

Managing diversification

Attilio Meucci introduces a diversification index that represents the effective number of bets in a portfolio. With this index, based on entropy and constrained principal component analysis, he performs mean-diversification management adjusted for transaction costs using a parsimonious set of optimal trades

The qualitative definition of diversification is very clear to portfolio managers: a portfolio is well diversified if it is not heavily exposed to individual shocks. However, oddly enough, there exists no broadly accepted, unique, satisfactory methodology to precisely quantify and manage diversification.

In the special case of systematic-plus-idiosyncratic factor models, diversification is measured as the percentage of risk explained by the systematic factors. However, 'idiosyncratic' shocks are actually correlated with each other. Furthermore, such a measure fails to analyse the degree of diversification within the pseudo-idiosyncratic component of the portfolio, but this becomes necessary in market-neutral and relative-value strategies such as equity pairs trading, where the systematic risk is hedged away.

In the special case of long-only portfolios, differential diversification is defined as the difference between the weighted sum of the volatilities of each position and the total portfolio volatility. This measure does not cover residual portfolios in such strategies as fixed-income immunisation, where the portfolio manager is interested in the diversification net of the parallel shifts of the curve, which are hedged away.

More diversification measures have been introduced, including measures based on portfolio weights, which do not account for correlations and volatilities (refer to appendix A1 online for a detailed list and a discussion). However, in addition to not applying in full generality, none of those measures highlights where diversification, or the lack thereof, arises in a given portfolio.

The contributions of this article are fourfold. First, instead of focusing on one single number, we introduce the diversification distribution, a tool to analyse the fine structure of a portfolio's concentration profile in fully general markets and with fully general long-short positions. Second, we introduce the effective number of uncorrelated bets, an actionable index of diversification based on the entropy of the diversification distribution. Third, we introduce the mean-diversification frontier, a quantitative framework to manage the trade-off between the expected return and the effective number of uncorrelated bets. Fourth, we provide an efficient algorithm to implement in practice mean-

diversification optimisation in the presence of impact and transaction costs, by trading only a few securities.

To achieve the above, we first express a generic portfolio in terms of the exposures to uncorrelated sources of risk. The most natural choice for such sources are the principal components, but these might not be manageable in constrained portfolios. Therefore, we generalise the principal components to the conditional principal portfolios, whose volatilities fully describe the concentration profile of the trader's positions. Then we interpret these volatilities as a set of probability masses, the diversification distribution. The dispersion of this distribution, as measured by its entropy, becomes a summary index of diversification with a very natural interpretation, namely the effective number of uncorrelated bets in the portfolio. Finally, with this index we build the mean-diversification efficient frontier, which quantifies precisely the trade-off between expected returns, transaction and impact costs, and diversification.

The article is arranged as follows. First, we motivate and construct the diversification distribution, which is the foundation of diversification analysis. Then we introduce the entropy-based diversification index and use it to tackle diversification management. We illustrate our approach by means of practical examples throughout. Several technical results are thoroughly discussed in the appendix of the extended version of this article, available online at www.symmys.com > Research > Working Papers. A fully documented implementation of our approach can be downloaded from www.symmys.com > Teaching > MATLAB.

Diversification analysis

Consider a generic market of N securities. We denote the returns of these securities over the given investment horizon by the N -dimensional vector \mathbf{R} . We represent a generic portfolio in this market by the vector \mathbf{w} of the weights of each security. The return $R_{\mathbf{w}}$ on the portfolio reads:

$$R_{\mathbf{w}} \equiv \mathbf{w}'\mathbf{R} \quad (1)$$

■ **Total risk analysis.** Since the ground-breaking work of Markowitz (1952), measuring risk in terms of variance has become standard practice in the financial industry. In uncorrelated markets, the individual securities in a generic portfolio (1) constitute additive sources of risk:

$$\text{Var}\{R_{\mathbf{w}}\} \equiv \sum_{n=1}^N \text{Var}\{w_n R_n\} \quad (2)$$

Therefore, in uncorrelated markets maximum diversification corresponds to equal variance-adjusted weights.

