Factors on Demand

Building a Platform for Portfolio Managers Risk Managers and Traders[1]

Attilio Meucci[2]
attilio_meucci@symmys.com

this version: July 28, 2010
latest version available at `www.ssrn.com/abstract=1565134`

Abstract

We introduce "factors on demand", a modular, multi-asset-class return decomposition framework that extends beyond the standard systematic-plus-idiosyncratic approach. This framework, which rests on the conditional link between flexible bottom-up estimation factor models and flexible top-down attribution factor models, attains higher explanatory power, empirical accuracy and theoretical consistency than standard approaches.

We explore applications stemming from factors on demand
- The joint use of a statistical model with non-idiosyncratic residual for return estimation and a cross-sectional model for return attribution
- The optimal hedge of a portfolio of options, even when the investment horizon is close to the expiry and thus the securities are heavily non-linear
- The "on demand" feature of FoD to extract a parsimonious set of few dominant attribution factors/hedges that change dynamically with time
- Accommodating in the same platform global and regional models that give rise to the same, consistent risk numbers
- Point-in-time style analysis, as opposed to the standard trailing regression
- Risk attribution to select target portfolios to track the effect of incremental alpha signals on the allocation process

Fully commented code supporting the above case studies is available at MATLAB Central File Exchange under the author's page

JEL Classification: C1, G11

Keywords: factor models, regression, estimation, attribution, copula, conditional distribution, random matrix theory, style analysis, optimal hedging, selection heuristics, GICS industry classification, Monte Carlo, cross-sectional models, time-series models, statistical models, factor analysis, systematic factors, idiosyncratic residuals.

[1] This article is published as Meucci, A. (2010) "Factors on Demand", *Risk*, **23**, 7, p. 84-89
[2] The author is grateful to Eva Chan, Luca Spampinato, Sridhar Gollamudi and Ninghui Liu for their helpful feedback

1 Introduction

Linear factor models are widely used by risk and portfolio managers to measure and manage the risk in their positions. In the standard framework, given a portfolio with weights $\mathbf{w}$, the model expresses the portfolio return $R_{\mathbf{w}}$ as a linear combination of a set of factors Z_k and a residual $\eta_{\mathbf{w}}$

$$R_{\mathbf{w}} \equiv \sum_{k=1}^{K} d_{\mathbf{w},k} Z_k + \eta_{\mathbf{w}}, \tag{1}$$

where $d_{\mathbf{w},k}$ are portfolio-specific coefficients that transfer the randomness of the K factors into the portfolio return. The decomposition (1) is obtained in two steps. First, the returns of all N securities are expressed in terms of the same K "systematic" factors Z_k, and security-specific "idiosyncratic" residuals

$$R_n \equiv \sum_{k=1}^{K} d_{n,k} Z_k + \eta_n, \quad n = 1, \ldots, N. \tag{2}$$

The parameters in (2) are estimated from historical observations by means of regression techniques and thus (2) is an *estimation* factor model. Then, the individual systematic-plus-idiosyncratic estimation models (2) are aggregated using the relationship $R_{\mathbf{w}} = \sum_{n=1}^{N} w_n R_n$ to provide the *attribution* factor model (1) for the portfolio return. As we see, in the standard framework, the attribution factor model is obtained from the bottom-up aggregation of the estimation factor models, see Figure 1A.

Two enhancements stem from the standard framework, one focused on attribution for portfolio management platforms, the other focused on estimation for risk management platforms, see Figure 1B-C.

Enhanced portfolio management platforms distinguish between an estimation factor model $R_n \equiv \sum_{l=1}^{L} b_{n,l} F_l + U_n$ with L systematic factors F_l and an attribution factor model $R_n \equiv \sum_{k=1}^{K} d_{n,k} Z_k + \eta_n$ with K attribution factors Z_k. Both factor models are applied to the single-security returns and the two models are connected by a linear transformation, see e.g. Meucci (2007) and Menchero and Poduri (2008). Then the single-security estimation and attribution models are aggregated to the portfolio level for risk estimation or attribution purposes respectively, see Figure 1B. Despite the enhancement, the estimation model fails to accurately measure the risk of non-linear instruments such as options, among other shortcomings.

Enhanced risk management platforms, which focus on estimation, do not fit the estimation factor model to the returns of the securities. Rather, an estimation factor model $X_s \equiv \sum_{l=1}^{L} b_{s,l} F_l + U_s$ is fitted to the S risk drivers of the securities, such as credit spreads and implied volatilities, which are non-linearly related to the portfolio return. This approach yields accurate measures for the risk of any instrument. However, the attribution step is not done. This makes it difficult to interpret and take action on the positions, see Figure 1C.

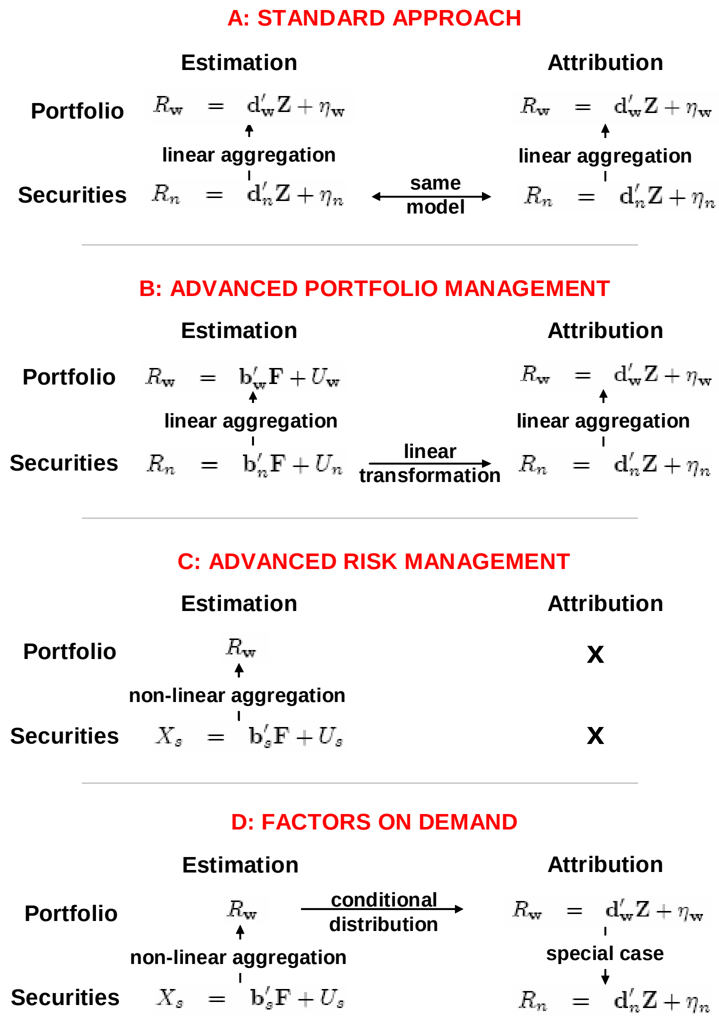

Figure 1: Factors on Demand versus other approaches to factor modeling

Factors on Demand (FoD) is a framework that incorporates the benefits of these enhancements without their shortcomings. FoD combines flexible estimation factor models that provide accurate return estimation for non-linear instruments with flexible attribution factor models, tailored to the specific portfolio under analysis. This is achieved by expanding on the above frameworks in three directions, see Figure 1D.

First, in FoD, flexible sets of estimation factors $\mathbf{F}$ are connected to another flexible set of attribution factors $\mathbf{Z}$ by a **conditional link**. This feature allows the user to perform the attribution (1) with different sets of factors $\mathbf{Z}$ without affecting the projected risk numbers of the portfolio. Furthermore, this feature yields an exact, easy to interpret, actionable linear attribution of the portfolio return $R_{\mathbf{w}}$ to the factors $\mathbf{Z}$, even when the attribution factors are intrinsically non-linearly related to the portfolio return.

Second, FoD computes the attribution (1) **top-down** directly on the portfolio return, instead of aggregating single-security attributions by means of the portfolio weights as in Figure 1A,B. In this context, the single-security attribution is a special case of the top-down attribution applied to a single-security portfolio. Due to the top-down approach, by construction the FoD attribution (1) has higher explanatory power than the bottom-up aggregation of other approaches.

Third, FoD does not rely on systematic-plus-idiosyncratic factor models, but rather on more general **dominant-plus-residual** models. This feature allows for a much more flexible choice of factors. Furthermore, the systematic-plus-idiosyncratic assumption is empirically and theoretically incorrect, as discussed in the companion article Meucci (2010), and thus the projected risk numbers that follow from such an assumption are inaccurate.

A combination of the above three features allows FoD to perform attribution with fully flexible criteria, such as r-square maximization for style analysis or CVaR minimization for hedging, and fully flexible constraints, such as the maximum number of factors allowed in the attribution.

In Section 2 we review the steps to build the advanced multi-asset-class risk management platform sketched in Figure 1 C.

