"The $\mathbb{P}$rayer"
Ten-Step Checklist for Advanced Risk and Portfolio Management

Attilio Meucci[1]
attilio_meucci@symmys.com

published[2]: April 1 2011
this revision: May 12 2011
latest revision available at http://symmys.com/node/63

Abstract

We present "the $\mathbb{P}$rayer", a recipe of ten sequential steps for all portfolio managers, risk managers, algorithmic traders across all asset classes and all investment horizons, to model and manage the P&L distribution of their positions.

For each of the ten steps of the $\mathbb{P}$rayer, we introduce all the key concepts with precise notation; we illustrate the key concepts by means of a simple case study that can be handled with analytical formulas; we point the readers toward multiple advanced approaches to address the non-trivial practical problems of real-life risk modeling; and we highlight a non-exhaustive list of common pitfalls.

JEL Classification: C1, G11

Keywords: Quest for Invariance, conservation law of money, estimation, projection, pricing, aggregation, attribution, evaluation, optimization, invariants, risk drivers, random walk, Levy process, autocorrelation, long memory, volatility clustering, non-parametric, Monte Carlo, panic copula, elliptical maximum likelihood, robustness, influence function, breakdown point, Bayesian, Fourier transform, full repricing, theta, delta, gamma, vega, carry, duration, convexity, liquidity, holding, portfolio weight, return, leverage, attribution, Factors on Demand, satisfaction, VaR, CVaR, Sharpe ratio, diversification, effective number of bets, utility, homogeneous risk, quadratic programming, cone programming, grid search, integer programming, combinatorial heuristics, Black-Litterman, entropy pooling, drawdown, dynamic programming, algorithmic trading, optimal execution, performance analysis, slippage, market impact, geometric linkage.

[1] The author is grateful to Garli Beibi, Arlen Khodadadi, Luca Spampinato, and an anonymous referee.

[2] This article appears as Meucci, A. (2011), *The Prayer: Ten-Step Checklist for Advanced Risk and Portfolio Management* - The Quant Classroom by Attilio Meucci - GARP Risk Professional, April/June 2011, p. 54-60/34-41

Contents

Introduction

The quantitative investment arena is populated by different players: portfolio managers, risk managers, algorithmic traders, etc. These players are further differentiated by the asset classes they cover, the different time horizons of their activities and a variety of other distinguishing features. Despite the many differences, all the above "quants" are united by the common goal of correctly modeling and managing the probability distribution of the prospective P&L of their positions.

Here we present "the $\mathbb{P}$rayer", a blueprint of ten sequential steps for quants across the board to achieve their common goal, see Figure 1. By following the $\mathbb{P}$rayer, quants can avoid common pitfalls and ensure that they are not missing important points in their models. Furthermore, quants are directed to areas of advanced research that extends beyond the traditional quant literature. We use the letter "$\mathbb{P}$" to signify the true probability space of the buy-side P&L, which stands in contrast to the risk-neutral probability space "$\mathbb{Q}$" used on the sell-side to price derivatives, see Meucci (2011b).

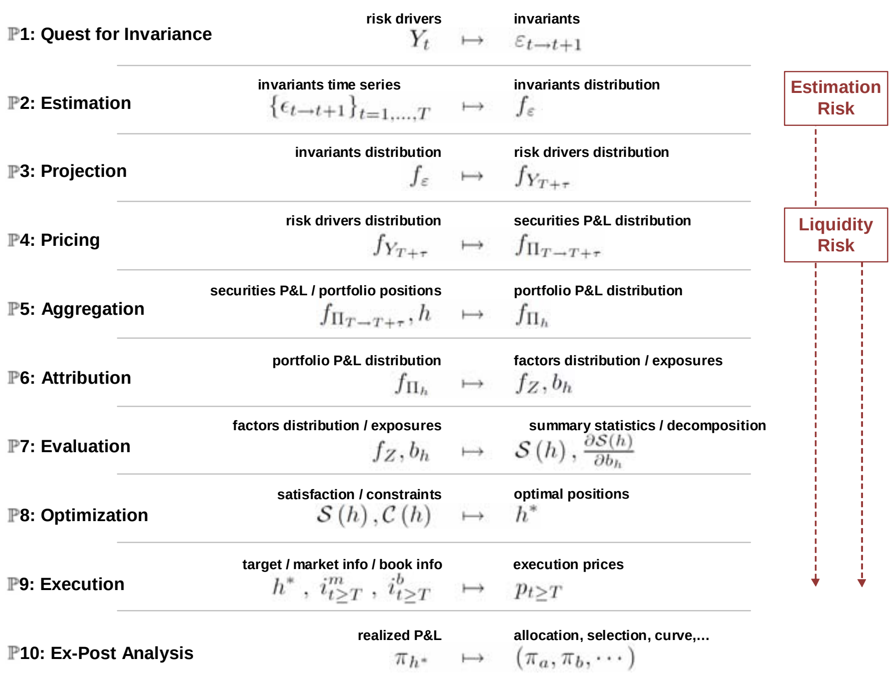

Figure 1: The "$\mathbb{P}$rayer": ten-step blueprint for risk and portfolio management

Below we discuss the ten steps of the $\mathbb{P}$rayer. Each step is concisely encapsulated into a definition with the required rigorous notation. Then a simple case study with a portfolio of only stocks and call options illustrates the steps with analytical solutions. Within each step, we prepare the ground for, and point to, advanced research that fine-tunes the models, or enhances the models' flexibility, or captures more realistic and nuanced empirical features. Each of these steps are deceptively simple at first glance. Hence, we highlight a few common pitfalls to further clarify the conceptual framework.

$\mathbb{P}$ 1 Quest for Invariance

The "quest for invariance" is the first step of the $\mathbb{P}$rayer, and the foundation of risk modeling. The quest for invariance is necessary for the practitioners to learn about the future by observing the past in a stochastic environment.

> **Key concept**. The Quest for Invariance step is the process of extracting from the available market data the "invariants", i.e. those patterns that repeat themselves identically and independently (i.i.d.) across time. The quest for invariance consists of two sub-steps: identification of the risk drivers and extraction of the invariants from the risk drivers.

The first step of the quest for invariance is to identify for each security the risk drivers among the market variables.

> **Key concept**. The risk drivers of a given security are a set of random variables
>
> $$Y_t \equiv (Y_{t,1}, \ldots, Y_{t,D})' \tag{1}$$
>
> that satisfy the following two properties: a) the risk drivers Y_t, together with the security terms and conditions, completely specify the security price at any given time t; b) the risk drivers Y_t, although not i.i.d., follow a stochastic process that is homogeneous across time, in that it is impossible to ascertain the sequential order of the realizations of the risk drivers from the study of the risk drivers past time series $\{y_t\}_{t=1,\ldots,T}$.

The risk drivers are variables that fully determine the price of a security, but in general they are not the price, because the price can be non-homogeneous across time: think for instance of a zero-coupon bond, whose price converges to the face value as the maturity approaches.

Homogeneity ensures that we can apply statistical techniques to the observed time series of the risk drivers $\{y_t\}_{t=1,\ldots,T}$ and project future distributions. Note that we use the standard convention where lower-case letters such as y_t denote realized variables, whereas upper-case letters such as Y_t denote random variables.

Illustration. Consider first the asset class of stocks. Denote by S_t the random price of one stock at the generic time t. The log-price of the stock $Y_t \equiv \ln S_t$, possibly adjusted by reinvesting the dividends, is not i.i.d. across time. However, the dynamics of the stock log-price is homogeneous across time: it is not possible to isolate any special period in the stock's future evolution that will distinguish its price pattern from a nearby period. Hence, to project into the future, the random variable $Y_t \equiv \ln S_t$ is a suitable candidate risk driver for the stock price S_t.

Next, consider a second asset class, namely stock options. Denote by $C_{t,k,e}$ the random price of a European call option on the stock, where k is a given strike and e is the given expiry date, or time *of* expiry. The call price, or its log-price, is not a risk driver, because the presence of the expiry date breaks the time homogeneity in the statistical behavior of the call option price.

In order to identify the risk drivers behind the call option, we transform the price into an equivalent, but homogeneous, variable, namely the implied volatility at a given time to expiry. More precisely, consider the Black-Scholes pricing formula

$$C_{t,k,e} \equiv c_{BS}\left(\ln S_t - \ln k, \Sigma_t, \upsilon_t\right), \tag{2}$$

where $\upsilon_t \equiv e - t$ is the time *to* expiry, Σ_t is the yet to be defined implied volatility for that time to expiry, and c_{BS} is the Black-Scholes formula

$$c_{BS}\left(m, \sigma, \upsilon\right) \equiv \frac{e^m}{k}\Phi(\frac{m + r\upsilon + \sigma^2\upsilon/2}{\sigma\sqrt{\upsilon}}) - e^{-r\upsilon}\Phi(\frac{m + r\upsilon - \sigma^2\upsilon/2}{\sigma\sqrt{\upsilon}}), \tag{3}$$

with Φ the standard normal cdf. At each time t, the price $C_{t,k,e}$ in (2) is observable, and so are S_t and υ_t. Therefore, the option formula (2) implies a value for Σ_t, which for this reason is called implied volatility.

