"The ℙrayer": The 10 Steps of Advanced Risk and Portfolio Management – Part 2

The second 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

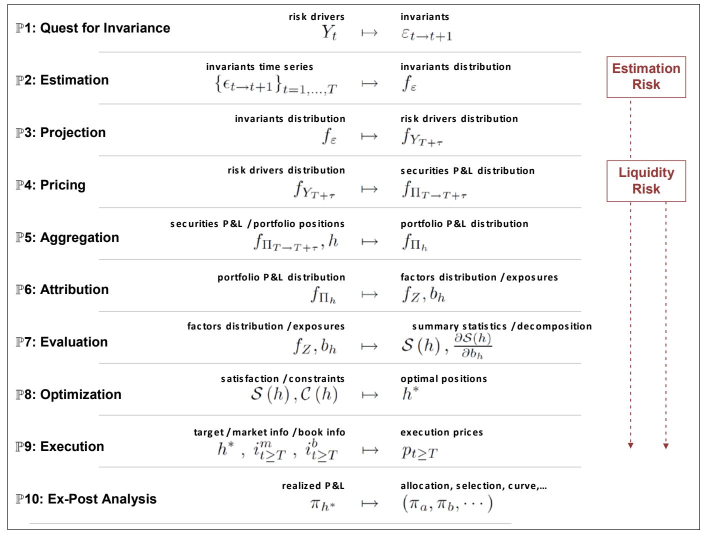

Here, we continue our journey through "The ℙrayer," a blueprint of 10 sequential steps for quants across the board to achieve their common goal. Steps 1-4 were discussed in the Quant Classroom published in the April issue of *Risk Professional* (also available at symmys.com). In this article, we discuss the remaining Steps 5-10. Please refer to Figure 1 (left) for a map of The ℙrayer.

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 Pricing Step ℙ 4. We drop from the notation of the P&L the subscript T →T+τ, 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, as follows:

$$\Pi_h = \sum_{n=1}^{N} h_n \Pi_n, \qquad (1)$$

THE QUANT CLASSROOM by ATTILIO MEUCCI

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

$$f_\Pi\,,\ h \quad \mapsto \quad f_{\Pi_h} \qquad (2)$$

Given one single scenario for the risk drivers $Y_{T+\tau}$, and thus for the securities P&L's $\Pi \equiv (\Pi_1, \ldots \Pi_N)'$, the computation of the portfolio P&L Π_h is immediately determined by the conservation law of money (1) as the sum of the single-security P&Ls in that scenario.

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

$$f_{\Pi_h}(x)\,dx = \int_{h'\pi \in dx} f_\Pi(\pi_1, \ldots, \pi_N)\, d\pi_1 \cdots d\pi_N, \qquad (3)$$

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 (1) is simply repeated in a discrete way, scenario-by-scenario.

Illustration. In our example with a stock and a call option, whose P&Ls 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 (1) to the stock P&L and the option P&L, which were obtained in the first part of this article in terms of their delta-vega pricing approximation – i.e.,

$$\begin{aligned} \Pi_h &\approx h_s s_T \ln \frac{S_{T+\tau}}{s_T} + h_c(\delta_{BS,T} \ln \frac{S_{T+\tau}}{s_T} + v_{BS,T} \ln \frac{\Sigma_{T+\tau}}{\sigma_T}) \\ &= (h_s s_T + h_c \delta_{BS,T}) \ln \frac{S_{T+\tau}}{s_T} + h_c v_{BS,T} \ln \frac{\Sigma_{T+\tau}}{\sigma_T}. \end{aligned} \qquad (4)$$

Thus, from the joint normal assumption for $\ln \frac{S_{T+\tau}}{s_T}$ and $\ln \frac{\Sigma_{T+\tau}}{\sigma_T}$, and from the fact that sums of jointly normal variables are normal, the total portfolio is normally distributed. Isolating the horizon τ, we obtain

$$\Pi_h \sim \mathrm{N}(\tau\mu_h, \tau\sigma_h^2), \qquad (5)$$

where

$$\mu_h \equiv (h_s s_T + h_c \delta_{BS,T})\,\mu_s + h_c v_{BS,T} \mu_\sigma \qquad (6)$$

$$\begin{aligned} \sigma_h^2 \equiv\ & (h_s s_T + h_c \delta_{BS,T})^2 \sigma_s^2 + h_c^2 v_{BS,T}^2 \sigma_\sigma^2 \\ & +2\,(h_s s_T + h_c \delta_{BS,T})\, h_c v_{BS,T} \rho \sigma_s \sigma_\sigma. \end{aligned} \qquad (7)$$

Previously, 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 (1) becomes easier to interpret in terms of returns and weights, as the total portfolio return is the following weighted average of the single-security returns:

$$R_h = \sum_{n=1}^{N} w_n R_n, \qquad (8)$$

where we dropped the horizon subscript for simplicity.

