A risk decomposition framework consistent with performance measurements

Marco Marchioro
www.statpro.com
Version 0.98

January 2017*

Abstract

We define a framework to compute the risk contributions of a portfolio consistently with a given performance-measurement schema. The framework has a wide array of applications, such as risk attribution, and matches the standard risk decomposition when the portfolio has a linear dependence on the factor exposures. We go beyond the traditional risk-decomposition approach, that uses Euler's theorem, and instead require the risk measure to belong to a newly-defined class of *severable risk measures*. As such the risk-decomposition framework defined here can also be used in presence of liquidity risk. We show how the framework can also be used to compute the risk-driver contributions for a non-linear portfolio, define a corresponding concept of risk-driver exposure, and show how to hedge portfolios with respect to a given risk measure.

Keywords: risk measures, distortion risk measures, standard deviation, severable risk measures, risk decomposition, risk contributions, risk attribution, portfolio hedging, risk exposure, non-linear portfolios, liquidity risk, performance attribution, performance contributions, performance-contribution schema

Contents

*Latest minor review January 10, 2017

1 Introduction and preliminary results

Innovation in the financial industry, caused by the investor's demand for gains and diversification, has been increasingly met by a more and more sophisticated offer of financial instruments (see, for example, reference [18]). As a consequence the complexity of managing *market risk* faced by the operators has increased significantly.

There are a number of quantitative tools available to risk managers such as the computation of different risk measures, stress tests, and sensitivity analysis (see reference [8] for more details). Among these tools the most regarded is possibly the computation of *risk measures*: to any given portfolio of assets we assign a real number ρ that can be interpreted as the amount of money needed to insure against tail losses. Typically a risk measure, or simply *risk*, is either expressed as a cash amount or as a percentage of the portfolio outstanding value. In this paper we focus on *market risk*, i.e. those adverse conditions caused by changes in the financial markets. Because of the generality of the framework formulation, other types of risks, such as credit risk or liquidity risk, could also be handled, however, they are not the focus of this work.

There are many different ways to compute risk and the various definitions vary according to the statistical assumptions behind them, the data utilised in the portfolio evaluations, and so on and so forth. Portfolio values depends in general from a large number of different sources, known as the *risk factors*, so that two different portfolios characterised by the same risk measure might behave very differently as a consequence of the same market changes. Therefore we prefer to complement the computation of a risk measure with the calculation of its risk contributions, so that we can better understand where the risk is coming from. Since, because of diversification, the most useful risk measures are not additive, i.e. the sum of the risk of each part is different from the risk of the whole, it is in general a non-trivial task to split a risk measure into separate contributions.

In this paper we use the term *risk contributions*, or equivalently the term *risk components*, to denote a set of terms so that their sum provides the portfolio risk measure. Similarly, we use the term *risk decomposition* as the methodology used to compute the risk contributions. Unfortunately the literature on this terminology is quite confusing and it is often found that the term risk attribution is also used to denote the risk decomposition as defined in this paper (see for example reference [20]). However, as it is a customary terminology in the asset-management industry, we reserve the term *risk attribution* for the specific case of risk decomposition when the risk components need to be associated to the different stages of the investment process.

When performing a portfolio risk decomposition we express a risk measure ρ as the sum of k contributions,

$$\rho = C_1 + C_2 + \ldots + C_k \,, \tag{1}$$

where each component C_i, with $i = 1, \ldots, k$, has a certain discernible financial meaning. In the specific case where the risk measure considered is the standard deviation and the portfolio value is given by the sum of k holdings, as shown for example in reference [9], each risk contribution can be computed as

$$C_i = \sigma_i \cdot \gamma_i \qquad \text{for } i = 1, \ldots, k \,, \tag{2}$$

where σ_i is the standard deviation of the i-th holding and γ_i is its correlation with the portfolio return. Expression (2) for the risk contributions is also convenient as it can be used in conjunction with the ex-post measures of portfolio performance (see, for example, reference [11]).

To compute the risk decomposition of other risk measures, those for example that go beyond the limitations of the standard deviation, the traditional approach is to rely on the fact that, as a function of the exposures, the risk measure satisfies the hypothesis of Euler's homogeneous-function theorem. Even thought, as shown by reference [12], this methodology goes a long way and provides useful risk-contribution results for a many interesting cases, it is still not applicable to a large number of non-linear portfolios. Specifically when liquidity risk is taken into account, doubling the exposure to a certain asset more than doubles the portfolio risk, so that the computation of risk contributions via Euler's theorem is not feasible.

The risk-decomposition methodologies described so far have some drawbacks. Firstly, when there is a strong non-linear dependency of risk from the risk exposures, and hence the hypothesis of Euler's homogeneous-function theorem are not satisfied, it is not clear how to compute the risk contributions. Secondly not all computed risk components are easily comparable with the ex-post portfolio performance measures.

In this paper we describe a risk-decomposition framework that can be used by itself or in conjunction with a performance-measurement technique. The technique we describe has the advantage to fall back to the traditional methods in the simple case of linear dependence from the risk exposures, with the possibility of computation of risk contributions in a much larger number of cases. Furthermore since the proposed method is based on the split of performance in each market outcome, it provides risk contributions that are easily compared with the ex-post results of a large-class of performance contribution methods.

1.1 Standard-deviation risk decomposition

We consider here the risk decomposition of the standard deviation, deriving expression (2) for the risk contributions. In terms of notation in the following discussions, in order to avoid confusions, we denote by the symbol Y the portfolio gain and with L the portfolio loss, so that

$$L = -Y\,. \tag{3}$$

Both the return and the loss can either be cash or computed as a percentage with respect to the portfolio market value (or any other reference value). Furthermore we always think at Y as the portfolio return over one period of time.

Consider a portfolio with a random gain Y that is obtained as the sum of the returns U_i of its k holdings, i.e.

$$Y = U_1 + \ldots + U_k\,. \tag{4}$$

We define the portfolio standard deviation $\sigma(Y)$ as the square root of the covariance of the portfolio with itself:

$$\sigma(Y) = \sqrt{\mathrm{Cov}\,(Y,Y)}\,.$$

Since we can write

$$\begin{aligned}\sigma(Y) &= \frac{1}{\sigma(Y)}\mathrm{Cov}\,(Y,Y)\\ &= \frac{\mathrm{Cov}(Y,U_1+\ldots+U_k)}{\sigma(Y)}\,,\end{aligned}$$

because of the additivity of the covariance in each of its arguments, it is easy to show that

$$\sigma(Y) = \frac{\mathrm{Cov}(Y,U_1)}{\sigma(Y)} + \ldots + \frac{\mathrm{Cov}(Y,U_k)}{\sigma(Y)}\,,$$

hence, using the standard definition of correlation, we have

$$\mathrm{Cov}(Y,U_i) = \gamma_i \cdot \sigma_i \cdot \sigma(Y)\,,$$

for each $i = 1,\ldots,k$, so that

$$\sigma(Y) = \sigma_1 \cdot \gamma_1 + \ldots + \sigma_k \cdot \gamma_k\,, \tag{5}$$

thus obtaining equation (2). As noted earlier this risk decomposition provides useful insights in the structure of the standard deviation and can even be applied consistently with performance measurements (again, see reference [11] for more details on the subject).

