OSSIAM RESEARCH TEAM

Impact of Covariance Cleaning Techniques on Minimum Variance Portfolios

March, 20, 2012

WHITE PAPER

Impact of Covariance Cleaning Techniques on Minimum Variance Portfolios

Pedro Jandrey-Natal and Ksenya Rulik

March, 20, 2012

Pedro Jandrey-Natal
Research assistant
pedro-fernando.jandrey-natal@polytechnique.org

Ksenya Rulik, PhD, CFA
Head of Quantitative research
ksenya.rulik@ossiam.com

Abstract

Minimum variance portfolio is constructed using an optimization procedure that takes as an input a covariance matrix. Inevitable estimation errors in the covariance lead to distortions of optimal allocation and should be addressed in the optimization. One solution is to use advanced estimation techniques that aim at reducing the error. We analyze here the impact of statistical covariance cleaning on the minimum variance portfolio. We build several minimum variance portfolios using different covariance estimators: two shrinkage estimators and two estimators based on Random Matrix theory, and compare volatility of these minimum variance portfolios with a portfolio constructed with a simple empirical covariance. Our results show that in presence of basic portfolio constraints, such as a long-only constraint, the effect of covariance cleaning on ex-post volatility is negligible. This result is in accordance with earlier research of Jagannathan, Ma [2003] and Roncalli [2011] who showed that portfolio constraints in their turn induce important transformations on covariance matrix. It turns out that an explicit cleaning of covariance matrix is dominated by the implicit cleaning induced by portfolio constraints. We also discuss a duality between the two-norm quadratic portfolio constraint and the James-Stein covariance shrinkage estimator and find and approximate relation between the degree of covariance shrinkage and the target for the two-norm constraint.

Minimum variance investing, an approach that significantly reduces equity portfolio risk, has recently gained particular attention. On the one hand, the uncertain market environment poses questions on whether the equity investments can be accessed with less risk. On the other hand, over the past decades evidence was gathered of long-term outperformance of low volatility stocks with respect to high volatility stocks ([1, 6, 7]) that undermines the widespread belief that higher risk is necessarily rewarded with higher returns.

Minimum variance portfolio provides in this regard an answer to both questions. It is a way to invest in equities with significantly less risk (30% of volatility reduction is the number that is often cited by diferent sources [1]). Moreover, by construction, the minimum variance portfolio has important exposure to low-volatility stocks that historically were shown to perform much better than predicted by the Capital Asset Pricing Model.

Estimation of covariance matrix of stock returns is at the heart of the minimum variance approach. The information contained in single-stock volatilities and pairwise correlations is translated into the portfolio positions via a technique minimizing expected portfolio variance. The resulting allocation gives the minimal possible ex-ante volatility of the overall portfolio under constraints: being long-only, fully invested, and satisfying additional constraints the investor might want to impose.

The quality of the covariance matrix estimation is thus an important challenge here. The most straightforward way to accomplish this task is to construct an empirical covariance matrix from the data on past returns. However, it is well known that the simple historical estimation contains errors coming from the finiteness of historical samples used for estimation, as well as from the changing nature of the market volatility. A number of advanced tools were proposed by financial econometrics to improve estimation efficiency of covariance matrix. In this paper we discuss the effect of some advanced statistical tools of covariance cleaning: shrinkage techniques and Random Matrix theory cleaning, and we assess whether these methods can improve the minimum variance portfolio constuction.

In the first section we introduce four different covariance matrix estimation techniques and their basic properties. In the second section, we first define the minimum variance optimization setup and then build a set of minimum variance portfolios using advanced covariance estimation methods introduced in the first section. To assess the impact of the covariance cleaning, we then compare realized volatility of different minimum variance portfolios corresponding to the different covariance cleaning techniques to that of the benchmark portfolio built with simple empirical covariance matrix. Covariance estimation errors always increase the ex-post portfolio variance, so a cleaning technique has an added value if it allows to reduce the ex-post portfolio variance, giving an indication that the estimation error after application of the technique is lower. In the third section we review the connection between the covariance matrix transformations and the use of additional constraints in portfolio optimization process. We consider in particular the equivalence between shrinkage techniques and 2-norm quadratic constraints, and find the implicit degree of covariance shrinkage that is implied by imposing quadratic equality constraints for a minimum variance portfolio. We conclude with the discussion of the results.

