of operational risk as defined by the New Basel Accord and as considered in our study. Also, these ex post risk contributions are different from the ex-ante weights implied by business mix.

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 the amount 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. 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 $(1-\mathbf{a})\%$ VaR estimate at time t for a k -period ahead return as $\text{VaR}_t(k, \mathbf{a})$, so that

$$(3.1) \quad \Pr(r_{t+k} < \text{VaR}_t(k, \mathbf{a})) = \mathbf{a}.$$

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. What standards should such a summary measure meet and how do the regulatory approaches currently in use measure up to these standards?

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 (greater returns come with greater risk), 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, 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, 2001). Crouhy, Galai and Mark (2001, p. 253) point out that this situation is common in credit risk with concentrated portfolios with large single exposures.

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

$$(3.2) \quad \text{ES}_t(\mathbf{a}, k) = E[r_{t+k} \mid r_{t+k} \leq \text{VaR}_t(\mathbf{a}, k)].$$

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 measuring the risk of a portfolio is Markowitz’s 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 $r_p = w_x r_x + w_y r_y + w_z r_z$ with mean $\mathbf{m}_p = w_x \mathbf{m}_x + w_y \mathbf{m}_y + w_z \mathbf{m}_z$ and variance

$$(3.3) \quad \mathbf{s}_p^2 = w_x^2 \mathbf{s}_x^2 + w_y^2 \mathbf{s}_y^2 + w_z^2 \mathbf{s}_z^2 + 2w_x w_y \mathbf{s}_{x,y} + 2w_x w_z \mathbf{s}_{x,z} + 2w_y w_z \mathbf{s}_{y,z},$$

where \mathbf{s}_i^2 is the variance of the i^{th} return and $\mathbf{s}_{i,j}$ is the covariance between return i and j . Let

$F_p^{-1}(\mathbf{a})$ be the \mathbf{a} -quantile of the standardized portfolio return distribution. Then, the portfolio

VaR can then be written as (see, for instance, Bradley and Taqqu, 2002)

$$(3.4) \quad \text{VaR}_p(\mathbf{a}) = \mathbf{m}_p + \mathbf{s}_p F_p^{-1}(\mathbf{a}) .$$

Solving for $F_p^{-1}(\mathbf{a})$ in equation (3.4) also provides the definition of the standardized quantile, $(\text{VaR}_p(\mathbf{a}) - \mathbf{m}_p) / \mathbf{s}_p$, which is the quantile measured in terms of the number of standard deviations from the mean.

Substituting the portfolio volatility equation (3.3) into the VaR equation (3.4), we have

(3.5)

$$\begin{aligned} \text{VaR}_p(\mathbf{a}) &= \mathbf{m}_p + \sqrt{w_x^2 \mathbf{s}_x^2 + w_y^2 \mathbf{s}_y^2 + w_z^2 \mathbf{s}_z^2 + 2w_x w_y \mathbf{s}_{x,y} + 2w_x w_z \mathbf{s}_{x,z} + 2w_y w_z \mathbf{s}_{y,z}} [F_p^{-1}(\mathbf{a})] \\ &= \mathbf{m}_p + \sqrt{w_x^2 \mathbf{s}_x^2 [F_p^{-1}(\mathbf{a})]^2 + w_y^2 \mathbf{s}_y^2 [F_p^{-1}(\mathbf{a})]^2 + w_z^2 \mathbf{s}_z^2 [F_p^{-1}(\mathbf{a})]^2 + 2w_x w_y \mathbf{s}_{x,y} [F_p^{-1}(\mathbf{a})]^2 + \dots} . \end{aligned}$$

Equation (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 CDF of the portfolio return density. We can think of the quantile $F_p^{-1}(\mathbf{a})$ 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. $F_p^{-1}(\mathbf{a}) = F_x^{-1}(\mathbf{a}) = F_y^{-1}(\mathbf{a}) = F_z^{-1}(\mathbf{a})$. For example, the family of elliptical distributions, of which the normal is a member, satisfies this condition. More broadly, this equality will hold when the portfolio density and the marginal risk densities come from the same density family. Given this equality, we may write portfolio VaR, which we denote H-VaR (H for hybrid), as

$$\begin{aligned}
(3.6) \quad \text{H-VaR}_p(\mathbf{a}) &= \mathbf{m}_p + \sqrt{w_x^2 \mathbf{s}_x^2 [F_x^{-1}(\mathbf{a})]^2 + w_y^2 \mathbf{s}_y^2 [F_y^{-1}(\mathbf{a})]^2 + w_z^2 \mathbf{s}_z^2 [F_z^{-1}(\mathbf{a})]^2} \\
&\quad + 2w_x w_y \mathbf{s}_{x,y} [F_x^{-1}(\mathbf{a})][F_y^{-1}(\mathbf{a})] + \dots \\
&= \mathbf{m}_p + \sqrt{w_x^2 [\text{VaR}_x(\mathbf{a}) - \mathbf{m}_x]^2 + w_y^2 [\text{VaR}_y(\mathbf{a}) - \mathbf{m}_y]^2 + w_z^2 [\text{VaR}_z(\mathbf{a}) - \mathbf{m}_z]^2} \\
&\quad + 2w_x w_y [(\text{VaR}_x(\mathbf{a}) - \mathbf{m}_x)(\text{VaR}_y(\mathbf{a}) - \mathbf{m}_y)] + \dots
\end{aligned}$$

Equation (3.6) says that portfolio VaR can be computed using the same formula as portfolio volatility, where VaR minus the mean replaces each volatility. If we calculate H-VaR when the marginals come from different density families, then some volatilities may 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.

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

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

When the risks are perfectly correlated ($\mathbf{r}_{i,j} = 1 \forall i, j$ where $\mathbf{r}_{i,j}$ is the correlation between i and j), equation (3.6) becomes:

$$\begin{aligned}
(3.7) \quad \text{Add-VaR}_p(\mathbf{a}) &= \mathbf{m}_p + w_x (\text{VaR}_x(\mathbf{a}) - \mathbf{m}_x) + w_y (\text{VaR}_y(\mathbf{a}) - \mathbf{m}_y) + w_z (\text{VaR}_z(\mathbf{a}) - \mathbf{m}_z) \\
&= w_x \text{VaR}_x(\mathbf{a}) + w_y \text{VaR}_y(\mathbf{a}) + w_z \text{VaR}_z(\mathbf{a})
\end{aligned}$$

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.⁸

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

$$(3.8) \quad \text{N-VaR}_p(\mathbf{a}) = \mathbf{m}_p + \sqrt{w_x^2 \mathbf{s}_x^2 [\Phi^{-1}(\mathbf{a})]^2 + w_y^2 \mathbf{s}_y^2 [\Phi^{-1}(\mathbf{a})]^2 + w_z^2 \mathbf{s}_z^2 [\Phi^{-1}(\mathbf{a})]^2 + 2w_x w_y \mathbf{s}_{x,y} [\Phi^{-1}(\mathbf{a})][\Phi^{-1}(\mathbf{a})] + \dots}$$

It is clear that N-VaR will be accurate only when the joint risk density is multivariate normal. N-VaR is most likely to fail when one or more marginal densities exhibit significant negative skewness or excess kurtosis. In that case, the normal quantile underestimates the actual quantile of the marginal density.

