

Risk Factor Contributions in Portfolio Credit Risk Models *

Dan Rosen[†] David Saunders[‡]

January 31, 2009

Abstract

Determining contributions to overall portfolio risk is an important topic in financial risk management. At the level of positions (instruments and subportfolios), this problem has been well studied, and a significant theory has been built, in particular around the calculation of marginal contributions. We consider the problem of determining the contributions to portfolio risk of **risk factors**, rather than positions. This problem cannot be addressed through an immediate extension of the techniques employed for position contributions, since, in general, the portfolio loss is not a linear function of the risk factors. We employ the Hoeffding decomposition of the loss random variable into a sum of terms depending on the factors. This decomposition restores linearity, at the cost of including terms that arise from the joint effect of more than one factor. The resulting cross-factor terms provide useful information to risk managers, since the terms in the Hoeffding decomposition can be viewed as best quadratic hedges of the portfolio loss involving instruments of increasing complexity. We illustrate the technique on multi-factor models of portfolio credit risk, where systematic factors may represent different industries, geographical sectors, etc.

*First version April 23, 2008. Previous title: Risk Contributions of Systematic Factors in Portfolio Credit Risk Models. Thanks to Dirk Tasche for helpful comments. D. Saunders is partially supported by an NSERC Discovery Grant. The authors are grateful to Fitch Solutions for financial support and data provided for this research.

[†]Fields Institute for Research in Mathematical Sciences and R^2 Financial Technologies, drosen@fields.utoronto.ca

[‡]Department of Statistics and Actuarial Science, University of Waterloo, dsaunders@uwaterloo.ca.

1 Introduction

Decomposing portfolio risk into its different sources is a fundamental problem in financial risk management. Once a risk measure has been selected, and the risk of a portfolio has been calculated, a natural question to ask is: *where does this risk come from?* Specifically, the risk manager may be interested in understanding contributions to portfolio risk of two types:

- **Positions:** individual instruments, counterparties and sub-portfolios
- **Risk Factors:** various systematic or idiosyncratic factors affecting portfolio losses (e.g. market risk factors such as interest rates, exchange rates, equity volatilities etc., macro-economic, geographic, or industry factors affecting market or credit risk)

The development of methodologies for the first type of risk contribution is of great importance for hedging, capital allocation, performance measurement and portfolio optimization. In this case, the portfolio losses can be written as the sum of losses of individual positions (instruments, counterparties, subportfolios). For such sums, there is an established theory for *additive* risk contributions based on the concept of *marginal contributions*, sometimes referred to as *Euler allocation*. The latter name derives from the fact that the formula for contributions to risk measures which are homogeneous functions of degree one of the portfolio weights (standard deviation, VaR, CVaR, etc.) follows directly from Euler's theorem (see Tasche (1999, 2008)). Position risk contributions in general, and the Euler allocation in particular, have received much attention in the recent literature. Koyluoglu and Stoker (2002) and Mausser and Rosen (2007) survey different contribution measures, with the latter reference focusing on applications to portfolio credit risk management. Kalkbrener (2005) considers a set of axioms for risk allocation methods, and demonstrates that they are only satisfied by the Euler allocation. Denault (2001) relates the Euler allocation to cooperative game theory and notions of fairness in allocation, while Tasche (1999) gives an interpretation in terms of optimization and performance measurement. Numerous authors have also considered computational issues arising in the calculation of position risk contributions. See, for example Emmer and Tasche (2005), Glasserman (2005), Glasserman and Li (2005), Huang et al. (2007), Kalkbrener et al. (2004), Merino and Nyfeler (2004), Tchistiakov et

al. (2004), and Mausser and Rosen (2007) and the references therein.

Just as fundamental for risk management, the development of methodologies for contributions of *risk factors* to portfolio risk has received comparatively little attention. In this case, portfolio losses cannot generally be written as a linear function of the individual risk factors. When each position depends (perhaps in a non-linear way) only on a single independent risk factor (or a small subset of risk factors), the problem can be addressed effectively by computing position contributions and transforming them to factor contributions. However, in many problems there are several factors which interact across large parts of the portfolio to drive potential losses, and the standard theory for deriving contributions cannot be directly applied. These factors might be systematic factors representing macro-economic variables, indices, or financial variables. Practical examples where such problems arise include:

- Portfolio credit risk, where multi-factor credit models are common.
- Portfolios with equity options (where equities are modeled using a multi-factor model), foreign exchange options or quanto options.
- Collateralized Debt Obligations and Asset Backed Securities.

Contributions of risk factors are important because they facilitate an understanding of the sources of risk in a portfolio; this is especially important for complex portfolios with many instruments, where individual instrument risk contributions may not be too enlightening, it is also useful in understanding the sources of risk for complicated derivative securities (e.g. portfolio credit derivatives). Furthermore, the current credit crisis has highlighted the need for tools that help us understand better the role of systematic risk factors in credit risk.

Recent papers which consider the problem of factor contributions directly include the following. Tasche (2008) shows how the Euler allocation can be extended to calculate contributions of individual names to CDO tranche losses in a consistent manner and proposes measures for the impact of risk factors in the nonlinear case; Cherny and Madan (2007) study position contributions of conditional losses given the risk factors (see the discussion in section 4 below); Rosen and

Saunders (2008) study the best hedge (in a quadratic sense) among linear combinations of the systematic factors in the context of the Vasicek model of portfolio credit risk; Bonti et al. (2006) conceptualize the risk of credit concentrations as the impact of stress in one or more systematic risk factors on the loss distribution of a credit portfolio.

In this paper, we develop an extension of the Euler allocation that applies to nonlinear functions of a set of risk factors. The technique is based on the Hoeffding decomposition, originally developed for statistical applications (see, for example, van de Vaart (1998) and Sobol (1993)). The intuition behind the methodology is simple: while we cannot write the portfolio loss function as a sum of functions of individual risk factors, the application of the Hoeffding decomposition allows us to express it as a sum of functions of all subsets of risk factors. The standard Euler allocation machinery can then be applied to the new loss decomposition. The price paid for this methodology is that we have to consider contributions not only from single risk factors, but also from the interaction of every possible collection of risk-factors.

The terms in the Hoeffding decomposition also have an important risk management interpretation. They can be thought of as portfolios of increasing complexity which depend on an increasing number of factors. Each term in the decomposition gives the best (quadratic) hedge of the residual risk in the portfolio that can be constructed with derivative securities depending on a given collection of factors, and which cannot be hedged by any subset of that collection (i.e. any previous term in the decomposition).

For a portfolio with k risk factors, the Hoeffding decomposition contains 2^k terms. For a small number of factors, we can compute the contributions of all of the terms explicitly, and then perhaps aggregate them in financially meaningful ways. When the number of factors is large, it may be more practical to compute the contributions for a smaller set, plus a residual contribution. It is often the case that only a handful of risk factors are significant. Alternatively, one may aggregate the factors into a smaller number of subsets and assess their contributions.