In Section 3 we introduce the FoD top-down conditional attribution that extends the risk management platform to an advanced flexible portfolio management platform. While illustrating the theory, we present some applications to highlight the advantages of FoD. One application relies on FoD to jointly use a statistical model with non-idiosyncratic residual for estimation and a cross-sectional model for attribution. Another application leverages FoD to optimally hedge a portfolio of options, even when the investment horizon is close to the expiry and thus the securities are heavily non-linear. Yet another application uses the "on demand" feature of FoD to extract a parsimonious set of few dominant attribution factors/hedges that change dynamically as time elapses.

In Section 4 we present a few more applications of FoD: accommodating global and regional models that give rise to the same, consistent risk numbers; point-in-time style analysis, as opposed to the standard trailing regression approach; and risk attribution to select target portfolios to track the effect of incremental alpha signals on the allocation process.

In Appendix A.1 we present a primer on the scenarios-probabilities representation of a distribution and the generation of Monte Carlo scenarios for FoD. in Appendix A.2 we discuss risk decomposition and analysis for distributions represented as scenarios-probabilities pairs.

Fully documented code for the applications is available for download at MATLAB Central File Exchange under the author's page as "Factors on Demand".

2 The traditional steps of risk modeling

Consider a security or a portfolio whose value at the generic time t is P_t, where we assume that any potential cash flow generated is reinvested in the same security or portfolio. We denote by T the current time when the investment decision is made and by τ the horizon of the investment: the horizon is typically one day for traders, of the order of a few weeks for portfolio managers, and possibly years for private investors. The current value P_T is observable in the market, but the value at the horizon $P_{T+\tau}$ is a random variable. We denote by R the forward-looking linear return from the current time to the investment horizon

$$R \equiv \frac{P_{T+\tau} - P_T}{P_T}. \tag{3}$$

More general definitions of return apply for leveraged instruments, which are also covered by FoD, but this issue is beyond the scope of this paper. Our ultimate goal is to estimate, analyze, stress-test, and optimize the distribution of the return (3). To achieve this, we perform the following steps, refer to Figure 2.

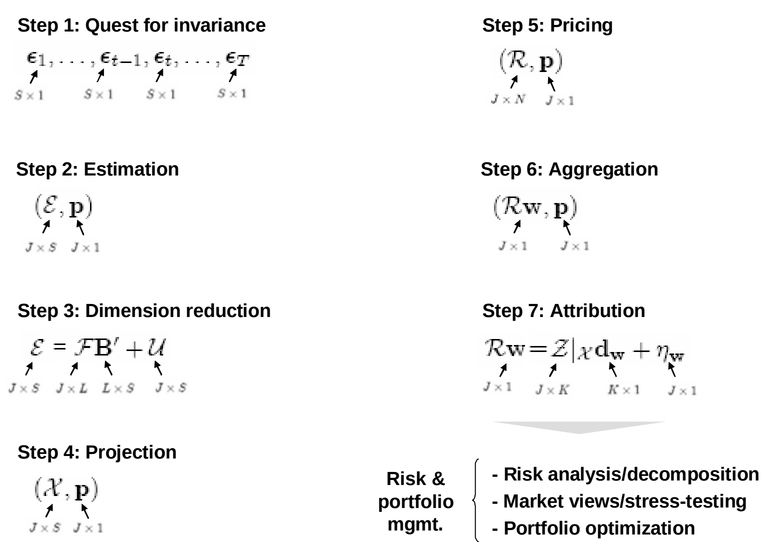

Figure 2: Road map for FoD implementation

Step 1: quest for invariance

The return of a security R is driven by one or more risk drivers, i.e. stochastic variables that determine its outcome. We denote by $\mathbf{X}_t \equiv (X_{t,1}, \ldots, X_{t,S})'$ the potentially large set of all the S risk drivers in a given market. In turn, these risk drivers are driven by the invariants, which are shocks $\boldsymbol{\epsilon}_t \equiv (\epsilon_{t,1}, \ldots, \epsilon_{t,S})'$ that are identically and independently distributed and that, due to these features, can be estimated from empirical observations, see Meucci (2005) for more details.

The most commonly assumed dynamic process that connects the risk drivers with the invariants is the random walk

$$h(\mathbf{X}_t) = h(\mathbf{X}_{t-1}) + \boldsymbol{\epsilon}_t, \tag{4}$$

where h is a suitable invertible deterministic function. Refer to Meucci (2009b) for an overview of more complex dynamics that include autocorrelations, stochastic volatility, long memory, etc.

To illustrate, consider one specific stock. Then the driver X_t is the price itself, and the invariant ϵ_t is the compounded return

$$\epsilon_t \equiv \ln X_t - \ln X_{t-1}, \tag{5}$$

which defines a random walk as in (4).

As a second example, consider a corporate bond. Then the drivers are the interest rates of a reference curve Rc_t^y, where y denotes the rate term such as one month, six months, one year, and the spreads Sp_t^y over those rates

$$\mathbf{X}_t \equiv \left(Rc_t^{1m}, \ldots, Rc_t^{30y}, Sp_t^{1m}, \ldots, Sp_t^{30y}\right)'; \tag{6}$$

and the invariants $\boldsymbol{\epsilon}_t$ are the changes in curve and spreads

$$\boldsymbol{\epsilon}_t \equiv \mathbf{X}_t - \mathbf{X}_{t-1}. \tag{7}$$

Again, this is a random walk as in (4).

Finally, consider a European call option with a given strike and expiry on a given underlying. Then the drivers are the underlying, which trades at the price U_t, and the entries $\sigma_t^{m,\tau}$ of the implied volatility surface for that underlying

$$\mathbf{X}_t \equiv \left(U_t, \sigma_t^{\underline{m},\underline{\tau}}, \ldots, \sigma_t^{\overline{m},\overline{\tau}}\right)', \tag{8}$$

where $[\underline{m}, \ldots, \overline{m}] \times [\underline{\tau}, \ldots, \overline{\tau}]$ denotes the points of a grid of moneyness and time-to-expiry values. In this case the invariants are the compounded returns of the underlying and the volatility surface

$$\boldsymbol{\epsilon}_t \equiv \ln \mathbf{X}_t - \ln \mathbf{X}_{t-1}, \tag{9}$$

which is a random walk as in (4).

Step 2: estimation

This step determines the distribution of the invariants, as represented by the pdf $f_{\boldsymbol{\epsilon}}$, which can be estimated by a variety of multivariate inference techniques, which include simple historical, nonparametric, maximum likelihood, Bayesian, robust, see Meucci (2005) for an in-depth overview.

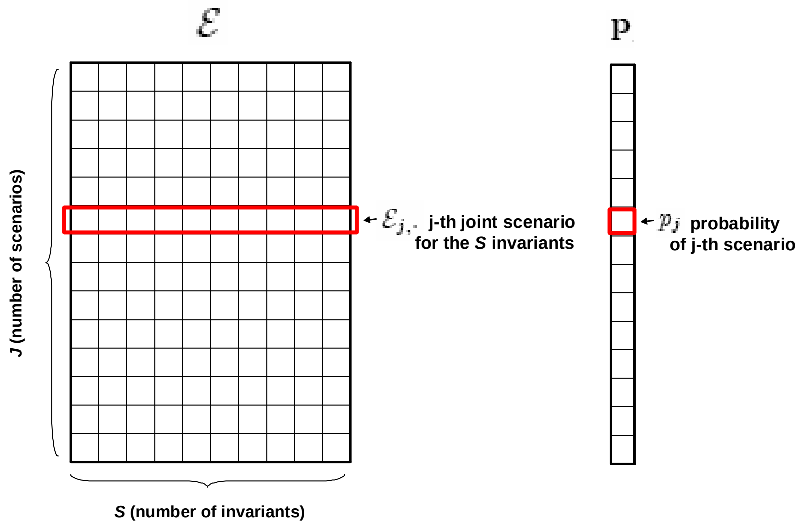

Figure 3: Scenarios-probabilities pair $(\mathcal{E}; \mathbf{p})$ correspoding to the distribution $f_{\boldsymbol{\epsilon}}$ of the invariants

The estimation process of the distribution $f_{\boldsymbol{\epsilon}}$ yields all the inputs necessary to generate a $J \times S$ panel $\mathcal{E}$ of J scenarios and a J-dimensional vector $\mathbf{p}$ of probabilities for the invariants

$$f_{\boldsymbol{\epsilon}} \Rightarrow (\mathcal{E}; \mathbf{p}) . \tag{10}$$

The generic j-th row $\mathcal{E}_{j,\cdot}$ of $\mathcal{E}$ represents a joint scenario for the S invariants and the generic j-th entry p_j of $\mathbf{p}$ represents the probability of that scenario, refer to Figure 3. Notice that the probabilities p_j are not necessarily all equal to each other.

To illustrate, $f_{\boldsymbol{\epsilon}}$ can be estimated as the empirical historical distribution. Then $\mathcal{E}_{j,\cdot} \equiv \boldsymbol{\epsilon}_t'$ are the historical observations of the invariants detected in the quest for invariance step, and $p_j \equiv 1/T$, where T is the total number of historical observations. We emphasize that the historical observations are used to represent the forward-looking distribution $f_{\boldsymbol{\epsilon}}$ of the invariants.

Alternatively, $(\mathcal{E}; \mathbf{p})$ can be Monte Carlo scenarios and probabilities corresponding to the copula and marginal distributions of the forward-looking estimated distribution $f_{\boldsymbol{\epsilon}}$ obtained as discussed in Appendix A.1.