1.2 Path to a general decomposition framework

A deep understanding of the steps that brought us to equation (5) provide us with a useful path to a general risk-decomposition framework. Therefore let us carefully analise the derivation of the standard-deviation split into components.

In order to obtain the result above we use two important ingredients: the linearity of the covariance functional (with respect to its second argument) and the split into additive components of the portfolio returns. Since to compute the covariance functional we need to put together the results of many outcomes, the first ingredient is statistical in nature as it depends on the underlying probability estimated for each future scenario. On the other hand, since the portfolio-return split into additive components is valid for each outcome, the second property is *strong* and independent from any probability assumption. We believe that the key to the computation of any risk contribution is to keep the probabilistic assumptions of the underlying risk factors independent from the split of the portfolio returns into additive terms.

Because of the linearity of the covariance in each of its arguments it is easy to split the standard deviation into additive components. Unfortunately there are not many risk measures that have this property, however, we can often extract a *useful* linear operator from a large number of risk measures.

Borrowing a term from contract law, we say that a risk measure ρ is *severable* with respect to a random variable Y if there exist an *associated linear functional* $A_Y[Z]$ so that in the specific case where $Z = -Y$ we have

$$\rho(Y) = A_Y[-Y]\,. \tag{6}$$

Note that, we could express the same equation in terms of the loss L as

$$\rho(Y) = A_Y[L]\,.$$

We added a subscript Y to the associated functional symbol to stress its dependence on the portfolio return. In order to define a severable risk measure it is sufficient to have a linear map $A_Y[Z]$ that satisfies equation (6), however not all associated functionals might be provide interesting results. In terms of notation, when the dependence from Y is clear from the context, we simply use the symbol A to denote the associated functional.

Recall that a linear functional satisfies both the additivity property, i.e. for any finite number of random variables $Z_1, Z_2, \ldots$ we have

$$A\,[Z_1 + Z_2 + \ldots] = A[Z_1] + A[Z_2] + \ldots\,, \tag{7}$$

and the homogeneity property, i.e. for any real number λ we have

$$A[\lambda \cdot Z] = \lambda \cdot A[Z]\,. \tag{8}$$

Although only the additivity property is strictly necessary in order to compute the risk contributions, in all practical cases the homogeneity property is also satisfied and therefore, for simplicity we also require it.

The first example of a severable risk measure is, obviously, the standard deviation $\sigma(Y)$ and, for a fixed portfolio random gain Y, the associated functional is simply given by

$$A_Y[Z] = \frac{1}{\sigma(Y)}\mathrm{Cov}(Y, Z)\,, \tag{9}$$

for any random variable Z. Note that in the above expression Z is the independent variable and that the portfolio return Y is only a fixed parameter of the functional A_Y. Also, in the standard-deviation case, the associated linear functional A_Y is an *average* of some sort, hence the choice of letter A to denote it.

The standard deviation is not an isolated case of a severable risk measure and there are many more interesting risk measures that satisfy the same property. In section 2 we show the severability of a large class of risk measures known as the distortion risk measures.

Proposal for a risk-decomposition framework Consider a risk measure ρ, severable with respect to a stochastic return Y. Suppose also that Y satisfies equation (4). We can generalise the computation of risk components by writing equation (1) with the risk contributions C_i given by

$$C_i = A_Y[-U_i]\,. \tag{10}$$

In order to obtain this equation we wrote $\rho(Y)$ in terms of the linear functional A_Y, we substituted Y with $U_1 + \ldots + U_k$, and then used the additivity property of A_Y. We thus obtain the risk contributions of a portfolio of holdings with respect to any severable risk measure.

2 Distortion risk measures

An important class of severable risk measures is given by the *distortion risk measures* (see for example reference [7]). Distortion risk measures are very popular and are among the most used measures of risk by practitioners. For example both expected shortfall and value at risk belong to this category. Furthermore all spectral risk measures are also distortion risk measures.

2.1 Definition of a generic risk measures

In the previous section we stressed the importance of risk computations as one of the important tools available in the hands of risk managers. In this subsection we define exactly what is a risk measure and describe its basic properties. Since the subject of risk measures is very wide, we provide only the basic results and refer to more specialised literature, such as reference [7], for more details on the subject.

We start by considering the financial markets in their present state and assume that there are a number, possibly an infinite number, of risk factors propelling the evolution of financial instruments. Assuming equilibrium, at least in principle, at the current time we can determine the quoted prices of all assets by the knowledge of the current market risk factors. We consider then the market evolution during one period of time, being that a day, a week, a month, or a year. For all practical purposes we assume the financial markets to evolve stochastically in the elapsed time.

We describe each market outcome, a.k.a. the market state, by a risk-factor vector ω. Technically, in order to use the plethora of results from the theory of probability spaces, we assume that Ω, the universe of possible market states at the end of the period, is a Borel set. For the purposes of our investments the market state is sufficient to determine the portfolio value. Hence, given the financial uncertainties, the best we can hope is a probabilistic representation of the market state at the end of the observation period.

Note that in this formulation we may or may not be able to directly observe the risk factors but only their effects on the market. In particular we assume that we can observe in the market a *finite number d* of random variables $X_1, \dots, X_d$, known as the *risk drivers*, that are sufficient to determine with a good approximation all market prices. Example of risk drivers are the return of stock prices and the increase in interest rates (see reference [4] for more details).

There are different ex-ante expectations of market movements depending on the chosen time horizon of our risk measure, so that if the time horizon is of one day or one year we have different probabilities of reaching certain events. For the given time horizon, we assume that we know the probability space[1] $(\Omega, \mathcal{F}, P)$ that describes the market state at the end of the period.

Since the portfolio return may have any type of dependence on the market state, we assume that it can be modelled by a bounded random variable Y on the given probability space. (Technically, we define the space $\mathcal{Y} = L^\infty(\Omega, \mathcal{F}, P)$ of essentially bounded functions and assume $Y \in \mathcal{Y}$.) The random variable $L = -Y$ then provides the portfolio loss, either in cash or as a percentage of the portfolio value.

Very often we are able to precisely estimate the portfolio loss in terms of the risk drivers by means of a deterministic *loss pricing function* f, so that we have

$$L = f(X_1, \dots, X_d)\,. \tag{11}$$

Note that we use the letter Y to denote the portfolio return, as opposed to the more traditional symbol X, in order to stress that Y is dependent on the risk drivers $X_1, \dots, X_d$. In turn the risk drivers may have any functional dependence on the risk factors, which may or may not be directly observable.

Given the definitions above we can, for example, compute the mathematical expectation of the loss L as

$$\rho_{\mathrm{ExpLoss}}(Y) = E[-Y] = E[L] = \int_\Omega L(\omega)\mathrm{d}P(\omega)\,. \tag{12}$$

The *expected loss* just defined is a functional, i.e. a map from $\mathcal{Y}$ to the real numbers, and provides the weighted average of the portfolio returns for the different market states according to the *natural* probability measure P. Since when $\rho_{\mathrm{ExpLoss}}(Y)$ is positive we expect a portfolio loss, we can regard the functional ρ_{ExpLoss} as the simplest risk measure we can define. According to this risk measure all market states are given the same importance.