In correlated markets, this is not the case. For instance, a duration-weighted portfolio that is long each node of a government curve is fully concentrated on the curve shifts (see Litterman & Scheinkman, 1991, and appendix A4 online for details). However, although the securities in the market are correlated, it is always possible to determine sources of risk that are uncorrelated, and therefore additive. The most natural choice of uncorrelated

risk sources is provided by the principal component decomposition of the returns covariance Σ :

$$\mathbf{E}'\Sigma\mathbf{E} \equiv \Lambda \quad (3)$$

In this expression, the diagonal matrix $\Lambda \equiv \text{diag}(\lambda_1^2, \dots, \lambda_N^2)$ contains the eigenvalues of Σ , sorted in decreasing order, and the columns of the matrix $\mathbf{E} \equiv (\mathbf{e}_1, \dots, \mathbf{e}_N)$ are the respective eigenvectors. The eigenvectors define a set of N uncorrelated portfolios, the principal portfolios, whose returns $\tilde{\mathbf{R}} \equiv \mathbf{E}^{-1}\mathbf{R}$ are decreasingly responsible for the randomness in the market. Indeed, the eigenvalues Λ correspond to the variances of these uncorrelated portfolios.

In figure 1, we display the geometrical interpretation of the decomposition (3): the eigenvectors $\mathbf{e}_1, \dots, \mathbf{e}_N$, that is, the principal portfolios, are the directions of the principal axes of an ellipsoid, and the square root of the eigenvalues $\lambda_1, \dots, \lambda_N$, that is, the volatilities of the principal portfolios, are the length of these axes (see Meucci, 2005, for further details). We emphasise that the decomposition (3), which is at the heart of the analysis to follow, holds for any market with a well-defined covariance, and not necessarily for normal markets only. For instance, the market in figure 1 is lognormal as in Black & Scholes (1973), but the same interpretation holds also for non-parametric distributions, or for t -distributed markets, possibly skewed, where no two portfolios are ever independent, etc.

A generic portfolio can be seen either as a combination of the original securities with weights \mathbf{w} as in (1), or as a combination of the uncorrelated principal portfolios with weights $\tilde{\mathbf{w}} \equiv \mathbf{E}^{-1}\mathbf{w}$ (see also Partovi & Caputo, 2004). In terms of the latter, we can introduce the variance concentration curve:

$$v_n \equiv \tilde{w}_n^2 \lambda_n^2, \quad n = 1, \dots, N \quad (4)$$

The generic entry v_n of this concentration curve represents the variance due to the n th principal portfolio: much like (2), since the principal portfolios are uncorrelated, the total portfolio variance is the sum of these terms:

$$\text{Var}\{R_{\mathbf{w}}\} \equiv \sum_{n=1}^N v_n \quad (5)$$

Equivalently, we can consider the volatility concentration curve:

$$s_n \equiv \frac{\tilde{w}_n^2 \lambda_n^2}{\text{Std}\{R_{\mathbf{w}}\}}, \quad n = 1, \dots, N \quad (6)$$

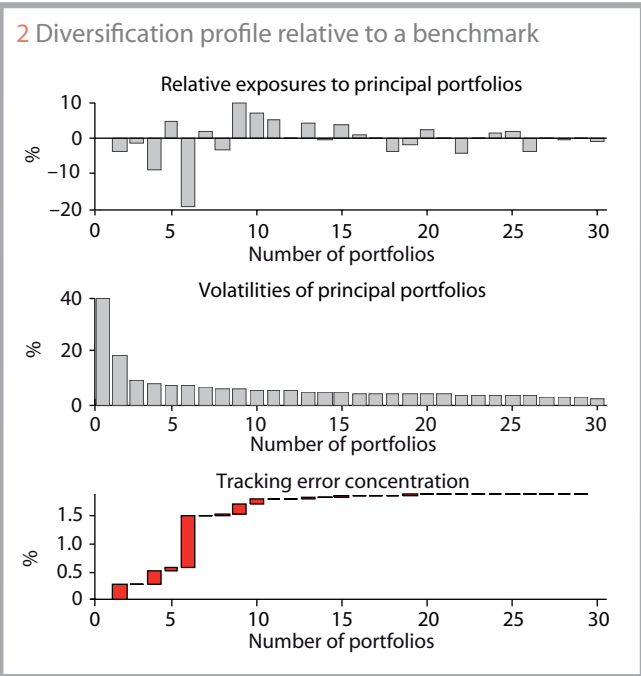
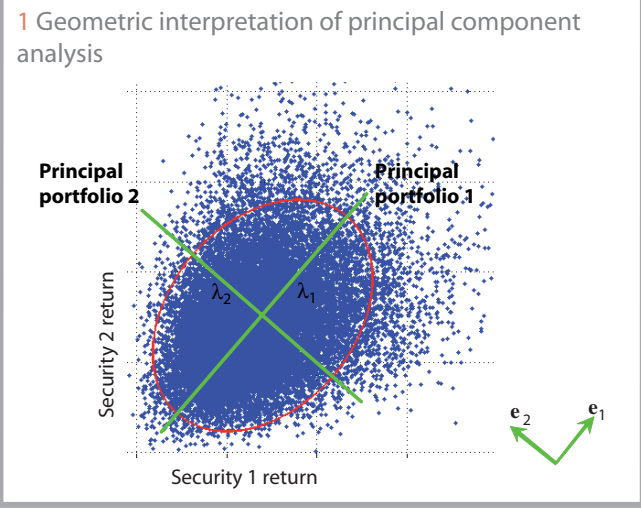
The volatility concentration curve does not simply represent a normalised decomposition of the variance concentration. As we prove in appendix A2 online, (6) represents the decomposition of volatility or tracking error into the contributions from each principal portfolio, as in Litterman (1996).