While the technique is general and can be applied to portfolio loss contributions in various financial problems (market, credit or operational risk), we develop the methodology in the context of multi-factor

credit portfolio models and derive specific results for calculating systematic factor contributions.¹ In this case, one obtains analytical expressions for conditional expectations of losses given subsets of systematic factors. Thus, the computational complexity of the methodology is equivalent to that of calculating standard position contributions. While there may be analytical expressions for volatility contributions (although perhaps too complex for many practical applications), VaR/CVaR contributions present computational challenges. However, advanced methods developed for position contributions and cited above may be applicable to calculating factor contributions as well.

Several papers have recently employed the Hoeffding decomposition in financial problems, although for different applications; Sobol (1993) has defined variance-based global sensitivity indices (see also Sobol and Kucherenko (2005)); Kucherenko and Shah (2007) consider applications to Monte-Carlo methods for options pricing; Lemieux and L'Ecuyer (2001) show how to use the decomposition to develop selection criteria for Quasi-Monte Carlo point sets, and present an application to derivatives pricing; Baur et al. (2004) use the decomposition to consider the sensitivity of portfolio credit risk models to model risk (mis-specification of model parameters and dependence structure) in a variance-based context.

The remainder of the paper is structured as follows. The second section discusses factor models for credit risk, setting notation and reviewing results that will be used in later sections. The third section briefly surveys the theory of risk contributions for positions, emphasizing marginal risk contributions and the Euler allocation principle. The fourth section presents Hoeffding decompositions and discusses their use to calculate factor contributions to portfolio credit risk, as well as various extensions and practical issues that arise when using the technique. The fifth section presents numerical examples applying the theory of factor contributions from section four to sample portfolios. The sixth section concludes.

¹In particular, we note that even in the context of credit risk, the methodology is not limited to so-called *factor models*, although this is the application on which we focus in this paper. The decomposition can be used to calculate the risk contribution of X_n whenever losses can be written as $L = L(X_1, \dots, X_n, \dots, X_N)$.

2 Factor Models for Credit Risk

We denote the portfolio loss function by L , and assume it can be written as the sum of losses due to individual positions, $L = \sum_{n=1}^N w_n L_n$. The random variable L arises from a factor model if there is a set of random variables $Z = (Z_1, \dots, Z_K)^T$ such that, conditional on Z , the random variables L_1, \dots, L_N are independent. In practical models the number of factors is much smaller than the number of positions in the portfolio, $K \ll N$. The factor structure can often be a significant computational advantage. First, in many cases when N is very large, and each term in the sum defining L is relatively small, we have that L is approximated well by (indeed converges to, as $N \rightarrow \infty$) $L_S = \mathbb{E}[L|Z]$, which is a random variable depending only on the joint distribution of the K systematic factors Z_k , rather than the (often much larger) set of risk factors determining L . Additionally, the factor structure can be of great use in simulating the portfolio loss distribution. For example, when pricing credit derivatives, one often runs a separate simulation on scenarios for the systematic factors, and then computes the distribution of portfolio losses given the systematic factors using conditional independence and a technique for computing convolutions efficiently.² A review with many results on factor models for credit risk can be found in McNeil et al. (2005).

An example of a factor model is the market benchmark Gaussian copula model, which in its single factor variant forms the basis for the credit risk capital charge in the Basel II accord. We briefly review the model to set notation. The credit state of the n th obligor in the portfolio is modelled by a random variable referred to as its creditworthiness index:

$$Y_n = \sum_{k=1}^K \beta_{nk} Z_k + \sigma_n \cdot \varepsilon_n \quad (1)$$

Here, the variables Z_k are i.i.d. standard normal random variables, referred to as the model's *systematic factors*, as they are common to all names in the portfolio. The random variables ε_n , called the idiosyncratic factors, also have the standard normal distribution, and are independent of each other, and of the systematic factors Z_k . The coefficients β_{nk} are the *factor loadings*, and represent the sensitivity

²Since the variables L_n are conditionally independent given Z , the distribution of L given Z is simply a convolution of the conditional distributions of each of the L_n .

of the n th obligor to the k th systematic factor. The residual standard deviations σ_n satisfy:

$$\sigma_n = \sqrt{1 - \sum_{k=1}^K \beta_{nk}^2}$$

which guarantees that each random variable Y_n has a standard normal distribution. The default probabilities are prescribed exogenously, with the probability of default over the time horizon for the n th name denoted by p_n . Obligor n defaults if its creditworthiness index falls below a given level, i.e. if $Y_n \leq H_n$. The threshold H_n is set to match the obligor's default probability, $H_n = \Phi^{-1}(p_n)$, where Φ is the standard cumulative normal distribution function and Φ^{-1} is its inverse.

The total portfolio loss over the time horizon is therefore:

$$L = \sum_{n=1}^N w_n \mathbf{1}_{\{Y_n \leq H_n\}}$$

where w_n is the loss-given-default adjusted exposure of the n th instrument. In general, w_n is a random variable, which may or may not be correlated with either the systematic or idiosyncratic factors, or both. In this paper, we generally assume that the coefficients w_n are constant. The analytic results employed in calculating factor contributions immediately generalize if w_n are random variables with known means that are independent of the systematic factors Z_k . In the case of correlated exposures and creditworthiness indices, the exact assumptions regarding the joint distribution of defaults and losses given default will determine whether the conditional expectations used in this paper are analytically tractable.

Conditional on a known value of the systematic factors $Z = (Z_1, \dots, Z_K)^T$, the default indicators are independent Bernoulli random variables with default probabilities:

$$p_n(Z) = \mathbb{E}[\mathbf{1}_{\{Y_n \leq H_n\}} | Z] = \Phi \left(\frac{H_n - \sum_{k=1}^K \beta_{nk} Z_k}{\sigma_n} \right)$$

So that the *systematic part* of the portfolio loss is defined to be:³

$$L_S = \mathbb{E}[L|Z] = \sum_{n=1}^N w_n p_n(Z)$$

It is well known that for large portfolios, where each loan makes a relatively small contribution to overall portfolio risk, L_S is a good approximation to the total portfolio loss L . For precise statements of this result, see, for example Gordy (2003), or McNeil et al. (2005).

3 Risk Contributions of Instruments and Sub-Portfolios

In this section, we briefly review the theory of risk contributions, with particular emphasis on marginal contributions (also known as the Euler allocation rule). For a more complete discussion of the theory of capital allocation, focusing in particular on credit risk management, see Mausser and Rosen (2007). For a survey of results on the Euler allocation rule, see Tasche (2008), or McNeil et al. (2005).