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 density 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 $F_p^{-1}(\mathbf{a}) \neq F_x^{-1}(\mathbf{a}) \neq F_y^{-1}(\mathbf{a}) \neq F_z^{-1}(\mathbf{a})$, we

⁸ Note that Add-VaR assumes an elliptical underlying distribution (the normal is a member of this family) and perfect correlation. If the actual distributions are not members of this family, then Add-VaR may not provide an

are left with the problem of computing (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 the joint return distribution. This can be used to calculate the portfolio return distribution and quantiles, since the portfolio returns are weighted averages of the individual returns.

4. Copulas

Copulas provide important theoretical insights and practical applications in multivariate modeling.⁹ 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 (mean, standard deviation, skewness, and kurtosis) are entirely determined by the marginals.

Copulas can be useful for combining risks when the marginal distributions are estimated individually. This is sometimes referred to as obtaining a joint density with “fixed marginals.” Using a copula, marginals that are initially estimated separately can then be combined in a joint density that preserves the characteristics of the marginals.

Copulas can also be used to obtain more realistic multivariate densities than the normal. For example, the normal dependence relation can be preserved using a normal copula, but

upper bound for VaR, which would be an example of the failure of sub-additivity.

⁹ An excellent reference is Nelsen (1999).

marginals can be entirely general (e.g. a normal copula with one Weibull marginal and one normal marginal). Or, normal marginals can be used with an alternative dependence structure such as that given by a Student-t or an Archimedian copula.

4.1. Representation of a joint density using a copula

One of the fundamental results concerning copulas is known as Sklar's Theorem. It states that any joint CDF can be written in terms of a copula and marginal CDFs. This representation shows that it is possible to separately specify the dependence between two variables and the marginal densities of each variable. The standard formulation is in terms of cumulative distribution functions. Specifically,

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

where $C(u,v)$ is the copula, F_x and F_y are the marginal CDFs, and F_{xy} is the joint CDF.

Alternatively there is the density representation:

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

where $c(u,v) = \left(\frac{\partial^2}{\partial u \partial v}\right) C(u,v)$ is the copula density.

In Equation (4.1), all of the marginal information is captured in $F_x(x)$ and $F_y(y)$, and all of the dependence information is summarized by $C(u,v)$. The copula itself is a joint CDF 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 the scaling of the variables. For example, the copula for X and Y is the same as the copula for $\ln(X)$ and $\exp(Y)$.

4.2. Constructing a joint density using a copula

The copula for any multivariate density function can be obtained using *the method of inversion*. This technique factors out the effects of the marginal densities on the dependence relation by substituting the arguments of the joint density with the quantile functions.

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

where $C(u, v)$ is the copula, F_x^{-1} and F_y^{-1} are the marginal quantile functions, $F_{x,y}$ is the joint CDF, and u and v are cumulative probabilities.

For example, we can determine the normal copula using the method of inversion. If $\Phi(x, y; \mathbf{r})$ is a bivariate standard normal cumulative density function and $\Phi^{-1}(u)$ and $\Phi^{-1}(v)$ are standard normal quantile functions, then the normal copula is simply $\Phi(\Phi^{-1}(u), \Phi^{-1}(v); \mathbf{r})$. The bivariate normal copula has a single parameter, the correlation coefficient \mathbf{r} .

A joint density with given marginals and a given copula can be created by plugging-in the marginal CDF's into the copula function:

$$(3.4) \quad F_{a,b}(a, b) = C(F_a(a), F_b(b)),$$

where $C(F_a(a), F_b(b))$ is the copula for x and y , F_a and F_b are the marginal CDFs, and $F_{a,b}$ is the joint CDF with the given marginals and copula.

So, to create a bivariate density with given marginals and a *normal* copula, one simply inserts the marginal CDFs into the normal copula as $\Phi(F_x^{-1}(u), F_y^{-1}(v); \mathbf{r})$. It is worth noting that the correlation between the risks is not equal to the rho parameter in the copula. The non-normality of the marginal densities drives a wedge between the parameter \mathbf{r} and the Pearson correlation coefficient. In fact, 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 density written in terms of a copula and marginal CDFs. For clarity, we illustrate how to simulate random variables A and B from a joint distribution generated using a bivariate normal copula and marginal CDF's given by F_x and F_y .

First, two uncorrelated standard normal vectors are generated and transformed to correlated standard normal vectors using the lower triangular Cholesky factor.¹⁰ Then, we have $X, Y \sim \Phi(x, y; \mathbf{r})$. X and Y are transformed into realizations from a normal copula by applying the normal CDF to each vector [$U = \Phi(X)$, $V = \Phi(Y)$] so that $U, V \sim \Phi(\Phi^{-1}(u), \Phi^{-1}(v); \mathbf{r})$. Finally, A and B are obtained by applying the inverse CDF for each marginal to U and V [$A = F_x^{-1}(U)$, $B = F_y^{-1}(V)$].

5. Marginal risk distributions: market, credit and operational

In this section we consider the problem of modeling the marginal risk-type distributions for a typical large, internationally active bank, the kind of bank regulators have in mind when

¹⁰ If the two uncorrelated standard normal random variables are Z_1 and Z_2 , then $X = Z_1$ and $Y = \mathbf{r}Z_1 + \sqrt{(1 - \mathbf{r}^2)}Z_2$. This method is easily extended to simulate Student-t variates with n degrees of freedom. The Student-t realization is obtained by multiplying each normal pair (X_i, Y_i) by $\sqrt{v/s_i}$ where s_i is drawn from a chi-square distribution with n

designing capital regulation (BCBS, 2001). 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 “the risk of direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events.”¹¹ This list is hardly exhaustive. For instance, underwriting risk is typically incurred by 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 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. We use this data to arrive at typical size and business mix characteristics. For example, 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

degrees of freedom. Then, the CDF for a Student-t(n) is applied instead of the inverse normal CDF to create a Student-t copula.

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%). 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) and operational risk (like custody).

Second is the units in which the risk type is measured and reported. In order 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 in credit risk, the output may be a loss or return distribution, typically in terms of percentage of exposure. In both cases we need to know the exposure at risk, e.g. trading or lending assets, in order to compute dollars at risk.

Finally, operational risk is not yet part of the regulatory framework – but will be under the New Basel Accord – so banks do not currently report operational risk in any form. Data for this risk type is only recently becoming available, and we will rely on external data 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 of the bank are subject to this risk type.

¹¹ Basel Committee on Banking Supervision (2001), §547.

The overall or portfolio 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. For market risk, it is trading assets plus liabilities (the trading book may contain both long and short positions), for credit risk it is lending assets, and for operational risk it is total assets; the “total book” is then just the sum of these three. The benchmark bank will thus be made up of 4.1% market risk-related activities, 26.4% credit risk- and 69.5% operational risk-related activities, where the proportions are relative to “total book.”