Finally, we can normalise the variance and the volatility concentration curves further into the diversification distribution:

$$p_n \equiv \frac{\tilde{w}_n^2 \lambda_n^2}{\text{Var}\{R_{\mathbf{w}}\}}, \quad n = 1, \dots, N \quad (7)$$

This is not simply a percentage version of the variance and the volatility concentration curves. As we show in appendix A3 online, the generic term p_n equals the r -square from a regression of the total portfolio return on the n th principal portfolio.

This approach generalises the methodology in Meucci (2007) to correlated markets. We emphasise that the above analysis also applies to management against a benchmark with weights \mathbf{b} .



Indeed, it suffices to replace the portfolio weights with the vector of the relative bets:

$$\mathbf{w} \mapsto \mathbf{w} - \mathbf{b} \quad (8)$$

Then (6) becomes the tracking error concentration curve and (7) becomes the relative diversification distribution.

To illustrate, we consider a simplified market of $N \equiv 30$ liquid mid-cap stocks in the Russell 3000 index. We assume that the manager tracks a benchmark in the same stocks with weights proportional to their market capitalisation. We estimate the covariance matrix of the returns Σ by exponential smoothing daily observations. Then we analyse the relative diversification of the equally weighted portfolio $\mathbf{w} \equiv \mathbf{1}/N$, which generates a tracking error of 1.89%. The top plot in figure 2 reports the relative exposures to the principal portfolios $\tilde{\mathbf{w}} \equiv \mathbf{E}^{-1}(\mathbf{w} - \mathbf{b})$; the middle plot reports the volatilities $\lambda_1, \dots, \lambda_N$ of the principal portfolios; and the bottom plot displays the tracking error concentration curve. Notice how the exposure to the

(2003) and Gouriéroux, Laurent & Scaillet (2000) (see appendix A6 online).

To analyse diversification in this context, instead of the standard principal portfolios (3), we decompose risk on to N conditional principal portfolios. As with the standard principal portfolios, these are uncorrelated portfolios that are decreasingly responsible for the randomness in the market. However, the conditional principal portfolios are adapted to the rebalancing constraints. More precisely, first we define recursively for $n = K + 1, \dots, N$ the portfolios that span the feasible reallocations:

$$e_n \equiv \arg \max_{e' e = 1} \{e' \Sigma e\} \quad (11)$$

such that $\begin{cases} \mathbf{A} e \equiv \mathbf{0} \\ e' \Sigma e_j \equiv 0 \text{ for all existing } e_j \end{cases}$

(see also figure 3). If \mathbf{A} is empty, this process generates the standard principal portfolios (3). For general matrices \mathbf{A} , the solutions e_{K+1}, \dots, e_N still represent combinations of securities that are mutually uncorrelated and decreasingly responsible for the randomness in the market. However, the conditional principal portfolios are adapted to the desired directions (10), in that they can be freely added to the current allocation. Indeed, we can immediately verify that a generic reallocation $\Delta \mathbf{w} \equiv \alpha_{K+1} e_{K+1} + \dots + \alpha_N e_N$ satisfies (10) for arbitrary coefficients α . Next, we complement the above set of feasible reallocations to span the whole market. Hence, we define recursively for $n = 1, \dots, K$ the following uncorrelated portfolios, which are decreasingly responsible for the randomness in the inaccessible directions:

$$e_n \equiv \arg \max_{e' e = 1} \{e' \Sigma e\} \quad (12)$$

such that $e' \Sigma e_j \equiv 0$, for all existing e_j

We refer the reader to appendix A7 online for the calculation of (11)–(12), which only involves iterations of standard principal component analysis.