We consider the total portfolio loss as a sum of the losses of individual positions (instruments or sub-portfolios):

$$L = \sum_{n=1}^N w_n L_n \tag{2}$$

where L_n is the random variable giving the loss per dollar of exposure in instrument n , and w_n is the amount of money invested in position n . The total risk of the portfolio is $\rho(L)$, where ρ is a risk measure mapping random variables to real numbers.

We are interested in defining a measure C_n of the contribution of the n th position to the total portfolio risk. Different methods of calculating risk contributions have been studied for different purposes. We present a brief list of the alternatives that are popular in practice.

Stand-Alone Contributions:

$$C_n = \rho(w_n L_n)$$

³This equation relies on our assumption that loss-given-default adjusted exposures are constant.

The stand-alone contribution of a position is simply its risk if it were held as a portfolio in isolation. It ignores the distributions of all other positions, and as such does not take into account any diversifying or hedging effects resulting from its inclusion in the institution's portfolio. It can be useful in measuring the reduction in risk due to diversification, and in measuring diversification factors for portfolios (see Garcia-Cespedes et al. (2006) and Tasche (2006)). It can also be considered as an upper bound on the contribution to the risk for any reasonable allocation rule. That is, for any allocation rule, we would expect to have $C_n \leq \rho(w_n L_n)$. This condition features in axiomatizations of capital contributions, e.g. Kalkbrener (2005), as well as the interpretation of the Euler allocation rule in terms of the theory of cooperative games, e.g. Denault (2001) or Koyluoglu and Stoker (2002). If the risk measure is *subadditive*⁴, then the sum of the standalone contributions provides an upper bound for the total portfolio risk:

$$\rho(L) \leq \sum_{n=1}^N \rho(w_n L_n)$$

Coherent risk-measures such as expected shortfall are subadditive. It is well known that Value-at-Risk and Economic Capital are not subadditive risk measures, and for them the above inequality can be violated.

Incremental Contributions: The incremental risk contribution of a position is the change in total risk arising from including the position in the portfolio.

$$C_n = \rho(L) - \rho\left(\sum_{m \neq n} w_m L_m\right)$$

This is a useful measure for traders considering adding the position L_n to their portfolio. When w_n is small, it may also be regarded as a finite difference approximation to the marginal risk contribution discussed below. It is typically not the case that incremental contributions of positions add up to the total portfolio risk (however, it should also be noted that this definition of risk contribution is motivated by applications where additivity is not necessarily desirable).

⁴A risk measure ρ is subadditive if for all X, Y , $\rho(X + Y) \leq \rho(X) + \rho(Y)$.

Marginal Contributions (Euler Allocation): We consider a risk measure that is positive homogeneous (i.e. $\rho(\lambda \cdot L) = \lambda \rho(L)$ for $\lambda \geq 0$). This includes measures such as standard deviation (σ_L), Value-at-Risk ($\text{VaR}(L)$), Economic Capital ($EC(L)$) and Conditional Value-at-Risk ($\text{CVaR}(L)$, also known as expected shortfall), and any risk measure satisfying the coherence axioms of Artzner et al. (1999). Under technical differentiability assumptions on ρ , Euler's theorem for positive homogeneous functions immediately implies:

$$\rho(L) = \sum_{n=1}^N C_n$$

where:

$$C_n = w_n \frac{d\rho}{d\varepsilon}(L + \varepsilon L_n)|_{\varepsilon=0} = w_n \frac{\partial \rho(L)}{\partial w_n}(w) \quad (3)$$

The n th term in the sum, C_n is then interpreted as the contribution of the n th position's loss (L_n) to the overall portfolio risk $\rho(L)$.

Explicit formulas for marginal risk contributions are available for some of the most important risk measures. For standard deviation:

$$C_n^\sigma = w_n \frac{\text{cov}(L_n, L)}{\sigma_L} \quad (4)$$

where σ_L is the standard deviation of L . For Value-at-Risk at the confidence level α , subject to technical conditions, Gouriéroux et al. (2001) and Tasche (1999) showed that:

$$C_n^{\text{VaR}} = w_n \mathbb{E}[L_n | L = \text{VaR}_\alpha(L)] \quad (5)$$

Finally, for CVaR, and again subject to technical conditions, Tasche (1999) showed that:

$$C_n^{\text{CVaR}} = w_n \mathbb{E}[L_n | L \geq \text{VaR}_\alpha(L)] \quad (6)$$

4 Hoeffding Decompositions and Systematic Risk Factor Contributions

In this section, we review the definition and elementary properties of the Hoeffding decomposition of a random variable. We then proceed to discuss the application of the Hoeffding decomposition to computing

factor contributions in portfolio credit risk models. Finally, we outline various issues that arise in applying the decomposition in practice.

In statistical applications, the term Hoeffding decomposition⁵ is usually reserved for the situation where the factors are independent. However, the general formula ((9), see below) is valid for correlated factors as well. We begin by discussing the decomposition, the corresponding factor contributions and their financial interpretation in the case of independent factors, and then consider the generalization to correlated factors.

4.1 Hoeffding Decompositions

In this section, we review the concept of the Hoeffding decomposition of a random variable, and give it a financial motivation. A more technical discussion of the decomposition is given in the appendix, and in van der Vaart (1998), sections 11.3-11.4.

We motivate the general decomposition by considering an example with a small number of factors. Suppose that the systematic credit loss L_S of a portfolio is driven by two independent factors, Z_1 and Z_2 . We can write the random variable L_S in the following way:

$$\begin{aligned}
 L_S &= \mathbb{E}[L] \\
 &+ \mathbb{E}[L|Z_1] - \mathbb{E}[L] \\
 &+ \mathbb{E}[L|Z_2] - \mathbb{E}[L] \\
 &+ L_S - (\mathbb{E}[L|Z_1] - \mathbb{E}[L]) - (\mathbb{E}[L|Z_2] - \mathbb{E}[L]) - \mathbb{E}[L] \quad (7)
 \end{aligned}$$

This decomposition is a tautology, but it conveys important financial information. The first, constant, term gives the best hedge of the loss (in the sense of quadratic hedging) that can be constructed using only a risk-free instrument. The second term gives the best hedge of the remaining risk that can be constructed using only instruments that depend on Z_1 , disregarding entirely the risk coming from Z_2 .⁶

⁵The Hoeffding decomposition sometimes goes by other names, including the Moebius decomposition and the Sobol decomposition. The decomposition is an important tool in applied statistics. See Van der Vaart (1998) and the references therein for numerous applications.

⁶We allow hedging using any instrument with a payoff of the form $f(Z_1)$. The problem with payoffs given by linear functions of the factors was considered in Rosen and Saunders (2008).

A similar interpretation is given to the third term (involving Z_2 in isolation, while ignoring the risk coming from Z_1). The final term gives the residual risk that cannot be hedged away with instruments which depend only on an individual risk factor, but must instead be hedged with instruments that depend on factor co-movements.

