"The ℙrayer": The 10 Steps of Advanced Risk and Portfolio Management

The first of a two-part article on the path from data analysis to optimal execution across all asset classes and investment styles.

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.

Figure 1: The "ℙrayer": A 10-Step Blueprint for Risk and Portfolio Management

ℙ1: Quest for Invariance	risk drivers Y_t	$\mapsto$	invariants $\varepsilon_{t\to t+1}$	
ℙ2: Estimation	invariants time series $\{\epsilon_{t\to t+1}\}_{t=1,\ldots,T}$	$\mapsto$	invariants distribution f_ε	Estimation Risk
ℙ3: Projection	invariants distribution f_ε	$\mapsto$	risk drivers distribution $f_{Y_{T+\tau}}$	
ℙ4: Pricing	risk drivers distribution $f_{Y_{T+\tau}}$	$\mapsto$	securities P&L distribution $f_{\Pi_{T\to T+\tau}}$	Liquidity Risk
ℙ5: Aggregation	securities P&L / portfolio positions $f_{\Pi_{T\to T+\tau}}, h$	$\mapsto$	portfolio P&L distribution f_{Π_h}	
ℙ6: Attribution	portfolio P&L distribution f_{Π_h}	$\mapsto$	factors distribution / exposures f_Z, b_h	
ℙ7: Evaluation	factors distribution / exposures f_Z, b_h	$\mapsto$	summary statistics / decomposition $\mathcal{S}(h), \frac{\partial \mathcal{S}(h)}{\partial b_h}$	
ℙ8: Optimization	satisfaction / constraints $\mathcal{S}(h), \mathcal{C}(h)$	$\mapsto$	optimal positions h^*	
ℙ9: Execution	target / market info / book info $h^*, i^m_{t\geq T}, i^b_{t\geq T}$	$\mapsto$	execution prices $p_{t\geq T}$	
ℙ10: Ex-Post Analysis	realized P&L π_{h^*}	$\mapsto$	allocation, selection, curve,... $(\pi_a, \pi_b, \cdots)$	

In this article, we present Steps 1 – 4 of "the ℙrayer," a blueprint of 10 sequential steps for quants across the board to achieve their common goal (see Figure 1, above). Steps 5 – 10 will be discussed in the next *Quant Classroom* column, which will be published in the June issue of *Risk Professional*.

By following the ℙrayer, 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 "ℙ" to signify the true probability space of the buy-side P&L, which stands in contrast to the risk-neutral probability space "ℚ" used on the sell-side to price derivatives (see Meucci [2011b]).

Each step of the ℙrayer 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.

ℙ1: Quest for Invariance

The "quest for invariance" is the first step of the ℙrayer, 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; and (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, let's 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

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

where $v_t \equiv e{-}t$ 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

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

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 v_t. Therefore, the option formula (2) defines 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 v_t are needed.

Since the time to expiry is deterministic, the call option requires two risk drivers to determine its price fully:

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

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; and (b) they become known at the end of the step – i.e., at time $t+1$.

THE QUANT CLASSROOM by ATTILIO MEUCCI

THE QUANT CLASSROOM by ATTILIO MEUCCI

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 (2005).

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 $\varepsilon_{t\rightarrow t+1}$. The most basic dynamics, yet the most statistically robust, which connects the invariants $\varepsilon_{t\rightarrow t+1}$ with the risk drivers Y_t, is the random walk

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

More advanced processes for the risk drivers account for such features as autocorrelations, stochastic volatility and long memory. We refer to Meucci (2009a) 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 also refer to Meucci (2009b) 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 aforementioned scatter-plot generally indicates that the compounded return $\ln (S_{t+1}/S_t)$ are approximately invariants

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

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

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

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 (2009a) 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 $\varepsilon_{t\rightarrow t+1}$ 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 [2009a]).
"...*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.

ℙ2: Estimation

As highlighted in the Quest for Invariance Step ℙ1, 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 $\{\boldsymbol{\epsilon}_t\}$ of the invariants $\boldsymbol{\varepsilon}$ and optionally additional information i_T that is available at the current time T,

$$\{\epsilon_t\}_{t=1,\ldots,T}\ ,\ i_T \quad \mapsto \quad f_\varepsilon. \tag{9}$$

Simple estimation approaches only process the time series of the invariants, but various advanced techniques also process information such as market-implied forward looking quantities, subjective Bayesian priors, etc.

The simplest of all estimators for the invariants distribution is the nonparametric 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 $\{\boldsymbol{\epsilon}_t\}_{t=1,\ldots,T}$ 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 $\boldsymbol{\varepsilon}$ 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 2x1 expectation vector $\boldsymbol{\mu} \equiv (\mu_s, \mu_\sigma)'$ and 2x2 covariance matrix, $\boldsymbol{\sigma}^2$ as below:

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

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 ℙ2, as well as, among others, in the Attribution Step ℙ4 and the Optimization Step ℙ8. We refer to Meucci (2010c) 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]).

