Risk Contributions from Generic User-Defined Factors[1]

Attilio Meucci
attilio_meucci@symmys.com

this version: March 2007

latest version available at
`symmys.com > Research > Working Papers`

Abstract

We draw on regression analysis to decompose volatility, VaR and expected shortfall into arbitrary combinations or aggregations of risk factors and we present a simple recipe to implement this approach in practice.

JEL Classification: C1, G11

Keywords: risk attribution, marginal contributions, Euler identity, risk factorization, volatility, tracking error, expected shortfall, positive homogeneous risk measures, principal component analysis, regression analysis

[1]This article is published in *The Risk Magazine*, June 2007, p. 84-88

1 Introduction

One of the main priorities of traders and portfolio managers is to analyze the risk of their prospective p&l. In general, the p&l can be expressed as the product of a vector of risk factors $\mathbf{F}$ times the respective exposures $\mathbf{b}$, which represent the practitioner's decision variables:

$$\Pi = \mathbf{b}'\mathbf{F}. \tag{1}$$

This formulation is very general. Indeed, it includes the case where $\mathbf{F} \equiv \Delta\mathbf{P}$ represents the price change of a set of securities (or asset classes) and $\mathbf{b} \equiv \mathbf{w}$ represents the respective portfolio weights. It also covers, among others, APT-like linear factor models with systematic and idiosyncratic components; the carry-duration-convexity approximation of bond trading; and the theta-delta-gamma-vega approximation routinely used on derivatives desks[2].

The risk of a position is assessed in terms of measures such as the standard deviation, the value at risk, and the expected shortfall. For such measures, risk can be expressed as the sum of the contributions from each of the factors $\mathbf{F}$, see below for a review and references.

However, managers often need to analyze risk according to new factors $\widetilde{\mathbf{F}}$ that are combinations of the preexisting factors $\mathbf{F}$. When the new factors $\widetilde{\mathbf{F}}$ span the risk of the whole market, there exists only one way to decompose risk according to the contributions from the new factors. Similarly, the risk contribution analysis is straightforward when the new factors $\widetilde{\mathbf{F}}$ represent aggregations of the original factors $\mathbf{F}$: for instance, the contributions from all the securities in a given industry are summed into one industry-specific contribution.

Nevertheless, practitioners typically are interested in the risk contributions from only a partial set of new factors $\widetilde{\mathbf{F}}$ that does not span the whole market risk. Computing the risk contributions in this case is not trivial, because, as it turns out, the exposures $\widetilde{\mathbf{b}}$ to the newly introduced factors are not well defined. Furthermore, even upon solving this problem, from an implementation perspective it is impossible to pursue different ad-hoc solutions for different cases, especially when arbitrary combinations of aggregations and partial specifications for the new factors are involved.

Here we propose a unified framework to deal with the above issues. On the one hand, we determine the most natural definition for the risk contributions under partial-factor specifications. To this aim, we draw on the risk attribution literature, see references below: the "natural" exposures $\widetilde{\mathbf{b}}$ are the regression coefficients of the p&l Π on the incomplete set of new factors $\widetilde{\mathbf{F}}$. On the other hand, we show how factor aggregation, partial-factor specification and full-factor specification are instances of the same process. Therefore, we can define and implement one routine that handles any possible scenario.

In Section 2 we review the main results regarding the decomposition of risk into its contributions from each of the original factors $\mathbf{F}$. In Section 3 we review

[2] In principle this approach also covers full repricing by adding higher order derivatives in the Taylor expansion of the p&l.

the risk decomposition process in the straightforward case where the new risk factors $\widetilde{\mathbf{F}}$ represent aggregations of the original factors $\mathbf{F}$: in this setting, risk contributions are defined by a bottom-up aggregation rule. In Section 4 we study the risk decomposition process in the less trivial case where the new risk factors $\widetilde{\mathbf{F}}$ fully span the risk in the market. In this setting, risk contributions are defined by a simple transformation rule. In Section 5 we tackle the problem of computing risk contributions from newly defined risk factors $\widetilde{\mathbf{F}}$ that do not span the whole market risk. In this setting, risk contributions are defined by regression analysis. In Section 6 we show how these seemingly different approaches are indeed one and the same. In Section 7 we conclude, presenting a routine to compute in practice risk contributions in any scenarios. To support intuition we illustrate every step of our discussion by means of a real-life example.

2 Risk contributions: a review

This section quickly summarizes results that were introduced by Litterman (1996) and Garman (1997) and then further developed by Tasche (1999), Mina (2002), Hallerbach (2003), Zhang and Rachev (2004), Scherer (2004). See also Meucci (2005) for a detailed review.