We now proceed to a more formal definition of the decomposition in the general multi-factor case. Let Z_1, \dots, Z_K be *independent* systematic factors with finite variances, and let $L = g(Z_1, \dots, Z_K)$ have finite variance. The Hoeffding decomposition gives an unique, canonical way of writing L as a sum of *uncorrelated* terms involving conditional expectations of g given sets of the factors Z . In particular, we have:

$$L = \sum_{A \subseteq \{1, \dots, K\}} g_A(Z_j; j \in A) \quad (8)$$

where

$$g_A(Z_j; j \in A) = \sum_{B \subseteq A} (-1)^{|A|-|B|} \cdot \mathbb{E}[L | Z_k, k \in B] \quad (9)$$

The sum (8) runs over all possible subsets $A \subseteq \{1, \dots, K\}$, of the collection of factors. Each term in the decomposition has a financial interpretation. The term $g_A(Z_j; j \in A)$ gives the best hedge (in the quadratic sense) of the residual risk driven by co-movements of the systematic factors $Z_j, j \in A$ that cannot be hedged by considering any smaller subset $B \subset A$ of the factors.

One can think of the decomposition as writing the total portfolio loss in terms of hedges involving instruments of increasing complexity. The first (constant) term $\mathbb{E}[L]$, corresponding to the empty set of factors, gives the best hedge possible using only a risk-free instrument. The ‘first-order’ terms $g_k = \mathbb{E}[L | Z_k] - \mathbb{E}[L]$ hedge the residual risk of the portfolio considering the k th factor in isolation. The ‘second-order’ terms $g_{k,j}$ hedge the remaining residual risk from joint moves in the factors Z_k and Z_j , and so on.

As noted above, the terms in the decomposition can be written more explicitly for small index sets:

$$\begin{aligned} g_\emptyset &= \mathbb{E}[L] \\ g_k &= \mathbb{E}[L | Z_k] - \mathbb{E}[L] \\ g_{k,j} &= \mathbb{E}[L | Z_k, Z_j] - \mathbb{E}[L | Z_k] - \mathbb{E}[L | Z_j] + \mathbb{E}[L] \end{aligned}$$

4.2 Systematic Risk Factor Contributions

In this section, we discuss the generalizations of the risk contribution measures presented for sub-portfolios in section 3 to risk factor contributions. In particular, we use Hoeffding decompositions to extend the theory of marginal risk contributions.

Stand-Alone Contributions: The natural extension of stand-alone contributions to the nonlinear case is:

$$C_k = \rho(\mathbb{E}[L|Z_k]) = \rho(\mathbb{E}[g(Z_1, \dots, Z_K)|Z_k])$$

These measures were studied under the term “factor risk contributions” by Cherny and Madan (2007). Under technical conditions, they show that for coherent, law-invariant risk measures ρ

$$B \subseteq A \Rightarrow \rho(\mathbb{E}[L|Z_k, k \in B]) \leq \rho(\mathbb{E}[L|Z_k, k \in A])$$

Intuitively, considering the risk driven by more factors results in a higher overall measure of portfolio risk. They also consider the case of a portfolio loss variable $L = \sum_n L_n$, where each term in the sum L_n depends on a (possibly multi-dimensional) factor Z_n and suggest approximating portfolio risk by:

$$\rho(L) \approx \sum_n \rho(\mathbb{E}[L_n|Z_n])$$

Incremental Contributions: The natural generalization of incremental risk contributions to nonlinear functions $L = g(Z_1, \dots, Z_K)$ is:

$$C_k = \rho(L_S) - \rho(\mathbb{E}[L|Z_{\sim k}])$$

where $Z_{\sim k} = \{Z_j, j \neq k\}$. This is the difference between the portfolio risk calculated including the systematic factor Z_k and the risk calculated when the k th factor is ignored. The results of Cherny and Madan (2007) cited above imply that in most meaningful cases this incremental risk will be positive.

Marginal Contributions: The Hoeffding decomposition formula (8) writes the loss variable as a sum of a set of random variables. Marginal risk contributions of each of these random variables can then be calculated using the results on the Euler allocation.

$$C_A = \frac{d\rho}{d\varepsilon}(L + \varepsilon g_A(Z_k; k \in A))|_{\varepsilon=0}$$

The interpretation of the contribution of the term g_A to the overall portfolio risk is that it is the residual contribution to risk arising from the interaction of the factors $Z_k, k \in A$ that is not captured by the effects of any subset of these factors. For example, the contribution of $g_k(Z_k)$ measures the impact of the factor Z_k that is not already accounted for in the expected loss. The contribution of $g_{k,j}(Z_k, Z_j)$ is the residual contribution to losses of the joint effect of the factors Z_k and Z_j , that is not already captured by the expected loss $\mathbb{E}[L]$ and the conditional expectations $\mathbb{E}[L|Z_k], \mathbb{E}[L|Z_j]$.

4.3 Application to Portfolio Credit Risk

We note that for the Gaussian copula credit risk model, one can compute each term in the Hoeffding decomposition explicitly. In order to do so, note that the expected loss in this model given a subset of the systematic factors can be computed explicitly. Indeed, simple analytical manipulations reveal that:

$$\begin{aligned} \mathbb{E}[L|Z_k, k \in A] &= \mathbb{E}[L_S|Z_k, k \in A] \\ &= \sum_{n=1}^N w_n \Phi\left(\frac{H_n - \sum_{k \in A} \beta_{nk} Z_k}{\sigma_{n,A}}\right) \end{aligned} \quad (10)$$

where

$$\sigma_{n,A} = \sqrt{1 - \sum_{k \in A} \beta_{n,k}^2}$$

When the decomposition (8) is applied to compute factor contributions to portfolio credit risk (see the examples below), the computations are greatly facilitated by the fact that the conditional expectations $\mathbb{E}[L|Z_k, k \in A] = h(Z_k, k \in A)$ are all known functions of Z with an explicit closed-form given by (10). This dramatically speeds up the calculation of contributions. For example, consider computing the ‘first-order’ contribution of a factor Z_k to CVaR. From the contribution formula (6), we have that this will be given by:

$$C_k^{\text{CVaR}} = \mathbb{E}[g_k(Z_k)|L \geq ES_\alpha] = \mathbb{E}[h_k(Z_k)|L \geq ES_\alpha] - \mathbb{E}[L]$$

Having an explicit formula for $h_k(Z_k) = \mathbb{E}[L|Z_k]$ makes computing the above conditional expectation much easier than if a numerical method had to be used each time h needed to be evaluated (for example when using a Monte-Carlo method to compute the conditional expectation).

Known analytical results (see, for example, Rosen and Saunders (2008)) can be employed to calculate the contributions to standard deviation for each term in the Hoeffding decomposition analytically. The resulting formulas are straightforward to derive, but rather cumbersome, and perhaps of less interest than contributions to Value-at-Risk and Conditional Value-at-Risk. Consequently, in the remaining sections, we focus instead on computing VaR and CVaR contributions.