Ideally we would like to model the return due to each risk-type as a function of observable risk factors such as equity returns or interest rate volatility. For bank i , returns from risk-type $j = \{\text{market, credit, operating}\}$ at time t would be driven by a set of observable risk factors $\mathbf{x}_{j,t}$ as in

$$(5.1) \quad r_{i,j,t} = \mathbf{a}_{i,j} + \mathbf{B}'_j \mathbf{x}_{j,t} + \mathbf{e}_{i,j,t}.$$

In equation (5.1) bank returns attributable to risk type j are said to be captured by the set of observable risk factors collected in $\mathbf{x}_{j,t}$. Naturally we are not able to capture all of the relevant risks, but so long as the omitted factors which appear in $\mathbf{e}_{i,j,t}$ are not correlated with observable factors, we suffer only loss of efficiency but no bias in estimating \mathbf{B}'_j . A similar argument can be made for GARCH-type dynamics in the errors.

Such a bottom-up approach would incorporate dependence across risks through the observable risk factors. This is easily done for market and credit risk but not for operational risk for the simple reason that it is far from clear which observable risk factors might drive this risk type. Hence for operational risk we work with the unconditional distribution, the details of which are provided below in Section 5.5. As a result we combine this bottom-up approach for

market and credit risk with a top-down approach for combining all three risks in the sense of imposing inter-risk dependence only via the copula.

One of the advantages of proceeding along the lines of (5.1) is that the time series available for bank data, $r_{i,j,t}$, is much shorter (1994-2002) than what is available for the risk factors $\mathbf{x}_{j,t}$ (for us, 1974-2002, though it could be longer in principle). So long as we may plausibly assume that the risk sensitivities \mathbf{b}_j have not changed significantly over time, we can generate a longer time series of bank-typical returns $\hat{r}_{j,t}$ (notice no bank-specific i -subscript) which then form the basis of our marginal risk distribution. In this way we benefit from a much richer history of risk factor movements than what we have seen in just the last nine years.

We are interested primarily in characterizing the higher moments (two through four) of those distributions as they influence risk. The average or expected return for each risk-type is ambiguously estimated, especially for operational risk where all we have are loss (negative return) outcomes. Thus we tie each marginal distribution to the overall market Sharpe ratio to account for expected gains associated with exposure to each risk type. Given the volatility of the marginal distribution for, say, credit activities, we can impute the mean of that distribution given the Sharpe ratio. We use the Sharpe ratio of annual returns for the S&P500 from 1974 to 2002, which is 0.45.

5.3. Market risk

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 most of the 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 plus liabilities, on annualized¹² returns and volatilities of risk factors from a broad set of asset classes in the capital markets.

Specifically we use the returns and volatilities of the S&P500, changes in the interest rate level (specifically, the 3-month treasury rate), slope (specifically 10Y minus 3M treasury rate) and volatility of the 10Y rate, and returns and volatilities of the trade weighted exchange rate. This isolates the trading return volatility that is driven by this set of risk factors which might be part of a bank's own market risk model. Our aim is to focus on the variability associated with risk factors and hence filter out variability due to fees and commissions.

The return on market risk-related activities for bank i at time t , $r_{i,mk,t}$, is a function of several market risk factors $\mathbf{x}_{mk,t}$ over the sample period 1994Q1 to 2002Q4 and a bank specific fixed effect:

$$(5.2) \quad r_{i,mk,t} = \mathbf{a}_{i,mk} + \mathbf{B}'_{mk} \mathbf{x}_{mk,t} + \mathbf{e}_{i,mk,t}$$

Results of this regression by OLS are provided in Table 1, with the fixed effects omitted for brevity. The regression is significant with an adjusted R-squared of 59%, and the signs of the regression coefficients are broadly in line with our expectations.

Equity returns, foreign exchange returns, and equity volatility are positively associated with trading returns. Interest rate volatility is negatively correlated with trading returns, which could reflect exposure from writing interest rate options or holding long positions in assets with negative convexity such as mortgages. A steepening of the yield curve has a positive impact on

¹² The annualized volatilities are the standard deviation of daily returns over one year scaled up by the number of trading days, namely by $\sqrt{252}$, and the annualized returns are quarterly returns times four.

market risk-related returns as does a decline in short-term interest rates but this coefficient is statistically insignificant.

5.4. Credit risk

Most credit risk related activities arise 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. The relevant credit risk factors are taken to be the overall firm default rate, measured by Moody's issuer weighted default rate on all corporate bonds, and the credit spread between Moody's Baa (seasoned) and the 10Y constant maturity Treasury rate. Thus, similar to the market return relation defined above in (5.2), we may define credit returns $r_{i,cr,t}$ as

$$(5.3) \quad r_{i,cr,t} = \mathbf{a}_{i,cr} + \mathbf{B}'_{cr} \mathbf{x}_{c,r,t} + \mathbf{e}_{i,cr,t}.$$

For consistency we use the same sample period as used in the market risk estimation above, namely 1994Q1 to 2002Q4 and include firm-specific fixed effects. The results are presented in Table 2. As expected, credit returns decrease when the overall default rate and the credit spread increase; the adjusted R-squared for the regression is 64%.

5.5. Operational risk

For our empirical work, we adopt the definition of operational risk set forth by the Basel Committee.¹³ Namely, these are losses due to failure of internal processes, people, systems and

¹³ See footnote 11.

external events. An alternative “residual” measure of operational risk is sometimes constructed by taking total risk net of market and credit risk.¹⁴

For our analysis of operational risk, we use results from de De Fournouville, DeJesus-Rueff, Jordan and Rosengren (2003), hereafter FJJR. They argue that the 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.” They estimate a size-bias corrected exponential density of log-losses with an exponential parameter of 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 dataset. We take the mid-point of their range to obtain a daily probability of an operational loss event of 65/365.

We construct the annual operational loss distribution for a 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. For each day when an operation 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 annual losses.¹⁵

¹⁴ Kuritzkes (2002), for instance, contrasts “event” risk, which conforms more closely with 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 this residual “business” risk. If we take these figures at face value it would imply that we are missing about 15-18% of “total” risk.

¹⁵ 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 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.

5.6. Marginal risk distributions

Given fitted market, credit and operational risk-related returns, we can now generate the inverse marginal CDFs needed as inputs to the copula to generate the total risk distribution. For market and credit risk, we use the longer time series of data from the risk factors to generate a series of returns that might be typical for a large, internationally active bank, $\{\hat{r}_{j,t}\}_{t=1974Q1}^{2002Q4}$. For operational risk, we use the 200,000 realizations created using the technique described in the previous section.