By collecting the conditional principal portfolios in a matrix $\mathbf{E} \equiv (e_1, \dots, e_N)$, we can express the returns covariance in the same format as (3):

$$\mathbf{E}' \Sigma \mathbf{E} \equiv \mathbf{\Lambda} \quad (13)$$

In this expression, $\mathbf{\Lambda}$ is a diagonal matrix whose non-zero entries represent the variances of the conditional principal portfolios:

$$\lambda_n^2 \equiv \text{Var}\{e'_n \mathbf{R}\}, \quad n = 1, \dots, N \quad (14)$$

In figure 4, we display the typical pattern of the generalised spectrum (14).


Proceeding as in the unconditional case, we can represent a generic allocation \mathbf{w} equivalently as a linear combination of the conditional principal portfolios with weights:

$$\tilde{\mathbf{w}} \equiv \mathbf{E}^{-1} \mathbf{w} \quad (15)$$

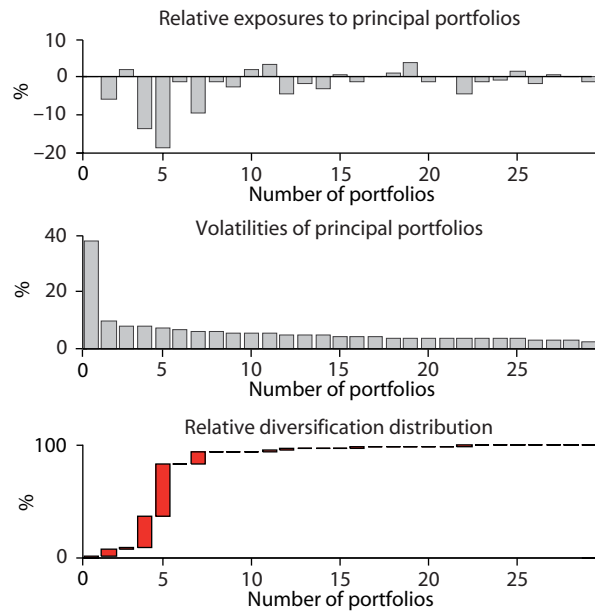
The total portfolio variance is the sum of the contributions from the conditional principal portfolios, since these are uncorrelated:

$$\text{Var}\{R_w\} = \sum_{n=1}^N \tilde{w}_n^2 \lambda_n^2 \quad (16)$$

4 Generalised PCA spectrum: variances of conditional principal portfolios



5 Diversification relative to a benchmark under budget constraint

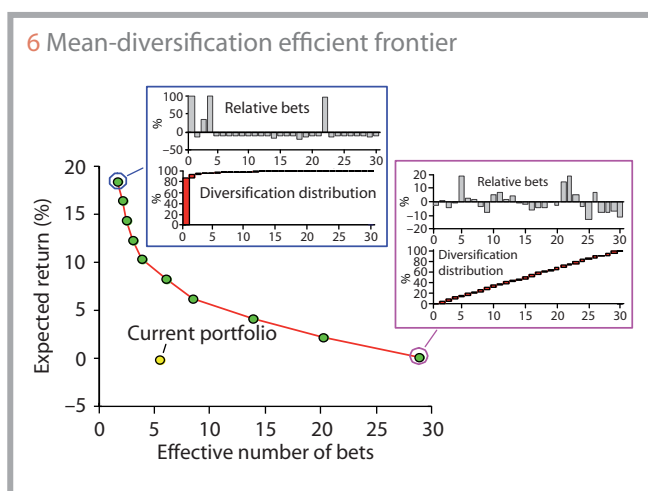


The conditioning equation (10) takes on a very simple form in the new co-ordinates (15). Indeed, the first K entries of $\tilde{\mathbf{w}}$ cannot be altered. Therefore, instead of (6) we focus on the conditional diversification distribution:

$$p_n | \mathbf{A} \equiv \frac{\tilde{w}_n^2 \lambda_n^2}{\sum_{m=K+1}^N \tilde{w}_m^2 \lambda_m^2}, \quad n = K + 1, \dots, N \quad (17)$$

where the notation emphasises that the conditioning is enforced by the matrix \mathbf{A} in (10).