[1] For example, the MSCI Minimum Volatility World index backtested back to 1995 had around 30% reduction in volatility with respect to the market-capitalization weighted index over the period 1995-2007, as documented by MSCI BARRA Research (Nielsen F., Aylursubramanian R., *Far From the Madding Crowd - Volatility Efficient Indices*, April 2008).

Covariance Matrix Estimation and Cleaning

Consider a universe of N stocks, with observed past returns over T days. We denote r_i^t the arithmetic return of the i-th stock at date t[2]. The *empirical covariance matrix* of stock returns ($\mathbf{S}$) is defined as follows:

$$\mathbf{S}_{i,j} = \frac{1}{T-1}\sum_{t=1}^{T}(r_i^t - \bar{r}_i)(r_j^t - \bar{r}_j), \quad (1)$$

where $\bar{r}_i = \frac{1}{T}\sum_{t=1}^{T} r_i^t$.

When the number of observations is not very large compared to the number of stocks (i.e. when T is comparable to N), this estimator is known to perform poorly (see e.g[10] or [2]).

Academic literature on efficient covariance estimation and forecasting techniques is vast. A widespread approach is based on factor models, re-constructing covariance matrix for stocks from a smaller covariance matrix among risk factors and the corresponding factor loadings. Another family of methods is purely statistical and aims at reducing the noise part of the estimated matrix. Finally, covariance forecasting techniques are very popular, that use the stylized facts such as volatility clustering to model the evolution of covariance structure in time (GARCH models for example). In this paper we focus on the statistical techniques of efficient covariance matrix estimation: *shrinkage* and *Random-Matrix-Theory cleaning.*

Shrinkage Techniques

Although the estimator introduced in the Eq.1 is unbiased, it may be very imprecise. This happens because the error of an estimator can be explained not only by its bias, but also by its variance, and it is the latter that can be very large for the empirical estimator. The aim of the so-called *shrinkage* techniques is to introduce an estimator that minimises estimation error as a whole, by making a combination (shrinking) of the original estimator with a biased, though more efficient (with smaller variance) estimator.

The simplest example of a shrinkage estimator is the *James-Stein shrinkage estimator* (**JSSE**), defined as a combination of the empirical covariance and an identity matrix:

$$\mathbf{JSSE} = \left[(1-\hat{\alpha}^*)\mathbf{S} + \hat{\alpha}^*\hat{\nu}^*\mathbf{Id}\right], \quad (2)$$

where $\mathbf{Id}$ is identity matrix, and $\hat{\alpha}^*$ and $\hat{\nu}^*$ are consistent estimators of the constants resulting from the following minimisation program:

$$\min_{\alpha,\nu} \mathbb{E}\left[\|(1-\alpha)\mathbf{S} + \alpha\nu\mathbf{Id} - \Sigma_{\mathbf{true}}\|^2\right], \quad (3)$$

where $\Sigma_{\mathbf{true}}$ is the true covariance matrix of the underlying statistical process and $\|A\|$ equals Trace(A^2) for a symmetric matrix A. For the exact expression of these parameters, see [11].

The **JSSE** estimator uses an identity matrix as its shrinkage target. This target corresponds to a covariance matrix where all pairwise correlations are zero and variances of all stocks are equal, that is quite unrealistic. But the great advantage of the indentity matrix is that it has no variance at all, leading to a significant reduction of the estimator variance. The parameter α is called the degree of shrinkage. When α is zero, the shrinkage is absent and the original empirical estimator is restored. When the parameter $\alpha = 1$ the weight is completely transfered onto the identity matrix, with no information from the empirical covariance left.

More realistic shrinkage targets are possible. One way to get a non-trivial target matrix is to consider the one-factor model introduced by Sharpe (1962), in which stock returns are generated by:

$$r = \alpha + \beta r_0 + \varepsilon, \quad (4)$$

[2]The arithmetic returns are defined as: $r_t = P_t/P_{t-1} - 1$, where P_t is stock price on the date t

where α is a constant intercept, r_0 represents market return vector, β is market sensitivity (stock beta) and the ε is a white noise uncorrelated with the market. In this model, the covariance matrix of stock returns $\mathbf{B}$ can be easily computed[3]. The *Ledoit-Wolf shrinkage estimator* (**LWSE**) consists in shrinking the empirical covariance matrix towards a consistent estimator of $\mathbf{B}$:

$$\mathbf{LWSE} = \left[(1-\hat{\alpha}^*)\mathbf{S} + \hat{\alpha}^*\hat{\mathbf{B}}\right], \qquad (5)$$

where $\hat{\alpha}^*$ is a consistent estimator of the constant resulting from the following minimisation program:

$$\min_{\alpha} \mathbb{E}\left[\|(1-\alpha)\mathbf{S} + \alpha\hat{\mathbf{B}} - \Sigma_{\mathbf{true}}\|^2\right]. \qquad (6)$$

For more information on the exact expression of these parameters, see [10].