Step 3: dimension reduction

In a multi-asset-class platform, the dimension S of the invariants can become very large. In order to properly estimate the distribution of the invariants $f_{\boldsymbol{\epsilon}}$ we reduce the dimension by imposing structure on the correlations of the invariants through an estimation linear factor model

$$\boldsymbol{\epsilon}_t \equiv \mathbf{B}_t \mathbf{F}_t + \mathbf{U}_t. \tag{11}$$

In this expression $\mathbf{F}_t$ is a L-dimensional vector of dominant estimation factors, where L is much smaller than S, the number of the invariants $\boldsymbol{\epsilon}_t$; $\mathbf{B}_t$ is a $S \times L$ matrix of coefficients that transfer the randomness of $\mathbf{F}_t$ to $\boldsymbol{\epsilon}_t$ and that may depend on time t; and $\mathbf{U}_t$ is a S-dimensional vector of residuals. This decomposition can be achieved by various time-series, cross-sectional, statistical/factor analysis, or hybrid models. In all models, including factor analysis, the residuals are correlated $\mathrm{Cor}\{U_{t,s}, U_{t,v}\} \neq 0$, see the proof in the companion article Meucci (2010). Therefore, (11) is a dominant-plus-residual rather than a systematic-plus-idiosyncratic linear factor model

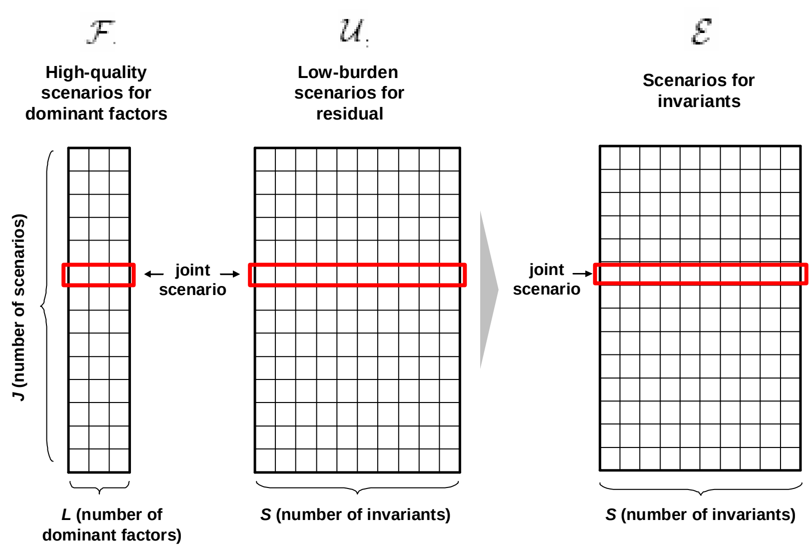

Figure 4: Scenarios-probabilities pair for invariants according to estimation factor model

Once we impose structure on the risk drivers by means of the estimation linear factor model (11), the $J \times S$ panel of scenarios $\mathcal{E}$ for the invariants (10) is efficiently obtained in terms of a $J \times L$ panel $\mathcal{F}$ that corresponds to the distribution of the non-noisy dominant estimation factors $\mathbf{F}$, and a $J \times S$ panel $\mathcal{U}$ that corresponds to the distribution of the residuals $\mathbf{U}$ as

$$\mathcal{E} \equiv \mathcal{F}\mathbf{B}' + \mathcal{U}, \tag{12}$$

The dominant panel $\mathcal{F}$ is low-dimensional and can be generated with high precision. The residual panel $\mathcal{U}$, though large, requires much less computational burden and sophistication, refer to Figure 4.

To illustrate, we estimate the distribution of the invariants by applying random matrix theory, which appears to outperform cross-sectional models for estimation purposes, see Gatheral (2008). We model the invariants $\boldsymbol{\epsilon}$, which is a S-dimensional vector, as the sum of $L \ll S$ dominant factors and a residual that in turn is the sum of $S-L$ residual factors. More precisely

$$\boldsymbol{\epsilon}_t \stackrel{d}{=} \sum_{l=1}^{L} \mathbf{b}_l F_{t,l} + \mathbf{U}_t. \tag{13}$$

In this expression $\mathbf{b}_1, \ldots, \mathbf{b}_S$ are the eigenvectors of the sample covariance of $\boldsymbol{\epsilon}$ and L is the cutoff which separates the significant non-noisy signals $\mathbf{F}_t \equiv (F_{t,1}, \ldots, F_{t,L})'$ from the noisy residual $\mathbf{U}_t \equiv (U_{t,1}, \ldots, U_{t,S})'$. For the non-noisy signals we assume marginal distributions fitted to historical observations by non-parametric kernel-smoothing and a normal copula. For the noisy residual we assume

$$\mathbf{U}_t \stackrel{d}{=} \overline{\lambda} \sum_{l=L+1}^{S} \mathbf{b}_l N_{t,l}, \tag{14}$$

where $\overline{\lambda}^2$ is the average of the smallest $S-L$ eigenvalues of the sample covariance of $\boldsymbol{\epsilon}$ and $N_{t,l}$ are standard univariate independent normal variables. Notice that the residuals are not idiosyncratic because they are correlated with each other. To generate the scenarios panels $\mathcal{F}$ and $\mathcal{U}$ and the vector of probabilities $\mathbf{p}$ please refer to Appendix A.1.

Step 4: projection

This step generates the scenarios-probabilities pair that correspond to the distribution $f_{\mathbf{X}}$ of the risk drivers $\mathbf{X}$ from the scenarios-probability pair of the invariants

$$(\mathcal{E}; \mathbf{p}) \Rightarrow (\mathcal{X}; \mathbf{p}). \tag{15}$$

In this expression $\mathcal{X}$ is a $J \times S$ panel and the generic j-th row of $\mathcal{X}$ represents one joint scenario for the S risk drivers $\mathbf{X}$. We can generate $\mathcal{X}$ because the risk drivers are driven by the invariants through dynamics such as the random walk (4).

To illustrate, consider a market of stocks and options and assume that the investment horizon is one step ahead, i.e. $\tau = 1$ in (3). Then from (5) and (9) we obtain for the risk drivers $\mathbf{X} \equiv \mathbf{X}_{T+1}$

$$\mathbf{X} = \mathbf{X}_T \star e^{\boldsymbol{\epsilon}_{T+1}}, \tag{16}$$

where $\star$ denotes the entry-by-entry multiplication and the exponential operates entry-by-entry. Therefore for the scenarios we obtain

$$\mathcal{X} = \left(\mathbf{1}\mathbf{X}_T'\right) \star e^{\mathcal{E}}, \tag{17}$$

where $\mathbf{1}$ is a $J \times 1$ vector of ones.

Step 5: pricing

This steps generates the scenarios-probabilities pair that corresponds to the distribution $f_{\mathbf{R}}$ of the securities returns from the scenarios-probabilities pair of the risk drivers

$$(\mathcal{X}; \mathbf{p}) \Rightarrow (\mathcal{R}; \mathbf{p}). \tag{18}$$

In this expression $\mathcal{R}$ is a $J \times N$ panel and the generic j-th row of $\mathcal{R}$ represents one joint scenario for the forward-looking returns $\mathbf{R}$ of the N securities in the market.

To obtain $\mathcal{R}$, we consider the security-specific deterministic pricing functions that map the S risk drivers $\mathbf{X}$ into the returns (3) of the security

$$R_n = R_n(\mathbf{X}), \quad n = 1, \ldots, N. \tag{19}$$

Then

$$\mathcal{R}_{j,n} = R_n(\mathcal{X}_{j,\cdot}), \tag{20}$$

where $\mathcal{X}_{j,\cdot}$ denotes the j-th row, i.e. the j-th joint scenario, in the risk drivers panel $\mathcal{X}$.

In our example, for a stock the single risk driver X is the price X_{T+1} at the horizon $T + \tau = T + 1$ and thus $\mathcal{R}$ is the following J-dimensional vector

$$\mathcal{R}_j = \frac{\mathcal{X}_j}{X_T} - 1, \tag{21}$$

For a call option

$$\mathcal{R}_j = \frac{C_{BS}(\mathcal{X}_{j,\cdot})}{C_T} - 1, \tag{22}$$

where C_{BS} is the Black-Scholes pricing formula for the call option, C_T is the currently traded price of the option, and $\mathcal{X}_{j,\cdot}$ is the j-th joint scenario of the risk drivers (8) at the horizon.

Step 6: aggregation

This step yields the scenarios-probabilities pair that corresponds to the distribution of the return $R_{\mathbf{w}}$ of a generic portfolio $\mathbf{w}$. Since $R_{\mathbf{w}} = \mathbf{w}'\mathbf{R}$ we readily obtain

$$\mathbf{w} \Rightarrow (\mathcal{R}\mathbf{w}; \mathbf{p}). \tag{23}$$

In this expression $\mathcal{R}\mathbf{w}$ is a J-dimensional vector. The generic j-th element of $\mathcal{R}\mathbf{w}$ represents one scenario for the forward-looking return of the portfolio $\mathbf{w}$, which occurs with probability p_j, the j-th entry of $\mathbf{p}$.

3 FoD top-down conditional attribution

In this section we describe how to perform step 7, attribution, in Figure 2.