The p&l (1) is a function of the exposures $\mathbf{b}$. Therefore, the risk of the portfolio must also be a function $\mathcal{R}(\mathbf{b})$ of the exposures. The most popular measure of risk is the standard deviation of the p&l, also known as tracking error for benchmark-driven allocations:

$$\mathcal{R}(\mathbf{b}) \equiv \sqrt{\mathbf{b}' \operatorname{Cov}\{\mathbf{F}\}\, \mathbf{b}}. \tag{2}$$

Alternative measures are the value at risk (VaR):

$$\mathcal{R}(\mathbf{b}) \equiv Q_{-\mathbf{b}'\mathbf{F}}(c), \tag{3}$$

where Q_X denotes the quantile of the random variable X and c is the confidence level, typically set as $c \approx 99\%$; and the expected shortfall (ES), also known as conditional value at risk:

$$\mathcal{R}(\mathbf{b}) \equiv \mathbb{E}\{-\mathbf{b}'\mathbf{F} | -\mathbf{b}'\mathbf{F} \geq Q_{-\mathbf{b}'\mathbf{F}}(c)\}. \tag{4}$$

All these measures are homogeneous: $\mathcal{R}(\mathbf{b})$ doubles if we double the exposures $\mathbf{b}$. Hence the following identity holds:

$$\mathcal{R}(\mathbf{b}) \equiv \sum_{n=1}^{N} b_n \frac{\partial \mathcal{R}(\mathbf{b})}{\partial b_n}, \tag{5}$$

where N is the dimension of the market $\mathbf{F}$. In other words, total risk can be expressed as the sum of the contributions from each factor, where the generic n-th contribution is the product of the "per-unit" marginal contribution $\partial\mathcal{R}/\partial b_n$ and the "amount" of the n-th factor in the portfolio, as represented by the exposure b_n.

If we measure risk by the volatility (2), the partial derivatives that appear in (5) read:

$$\frac{\partial \mathcal{R}(\mathbf{b})}{\partial \mathbf{b}} = \frac{\mathrm{Cov}\{\mathbf{F}\}\,\mathbf{b}}{\sqrt{\mathbf{b}'\,\mathrm{Cov}\{\mathbf{F}\}\,\mathbf{b}}}, \tag{6}$$

If we measure risk in terms of the VaR (3) we can express the partial derivatives as in Hallerbach (2003), Gourieroux, Laurent, and Scaillet (2000), Tasche (2002) as conditional expectations:

$$\frac{\partial \mathcal{R}(\mathbf{b})}{\partial \mathbf{b}} \equiv -\mathbb{E}\{\mathbf{F}| - \mathbf{b}'\mathbf{F} \equiv Q_{-\mathbf{b}'\mathbf{F}}(c)\}. \tag{7}$$

Similarly, also when risk is measured by the ES (4) the partial derivatives can be expressed as conditional expectations:

$$\frac{\partial \mathcal{R}(\mathbf{b})}{\partial \mathbf{b}} \equiv -\mathbb{E}\{\mathbf{F}| - \mathbf{b}'\mathbf{F} \geq Q_{-\mathbf{b}'\mathbf{F}}(c)\}. \tag{8}$$

In an elliptical market, the derivatives (6)-(8), and therefore the risk contributions (5), can be computed analytically. In fully general, highly non-normal markets we can represent the joint distribution of $\mathbf{F}$ in terms of a $J \times N$ panel $\mathcal{F}$ of Monte Carlo simulations: the generic j-th row represents a joint scenario for the factors and the generic n-th column represents the marginal distribution of the n-th factor. Generating $\mathcal{F}$ is often a simple task even under very complex joint distributional assumptions for the factors. Then the covariance in (6) can be approximated by the sample covariance of $\mathcal{F}$. As for the VaR, the expectations in (7) can be approximated as in Mausser (2003), see also Epperlein and Smillie (2006) and Meucci, Gan, Lazanas, and Phelps (2007):

$$\frac{\partial \mathcal{R}(\mathbf{b})}{\partial \mathbf{b}} \approx -\mathbf{k}_c' \mathcal{S}_{\mathbf{b}}. \tag{9}$$

In this expression $\mathcal{S}_{\mathbf{b}}$ is a $J \times N$ panel, whose generic j-th column is the j-th column of the panel $\mathcal{F}$, sorted as the order statistics of the J-dimensional vector $-\mathcal{F}\mathbf{b}$; and $\mathbf{k}_c$ is a smoothing kernel, peaked around the rescaled confidence level cJ. Similarly, for the ES we can approximate the expectations in (8) as

$$\frac{\partial \mathcal{R}(\mathbf{b})}{\partial \mathbf{b}} \approx -\mathbf{q}_c' \mathcal{S}_{\mathbf{b}}, \tag{10}$$

where $\mathbf{q}_c$ is a step function that jumps from 0 to $1/cJ$ at the rescaled confidence level cJ of the ES.