Covariance Cleaning with Random Matrix Theory

Another technique used in covariance matrix estimation is the *noise filtering with Random Matrix Theory* (RMT). Although the mathematical tools necessary to achieve a complete understanding of this approach are quite technical, its general idea is rather simple. Think about a no-structure covariance matrix: all correlations are zero and all variances are equal to 1. If a sample of random variables having this covariance structure is generated along with estimations of the covariance matrix on the finite sample, the resulting empirical covariance matrix will not be a pure identity, as the true covariance is. There will be non-zero diagonal terms, corresponding to false negative or positive correlations and the diagonal elements - variance estimations - will indicate that some series are more volatile than others, that is not true. All this false structure comes from estimation noise, and the magnitude of the estimation noise can be quantified based on the sample size and the number of observations. The empirical covariance matrix can be compared to a matrix generated randomly and factor out the noise from the statistically significant information. The specific amount of structure imposed on the random matrix is computed according to a certain prior belief we have for the true underlying statistical process. In other words, the greater the agreement between empirical data and model predictions is, the more structure we impose. Obviously, this approach will yield good results only if the model is accurate enough.

The simplest estimator of this family, *Marĉenko-Pastur estimator* (**MP**), performs the comparison of empirical covariance matrix to the simplest no-structure noisy covariance. More precisely, the cleaned covariance matrix is obtained by cutting off the eigenvalues of the empirical covariance that may be considered as meaningless noise. Intuitively, with this method, the structure of the true covariance matrix is very similar to the identity matrix, apart from few most significant eigenvalues that correspond to major common risk factors. The method is thoroughly explained in [3].

An evolution of the previous method, the *Power-Law estimator* (**PL**), is obtained by replacing all the eigenvalues of the empirical covariance matrix by a set of new, synthetic eigenvalues. These eigenvalues are distributed according to a Power Law distribution (also known as Pareto distribution). For a complete discussion on RMT cleaning recipes, see [4].

The methods of Random Matrix Theory are applied to correlation matrix rather than to the full covariance matrix, as the tools designed in this field allow to quantify the noise in covariance estimated on normalized random vari-

[3] the exact expression of this estimator is $\mathbf{B}_{i,j} = \sigma_m^2\beta_i\beta_j + \Sigma_\epsilon$, where σ_m^2 is variance of market returns and Σ_ϵ is covariance of the white noise components ϵ_i

ables. To construct a cleaned covariance matrix it is necessary first to decompose the empirical covariance matrix in a product of a variance vector and correlation matrix, then clean the correlation matrix with an Random Matrix Theory technique, and finally reconstruct the full covariance matrix using the original variance vector and the cleaned correlation matrix.

Minimum variance setup

Minimum variance allocation is a solution of the following optimization problem:

$$\min_{w \in \mathcal{C}} \mathbf{w}'\mathbf{\Sigma}\mathbf{w}, \tag{7}$$

where $\mathbf{w}$ denotes a vector of portfolio weights and $\mathcal{C}$ represents a set of constraints. We impose here a budget constraint, i.e. $\sum_i w_i = 1$, and a no-short-sale constraint, i.e. $w_i \geq 0, i = 1, 2, ..., N$.

Note that the optimization scheme takes Σ, the covariance matrix of stock returns, as an input. Different covariance matrices would yield different optimal allocations, each minimizing expected (ex-ante) portfolio variance for the given covariance among the stocks. There is no reference to expected stock returns in this setup, that amounts to a silent assumption of equal expected returns for all stocks in the selected group.