While building the risk-management platform in steps 1-6 we introduced a few key stochastic variables, namely the invariants $\boldsymbol{\epsilon}$, which drive the dynamics of the risk drivers $\mathbf{X}$; the dominant factors $\mathbf{F}$ in the estimation factor model for the invariants $\boldsymbol{\epsilon} \equiv \mathbf{BF} + \mathbf{U}$; the securities returns $\mathbf{R}(\mathbf{X})$, which are determined by the risk drivers through pricing functions; and the return $R_{\mathbf{w}} = \mathbf{w}'\mathbf{R}(\mathbf{X})$ of an arbitrary portfolio $\mathbf{w}$.

Assume now that we want to attribute the portfolio return to K arbitrary attribution factors $\mathbf{Z} \equiv (Z_1, \ldots, Z_K)'$, i.e. yet-to-be realized random variables that are correlated with the portfolio return. We aim at performing the attribution as in formula (1), which we report here for convenience

$$R_{\mathbf{w}} \equiv \sum_{k=1}^{K} d_{\mathbf{w},k} Z_k + \eta_{\mathbf{w}}. \tag{24}$$

For instance, assume that we hold a portfolio of options and that we have used a statistical model $\boldsymbol{\epsilon} \equiv \mathbf{BF}+\mathbf{U}$ for the estimation of the invariants followed by full repricing to obtain the distribution of the portfolio return.

Suppose that we wish to attribute the portfolio return to cross-sectional industry factors. One could be tempted to run a historical regression of the options returns on the cross-sectional factors $R_{n,t} \equiv \sum_{k=1}^{K} d_{n,k} Z_{k,t} + \eta_{n,t}$ and then aggregate these single-security models to yield the portfolio-level attribution (24). However, this approach is not viable for options, because the returns of the options are not invariants. Furthermore, even if the options returns were invariants as it is the case for stocks, the newly fitted model would also create new numbers for the risk of the portfolio, different from the numbers obtained with the statistical model, which we elected as the option of choice for estimation. It would not be optimal to have two sets of VaR, standard deviation, etc. How can we reconcile such numbers?

Now consider again our portfolio of options, but suppose that we wish to hedge it with a set of products. If the products chosen are among the underlyings, a simplistic approach would compute the "deltas" of the securities and aggregate these deltas according to the portfolio weights. However, what if the hedging products are not the underlyings? And how can we take advantage of the correlations among the hedging products?

As a third challenge, suppose again that we wish to attribute our portfolio of options to industry factors, but assume that the portfolio contains only a minuscule portion of a given industry. Do we really want to see that industry among the attribution factors?

FoD solves all the above and related problems, by relying on three pillars: top-down attribution, conditional link between attribution factors and portfolio return, and dominant-plus-residual models.

In practice, let us start with a given portfolio $\mathbf{w}$ and let us focus on the aggregate portfolio residual, which depends on the attribution coefficients $\mathbf{d}$ as follows

$$\eta_{\mathbf{w}}^{\mathbf{d}} \equiv \mathbf{w}'\mathbf{R}(\mathbf{X}) - \mathbf{d}'\mathbf{Z}. \tag{25}$$

In order to perform the attribution (24) for the specific portfolio $\mathbf{w}$, we select the coefficients $\mathbf{d}$ top-down to give the distribution of the aggregate portfolio residual $f_{\eta_{\mathbf{w}}^{\mathbf{d}}}$ the most desirable features according to a fully flexible attribution criterion $\mathcal{T}$ and fully flexible constraints $\mathcal{C}$ on the coefficients

$$\mathbf{d_w} \equiv \underset{\mathbf{d}\in\mathcal{C}}{\operatorname{argmax}}\, \mathcal{T}\left(f_{\eta_{\mathbf{w}}^{\mathbf{d}}}\right). \tag{26}$$

To illustrate, we can aim at maximizing the r-square of the attribution factors by making the residual as small as possible. Then

$$\mathcal{T}\left(f_{\eta_{\mathbf{w}}^{\mathbf{d}}}\right) \equiv -\operatorname{E}\{\left(\eta_{\mathbf{w}}^{\mathbf{d}}\right)^2\}. \tag{27}$$

Furthermore, we can aim at using only the best $\widetilde{K}$ factors out of the total pool of K factors $\mathbf{Z}$, i.e. we apply the constraint $\operatorname{card}(\mathbf{d}) \leq \widetilde{K}$, where $\operatorname{card}(\mathbf{d})$ denotes the cardinality of $\mathbf{d}$, i.e. the number of non-zero entries in the vector $\mathbf{d}$. Then the attribution optimization (26) becomes

$$\mathbf{d_w} \equiv \underset{\operatorname{card}(\mathbf{d})\leq\widetilde{K}}{\operatorname{argmin}} \operatorname{E}\{\left(\mathbf{w}'\mathbf{R}\left(\mathbf{X}\right) - \mathbf{d}'\mathbf{Z}\right)^2\}. \tag{28}$$

Notice that the distribution of the residual (25) is determined by the joint distribution $f_{\mathbf{X},\mathbf{Z}}$ of the risk drivers and of the attribution factors. Therefore, regardless which criterion $\mathcal{T}$ and constraints $\mathcal{C}$ we choose, in order to perform the attribution optimization (26) we need $f_{\mathbf{X},\mathbf{Z}}$ and the scenarios-probability pair that corresponds to this distribution

$$f_{\mathbf{X},\mathbf{Z}} \Rightarrow \left(\mathcal{X}, \mathcal{Z}|_{\mathcal{X}}; \mathbf{p}\right). \tag{29}$$

In this expression $\mathcal{X}$ is the $J \times S$ panel of scenarios for the S risk drivers $\mathbf{X}$ that was generated in the projection step (15) together with the vector of the respective probabilities $\mathbf{p}$; and $\mathcal{Z}|_{\mathcal{X}}$ is a yet-to-be generated $J \times K$ panel of scenarios for the K attribution factors $\mathbf{Z}$, where the notation highlights that $\mathcal{Z}|_{\mathcal{X}}$ is generated after $\mathcal{X}$ is given. We discuss later in this section the methodology to generate $\mathcal{Z}|_{\mathcal{X}}$ in such a way that the generic j-th row of the joint panel $(\mathcal{X}, \mathcal{Z}|_{\mathcal{X}})$ represents a joint scenario for $\mathbf{X}$ and $\mathbf{Z}$ that occurs with probability p_j, the j-th entry of $\mathbf{p}$.

Assuming for now that $\mathcal{Z}|_{\mathcal{X}}$ has been generated, it becomes trivial to generate the scenarios-probabilities pair that correspond to the distribution of the residual. Indeed, from (25) we obtain

$$f_{\eta_{\mathbf{w}}^{\mathbf{d}}} \Rightarrow \left(\mathcal{R}\mathbf{w} - \mathcal{Z}|_{\mathcal{X}}\mathbf{d}; \mathbf{p}\right). \tag{30}$$

Then the top-down attribution optimization (26) can be computed numerically using the scenarios-probabilities pair (30).

This process yields the ultimate FoD attribution, which mirrors (24)

$$\mathcal{R}\mathbf{w} = \mathcal{Z}|_{\mathcal{X}}\mathbf{d_w} + \eta_{\mathbf{w}}, \tag{31}$$

where with minor abuse of notation $\eta_{\mathbf{w}}$ here also indicates the J-dimensional vector of the residual scenarios.

By construction, the FoD top-down approach attains a higher value for the attribution criterion $\mathcal{T}$ than the bottom-up approaches in Figure 1A-B-C.

Furthermore, the bottom-up approaches in general fail to satisfy the constraints $\mathcal{C}$ in the attribution optimization, as we will see in the example that follows.

Also, the portfolio return $\mathcal{R}\mathbf{w}$ on the left hand side in (31) is unaffected by the attribution factors on the right hand side: the user can choose two different attribution factor models to interpret and manage the portfolio, without affecting the portfolio risk numbers. This is possible because the generation of the scenarios for the attribution factors $\mathcal{Z}|_{\mathcal{X}}$ is conditioned on having first generated scenarios for the portfolio return, which are fully driven by the risk drivers scenarios $\mathcal{X}$.

Finally, the FoD attribution (31) is a dominant-plus-residual factor model, rather than a more restrictive systematic-plus-idiosyncratic model. The portfolio return, namely the left hand side in (31), correctly accounts for all the correlations hidden on the right hand side in the attribution factors and in the residual. The risk numbers for the portfolio such as the standard deviation are based on the left hand side only and do not need to rely on theoretically and empirically incorrect systematic-plus-idiosyncratic assumptions on the covariance of the attribution factors and of the single-security residuals.