To illustrate, we consider a portfolio of US government bonds. The p&l of this portfolio is well described in the form (1) by a quadratic duration-convexity approximation. More explicitly, there exist $N \equiv 7$ factors, namely the changes of six key-rates of the Treasury yield curve (6-month, 2,5,10,20 and 30 year) and an average quadratic term for the convexity:

$$\mathbf{F} \equiv \left(\Delta y_{6m}, \Delta y_{2y}, \Delta y_{5y}, \Delta y_{10y}, \Delta y_{20y}, \Delta y_{30y}, \Delta y^2\right)', \tag{11}$$

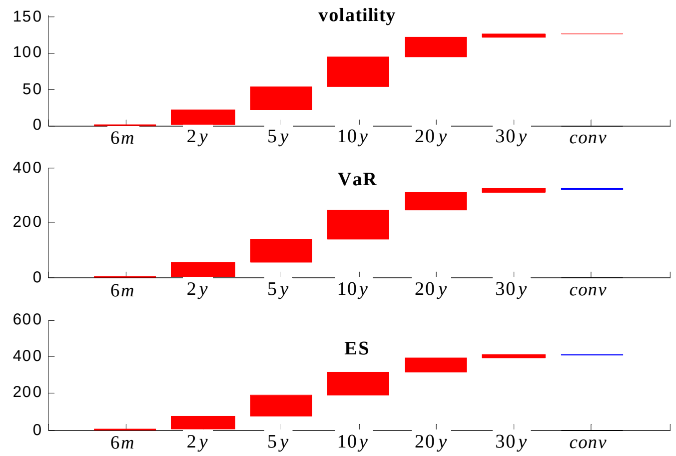

Figure 1: Contributions to portfolio risk from a pre-specified set of factors

where

$$\Delta y^2 \equiv \frac{1}{6}\left(\Delta y_{6m}^2 + \Delta y_{2y}^2 + \Delta y_{5y}^2 + \Delta y_{10y}^2 + \Delta y_{20y}^2 + \Delta y_{30y}^2\right). \tag{12}$$

The exposures $\mathbf{b}$ are the respective key rate durations and the average convexity. In our example, on a specific day these sensitivities read:

b_{6m}	b_{2y}	b_{5y}	b_{10y}	b_{20y}	b_{30y}	b_{y2}
0.091	0.752	1.059	1.516	1.223	0.266	0.481

(13)

where we assume that the p&l is represented as a return and measured in basis points.

We estimate the joint distribution of the six rate changes from a dataset of monthly realizations over a time span of ten years. In particular, in a two-step approach, we first fit each rate to a different Student t distribution: the location parameters are null, the degrees of freedom are 3, 4, 5, 7, 10 and 15 respectively, and the dispersion parameters can be implied from (14) below. Then we fit the joint structure to a normal copula, whose correlation matrix can also be implied from (14) below. From the estimated distribution we generate a $J \times N$ panel $\mathcal{F}$ of $J \equiv 10^6$ Monte Carlo simulations. The first $N-1$ columns are generated according to the above marginal-copula decomposition; the last column is computed deterministically from the first $N-1$ columns as in (12).

We compute the covariance in (6) as the sample covariance of the panel:

Cov $\{\mathbf{F}\}$	Δy_{6m}	Δy_{2y}	Δy_{5y}	Δy_{10y}	Δy_{20y}	Δy_{30y}	Δy^2
Δy_{6m}	593	555	440	311	226	206	0
Δy_{2y}	·	904	862	669	508	464	0
Δy_{5y}	·	·	942	787	622	577	0
Δy_{10y}	·	·	·	729	609	574	0
Δy_{20y}	·	·	·	·	543	516	0
Δy_{30y}	·	·	·	·	·	498	0
Δy^2	·	·	·	·	·	·	29

(14)

Given the exposures (13) and the per-unit contributions (6) we can compute the total volatility and the contributions from each factor. Similarly, we compute the VaR sensitivities (9) and hence the VaR and the contributions from each factor. Finally, we compute the ES sensitivities (10) and hence the ES and the contributions from each factor. We display these contributions in Figure 1 and we report them below:

	Total	C_{6m}	C_{2y}	C_{5y}	C_{10y}	C_{20y}	C_{30y}	C_x
SDev	126	1.2	20.3	31.8	40.4	27.0	5.5	0.1
VaR	320	3.2	53.1	83.0	102.5	66.3	13.3	-1.9
ES	406	4.3	73.6	109.6	127.5	79.2	15.5	-3.8

(15)

Notice that the VaR $\approx$ 320 b.p. and the ES $\approx$ 406 b.p. are not consistent with a normal assumption with standard deviation $\approx$ 126 b.p., where VaR $\approx$ 294 b.p. and ES $\approx$ 337 b.p. Furthermore, notice from Figure 1 that the relative importance of the different contributions changes for different measures, a consequence of the non-elliptical joint distribution of the factors.