For market and credit risk, we then fit a parametric density function to this projected return series. This procedure, of course, also provides the inverse CDF. Market risk-related projected returns are fitted to a Student-t distribution so that the degrees of freedom (df) matches the observed kurtosis.¹⁶ This turns out to be $df = 11$, with $\mathbf{m} = 0.29\%$ and $\mathbf{s} = 0.63\%$. For credit risk-related projected returns, we fit a three-parameter Weibull distribution using maximum likelihood estimation. The parameter values are $\mathbf{q} = -2.24$, $\mathbf{s} = 2.24$ and $c = 709.78$. Standard goodness of fit tests (Kolmogorov-Smirnov, Cramer-von Mises) can not reject this specification.

In the case of operational risk, it is difficult to find a parametric family that adequately approximates the simulated distribution. Instead, we choose to fit a cubic spline to the empirical quantile function. The empirical quantile function defines the relationship between each sorted loss realization and its rank in the dataset (divided by the number of observations). We use this spline estimate as our inverse CDF for operational risk.

¹⁶ Using the Jarque-Bera statistic, normality is rejected at the 5% (but not 10%) level.

The results for 200,000 simulated bank-years are presented numerically in Table 3 (columns two and three) and visually in Figure 1. All marginal densities are drawn on the same scale on the x-axis (returns). Note that means are imputed from the standard deviations implied by a Sharpe ratio of 0.45. Market risk has highest volatility ($\mathbf{s}_{mk} = 0.64\%$) and thinnest tails ($\mathbf{k}_{mk} = 3.9$), while operational risk has lowest volatility ($\mathbf{s}_{op} = 0.04\%$) and fattest tails ($\mathbf{k}_{op} = 35.3$). Credit risk is in-between with $\mathbf{s}_{cr} = 0.41\%$ and $\mathbf{k}_{cr} = 5.2$. The market risk distribution is symmetric while the credit risk distribution is moderately left-skewed at -1.1 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 (0.04%), but both skewness (-4.5) and especially kurtosis (35.3) are quite high. Thus, typically operational exposure 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).

6. Aggregating the risks

To aggregate the three risk types, we need to know their marginal distributions, specify a copula, assign correlations and have a view on the relative weight or contribution each marginal would have on the joint risk distribution. First, however, we have to make sure that the horizons over which the risks are measured and modeled are harmonized.

Credit risk is typically calibrated to a one-year horizon, as is operational risk. The New Basel Accord, which focuses on these two risk types, is also tied to this horizon. Beyond regulation, however, the logic to a one-year horizon is that it typically ties to reported business

line P&Ls, corresponds to the internal capital allocation and budgeting cycle, and is a period in which an institution can access the markets for additional capital.

Market risk modeling lies at the short horizon end of the spectrum for good reason: the horizon over which market risk is measured should reflect the underlying business activity and decision horizon, such as the frequency of trading (Joint Forum, 2001). The requirement under the market risk amendment to the Basel Accord (BCBS, 1996) is that the market risk capital charge be sufficient to cover a ten-day holding period. Typically banks measure market risk, e.g. using VaR, at a daily frequency, and the conversion from daily to ten-day (time scaling) is done using the root- t rule: multiply the daily standard deviation by the square root of the horizon. Thus for a ten-day horizon, the daily VaR estimates are increased by $\sqrt{10} \approx 3.16$.¹⁷ This logic can be extended to any other horizon, such as one year. There are about 252 trading days in the year, resulting in a multiple of $\sqrt{252} \approx 15.9$. We follow this procedure so that in the end, all risk types are pegged to a one-year horizon.¹⁸

The normal copula is a useful benchmark since we can consider the multivariate normal distribution as a base case for combining risks. It is also commonly seen in the literature (for instance, Li, 2000, Ward and Lee, 2002). We know that the marginals can deviate substantially from normality, so a first move away from multivariate normality is to use more realistic marginals with a normal copula. More recently, the Student-t copula has become prominent (see, for instance, Frey and McNeil, 2001, and Schönbucher, 2002) because it can capture “tail dependence,” which is controlled by the degrees of freedom parameter. In contrast, a normal

¹⁷ Diebold, Hickman, Inoue and Schuermann (1998) show that such scaling is inappropriate (it is only strictly appropriate for *iid* data) and may overestimate volatility at longer horizons.

¹⁸ It turns out that regulatory capital is actually a multiple of 10-day VaR; the multiple ranges from 3 to 4. Stahl (1997), using Chebychev’s inequality, shows that scaling this range is consistent with accounting for distributional

copula has tail independence. We report 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

Our benchmark large, internationally active bank is defined by the median risk exposures from our sample of 17 banks. This drives the weighting of market, credit and operational risk-related distributions to arrive at the joint distribution. Those weights were determined in Section 5.2 to be 4.1% for market risk-related activities, 26.4% credit risk and 69.5% operational risk related activities.

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 may have less exposure to traditional sources of operational risk such as processing or custody related activities. A “processing bank,” conversely, would likely to 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 own 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 4. For our benchmark institution we take a midpoint of those values such that the correlation between market and credit is 50%, and the correlation of each of those two risk types with operational risk is 20%.¹⁹

uncertainty. Specifically, since we don't know the true distribution, we should inflate the volatility to get to the desired quantile. Stahl (1997) shows that for 95% VaR this should be 2.7, and for 99% VaR it should be 4.3.

¹⁹ 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 theoretical

6.2. Results

The total risk/return distribution using the benchmark weightings and correlations with a normal copula is displayed graphically in Figure 2 and numerically in the last column of Table 3. The overall volatility is 0.13% on total book (defined as the sum of trading assets plus liabilities, lending assets and total assets), and, using a Sharpe ratio of 0.45, the mean is 0.07%. The overall distribution inherits features from the marginals, and so the skewness, at -1.0 , of the credit and operational risk densities is reflected here, as is the kurtosis at 4.8.

We start our discussion by focusing on the copula-based results across the two sets of experiments, varying business mix (Figure 3 and Table 5) and inter-risk correlations (Figure 4 and Table 6). The line labeled Copula uses a normal copula with the benchmark inter-risk correlations, and the next three are approximations. Following the legend, the first is the total risk using a normal copula. The second is the hybrid approach as described in equation (3.6), which is based on an elliptical density formula for portfolio VaR. The third is “Additive,” the VaR obtained by simply adding up the VaRs of the marginals. Finally, the line “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 Figure 3 we see the impact of a change in business mix on the 0.1st percentile of the total risk/return distribution under different calculations of VaR of the joint risk distribution given the default correlations specified above in Section 6.1.²⁰ Table 5 reports the standardized quantiles (volatility multiples), volatilities and Copula-VaRs for the total risk distribution as

values’ that are used for the copula. For our analysis across business mix with fixed correlations, the market versus credit correlation at $\rho = 0.50$ is equal to 0.48 and the market/credit versus operational correlation at $\rho = 0.20$ is 0.16. See also Table 7.

²⁰ This percentile is commonly used in industry and conforms to the tolerance level of the New Basel Accord.

business mix shifts. These quantiles are useful in answering how many sigmas (of size given in the same row), relative to the mean, are needed to reach the 0.1st percentile (99.9% VaR).