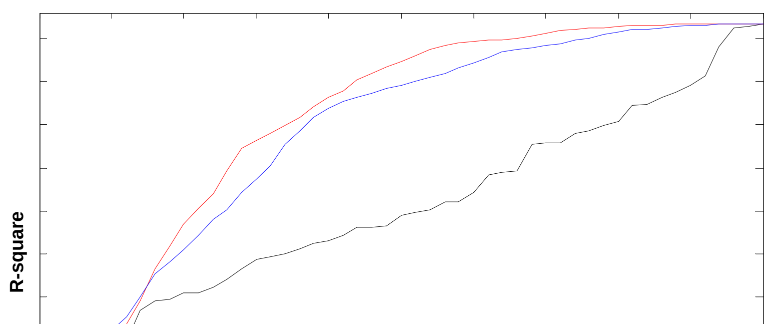

Figure 5: Performance of top-down FoD attribution as function of the number of factors

To illustrate the practical implementation of FoD, we continue with our example (27)-(28). Using the scenarios-probabilities pair (29), we reformulate the problem as follows

$$\mathbf{d}_{\mathbf{w}} \equiv \underset{\operatorname{card}(\mathbf{d}) \leq \widetilde{K}}{\operatorname{argmin}} \left\{ \mathbf{d}'\mathbf{H}\mathbf{d} - 2\mathbf{a}'\mathbf{d} \right\}. \tag{32}$$

The $K \times K$ matrix $\mathbf{H}$ in the quadratic term and the K-dimensional vector $\mathbf{a}$ in the linear term are defined in terms of the scenarios-probabilities pair as

$$H_{k,k'} \equiv \sum_{j=1}^{J} p_j \mathcal{Z}_{j,k} \mathcal{Z}_{j,k'} \tag{33}$$

$$a_k \equiv \sum_{j=1}^{J} p_j \mathcal{Z}_{j,k} \sum_{n=1}^{N} \mathcal{R}_{j,n} w_n, \tag{34}$$

where the returns scenarios $\mathcal{R}$ are a function of the scenarios for the risk drivers $\mathcal{X}$ from the pricing step (18). The cardinally constrained quadratic program (32) can be solved efficiently with heuristics such as recursive rejection routine discussed in Meucci (2005). In Figure 5 we display a case study with a potential pool of $K \equiv 50$ factors, see the FoD code available at MATLAB Central File Exchange under the author's page for more details.

Notice that by construction the top-down approach (28) or (32) has a higher r-square than the bottom-up approaches. Also, the bottom-up approaches cannot satisfy the cardinality constraint in (28) or (32) . Finally, the constraint correlates the attribution factors with the residual, and thus the resulting attribution model is not systematic-plus-idiosyncratic, but rather dominant-plus-residual.

Now we return to the generation of the scenarios $\mathcal{Z}|_{\mathcal{X}}$ that complete the scenarios-probability pair (29) which corresponds to the joint distribution of risk drivers and attribution factors $f_{\mathbf{X},\mathbf{Z}}$. We distinguish two methods, where the first is a special case of the second.

The first method applies when the attribution factors are fully determined by the S risk drivers $\mathbf{X}$

$$Z_k = Z_k(\mathbf{X}), \quad k = 1, \ldots, K. \tag{35}$$

Then $\mathcal{Z}|_{\mathcal{X}}$ is easily obtained similarly to the panel of the securities returns (18): for each factor $k = 1, \ldots, K$ we feed the J rows of the $J \times S$ panel of risk drivers $\mathcal{X}$ through the function (35), thereby obtaining the J entries of the k-th column of $\mathcal{Z}|_{\mathcal{X}}$.

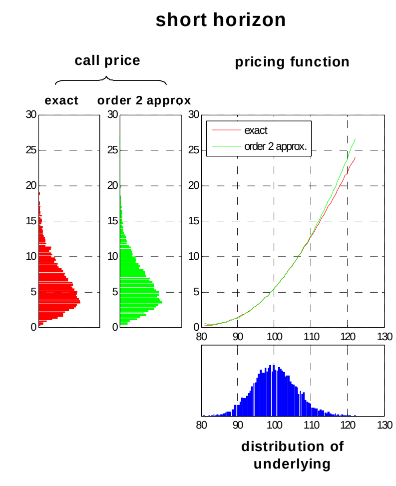

Figure 6: Delta-gamma option payoff approximation with different investment horizons

To illustrate this first method, we use FoD to hedge a portfolio of options on one stock with the stock. To solve this problem, one could be tempted to rely on the "deltas", computed according to the Black-Scholes pricing formula. However, the approximation of the option payoff provided by the delta and higher order analytics such as the gamma becomes incorrect for long investment horizons, due to the square root rule propagation of risk, see Figure 6.

Instead, we rely on a combination of FoD and full-repricing. In this context, the options returns panel on the left hand side in (31) follows from (22) and the single attribution factor on the right hand side in (31) is the return of the stock (21)

$$\mathcal{Z}|_{\mathcal{X}} \equiv \frac{\mathcal{X}_{\cdot,1}}{X_T} - 1, \tag{36}$$

where $\mathcal{X}_{\cdot,1}$ denotes the first column of the panel $\mathcal{X}$, which corresponds to the underlying (8). The single attribution coefficient d in (31) represents the weight of the stock; and the residual is the return of the hedged portfolio.

Given that we intend to control the downside of the hedged position without limiting its upside, we use the flexibility of FoD to specify the target in the attribution optimization (26) as the conditional value at risk

$$\mathcal{T}\left(f_{\eta^d_{\mathbf{w}}}\right) \equiv -\,\mathrm{CVaR}\left\{\eta^d_{\mathbf{w}}\right\}. \tag{37}$$

Therefore, we solve

$$d_{\mathbf{w}} \equiv \underset{d}{\operatorname{argmin}}\,\mathrm{CVaR}\{\mathcal{R}\mathbf{w} - d\mathcal{Z}|_{\mathcal{X}}\}, \tag{38}$$

This problem can be solved by linear programming as in Rockafellar and Uryasev (2000), to yield the FoD attribution (31). Notice that (38) and (31) provide an exact, actionable linear hedge for the portfolio return even when the options in the portfolio are close to expiry and thus heavily non-linear. Also, this methodology can be readily extended to multiple hedging products with cardinality constraint on the maximum number of hedging products.

In Figure 7 we report the number of stocks necessary to hedge single call options with different expiries using the Black-Scholes delta and using the FoD attribution (38) and (31). The code for this case study is available at MATLAB Central File Exchange under the author's page. For more details on the computations, refer to Meucci (2009a).

	100 days	150 days	200 days	250 days	300 days
FOD	5.8	5.3	5.0	4.9	4.8
BS	5.7	5.4	5.2	5.1	5.0

Figure 7: Number of stocks to hedge call options with long investment horizons

The second method to generate the scenarios $\mathcal{Z}|_{\mathcal{X}}$ that complete the scenarios-probability pair (29) for $f_{\mathbf{X},\mathbf{Z}}$ is more general, as it applies when the attribution factors $\mathbf{Z}$ are not a simple direct function of the risk drivers $\mathbf{X}$, but are statistically correlated with them. Using the identity $f_{\mathbf{X},\mathbf{Z}} = f_{\mathbf{X}} f_{\mathbf{Z}|\mathbf{X}}$, where $f_{\mathbf{Z}|\mathbf{X}}$ is the conditional distribution of the attribution factors $\mathbf{Z}$ given the risk drivers $\mathbf{X}$, we realize that in order to obtain $f_{\mathbf{X},\mathbf{Z}}$ we need to estimate $f_{\mathbf{Z}|\mathbf{X}}$. To do so, we pursue the quest for invariance for the risk drivers and then we apply conditional estimation techniques similar to those discussed in Step 2. Then for each row in the panel $\mathcal{X}$, i.e. for each scenario of $\mathbf{X}$, we generate one conditional scenario for $\mathbf{Z}$ from $f_{\mathbf{Z}|\mathbf{X}}$, thereby obtaining the $J \times K$ conditional panel $\mathcal{Z}|_{\mathcal{X}}$.

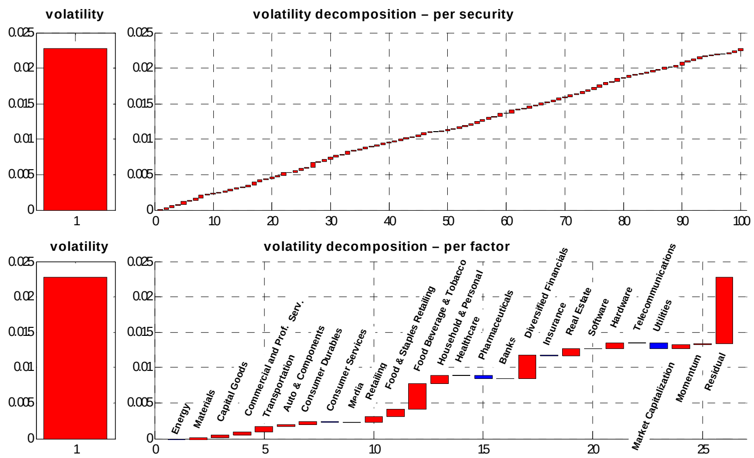

Figure 8: Risk analysis of equity portfolio with two underlying factor models: non-idio statistical and cross-sectional

To illustrate this second method, we use FoD to attribute the portfolio return to a cross-sectional factor model after the portfolio return has been estimated as in Steps 1-6 using a statistical factor model for dimension reduction.