3 Risk contributions from aggregate factors

When the number of risk factors in (1) is large, practitioners tend to analyze risk at an aggregate level. Formally we consider K buckets $N_1, \ldots, N_K$ that exhaustively and mutually exclusively span all the N factors $\mathbf{F}$.

In our example, we might be interested in $K \equiv 3$ buckets: the short end of the curve, represented by the 6m, 2y and 5y key rates; the long end of the curve, represented by the 10y, 20y and 30y key rates; and the convexity. Therefore the buckets read:

$$N_1 \equiv \{1, 2, 3\}, \quad N_2 \equiv \{4, 5, 6\}, \quad N_3 \equiv \{7\}. \tag{16}$$

It seems natural to define the contribution to risk $\widetilde{C}_k$ from the generic k-th bucket as the sum of the individual contributions from each factor in the bucket.

From (5) we obtain:

$$\widetilde{C}_k \equiv \sum_{n \in N_k} \frac{\partial \mathcal{R}(\mathbf{b})}{\partial b_n} b_n. \tag{17}$$

In our example, by adding the entries in (15) we obtain the contributions to risk from each bucket:

	Total	$\widetilde{C}_1$	$\widetilde{C}_2$	$\widetilde{C}_3$
SDev	126	53.3	72.9	0.1
VaR	320	139.4	182.0	-1.9
ES	406	187.6	222.2	-3.8

(18)

Apparently, the intuitive rule (17) bears no connection with the problem of computing the risk contributions from combining the original factors $\mathbf{F}$ into newly defined factors $\widetilde{\mathbf{F}}$. As we see in Section 6, this is not the case.

4 Risk contributions from a full set of new factors

In addition to aggregating the risk of the building blocks $\mathbf{F}$, practitioners typically need to rearrange these sources of risk into new risk factors $\widetilde{\mathbf{F}}$ that are linear combinations of the original factors:

$$\widetilde{\mathbf{F}} \equiv \mathbf{PF}. \tag{19}$$

In this expression each row of the "pick" matrix $\mathbf{P}$ represents a linear combination that defines the respective new factor.

In this section we assume that the new factors completely span the risk in the market. In other words, we assume that the pick matrix $\mathbf{P}$ is invertible. To compute the exposures $\widetilde{\mathbf{b}}$ to the new factors we can write the p&l (1) as $\Pi = \mathbf{b}'\mathbf{P}^{-1}\mathbf{PF} \equiv \widetilde{\mathbf{b}}'\widetilde{\mathbf{F}}$. Therefore

$$\widetilde{\mathbf{b}} \equiv \mathbf{P}'^{-1}\mathbf{b}. \tag{20}$$

The per-unit risk contributions from the new factors are also a simple transformation of the per-unit risk contributions from the original factors. Indeed, as we show in a technical appendix available upon request:

$$\frac{\partial \mathcal{R}}{\partial \widetilde{\mathbf{b}}} = \mathbf{P} \frac{\partial \mathcal{R}}{\partial \mathbf{b}}. \tag{21}$$

With (20) and (21) it is immediate to compute the risk contribution from the generic k-th new factor as the product of the per-unit marginal contribution $\partial \mathcal{R} / \partial \widetilde{b}_k$ times the respective exposure $\widetilde{b}_k$.

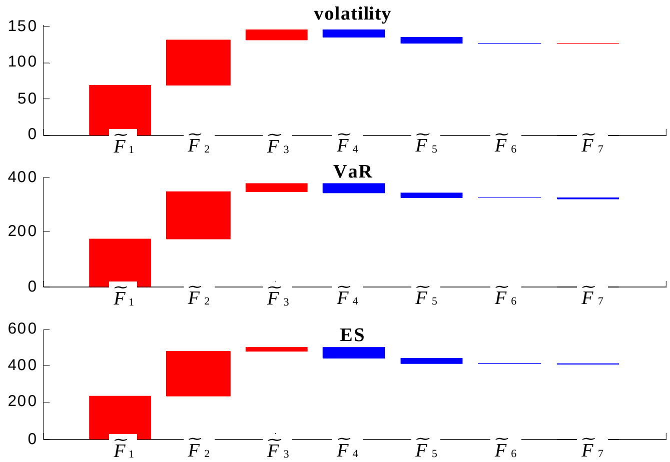

Figure 2: Contributions to portfolio risk from an exhaustive set of new factors

For instance, suppose that the portfolio manager is interested in the exposure to the forward curve, as proxied by the difference between adjacent key rates. In this case the pick matrix reads:

$\mathbf{P}$	Δy_{6m}	Δy_{2y}	Δy_{5y}	Δy_{10y}	Δy_{20y}	Δy_{30y}	Δy^2
$\widetilde{F}_1$	1	0	0	0	0	0	0
$\widetilde{F}_2$	-1	1	0	0	0	0	0
$\widetilde{F}_3$	0	-1	1	0	0	0	0
$\widetilde{F}_4$	0	0	-1	1	0	0	0
$\widetilde{F}_5$	0	0	0	-1	1	0	0
$\widetilde{F}_6$	0	0	0	0	-1	1	0
$\widetilde{F}_7$	0	0	0	0	0	0	1