It is interesting to note that not all factor models result in conditional expectations that are nonlinear functions of the systematic factors. For example, the CreditRisk+ model (see Credit Suisse Financial Products (1997)) assumes that losses for each counterparty have a Poisson mixture distribution with conditional mean given by a linear function of a set of systematic factors having gamma distributions. The conditional expectation of the portfolio loss is then a linear function of the systematic factors, and the standard theory of position contributions can be used to compute risk factor contributions.

4.4 Practical Issues

There are several practical issues that arise with the use of the Hoeffding decomposition for marginal risk factor contributions.

1. It is important to recognize that the formula (8) gives a *decomposition* of the random variable L , rather than an *expansion* with an error estimate. In particular, truncating the decomposition by considering, for example, only the expected loss, $\mathbb{E}[L]$ and the first order effects $g_k(Z_k)$ may not result in a good approximation to the portfolio risk. This sum of first-order terms is called the Hájek projection (van der Vaart (1998), Hájek (1968)). It is the projection of L on the space of all (square-integrable) random variables of the form $\sum_{k=1}^K g_k(Z_k)$. Whether or not it gives a good approximation to the loss random variable L depends on the particular choice of L and the nature of its dependence on the factors Z_k .

As a simple example, consider the case where all the factors Z_k are independent and have mean zero, and $L = \prod_{k=1}^K Z_k$. Then all conditional expectations of the form $\mathbb{E}[L|Z_k, k \in A]$ with $A \neq \{1, \dots, K\}$ are equal to zero. The Hoeffding decom-

position is trivial, with $g_A \equiv 0$ for all A except $\{1, \dots, K\}$ and $g_{\{1, \dots, K\}}(Z) = L$.

2. There are 2^K terms in the decomposition. This limits the applicability of the technique to models with a moderate number of systematic factors (or cases where the decomposition can be truncated after a small number of terms, see the previous point). This situation can be remedied somewhat because the Hoeffding decomposition still works if we consider collections of variables rather than single variables.
3. The contribution of the constant term in the Hoeffding decomposition (8) contains effects from all the factors, and it can be difficult to separate these effects out. For example, consider the case where the systematic loss actually is a linear function of the systematic factors (as is the case for CreditRisk+), or more generally any case where L equals its Hájek projection:⁷

$$L = \sum_{k=1}^K g_k(Z_k) = \sum_{k=1}^K Y_k$$

It seems natural to take the risk contribution of factor k to be that given to Y_k by the Euler allocation rule. However, the Hoeffding decomposition given above consists of $K + 1$ terms, the expected loss $\mathbb{E}[L]$ and the K residual factor terms

$$\begin{aligned} g_k(Z_k) &= \mathbb{E}[L|Z_k] - \mathbb{E}[L] \\ &= Y_k + \sum_{j \neq k} \mathbb{E}[Y_j] - \mathbb{E}[L] \\ &= Y_k - \mathbb{E}[Y_k] \end{aligned}$$

Risk contributions calculated using the Hoeffding rule thus give a contribution from the expected loss, and K contributions from the ‘surprises’ in the transformed factors Y_k . These contributions are different from those calculated for the variables Y_k using the Euler allocation rule. We note that this issue only arises when the random variable L has non-zero mean and the risk-measure depends on the mean (for example, it occurs when considering

⁷This situation may also arise when the portfolio can be divided into sub-portfolios which only depend on one of the systematic factors, see Cherny and Madan (2007) for a related discussion.

risk contributions to VaR and CVaR but not when considering contributions to economic capital and standard deviation, or in general when working with unexpected losses).

4. The numerical difficulties that beset the computation of risk contributions for instruments and sub-portfolios (see, for example, Mausser and Rosen (2007)) carry over to contributions of each term in the Hoeffding decomposition. Techniques that have been applied for calculating position risk contributions in factor models of portfolio credit risk (see Glasserman (2005), Huang et al. (2007), Merino and Nyfeler (2004) and the other references given in the introduction) can potentially be used for computing factor risk contributions as defined here. However, to date we have not run any computational experiments to test this conjecture.

4.5 Interpreting Risk Contributions for Correlated Factors

The Hoeffding decomposition presented above is usually applied to independent factors, in which case the terms in the decomposition are orthogonal (see appendix) and straightforward to interpret as best hedges. The general decomposition formula (8) is still valid for dependent factors.⁸ For example, any random variable $L = g(Z_1, Z_2)$ can be trivially decomposed as:

$$L = \mathbb{E}[L] + (\mathbb{E}[L|Z_1] - \mathbb{E}[L]) + (\mathbb{E}[L|Z_2] - \mathbb{E}[L]) + (L - \mathbb{E}[L|Z_1] - \mathbb{E}[L|Z_2] + \mathbb{E}[L])$$

Indeed, considering dependent factors may be more useful for risk management purposes, where financially meaningful factors are typically correlated and independent (uncorrelated) factors can only be achieved by transforming the data (for example using principal components). However, one must be careful when interpreting the above decomposition for correlated factors, as the following example illustrates.

Let (Z_1, Z_2) have a joint normal distribution with standard normal

⁸We note that, in general, each term in the sum (9) can depend on the joint distribution of the factors.

marginals and correlation ρ . Define the loss random variable:

$$L = w_1 Z_1 + w_2 Z_2$$

with Value-at-Risk:

$$\text{VaR}_\alpha(L) = \sigma_L N^{-1}(\alpha) = \sqrt{w_1^2 + w_2^2 + 2\rho w_1 w_2} \cdot N^{-1}(\alpha)$$

Marginal risk contributions calculated by treating Z_1 and Z_2 as sub-portfolios then give:

$$C_1^E = w_1 \frac{\partial \text{VaR}_\alpha(L)}{\partial w_1} = \frac{w_1^2 + \rho w_1 w_2}{\sigma_L} N^{-1}(\alpha) = w_1 \mathbb{E}[Z_1 | L = \text{VaR}_\alpha(L)]$$

$$C_2^E = w_2 \frac{\partial \text{VaR}_\alpha(L)}{\partial w_2} = \frac{w_2^2 + \rho w_1 w_2}{\sigma_L} N^{-1}(\alpha) = w_2 \mathbb{E}[Z_2 | L = \text{VaR}_\alpha(L)]$$