The top panel in Figure 3 displays the impact of changing the share of book devoted to market risk-related activities vs. credit book while holding operational risk book constant at its benchmark level of 69.5%. Thus, the leftmost set of points corresponds to 0% market and 30.5% credit risk, the next is 3.0% market and 27.5% credit, and at the other extreme we have 30.5% market and 0% credit. Naturally, the extreme points are rather unrealistic, but in-between we may view these weights as reasonably reflecting different business mixes from, for example, little to more intensive trading activity.

Initially, at 100% credit exposure, the volatility of the total risk distribution is 0.13% and the 0.1% standardized quantile is -4.90 standard deviations from the mean (Table 5, left panel). As holdings are shifted to market exposure, VaR at first declines. The initial decline in tail thickness (market has thinner tails than credit) outpaces the volatility increase (market has higher volatility than credit). At market risk weights larger than 50%, the volatility increase dominates giving the VaR curve a parabolic shape. VaR ranges from low of -0.51% to a high of -0.65% with the lowest risk achieved at equal weights of market and credit exposure.

The bottom panel of Figure 3 and the right panel of Table 5 show the impact of a *change* in share of book devoted to operational risk-related activities, while holding the relative proportion of the market and credit book constant. Thus, for the bottom panel, the leftmost points represent the benchmark weighting, 4.1% market, 26.4% credit and 69.5% operational. For the second set of points, a shift of 10% from market and credit towards operational would result in (same order) 3.7%, 23.7% and 72.6%, and so on. Since it is conceptually not reasonable

to assign a weight of zero to operational risk activities, we adopt this approach based on deviations from the benchmark weighting.

As assets are shifted from market and credit to operational exposure, there is an increase in tail thickness (Table 5, right panel). The standardized quantile at baseline holdings is -5.15 standard deviations, and it increases to -9.64 standard deviations for a bank with mostly operational exposure. At the same time, as operational exposure grows, the volatility of total risk declines (from 0.13% to 0.04%). The net effect is that the VaR curve first declines as volatility decreases more quickly than tail-thickness increases. At operational exposures greater than 50% , the curve flattens reflecting the offsetting effects of decreased volatility and fatter tails.

6.2.2. Impact of inter-risk correlation

In Figure 4 we explore the impact of inter-risk correlation on the 0.1^{st} percentile (99.9% VaR) of the total risk/return distribution using the benchmark weightings. Its companion table (Table 6) shows standardized quantiles, volatilities, Copula-VaRs and the approximation error for N-VaR and Add-VaR relative to Copula-VaR. The top panels of Figure 4 and Table 6 display the impact of changing correlation between market and credit risk while correlation of both risk types to 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% .²¹

²¹ 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 r_{ij} for any risk type i or j to be arbitrarily high. See, also footnote 19.

As expected, higher levels of market and credit risk correlation lead to higher levels of VaR (Figure 4, top panel).²² VaR increases from -0.50% at a correlation of zero to -0.57% for correlation of 0.9. In Table 6, top panel, we see that higher market and credit correlation leads to higher volatility and slightly thinner tails.

Looking now to the bottom panel in Figure 4 and Table 6, VaR increases as we raise the correlation between operational risk and market/credit risk. From initial level of -0.51% at zero correlation, VaR increases to -0.72% at correlation of 0.8. In this case, higher correlation results in higher volatility and fatter tails. 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.

We also compute the approximation error of N-VaR and Add-VaR relative to Copula-VaR (last two columns Table 6). Naturally the results are specific to these set of weights and the specification of the normal copula. Still, for mid-level correlations we see approximation errors of 30% to 50%. The average relative difference of N-VaR to Copula-VaR is about 41% and of Add-VaR to Copula-VaR about 47%.

Note that this is not the same as a diversification benefit which would be relative to Add-VaR. Since Add-VaR is never smaller than Copula-VaR, the corresponding diversification benefits are smaller than the approximation error, namely around 25% to 35%. These results are somewhat higher than reported by the 15% reported by KSW, although they use the 99% VaR level and assume joint normality. Our results are closer to the 20+% reported by Dimakos and Aas (2002), who also use a normal copula and look at the far tail (99.97% in their case).

²² Embrechts, McNeil and Straumann (2002) give examples of increased rho leading to lower VaR.

Deviation from normality in the marginal distributions drives a wedge between \mathbf{r} and the correlation coefficient. This is illustrated in Table 7 where we display the relation between \mathbf{r} and the correlation coefficient from simulated returns using a normal copula. The first column corresponds to the bottom panel of Figure 4 while the second column corresponds to the top panel in Figure 4.

While the market and credit correlation (column 2) is always quite close to the specified rho (within 0.02), the operational versus market/credit correlation is much lower (column 1). As rho is varied from 0.10 to 0.80, market and credit correlation moves from 0.10 to 0.78 but operational correlation moves from 0.08 to 0.68. Even though operational 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 \mathbf{r} 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, we see several patterns. First, Add-VaR always provides the largest VaR estimate, and N-VaR always provides the smallest estimate. Add-VaR is the largest, since it fixes the correlation matrix at unity, when in fact the empirical correlations are much lower. N-VaR is smallest since it uses the lowest standardized quantiles for the marginals, i.e. $\Phi^{-1}(0.001) = -3.09$ versus actual quantiles of -3.64 (market), -4.87 (credit), -9.09 (operational). For this reason, Add-VaR and N-VaR are quite different from the copula measure, indicating that they are both significantly biased. The hybrid approach (H-VaR) tracks Copula-VaR well, but it is also upwardly biased (and thus conservative).

In the business mix experiment (Figure 3), Add-VaR exhibits no diversification benefits, of course, and is monotonically decreasing as we shift to a lower volatility asset. This pattern is more pronounced in the bottom panel, where activities are shifted towards operational risk and away from both market and credit. If we assume the joint distribution is normal (N-VaR), then overall risk increases modestly as increased variance is somewhat offset by the diversification effect. Indeed, in the top panel N-VaR initially underestimates VaR by about 0.25% (percentage points), but comes closer as market risk exposure grows (which is the distribution that is closest to normal). H-VaR is within 0.07% for all weights and has the same parabolic shape as the copula measurement.

In the bottom panel of Figure 3, as the business mix shifts entirely towards operational risk intensive activities, 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 (Figure 4), it is striking how close H-VaR is to Copula-VaR, though always more conservative. Since the normal distribution is not as fat-tailed as the actual distributions, N-VaR consistently underestimates Copula-VaR. N-VaR also changes linearly with ρ , while Copula-VaR and H-VaR reflect non-linear interactions of the tails with the correlation. Once again, H-VaR is a reasonably good (but biased) approximation with a maximum deviation from Copula-VaR of 0.06%.

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 fails to capture any of these features and systematically over-estimates the risk of the portfolio

distribution. The closest approximation to the copula VaR is H-VaR, though it always overestimates risk.