(22)

Applying (20)-(21) to the numbers computed in Section 2 we obtain the contributions to risk due to the forward factors:

	Total	$\widetilde{C}_1$	$\widetilde{C}_2$	$\widetilde{C}_3$	$\widetilde{C}_4$	$\widetilde{C}_5$	$\widetilde{C}_6$	$\widetilde{C}_7$
SDev	126	67.6	63.4	12.6	-10.1	-6.9	-0.4	0.1
VaR	320	174.0	169.5	31.5	-32.5	-20.0	-1.2	-1.9
ES	406	234.4	241.5	22.8	-58.2	-28.9	-1.7	-3.8

(23)

In Figure 2 we display these contributions. Again, a simple check shows that the normal, or even elliptical assumption, is not viable.

5 Risk contributions from a partial set of new factors

In general, the number N of the risk factors $\mathbf{F}$ that drive the p&l (1) is large. In this case, practitioners typically only wish to aggregate risk as discussed in Section 3, or focus on a small set $K \leq N$ of important user-defined factors $\widetilde{\mathbf{F}}$, or consider a combination of the two approaches. To tackle this problem, we generalize the factor specification (19) as follows:

$$\widetilde{\mathbf{F}} \equiv \mathbf{P}\mathbf{F}, \tag{24}$$

where the pick matrix $\mathbf{P}$, which we can assume to have full rank, has now only $K \leq N$ independent rows.

To illustrate, assume that the portfolio manager is interested in the risk contributions from the first three principal movements of the curve. We perform the principal component decomposition of the first six north-west entries of the matrix (14) and with the eigenvectors corresponding to the three largest eigenvalues we build the three rows of the pick matrix:

$\mathbf{P}$	Δy_{6m}	Δy_{2y}	Δy_{5y}	Δy_{10y}	Δy_{20y}	Δy_{30y}	Δy^2
$\widetilde{F}_1$	0.28	0.48	0.51	0.44	0.36	0.34	0
$\widetilde{F}_2$	-0.71	-0.38	0.0	0.28	0.36	0.37	0
$\widetilde{F}_3$	-0.59	0.41	.46	-0.1	-0.33	-0.41	0

(25)

As in Litterman and Scheinkman (1991) the factor $\widetilde{F}_1$ corresponds approximately to a parallel shift; the factor $\widetilde{F}_2$ corresponds to a steepening/flattening; and the factor $\widetilde{F}_3$ corresponds to a butterfly twist of the curve.

The new factors $\widetilde{\mathbf{F}}$ drive the randomness in the p&l through some exposures $\widetilde{\mathbf{b}}$. However, these exposures are not defined unequivocally. Different choices for the exposures give rise to different residuals in the p&l:

$$\Pi = \widetilde{\mathbf{b}}'\widetilde{\mathbf{F}} + \epsilon. \tag{26}$$

Consequently, different choices for $\widetilde{\mathbf{b}}$ lead to different contributions from the factors $\widetilde{\mathbf{F}}$.

We can see this phenomenon in our example. Consider the first factor in (11), namely the 6m key rate change. The contribution due to this factor depends on the remaining factors, i.e. it depends on the definition of the residual. Indeed, when the remaining factors are the other key rates the contribution of the 6m key rate is minimal, see Figure 1. However, when the remaining factors are the forward rates the contribution of the 6m key rate is large, see Figure 2.

To define the "best" exposures we turn to risk attribution techniques, pioneered by Sharpe (1992) and later studied and applied in a variety of contexts, see e.g. Fung and Hsieh (1997) and Brown and Goetzmann (2003). Much like in (26), in risk attribution the p&l is expressed as a linear combination of factors, where the exposures $\widetilde{\mathbf{b}}$ are set to minimize the variance of the regression residual, i.e. to best explain the p&l. As we show in a technical appendix available upon request, the solution $\widetilde{\mathbf{b}}$ is the standard regression coefficient

$$\widetilde{\mathbf{b}} \equiv \left(\mathbf{P}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\mathbf{P}'\right)^{-1}\mathbf{P}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\mathbf{b}, \tag{27}$$

which decorrelates the new factors $\widetilde{\mathbf{F}}$ from the residual ϵ[3]. With this specific choice of exposures we can compute the risk contribution $\widetilde{C}_k$ from the generic k-th factor $\widetilde{F}_k$ as the product of the exposure $\widetilde{b}_k$ times the respective per-unit contribution $\partial\mathcal{R}/\partial\widetilde{b}_k$, which can be computed easily in terms of the original per-unit contributions $\partial\mathcal{R}/\partial\mathbf{b}$. Indeed, a few algebraic manipulations detailed in the technical appendix show that (21) holds true even when the matrix $\mathbf{P}$ is not invertible:

$$\frac{\partial\mathcal{R}}{\partial\widetilde{\mathbf{b}}} = \mathbf{P}\frac{\partial\mathcal{R}}{\partial\mathbf{b}}. \tag{28}$$