Notice that the expected loss, being itself linear, is severable for any portfolio return Y with the same associated linear functional

$$A[Z] = E[Z]\,.$$

The risk contributions computed using this functional provide a risk decomposition of the expected loss

[1] As usual, a *probability space* is defined by a set Ω, a σ-algebra on Ω, and a probability measure P.

ρ_{ExpLoss} in terms of the sum of the expected losses of the portfolio holdings (see subsection 1.2).

More in general, we define a *risk measure* ρ to be a functional on $\mathcal{Y}$ and write $\rho : \mathcal{Y} \to \mathbb{R}^+$. In order for ρ to actually represent a financial risk measure there are a number of assumptions that should be satisfied. For example we are comfortable to assume that a risk measure should satisfy *positive homogeneity*:

$$\rho(\lambda \cdot Y) = \lambda \cdot \rho(Y)\,, \tag{13}$$

for any given real $\lambda \geq 0$. Positive homogeneity is just one possible assumption that we can make on the risk measures and it is not necessary for our framework.

Depending on the mathematical assumptions that we require for the functional ρ we end up with different classes of risk measures. Also note that the above definition of risk measure is fully compatible with the inclusion of the portfolio liquidity risk (see reference [5] for more details). The literature on this subject is very vast, see for example reference [7]. In this paper we focus on the so called *distortion risk measures*.

2.2 Definition of distortion risk measures

In order to define a distortion risk measure we start from the definition of the average portfolio loss:

$$E[Y] = \int_{\Omega} L(\omega)\mathrm{d}P(\omega) = -\int_0^1 F_Y^{-1}(p)\mathrm{d}p\,, \tag{14}$$

where the second equality is a well-known result in probability theory, given that $F_Y^{-1}(y)$ is the inverse cumulative return distribution of the portfolio returns. The last term on the right-hand side of this equation can be interpreted as the average of the term $F_Y^{-1}(p)$ uniformly over the range $[0, 1]$. In order to have a usable definition of risk measure we need to allow the user to express an opinion on the importance of different losses and gains. Therefore to obtain a more adeguate definition of risk we skew the probability in the range $[0, 1]$ so that different losses carry different weights. In other words we increase the weight of outcomes that we deem more damaging while, at the same time, lowering that of harmless events.

Mathematically we define a *distortion function* g, $g : [0, 1] \to [0, 1]$, as a non-decreasing right-continuous function that satisfies $g(0)=0$ and $g(1)=1$. Then we note that for any event S in the σ-algebra $\mathcal{F}$, the set function Q_g, so that

$$Q_g(S) = g(P(S))\,,$$

is *not* a probability measure but a so-called *distorted probability* (see reference [6] for more details). Even though we cannot define the standard Lebesgue integral for a distorted probability, it is still possible to define a non-linear expectation with respect to Q_g. Technically this is performed by computing the Choquet integral,

$$\rho_g(Y) = \int_{\mathcal{C}} L\mathrm{d}Q_g\,, \tag{15}$$

that does not satisfy the usual linearity property of the standard integrals. We defer to reference [19] for more results on the topic of distortion risk measures defined as non-linear expectations and recall here only the most important results.

As shown in reference [7], as a consequence of definition (15), any given distortion function g provides a distortion risk measure ρ_g that can be computed as the *distorted loss expectation*:

$$\rho_g(Y) = -\int_0^1 F_Y^{-1}(p)\mathrm{d}H(p)\,, \tag{16}$$

where the function H is defined as

$$H(u) = 1 - g(1 - u)\,.$$

Hence the distortion function g allows us to state the investor risk profile by specifying how to allocate the weights of different returns and losses.

As shown among others by reference [2], there is an intimate connection between distortion risk measures and *spectral risk measures*. Indeed for any *spectrum function*[2] $\phi(p)$ so that

$$g(p) = \int_0^p \phi(u)\mathrm{d}u\,,$$

[2] In principle the spectrum function could be generalised to an *admissible risk aversion function* as shown, e.g., in reference [1].

we have

$$\rho_g(Y) = -\int_0^1 \phi(p) F_Y^{-1}(p) \mathrm{d}p \, .$$

Hence, if ϕ is increasing, which happens if and only if g is concave, the risk measure ρ_g is also coherent. Coherent risk measures are very important as they can be used to determine optimal portfolios (see, for example, reference [1]).

Properties of distortion risk measures As mentioned in reference [7] distortion risk measures have many interesting properties. The most important ones are

1. Monotonicity: if $L \geq 0$ then $\rho_g(Y) \geq 0$
2. Positive homogeneity, already defined by equation (13)
3. Translation invariance: for all $c \in \mathbb{R}$ we have $\rho_g(c+Y) = c + \rho_g(Y)$

In general distortion risk measures are not additive, i.e.

$$\rho_g(Y+Z) \neq \rho_g(Y) + \rho_g(z) \, ,$$

however if the distortion function g is concave the risk measure becomes subadditive:

$$\rho_g(Y+Z) \leq \rho_g(Y) + \rho_g(z) \, .$$

Subadditivity is a very desirable property for risk measures since it states that it is possible to diversify portfolio risk by adding to the portfolio new sources of uncertainties.

2.3 Examples of distortion risk measures

We briefly introduce the two most common distortion risk measures, namely the value at risk and the expected shortfall, and refer to citation [7] for more details on the subject.

The first text-book example of a distortion risk measure is given by the well-known *value at risk*. Qualitatively value at risk is the maximum loss that we can expect for any given confidence level α, for example α=99%. We define *value at risk*, also known as VaR, and denote it with the symbol ρ_{VaR}, as the distortion risk measure generated by the function

$$g_{\mathrm{VaR}}(p) = \begin{cases} 0 & \text{if } p < 1-\alpha \, , \\ 1 & \text{if } p \geq 1-\alpha \, . \end{cases}$$

It can be shown that equation (16) can be written in this case as

$$\rho_{\mathrm{VaR}}(Y) = -F_Y^{-1}(1-\alpha) \, .$$

Similarly for any given confidence level α we define the risk measure *expected shortfall*, also known as conditional VaR, and denote it with the symbol ρ_{ES}, as the distortion measure generated by

$$g_{\mathrm{ES}}(p) = \begin{cases} \dfrac{p}{1-\alpha} & \text{if } p < 1-\alpha \, , \\ 1 & \text{when } p \geq 1-\alpha \, . \end{cases}$$

In the case of expected shortfall equation (16) simply becomes

$$\begin{aligned} \rho_{\mathrm{ES}}(Y) &= E\left[-Y | Y \leq F_Y^{-1}(1-\alpha)\right] \\ &= E\left[L | L \geq \rho_{\mathrm{VaR}}(\alpha)\right] \, . \end{aligned}$$

In other words, the expected shortfall is the average of all losses that are larger than the value at risk at the given confidence level.