In order to test the estimation techniques for the covariance matrix presented above, we construct for every different covariance estimation method an associated minimum variance portfolio. We also build a benchmark minimum variance portfolio using the empirical covariance matrix (no cleaning applied). Then we analyse the *ex-post* volatility of these minimum variance portfolios over a time period subsequent to that of the covariance estimation period. As was discussed above, if one estimator happens to be more efficient than the others, we expect a smaller ex-post variance of the minimum variance portfolio associated to this estimator.

We use a dataset of $N = 350$ stocks that were present in the STOXX Europe 600 index between November 2006 and April 2011[4]. The period considered (from January 2005 to April 2011) consists of 1680 days, that we subdivide in overlapping periods of $T = 500$ days to run separate covariance estimations. Each 500-day period is shifted by 100 days with respect to the period preceeding it. Then we take 100 days after each covariance estimation period to compute realized volatility of the minimum variance portfolios. This test is equivalent to backtesting a set dynamic minimum variance strategies, one for each covariance estimation technique. The strategies are rebalanced every 100 days, with new covariance matrices estimated over 500-days window preceeding the rebalancing date. Finally, we compute the volatilities of these strategies on 100-day rolling window.

Fig.1 shows the results of this exercise. For comparison, we plot on the graph the 100-day roling volatility of the STOXX Europe 600 index over the same period. The comparison to the STOXX 600 shows that all the minimum variance strategies achieved significant volatility reduction with respect to the market capitalization weighted index. Still, it can be noted that all covariance estimation techniques yield very similar ex-post volatilities and the more sophisticated techniques employing covariance cleaning produced no significant change with respect to the empirical covariance matrix.

The failure of the covariance cleaning tech-

[4]The data is a courtesy of STOXX. The STOXX indices are the intellectual property (including registered trademarks) of STOXX Limited, Zurich, Switzerland and/or its licensors ("Licensors"), which is used under license. None of the products based on those Indices are sponsored, endorsed, sold or promoted by STOXX and its Licensors and neither of the Licensors shall have any liability with respect thereto.

niques to bring additional ex-post risk reduction for minimum variance portfolio is surprising. We know that all these cleaning recipes depend on a parameter determining the degree of cleaning that needs to be estimated. By cleaning parameter we mean here the shrinking constant α (see Eq.2 and Eq.5) in the shrinkage case and the proportion of eigenvalues that are cut off from the empirical covariance matrix in the Random Matrix cleaning case. What if the cause of the apparent inefficiency of cleaning is precisely the misestimation of the cleaning parameter? To answer this question, we analyse the changes in realized volatility that are produced by varying the cleaning parameter value.

Fig.2 shows the results of this analysis. The vertical axis represents the ratio of ex-post volatility of a minimum variance portfolio corresponding to a given cleaning technique to volatility of a minimum variance portfolio corresponding to the empirical covariance matrix. If the ratio is smaller than 1, it means that the cleaning technique in question yields lower ex-post volatility of the minimum variance portfolio than the empirical covariance matrix. The horizontal axis represents the value of cleaning parameter used in covariance estimation with cleaning. Thus each line on the plot represents volatility reduction potential of one cleaning technique, applied with different values of the cleaning parameter. The volatility ratios are estimated as averages over the 100-day non-overlapping subperiods defined above.

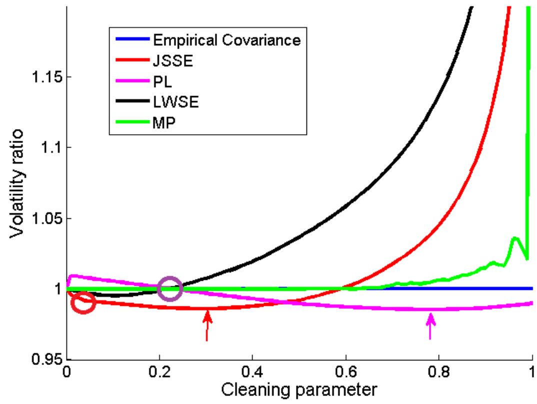

Figure 2: Ratio of ex-post volatility of minimum variance portfolios with cleaning to the volatility of the portfolio without cleaning, as a function of cleaning parameter.