The implied volatility for a given time *to* expiry, or better, the logarithm of the implied volatility $\ln \Sigma_t$, displays a homogeneous behavior through time and thus it is a good candidate risk driver for the option. From the option formula (2) we observe that the implied volatility alone is not sufficient to determine the call price in the future, as, in addition, the log-price $\ln S_t$ and the time to expiry υ_t are needed. Since the time to expiry is deterministic, the call option requires two risk drivers to fully determine its price

$$\begin{pmatrix} Y_{s,t} \\ Y_{\sigma,t} \end{pmatrix} \equiv \begin{pmatrix} \ln S_t \\ \ln \Sigma_t \end{pmatrix}. \tag{4}$$

The second step of the quest for invariance is the extraction of the invariants, i.e. the repeated patterns, from the homogeneous series of the risk drivers.

Key concept. The invariants are shocks that steer the stochastic process of the risk drivers Y_t over a given time step $t \to t+1$.

$$\varepsilon_{t\to t+1} \equiv (\varepsilon_{1,t\to t+1}, \ldots, \varepsilon_{Q,t\to t+1})' \tag{5}$$

The invariants satisfy the following two properties: a) they are identically and independently distributed (i.i.d.) across different time steps; b) they become known at the end of the step, i.e. at time $t+1$.

Note that each of the D risk drivers (1) can be steered by one or more invariants, therefore $Q \geq D$.

To determine whether a variable is i.i.d. across time, the easiest test is to scatter-plot the series of the variable versus its own lags. If the scatter-plot, or better, its location-dispersion ellipsoid, is a circle, then the variable is a good candidate for an invariant. For more on this and related tests see Meucci (2005a).

Being able to identify the invariants that steer the dynamics of the risk drivers is of crucial importance because it allows us to project the market randomness to the desired investment horizon. Often, practitioners make the mistake of projecting variables they have on hand, most notably returns, instead of the invariants. This, of course, leads to incorrect measurement of risk at the horizon, and thus to suboptimal trading decisions.

The stochastic process for the risk drivers Y_t is steered by the randomness of the invariants $\varepsilon_{t\to t+1}$. The most basic dynamics, yet the most statistically robust, which connects the invariants $\varepsilon_{t\to t+1}$ with the risk drivers Y_t is the random walk

$$Y_{t+1} = Y_t + \varepsilon_{t\to t+1}. \tag{6}$$

More advanced processes for the risk drivers account for such features as autocorrelations, stochastic volatility, and long memory. We refer to Meucci (2009b) for a review of these more general processes and how they related to random walk and invariants both in discrete and in continuous time, with theory, case studies, and code. We refer to Meucci (2009c) for the multivariate case, and how it relates to cointegration and statistical arbitrage.

Illustration. Consider our first asset class example, the stock. As discussed, the only risk driver is the log-price $Y_t \equiv \ln S_t$. The above scatter-plot generally indicates that the compounded return $\ln(S_{t+1}/S_t)$ are approximately invariants

$$\varepsilon_{t\to t+1} \equiv \ln S_{t+1} - \ln S_t. \tag{7}$$

Therefore the risk driver $Y_t \equiv \ln S_t$ follows a random walk, as in (6).

Now, consider our second asset class, the call option example. The empirical scatter-plot shows that the changes of the log-implied volatility are approximately i.i.d. across time. Furthermore, our analysis of the stock example (7) implies that the changes of the log-price are invariants. Therefore, using notation similar to (4), we obtain

$$\begin{pmatrix} \varepsilon_{s,t\to t+1} \\ \varepsilon_{\sigma,t\to t+1} \end{pmatrix} \equiv \begin{pmatrix} \ln S_{t+1} \\ \ln \Sigma_{t+1} \end{pmatrix} - \begin{pmatrix} \ln S_t \\ \ln \Sigma_t \end{pmatrix}. \tag{8}$$

This is also a random walk as in (6). Notice that this is a multivariate random walk.

The outcome of the quest for invariance, i.e. the set of risk drivers and their corresponding invariants, depends on the asset class and on the time scale of our analysis. For instance, for interest rates a simple random walk assumption (6) can be viable for time steps of one day, but for time steps of the order of one year mean-reversion becomes important. Similarly, for stocks at high frequency steps of the order of fractions of a second, the very time step becomes a random variable, which calls for its own invariant. We refer to Meucci (2009b) for a review.

Pitfalls. "*...The random walk is a stationary process...*". A random walk, such as Y_t in (6) is not stationary. The steps of the random walk $\varepsilon_{t\to t+1}$ are stationary, and actually they are the most stationary of processes, namely invariants.

"*...The random walk is too crude an assumption...*". Once the data is suitably transformed into risk drivers, the random walk assumption is very hard to beat in practice, see Meucci (2009b).

"*...Returns are invariants ...*". Returns are not invariants in general. In our call option example, the past returns of the call option price do not teach us anything about the future returns of the option.

$\mathbb{P}$ 2 Estimation

As highlighted in the Quest for Invariance Step $\mathbb{P}$ 1, the stochastic behavior of the risk drivers is steered by the "invariants". Once the invariants have been identified, their distribution can be estimated from empirical analysis and from other sources of information.

Because of the invariance property, the distribution of the invariants does not depend on the specific time t. We represent this distribution in terms of its probability density function (pdf) f_ε. Note that, although the invariants are distributed independently across time, multiple invariants can be correlated with each other over the same time step. Therefore f_ε needs to be modeled as a multivariate distribution.

Key concept. The Estimation Step is the process of fitting a distribution f_ε to both the observed past realizations $\{\epsilon_{t\to t+1}\}$ of the invariants ε and optionally additional information i_T that is available at the current time T

$$\{\epsilon_{t\to t+1}\}_{t=1,\ldots,T} \;,\; i_T \quad \mapsto \quad f_\varepsilon. \tag{9}$$

Simple estimation approaches only process the time series of the invariants $\{\epsilon_{t\to t+1}\}$, but various advanced techniques also process information i_T such as market-implied forward looking quantities, subjective Bayesian priors, etc.

The simplest of all estimators for the invariants distribution is the non-parametric empirical distribution, justified by the law of large numbers, i.e. "i.i.d. history repeats itself". The empirical distribution assigns an equal probability $1/T$ to each of the past observations in the series $\{\epsilon_t\}_{t=1,\ldots,T}$ of the historical scenarios.

Alternatively, for the distribution of the invariants, one can make parametric assumptions such as multivariate normal, elliptical, etc.

Illustration. To illustrate the parametric approach, we consider our example (8), where the invariants ε are changes in moneyness and changes in log-implied volatility from t to $t+1$. We can assume that the distribution f_ε is bivariate normal with 2×1 expectation vector $\mu \equiv (\mu_s, \mu_\sigma)'$ and 2×2 covariance matrix σ^2 as below

$$\begin{pmatrix} \varepsilon_{s,t\to t+1} \\ \varepsilon_{\sigma,t\to t+1} \end{pmatrix} \sim \mathrm{N}(\begin{pmatrix} \mu_s \\ \mu_\sigma \end{pmatrix}, \begin{pmatrix} \sigma_s^2 & \rho\sigma_s\sigma_\sigma \\ \rho\sigma_s\sigma_\sigma & \sigma_\sigma^2 \end{pmatrix}). \tag{10}$$

The expectation can be estimated with the sample mean $\mu \equiv \frac{1}{T}\sum_t \epsilon_t$, and the covariance with the sample covariance $\sigma^2 \equiv \frac{1}{T}\sum_t (\epsilon_t - \mu)(\epsilon_t - \mu)'$, where $'$ denotes the transpose.

In large multivariate markets it is important to impose structure on the correlations of the distribution of the invariants f_ε. This is often achieved in practice by means of linear factor models. Linear factor models are an essential tool of risk and portfolio management, as they play a key role in the Estimation Step $\mathbb{P}$ 2, as well as, among others, in the Attribution Step $\mathbb{P}$ 6 and the Optimization Step $\mathbb{P}$ 8. We refer to Meucci (2010h) for a thorough review of theory, code, empirical results, and pitfalls of linear factor models in these three and other contexts.

A highly flexible approach to estimate and model distributions is provided by the copula-marginal decomposition, see e.g. Cherubini, Luciano, and Vecchiato (2004). In order to use this decomposition in practice, as well as any alternative outcome of the estimation process, the only feasible option is to represent distributions in terms of historical scenarios similar to the above, or Monte Carlo generated scenarios, see Meucci (2011a).