Consider a portfolio of stocks, where the returns of the securities $\mathbf{R}$ are invariants. Assume that we use the random matrix theory statistical factor model (11) to model their distribution as $\mathbf{R} = \mathbf{BF} + \mathbf{U}$, where $\mathbf{F}$ are L dominant non-noisy signals and $\mathbf{U}$ are noisy residuals. As in (12) we generate a scenarios-probabilities pair $(\mathcal{R} \equiv \mathcal{F}\mathbf{B}' + \mathcal{U}; \mathbf{p})$ for the stock returns that corresponds to the returns distribution $f_{\mathbf{R}}$, where $\mathcal{F}$ are scenarios for the dominant non-noisy signals and $\mathcal{F}$ are scenarios for the noisy residuals. With these inputs we can compute the standard deviation of the portfolio $\sigma_{\mathbf{w}}$ and decompose it into the additive contributions from each security using the Euler identity, see e.g. Meucci (2005)

$$\sigma_{\mathbf{w}} = \sum_{n=1}^{N} w_n \frac{\partial \sigma_{\mathbf{w}}}{\partial w_n}. \tag{39}$$

In the top portion of Figure 8 we display such a decomposition in a case study that follows the above steps, refer to Appendix A.2 and to the FoD code available at MATLAB Central File Exchange under the author's page.

Now, assume that we want to attribute the portfolio return to a set of K cross-sectional factors $\mathbf{Z}$, such as the GICS industry factors. To this purpose we can apply the specific instance of FoD attribution optimization (32), which in this context maximizes the r-square of a constrained number of cross-sectional factors.

To implement FoD we need to generate the conditional scenarios $\mathcal{Z}|_{\mathcal{R}}$, where we used the fact that in the case of stocks the returns $\mathbf{R}$ are also equivalent to the risk drivers, see (21). Therefore, we would need to estimate the conditional distribution $f_{\mathbf{Z}|\mathbf{R}}$. However, assuming independence of the cross-sectional factors $\mathbf{Z}$ from the pure noise $\mathbf{U}$ we only need to estimate the conditional distribution $f_{\mathbf{Z}|\mathbf{F}}$ of the K cross-sectional factors on the few L non-noisy signals, because $f_{\mathbf{Z}|\mathbf{R}} = f_{\mathbf{Z}|\mathbf{F}}$. We apply simple regression techniques to estimate $f_{\mathbf{Z}|\mathbf{F}}$ and we generate the conditional scenarios $\mathcal{Z}|_{\mathcal{F}}$ accordingly, see more details in Appendix A.1.

With the above inputs we can solve the FoD attribution optimization (32) and obtain the FoD attribution (31). Then we can decompose the *same* $\sigma_{\mathbf{w}}$ as (39) into the additive contributions from the attribution factors

$$\sigma = \sum_{k=1}^{\widetilde{K}+1} d_k \frac{\partial \sigma}{\partial d_k}, \tag{40}$$

where the $(\widetilde{K}+1)$-st factor is the residual. In the bottom portion of Figure 8 we display this decomposition in our case study, refer to Appendix A.2 and to the FoD code available at MATLAB Central File Exchange under the author's page.

We emphasize that in this application of FoD we obtained a portfolio-specific, parsimonious cross-sectional factor model where the risk numbers follow from random matrix theory, which is a statistical factor model with non-idiosyncratic residual.

4 Further applications of FoD

In this section we present a few more applications of FoD. All the applications follow the risk modeling steps 1-6 discussed in Section 2 and run the top-down attribution optimization (26) to obtain the FoD attribution (31), which is step 7, refer to Figure 2.

4.1 Global versus regional factor models

In the equity market the returns of the stocks $\mathbf{R}$ can in first approximation be considered as the invariants. This covers Step 1 in the road map to FoD in Figure 2.

Suppose that we are focusing on one specific region, say the US. Suppose that we rely on a cross-sectional model for estimation and dimension reduction, Steps 2-3 in Figure 2. As in (11), we model

$$\mathbf{R} \equiv \mathbf{B}\mathbf{F} + \mathbf{U}, \tag{41}$$

where $\mathbf{F}$, the dominant estimation factors, are the GICS industries factor returns.

Now, suppose that we need to expand the platform to provide global coverage. When analyzing their positions, global equity portfolio managers do not appreciate the fine granularity provided by regional models. For instance, portfolio managers needs to know about exposure to "Utilities" as a whole, and not "Utilities US", "Utilities Country ABC", etc.

A possible approach to cater to both the regional and global portfolio manager would be to estimate two different factor models: a global equity model, based on a parsimonious set of aggregate factors to cater to the global manager; and a set of regional models, based on much more granular regional factors, to cater to the local managers.

This approach is suboptimal for several reasons, most notably the cost of maintaining different systems and the inconsistency of the risk analysis. To illustrate the latter, consider a US-only portfolio. We can compute the risk of this portfolio based both on the regional model and on the global model, obtaining two different sets of risk numbers.

FoD solves this dichotomy. First, we the use granular regional models for estimation: given a set $(\alpha, \ldots, \omega)$ of regions, we fit the respective regional models individually

$$\begin{aligned} \mathbf{R}^{(\alpha)} &\equiv \mathbf{B}^{(\alpha)}\mathbf{F}^{(\alpha)} + \mathbf{U}^{(\alpha)} \\ &\ \ \vdots \\ \mathbf{R}^{(\omega)} &\equiv \mathbf{B}^{(\omega)}\mathbf{F}^{(\omega)} + \mathbf{U}^{(\omega)}, \end{aligned} \tag{42}$$

and then we join these models by imposing structure on the cross-correlations. Skipping the unnecessary Steps 4-5 of projection and pricing in Figure 2, we generate the scenarios-probabilities pair $(\mathcal{R}; \mathbf{p})$ that corresponds to the distribution of the returns. Then as in (23) we obtain the scenarios-probabilities pair $(\mathcal{R}\mathbf{w}; \mathbf{p})$ that corresponds to the aggregate distribution of the portfolio return, Step 6 in Figure 2. This distribution is used to compute and analyze the risk in the portfolio.

Next, we perform Step 7 in Figure 2, namely the FoD top-down attribution (31), which we report here

$$\mathcal{R}\mathbf{w} = \mathcal{Z}|_{\mathcal{X}}\mathbf{d}_{\mathbf{w}} + \eta_{\mathbf{w}}. \tag{43}$$

In this expression $\mathcal{Z}|_{\mathcal{X}}$ is the conditional panel that represents the distribution of the attribution factors, which we set as either the global, or the regional factors. To this purpose, we notice that, as in (35), the global factors are a deterministic function of the regional factors

$$\mathbf{Z}^{(Gl)} \equiv \mathbf{A} \begin{pmatrix} \mathbf{F}^{(\alpha)} \\ \vdots \\ \mathbf{F}^{(\omega)} \end{pmatrix}, \tag{44}$$

where $\mathbf{A}$ is a matrix that aggregates the granular regional factors into global factors. Similarly, the regional factors are trivially a deterministic function (35) of themselves

$$\mathbf{Z}^{(Rg)} \equiv \begin{pmatrix} \mathbf{F}^{(\alpha)} \\ \vdots \\ \mathbf{F}^{(\omega)} \end{pmatrix}. \tag{45}$$

Therefore it is immediate to generate the conditional panel $\mathcal{Z}|_{\mathcal{X}}$ as in (35) and comments thereafter.

To determine the exposures $\mathbf{d}_{\mathbf{w}}$ in (43) we maximize the r-square of the attribution factors to the portfolio, i.e. we set in (26) as target

$$\mathcal{T}\left(f_{\eta_{\mathbf{w}}^{\mathbf{d}}}\right) \equiv -\mathrm{E}\{\left(\eta_{\mathbf{w}}^{\mathbf{d}}\right)^{2}\}, \tag{46}$$

obtaining a quadratic program such as (32). As in (32), we can, but we do not need to, impose cardinality constraints to only select the most representative factors for the specific portfolio $\mathbf{w}$.

To summarize, using the global attribution factors (44) on the right hand side in the FoD top-down attribution (43) we can express the portfolio return on the left hand side, and all the risk numbers related to it, in a way suitable to the global portfolio manager. On the other hand, using the regional attribution factors (45) to perform the FoD attribution, we can express the same portfolio return and risk numbers, in a way that suits the regional portfolio manager.

4.2 Style analysis

Consider a generic time t and a portfolio $\mathbf{w}^{t}$. Style analysis amounts to expressing the portfolio return $R_{\mathbf{w}^{t}}$ as a linear combination of the returns Z_k of a set of style indices

$$R_{\mathbf{w}^{t}} \approx \sum_{k=1}^{K} d_k^t Z_k, \tag{47}$$

where in the notation we emphasized the time dependence of the exposures d_k^t. Typically, the exposures d_k^t of the style returns are constrained to be positive and sum to one, in such a way that they can be interpreted as the weights of a replication portfolio.

In the pathbreaking, now standard, approach by Sharpe (1992), the coefficients d_k^t are computed by running a constrained time series regression of the realizations of the portfolio returns on the realizations of the style returns. This approach is suboptimal because the regression approach assumes no dynamical reallocation of the portfolio weights. Therefore, the result is a depiction of the style over the estimation period, rather than a true point-in-time analysis of the style of the current portfolio.