6.3. Robustness tests

Since VaR is not a coherent risk measure, we also repeated the set of experiments using the expected shortfall measure as defined in equation (3.2). The general conclusions are no different from the VaR measure. By way of illustration, we present in Figure 5 two experiments focusing on the risk sensitivity to operational exposure and operational correlation.

The top panel of Figure 5 shows the impact on ES at 0.1% of a *change* in share of book devoted to operational risk-related activities, while holding the relative proportion of the market and credit book constant. This is analogous to VaR displayed the bottom panel in Figure 3. Similarly, the correlation experiment in the bottom panel of Figure 5 examines the impact on ES at 0.1% 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 Figure 4.

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 Figure 5, copula ES follows the same pattern as copula VaR (bottom panel of Figure 4), but is consistently about 0.10% larger (again, percentage points). Remarkably, even for expected shortfall, the hybrid ES comes quite close to the copula ES (within 0.07%), albeit always with a positive bias.

To measure VaR sensitivity to the choice of copula, we repeat the experiments forming joint distributions using Student-t copulas with 5 and 10 degrees of freedom instead of the normal. As noted earlier, lower degrees of freedom for the t-copula increase tail dependence and

increase tail risk. The normal copula represents the limiting case of infinite degrees of freedom and tail independence.

The results are presented in Figure 6. In the bottom panel, we see that copula VaR is typically larger by about 0.03% for a Student-t(10) copula than for a normal copula. The difference is 0.06% for a Student-t(5) copula compared to a normal copula. These are fairly representative differences for all experiments.²³ We find the average proportional increase in VaR is about 11% when the normal copula is replaced with Student-t(5). This is much less than the sensitivity to changes in operational exposure or correlation, but fairly close to the sensitivity of 11% resulting from variation in market and credit correlation.

In the case of strong tail dependence (e.g., Student-t(5) copula), the hybrid approach provides a fairly accurate estimate of copula VaR. The maximum approximation error is 0.03% (percentage points) for the business mix experiment and 0.01% for the correlation experiment ($\rho < 0.7$). While hybrid approach overestimates risk under tail independence (normal copula), it is more accurate when the risks exhibit tail dependence.

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 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 and are presented in Table 8. 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 -0.72% at a correlation of 0.8 and the smallest VaR

²³ For the expected shortfall measure the results are very similar.

of -0.51% at a correlation of zero. Using the average VaR of -0.60%, the relative sensitivity is 36%.

Unsurprisingly we see the greatest sensitivities are to the approximation choice relative to the copula approach. Assumptions of joint normality or simply adding up the marginal risks result in approximation errors averaging 43% but ranging from 13% to 64%. The second largest sensitivities are due to operational risk exposure (42%) and operational risk correlation (36%). In contrast, the market versus credit exposure experiment results in sensitivity of 26% and the market versus credit correlation experiment results in sensitivity of 14%. Thus, assumptions about operational exposures and especially correlations seem to be more important for accurate risk estimates than assumptions about relative exposures or correlations of market and credit risk. Finally, the choice of copula seems to have a more modest impact; the average sensitivity across all experiments was only 13%.

7. Final Comments

How would one aggregate market, credit, and operational risk distributions to arrive at an overall integrated joint 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 the method of 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 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 may 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 approximates 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.

References

1. Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1999, "Coherent Measures of Risk," *Mathematical Finance* 9, 203-228.
2. Artzner, P., Delbaen, F., Eber, J.M., Heath, D., 1997, "Thinking coherently," *Risk* 10 (11), 68-71.
3. Basel Committee on Banking Supervision, 2001, *The New Basel Capital Accord*, <http://www.bis.org/publ/bcbsca03.pdf>, January.
4. Basel Committee on Banking Supervision, 1996, "Amendment to the Capital Accord to Incorporate Market Risks (No. 24)", <http://www.bis.org/publ/bcbs24a.htm>, January.
5. Basel Committee on Banking Supervision, 1988, "Internal Convergence of Capital Measurement and Capital Standards, <http://www.bis.org/publ/bcbs04A.pdf>, July.
6. Berkowitz, Jeremy, 2001, "Testing the Accuracy of Density Forecasts, applications to Risk Management," *Journal of Business and Economic Statistics*, 19, 465-474.
7. Bouyé, Eric, 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.
8. Bouyé, Eric, V. Durrleman, A. Nikeghbali G. Riboulet and T. Roncalli, 2001, "Copulas: An Open Field for Risk Management," Working paper, Group de Recherche Operationelle, Credit Lyonnais, (March). Available at <http://gro.creditlyonnais.fr/content/wp/copula-rm.pdf>.
9. Bradley, B.O. and M.S. Taqqu, 2002, "Financial Risk and Heavy Tails," in S.T. Rachev (ed.), *Heavy-Tailed Distributions in Finance*, Amsterdam, NL: North Holland.
10. Christoffersen, Peter and Francis X. Diebold, 2000, "How Relevant is Volatility Forecasting for Financial Risk Management?" *Review of Economics and Statistics* 82, 12-22.
11. Crouhy, Michel, Dan Galai and Robert Mark, 2001, *Risk Management*, New York: McGraw Hill.
12. De Fountnouvelle, Patrick, Virginia DeJesus-Rueff, John Jordan and Eric Rosengren, 2003, "Using Loss Data to Quantify Operational Risk", Federal Reserve Bank of Boston Working Paper.
13. Diebold, Francis X., T.A. Gunther and A.S. Tay, 1998, "Evaluating Density Forecasts", *International Economic Review* 39, 863-883.
14. Diebold, Francis X., Andrew Hickman, Atsushi Inoue and Til Schuermann, 1998, "Scale Models", *Risk* 11, 104-107.
15. Dimakos, Xenia K. and Kjersti Aas, 2002, "Integrated Risk Modeling", The Norwegian Computing Center Working Paper.
16. Durrleman, V., A. Nikeghbali and T. Roncalli, 2000, "Which copula is the right one?," Working paper, Group de Recherche Operationelle, Credit Lyonnais, (August). Available at <http://gro.creditlyonnais.fr/content/wp/copula-choice.pdf>.