An important advanced topic is **estimation risk**. It is important to emphasize that, regardless how advanced an estimation technique is applied to model the joint distribution of the invariants, the final outcome will be an estimate, i.e.

only an approximation, of the true, unknown, distribution of the invariants f_{ε}. Estimation risk is the risk stemming from using an estimate of the invariants distribution in the process of managing the portfolio's positions, instead of the true, unknown distribution of the invariants.

Estimation risk, which first appears here in the context of the Estimation Step $\mathbb{P}$ 2, affects the cornerstones of risk and portfolio management, namely the Evaluation Step $\mathbb{P}$ 7, the Optimization Step $\mathbb{P}$ 8, and the Execution Step $\mathbb{P}$ 9. We will explore in those steps techniques that attempt to address estimation risk, which include scenario analysis, Fully Flexible Probabilities, robust estimation and optimization, multivariate Bayesian statistics, etc.

Pitfall. "*...In order to estimate the return of a bond I can analyze the time series of the past bond returns...*". The price of bonds with short maturity will soon converge to its face value. As a result, the returns are not invariants, and thus their past history is not representative of their future behavior. Estimation must always be linked to the quest for invariance.

"*...In markets with a large number Q of invariants I use a cross-sectional linear factor model on returns with $K \ll Q$* factors*. This reduces the covariance parameters to be estimated from* $\approx Q^2/2$ to $\approx K^2/2 + Q$.". A cross-sectional factor model has the same number of unknown quantities as a time-series model. The cross-sectional factors are typically autocorrelated. The residuals in both cross-sectional and time-series models are not truly idiosyncratic, as they display non-zero correlation with each other. For more on these and related pitfalls for linear factor models, see Meucci (2010h).

$\mathbb{P}$ 3 Projection

Ultimately we are interested in the value of our positions at the investment horizon. In order to determine the distribution of our positions, we must first determine the distribution of the risk drivers at the investment horizon. This distribution, in turn, is obtained by projecting to the horizon the invariants distribution, obtained in the Estimation Step $\mathbb{P}$ 2.

We denote the current time as $t \equiv T$ and the generic investment horizon $t \equiv T + \tau$, where τ is the distance to the horizon, say, one week.

Key concept. The Projection Step is the process of obtaining the distribution at the investment horizon $T + \tau$ of the relevant risk drivers Y_t from the distribution of the invariants and additional information i_T available at the current time T

$$f_{\varepsilon}\ ,\ i_T \quad \mapsto \quad f_{Y_{T+\tau}}. \tag{11}$$

In order to project the risk drivers we must go back to their connection with the invariants analyzed in the Quest for Invariance Step $\mathbb{P}$ 1.

If the drivers evolve as a random walk (6), then by recursion of the random walk definition $Y_{t+2} = Y_{t+1} + \varepsilon_{t+1\to t+2} = Y_t + \varepsilon_{t\to t+1} + \varepsilon_{t+1\to t+2}$ we obtain that the risk drivers at the horizon $Y_{T+\tau}$ are the current observable value y_T plus the sum of all the intermediate invariants

$$Y_{T+\tau} = y_T + \varepsilon_{T\to T+1} + \cdots + \varepsilon_{T+\tau-1\to T+\tau}. \tag{12}$$

Using the independence of the invariants, (12) yields for the variance

$$\mathrm{V}\{Y_{T+\tau}\} = \mathrm{V}\{\varepsilon_{T\to T+1}\} + \cdots + \mathrm{V}\{\varepsilon_{T+\tau-1\to T+\tau}\}. \tag{13}$$

Since all the ε's in (12) are i.i.d., all the variances in (13) are equal, and thus we obtain the well-known "square-root rule" for the projection of the standard deviation $\mathrm{Sd}\{Y_{T+\tau}\} = \sqrt{\tau}\,\mathrm{Sd}\{\varepsilon\}$. Note that we did not make any distributional assumption such as normality to derive the square-root rule.

Simple results to project other moments under the random walk assumption (6), such as skewness and kurtosis, can be found in Meucci (2010a) and Meucci (2010i). Projecting the whole distribution is more challenging, but can still be accomplished using Fourier transform techniques, see Albanese, Jackson, and Wiberg (2004).

In the more general case where the drivers do not evolve as a random walk (6), the projection can be obtained by redrawing scenarios according to the given dynamics, see e.g. Meucci (2010b) for the parametric case and Paparoditis and Politis (2009) for the empirical distribution.

Illustration. In our oversimplified normal example the projection can be performed analytically. Indeed, from the normal distribution of the invariants (10) it follows, from the preservation of normality with the sum of independent normal variables, that the sum of the invariants is normal $\varepsilon_{T\to t+\tau} \sim \mathrm{N}(\tau\mu, \tau\sigma^2)$. Thus we obtain for the distribution of the two risk drivers at the horizon

$$\begin{pmatrix} \ln S_{T+\tau} \\ \ln \Sigma_{T+\tau} \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} \ln s_T \\ \ln \sigma_T \end{pmatrix} + \tau \begin{pmatrix} \mu_s \\ \mu_\sigma \end{pmatrix}, \tau \begin{pmatrix} \sigma_s^2 & \rho\sigma_s\sigma_\sigma \\ \rho\sigma_s\sigma_\sigma & \sigma_\sigma^2 \end{pmatrix}\right). \tag{14}$$

Pitfall. "*...To project the market I assume normality and therefore I multiply the standard deviation by the square root of the horizon...*". The square root rule is true for all random walks with finite-variance invariants, regardless of their distribution. However, the square-root rule only applies to the standard deviation and does not allow to determine all the other moments of the distribution, unless the distribution is normal.

ℙ 4 Pricing

Now that we have the distribution of the risk drivers at the horizon $Y_{T+\tau}$ from the Projection Step ℙ 3, we are ready to compute the distribution of the security prices in our book. Recall that the value of the securities at the investment horizon, by design, is fully determined by a) risk drivers at the horizon $Y_{T+\tau}$ and b) non-random information i_T known at the current time, such as terms and conditions

$$P_{T+\tau} = p\left(Y_{T+\tau}; i_T\right). \tag{15}$$

Then, given the security price at the horizon $P_{T+\tau}$, the security P&L from the current date to horizon $\Pi_{T\to T+\tau}$ is the difference between the horizon value (15), which is a random variable, and the current value, which is observable and thus part of the available information set i_T. Thus the horizon profit function reads

$$\Pi_{T\to T+\tau} = p\left(Y_{T+\tau}; i_T\right) - p_T. \tag{16}$$

Note that the P&L must be adjusted for coupons and dividends, either by reinvesting them in the pricing function (15), or by an additional cash flow term in (16).

Key concept. The Pricing Step is the process of obtaining the distribution of the securities P&L's over the investment horizon from the distribution of the risk drivers at the horizon and current information such as terms and conditions, by means of the pricing function

$$f_{Y_{T+\tau}}\ ,\ i_T \quad \mapsto \quad f_{\Pi_{T\to T+\tau}} \tag{17}$$

At times it is convenient to approximate the pricing function (15) by its Taylor expansion

$$p\left(y; i_T\right) = p\left(\overline{y}; i_T\right) + \left(y - \overline{y}\right)' \partial_y p\left(\overline{y}; i_T\right) + \left(y - \overline{y}\right)' \frac{\partial_{yy} p\left(\overline{y}; i_T\right)}{2}\left(y - \overline{y}\right) + \cdots \tag{18}$$

where $\overline{y}$ is a significative value of the risk drivers, often the current value $\overline{y} \equiv y_T$; $\partial_y p\left(\overline{y}; i_T\right)$ denotes the vector of the first derivatives; and $\partial_{yy} p\left(\overline{y}; i_T\right)$ denotes the matrix of the second cross-derivatives.

Depending on its end users, the coefficients in the Taylor approximation (18) are known under different names. In the derivatives world, they are called the "Greeks": theta, delta, gamma, vega, etc. In the fixed-income world the coefficients are called carry, duration and convexity.