In our example, when risk is measured in terms of the volatility, we obtain the following contributions:

	Total	$\widetilde{C}_1$	$\widetilde{C}_2$	$\widetilde{C}_3$	residual
SDev	126	123.9	2.1	0.1	0.1
VaR	320	316.3	3.5	1.0	-1.2
ES	406	407.4	-1.1	2.6	-2.9

(29)

As expected, a long-only portfolio of bonds is mainly exposed to parallel shifts. In other words, the portfolio duration is an accurate representation of the portfolio risk. The effect of curve flattening and curve twisting is of orders of magnitude smaller than the parallel shift. Notice that the multiple sorting involved in the direct computation of the sensitivities (9) and (10) is a costly operation. Using (28) we only need to compute this operation once for a given portfolio.

6 A generalized framework

We have discussed above three apparently different problems for risk attribution: aggregation, accounted for by means of the simple aggregation rule (17); full factor specification, covered by the transformation rules (20)-(21); and partial factor specification, solved by the regression approach (27)-(28). Since any combination of the above problems might arise in day-by-day applications, from an implementation perspective it is not clear how to solve the risk attribution process in a non-ad-hoc way. However, as it turns out, all the above problems, as well as any combinations thereof, can be cast in a unified framework that is also computationally straightforward.

[3] This is not true for very thick-tailed distributions such as the Cauchy, for which the regression coefficient is not defined

Regarding the full-factor specification, when $\mathbf{P}$ is invertible the regression solution (27) trivially becomes (20). Therefore the full-factor specification represents a special case of the regression approach.

As far as the aggregation rule (17) is concerned, we can express the p&l (1) in terms of new bucket-specific factors as follows:

$$\Pi = \sum_{k=1}^{K} \widetilde{b}_k \widetilde{F}_k. \tag{30}$$

In this expression the exposures are trivially defined as $\widetilde{\mathbf{b}} \equiv \mathbf{1}$, a vector of ones, and the new factors are defined as the portfolio-weighted sum of all the original factors in a given bucket:

$$\widetilde{F}_k \equiv \sum_{n \in N_k} F_n b_n. \tag{31}$$

This set of factors corresponds to a pick matrix in the partial factor specification (24) which is defined entry-wise as:

$$P_{nk} \equiv \begin{cases} b_n & \text{if the factor } F_n \text{ is in the } k\text{-th bucket} \\ 0 & \text{otherwise.} \end{cases} \tag{32}$$

Notice that in (30) the p&l is fully described by the new factors. In other words, $\widetilde{\mathbf{b}} \equiv \mathbf{1}$ minimizes the residual, which is null. Hence, $\widetilde{\mathbf{b}} \equiv \mathbf{1}$ represents the regression solution for the partial factor specification defined by the pick matrix (32) and the aggregation rule (17) coincides with the product of the per-unit marginal contribution $\partial \mathcal{R} / \partial \widetilde{b}_k$ times the respective trivial exposure $\widetilde{b}_k \equiv 1$.

Therefore, both full-factor specification and factor aggregation can be cast within the partial factor framework, which is solved by means of regression analysis.

7 Conclusions

We presented a unified approach to compute the contributions to risk from generic user-defined factors which are aggregations and/or linear combinations of the risk factors that drive the p&l. The algorithm to implement this approach proceeds as follows:

- Start with the p&l Π as a function of given factors $\mathbf{F}$ and their exposures $\mathbf{b}$ as in (1); and the risk-contribution analysis in term of those factors as in (5).

- Determine new factors $\widetilde{\mathbf{F}}$ as linear combinations of the existing factors. In particular, if considering risk aggregations, the respective linear combinations are defined as in (32).

- Stack the coefficients of these linear combinations to form the pick matrix $\mathbf{P}$ in (24). If the rank of $\mathbf{P}$ is not full, delete the redundant rows (and thus the respective factors) until the rank of $\mathbf{P}$ is full.

- Compute the exposures $\widetilde{\mathbf{b}}$ of the new factors as in (27).
- Compute the per-unit contributions to risk from the new factors $\partial\mathcal{R}/\partial\widetilde{\mathbf{b}}$ as in (28).
- Compute the contribution to risk from the generic k-th factor as the product of the k-th entry of $\widetilde{\mathbf{b}}$ times the k-th entry of $\partial\mathcal{R}/\partial\widetilde{\mathbf{b}}$.

The above routine can be easily coded and all the computations can be performed in a fraction of a second.

Acknowledgments: the author is grateful to Robert Durie, Anthony Lazanas, Antonio Silva and two anonymous referees for their feedback

References

Brown, S. J., and W. N. Goetzmann, 2003, Hedge funds with style, *Journal of Portfolio Management* 29, 101–112.