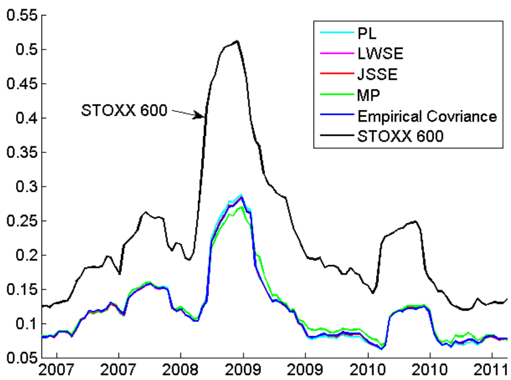

Figure 1: Out-of-sample annualized volatility of minimum variance portfolios. Note that all estimation techniques yield very similar results.

We denote by empty circles the theoretical cleaning parameter values that are automatically determined by the respective cleaning schemes. On the left extreme of the graph, when the cleaning parameter α is equal to zero, all the covariance matrices coincide with the empirical covariance matrix. At the opposite extreme, when the parameter goes towards 1, the cleaning techniques approach their targets (in the case of shrinkage methods) or noise priors (in the case of Random Matrix cleaning). Three of the four techniques (**JSSE**, **LWSE**, **MP**) yield a volatility ratio that can be greater than 1 and grows rapidly when approaching the right edge of the graph, indicating that the resulting allocation is not of the minimum variance type anymore and important information from the covariance matrix has been cleaned away. The fourth method (**PL**) has more consistent behaviour, and the respective volatility ratio stays smaller than 1 over almost all the range of the values of cleaning parameter.

The maximal volatility reduction with respect to the empirical covariance matrix is achieved by the **PL** and **JSSE** estimators. The magnitude of the maximal reduction is of 1.45% for **PL** with cleaning factor $\alpha = 0.78$, and of 1.39% for the **JSSE** estimator with cleaning factor $\alpha = 0.31$. The two remaining cleaning methods, **LWSE** and **MP**, bring virtually no volatility reduction across the whole range of cleaning parameter.

This analysis shows that even changing the value of cleaning parameter, the covariance cleaning methods still do not lead to substantial *ex-post* volatility improvement of minimum variance portfolio. For comparison, the minimum variance portfolio based on the empirical covariance matrix has the volatility reduced by roughly 30% with respect to the market capitalization weighted portfolio. The additional reduction of 1.45% is hardly a significant result. The cause of this limited noise correction potential of cleaning techniques in the minimum variance framework is thus not in the estimator construction itself.

Role of Portfolio Constraints

In this section we argue that the surprising result of the previous section, namely that there is no additional ex-post volatility reduction in minimum variance portfolio when using covariance cleaning techniques, can be explained by the presence of constraints in the portfolio optimization problem.

First we review some recent results that demonstrate that imposing additional constraints in the optimization problem described by Eq.7 is equivalent to solving an unconstrained problem with a modified covariance matrix. Further we discuss a direct correspondence between covariance shrinkage (of the JSSE type) and optimization with a quadratic (norm) constraint.

Duality between constraints and covariance transformations

Equivalence between imposing constraints and transforming the input covariance matrix was first pointed out in [8], where the authors used as example the non-negativity constraint. A recent overview of the impact of different types of constraints on covariance matrix in the optimization framework can be found in [12].

First consider a problem of finding a global minimum variance portfolio subject only to a budget constraint(i.e. $\sum_i w_i = 1$). Mathematically speaking, we look for a vector of weights w that satisfies

$$\min_{w \in \mathbb{R}^N, \mathbf{1}'w=1} w'\Sigma w, \tag{8}$$

where $\mathbf{1}$ is a vector of ones. Using Lagrange multipliers, it can be easily demonstrated that the solution $w^*(\Sigma)$ reads[5] :

[5] In this section, we use the matrix notation to stress

$$w^*(\Sigma) = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}'\Sigma^{-1}\mathbf{1}}. \qquad (9)$$

The budget constraint is an essential ingredient of the minimum variance optimization. Omitting this constraint results in an empty solution, corresponding to a zero variance absolute minimum. Therefore we will call the problem given by the Eq.8 and the solution Eq.9 *unconstrained*. This solution in general represents a long-short portfolio, with 100% *net* long exposure to stocks.