On the other hand, the Hoeffding decomposition gives:

$$L = \mathbb{E}[L|Z_1] + \mathbb{E}[L|Z_2] + (L - \mathbb{E}[L|Z_1] - \mathbb{E}[L|Z_2])$$

$$= (w_1 + w_2 \rho) Z_1 + (w_1 \rho + w_2) Z_2 - \rho(w_2 Z_1 + w_1 Z_2)$$

This gives:

$$C_1^H = (w_1 + w_2 \rho) \mathbb{E}[Z_1 | L = \text{VaR}_\alpha(L)] = \frac{(w_1 + w_2 \rho)^2}{\sigma_L} N^{-1}(\alpha)$$

$$C_2^H = (\rho w_1 + w_2) \mathbb{E}[Z_2 | L = \text{VaR}_\alpha(L)] = \frac{(\rho w_1 + w_2)^2}{\sigma_L} N^{-1}(\alpha)$$

$$C_{12}^H = -\rho(w_2 \mathbb{E}[Z_1 | L = \text{VaR}_\alpha(L)] + w_1 \mathbb{E}[Z_2 | L = \text{VaR}_\alpha(L)])$$

$$= -\frac{\rho^2 w_1^2 + 2\rho w_1 w_2 + \rho^2 w_2^2}{\sigma_L} N^{-1}(\alpha)$$

When Z_1 and Z_2 are independent, the Euler and Hoeffding contributions coincide. Otherwise, they are different. The Hoeffding decomposition results in terms driven by “pure Z_1 risk” and “pure Z_2 risk”, as well as a term that adjusts for the joint influence of Z_1 and Z_2 . Notice that if $w_1 \cdot w_2 > 0$ then $2\rho^2 w_1 w_2 < 2\rho w_1 w_2$ and

$$\rho^2 w_1^2 + \rho^2 w_2^2 + 2\rho w_1 w_2 > \rho^2 w_1^2 + \rho^2 w_2^2 + 2\rho^2 w_1 w_2 = \rho^2 (w_1 + w_2)^2$$

So in this case, there is a reduction due to ‘risk-factor diversification’. However, if $w_1 \cdot w_2 < 0$ and

$$\rho < \frac{(w_1 + w_2)^2}{w_1^2 + w_2^2} - 1$$

then there is an amplification of risk when both factors are considered.

The financial interpretation of the decomposition clarifies the meaning of the above calculations. Recall that the term $g_1(Z_1) = \mathbb{E}[L|Z_1] = w_1 Z_1 + w_2 \mathbb{E}[Z_2|Z_1]$ represents the best hedge possible with instruments depending solely on Z_1 (here we have simplified the decomposition using the fact that the means of both variables are zero). In the case of independence, $\mathbb{E}[Z_2|Z_1] = 0$, and the Euler contributions are recovered. However, when Z_1 and Z_2 are not independent, *the term $\mathbb{E}[Z_2|Z_1]$ gives the portion of the risk driven by Z_2 that can be hedged with instruments that depend on Z_1 .* A similar effect occurs in hedging Z_1 risk with instruments depending on Z_2 in the computation of the term g_2 . This results in some ‘over-hedging’ owing to the double counting (Z_2 risk is hedged by both g_1 and g_2 , and Z_1 risk is hedged by both g_2 and g_1). The effect of this ‘over-hedging’ is then removed by the term g_{12} .

5 Numerical Examples

We consider two numerical examples calculating marginal systematic factor risk contributions using the Hoeffding decomposition in the Gaussian copula model. The first is a stylized example designed to illustrate important aspects of the method. The second is a systematic risk factor decomposition of the risk in the CDXIG index.

5.1 Two Factor Examples

The examples in this section are based on those in Pykhtin (2004) and Rosen and Saunders (2008). Consider a simple credit risk model with two systematic factors and a portfolio with only two names.⁹ The two names in the portfolio are A and B , their (loss-given-default weighted) exposures are w_A and w_B and their default probabilities are p_A and p_B , with corresponding default thresholds $H_A = \Phi^{-1}(p_A)$ and $H_B = \Phi^{-1}(p_B)$ respectively. The creditworthiness indices of A and B

⁹Since we are ignoring idiosyncratic risk, we could equivalently consider two homogeneous sub-portfolios.

are given by:

$$Y_A = \beta_A Z_1 + \sqrt{1 - \beta_A^2} \cdot \varepsilon_A$$

$$Y_B = \beta_B \lambda Z_1 + \beta_B \sqrt{1 - \lambda^2} Z_2 + \sqrt{1 - \beta_B^2} \cdot \varepsilon_B$$

where Z_1, Z_2 are systematic factors and $\varepsilon_A, \varepsilon_B$ are idiosyncratic factors; all factors are independent and have the standard normal distribution. The parameter $\lambda \in [0, 1]$ controls the degree to which name B depends on the first or second systematic factor, with $\lambda = 1$ corresponding to B (and thus the entire portfolio) only depending on the first systematic factor Z_1 , and $\lambda = 0$ corresponding to the creditworthiness index of name B depending solely on the second systematic factor (so that the creditworthiness indices of names A and B are independent). The parameters β_A, β_B give the relative importance of systematic risk and idiosyncratic risk for names A and B respectively. The total portfolio loss random variable thus becomes:

$$L = w_A \mathbf{1}_{\{Y_A \leq \Phi^{-1}(p_A)\}} + w_B \mathbf{1}_{\{Y_B \leq \Phi^{-1}(p_B)\}}$$

with the systematic loss variable being:¹⁰

$$L_S = \mathbb{E}[L|Z]$$


$$= w_A \Phi \left(\frac{\Phi^{-1}(p_A) - \beta_A Z_1}{\sqrt{1 - \beta_A^2}} \right) + w_B \Phi \left(\frac{\Phi^{-1}(p_B) - \beta_B (\lambda Z_1 + \sqrt{1 - \lambda^2} \cdot Z_2)}{\sqrt{1 - \beta_B^2}} \right)$$

We consider risk contributions to both VaR and CVaR, at the confidence levels $\alpha = 0.9, 0.95, 0.99, 0.995, 0.999$.


5.1.1 Base Case, $p_A = p_B = 0.005, w_A = w_B = 0.5, \beta_A = \beta_B = 0.5$

Assume the two names (sub-portfolios) have identical default probabilities, $p_A = p_B = 0.05$, (loss given default adjusted) exposures $w_A = w_B = 0.5$ and systematic risk factor loadings $\beta_A = \beta_B = 0.5$. Risk contributions for both VaR and CVaR as a function of λ for all confidence levels are given in figures a-j.


¹⁰This is also the systematic loss if A and B are homogeneous sub-portfolios with weights w_A and w_B respectively.




(a) Base Case VaR Contributions:
90%




(b) Base Case VaR Contributions:
95%




(c) Base Case VaR Contributions:
99%




(d) Base Case VaR Contributions:
99.5%




(e) Base Case VaR Contributions:
99.9%




(f) Base Case CVaR Contributions:
90%




(g) Base Case CVaR Contributions:
95%



(h) Base Case CVaR Contributions:
99%



(i) Base Case CVaR Contributions:
99.5%



(j) Base Case CVaR Contributions:
99.9%

Contributions are calculated based on a Monte-Carlo simulation using ten million scenarios.¹¹ VaR contributions are calculated using a kernel estimator for the conditional expectation with equal weights. Observe that even with a large number of scenarios, VaR contributions are subject to significant estimation error, which could likely be reduced using importance sampling (e.g. Glasserman (2005), Merino and Nyfeler (2004)).