It is easy to check that if the rows of \mathbf{A} represent some unconditional principal portfolios then conditional and unconditional principal portfolios coincide. This is the case in the single-sector stock market with the budget constraint, where indeed $\mathbf{A} \equiv \mathbf{1}'$ and, as shown in appendix A5 online, $\mathbf{1}$ is a principal portfolio.



Then the conditional analysis $p_n|\mathbf{1}'$ in (17) is equivalent to the simple adjustment to the unconditional analysis (9).

The conditional diversification distribution provides a clear picture of the diversification structure of a portfolio when only specific rebalancing directions are allowed. By controlling the terms of this distribution, a manager can efficiently manage diversification, hedging or leveraging the conditional principal portfolios based on his views on the market.

As in the unconditional case, the above conditional analysis also applies to management against a benchmark with weights \mathbf{b} : by replacing as in (8) the portfolio weights with the vector of the relative bets, (17) yields the relative conditional diversification distribution.

To illustrate, we consider our example of 30 liquid mid-cap stocks in the Russell 3000 index managed against a 30-stock benchmark. First, we calculate the relative diversification profile as it follows from conditioning the large 3000×3000 global matrix to allow for trades only in those 30 liquid stocks. As expected, the result is the same as the analysis in figure 2.

Then we calculate and display in figure 5 the relative diversification profile that follows from further conditioning the analysis on the budget constraint. The load is now more evenly spread over two principal portfolios. Depending on their views on the market, the portfolio manager can proceed to either hedge these risks or to increase these exposures by means of a parsimonious set of trades, as discussed further below.

Diversification management

Both the simple unconditional diversification distribution (7) and the conditional diversification distribution (17), whether absolute or relative to a benchmark, are always defined, are always positive and always add up to one. Therefore the diversification distribution can be interpreted as a set of probability masses associated with the principal portfolios.

The shape of the ensuing probability density provides a clear picture of the level of diversification of a given allocation. For a well-diversified portfolio, the probability masses p_n are approximately equal and thus the diversification distribution is close to uniform. On the other hand, if the portfolio is concentrated in a specific conditional principal direction, the diversification distribution displays a sharp peak.

Therefore, portfolio diversification can be represented, as in Meucci (2007), by the dispersion of the diversification distribution,

as summarised by its entropy, or equivalently by its exponential:

$$\mathcal{N}_{Ent} \equiv \exp\left(-\sum_{n=K+1}^N p_n \ln p_n\right) \quad (18)$$

where the unconditional case corresponds to $K = 0$.

The interpretation of \mathcal{N}_{Ent} is very intuitive. Indeed, it is easy to verify that $\mathcal{N}_{Ent} = 1$ in fully concentrated positions, that is, when all the risk is completely due to one single principal portfolio. On the other hand, \mathcal{N}_{Ent} achieves its maximum value $N - K$ in portfolios whose risk is homogeneously spread among the $N - K$ available principal portfolios. In other words, \mathcal{N}_{Ent} represents the true number of uncorrelated bets in a fully general, potentially long-short, portfolio in a fully general market.

We remark that entropy has already been used to measure diversification (see, for example, Bera & Park, 2004, discussed in appendix A1 online). However, such a definition acts on the portfolio weights, which are not always positive, do not necessarily add up to one, and do not account for volatilities and correlations.

To illustrate, we consider our example of an equally weighted portfolio of $N \equiv 30$ liquid mid-cap stocks benchmarked to a market-capitalisation weighted index of the same stocks.

We focus on the unconditional diversification number, that is, $K \equiv 0$ in (18). Since the portfolio is managed against a benchmark, we consider the relative diversification analysis, which corresponds to replacing in the construction of the index (18) the relative bets whenever the portfolio weights appear, as specified by (8). The result is $\mathcal{N}_{Ent} \approx 5.5$ effective relative bets, which is consistent with the bottom plot in figure 2 (see also figure 6).