In "The ℙrayer," 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 both the return $R \equiv \Pi/d$ and the weight (see Meucci, 2010d) are always defined.

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

Third, the linear returns $(p_{T+\tau} - p_T)/\,p_T$, which appear in the aggregation rule (8), are often confused with the compounded returns $\ln(p_{T+\tau}/p_T)$, 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 (see Meucci, 2010b).

ℙ6 Attribution

With the Aggregation Step 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

THE QUANT CLASSROOM by ATTILIO MEUCCI

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 $Z \equiv (Z_1, \dots Z_K)'$ that drive the portfolio P&L and that have a known joint distribution with the portfolio P&L $f_{\Pi_h, Z}$.

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 $b_h \equiv (b_{h,1}, \dots b_{h,K})'$, as follows:

$$\Pi_h = \sum_{k=1}^{K} b_{h,k} Z_k. \tag{9}$$

Note that the attribution to arbitrary factors in general gives rise to a portfolio-specific residual. The formulation (9) covers this case, by setting such a 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

$$f_{\Pi_h, Z} \quad \mapsto \quad b_h \tag{10}$$

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 $K \equiv N$ (the number of securities). Then the attribution equation (9) becomes the conservation law of money (1).

If on the other hand the trader wishes to hedge a given risk — say, volatility risk — then he or she will choose as a factor Z_k the projected P&L Π_k of a truly actionable instrument, such as a variance swap, that 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, then he or she will use changes in implied volatility as one of the risk factors.

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

Notice that (9) 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 that define this linear factor model — along with its connections with the linear factor models introduced in the Estimation Step P2 — is the spirit of the "Factors on Demand" approach in Meucci (2010a).

Illustration. In our stock and option example, we look at a simple attribution (9) 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 (4), we immediately obtain

$$\Pi_h = b_{h,s} Z_s + b_{h,\sigma} Z_\sigma, \tag{11}$$

where the total exposures to Z_s and Z_σ read, respectively,

$$b_{h,s} \equiv h_s s_T + h_c \delta_{BS,T}, \quad b_{h,\sigma} \equiv h_c v_{BS,T}. \tag{12}$$

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.

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 $\tilde{h}$. 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, 2009).

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 S for the projected distribution of the given portfolio P&L Π_h with holdings h

$$f_{\Pi_h} \quad \mapsto \quad \mathcal{S}(h). \tag{13}$$

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

$$f_{\Pi_h, Z}, \; b_h \quad \mapsto \quad \mathcal{S}(h) = \sum_{k=1}^{K} b_{h,k} \mathcal{S}_k, \tag{14}$$

where $b_{h,k}$ represents the "amount" of the factor Z_k in the portfolio-projected P&L and S_k represents the "per-unit" contribution to the statistic $S(h)$ 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 (13) 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 (5), we obtain

$$\mathcal{S}(h_s, h_c) \equiv \mathrm{E}\{\Pi_h\} = \tau\mu_h = \tau h_s s_T \mu_s + \tau h_c (\delta_{BS,T}\mu_s + v_{BS,T}\mu_\sigma). \tag{15}$$

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

$$\mathcal{S}(h_s, h_c) \equiv \mathrm{Sd}\{\Pi_h\} = \sqrt{\tau}\sigma_h, \tag{16}$$

where σ_h is defined in (7).

For the optional summary statistics attribution sub-step (14), a simple linear decomposition that mirrors the attribution equation (9) 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:

$$\mathcal{S}(h) = \sum_{k=1}^{K} b_{h,k} \frac{\partial \mathcal{S}(h)}{\partial b_{h,k}}. \tag{17}$$

Therefore, if the summary statistics are homogeneous, we can take advantage of Euler's identity (17) to perform the summary statistics attribution sub-step (14), which becomes (17).