For VaR, in all cases we observe the same general behaviour of increasing importance of the first factor and decreasing importance of the second factor as λ increases, which is expected from the specification of the loss variable. Contributions from the expected loss term are more significant for lower confidence levels, whereas the unexpected loss terms corresponding to the two systematic factors, as well as their joint influence, become more prominent as confidence levels further in the tail are considered. The CVaR contributions follow the same basic pattern as the VaR contributions. However, certain effects that are observed further in the tail for VaR (increased importance of the factors compared to the expected loss term, increased importance of the joint effect term C_{12}) become apparent for lower confidence levels with CVaR.

It is particularly interesting to study the behaviour of the *joint factor contribution* C_{12} . Near $\lambda = 0.65$ and especially for risk measures sensitive to the extreme tail of the loss distribution (CVaR, or VaR for high confidence levels), C_{12} takes on additional importance. This is also illustrated in figures k-1, where we fix $\lambda = 0.5$, and consider CVaR contributions for different confidence levels.


5.1.2 Two Rating Case, $p_A = 0.001$, $p_B = 0.02$, $\beta_A = 0.5$, $\beta_B = 0.2$, $\lambda = 0.5$

We keep the same general two-factor specification. In this case, sub-portfolio A has a higher credit quality ($p_A = 0.001$), and a greater weight on the systematic factor ($\beta_A = 0.5$). It can be interpreted as representing the investment grade component of an institution's portfolio. In contrast, sub-portfolio B has lower credit quality ($p_B = 0.02$)

¹¹We thank (name withheld) for pointing out to us that in this simple case Monte-Carlo simulation is, strictly speaking, not necessary, as the expectations can be computed using more efficient numerical methods.



(k) Base Case CVaR Contributions: Different Confidence Levels



(l) Base Case CVaR Contributions: Different Confidence Levels

and a lower degree of systematic dependence (i.e. higher idiosyncratic risk, $\beta_B = 0.2$), and can be interpreted as the high-yield component of the portfolio. We vary the weight w_A placed on the investment grade portion of the portfolio $w_A = 0.2, 0.3, 0.5, 0.7$, with $w_B = 1 - w_A$, so that the weights may be interpreted as the (exposure-weighted) percentages of wealth invested in each asset class. CVaR contributions for different asset allocations and different confidence levels are given in figures m-v.¹² We observe that the contribution of Z_1 in isolation (C_1) grows quickly in the portfolio weight w_A , principally at the expense of the contribution C_2 , and that this effect is more pronounced at higher confidence levels. The residual contribution of factor co-movements C_{12} also tends to decline as the weight w_A is increased. This contribution is highest when w_A is small ($w_A = 0.2$), and the confidence level is high (e.g. 99.9%).

5.2 Multi-Factor Example

We now consider as the underlying portfolio the CDXIG index, using data from March 31, 2006. One year default probabilities are implied from credit default swap spreads.¹³ We assume a multi-factor model with a single global factor Z_G and a set of sector factors Z_S representing the systematic risk of industry sectors which is not already captured by the global factor Z_G . The creditworthiness index of a name in sector S is given by:


$$Y = \sqrt{\rho_G} \cdot Z_G + \sqrt{\rho_S - \rho_G} \cdot Z_S + \sqrt{1 - \rho_S} \cdot \varepsilon$$

where $\rho_G = 0.17$ and $\rho_S = 0.23$ are parameters that are the same for all instruments. This specification ensures a correlation of ρ_G for the creditworthiness indices of two names not in the same sector (inter-sector correlation) and a correlation of ρ_S for the creditworthiness indices of two names in the same sector (intra-sector correlation). These values are selected to match estimated correlations from Akhavein et al. (2005).


We use a total of seven industrial sectors, consolidated from 25 Fitch sectors. Each name in the index is mapped to one industry sector.

¹²Again, contributions are calculated using Monte-Carlo without importance sampling, with ten million scenarios.


¹³Technically then, we are using risk-neutral default probabilities.




(m) CVaR Contributions: 90%




(n) CVaR Contributions: 90%




(o) CVaR Contributions: 95%




(p) CVaR Contributions: 95%




(q) CVaR Contributions: 99%




(r) CVaR Contributions: 99%




(s) CVaR Contributions: 99.5%



(t) CVaR Contributions: 99.5%



(u) CVaR Contributions: 99.9%



(v) CVaR Contributions: 99.9%

Aggregate Sector	Weight (Exposure)	Weight (EL)
TECH	20.8%	20.4%
SERVICE	9.6%	9.9%
PHARMA	5.6%	3.7%
RETAIL	20.0%	27.8%
FINANCE	19.2 %	15.0 %
INDUSTRIAL	9.6%	9.8%
ENERGY	15.2%	13.4%

Table 1: CDXIG Index Sector Concentrations

For the details of the mapping, see Rosen and Saunders (2007). The industry concentrations for the index are given in Table 1, weighted both by notional and expected loss. The average one year default probability of the index is 0.19%.

It is easy to see that for this specification of the Gaussian copula model, the only nonzero terms in the Hoeffding decomposition correspond to the expected loss term, the single factor terms (both global and sector), and the joint influence terms involving the global factor and one of the sector factors. Percentage CVaR contributions for these terms for several confidence levels are given in Table 2. Contributions are calculated using Monte-Carlo simulation with one million scenarios. Notice that the global factor dominates the risk contributions. This is not surprising, given that it is the only factor that influences all of the names. Again, the contribution of the expected loss term is more significant for lower quantiles, and less significant as one moves further in the tail. In this region one also sees a higher contribution from the terms giving the joint influence of the global factor and one of the sectors. To examine the impact of the parameters on the risk contributions calculated by the model, we repeat the calculations with an intra-sector correlation of 70% and an inter-sector correlation of 20%. This leads to the results in Table 3. Observe that in this case we see substantially higher contributions for the industry factors (in particular TECH and RETAIL).