Illustration. In our stock example, the single risk driver is the log-price $Y_t \equiv \ln S_t$. Therefore the horizon pricing function (15) reads $p\left(y\right) = e^y$. Its Taylor approximation reads $p\left(y\right) \approx e^{y_T}\left(1 + y - y_T\right)$. Then the P&L of the stock (16) reads

$$\Pi_{s,T\to T+\tau} \approx s_T\left(\ln S_{T+\tau} - \ln s_T\right). \tag{19}$$

Hence, from the distribution of the risk drivers (14), it follows that the distribution of the stock P&L is approximately normal

$$\Pi_{s,T\to T+\tau} \sim \mathrm{N}(\tau s_T \mu_s, \tau s_T^2 \sigma_s^2). \tag{20}$$

For our call option with strike k and expiry e, the risk drivers are the log-price $Y_{s,t} \equiv \ln S_t$ and the log-implied volatility $Y_{\sigma,t} \equiv \ln \Sigma_t$, as in (4). The horizon pricing function (15) follows from the Black-Scholes formula (2) and reads

$$p_{BS}\left(y_s, y_\sigma; i_T\right) = c_{BS}\left(y_s - \ln k, e^{y_\sigma}, e - T - \tau\right). \tag{21}$$

When the investment horizon is much shorter than the time to expiry of the option, i.e. $\tau \ll e - T$, the following first-order Taylor approximation suffices to proxy the price $p_{BS}\left(y_s, y_\sigma; i_T\right) \approx p_{BS}\left(y_{s,T}, y_{\sigma,T}; i_T\right) + \delta_{BS,T} \cdot \left(y_s - y_{s,T}\right) + v_{BS,T} \cdot \left(y_\sigma - y_{\sigma,T}\right)$, where $\delta_{BS,T} \equiv \partial p_{BS}\left(y_{s,T}, y_{\sigma,T}\right) / \partial y_s$ is the option's current Black-Scholes "delta" and $v_{BS,T} \equiv \partial p_{BS}\left(y_{s,T}, y_{\sigma,T}\right) / \partial y_\sigma$ is the option's current Black-Scholes "vega". Then the P&L of the call option (16) reads

$$\Pi_{c,T \to T+\tau} \approx \left(\ln S_{T+\tau} - \ln s_T\right) \delta_{BS,T} + \left(\ln \Sigma_{T+\tau} - \ln \sigma_T\right) v_{BS,T}. \tag{22}$$

We stated in the distribution of the risk drivers (14) that the log-changes in (22) are jointly normal. Therefore, the distribution of the P&L is normal, because the linear combination of jointly normal variables is normal

$$\Pi_{c,T \to T+\tau} \sim \mathrm{N}(\tau \mu_c, \tau \sigma_c^2), \tag{23}$$

where

$$\mu_c \equiv \delta_{BS,T} \mu_s + v_{BS,T} \mu_\sigma \tag{24}$$
$$\sigma_c^2 \equiv \delta_{BS,T}^2 \sigma_s^2 + v_{BS,T}^2 \sigma_\sigma^2 + 2 \delta_{BS,T} v_{BS,T} \rho \sigma_s \sigma_\sigma. \tag{25}$$

Notice how the expectation of the call option's P&L depends on the expectations of the stock compounded return and the expectation of the log-changes in implied volatility, multiplied by the horizon τ. A similar relationship holds for the standard deviation of the call's P&L.

It is worth noticing that pricing becomes a surprisingly easy task when the distribution of the risk drivers is represented in terms of scenarios, as (16) is simply repeated scenario-by-scenario by inputting discrete realized risk drivers values.

We conclude the Pricing Step by highlighting two problems. First, a data and analytics problem: in many companies there might not be readily available pricing functions with all terms and conditions.

Second, the problem of **liquidity risk**. The pricing step assumes the existence of one single price, which is fully determined by the risk drivers $P_t = p\left(Y_t; i_T\right)$ as in (15). In reality, for any security there exist multiple possible prices, which represent supply and demand. The actual execution price depends on supply and demand, on the size and style of the transaction, and on other factors. As we will see, liquidity risk, which first comes to the surface here in the Pricing Step, affects with increasing intensity the Evaluation Step $\mathbb{P}$

7, the Optimization Step ℙ 8, and the Execution Step ℙ 9. We will discuss in those steps methodologies to address liquidity risk.

Pitfall. *"...The delta approximation gives rise to parametric risk models that assume normality..."*. The Taylor approximation of the pricing function can be paired with any distributional assumption, not necessarily normal, on the risk drivers.

"...The goodness of the Taylor approximation depends on the specific security...". The goodness of the Taylor approximation depends on the security and on the investment horizon: due to the square-root propagation of the standard deviation (13), the longer the horizon, the wider the distribution of the risk drivers. Therefore the approximation worsens with longer horizons.

ℙ 5 Aggregation

The Pricing Step ℙ 4 yields the projected P&L's of the single securities. The Aggregation Step generates the P&L distribution for a portfolio with multiple securities.

Consider a market of N securities, whose P&L's $\Pi \equiv (\Pi_1, \ldots, \Pi_N)'$ are obtained from the horizon pricing function (16). Notice that for simplicity we drop the subscript $T \to T + \tau$, because it is understood that from now on the ℙrayer focuses on the projected P&L between now and the future investment horizon.

Consider a portfolio, which is defined by the holdings in each position at the beginning of the period $h \equiv (h_1, \ldots, h_N)'$. The holdings are the number of shares for stocks, the number of standardized-notional contracts for swaps, the number of standardized-face-value-bond for bonds, etc.

The portfolio P&L is determined by the "conservation law of money": the total portfolio P&L is the sum of the holding in each security times the P&L generated by each security

$$\Pi_h = \sum_{n=1}^{N} h_n \Pi_n, \tag{26}$$

where we have assumed no rebalancing during the period.

Key concept. The Aggregation Step is the process of computing the distribution of the portfolio P&L Π_h by aggregating the joint distribution of the securities P&L with the given holdings

$$f_\Pi \, , \, h \quad \mapsto \quad f_{\Pi_h} \tag{27}$$

Given one single scenario for the risk drivers $Y_{T+\tau}$ and thus for the securities P&L's in (16), the computation of the portfolio P&L Π_h is immediately determined by the conservation law of money (26) as the sum of the single-security P&L's in that scenario.

However, to arrive at the whole continuous distribution of the portfolio P&L f_{Π_h} we must compute multiple integrals

$$f_{\Pi_h}(x)\,dx = \int_{h'\pi \in dx} f_{\Pi}(\pi_1, \ldots, \pi_N)\,d\pi_1 \cdots d\pi_N, \tag{28}$$

which is in general a very challenging operation. On the other hand, the computation of the aggregate P&L distribution becomes trivial when the market is represented in terms of scenarios, as the conservation law of money (26) is simply repeated in a discrete way scenario-by-scenario.

Illustration. In our example with a stock and a call option, whose P&L's are normally distributed, suppose we hold a positive or negative number h_s of shares of the stock and a positive or negative number h_c of the call. Then the total P&L follows from applying the aggregation rule (26) to the stock P&L (19) and the option P&L (22) and thus reads

$$\begin{aligned} \Pi_h &\approx h_s s_T \ln \tfrac{S_{T+\tau}}{s_T} + h_c(\delta_{BS,T} \ln \tfrac{S_{T+\tau}}{s_T} + v_{BS,T} \ln \tfrac{\Sigma_{T+\tau}}{\sigma_T}) \\ &= (h_s s_T + h_c \delta_{BS,T}) \ln \tfrac{S_{T+\tau}}{s_T} + h_c v_{BS,T} \ln \tfrac{\Sigma_{T+\tau}}{\sigma_T}. \end{aligned} \tag{29}$$

Thus, from the joint normal assumption (14) and the fact that sums of jointly normal variables are normal, the total portfolio is normally distributed. Isolating the horizon τ we obtain

$$\Pi_h \sim \mathrm{N}(\tau\mu_h, \tau\sigma_h^2), \tag{30}$$

where

$$\mu_h \equiv (h_s s_T + h_c \delta_{BS,T})\,\mu_s + h_c v_{BS,T} \mu_\sigma \tag{31}$$

$$\begin{aligned} \sigma_h^2 \equiv\; & (h_s s_T + h_c \delta_{BS,T})^2 \sigma_s^2 + h_c^2 v_{BS,T}^2 \sigma_\sigma^2 \\ & + 2\,(h_s s_T + h_c \delta_{BS,T})\, h_c v_{BS,T} \rho \sigma_s \sigma_\sigma \end{aligned} \tag{32}$$

Notice how both expectation and variance follow from the projection to the horizon of the invariants distribution (10).

Above we described in full the aggregation step. However, this topic is not complete without comparing the aggregation of the P&L with an equivalent, more popular, yet more error-prone, formulation in terms of returns.

The reader is probably familiar with the notion of returns, which allow for performance comparisons across different securities, and portfolio weights. The return is the ratio of the P&L over the current price $R_{T \to T+\tau} \equiv \Pi_{T \to T+\tau}/p_T$. The weight of a security is its relative market value within the portfolio $w_n \equiv h_n p_{n,T} / \sum_m h_m p_{m,T}$ and satisfies the "pie-chart" rule $\sum_n w_n = 1$.

The conservation law of money (26) becomes easier to interpret in terms of returns and weights, as the total portfolio return is the weighted average of the single-security returns

$$R_h = \textstyle\sum_{n=1}^N w_n R_n, \tag{33}$$

where we dropped the horizon subscript for simplicity.

In the $\mathbb{P}$rayer, we refrain from conceptualizing the aggregation and the subsequent steps in terms of returns, and we rely on returns only for interpretation purposes, for the following reasons.