Now, what happens to the solution when we impose additional constraints? For the sake of simplicity, we consider only lower bounds on weights as additional constraints. A particular case of the lower bound constraint is a non-negativity, or no-short-sale constraint where the lower bound is equal to zero. Let's denote by $\widetilde{w}(\Sigma)$ the solution of the minimization problem in Eq. 7, where

$$\mathcal{C} = \{w \in \mathbb{R}^N : \mathbf{1}'w = 1 \text{ and } \forall i, w_i \geq w_i^+\}. \qquad (10)$$

It was demonstrated that for a *constrained* problem of the type Eq.10 there exists an associated *unconstrained* problem of the type Eq.8 with a transformed covariance matrix $\widetilde{\Sigma}$, such that $w^*(\widetilde{\Sigma}) = \widetilde{w}(\Sigma)$, i.e. the solution of Eq.10 using Σ is the solution of Eq.8 using $\widetilde{\Sigma}$.

It is possible to prove that such matrix $\widetilde{\Sigma}$ exists, and moreover it can be related to the original Σ matrix. Following [12], $\widetilde{\Sigma}$ can be decomposed as:

$$\widetilde{\Sigma} = \Sigma + \Delta. \qquad (11)$$

The matrix Δ is called perturbation matrix, and its elements depend both on the initial covariance Σ and the vector of lower bounds. No explicit expression of the Δ matrix has been given in the literature, but we believe that this implicit transformation of covariance matrix is significant and strongly alters the initial matrix. Actually, the transformation is so profound that it converts a long-short allocation into a long-only allocation, when added to the minimum variance optimizer. The cleaning techniques applied to the covariance matrix have much weaker effect on the background of this implicit constraints-implied transformation, and thus do not lead to significant improvements in the ex-post volatility of the optimized portfolio.

Below we give an example of a portfolio constraint that produces a transformation of covariance matrix that can be given explicitly, in the form of a covariance shrinkage already discussed above.

The norm constraint

Following the same lines, De Miguel et al. in [5] showed that imposing quadratic equality constraint to an optimization scheme is equivalent to modifying the input matrix Σ in a very specific way. The quadratic constraint reads:

$$\sum_i w_i^2 = \frac{1}{H}, \qquad (12)$$

where H is a constant. Using Lagrange multipliers, it can be easily verified that imposing the constraint Eq.12 amounts to replacing the original covariance matrix by its shrinkage estimate that has as a shrinkage target an identity matrix (in other words, the **JSSE** estimator):

$$\widetilde{\Sigma} = \Sigma + \alpha \mathrm{Id}. \qquad (13)$$

Thus the quadratic constraint has the same effect on the minimum variance allocation as the application of the **JSSE** shrinkage estimator.

Now, for a fixed constant H, *what degree* of implicit shrinkage are we imposing via this

the fact that the solution of the minimisation problem (8) depends on the input matrix.

constraint? In other words, given the duality between the covariance shrinkage and the quadratic constraint, what value of the shrinkage parameter α is equivalent to a given value of the parameter H? To answer this question we run two optimization algorithms, one is a long-only constrained problem given by Eq.7 and Eq.10 with a covariance matrix shrunk with the **JSSE** estimator, and a long-only constrained problem with an additional quadratic constraint 12 and a simple empirical covariance matrix (no shrinkage applied). Comparing distances among the allocations resulting from the two algorithms with varying values of the shrinkage parameter α and the quadratic constraint right-hand-side H, we build an empirical correspondence among the different α and H values. This exercise was repeated on overlapping 500-day periods of our sample. We have been using the same set of stocks and the same period as in the previous section. The results are shown in the Fig.3.

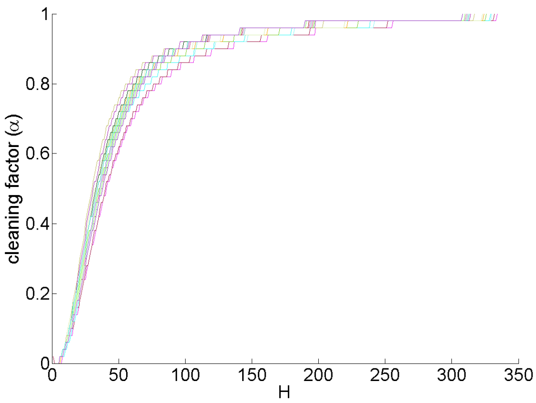

Figure 3: Relation between the constant H and the equivalent shrinkage coefficient α. Note that there are different curves for different samples.