Factor	CVaR-90%	CVaR-95%	CVaR-99%	CVaR-99.5%	CVaR-99.9%
$E[L]$	19.13%	13.63%	7.19%	5.71%	3.59%
GLOBAL	73.26%	76.99%	80.46%	81.13%	82.07%
TECH	0.78%	0.67%	0.50%	0.44%	0.33%
SERVICE	0.19%	0.15%	0.09%	0.07%	0.05%
PHARMA	0.03%	0.03%	0.02%	0.01%	0.00
RETAIL	1.35%	1.10%	0.75%	0.64%	0.47%
FINANCE	0.47%	0.41%	0.31%	0.26%	0.17%
INDUSTRIAL	0.18%	0.16%	0.11%	0.10%	0.07%
ENERGY	0.37%	0.32%	0.21%	0.19%	0.14%
GLOBAL+TECH	0.95%	1.52%	2.60%	2.92%	3.44%
GLOBAL+SERVICE	0.24%	0.35%	0.48%	0.45%	0.61%
GLOBAL+PHARMA	0.04%	0.07%	0.10%	0.11%	-0.02%
GLOBAL+RETAIL	1.53%	2.33%	3.59%	3.97%	4.43%
GLOBAL+FINANCE	0.68%	1.06%	1.78%	1.95%	2.18%
GLOBAL+INDUSTRIAL	0.29%	0.42%	0.64%	0.73%	0.85%
GLOBAL+ENERGY	0.51%	0.79%	1.19%	1.34%	1.62%

Table 2: CDXIG Systematic Factor CVaR Contributions

Factor	CVaR-90%	CVaR-95%	CVaR-99%	CVaR-99.5%	CVaR-99.9%
$E[L]$	11.63%	6.99%	2.87%	2.17%	1.34%
GLOBAL	32.76%	27.13%	19.27%	17.54%	16.15%
TECH	8.42%	7.80%	7.04%	6.91%	7.95%
SERVICE	3.09%	2.44%	1.29%	0.90%	0.15%
PHARMA	0.79%	0.57%	0.15%	0.09%	0.05%
RETAIL	11.94%	10.89%	9.47%	9.20%	9.11%
FINANCE	5.69%	5.06%	4.37%	4.25%	4.04%
INDUSTRIAL	3.06%	2.45%	1.33%	0.99%	0.23%
ENERGY	4.65%	4.08%	3.16%	2.92%	1.88%
GLOBAL+TECH	3.58%	7.02%	12.75%	14.74%	18.52%
GLOBAL+SERVICE	1.97%	2.97%	2.88%	2.27%	0.91%
GLOBAL+PHARMA	0.73%	0.92%	0.50%	0.36%	0.25%
GLOBAL+RETAIL	4.41%	9.01%	16.02%	18.05%	21.29%
GLOBAL+FINANCE	2.91%	5.38%	9.28%	10.27%	11.66%
GLOBAL+INDUSTRIAL	1.94%	2.93%	3.07%	2.55%	1.15%
GLOBAL+ENERGY	2.42%	4.36%	6.54%	6.80%	5.32%

Table 3: CDXIG Systematic Factor CVaR Contributions

6 Concluding Remarks

This paper extends the techniques developed for marginal risk contributions of portfolio positions (instruments and sub-portfolios) to the calculation of risk contributions from risk factors, with a focus on applications in portfolio credit risk management. This enables a fuller risk analysis of any portfolio, where loss contributions can be calculated on both an instrument/sub-portfolio and a risk factor basis to give risk managers' a better impression of the key drivers of potential portfolio losses. We employ the Hoeffding decomposition of the loss random variable and calculate contributions for each term in this decomposition. We further develop several examples of applications of this decomposition technique to factor models of portfolio credit risk. We demonstrate the financial interpretation of the risk contributions arising from the terms in the Hoeffding decomposition of portfolio losses. They represent the quadratic "best hedges" involving instruments of increasing complexity.

The ideas and examples presented in this paper raise a number of questions that we believe are worthy of consideration in future research. Some of these include:

- From a practical perspective, the development of efficient algorithms for the calculation of factor contributions is important. For example, the application of importance sampling and other techniques used for position contributions (e.g. Glasserman (2005), Huang et al. (2007), and Merino and Nyfeler (2004)) needs to be investigated.
- The development of efficient algorithms for computing the terms of the Hoeffding decomposition, conditional on realizations of the factor values, is necessary when explicit analytical formulas are not available. This is crucial to areas such as the risk management of structured credit portfolios, as well as other non-linear portfolios including equity, fixed-income and FX options.
- From a theoretical perspective, it is important to develop an axiomatic characterization of factor contribution measures in a manner similar to Kalkbrener et al. (2004).
- Finally, we have related the Hoeffding decomposition to quadratic hedges involving instruments of increasing complexity. This technique may be extended to consider best hedges with respect to

other risk measures. In particular, it seems natural to consider best hedges defined by the risk measure ρ whose value for a particular portfolio we are seeking to allocate. Such hedges are unlikely to be analytically tractable and hence require the development of efficient numerical algorithms.

A Hoeffding Decompositions

This appendix presents a more technical review of Hoeffding decompositions than that which appears in the main body of the text. For a fuller discussion, see van der Vaart (1998), sections 11.3-11.4, which we follow closely.

Consider a fixed probability space $(\Omega, \mathcal{F}, \mathbb{Q})$. Let L be a random variable with finite variance, and let Z_1, \dots, Z_K be independent random variables with finite variances. For a given index set $A \subseteq \{1, \dots, K\}$, define a closed subspace H_A of $L^2(\mathbb{Q})$ in the following way. H_A consists of all square-integrable random variables of the form $g(Z_i; i \in A)$ for some Borel measurable g (that is, all variables in $L^2(\mathbb{Q})$ that are measurable with respect to the σ -algebra $\mathcal{F}_A = \sigma(Z_k, k \in A)$) such that additionally:

$$\mathbb{E}[g(Z_k, k \in A) | Z_j, j \in B] = 0 \quad (11)$$

for every B such that $|B| < |A|$. Then we have the direct sum:

$$L^2(\Omega, \mathcal{F}_A, \mathbb{Q}) = \bigoplus_{B \subseteq A} H_B \quad (12)$$

Denoting by P_B the projection onto H_B , we then obtain the (orthogonal) decomposition:

$$\mathbb{E}[L | \mathcal{F}_A] = \mathbb{E}[L | Z_k, k \in A] = \sum_{B \subseteq A} P_B L \quad (13)$$

The projections P_A can be evaluated more explicitly (for a proof, see van der Vaart (1998)):

$$P_A L = \sum_{B \subseteq A} (-1)^{|A|-|B|} \cdot \mathbb{E}[L | Z_k, k \in B] = g_A(Z_k; k \in A) \quad (14)$$

In particular, if $L = g(Z_1, \dots, Z_K)$, then by taking A to be the full index set, $A = \{1, \dots, K\}$ in (12) we obtain:

$$\begin{aligned} L &= \sum_{B \subseteq \{1, \dots, K\}} g_B(Z_k; k \in B) \\ &= \sum_{B \subseteq \{1, \dots, K\}} \sum_{B' \subseteq B} (-1)^{|B|-|B'|} \cdot \mathbb{E}[L | Z_k, k \in B'] \end{aligned} \quad (15)$$

this is called the Hoeffding decomposition of the random variable L .