In particular, for the VaR, the decomposition (17) amounts to the classical definition of marginal contributions to VaR (see, e.g., Garman, 1997), and, for the standard deviation, the decomposition (17) 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 (9) is the bottom-up attribution to the individual securities through the conservation law of money (1). Similarly, the simplest case of attribution of the summary statistics (14) is the attribution of the summary statistics $S(h)$ to the individual securities

$$\mathcal{S}(h) = \sum_{n=1}^{N} h_n \frac{\partial \mathcal{S}(h)}{\partial h_n}. \tag{18}$$

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 (16).

The exposure $b_{h,s}$ of the projected portfolio P&L (9) to the stock factor $Z_s \equiv \ln(S_{T+\tau}/s_T)$ and the exposure $b_{h,\sigma}$ to the implied volatility factor $Z_\sigma \equiv \ln(S_{T+\tau}/\sigma_T)$ were calculated in (12). Then the attribution (17) to each of the two risk drivers of the standard deviation of the projected portfolio P&L becomes

$$\begin{pmatrix} \frac{\partial \, \mathrm{Sd}\{\Pi_h\}}{\partial b_{h,s}} \\ \frac{\partial \, \mathrm{Sd}\{\Pi_h\}}{\partial b_{h,\sigma}} \end{pmatrix} = \frac{\sqrt{\tau}}{\sigma_h} \begin{pmatrix} \sigma_s^2 & \rho\sigma_s\sigma_\sigma \\ \rho\sigma_s\sigma_\sigma & \sigma_\sigma^2 \end{pmatrix} \begin{pmatrix} h_s s_T + h_c \delta_{BS,T} \\ h_c v_{BS,T} \end{pmatrix} \tag{19}$$

where σ_h is defined in (7) (see the proof in the technical appendix available at http://symmys.com/node/63). The total contributions to risk from the factors follow by multiplying the entries on the left-hand side of (19) by the respective exposures.

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

THE QUANT CLASSROOM by ATTILIO MEUCCI

$$\begin{pmatrix} \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial h_s} \\ \frac{\partial \operatorname{Sd}\{\Pi_h\}}{\partial h_\sigma} \end{pmatrix} = \frac{\sqrt{\tau}}{\sigma_h} \begin{pmatrix} s_T^2 \sigma_s^2 & \sigma_{\Pi_s,\Pi_c} \\ \sigma_{\Pi_s,\Pi_c} & \sigma_{\Pi_c}^2 \end{pmatrix} \begin{pmatrix} h_s \\ h_c \end{pmatrix} \tag{20}$$

where

$$\sigma_{\Pi_c}^2 \equiv \sigma_s^2 \delta_{BS,T}^2 + \sigma_\sigma^2 v_{BS,T}^2 + 2\sigma_\sigma \sigma_s \rho \delta_{BS,T} v_{BS,T} \tag{21}$$

$$\sigma_{\Pi_s,\Pi_c} \equiv \delta_{BS,T} s_T \sigma_s^2 + s_T v_{BS,T} \sigma_\sigma \sigma_s \rho, \tag{22}$$

see the proof in the technical appendix available at http://symmys.com/node/63. The total contributions to risk from the stock and the call option follow by multiplying the entries on the left hand side of (20) by the respective holdings h_s and h_c.

The computation of the summary statistics $S(h)$ is hard to perform in practice, unless the market is normal as in our example (16), because complex multiple integrals are involved. For instance, using the same notation as in (3), the VaR with confidence c is defined by

$$1 - c \equiv \int_{h'\pi \leq VaR} f_\Pi (\pi_1, \dots, \pi_N) \, d\pi_1 \cdots d\pi_N. \tag{23}$$

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 (17) of the summary statistics is even harder, unless the market is normal, as in our example (19)-(20). Fortunately, these computations become simple when the market distribution is represented in terms of scenarios (see Meucci, 2010a).

Before concluding, we must address two key problems of risk and portfolio management: estimation risk, introduced in the Estimation Step ℙ 2, and liquidity risk, introduced in the Pricing Step 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 tests 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, 2010c), 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 Prayer 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.

ℙ 8 Optimization

In the Evaluation Step ℙ 7, the risk manager or portfolio manager obtains a set of computed summary statistics 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 $\mathcal{S}(h) \equiv \operatorname{E}\{\Pi_h\} - \gamma \operatorname{Sd}\{\Pi_h\}$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 $\mathcal{S}(h) \equiv \operatorname{E}\{u(\Pi_h)\}$.

More generally, we call "index of satisfaction" the function that translates the P&L distribution of the portfolio with holdings $h \equiv (h_1, \dots h_N)'$ into a personal preference ranking. We denote the index of satisfaction by the general notation $S(h)$ used in (13) for the evaluation summary statistics, because any index of satisfaction is also a summary statistic.