The last 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.

Advanced estimation techniques that attempt to address this issue include multivariate robust estimation with low influence function and high breakdown point, and multivariate Bayesian estimation. We refer to Meucci (2005) for an in-depth review.

Alternatively, to address estimation risk, practitioners rely on scenario analysis, where one joint scenario for the risk drivers, plausible or extreme, is isolated, and its effect on the P&L is evaluated, as we will see in Step ℙ4.

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* $K^2/2+Q$" 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 (2010c).

ℙ3: 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 ℙ2.

We denote the current time as $t \equiv T$ and the generic invest-

THE QUANT CLASSROOM by ATTILIO MEUCCI

ment 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 (2010d). 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 $\varepsilon_{T \to t+\tau} \sim \mathrm{N}(\tau\mu, \tau\sigma^2)$. Therefore, 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.

$\mathbb{P}$4: Pricing

Now that we have the distribution of the risk drivers at the horizon $Y_{T+\tau}$ from the Projection Step $\mathbb{P}3$, 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. Consequently, 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 $\overline{y} \equiv y_T$; $\partial_y p\ (\overline{y}; i_T)$ denotes the vector of the first derivatives; and $\partial_{yy} p\ (\overline{y}; i_T)$ 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 $p(y) = e^y$. 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

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

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

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

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 $\delta_{BS,T} \equiv \partial p_{BS}\left(y_s, y_\sigma\right) / \partial y_s$ is the option's current Black-Scholes "delta" and $\delta_{BS,T} \equiv \partial p_{BS}\left(y_s, y_\sigma\right) / \partial y_s$ 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, as follows:

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

where

$$\mu_c \equiv \delta_{BS,T}\mu_s + v_{BS,T}\mu_\sigma \tag{24}$$

$$\sigma_c^2 \equiv \delta_{BS,T}^2\sigma_s^2 + v_{BS,T}^2\sigma_\sigma^2 + 2\delta_{BS,T}v_{BS,T}\rho\sigma_s\sigma_\sigma. \qquad (25)(25)$$

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.

By following the $\mathbb{P}$rayer, quants can avoid common pitfalls and ensure that they are not missing important points in their models.

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, we have the problem of liquidity risk. The pricing step assumes the existence of one single price, which is fully determined by the risk drivers $P_t = p(Y_t; i_T)$, 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.

Techniques to model liquidity risk are very different from other types of market risk. We will discuss in Steps $\mathbb{P}7$, $\mathbb{P}8$ and $\mathbb{P}9$ methodologies to address liquidity risk.

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.

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.

...To be continued in the next "classroom."

Attilio Meucci is the chief risk officer at Kepos Capital LP. He runs the 6-day "Advanced Risk and Portfolio Management Bootcamp," see symmys.com. *The author is grateful to Garli Beibi, Luca Spampinato and an anonymous referee.*

REFERENCES

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

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

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

Ibid, 2009a. "Review of discrete and continuous processes in finance: Theory and applications," Working Paper, available at http://ssrn.com/abstract=1373102.

Ibid, 2009b. "Review of statistical arbitrage, cointegration, and multivariate Ornstein-Uhlenbeck," Working Paper, available at http://ssrn.com/abstract=1404905.

Ibid, 2010a. "The Quant Classroom by Attilio Meucci: Annualiza tion and general projection of skewness, kurtosis and all summary statistics," *Risk Professional*, August, 55—56. Available at http://ssrn.com/abstract=1635484.

Ibid, 2010b. "The Quant Classroom by Attilio Meucci: Common misconceptions about 'beta' - hedging, estimation and horizon effects," *Risk Professional* , June, 42—45. Available at http://ssrn.com/abstract=1619923.

Ibid, 2010c. "Review of linear factor models: Unexpected common features and the systematic-plus-idiosyncratic myth," Working Paper, available at http://ssrn.com/abstract=1635495.

Ibid, 2010d, "The Quant Classroom by Attilio Meucci: Square-root rule, covariances and ellipsoids - how to analyze and visualize the propagation of risk," *Risk Professional*, February, 52—53. Available at http://ssrn.com/abstract=1548162.

Ibid, 2011a. "Flexible copulas for buy-side and sell-side," Working Paper, available at http://ssrn.com/abstract=1752702.

Ibid, 2011b. "The Quant Classroom by Attilio Meucci: "P" versus "Q": Differences and commonalities between two areas of quantitative finance," *Risk Professional*, February, 43—44. Available at http://ssrn.com/abstract=1717163.

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