2.4 Severability of distortion risk measures

In order to be able to compute the risk components with respect to a distortion risk measure we need to show how to create an associated linear functional for any given portfolio gain Y. Let us proceed as follows: for any given percentile p, with $0 \leq p \leq 1$, we define the measurable set Ω_Y^p, subset of Ω, counter image of p as

$$\Omega_Y^p = Y^{-1}(F_Y^{-1}(p)) \, .$$

In other words any risk-factor vector ω in Ω_Y^p has a return of $F_Y^{-1}(p)$, i.e.

$$F_Y^{-1}(p) = Y(\omega) \qquad \text{for all } \omega \in \Omega_Y^p \, .$$

As a consequence we also have that

$$F_Y^{-1}(p) = \frac{1}{P(\Omega_Y^p)} \int_{\Omega_Y^p} Y(\omega) \mathrm{d}P(\omega)$$
$$= E\left[Y | Y = F_Y^{-1}(p)\right] ,$$

where, as shown in the first of these equations, the conditional expectation is computed with respect to the P probability measure. Using this last equation we can rewrite expression (16) for the distortion risk measure ρ_g as

$$\rho_g(Y) = \int_0^1 E\left[L | Y = F_Y^{-1}(p)\right] \mathrm{d}H(p) ,$$

which, again, emphasises how any distortion risk measure can be seen as an average of the expected loss on a distorted probability. In view of this equation we define the associated linear functional A_Y as

$$A_Y[Z] = \int_0^1 E\left[Z | Y = F_Y^{-1}(p)\right] \mathrm{d}H(p) , \qquad (17)$$

for any random variable $Z \in \mathcal{Y}$. Of course, by definition we have that

$$\rho_g(Y) = A_Y[L]$$

and that A_Y is linear because of the linearity of conditional-expectations and that of regular integrals. As a consequence the functional A_Y just defined can be used in computing the risk contributions of $\rho(Y)$.

In the case of distortion risk measures $A_Y[Z]$ represents the distorted weighted average of Z over the portfolio losses that determine the risk measure ρ_g. Since the average in equation (17) is biased towards those returns that are most significant for the risk measure, we can say that $A_Y[Z]$ is the average *at* the risk measure.

Linear functionals for VaR and expected shortfall In the case of value at risk at a confidence level α equation (17) naturally leads to

$$A_{\mathrm{VaR}}[Z] = E\left[Z | L = \rho_{\mathrm{VaR}}(\alpha)\right] , \qquad (18)$$

i.e. the average Z conditional to have a loss equal to the value at risk. Similarly for the expected shortfall at the same confidence level, equation (17) becomes

$$A_{\mathrm{ES}}[Z] = E\left[Z | L \geq \rho_{\mathrm{VaR}}(\alpha)\right] , \qquad (19)$$

that is the average value of Z conditional to a loss that is not smaller then the value at risk for the same confidence level.

Non uniqueness of the associated linear functional Even though the functional defined by equation (17) is very reasonable, and as we shall see later is the best choice that can be made, there is room for other definitions. For example take a bounded positive random variable W, a weight of some sort, and define for each p the normalisation factor

$$N_Y(p) = E\left[W | Y = F_Y^{-1}(p)\right] .$$

Then we can create a new linear functional A_Y^W as the weighted average of Z with respect to the random variable W:

$$A_Y^W[Z] = \frac{1}{N_Y(p)} \int_0^1 E\left[W \cdot Z | Y = F_Y^{-1}(p)\right] \mathrm{d}H(p) . \qquad (20)$$

Obviously A_Y^W is also a linear functional and it is easy to show that

$$A_Y^W[L] = \rho(Y) ,$$

so that it also implies the severability of the distortion risk measures. Note that, differently to the definition of distortion risk measure, the weight here is not applied to portfolio loss, the dependent variable, instead it is defined on market outcomes i.e. in the risk-factor space.

One motivation behind definition (20) is that we can compute the risk contributions providing different weights to different risk drivers. For example, in this way we can compute the risk contributions conditional to certain views in a way that is similar in spirit to the motivations of reference [13]. We could for example define

$$W = \begin{cases} 1 & \text{if } X_1 > \eta, \\ 0 & \text{otherwise} , \end{cases}$$

In this case we would obtain a risk decomposition biased for the risk driver X_1 to be bigger than a certain value η, with the consequence of embedding this view into the computation of the risk contributions.

Since for any distortion risk measure we may have multiple associated linear functionals for the same pricing function, it is possible to compute risk contributions in different ways, a remarkable feature of the risk-decomposition framework proposed in this paper.

3 The risk-decomposition framework

3.1 Framework definition

In subsection 1.2 we briefly outline the basic concepts behind the risk-decomposition framework and we derive an informal representation of the risk components in the special case of the standard deviation. In this section we provide a generalisation of that concept and give a formal definition of a risk-decomposition framework that can be used in a large number of practical cases of interest.

When computing risk, as already described in subsection 2.1, we consider the market transitioning a period that starts in the past, or the present, at t_0 and ends at a future date t_1. We then assume the market state at t_1 to be described by a random risk-factor vector ω belonging to the universe of possible outcomes Ω. Even thought at t_0 we do not know the market outcome ω, we assume that there is a preferred state ω^0, known as the *equilibrium state* (sometimes called the *reference market state*), that would theoretically happen if there were no news in the market. Thus ω^0 describes the risk-factor evolution as the sole consequence of the passage of time. Also note that ω^0 is not necessarily the market state at t_0 since a deterministic evolution of the risk factors is sometimes presumed.

We define the random variable Y to be the portfolio return, or the cash profit and loss, at time t_1 with respect to its value at time t_0. At time t_0 the *ex-ante* return Y is a random variable and we might use the properties of the probability space $(\Omega, \mathcal{F}, P)$ to approximate financial quantities, such as the risk measures defined in subsection 2.1.

At the end of the evaluation period, i.e. at time t_1, we transitioned from the ex-ante world to the *ex-post* one and the actual market outcome ω^1 becomes known. As a consequence we can observe the realisation of all random variables such as the realised return y, or the realised risk-drivers values $x_1, \ldots, x_d$.

Consider now the transition from t_0 to t_1 from the point of view of performance measurements in the ex-post case. The single number y represents the ex-post return from t_0 to t_1 and is the major quantity of interest. However usually, at least in asset management, we are not satisfied with the knowledge of the return and use one of the many different techniques (see reference [3] for more details on the subject) to split this ex-post return into additive components.

Typically we write the realised return y as the sum of k *performance contributions* $z_1, \ldots, z_k$:

$$y = z_1 + \ldots + z_k \,. \tag{21}$$

There are many different methods to define portfolio components, in all of them the sum of the return contributions equals the portfolio return. We use the term (ex-post) *performance-contribution schema* (see, again, reference [3]) to denote any procedure that allows us to split the ex-post performance into additive components as in equation (21).