Epperlein, E., and A. Smillie, 2006, Cracking VaR with kernels, *Risk Magazine* 19, 70–74.

Fung, W., and D. A. Hsieh, 1997, Empirical characterizations of dynamic trading strategies: The case of hedge funds, *Review of Financial Studies* 10, 275–302.

Garman, M., 1997, Taking VaR to pieces, *Risk* 10, 70–71.

Gourieroux, C., J. P. Laurent, and O. Scaillet, 2000, Sensitivity analysis of values at risk, *Journal of Empirical Finance* 7, 225–245.

Hallerbach, W., 2003, Decomposing portfolio value-at-risk: A general analysis, *Journal of Risk* 5, 1–18.

Litterman, R., 1996, Hot spots and hedges, *Goldman Sachs and Co., Risk Management Series.*

——— , and J. Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54–61.

Mausser, H., 2003, Calculating quantile-based risk analytics with L-estimators, *Jorunal of Risk Finance* pp. 61–74.

Meucci, A., 2005, *Risk and Asset Allocation* (Springer).

——— , 2006, Beyond Black-Litterman in practice: A five-step recipe to input views on non-normal markets, *Risk* 19, 114–119.

——— , Y Gan, A. Lazanas, and B. Phelps, 2007, A portfolio mangers guide to Lehman Brothers tail risk model, *Lehman Brothers Publications*.

Mina, J., 2002, Risk attribution for asset managers, *RiskMetrics group working paper series*.

Scherer, B., 2004, *Portfolio Construction and Risk Budgeting* (Risk Books).

Sharpe, W. F., 1992, Asset allocation: Management style and performance measurement, *Journal of Portfolio Management* pp. 7–19.

Tasche, D., 1999, Risk contributions and performance measurement, *Working Paper, Technische Universitaet Muenchen.*

——— , 2002, Expected shortfall and beyond, *Journal of Banking and Finance* 26, 1519–1533.

Zhang, Y., and S. Rachev, 2004, Risk attribution and portfolio performance measurement - an overview, *working paper*.

A Appendix

A.1 Formulas for per-unit contributions

Let us complete the partial factor specification (24) by means of $N-K$ ancillary factors $\widehat{\mathbf{F}} \equiv \widehat{\mathbf{P}}\mathbf{F}$ defined by a matrix $\widehat{\mathbf{P}}$ such that

$$\overline{\mathbf{P}} \equiv \begin{pmatrix} \mathbf{P} \\ \widehat{\mathbf{P}} \end{pmatrix} \tag{33}$$

is invertible. In this scenario the new exposures $\widetilde{\mathbf{b}}$ follow from (20) as the first K entries of the vector $\overline{\mathbf{P}}'^{-1}\mathbf{b}$, whereas the remaining $N-K$ entries $\widehat{\mathbf{b}}$ of the vector $\overline{\mathbf{P}}'^{-1}\mathbf{b}$ combine with the ancillary factors $\widehat{\mathbf{F}}$ to give rise to a residual $\epsilon \equiv \widehat{\mathbf{b}}'\widehat{\mathbf{F}}$ in (26), which reads:

$$\Pi \equiv \widetilde{\mathbf{b}}'\widetilde{\mathbf{F}} + \widehat{\mathbf{b}}'\widehat{\mathbf{F}}. \tag{34}$$

Then

$$\mathcal{R} = \widetilde{\mathbf{b}}' \frac{\partial \mathcal{R}}{\partial \widetilde{\mathbf{b}}} + \widehat{\mathbf{b}}' \frac{\partial \mathcal{R}}{\partial \widehat{\mathbf{b}}}. \tag{35}$$

Recalling that

$$\begin{pmatrix} \widetilde{\mathbf{b}} \\ \widehat{\mathbf{b}} \end{pmatrix} \equiv \overline{\mathbf{P}}'^{-1}\mathbf{b}, \tag{36}$$

from the chain rule of calculus we obtain:

$$\begin{pmatrix} \partial \mathcal{R}/\partial \widetilde{\mathbf{b}} \\ \partial \mathcal{R}/\partial \widehat{\mathbf{b}} \end{pmatrix} = \begin{pmatrix} \mathbf{P} \\ \widehat{\mathbf{P}} \end{pmatrix} \frac{\partial \mathcal{R}}{\partial \mathbf{b}}. \tag{37}$$

In particular

$$\frac{\partial \mathcal{R}}{\partial \widetilde{\mathbf{b}}} = \mathbf{P} \frac{\partial \mathcal{R}}{\partial \mathbf{b}}. \tag{38}$$