First, P&L and holdings are always unequivocal, whereas returns and weights are subjective. Indeed, for leveraged securities, such as swaps and futures, the definition of returns and weights is not straightforward. In these circumstances we need to introduce a subjective "basis" denominator d known at the beginning of the return period, such that the return $R \equiv \Pi/d$ is always defined, and so is the weight, see Meucci (2010f).

Second, returns are often confused with the invariants, and thus incorrectly used for estimation.

Third, the linear returns $\left(p_{T+\tau} - p_T\right)/p_T$ which appear in the aggregation rule (33) are often confused with the compounded returns $\ln\left(p_{T+\tau}/p_T\right)$, which do *not* satisfy the aggregation rule.

Pitfall. *"...Returns are invariants. Therefore we can estimate their distribution from their past realizations and aggregate this distribution to the portfolio level using the weights..."*. Only in asset classes such as stocks do the concepts of invariant and return dangerously overlap. Furthermore, even for stocks, the projection does not apply directly to the returns, and thus one has to follow all the steps of the $\mathbb{P}$rayer, see Meucci (2010e).

$\mathbb{P}$ 6 Attribution

With the Aggregation Step $\mathbb{P}$ 5, we have arrived at the projected portfolio P&L distribution. In order to assess, manage, and hedge a portfolio with $h \equiv (h_1, \ldots, h_N)'$ holdings, it is important to ascertain the sources of risk that affect it. Given the distribution of the projected portfolio P&L, we would like to identify a parsimonious set of relevant factors $Z \equiv (Z_1, \ldots, Z_K)'$ that drive the portfolio P&L and whose joint distribution with the portfolio P&L $f_{\Pi_h, Z}$ is known.

More specifically, because the identification of the factors should be actionable and easy to interpret, the attribution should be linear. Thus, the attribution is defined by coefficients $b_h \equiv (b_{h,1}, \ldots, b_{h,K})'$, as follows

$$\Pi_h = \sum_{k=1}^{K} b_{h,k} Z_k. \qquad (34)$$

Note that the attribution to arbitrary factors in general gives rise to a portfolio-specific residual. The formulation (34) covers this case, by setting such residual as one of the factors Z_k, with attribution coefficient $b_{h,k} \equiv 1$.

Key concept. The Attribution Step decomposes the projected portfolio P&L linearly into a set of K relevant risk factors Z, yielding the K portfolio-specific exposures b_h

$$f_{\Pi_h,Z} \quad \mapsto \quad b_h \tag{35}$$

The relevant question is which attribution factors Z to use. Naturally, different intentions of the trader or portfolio manager call for different choices of attribution factors.

The most trivial attribution assigns the projected portfolio P&L back to the contributions from each security, i.e. $Z_k \equiv \Pi_k$ is the projected P&L from the generic k-th security, $b_{h,k} \equiv h_k$ are the holdings of the k-th security in the portfolio, and the number of factors is $K \equiv N$, i.e. the number of securities. Then the attribution equation (34) becomes the conservation law of money (26).

If on the other hand the trader wishes to hedge a given risk, say volatility risk, then he will choose as a factor Z_k the projected P&L Π_k of a truly actionable instrument, such as a variance swap, which might or might not have been part of the original portfolio.

Alternatively, the portfolio manager might wish to monitor the exposure to a given risk factor, without the need to hedge it. If for instance the manager is interested in the total "vega" of its portfolio for example, then he will use changes in implied volatility as one of the risk factors.

Furthermore, in case there exist too many possible factors or hedging instruments, the manager will want to express his portfolio as a function of only those few factors that truly affect the P&L.

Notice that (34) is a portfolio-specific top-down linear factor model. The flexible choice of the optimal attribution factors Z and optimal exposures b_h with flexible constraints which define this linear factor model, along with its connections with the linear factor models introduced in the Estimation Step $\mathbb{P}$ 2, is the spirit of the "Factors on Demand" approach in Meucci (2010c).

Illustration. In our stock and option example, we look at a simple attribution (34) to the original sources of risk. Accordingly, we set as attribution factors the stock compounded return $Z_s \equiv \ln(S_{T+\tau}/s_T)$ and the implied volatility log-change $Z_\sigma \equiv \ln(\Sigma_{T+\tau}/\sigma_T)$. Thus, we have $K \equiv 2$ factors. From the expression of the portfolio P&L (29) we immediately obtain

$$\Pi_h = b_{h,s} Z_s + b_{h,\sigma} Z_\sigma, \tag{36}$$

where the total exposures to Z_s and Z_σ read respectively

$$b_{h,s} \equiv h_s s_T + h_c \delta_{BS,T}, \quad b_{h,\sigma} \equiv h_c v_{BS,T}. \tag{37}$$

Pitfall. *"...If I use a factor model to estimate the returns distribution of some stocks and I want my portfolio to be neutral to a given factor, I simply make sure that the exposure to that factor is zero in my portfolio..."*. Ensuring a null-exposure coefficient for one factor does not guarantee immunization, because the given factor is in general correlated with other factors. To provide full immunization we must resort to Factors on Demand.

$\mathbb{P}$ 7 Evaluation

Up to this step, we have obtained the projected distribution of the P&L Π_h of a generic portfolio with holdings h and attributed it to relevant risk factors Z. In the evaluation step, the goal is to compare the P&L distribution of the current portfolio h with the P&L distribution of a different potential portfolio $\widetilde{h}$. Evaluation is one of the risk and portfolio manager's primary tasks.

Since each portfolio is represented by the whole projected distribution of its P&L, it is not possible to compare two portfolios in terms of which P&L is higher. To obviate this problem, typically practitioners rely on one or more summary statistics for the projected P&L distribution.

The most standard statistics are the expected value, the standard deviation and the Sharpe ratio, also known respectively as expected outperformance, tracking error and information ratio in the case of benchmarked portfolio management. Other measures include the value at risk (VaR), the expected shortfall (ES or CVaR), skewness, kurtosis, etc. More innovative statistics include coherent measures of risk aversion, see Artzner, Delbaen, Eber, and Heath (1997); spectral measures of risk aversion, see Acerbi (2002); and measures of diversification, such as the "effective number of bets", see Meucci (2009a).

We emphasize that, in this context, all the above are ex-ante measures of risk for the projected portfolio P&L Π_h, rather than ex-post measures of performance.

Key concept. The Evaluation Step consists of two sub-steps. The first sub-step is the computation of one or more summary statistics $\mathcal{S}$ for the projected distribution of the given portfolio P&L $\Pi_h$with holdings h

$$f_{\Pi_h} \quad \mapsto \quad \mathcal{S}(h). \tag{38}$$

The second, optional, sub-step is the attribution of the summary statistics $\mathcal{S}(h)$ to the fully flexible attribution factors Z utilized in the Attribution Step

$$f_{\Pi_h,Z}\ ,\ b_h \quad \mapsto \quad \mathcal{S}(h) = \sum_{k=1}^{K} b_{h,k}\mathcal{S}_k, \tag{39}$$

where $b_{h,k}$ represents the "amount" of the factor Z_k in the portfolio projected P&L and $\mathcal{S}_k$ represents the "per-unit" contribution to the statistic $\mathcal{S}(h)$ from the factor Z_k.

Illustration. In our simple normal market of one stock and one option, any portfolio is determined by the holdings $h \equiv (h_s, h_c)'$. Let us focus on the first sub-step (38) and let us compute the most basic summary statistics of the P&L, namely its expected value. Then from the distribution of a generic portfolio P&L (30) we obtain

$$\mathcal{S}(h_s, h_c) \equiv \mathrm{E}\{\Pi_h\} = \tau\mu_h = \tau h_s s_T \mu_s + \tau h_c (\delta_{BS,T}\mu_s + v_{BS,T}\mu_\sigma). \quad (40)$$

Similarly, if the manager cares about a measure of volatility, a suitable measure is the standard deviation

$$\mathcal{S}(h_s, h_c) \equiv \mathrm{Sd}\{\Pi_h\} = \sqrt{\tau}\sigma_h, \quad (41)$$

where σ_h is defined in (32).

For the optional summary statistics attribution sub-step (39), a simple linear decomposition that mirrors the attribution equation (34) is not feasible. For instance, for the standard deviation it is well known that $\mathrm{Sd}\{\Pi_h\} \neq \sum_{k=1}^{K} b_{h,k} \mathrm{Sd}\{Z_k\}$. However, notice that numerous summary statistics such as expectation, standard deviation, VaR, ES, and spectral measures display an interesting feature: they are homogeneous, i.e. by doubling all the holdings in the portfolio, those summary statistics also double. As proved by Euler, for homogeneous statistics the following identity holds true

$$\mathcal{S}(h) = \sum_{k=1}^{K} b_{h,k} \frac{\partial \mathcal{S}(h)}{\partial b_{h,k}}. \quad (42)$$

Therefore, if the summary statistics is homogeneous, we can take advantage of Euler's identity (42) to perform the summary statistics attribution sub-step (39), which becomes (42).