To optimise diversification, portfolio managers can calculate the mean-diversification efficient frontier:

$$\mathbf{w}_\phi \equiv \arg \max_{\mathbf{w} \in \mathcal{C}} \{ \phi \boldsymbol{\mu}' \mathbf{w} + (1 - \phi) \mathcal{N}_{Ent}(\mathbf{w}) \} \quad (19)$$

where $\boldsymbol{\mu}$ denotes the estimated expected returns and \mathcal{C} is a set of investment constraints. The parameter ϕ spans the interval $[0, 1]$. For small values of ϕ , diversification is the main concern, whereas as ϕ approaches one, the focus shifts on the expected returns.

For example, we consider our market of $N \equiv 30$ liquid stocks where as above \mathcal{N}_{Ent} in (19) is the unconditional number of relative bets. We set the expected returns $\boldsymbol{\mu}$ using a risk-premium argument $\boldsymbol{\mu} \equiv 0.5\boldsymbol{\sigma}$, where $\boldsymbol{\sigma}$ are the square roots of the diagonal elements of $\boldsymbol{\Sigma}$. We assume that stock weights are bounded between -10% and 100% .

As we see in figure 6, for small values of ϕ , long-short positions alternate to increase diversification. The relative diversification distribution is highly uniform and the number of effective uncorrelated relative bets \mathcal{N}_{Ent} approaches 30, that is, the number of securities in this market. On the other hand, for large values of ϕ , the relative diversification distribution concentrates on very few effective bets with higher expected returns. Notice how the interplay of positions and correlations heavily concentrates risk, although the relative bets appear to be rather well distributed. Also notice that the current equally weighted allocation is heavily suboptimal.

In practice, the expected returns should be adjusted for transaction and impact cost:

$$\boldsymbol{\mu}'\mathbf{w} \mapsto \boldsymbol{\mu}'\mathbf{w} - \mathcal{T}(\mathbf{w}, \mathbf{w}_{cur}) \quad (20)$$

where \mathcal{T} is an empirically fitted function of the current allocation \mathbf{w}_{cur} and the target portfolio \mathbf{w} (see, for example, Torre & Ferrari, 1999, and Almgren *et al.*, 2005). Since \mathcal{T} is a discontinuous function of the weights, and thus is not even convex, the diversification frontier cannot be calculated exactly, because the optimisation (19)–(20) can be solved numerically only in trivially small markets. Furthermore, it is operationally impractical to monitor the progress of a large number of trades.

Instead, it is better to pursue a parsimonious reallocation that gives rise to the largest possible increase in the mean-diversification trade-off with only a small number of trades. The exact combinatorial search for the best combination of trades is unfeasible, but it can be replaced by a recursive-acceptance selection algorithm, which converges in a matter of seconds. Intuitively, this routine adds the best ‘team player’ one at a time from among all the securities available in the market. Each time a player is tested, the optimisation (19)–(20) runs on a very small number of variables and thus can be solved. We discuss the details of this algorithm in appendix A8 online.

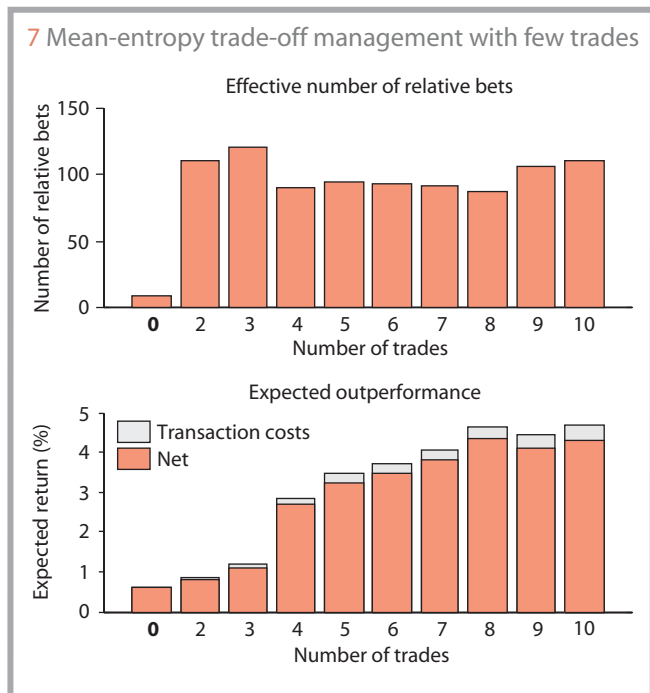
To illustrate, we consider a mutual fund benchmarked to the full Russell 3000 index. Again, we consider \mathcal{N}_{Ent} in (19) to be the unconditional number of relative bets and we set $\phi \approx 0.5$. We assume a transaction cost model:

$$\mathcal{T}(\mathbf{w}; \mathbf{w}_{cur}) \equiv \sum_{n=1}^N \mathcal{T}_n(w_n - w_{n,cur}) \quad (21)$$

where \mathbf{w}_{cur} denotes the current equally weighted portfolio and:

$$\mathcal{T}_n(v) \equiv \begin{cases} 0 & \text{if } v = 0 \\ \alpha_n & \text{if } 0 < |v| \leq \gamma_n \\ \beta_n |v|^\zeta & \text{if } |v| > \gamma_n \end{cases} \quad (22)$$

In this expression, $\zeta \approx 1.5$ and the security-specific coefficients α_n , β_n and γ_n are fitted empirically. Refer to the above references for the rationale behind this specification. Notice the discontinuity in the origin, which makes (19)–(20) intractable with standard optimisation techniques.



Instead, we run our optimisation routine and we display the results in figure 7. The current equally weighed portfolio gives rise to only $\mathcal{N}_{Ent} \approx 9$ effective relative bets and less than a 1% expected return over the benchmark. Given the budget constraint, the minimum number of trades allowed is two. With only two trades, the number of effective relative bets jumps to $\mathcal{N}_{Ent} \approx 110$. Three trades marginally increase both diversification and expected outperformance. With four trades, a dramatic improvement of expected outperformance is achieved at the expense of a little diversification. Adding trades beyond this point is not advisable. ■

Attilio Meucci is head of research at Bloomberg ALPHA, Portfolio Analytics and Risk, in New York. He is grateful to Gianluca Fusai, Sridhar Gollamudi, Harvey Stein and two anonymous referees for their helpful feedback. Email: a.meucci@bloomberg.net

References

- | | | | |
|---|--|--|---|
| <p>Almgren R, C Thum, E Hauptmann and H Li, 2005
<i>Equity market impact</i>
Risk July, pages 57–62</p> <p>Bera A and S Park, 2004
<i>Optimal portfolio diversification using maximum entropy</i>
Working paper</p> <p>Black F and M Scholes, 1973
<i>The pricing of options and corporate liabilities</i>
Journal of Political Economy 81, pages 637–654</p> <p>Gourieroux C, J-P Laurent and O Scaillet, 2000
<i>Sensitivity analysis of values at risk</i>
Journal of Empirical Finance 7, pages 225–245</p> | <p>Hallerbach W, 2003
<i>Decomposing portfolio value-at-risk: a general analysis</i>
Journal of Risk 5, pages 1–18</p> <p>Hannah L and J Kay, 1977
<i>Concentration in modern industry. Theory, measurement and the UK experience</i>
Macmillan</p> <p>Litterman R, 1996
<i>Hot spots and hedges</i>
Goldman Sachs and Company, Risk Management Series</p> <p>Litterman R and J Scheinkman, 1991
<i>Common factors affecting bond returns</i>
Journal of Fixed Income 1, pages 54–61</p> | <p>Markowitz H, 1952
<i>Portfolio selection</i>
Journal of Finance 7, pages 77–91</p> <p>Meucci A, 2005
<i>Risk and asset allocation</i>
Springer</p> <p>Meucci A, 2007
<i>Idiosyncratic portfolio risk in non-normal markets</i>
Lehman Publications</p> <p>Partovi M and M Caputo, 2004
<i>Principal portfolios: recasting the efficient frontier</i>
Economics Bulletin 7, pages 1–10</p> | <p>Potters M, J-P Bouchaud and L Laloux, 2005
<i>Financial applications of random matrix theory: old laces and new pieces</i>
arXiv:physics/0507111v1</p> <p>Tasche D, 2006
<i>Measuring sectoral diversification in an asymptotic multifactor framework</i>
Journal of Credit Risk 2, pages 33–55</p> <p>Torre N and M Ferrari, 1999
<i>Barra – the market impact model</i>
Barra Research Insights</p> |
|---|--|--|---|

The online appendices referred to in this article can be found at www.symmys.com > Research > Working Papers