A.2 Formulas for optimal exposures

First we derive the expression of the regression coefficients. Imposing that the residual be uncorrelated with the factors (24) we obtain:

$$\begin{aligned} \mathbf{0} &= \operatorname{Cov}\left\{\widetilde{\mathbf{F}}, \Pi - \widetilde{\mathbf{F}}'\widetilde{\mathbf{b}}\right\} \\ &= \operatorname{Cov}\left\{\widetilde{\mathbf{F}}, \Pi\right\} - \operatorname{Cov}\left\{\widetilde{\mathbf{F}}\right\}\widetilde{\mathbf{b}} \\ &= \operatorname{Cov}\left\{\mathbf{PF}, \mathbf{F}'\mathbf{b}\right\} - \operatorname{Cov}\left\{\mathbf{PF}\right\}\widetilde{\mathbf{b}} \\ &= \mathbf{P}\operatorname{Cov}\left\{\mathbf{F}\right\}\mathbf{b} - \mathbf{P}\operatorname{Cov}\left\{\mathbf{F}\right\}\mathbf{P}'\widetilde{\mathbf{b}}. \end{aligned} \tag{39}$$

Therefore

$$\widetilde{\mathbf{b}} = \left(\mathbf{P}\operatorname{Cov}\left\{\mathbf{F}\right\}\mathbf{P}'\right)^{-1}\mathbf{P}\operatorname{Cov}\left\{\mathbf{F}\right\}\mathbf{b}. \tag{40}$$

Now we show that the regression solution (40) can be obtain by factor completion as in (33). Consider the principal component factorization of the factor covariance:

$$\mathrm{Cov}\left\{\mathbf{F}\right\} \equiv \mathbf{E}\boldsymbol{\Lambda}^{1/2}\boldsymbol{\Lambda}^{1/2}\mathbf{E}, \tag{41}$$

where $\mathbf{E}$ is the orthogonal matrix of the eigenvectors and $\boldsymbol{\Lambda}$ is the diagonal matrix of the positive eigenvalues.

Define $\widehat{\mathbf{P}}$ in (33) as follows:

$$\widehat{\mathbf{P}} \equiv \mathrm{nul}\left(\mathbf{P}\mathbf{E}\boldsymbol{\Lambda}^{1/2}\right)' \boldsymbol{\Lambda}^{-1/2}\mathbf{E}', \tag{42}$$

where $\mathrm{nul}\left(\mathbf{A}\right)$ is any matrix in the null space of the generic $K \times N$ matrix $\mathbf{A}$, i.e. any $N \times (N-K)$ matrix such that

$$\mathrm{nul}\left(\mathbf{A}\right)' \mathbf{A}' = \mathbf{0}_{(N-K)\times K}. \tag{43}$$

We want to prove that $\widetilde{\mathbf{F}} \equiv \mathbf{P}\mathbf{F}$ and $\widehat{\mathbf{F}} \equiv \widehat{\mathbf{P}}\mathbf{F}$ are uncorrelated. In other words, we want to prove that the following matrix is block diagonal:

$$\begin{aligned} \mathrm{Cov}\left\{\overline{\mathbf{P}}\mathbf{F}\right\} &= \begin{pmatrix} \mathbf{P} \\ \widehat{\mathbf{P}} \end{pmatrix} \mathrm{Cov}\left\{\mathbf{F}\right\} \left(\mathbf{P}', \widehat{\mathbf{P}}'\right) \\ &= \begin{pmatrix} \mathbf{P}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\mathbf{P}' & \mathbf{P}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\widehat{\mathbf{P}}' \\ \widehat{\mathbf{P}}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\mathbf{P}' & \widehat{\mathbf{P}}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\widehat{\mathbf{P}}' \end{pmatrix} \end{aligned} \tag{44}$$

Indeed using (42) and (43) we obtain:

$$\begin{aligned} \widehat{\mathbf{P}}\,\mathrm{Cov}\left\{\mathbf{F}\right\}\mathbf{P}' &= \mathrm{nul}\left(\mathbf{P}\mathbf{E}\boldsymbol{\Lambda}^{1/2}\right)' \boldsymbol{\Lambda}^{-1/2}\mathbf{E}'\mathbf{E}\boldsymbol{\Lambda}^{1/2}\boldsymbol{\Lambda}^{1/2}\mathbf{E}'\mathbf{P}' \\ &= \mathrm{nul}\left(\mathbf{P}\mathbf{E}\boldsymbol{\Lambda}^{1/2}\right)' \boldsymbol{\Lambda}^{1/2}\mathbf{E}'\mathbf{P}' = \mathbf{0} \end{aligned} \tag{45}$$

The proof that the contributions to risk are not affected by a specific choice of representative in $\mathrm{nul}\left(\mathbf{A}\right)$ follows as in Meucci (2006).

Notice that the role of $\widehat{\mathbf{P}}$ is not apparent in (40) because $\widehat{\mathbf{P}}$ as defined in (42) is in turn a function of $\mathrm{Cov}\left\{\mathbf{F}\right\}$ and $\mathbf{P}$.