If the portfolio weights $\mathbf{w}^{t}$ are unknown and only the portfolio returns are observable, a partial improvement is suggested in Corielli and Meucci (2004). However, if the portfolio weights are known, FoD provides a substantial improvement. First, we proceed as in Steps 1-6 in Figure 2 to obtain the scenarios-probabilities pair $\left(\mathcal{R}^{t}\left(\mathcal{X}^{t}\right);\mathbf{p}^{t}\right)$ for the returns of the securities from the risk drivers $\mathcal{X}^t$, all of which are *forward-looking at time t*. Then we use the returns of the style indices as attribution factors in the FoD top-down attribution (31). To this purpose, we obtain as in Step 7 the conditional scenarios $\mathcal{Z}|^t_{\mathcal{X}^t}$ for the returns of the style indices, which are *forward-looking at time t*. Then we compute the exposures $\mathbf{d}_{\mathbf{w}^t}$ to maximize the r-square of the style factors, i.e. we set in (26) as target $\mathcal{T}(f_{\eta^{\mathbf{d}}_{\mathbf{w}}}) \equiv -\operatorname{E}\{(\eta^{\mathbf{d}}_{\mathbf{w}})^2\}$. Rearranging, we obtain as in (32) a quadratic program, with a different set of constraint

$$\mathbf{d}_{\mathbf{w}^t} \equiv \underset{\mathbf{d}'\mathbf{1}=1,\mathbf{d}\geq\mathbf{0}}{\operatorname{argmin}} \{\mathbf{d}'\mathbf{H}_t\mathbf{d} - 2\mathbf{a}_t'\mathbf{d}\}. \tag{48}$$

The exposures (48) reflect the portfolio composition $\mathbf{w}^t$ at time t, as well as the forward-looking risk in the market at that time, instead of the average portfolio composition and market structure during the rolling estimation window that precedes t.

4.3 Portfolio-based risk-attribution

Another application of FoD is the risk-based representation of a given portfolio $\mathbf{w}$ in terms of a select set of portfolios

$$\mathbf{w} \approx d_1\mathbf{w}_1 + \cdots + d_K\mathbf{w}_K. \tag{49}$$

For instance, the portfolios $\mathbf{w}_1, \ldots, \mathbf{w}_K$ can be defined as the weights implied by a set of "incremental alpha signals", see Grinold (2006).

To perform this attribution, first we generate the scenarios-probabilities pair $(\mathcal{R}(\mathcal{X}); \mathbf{p})$ that corresponds to the joint distribution of the securities returns from the risk drivers $\mathcal{X}$ as in Steps 1-6 in Figure 2. Then we define as attribution factors the returns of the portfolios (49) which we intend to use as a basis for the attribution

$$\mathcal{Z}|_{\mathcal{X}} \equiv (\mathcal{R}(\mathcal{X})\mathbf{w}_1, \ldots, \mathcal{R}(\mathcal{X})\mathbf{w}_K). \tag{50}$$

Now we perform the FoD top-down attribution (31). If we compute the exposures $\mathbf{d}_{\mathbf{w}^t}$ to minimize the variance of the residual i.e. we set in (26) as target $\mathcal{T}(f_{\eta^{\mathbf{d}}_{\mathbf{w}}}) \equiv -\operatorname{Var}\{(\eta^{\mathbf{d}}_{\mathbf{w}})^2\}$ and we do not impose constraints we obtain the analytical solution that appears in Grinold (2006)

$$\mathbf{d}_{\mathbf{w}} = \left(\mathbf{W}'\widehat{\boldsymbol{\Sigma}}\mathbf{W}\right)^{-1}\mathbf{W}'\widehat{\boldsymbol{\Sigma}}\mathbf{w}, \tag{51}$$

where $\mathbf{W} \equiv (\mathbf{w}_1|\cdots|\mathbf{w}_K)$ is the $N \times K$ matrix of the juxtaposition of the holdings and $\widehat{\boldsymbol{\Sigma}}$ is the covariance of the returns panel $\mathcal{R}$. More in general, we can perform portfolio based risk attribution with general targets and general constraints.

References

Corielli, F., and A. Meucci, 2004, Linear models for style analysis, *Statistical Methods and Applications* 13, 105–129.

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

Gatheral, J., 2008, Random matrix theory and covariance estimation, *Working Paper*.

Gollamudi, S., 2009, Conditional simulations with exact means and covariances, *Working Paper*.

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

Grinold, R., 2006, Attribution: Modeling asset characteristics as portfolios, *Journal of Portfolio Management* 32, 9–24.

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

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

Menchero, J., and V. Poduri, 2008, Custom factor attribution, *Financial Analysts Journal* 64, 81–92.

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 Extended version available at http://ssrn.com/abstract=872577.

———, 2007, Risk contributions from generic user-defined factors, *Risk* 20, 84–88 Extended version available at http://ssrn.com/abstract=930034.

———, 2009a, Exercises in advanced risk and portfolio management - with step-by-step solutions and fully documented code, *Free E-Book* available at http://ssrn.com/abstract=1447443.

———, 2009b, Review of discrete and continuous processes in finance: Theory and applications, *Working paper* Available at http://ssrn.com/abstract=1373102.

———, 2009c, Simulations with exact means and covariances, *Risk* 22, 89–91 Extended version available at http://ssrn.com/abstract=1415699.

———, 2010, Review of linear factor models: Unexpected common features and the systematic-plus-idiosyncratic myth, *Working paper* Available at http://ssrn.com/abstract=1635495.

Rockafellar, R.T., and S. Uryasev, 2000, Optimization of conditional value-at-risk, *Journal of Risk* 2, 21–41.

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

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

A Appendix

In this appendix we discuss technical results that can be skipped at first reading.

A.1 Monte Carlo scenarios from marginal-copula distributions

Consider an arbitrary set of K random variables, which in this section we denote by the K-dimensional vector $\mathbf{X}$.

We represent the marginal distribution of each factor X_k by means of its cumulative distribution function:

$$F_k(x) \equiv \operatorname{Prob}\{X_k \leq x\}, \quad k = 1, \ldots, K. \tag{52}$$

We use the information available in the time series of each factor to estimate all the cdf's (52).

Since the marginal distributions are determined in (52), the full joint distribution of the systematic factors $\mathbf{X}$ is completely determined by the choice of a dependence structure, also known as copula, see e.g. Meucci (2005). For instance, we can model the dependence among the factors by means of a normal copula. In this case, consider a normal vector with correlation matrix $\mathbf{\Gamma}$

$$\mathbf{Y} \sim \mathrm{N}(\mathbf{0}, \mathbf{\Gamma}). \tag{53}$$

This amounts to modeling the joint distribution of the factors as follows

$$\begin{pmatrix} X_1 \\ \vdots \\ X_K \end{pmatrix} \stackrel{d}{=} \begin{pmatrix} F_1^{-1}(\Phi(Y_1)) \\ \vdots \\ F_K^{-1}(\Phi(Y_K)) \end{pmatrix}, \tag{54}$$

where Φ denotes the cdf of the standard normal distribution. This joint structure is consistent with the marginal specification described above: indeed, it turns out that the cdf of the generic k-th factor implied by (54) is precisely (52).

To ensure flexibility and modularity, we represent the joint distribution of the factors $\mathbf{X}$ in terms of a $J \times K$ panel $\mathcal{X}$ of J joint Monte Carlo simulations: the generic j-th row represents a joint scenario for the factors $\mathbf{X}$ and the generic k-th column represents the marginal distribution of the k-th factor X_k. The quality of the simulations is roughly independent of the number K of risk factors. On the other hand, the quality improves with the number of simulations, but so does the computational cost: we choose the number J of simulations appropriately to achieve a balance between quality and computational cost.

To produce $\mathcal{X}$, first we generate a $J \times K$ panel $\mathcal{C}$ for the desired copula. For instance, for a normal copula, first we generate a $J \times K$ panel $\mathcal{Y}$ of joint Monte Carlo simulations from the normal distribution (53) by matching moments as in Meucci (2009c) or Gollamudi (2009), if we also need to match the joint moments with another normal copula, as in the generation of conditional scenarios. Then we apply the standard normal cdf Φ to each entry of the panel $\mathcal{Y}$, thereby obtaining the $J \times K$ panel $\mathcal{C} \equiv \Phi(\mathcal{Y})$. The columns of this panel have a uniform distribution and represent the copula.