17. Embrechts Paul, F. Lindskog, and Alexander McNeil, 2003, "Modeling Dependence with Copulas and Applications to Risk Management," in *Handbook of Heavy Tailed Distributions in Finance*, ed. S. Rachev, Rotterdam, NL: Elsevier, Chapter 8, 329-384.
18. Embrechts, Paul, Alexander McNeil and Daniel Straumann, 2002, "Correlation and Dependence in Risk Management: Properties and Pitfalls," In: *Risk Management: Value at Risk and Beyond*, ed. M.A.H. Dempster, Cambridge, UK: Cambridge University Press, 176-223.
19. Embrechts, Paul, Alexander McNeil and Daniel Straumann, 1999, "Correlation: Pitfalls and Alternatives," *Risk* (May), 69-71.
20. Fermanian, Jean-David and Olivier Scaillet, 2003, "Nonparametric Estimation of Copulas for Time Series," *Journal of Risk* 5 (4), 25-54.
21. Frey, Rüdiger and Alexander J. McNeil, 2001, "Modelling Dependent Defaults," available at <http://e-collection.ethbib.ethz.ch/show?type=bericht&nr=273>.
22. Glasserman, Paul, Philip Heidelberger and Perwez Shahabuddin, 2002, "Portfolio Value-at-Risk with Heavy-Tailed Risk Factors," *Mathematical Finance* 12 (July), 239-270.
23. Hirtle, Beverly, 2003, "What Market Risk Capital Reporting Tells Us About Bank Risk", *Federal Reserve Bank of New York Economic Policy Review* 9 (3), 37-54.
24. Johnson, N.L., S. Kotz and N. Balakrishnan, 1994, *Discrete Multivariate Distributions*, New York: Wiley.
25. Joint Forum, 2003, *Trends in Risk Integration and Aggregation*, August; available at <http://www.bis.org/publ/joint07.pdf>.
26. Joint Forum, 2001, *Risk Management Practices and Regulatory Capital, Cross-Sectoral Comparison*, November; available at <http://www.bis.org/publ/joint03.pdf>.
27. Jorion, Philippe, 2001, *Value at Risk: The Benchmark for Managing Financial Risk* 2nd Edition, New York: McGraw Hill.
28. Koyluoglu, H. Ugur and James Stoker, 2002, "Honour your Contribution," *Risk*, April, 90-94.
29. Kuritzkes, Andrew, Til Schuermann and Scott M. Weiner, 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*, 141-194.
30. Kuritzkes, Andrew, 2002, "Operational Risk Capital: A Problem of Definition," *Journal of Risk Finance*, Fall, 1-10.
31. Li, D. X., 2000, "On Default Correlation: A Copula Function Approach," *Journal of Fixed Income*, March, 43 - 54.
32. Longin, F. and B. Solnik, 2001, "Extreme Correlation of International Equity Markets," *Journal of Finance* 56 (2), 649-676.
33. Nelsen, Roger B., 1999, *An Introduction to Copulas*. New York: Springer-Verlag.
34. Patton, Andrew, "Modelling Time-Varying Exchange Rate Dependence Using the Conditional Copula," July 2002, UCSD Discussion Paper 01-09. Available at <http://www.econ.ucsd.edu/papers/files/2001-09.pdf>.

35. Poon, Ser-Huang, Michael Rockinger and Jonathan Tawn, 2004, "Extreme Value Dependence in Financial Markets: Diagnostics, Models and Financial Implications," *Review of Financial Studies* 17 (2), 581-610.
36. Rosenberg, Joshua, 2003, "Nonparametric pricing of multivariate contingent claims," *Journal of Derivatives*, 10 (3), 9-26.
37. Schönbucher, Philipp J., 2002, "Taken to the Limit: Simple and Not-so-Simple Loan Loss Distributions," Working Paper, Department of Statistics, Bonn Univ; available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=378640.
38. Stahl, Gerhard, 1997, "Three Cheers," *Risk* 10 (5), 67-69.
39. 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.
40. Wang, Shaun S., 2002, "A Set of New Methods and Tools for Enterprise Risk Capital Management and Portfolio Optimization", CAS Forum Summer 2002; available at <http://www.casact.org/pubs/forum/02sforum/02sf043.pdf>.
41. Wang, Shaun S., 1999, "Aggregation of Correlated Risk Portfolios: Models & Algorithms", Casual Actuarial Society Theory of Risk Working Paper; available at <http://casact.org/cofor/Wang.pdf>.
42. Ward, Lisa S. and David H. Lee, 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>.

Tables

Variable	Parameter Estimate	Standard Error	P-value
constant	0.0306	0.0082	0.0002
Equity return	0.0057	0.0036	0.1156
Equity volatility	0.0369	0.0188	0.0496
Δ interest rate	-0.0203	0.0823	0.8049
Δ slope	0.1217	0.0691	0.0786
Interest rate volatility	-1.1888	0.5859	0.0429
FX returns	0.0256	0.0086	0.0031
FX volatility	-0.0296	0.0976	0.7615
F-statistic	38.5		<0.0001
\bar{R}^2	0.585		
number of obs.	612		

Table 1: Market risk regression with bank fixed effects. Dependent variable is as annual trading revenues divided by the sum of trading assets and trading liabilities. Data is from Y9-C forms 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 (Source: Datastream), Equity volatility is the standard deviation of daily returns over one year scaled up by the number of trading days, namely by $\sqrt{252}$. Δ interest rate is the change in the 3M treasury rate, Δ slope is the change in the slope of the yield curve, the slope being defined as the difference between the 10Y and 3M treasury rate, and Interest rate volatility is the volatility of the change in the 10Y treasury rate (Source: Board of Governors of the Federal Reserve System/FRED II, Release: H.15 Selected Interest Rates). FX returns and FX volatility is the annualized log-return and volatility of the trade-weighted currency index. (Source: Board of Governors of the Federal Reserve System H.10 release Foreign exchange rates).

Variable	Parameter Estimate	Standard Error	P-value
constant	0.0687	0.0027	<0.001
Default rate	-0.3414	0.0567	<0.001
Credit spread	-0.3997	0.1815	0.0281
\bar{R}^2	0.640		
number of obs.	612		

Table 2: Credit risk regression with bank fixed effects. Dependent variable is net interest income less provisions. Data is from Y9-C forms over the period 1994Q1-2002Q4 using the bank holding company sample discussed in Section 5.2. Explanatory variables are Moody's issuer-weighted default rate on all corporate bonds and the credit spread between Moody's Baa (seasoned) and the 10Y constant maturity Treasury rate. Sources: FRED, Moody's.

	Market Risk	Credit Risk	Operational Risk	Total Risk
Mean	0.26%	0.18%	0.02%	0.07%
Std. dev.	0.64%	0.41%	0.04%	0.13%
Skewness	0.0	-1.1	-4.5	-1.0
Kurtosis	3.9	5.2	35.3	4.8
0.1 percentile	-2.05%	-1.79%	-0.37%	-0.54%
0.5 percentile	-1.52%	-1.28%	-0.27%	-0.38%
1 st percentile	-1.30%	-1.08%	-0.16%	-0.32%

Table 3: Characteristics of simulated marginal risk distributions for bank-years = 200,000 using methods described in Section 5.6. All values are annualized. Mean imputed from standard deviation using a Sharpe Ratio of 0.45. The total risk distribution is generated by simulating the weighted average of market (4.1%), credit (26.4%), and operational (69.5%) returns from a joint density defined using a normal copulas with the benchmark correlations (market vs. credit = 0.5, market/credit vs. operational = 0.2).

	Credit	Market/ALM	Operational & Other
Credit	100%	30%	20%
		30%	44%
		80%	40%
Market/ALM	30%		20%
	30%	100%	13%
	80%		40%
Operational & Other	20%	20%	
	44%	13%	100%
	40%	40%	

Table 4: Inter-risk correlations (Table 4 from Kuritzkes, Schuermann and Weiner, 2003).
Range of inter-risk correlations reported in academic and industry studies.