In particular, for the VaR, the decomposition (42) amounts to the classical definition of marginal contributions to VaR, see e.g. Garman (1997), and, for the standard deviation, the decomposition (42) amounts to the "hot-spots", see Litterman (1996).

We recall that the simplest case of the flexible, top-down, Factors on Demand attribution of the portfolio P&L (34) is the bottom-up attribution to the individual securities through the conservation law of money (26). Similarly, the simplest case of attribution of the summary statistics (39) is the attribution of the summary statistics $\mathcal{S}(h)$ to the individual securities

$$\mathcal{S}(h) = \sum_{n=1}^{N} h_n \frac{\partial \mathcal{S}(h)}{\partial h_n}. \quad (43)$$

Illustration. To illustrate the attribution to a summary statistic of the portfolio projected P&L, we rely on our example of a stock and a call option. We focus on the standard deviation (41).

The exposure $b_{h,s}$ of the projected portfolio P&L (34) to the stock factor $Z_s \equiv \ln(S_{T+\tau}/s_T)$ and the exposure $b_{h,\sigma}$ to the implied volatility factor

$Z_\sigma \equiv \ln(\Sigma_{T+\tau}/\sigma_T)$ were calculated in (37). Then the attribution (42) to each of the two risk drivers of the standard deviation of the projected portfolio P&L becomes

$$\begin{pmatrix} \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial b_{h,s}} \\ \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial b_{h,\sigma}} \end{pmatrix} = \frac{\sqrt{\tau}}{\sigma_h} \begin{pmatrix} \sigma_s^2 & \rho\sigma_s\sigma_\sigma \\ \rho\sigma_s\sigma_\sigma & \sigma_\sigma^2 \end{pmatrix} \begin{pmatrix} h_s s_T + h_c \delta_{BS,T} \\ h_c v_{BS,T} \end{pmatrix}, \tag{44}$$

where σ_h is defined in (32), see the proof in the appendix. The total contributions to risk from the factors follow by multiplying the entries on the left hand side of (44) by the respective exposures (37).

For the attribution to the individual securities, i.e. the stock and the call option, a similar calculation yields

$$\begin{pmatrix} \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial h_s} \\ \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial h_\sigma} \end{pmatrix} = \frac{\sqrt{\tau}}{\sigma_h} \begin{pmatrix} s_T^2\sigma_s^2 & \sigma_{\Pi_s,\Pi_c} \\ \sigma_{\Pi_s,\Pi_c} & \sigma_{\Pi_c}^2 \end{pmatrix} \begin{pmatrix} h_s \\ h_c \end{pmatrix}, \tag{45}$$

where

$$\sigma_{\Pi_c}^2 \equiv \sigma_s^2\delta_{BS,T}^2 + \sigma_\sigma^2 v_{BS,T}^2 + 2\sigma_\sigma\sigma_s\rho\delta_{BS,T}v_{BS,T} \tag{46}$$

$$\sigma_{\Pi_s,\Pi_c} \equiv \delta_{BS,T}s_T\sigma_s^2 + s_T v_{BS,T}\sigma_\sigma\sigma_s\rho, \tag{47}$$

see the proof in the appendix. The total contributions to risk from the stock and the call option follow by multiplying the entries on the left hand side of (45) by the respective holdings h_s and h_c.

The computation of the summary statistics $\mathcal{S}(h)$ is hard to perform in practice, unless the market is normal as in our example (41), because complex multiple integrals are involved. For instance, using the same notation as in (28), the VaR with confidence c is defined by

$$1 - c \equiv \int_{h'\pi \leq VaR} f_\Pi(\pi_1, \ldots, \pi_N)\, d\pi_1 \cdots d\pi_N. \tag{48}$$

To address this problem, one can rely on approximation methods such as the Cornish-Fisher expansion, or the elliptical assumption, see Meucci (2005a) for a review. The computation of the partial derivatives for the decomposition (42) of the summary statistics is even harder, unless the market is normal as in our example (44)-(45). Fortunately, these computations become simple when the market distribution is represented in terms of scenarios, see Meucci (2010c).

Before concluding, we must address two key problems of risk and portfolio management: estimation risk, introduced in the Estimation Step $\mathbb{P}$ 2, and liquidity risk, introduced in the Pricing Step $\mathbb{P}$ 4.

As far as estimation risk is concerned, the projected distribution of the P&L Π_h that we are evaluating is only an estimate, not the true projected distribution, which is unknown. Therefore, estimation risk affects the Evaluation Step. As a simple, effective way to address this issue, risk managers perform stress-test or scenario analysis, which amounts to evaluating the P&L under specific, typically extreme or historical, realizations of the risk drivers. A more advanced general approach to stress testing is "Fully Flexible Probabilities", see Meucci (2010d), which allows the portfolio manager to assign non-equal probabilities to the historical scenarios, according to such criteria as exponential smoothing, rolling window, kernel conditioning and, more flexibly, the generalized Bayesian approach "entropy pooling".

As far as liquidity risk is concerned, the projected distribution of the P&L Π_h that we are evaluating does not account for the effect of our own trading. A theory to correct for this effect in the context of risk management was developed in Cetin, R., and Protter (2004) and Acerbi and Scandolo (2007). For an easy to implement liquidity adjustment to the P&L distribution refer to Meucci and Pasquali (2010).

Pitfall. *"...To compute the volatility of the P&L we can simply run the sample standard deviation of the past P&L realizations..."*. The history of the past P&L can be informative only if the P&L is an invariant. This seldom happens, consider for instance the P&L generated by a buy-and-hold strategy in one call option. In general, one has to follow all the steps of the $\mathbb{P}$rayer to compute risk numbers.

Pitfall. *"...To compute the VaR I can multiply the standard deviation by a threshold number such as 1.96..."*. This calculation is only correct with very specific, unrealistic, typically normal, models for the market distribution.

$\mathbb{P}$ 8 Optimization

In the Evaluation Step $\mathbb{P}$ 7, the risk manager or portfolio manager obtains a set of computed summary statistics $\mathcal{S}$ to assess the goodness of a portfolio with holdings h. These statistics can be combined in a subjective manner to give rise to new statistics. For instance, a portfolio with expected return of 2% and standard deviation of 5% could be good for an investor with low risk tolerance, but bad for an aggressive trader. In this case, a trade-off statistic $\mathcal{S}(h) \equiv \mathrm{E}\{\Pi_h\} - \gamma \,\mathrm{Sd}\{\Pi_h\}$ can rank the portfolios according to the preferences of the investor, reflected in the parameter γ. Alternatively, we can use a subjective utility function u and rank portfolios based on expected utility $\mathcal{S}(h) \equiv \mathrm{E}\{u(\Pi_h)\}$.

More in general, we call index of satisfaction the function that translates the P&L distribution of the portfolio with holdings $h \equiv (h_1, \ldots, h_N)'$ into a personal preference ranking. We denote the index of satisfaction by the general notation $\mathcal{S}(h)$ used in (38) for the evaluation summary statistics, because any index of satisfaction is also a summary statistic.

Given an index of satisfaction $\mathcal{S}(h)$, it is now possible to optimize the holdings h accordingly. Portfolio optimization is the primary task of the portfolio manager.

Clearly, the optimal allocation should not violate a set of hard constraints, such as the budget constraint, or soft constraints, such as constraints on leverage, risk, etc. We denote by $\mathcal{C}$ the set of all such constraints and by "$h \in \mathcal{C}$" the condition that the allocation h satisfies the given constraints.

Key concept. The Optimization Step is the process of computing the holdings that maximize satisfaction, while not violating a given set of investment constraints

$$h^* \equiv \underset{h \in \mathcal{C}}{\operatorname{argmax}} \{\mathcal{S}(h)\}. \tag{49}$$

We emphasize that the choice of the most suitable index of satisfaction $\mathcal{S}$, as well as the specific constraints $\mathcal{C}$, vary widely depending on the profile of the securities P&L distribution, the investment horizon, and other features of the market and the investor.

Illustration. In our stock and option example we can compute the best hedge for one call option. In this context, the general framework (49) becomes

$$(h_s, h_c)^* \equiv \underset{h_c \equiv 1}{\operatorname{argmax}} \{-\operatorname{Sd}\{\Pi_h\}\}. \tag{50}$$

Then the first order condition on the P&L standard deviation, computed in (30)-(32), yields

$$h_s \equiv -\frac{\delta_{BS,T}}{s_T} - \frac{v_{BS,T}}{s_T}\rho\frac{\sigma_\sigma}{\sigma_s}. \tag{51}$$

If the correlation ρ between implied volatility and underlying were null, the best hedge would consist in shorting a "delta" amount of underlying. In general ρ is substantially negative: for instance, the sample correlation between VIX and S&P 500 is $\rho \approx -0.7$. Therefore, a correction to the simplistic delta hedge must be applied.