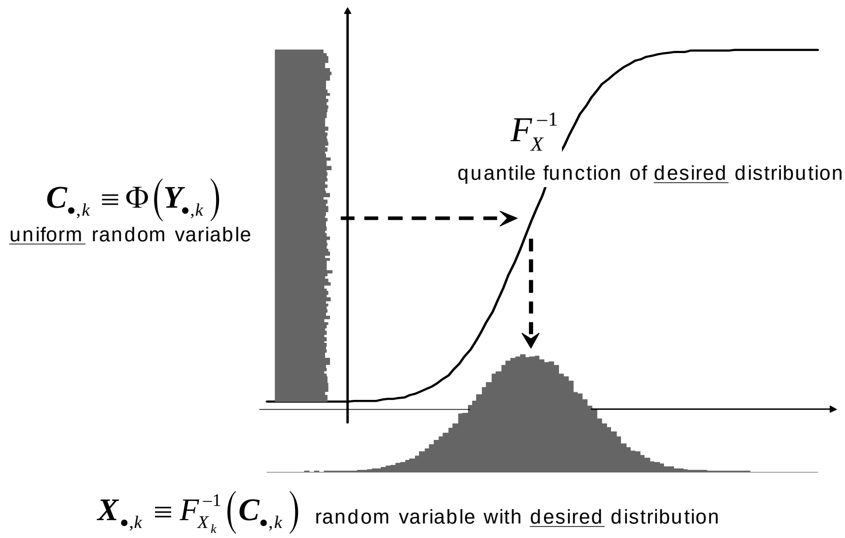

Figure 9: Uniform distribution is transformed into desired marginal distribution with desired joint structure

Once we have the copula panel $\mathcal{C}$, we apply the suitable quantile function F_k^{-1} to each column of the copula panel $\mathcal{C}$, by linear interpolation of the cdf grid as in Meucci (2006), see Figure 9. More precisely, consider the linear interpolator of a function $\mathbf{z}_{grid} \equiv f(\mathbf{y}_{grid})$ evaluated at a given grid of points $\mathbf{y}_{grid}$. We denote the interpolator as follows:

$$y \mapsto z = \operatorname{interp}(y; \mathbf{y}_{grid}, \mathbf{z}_{grid}). \tag{55}$$

We evaluate each generic k-th marginal (52) on a grid of $S \approx 10^3$ points $x_{s,k}$ that include extreme scenarios

$$F_{s,k} \equiv F_k(x_{s,k}), \quad s = 1, \ldots, S. \tag{56}$$

Then we can produce the $J \times K$ panel $\mathcal{X}$ entry-wise

$$\mathcal{X}_{j,k} \equiv \operatorname{interp}(\mathcal{C}_{j,k}; F_{\cdot,k}, x_{\cdot,k}), \tag{57}$$

where $\mathcal{C}_{j,k}$ is the (j,k)-th entry of the copula panel $\mathcal{C}$, $F_{\cdot,k}$ denotes the S-dimensional vector of the cdf of the k-th marginal (56) and $x_{\cdot,k}$ denotes the S-dimensional vector of the respective grid values. Indeed, it follows that $x_{\cdot,k}$ represents the quantile function F_k^{-1} evaluated at the grid of points $F_{\cdot,k}$ and $\mathcal{C}_{j,k}$ represents a simulation of the k-th grade. Notice that we do not need to invert the cdf explicitly, i.e. to compute the quantile function.

A.2 Risk decomposition: a primer

Here we present a quick primer on risk decomposition and risk computation techniques. For more details see Meucci (2005).

We start from the expression of the portfolio return as the product of a vector of S risk sources $\mathbf{F}$ times the respective exposures $\mathbf{b}$, which represent the practitioner's decision variables:

$$R = \sum_{s=1}^{S} d_s F_s. \tag{58}$$

This formulation includes $R = \mathbf{w}'\mathbf{R}$, where $S \equiv N$, the number of securities; $F_n \equiv R_n$ represents the returns of the securities; and $\mathbf{d} \equiv \mathbf{w}$ represents the respective portfolio weights as well as $R = \mathbf{d}'_{\mathbf{w}}\mathbf{Z} + \eta_{\mathbf{w}}$ as obtained in the top-down FoD attribution (24), where the last term in (58) would account for the residual.

The standard deviation is defined as follows

$$Sd_R \equiv \sqrt{\mathbb{E}\left\{(R - \mathbb{E}\{R\})^2\right\}}. \tag{59}$$

In benchmark-relative allocations the standard deviation is known as tracking error. Intuitively, the standard deviation is a measure of the oscillations of the return in normal market conditions.

The VaR is defined as a quantile of the loss:

$$VaR_c \equiv Q_{-R}(c), \tag{60}$$

where $Q_X(c)$ denotes the $c \times 100$-quantile of the distribution of X, where the confidence c is typically set very high, of the order of $c \approx 99\%$. Intuitively, in a set of, say, $100,000$ simulations, the 99%-confidence VaR is the best among the worst $1,000$ scenarios.

Since the VaR is insensitive to the distribution of the remaining 999 worst-case scenarios, one introduces the ES, defined as the expected loss, conditioned on the loss exceeding the VaR:

$$ES_c \equiv \mathbb{E}\{-R| - R \geq VaR_c\}. \tag{61}$$

Intuitively, in a set of, say, $100,000$ simulations, the 99%-confidence ES is the average among the worst $1,000$ scenarios.

Ideally, we would like to write the risk of the portfolio, as measured by volatility, VaR, or ES, as the product of the exposures times the factor-specific "isolated" volatility, VaR, or ES of the individual sources of risk, in a way fully symmetrical to (58). Unfortunately, such an identity does not hold. Let us consider for instance the standard deviation (59). It is well known that

$$Sd_R \neq \sum_{s=1}^{S} d_s Sd_s. \tag{62}$$

The theory behind risk contributions rests on the observation that volatility, VaR and ES are homogenous: by doubling the exposures $\widetilde{\mathbf{b}}$ in (58) we double the risk in the portfolio, see Meucci (2005) for a detailed review. Although the decomposition (62) is not feasible, the following identity holds true because the volatility is homogeneous, as proved by Euler

$$Sd_R \equiv \sum_{s=1}^{S} d_s \frac{\partial Sd}{\partial d_s}, \tag{63}$$

where

$$\frac{\partial Sd_R}{\partial \mathbf{d}} = \frac{\mathrm{Cov}\{\mathbf{F}\}\,\mathbf{d}}{\sqrt{\mathbf{d}'\,\mathrm{Cov}\{\mathbf{F}\}\,\mathbf{d}}}. \tag{64}$$

Notice that (63) is an exact identity, not a first-order approximation. In words, total risk can still be expressed as the sum of the contributions from each factor, where the generic s-th contribution is the product of the "per-unit" marginal contribution $\partial Sd/\partial d_s$ times the "amount" of the s-th factor in the portfolio, as represented by the exposure d_s. Unfortunately, the per-unit marginal contribution $\partial Sd/\partial d_s$ is not a truly "isolated", factor-specific quantity, as it depends on the whole portfolio. However, within the scope of the given portfolio, (63) does indeed provide an additive decomposition of risk.

We start with the $J \times S$ panel of joint scenarios $\mathcal{F}$ of the factors $\mathbf{F}$ in (58). Due to the remarks after (58), $\mathcal{F}$ is either the panel of joint returns scenarios $\mathcal{R}$ of the securities returns or, more generally

$$\mathcal{F} \equiv (\mathcal{Z}|_{\mathcal{X}}, \eta_{\mathbf{w}}), \tag{65}$$

where $\mathcal{Z}|_{\mathcal{X}}$ is the conditional panel of the attribution factors and $\eta_{\mathbf{w}}$ are the residual scenarios in (31). Then, the covariance of $\mathbf{F}$ that appears in the partial derivatives (64) is provided by the sample covariance of the panel $\mathcal{F}$.

Just like for the standard deviation (62), the VaR is not the weighted average of the isolated VaR's:

$$VaR_c \neq \sum_{s=1}^{S} d_s VaR_s. \tag{66}$$

However, the VaR is homogeneous, too, and therefore we can write it as the sum of the contributions from each factor:

$$VaR_c \equiv \sum_{s=1}^{S} d_s \frac{\partial VaR_c}{\partial d_s}. \tag{67}$$

Again, total risk can still be expressed as the sum of the contributions from each factor, where the generic s-th contribution is the product of the "per-unit" marginal contribution $\partial VaR_c/\partial d_s$ times the "amount" of the s-th factor in the portfolio, as represented by the exposure d_s.

In non-normal markets the volatility does not fully determine the VaR. However, the partial derivatives that appear in (67) can be expressed conveniently as in Hallerbach (2003), Gourieroux, Laurent, and Scaillet (2000), Tasche (2002):

$$\frac{\partial VaR_c}{\partial \mathbf{d}} \equiv -\mathbb{E}\left\{\mathbf{F}|R \equiv -VaR_c\right\}. \tag{68}$$

In turn, these expectations can be approximated numerically as in Mausser (2003), Epperlein and Smillie (2006):

$$\frac{\partial VaR_c}{\partial \mathbf{d}} \approx -\mathbf{k}_c' \mathcal{S}_{\mathbf{d}}. \tag{69}$$

In this expression $\mathcal{S}_{\mathbf{d}}$ is a $J \times S$ 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 $-\mathcal{F}\mathbf{d}$; and $\mathbf{k}_c$ is a Gaussian smoothing kernel peaked around the re-scaled confidence level cJ.

Finally, also the ES is homogeneous and thus we can write also the ES as the sum of the contributions from each factor:

$$ES_c \equiv \sum_{s=1}^{S} d_s \frac{\partial ES_c}{\partial d_s}. \tag{70}$$

In non-normal markets the volatility does not fully determine expected shortfall. However, the partial derivatives that appear in (70) can be expressed as

$$\frac{\partial ES_c}{\partial \mathbf{d}} \equiv -\mathbb{E}\left\{\mathbf{F}|R \leq -Q_{-\mathbf{d}'\mathbf{F}}\left(c\right)\right\}. \tag{71}$$

In turn, we can approximate numerically these expectations as

$$\frac{\partial ES_c}{\partial \mathbf{d}} \approx -\mathbf{q}_c' \mathcal{S}_{\mathbf{d}}, \tag{72}$$

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