Given an index of satisfaction $S(h)$, 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 C the set of all such constraints and by "$h \in C$" 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, as follows:

$$h^* \equiv \underset{h \in \mathcal{C}}{\operatorname{argmax}} \{\mathcal{S}(h)\}. \tag{24}$$

We emphasize that the choice of the most suitable index of satisfaction S, as well as the specific constraints 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 (24) becomes

$$(h_s, h_c)^* \equiv \underset{h_c \equiv 1}{\operatorname{argmax}} \{-\operatorname{Sd}\{\Pi_h\}\}. \tag{25}$$

Then the first order condition on the P&L standard deviation, computed in (5)-(7), yields

$$h_s \equiv -\frac{\delta_{BS,T}}{s_T} - \frac{v_{BS,T}}{s_T} \rho \frac{\sigma_\sigma}{\sigma_s}. \tag{26}$$

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 (24) 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, as follows:

$$h_\lambda \equiv \underset{h \in \mathcal{C}}{\operatorname{argmax}} \{\operatorname{E}\{\Pi_h\} - \lambda \operatorname{Vr}\{\Pi_h\}\}, \quad \lambda \in \mathbb{R}. \tag{27}$$

This step reduces the dimension of the problem from N, and the dimension of the market, to 1, the value of λ. The optimization (27) 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&Ls sparse (see Meucci, 2010f).

In the second step of the mean-variance heuristic, the optimal portfolio is selected by a one-dimensional search

$$h^* \equiv \underset{\lambda \in \mathbb{R}}{\operatorname{argmax}} \{\mathcal{S}(h_\lambda)\}. \tag{28}$$

The optimization (28) can be performed by a simple grid-search.

As it was the case for the Evaluation Step ℙ7, we must address estimation risk, introduced in the Estimation Step ℙ2: 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 ℙ8.

To address the issue of estimation risk, portfolio managers rely on more advanced approaches than the simple two-step mean-variance heuristic (27)-(28). 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, and Browne and Kosowski, 2010. Refer to Meucci, 2010e, for a review and code).

Finally, liquidity risk, discussed in the Pricing Step ℙ4, 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 (27)-(28).

ℙ9 Execution

The Optimization Step ℙ8 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^m_t and internal book information i^b_t to attain the target portfolio h^* by a sequence of transactions at given prices $p_t \equiv (p_{t,1 \ldots,} p_{t,N})'$

$$h^* \,,\ \{i^m_t\}_{t \geq T} \,,\ \{i^b_t\}_{t \geq T} \quad \mapsto \quad \{p_t\}_{t \geq T} \,. \tag{29}$$

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 ℙ8 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 ℙ4. 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 ℙ9 should be embedded into the Optimization Step ℙ9..."* 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 ℙ9, 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, as follows:

$$\pi_{h^*} \quad \mapsto \quad (\pi_a, \pi_b, \cdots) \,. \tag{30}$$

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 ℙ8 and slippage performance from the Execution Step ℙ9.

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 ℙ6.

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 (11), and the implied volatility component reads $b_{h,s} \ln(\sigma_{T+\tau}/s_T)$. The residual is the plugin 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.

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

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.

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. Ibid, 2005b. "Robust Bayesian asset allocation," Working Paper, available at http://ssrn.com/abstract=681553.

Ibid, 2008. "Fully flexible views: Theory and practice," Risk 21, 97—102. Extended version available at http://ssrn.com/abstract=1213325.

Ibid, 2009. "Managing diversification," *Risk* 22, 74—79. Extended version available at http://ssrn.com/abstract=1358533.

Ibid, 2010a. "Factors on Demand - building a platform for portfolio managers risk managers and traders," *Risk* 23, 84—89. Available at http://ssrn.com/abstract=1565134.

Ibid, 2010b. "Linear vs. compounded returns - common pitfalls in portfolio management," *Risk Professional*, "The Quant Classroom by Attilio Meucci," April, 52—54. Available at http://ssrn.com/abstract=1586656.

Ibid, 2010c. "Personalized risk management: Historical scenarios with fully flexible probabilities," *Risk Professional*, "The Quant Classroom by Attilio Meucci," December, 40—43. Available at http://ssrn.com/abstract=1696802.

This is an abbreviated list. To read all of the references, please go to http://symmys.com/node/63.

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

THE QUANT CLASSROOM by ATTILIO MEUCCI