Holdings	Market vs. Credit			Operational vs. Market & Credit		
	Standardized Quantile	Volatility	VaR (99.9%)	Standardized Quantile	Volatility	VaR (99.9%)
0%	-4.90	0.13%	-0.58%	-5.15	0.13%	-0.54%
10%	-4.74	0.13%	-0.55%	-5.17	0.12%	-0.50%
20%	-4.60	0.13%	-0.53%	-5.25	0.11%	-0.46%
30%	-4.38	0.13%	-0.51%	-5.37	0.10%	-0.42%
40%	-4.22	0.14%	-0.51%	-5.58	0.09%	-0.40%
50%	-4.02	0.15%	-0.51%	-5.95	0.08%	-0.38%
60%	-3.90	0.15%	-0.52%	-6.53	0.07%	-0.37%
70%	-3.83	0.16%	-0.55%	-7.30	0.06%	-0.37%
80%	-3.78	0.18%	-0.58%	-8.24	0.05%	-0.37%
90%	-3.73	0.19%	-0.61%	-9.08	0.05%	-0.37%
100%	-3.72	0.20%	-0.65%	-9.64	0.04%	-0.37%

Table 5: Standardized quantile (volatility multiple) and volatility for 0.1% tail of the total risk/return distribution as business mix changes. Each row represents the share of book devoted to market risk-related activities vs. credit book while holding operational risk book constant. The standardized quantile for the market, credit and operational risk/returns distributions are -3.64, -4.87 and -9.09 respectively. The marginal volatilities in turn are (same order) 0.64%, 0.41% and 0.04%. See also Figure 3.

Market vs. Credit

Rho	Standardized Quantile	Volatility	VaR (99.9%)	Approximation Error	
				N-VaR to Copula	Add-VaR to Copula
0	-4.76	0.12%	-0.50%	40%	64%
0.1	-4.76	0.12%	-0.51%	40%	61%
0.2	-4.77	0.12%	-0.52%	40%	57%
0.3	-4.76	0.13%	-0.53%	40%	54%
0.4	-4.73	0.13%	-0.53%	39%	52%
0.5	-4.71	0.13%	-0.54%	39%	51%
0.6	-4.71	0.13%	-0.55%	39%	49%
0.7	-4.65	0.13%	-0.55%	38%	47%
0.8	-4.63	0.14%	-0.56%	37%	45%
0.9	-4.64	0.14%	-0.57%	38%	43%

Operational vs. Market & Credit

Rho	Standardized Quantile	Volatility	VaR (99.9%)	Approximation Error	
				N-VaR to Copula	Add-VaR to Copula
0	-4.63	0.13%	-0.51%	38%	61%
0.1	-4.67	0.13%	-0.52%	39%	56%
0.2	-4.71	0.13%	-0.54%	39%	51%
0.3	-4.77	0.13%	-0.56%	40%	45%
0.4	-4.88	0.14%	-0.59%	41%	39%
0.5	-5.02	0.14%	-0.62%	43%	31%
0.6	-5.15	0.14%	-0.65%	44%	25%
0.7	-5.29	0.14%	-0.68%	46%	19%
0.8	-5.46	0.15%	-0.72%	48%	13%

Table 6: Standardized quantile (volatility multiple) and volatility for 0.1% tail of the total risk/return distribution as inter-risk correlation changes. One 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 21 and Table 7 below. Both sets of results: default weights are 4.1% for market risk-related activities, 26.4% credit risk and 69.5% operational risk related activities. Approximation error is the absolute relative difference to copula-VaR. This is not the same as diversification benefit which would be relative to Add-VaR. Note that the standardized quantile for the market, credit and operational risk/returns distributions are -3.64 , -4.87 and -9.09 respectively. The marginal volatilities in turn are (same order) 0.64%, 0.41% and 0.04%. See also Figure 4.

rho	corr (operational, {market, credit})	corr (market, credit)
0.1	0.08	0.09
0.2	0.16	0.19
0.3	0.24	0.29
0.4	0.33	0.38
0.5	0.41	0.48
0.6	0.50	0.58
0.7	0.59	0.68
0.8	0.68	0.78

Table 7: Relation between rho and the correlation coefficient from simulated returns using a normal copula. First column corresponds to the bottom panel of Figure 4 where the correlation between market and credit stays fixed at the benchmark level of $\mathbf{r} = 50\%$ as it changes with respect to operational risk and the other two risk types. The second column corresponds to the top panel in Figure 4 where the correlation between market and credit is allowed to change while each of their correlation with operational risk remains fixed at the benchmark level of 20%.

Factor	Average Risk Sensitivity
Approximation choice*	43%
Operational risk exposure	42%
Operational risk correlation	36%
Market vs. credit risk exposure	26%
Market vs. credit risk correlation	14%
Copula choice**	13%

Table 8: Risk sensitivities. Range of VaRs across different factors and experiments in descending order. Risk sensitivity = $\frac{|\text{Max VaR} - \text{Min VaR}|}{\text{Average VaR}}$.

* Average difference for normal copula versus N-VaR and Add-VaR.

** Average difference for normal copula versus Student-t(5) copula.

Figures




Figure 1: Kernel densities of simulated marginal risk distributions for bank-years = 200,000 using methods described in Section 5.6. See also Table 3. For market and credit risk, projected returns are fitted to Student- t (for market) and Weibull (for credit) distributions.




Figure 2: Kernel density plot of the total risk/return distribution. Default weights are 4.1% for market risk-related activities, 26.4% credit risk and 69.5% operational risk related activities using a normal copula with the default inter-risk correlations of 50% for market and credit, and 20% of those two risk types with operational risk.





Figure 3 : Impact of business mix; 0.1% tail of the total risk/return distribution. Top panel: Impact of change in share of book devoted to market risk-related activities vs. credit book while holding operational risk book constant. Bottom panel: Impact of change in share of book devoted to operational risk-related activities while holding the sum of the market and credit book constant. Both panels: normal copula with the default inter-risk correlations of 50% for market and credit, and 20% of those two risk types with operational risk.





Figure 4: Impact of inter-risk correlation; 0.1% tail of the total risk/return distribution. 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: default weights are 4.1% for market risk-related activities, 26.4% credit risk and 69.5% operational risk related activities.





Figure 5: Impact of business mix and inter-risk correlation on the 0.1% expected shortfall (ES) of the total risk/return distribution. Top panel: Impact of change in share of book devoted to operational risk-related activities while holding the sum of the market and credit book constant; analogous to bottom panel of Figure 3. 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 Figure 4.




Figure 6: Sensitivity to choice of copula: normal, Student- $t(n)$ for $n = 5, 10$. Impact of business mix and inter-risk correlation on the 0.1% VaR of the total risk/return distribution (analogous to Figure 5 above). Top panel: Impact of change in share of book devoted to operational risk-related activities while holding the sum of the market and credit book constant; analogous to bottom panel of. 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%.