In general, the numerical optimization (49) is a challenging task. To address this issue one can resort to the two-step mean-variance heuristic. First, the mean-variance efficient frontier is computed

$$h_\lambda \equiv \underset{h \in \mathcal{C}}{\operatorname{argmax}} \{\operatorname{E}\{\Pi_h\} - \lambda \operatorname{Vr}\{\Pi_h\}\}, \quad \lambda \in \mathbb{R}. \tag{52}$$

This step reduces the dimension of the problem from N, the dimension of the market, to 1, the value of λ. The optimization (52) can be solved by variations of quadratic programming. The optimization becomes particularly efficient when a linear factor model makes the covariance of the securities P&L's sparse, see Meucci (2010h).

Second, the optimal portfolio is selected by a one-dimensional search

$$h^* \equiv \underset{\lambda \in \mathbb{R}}{\operatorname{argmax}} \{\mathcal{S}(h_\lambda)\}. \tag{53}$$

The optimization (53) can be performed by a simple grid-search.

As it was the case for the Evaluation Step $\mathbb{P}$ 7, we must address estimation risk, introduced in the Estimation Step $\mathbb{P}$ 2: the projected distribution of the P&L that we are optimizing is only an estimate, not the true projected distribution, which is unknown. As it turns out, the optimal portfolio is extremely sensitive to the input estimated distribution, which makes estimation risk particularly relevant for the Optimization Step $\mathbb{P}$ 8.

To address the issue of estimation risk, portfolio managers rely on more advanced approaches than the simple two-step mean-variance heuristic (52)-(53). These advanced approaches include robust optimization, which relies on cone programming, see Ben-Tal and Nemirovski (2001) and Cornuejols and Tutuncu (2007); Bayesian allocation, see Bawa, Brown, and Klein (1979); robust Bayesian allocation, see Meucci (2005b); and resampling, see Michaud (1998). We refer to Meucci (2005a) for an in-depth review.

Since estimation is imperfect, tactical portfolio construction enhances performance by blending market views and predictive signals into the estimated market distribution. Well-known techniques to perform tactical portfolio construction are the approach by Grinold and Kahn (1999), which mixes signals based on linear factor models for returns; the Bayesian inspired methodology by Black and Litterman (1990); and the generalized Bayesian approach "Entropy Pooling" in Meucci (2008).

Due to the rapid decay of the quality of predictive tactical signals, managers separate tactical portfolio construction from strategic rebalancing, which takes into account shortfall and drawdown control and is optimized based on techniques that range from dynamic programming to heuristics, see e.g. Merton (1992), Grossman and Zhou (1993), Browne and Kosowski (2010), and refer to Meucci (2010g) for a review and code.

Finally, liquidity risk, discussed in the Pricing Step $\mathbb{P}$ 4, impacts the Optimization Step: transaction costs must be paid to reallocate capital and the process of executing a transaction impacts the execution price. Therefore, market impact models must be embedded in the portfolio optimization process. The standard approach in this direction is a power-law impact model, see e.g. Keim and Madhavan (1998).

Pitfall. *"...Mean-variance assumes normality..."*. The mean-variance approach does not assume normality: any market distribution can be fed into the two-step process (52)-(53).

$\mathbb{P}$ 9 Execution

The Optimization Step $\mathbb{P}$ 8 delivers a desired allocation $h^* \equiv (h_1^*, \ldots, h_N^*)'$. To achieve the desired allocation, it is necessary to rebalance the positions from the current allocation $h_T \equiv (h_{1,T}, \ldots, h_{N,T})'$. This rebalancing is not executed immediately. As time evolves, the external market conditions change. Simultaneously, the internal state of the book, represented by the updated allocation, the updated constraints, etc., changes dynamically. To execute a rebalancing trade, this information must be optimally processed.

Key concept. The Execution Step processes the evolving external market information i_t^m and internal book information i_t^b to attain the target portfolio h^* by a sequence of transactions at given prices $p_t \equiv (p_{t,1}, \ldots, p_{t,N})'$

$$h^* \; , \; \{i_t^m\}_{t \geq T} \; , \; \{i_t^b\}_{t \geq T} \quad \mapsto \quad \{p_t\}_{t \geq T} \, . \tag{54}$$

Note that often the execution step is implemented in aggregate across different books. This aggregation is particularly useful, as it allows for netting of conflicting buy-sell orders from different traders or managers. Performing this netting in the Optimization Step ℙ 8 would be advisable, see e.g.O'Cinneide, Scherer, and X. (2006) and Stubbs and Vandenbussche (2007). However, this can be hard in practice.

Execution is closely related to liquidity risk, first introduced in the Pricing Step ℙ 4. The literature on liquidity, market impact, algorithmic trading and optimal execution is very broad, see e.g. Almgren and Chriss (2000) and Gatheral (2010).

Illustration. For illustrative purposes, we mention the simplest execution algorithm, namely "trading at all costs". This approach disregards any information on the market or the book and delivers immediately the desired final allocation by depleting the cash reserve. We emphasize that trading at all costs can be heavily suboptimal.

Pitfall. *"...The Execution Step ℙ 9 should be embedded into the Optimization Step ℙ 9..."*. In practice it is not possible to process simultaneously real-time information and all the previous steps of the ℙrayer. Furthermore, execution works best across all books, whereas optimization is specific to each individual manager.

ℙ 10 Ex-Post Analysis

In the Execution Step ℙ 9 we implemented the allocation $h^* \equiv (h_1^*, \ldots, h_N^*)$ for the period between the current date T and the investment horizon $T + \tau$. Upon reaching the horizon, we must evaluate the P&L π_{h^*} realized over the horizon by the allocation, where the lower-case notation emphasizes that the P&L is no longer a random variable, but rather a number that we observe ex-post.

Key concept. The Ex-Post Analysis Step identifies the contributions to the realized P&L from different decision makers and market factors

$$\pi_{h^*} \quad \mapsto \quad (\pi_a, \pi_b, \cdots). \tag{55}$$

Ex-post performance analysis is a broad subject that attracts tremendous attention from practitioners, as their compensation is ultimately tied to the results of this analysis. Ex-post performance can be broken down into two components: performance of the target portfolio from the Optimization Step $\mathbb{P}$ 8 and slippage performance from the Execution Step $\mathbb{P}$ 9.

To analyze the ex-post performance of the target portfolio, the most basic framework decomposes this performance into an allocation term and a selection term, see e.g. Brinson and Fachler (1985). More recent work attributes performance to different factors, such as foreign exchange swings or yield curve movements, consistently with the Attribution Step $\mathbb{P}$ 6.

The slippage component can be decomposed into unexecuted trades and implementation shortfall attributable to market impact, see Perold (1988).

Furthermore, performance must be fairly decomposed across different periods, see e.g. Carino (1999) and Menchero (2000).

Illustration. In our stock and option example, we can decompose the realized P&L into the cost incurred by the "trading at all costs" strategy, a stock component, an implied volatility component, and a residual. In particular, the stock component reads $b_{h,s} \ln(s_{T+\tau}/s_T)$ as in (36), and the implied volatility component reads $b_{h,s} \ln(\sigma_{T+\tau}/s_T)$ The residual is the plug-in term that makes the sum of all components add up to the total realized P&L.

Pitfall. *"...I prefer geometric performance attribution, because it can be aggregated exactly across time and across currencies..."*. The geometric, or multiplicative approach to ex-post performance is arguably less intuitive, because it does not accommodate naturally a linear decomposition in terms of different risk or decisions factors.

References

Acerbi, C., 2002, Spectral measures of risk: A coherent representation of subjective risk aversion, *Journal of Banking and Finance* 26, 1505–1518.

———, and G. Scandolo, 2007, Liquidity risk theory and coherent measures of risk, *Working Paper*.

Albanese, C., K. Jackson, and P. Wiberg, 2004, A new Fourier transform algorithm for value at risk, *Quantitative Finance* 4, 328–338.

Almgren, R., and N. Chriss, 2000, Optimal execution of portfolio transactions, *Journal of Risk* 3, 5–39.

Artzner, P., F. Delbaen, J. M. Eber, and D. Heath, 1997, Thinking coherently, *Risk Magazine* 10, 68–71.

Bawa, V. S., S. J. Brown, and R. W. Klein, 1979, *Estimation Risk and Optimal Porfolio Choice* (North Holland).

Ben-Tal, A., and A. Nemirovski, 2001, *Lectures on modern convex optimization: analysis, algorithms, and engineering applications* (Society for Industrial and Applied Mathematics).

Black, F., and R. Litterman, 1990, Asset allocation: combining investor views with market equilibrium, *Goldman Sachs Fixed Income Research*.