Let us now look at the point of view of the risk manager: the ex-ante world. Even thought the performance measurement is typically computed after the time t_1, at t_0 we can altogether apply it to all the ex-ante realisations of the random return Y. In this way, using the given performance-contribution schema, for a given market outcome $\bar{\omega}$ we compute the split of $Y(\bar{\omega})$ into k contributions $Z_1(\bar{\omega}), \ldots, Z_k(\bar{\omega})$, i.e. we have

$$Y(\bar{\omega}) = Z_1(\bar{\omega}) + \ldots + Z_k(\bar{\omega}) \,. \tag{22}$$

Since we can write equation (22) for every $\bar{\omega} \in \Omega$ then the portfolio return contributions $Z_1, \ldots, Z_k$, defined above are random variables.

Formally, given a random portfolio return Y we define an associated (ex-ante) *performance-contribution schema* $\mathcal{Z}$ as a set of random variables $Z_1, \ldots, Z_k$, so that

$$Y = Z_1 + \ldots + Z_k \,. \tag{23}$$

Note that while in equation (4) the random variables U_i represented the portfolio holdings, in the component-return schema (23) we may have random variables that have nothing to do with the holding returns.

At this point it is effortless to derive the risk contributions for a severable risk measure, according to given performance-contribution schema, following the same reasoning as subsection 1.2. For any portfolio return Y, a severable risk measure ρ and a performance-contribution schema $\mathcal{Z}$, we can write the risk-decomposition equation as

$$\rho(Y) = A[L] = C_1 + \ldots + C_k\,, \tag{24}$$

where A is the associated linear functional and the risk contributions C_i's are computed as

$$C_i = -A[Z_i] \qquad \text{for } i = 1, \ldots, k\,. \tag{25}$$

In the interpretation where the operator A is the average at the risk measure, the risk contributions of equation (25) have the meaning of the average losses associated to the performance-contribution schema $\mathcal{Z}$.

The two main ingredients to the risk contribution framework identified in subsection 1.2 are now well defined: they are the linear functional associated to a severable risk measure and the ex-ante performance-contribution schema. Given these two elements it is always possible to split risk into additive components. Nonetheless we need relevant associated linear functionals and pertinent performance-contribution schemata to obtain meaningful risk contributions.

In subsections 1.1 and 1.2 we show how the standard deviation is severable with respect to any random return obtained as the sum of the holding returns. Similarly in subsection 2.4 we constructed at least two examples of fitting associated linear functionals for any distortion risk measure with respect to a given random return. In the remainder of this section we describe two important performance-contribution schemata that allow us to define two very useful, although different, types of risk decomposition. Finally in section 4 we delineate a performance-contribution schema to be used when defining the risk components of non-linear portfolios.

3.2 Performance and risk attribution

There are many different ways to compute performance attribution, all of them quantify the excess return of a portfolio against its benchmark and attribute it to the different stages of the investment process. Risk attribution is similar in kind: given the portfolio risk we identify its different sources and assign them to the appropriate stages of the investment process.

In the environment of asset management a fund performance is usually compared with that of a benchmark (see reference [3] for more details), hence at the end of the period we consider the ex-post active return y, i.e. the return of the portfolio minus that of the benchmark. Depending on the circumstances there may be different people, or departments, involved into different the stages of the investment process. In *performance attribution* one seeks to split the active return into components and assign them to the actors that were responsible for them.

For example, very often, we may have three separate departments: one in charge of sector allocation, another of equity selection and they may all invest in foreign assets so that there is a third group handling the currency-exchange rates. In this case we assume to have available some performance-attribution techniques to write the ex-post active return y as

$$y = z_{\text{alloc}} + z_{\text{select}} + z_{\text{currency}}\,, \tag{26}$$

i.e. the sum of an allocation component z_{alloc}, a selection contribution z_{select}, and the performance z_{currency} that can be attributed to the exchange-rate evolution. The exact details on how the different performance components are computed in equation (26) are not relevant here and we can just assume that we have a procedure to calculate these contributions.

Computation of active risk For each past time period equation (26) provides a way to compute the portfolio-return attribution. At the current date we do not know what the active return on the next period will be so that we use the random variable Y to model it. However, given a market outcome $\bar{\omega}$, using the performance-attribution method that lead us to equation (26), we can write the loss as

$$L(\bar{\omega}) = -\left[Z_{\text{alloc}}(\bar{\omega}) + Z_{\text{select}}(\bar{\omega}) + Z_{\text{currency}}(\bar{\omega})\right] . \tag{27}$$

Since the previous equation holds for every $\bar{\omega}$, the ex-ante attribution components Z_{alloc}, Z_{select}, and Z_{currency} are random variables. Hence given a severable risk measure ρ we can apply the risk-decomposition framework outlined in the previous subsection and write the risk attribution of the portfolio for the risk measure ρ as

$$\begin{aligned}\rho(Y) &= A[L] \\ &= A[-Z_{\text{alloc}}] + A[-Z_{\text{select}}] + A[-Z_{\text{currency}}] \\ &= C_{\text{alloc}} + C_{\text{select}} + C_{\text{currency}} ,\end{aligned}$$

with the obvious definitions of the attribution components C_{alloc}, C_{select}, and C_{currency}. Using the risk-attribution figures thus computed, we can complement the performance-attribution analysis with the risk-attribution components and determine which part of risk is taken in the different stages of the investment process.

3.3 Linear exposure schema

Similarly to the risk-attribution example, the risk-decomposition framework works very well together with any performance-contribution schema. In this subsection we derive the *classical* risk-contribution results and show them to be a particular case of our risk-decomposition framework.

Consider a portfolio with a stochastic return Y that has a linear dependence from k market factors U_i's, i.e. we can write

$$Y = b_1 \cdot U_1 + \ldots + b_k \cdot U_k , \tag{28}$$

where $b_1, \ldots, b_k$ are the linear-factor exposures.

As simple as expression (28) may appear, as discussed in reference [12], it represents a number of interesting cases:

- the portfolio return seen as the sum of sectors, or constituents, when the exposure b_k are appropriates weights and the factors U_k's are the returns of sectors or the constituents;
- the portfolio performance approximated by linear-factor models;
- the bond return approximated using carry, duration, and convexity;
- the approximation of an option-premium performance with the standard greeks (delta, gamma, and so on).

Notice that when equation (28) is obtained as the Taylor expansion of a non-linear pricing function the meaning of the exposures b_i's may become blurred. For example in the case of the return of a plain/vanilla equity option, linearised using the standard delta/gamma approximation, when U_1 represents the delta contribution to the return and U_2 provides the gamma part, both the exposures b_1 and b_2 are representative of the underlying stock-price variations. Hence the actual exposure to the stock-price changes is unclear.

Traditional computation of risk contributions According to the traditional risk-decomposition approach for a portfolio with a return given by equation (28) (see for example reference [20]), we can compute the risk contributions for a risk measure R that is positively homogeneous when viewed as a function of the risk exposures b_i's. Indeed, according to the Euler's homogeneous function theorem, under this hypothesis the risk measure R satisfies[3]

$$R(b_1, \ldots, b_k) = b_1 \frac{\partial R}{\partial b_1} + \ldots + b_d \frac{\partial R}{\partial b_d} , \tag{29}$$

which provides the basis to write the risk contributions C_i's as

$$C_i = b_i \frac{\partial R}{\partial b_i} = b_i M_i . \tag{30}$$

The risk components computed in this way have been used successfully by academics and practitioners alike.