Brinson, G., and N. Fachler, 1985, Measuring non-US equity portfolio performance, *Journal of Portfolio Management* pp. 73–76.

Browne, S., and R. Kosowski, 2010, Drawdown minimization, *Encyclopedia of Quantitative Finance, Wiley*.

Carino, D., 1999, Combining attribution effects over time, *Journal of Performance Measurement* pp. 5–14.

Cetin, U, Jarrow R., and P. Protter, 2004, Liquidity risk and arbitrage pricing theory, *Finance and Stochastics* 8, 311–341.

Cherubini, U., E. Luciano, and W. Vecchiato, 2004, *Copula Methods in Finance* (Wiley).

Cornuejols, G., and R. Tutuncu, 2007, *Optimization Methods in Finance* (Cambridge University Press).

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

Gatheral, J., 2010, No-dynamic-arbitrage and market impact, *Quantitative Finance* 10, 749–759.

Grinold, R. C., and R. Kahn, 1999, *Active Portfolio Management. A Quantitative Approach for Producing Superior Returns and Controlling Risk* (McGraw-Hill) 2nd edn.

Grossman, S., and Z. Zhou, 1993, Optimal investment strategies for controlling drawdowns, *Mathematical Finance* 3, 241–276.

Keim, D. B., and A. Madhavan, 1998, The cost of institutional equity trades: An overview, *Rodney L. White Center for Financial Research Working Paper Series.*

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

Menchero, J., 2000, An optimized approach to linking attribution effects over time, *Journal of Performance Measurement* 5.

Merton, R. C., 1992, *Continuous-Time Finance* (Blackwell).

Meucci, A., 2005a, *Risk and Asset Allocation* (Springer) Available at http://symmys.com.

———, 2005b, Robust Bayesian asset allocation, *Working Paper* Article and code available at http://symmys.com/node/102.

———, 2008, Fully flexible views: Theory and practice, *Risk* 21, 97–102 Article and code available at http://symmys.com/node/158.

———, 2009a, Managing diversification, *Risk* 22, 74–79 Article and code available at http://symmys.com/node/199.

———, 2009b, Review of discrete and continuous processes in finance: Theory and applications, *Working paper* Article and code available at http://symmys.com/node/131.

———, 2009c, Review of statistical arbitrage, cointegration, and multivariate Ornstein-Uhlenbeck, *Working Paper* Article and code available at http://symmys.com/node/132.

———, 2010a, Annualization and general projection of skweness, kurtosis, and all summary statistics, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* August, 55–56 Article and code available at http://symmys.com/node/136.

———, 2010b, Common misconceptions about 'beta' - hedging, estimation and horizon effects, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* June, 42–45 Article and code available at http://symmys.com/node/165.

——— , 2010c, Factors on Demand - building a platform for portfolio managers risk managers and traders, *Risk* 23, 84–89 Article and code available at http://symmys.com/node/164.

——— , 2010d, Historical scenarios with fully flexible probabilities, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* December, 40–43 Article and code available at http://symmys.com/node/150.

——— , 2010e, Linear vs. compounded returns - common pitfalls in portfolio management, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* April, 52–54 Article and code available at http://symmys.com/node/141.

——— , 2010f, Return calculations for leveraged securities and portfolios, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* October, 40–43 Available at http://symmys.com/node/140.

——— , 2010g, Review of dynamic allocation strategies: Convex versus concave management, *Working Paper* Article and code available at http://symmys.com/node/153.

——— , 2010h, Review of linear factor models: Unexpected common features and the systematic-plus-idiosyncratic myth, *Working paper* Article and code available at http://www.symmys.com/node/336.

——— , 2010i, Square-root rule, covariances and ellipsoids - how to analyze and visualize the propagation of risk, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* February, 52–53 Article and code available at http://symmys.com/node/137.

——— , 2011a, New copulas for risk and portfolio management, *Risk* 24, 83–86 Article and code available at http://symmys.com/node/335.

——— , 2011b, "P" versus "Q": Differences and commonalities between the two areas of quantitative finance, *GARP Risk Professional - "The Quant Classroom by Attilio Meucci"* February, 43–44 Article and code available at http://symmys.com/node/62.

——— , and S. Pasquali, 2010, Liquidity-adjusted portfolio distribution and liquidity score, *Working Paper* Article and code available at http://symmys.com/node/350.

Michaud, R. O., 1998, *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation* (Harvard Business School Press).

O'Cinneide, C., B. Scherer, and Xu X., 2006, Pooling trades in quantitative invest- ment process, *Journal of Investing* 32, 33–43.

Paparoditis, E., and D. N. Politis, 2009, Resampling and subsampling for financial time series, *in Handbook of Financial Time Series, (Andersen, T., Davis, R., Kreiss, J.-P., and Mikosch, T., Eds.), Springer, Berlin-Heidelberg* pp. 983–999.

Perold, A. F., 1988, The implementation shortfall: Paper vs. reality, *Journal of Portfolio Management* 14, 4–9.

Stubbs, R. A., and D. Vandenbussche, 2007, Multi-portfolio optimization and fairness in allocation of trades, *Axioma Research Paper No. 013*.

A Appendix

First, consider the following rule, which holds for any square matrix a and conformable vector x

$$\frac{\partial\sqrt{x'ax}}{\partial x} = \frac{ax}{\sqrt{x'ax}}. \tag{56}$$

To prove (44) we recall from the attribution to the risk factors (37) that the standard deviation (41) reads

$$\tau\sigma_h^2 = (\mathrm{Sd}\{\Pi_h\})^2 = \tau \begin{pmatrix} b_{h,s} & b_{h,\sigma} \end{pmatrix} \begin{pmatrix} \sigma_s^2 & \sigma_\sigma\sigma_s\rho \\ \sigma_\sigma\sigma_s\rho & \sigma_\sigma^2 \end{pmatrix} \begin{pmatrix} b_{h,s} \\ b_{h,\sigma} \end{pmatrix}. \tag{57}$$

Then (44) follows from (56).

To prove (45), we recall that the stock P&L (19) reads

$$\Pi_s \approx s_T (\ln S_{T+\tau} - \ln s_T) \tag{58}$$

and the call option P&L (22) reads

$$\Pi_c \approx (\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T} + (\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T} \tag{59}$$

We recall from (14) that all the log-changes above are jointly normal. Therefore the entries of covariance matrix read

$$\begin{aligned}
\tau\sigma_{\Pi_s}^2 &\equiv \mathrm{V}\{\Pi_s\} = \mathrm{V}\{s_T(\ln S_{T+\tau} - \ln s_T)\} = \tau s_T^2\sigma_s^2 && (60)\\
\tau\sigma_{\Pi_c}^2 &\equiv \mathrm{V}\{\Pi_c\} = \mathrm{V}\{(\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T} + (\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T}\} && (61)\\
&= \mathrm{V}\{(\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T}\} + \mathrm{V}\{(\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T}\}\\
&\quad +2\,\mathrm{Cv}\{(\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T}, (\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T}\}\\
&= \tau\left(\sigma_s^2\delta_{BS,T}^2 + \sigma_\sigma^2 v_{BS,T}^2 + 2\sigma_\sigma\sigma_s\rho\delta_{BS,T}v_{BS,T}\right)\\
\tau\sigma_{\Pi_s,\Pi_c} &\equiv \mathrm{Cv}\{\Pi_s,\Pi_c\} = \mathrm{Cv}\{s_T(\ln S_{T+\tau} - \ln s_T), && (62)\\
&\quad (\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T} + (\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T}\}\\
&= \mathrm{Cv}\{s_T(\ln S_{T+\tau} - \ln s_T), (\ln S_{T+\tau} - \ln s_T)\,\delta_{BS,T}\}\\
&\quad + \mathrm{Cv}\{s_T(\ln S_{T+\tau} - \ln s_T), (\ln \Sigma_{T+\tau} - \ln \sigma_T)\,v_{BS,T}\}\\
&= \tau\left(\delta_{BS,T}s_T\sigma_s^2 + s_T v_{BS,T}\sigma_\sigma\sigma_s\rho\right),
\end{aligned}$$

where $\mathrm{Cv}\{X,Y\}$ denotes the covariance between X and Y.

As in (29), the portfolio P&L reads

$$\Pi_h = h_s\Pi_s + h_c\Pi_c \tag{63}$$

Thus the standard deviation (41) reads

$$(\mathrm{Sd}\{\Pi_h\})^2 = \tau \begin{pmatrix} h_s & h_c \end{pmatrix} \begin{pmatrix} \sigma_{\Pi_s}^2 & \sigma_{\Pi_s,\Pi_c} \\ \sigma_{\Pi_s,\Pi_c} & \sigma_{\Pi_c}^2 \end{pmatrix} \begin{pmatrix} h_s \\ h_c \end{pmatrix} \tag{64}$$

Then (45) follows from (56).