[3] We use the letter R for the risk measure as a function of the exposures to distinguish it clearly from the functional ρ.

The term

$$M_i = \frac{\partial R}{\partial b_i}\,, \qquad (31)$$

known as the *marginal risk* of the i-th market factor, can be viewed as the infinitesimal increment of risk with respect to the exposure b_i. In this way marginal risk denotes what is the expected additional risk that we would take if we were to increase infinitesimally the b_i portfolio exposure.

Also note that, differently from equation (24), the Euler's decomposition given by equation (29) does not provide a straightforward link between the ex-post performance measurement and the ex-ante risk.

Finally, in order to obtain equation (29) from Euler's theorem, we need to assume the homogeneity of the risk measure with respect to the exposures b_i's. This hypothesis might not be satisfied when liquidity risk is taken into account (see, for example, reference [5]). On the other hand, as noted earlier, the risk-decomposition framework defined in this paper is free from such an assumption and can be used also in presence of liquidity risk.

Canonical severable risk measures In the risk-decomposition framework defined earlier we compute contributions without assuming the a-priori homogeneity of the risk measure. Indeed for the linear exposure schema (28), given a severable risk measure ρ with an associated linear functional A, we can compute the risk contributions C_i as

$$C_i = -b_i\, A\left[U_i\right]\,.$$

Of course the contributions computed for a general functional A are typically different from those computed using equation (30). We call a positive-homogeneous risk measure *suitably severable* if there exist an associated functional $\tilde{A}$ that satisfies

$$\frac{\partial R}{\partial b_i} = -\tilde{A}\left[U_i\right] \qquad \text{for } i = 1, \ldots, d\,. \qquad (32)$$

Since $\tilde{A}$ is generally accepted by the market as the *best'* candidate for the job, we call it the *canonical functional*. The term *canonical* used here has a similar meaning as the term *suitable* used in reference [16].

Combining equations (31) and (32), for each risk driver indexed by i, with $i = 1, \ldots, d$, we have

$$M_i = -\tilde{A}\left[U_i\right]\,, \qquad (33)$$

i.e. a relationship of marginal risk and the canonical functional.

By definition the risk contributions computed for a suitably severable risk measure matches those computed using the traditional approach. For this reason the canonical functional should be regarded as distinctive among all possible associated linear functionals. Regardless one may still want, in special circumstances, to compute the risk contributions with a functional that is not canonical.

In reference [17] it is shown for spectral risk measures, however the derivation can be easily adapted to distortion risk measures, that marginal risk can be computed as

$$\frac{\partial R}{\partial b_i} = -\int_0^1 E\left[U_i | Y = F_Y^{-1}(p)\right] \mathrm{d}H(p)\,.$$

As a consequence the canonical functional for distortion risk measures is that already defined in equation (17). Similarly it is also easy to show that the standard-deviation is suitably severable with a canonical functional given by equation (9).

4 Non-linear portfolios

In subsection 3.3 we consider the case of a pricing function that has a linear dependence from the market factors. In this section we examine a non-linear pricing function, hence we assume the portfolio return Y, or equivalently the loss L, to have a non-linear exposure to the risk drivers. Recall that the risk drivers $X_1, \ldots, X_d$, are d random variables that can be used in conjunction with a deterministic loss pricing function f to obtain the random loss L as in equation (11).

Consider, for example, a portfolio composed by a single plain/vanilla equity option. In this case we can choose the pricing function f to be the Black-Scholes-Merton formula and, neglecting the effect of interest rates, we might just have two risk drivers:

1. The equity relative return X_s across the observation period
2. and the volatility increment X_σ.

Another instance that one can take into consideration is a fixed-income portfolio of exotic bonds where the non-linear terms are especially important. The most important case to consider, however, is that of a linearly-hedged portfolio, i.e. a portfolio in which all the linear exposures have been removed by carefully chosen hedging instruments. In this last case it is crucially important to keep track of the non-linear terms and to properly account for their contributions to risk.

In the next subsection we propose to use a the *first-degree projection schema* defined in reference [10]. This performance-contribution schema splits the ex-ante portfolio loss into risk-driver contributions. We concede that more than one schema may be used for the same portfolio to provide as much actionable information as possible to the portfolio manager. For example one could also use the complete projection schema (also defined in reference [10]) or machine-learning schema such as that described in paper [4].

4.1 Projection schema for a non-linear portfolio

Recall that in section 3 we define the equity market state ω^0 as the outcome that would occur if there was no new information in the financial markets, i.e. due purely to the passage of time. The risk-driver values corresponding to the risk factors ω^0 have the meaning of the equilibrium risk drivers. Hence at t_0 we can determine the reference risk drivers[4] as

$$x_1^0 = X_1(\omega^0)\,, \qquad \dots \quad , \; x_d^0 = X_d(\omega^0)\,.$$

In general the equilibrium risk drivers are different from the expected one according to the P probability measure. For simplicity of the following arguments we assume a null *carry loss*, i.e.

$$f(x_1^0, \dots, x_s^0) = 0\,, \tag{34}$$

as the loss that occurs when the market does not deviate from equilibrium. When the calendar term is not null we can compute its value and then redefine the loss function f so that equation (34) is satisfied (see, for example, reference [4]).

We need to define a performance-contribution schema that helps us to identify the loss contributions coming from the different risk drivers. Again there might be different ways of achieving this result, for example by computing the average loss conditional to the risk drivers being the equilibrium one. Using an approach of this kind we could, compute the risk-driver contributions by using the Hoeffding decomposition, as it is performed in reference [15] for the case of credit risk. While such an approach would provide in some way a reasonable definition of risk contributions, we decide to discard it because it is not compatible with a performance-contribution schema that is independent from the ex-ante probability space. Using such an approach indeed would be more problematic to compare the ex-ante risk contributions with the ex-post performance components.

Consistently with the first-degree projection schema, we propose to define the *projected risk-driver losses* as the losses that would occur if the market moved only for each single risk driver separately:

$$\begin{aligned} Z_1 &= f(X_1, x_2^0, \dots, x_d^0)\,,\\ Z_2 &= f(x_1^0, X_2, \dots, x_d^0)\,,\\ \dots &= \dots\\ Z_d &= f(x_1^0, x_2^0, \dots, X_d) \end{aligned}$$

(again, see reference [10] for a more formal derivation of this schema). Note that the risk-diver loss Z_i depends only on the evolution of the i-th risk driver X_i, however in a totally non-linear way. We then define the *cross-driver loss*

$$Z_* = L - (Z_1 + \dots + Z_d)\;,$$

as the loss that is not explained by the movement of any single risk driver by itself and that can be attributed to the co-movement of multiple risk-drivers at the same time. As shown in reference [10], we can show that the cross-driver term Z_* is given by the sum of all the losses

[4] Note the standard use of lower-case letters for deterministic values and that of upper-case letters for the random variables

due to the displacements of two or more risk drivers at the same time.

Because of the above definitions we have

$$L = Z_1 + \ldots + Z_d + Z_* ,$$

so that the set of random variables $\mathcal{Z}$=(-Z_1, ..., -Z_d, -Z_*) constitutes a valid ex-ante performance-contribution schema for the loss L.

Using the results of section 3, for any severable risk measure ρ, with an associate functional A, we can write the risk decomposition as

$$\begin{aligned} \rho(Y) &= A[L] \\ &= A[L_1] + \ldots + A[L_d] + A[L_*] \\ &= \ C_1 \ + \ldots + \ C_d \ + \ C_* , \end{aligned}$$

with the obvious definitions of $C_1, \ldots, C_d$, and C_*. To this equation we can give the following interpretation: risk is split into the sum of the average projection losses at the risk measures for each risk driver, plus the average loss at risk that cannot be explained by the movement of any single risk driver by itself.

In the special case of the canonical functional for the expected shortfall, for example for the first risk driver, we have

$$C_1 = E\left[f(X_1, x_2^0, \ldots, x_d^0) \,|\, L \geq \rho_{\mathrm{VaR}}\right] .$$

Hence the risk component for the first risk driver is given by the average, for all events that have a loss larger than the corresponding VaR, of the projected loss that would happen if all risk drivers but the first were to remain at equilibrium. This is one of the significant results of the current paper as provides an intuitive definition of risk-driver contributions for non-linear portfolios.

4.2 Risk-driver exposures

In equation (28) we write the portfolio return as a linear combination of the market factors and identify their corresponding exposures by the factor coefficients. A non-linear portfolio however cannot be written as a linear combination of the risk drivers so that we do not have a straightforward definition of exposure. In this subsection we define a concept of risk-driver exposure so that we can write an equation similar to expression (28) for the portfolio risk drivers.

Marginal risk drivers Inspired by equation (33), for each risk driver indexed by i, with $i = 1, \ldots, d$, given a severable risk measure with an associated functional A, we define the *marginal risk driver* as

$$M_i^X = A[X_i] .$$

The above definition is a generalisation of the concept of marginal risk, that is however only available for portfolios that have a linear dependence from the risk-driver exposures. Differently from the traditional definition of marginal risk, the concept of marginal risk driver defined above can be applied to any non-linear portfolio and it can also be used when we are including liquidity risk in the computations.

In the example considered at the beginning of this section, where X_S is the percentage increase of the equity underlying a plain/vanilla option, $M_S^X = A[X_S]$ is the average equity return at the risk measure. Similarly, in the same example, $M_\sigma^X = A[X_\sigma]$ represents the average volatility increment at risk.

For a given risk driver X, when the risk measure considered is the expected shortfall and A is the canonical functional, we have

$$M^X = E\left[X \,|\, L \geq \rho_{\mathrm{VaR}}\right] ,$$

which has a very intuitive interpretation as the average risk driver for portfolio losses higher than the corresponding value at risk.

Risk-driver exposures Assuming that $M_i^X \neq 0$ for all $i = 1, \ldots, d$, using the definition of marginal risk drivers we can express the risk measure as

$$\begin{aligned} \rho = &\ \ C_1 \ + \ \ldots \ + \ C_d \ + \ C_* \\ &= B_1 \cdot M_1^X + \ldots + B_d \cdot M_j^X + C_* , \end{aligned}$$

where the *risk-driver exposures* B_i's are defined as

$$B_i = \frac{C_i}{M_i^X} = \frac{A[Z_i]}{A[X_i]} , \tag{35}$$

for all $i = 1, \ldots, d$. In the above equation the term C_* contains all contributions coming from different risk drivers, hence it cannot be expressed in terms of a single risk-driver exposure.

Notice that the risk-driver exposures B_i's are real numbers. Also, since the associated functional A depends on the risk measure, we are going to have different risk-driver exposures for different risk measures. When risk is viewed as a function $\tilde{R}$ of the risk-driver exposures B_i's,

$$\tilde{R}(B_1, \ldots, B_d, C_*) = B_1 \cdot M_1^X + \ldots + B_d \cdot M_j^X + C_* ,$$

we also have

$$M_j^X = \frac{\partial \tilde{R}}{\partial B_j} \qquad \text{for } j = 1, \ldots, d ,$$

to be compared with equation (31).

The risk-driver exposures look like good candidates for approximating the loss around the risk measure. Thus we define the *linearised loss at the risk measure* $\tilde{L}$ as

$$\tilde{L} = B_1 \cdot X_1 + \ldots + B_d \cdot X_d + C_* ,$$

i.e. the loss that one would expect for a linear dependence of the portfolio from the risk drivers with coefficients B_i's. Notice that, while in general we have $\tilde{L} \neq L$, however it is easy to show that

$$\rho(L) = A[L] = A[\tilde{L}] \ \neq \rho(\tilde{L}) ,$$

i.e. the linearised loss at the risk measure has the same average at risk as the original loss itself (however it may have a different risk).

Risk-driver exposures not only can be used to quantitatively understand how a certain risk measure depends from a specific risk driver, but also they are especially useful when there is the need to effectively hedge a non-linear portfolio.

4.3 Hedging of non-linear portfolios

We consider in this subsection one of the main reasons to compute the risk-driver contributions: portfolio hedging. For simplicity we start by considering the portfolio hedging with respect to a single risk driver and consider later the case of multiple risk drivers. Suppose we have a portfolio with a loss function $f(x_1, x_2, \ldots, x_d)$ and a severable risk measure ρ. Assume that the risk contribution of the first risk driver C_1 is positive and that we want to hedge it. In the case where C_1 is not positive the first risk driver is already diversifying the portfolio and there is no need for extra hedging. In some other cases we do have a positive risk contribution from a certain risk driver, however we may decide not to hedge it since the same contribution is also an important source of return.

In order to hedge the given portfolio we purchase on the market an instrument that depends mainly from the first risk driver. Let's assume that there is such an instrument and denote with $h(x_1)$ the return pricing function of the hedging instrument so that $h(X_1)$ is its random return. Given the number of shares r_h of the hedge instrument, neglecting liquidity risk, we can write the hedged-portfolio return as

$$Y_h = Y + r_h \cdot h(X_1) .$$

In order to obtain the value for the hedge ratio r_H we first compute the hedge-instrument exposure at risk:

$$B_h = \frac{A[h(X_1)]}{M_1^X} ,$$

where A is the linear functional associated to the given risk measure. In other words B_h is the exposure that a unitary hedge instrument would have if it were already present in the portfolio in negligible quantities.

As described in the previous subsection we compute the (non-linear) exposure B_1 with respect to the first risk driver and the corresponding marginal risk M_1^X. When r_h is small, so that including r_h hedging instruments in the portfolio doesn't alter significantly the return profile, we can use the functional A to approximate the hedge-portfolio risk:

$$\begin{aligned}\rho(Y_H) &= \rho(Y + r_h \cdot h(X_1)) \simeq A\,[Y + r_h \cdot h(X_1)] \\ &= B_1 \cdot M_1^X + \ldots + B_d \cdot M_d^X + C_* + r_h \cdot B_h \cdot M_1^X .\end{aligned}$$

At this point we should be able to reduce the risk contribution coming from the first risk driver by setting

$$r_h = -\frac{B_1}{B_h} .$$

This hedge ratio is similar, at least in principle, to the standard hedge ratio that one computes using the linear sensitivity to the risk drivers. There are however two major differences with respect to the linear case. The first one is that the risk contribution is in principle present also for portfolios that are perfectly delta hedged. The second one is that the hedge ratio is specific to the given risk measure so that, for example in the case of the expected shortfall, it would typically vary with the confidence level.

Hedging process and portfolio risk management The portfolio hedging technique described for a single risk driver can be used as one of the steps in a more comprehensive hedging process that has the final goal to reduce the overall portfolio risk.

Indeed given an un-hedged, or partially hedged, portfolio consider the following steps:

1. compute the risk contributions from the different risk drivers
2. identify the risk driver that provides the largest risk contribution,
3. choose an appropriate hedge instrument and perform the single-risk-driver hedging as described earlier;
4. given the new hedged portfolio go to the first step and identify the next important risk contribution,
5. repeat the process until all risk contributions are reasonably under check.

Since at each iteration we approximated the hedge-portfolio risk using the linear functional, the hedging procedure just described will not be exact and would probably require some manual tweaking. In any event it is a good starting point for understanding what hedging instruments should be added to the portfolio.

5 Summary and conclusions

The computation of risk contributions is useful to complement the computation of a portfolio risk measure. Even though the theory of risk measures for a general probability distribution has been well established, so far there seems to be no agreement on the computation of risk contributions should for a portfolio with non-linear exposure to the risk drivers. Furthermore the breakdown of the risk-measure homogeneity with respect to the asset exposures, when liquidity risk is taken into considerations, undermines the standard risk-decomposition approach based on Euler's theorem.

In this paper we describe a risk-decomposition framework that generalises previous methods and provides useful results in a wide range of applications, from risk attribution to the risk decomposition of non-linear portfolios. Given a random portfolio return, the two main method ingredients are the linear functional associated to a severable risk measure and an ex-ante performance-contribution schema. More precisely given a severable risk measure, see section 1.2, and an ex-ante performance-contribution schema, see section 3.1, we can compute the risk contributions as in equation (25). For any given risk measure that is also positively homogenous with respect to the market factors, as shown in subsection 3.3, the results of our framework, when the canonical functional is used, are compatible with the standard homogeneous-function theorem results.

In section 2 we show that all distortion risk measures are severable with respect to any given portfolio return. The actual class of severable risk measures is even larger and includes, for example, the standard deviation. It is still an open question, to be investigated in future works, to precisely understand how large is the set of severable risk measures.

Another important ingredient of the risk-decomposition framework is the availability of a performance-contribution schema (see subsection 3.2 for more details). In this paper we advocate the use of the first degree projection schema as defined in subsection 4.1. Other possible schemata include the portfolio return linearly dependent on certains market factors, as shown in subsection 3.3; any performance attribution schema, as defined in 3.2; also, possibly for portfolios with a high non-linear dependence on the risk drivers, the second-degree projection schema described in reference [10]. Finally when it is necessary to compute the complete dependence of risk from the risk drivers (and their compounding effects) we should use the complete projection schema described in reference [10].

Even though all the contribution schemata analysed in this paper were designed with performance in mind, we could also use a schema that is created just for the risk-decomposition purpose. For example reference [4] describes an ex-ante schema that can also be used to compute risk contributions for complex non-linear portfolios.

In this paper we describe the risk-decomposition framework in the environment of market risk, however, the same methodology could also be used for credit risk. For example in reference [15] the random credit loss is split into contributions using the Hoeffding decomposition. Given a severable risk measure we could use that decomposition as a performance schema and apply our framework to compute the credit risk contributions.

The risk-contribution framework described in section 3 allows us to go beyond the traditional concept that there is only one way to perform the risk decomposition of a portfolio. Instead, given a portfolio and a risk measure we can compare the results obtained by using more than one performance-contribution schema and different associated linear functionals; in this way we obtain many different ways of how risk can be split into contributions.

Finally, since the risk-decomposition framework defined here does not rely on the homogeneity of the risk measure, it can also be used when market risk and liquidity risk provide a common risk measure (see for example reference [14]).

References

[1] Carlo Acerbi. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking and Finance*, 26, p1505-1518, 2002. 6, 7

[2] Alexandre Adam, Mohamed Houkari, and Jean-Paul Laurent. Spectral risk measures and portfolio selection. *Journal of Banking & Finance*, Volume 32, Issue 9:1870–1882, 2008. 6

[3] Carl Bacon. *Practical Portfolio Performance Measurement and Attribution*. Wiley Finance, second edition, 2008. 9, 10

[4] Davide Borrello. Non-linear risk-driver contribution of portfolio scenarios. In *Risk and Performance Attribution*, Quantitative Research Series. StatPro website, December 2015. Previews not yet available. 5, 13, 17

[5] Acerbi Carlo and Scandolo Giacomo. Liquidity risk theory and coherent risk mea- sures. *Quantitative Finance*, 8/7:681–692, 2008. 6, 12

[6] Dieter Denneberg. *Non-additive measure and integral*. Kluwer Academic Publishers, 1994. 6

[7] Ekaterina N. Sereda et al. Distortion risk measures in portfolio optimization. In John B. Guerard Jr., editor, *Handbook of Portfolio Construction*. Springer Link, 2010. 4, 5, 6, 7

[8] John Hull. *Risk Management and Financial Institutions*. Wiley Finance, third edition, 2012. 2

[9] Philippe Jorion. *Value at Risk*. Mc Graw Hill, 2000. 2

[10] Marco Marchioro. Projection performance contributions of non-linear portfolios. In *Risk and Performance Attribution*, Quantitative Research Series. StatPro website, January 2017. Permanent link marchioro.org/papers/. 13, 17

[11] Jose Menchero and Ben Davis. Risk contribution is exposure times volatility times correlation: Decomposing risk using the x-sigma-rho formula. *The Journal of Portfolio Management*, 37(2):97–106, 2011. 2, 3

[12] Attilio Meucci. Contributions from generic user-defined factors. *Risk Magazine*, June, 2007. 2, 11

[13] Attilio Meucci. Fully flexible views: Theory and practice. *Risk*, 21(10), 2008. 8

[14] Attilio Meucci. A fully integrated liquidity and market risk model. *Financial Analysts Journal*, 68(6), 2012. 17

[15] Dan Rosen and David Saunders. Risk factor contributions in portfolio credit risk models. *Journal of Banking & Finance*, 34(2):336–349, 2010. 13, 17

[16] Dirk Tasche. Risk contributions and performance measurement. 1999. 12

[17] Dirk Tasche. Expected shortfall and beyond. *Journal of Banking & Finance*, 26:1519–1533, 2002. 12

[18] Stephen Valdez and Philip Molyneux. *An Introduction to Global Financial Markets*. MacMillman, sixth edition, 2010. 2

[19] Xianyi Wu and Jinglong Wang. On characterization of distortion premium principle. *ASTIN Bulletin*, 33(1):1–10, 2003. 6

[20] Yongli Zhang and Svetlozar Rachev. Risk attribution and portfolio performance measurement-an overview. *Journal of Applied Functional Analysis*, 1(4):373, 2006. 2, 11