The graph shows the correspondence curves between the cleaning parameter values (α) and the constant in the right-hand side of quadratic portfolio constraint (H). Different curves are obtained by using different 500-days subsamples of our observation periods. The curves are grouped quite close to each other, indicating that the correspondence is sufficiently stable, but depends on the covariance matrix details that change from one subperiod to another.

Note that even modest values of the H-parameter imply a significant degree of shrinkage ($alpha$ rapidly approaching 1). To give the reader a better understanding about the above result, we recall that the value of the H-parameter can be seen as a measure of portfolio diversification. For example, for an equi-weighted portfolio of N assets ($w_i = 1/N$), the inverse of the sum of square weights (inverse of the Eq. 12) by construction yields $H = N$. Talking about a non-equal weighted portfolio, it could be said that a (generic) portfolio with the inverse sum of square weights equal to H is as diversified as an equiweighted portfolio of H assets. This means that the values of H that are much smaller than the overall number of stocks in the eligible universe can be associated with quite concentrated portfolios.

With this interpretation in mind, by varying the value of H (or, quivalently, the shrinkage parameter α) the allocation starts from a very concentrated no-shrinkage one ($H = 5$, $\alpha = 0$), and goes towards an equi-weighted portfolio (limit of $H = 350$, $\alpha = 1$). This means that the initial portfolio (no covariance shrinkage, or no quadratic constraint applied) is concentrated in just a few stocks and significant amounts of shrinkage (or the quadratic norm constraint on the weights with higher values of H) needs to be imposed to make the allocation more diversified. Note that the combined results of Fig.2 and Fig.3 indicate that an optimal value for H would be around 30, that gives an implicit shrinkage corresponding to $\alpha = 0.3$.

Conclusion

Given the difficulties arising in mean-variance optimization because of estimation errors, particular attention should be paid to the design of the optimization problem. As the recent research shows, there might be little scope in simultaneous application of both advanced noise cleaning of covariance matrix and stringent portfolio constraints. Indeed, the constraints that bound the optimized allocation may be seen as implicit transformations of covariance matrix, altering the latter in very profound ways.

We show here that in minimum variance optimizaton with long-only constraints the use of statistical covariance cleaning techniques, such as skrinkage and cleaning with Random Matrix Theory, have very marginal effect on the properties of the resulting portfolios. Only two of the four considered techniques, namely the Power Law cleaning and James-Stein covariance shrinkage, brought slight additional reduction in the ex-post variance with respect to the use of empirical covariance matrix. The other two methods, Ledoit-Wolf covariance shrinkage and Marcenko-Pastur eigenvalue clipping, had virtually no effect.

The very same cleaning methods are known to be particularly efficient in the absence of non-negativity constraints. Ledoit and Wolf [10] showed that **JSSE** estimator can reduce the volatility of minimum variance long-short portfolios by up to 17% compared to the empirical covariance matrix. Bouchaud and Potters in [4] showed that **PL** noise cleaning technique reduces realized volatility of minimum variance Long-Short portfolios by up to 23%.

Why do the cleaning recipes yield so different results in the presence of non-negativity constraints? Intuitively, unconstrained minimum variance portfolios (i.e., portfolios whose weights are allowed to be negative) tend to allocate substantial weight to both long and short positions, thus neutralizing the most important risk factors and leaving the exposure to the weakest noise components of the covariance matrix. But this is exactly the part of the covariance matrix that tends to be more misestimated. That is why noise filtering is extremely important when dealing with long-short unconstrained portfolios.

On the other hand, errors in covariance estimation do not have significant impact on *constrained* minimum variance portfolios, where, for example, negative weights are not allowed, or where particular diversification targets are given. This comes from the fact that portfolio constraints act as powerful implicit transformations of the covariance matrix, thus dwarfing the effect of the additional cleaning applied. For example, imposing a non-negativity constraint on the weights of the minimum variance portfolio is *strictly* equivalent to solving an unconstrained optimization problem with a transformed covariance matrix as an input. In other words, by simply imposing non-negativity constraints, we are implicitly implementing a cleaning procedure.

One particular case when explicit dual covariance transformation can be established, is the case of quadratic constraint in portfolio optimization. This constraint is very useful because it allows to target a specific level of portfolio diversification, expressed by the value of the constant in the right-hand-side of the constraint. This type of constraint is strictly equivalent to the well-known covariance shrinkage, and the shrinkage degree is directly related to the diversification target set by the quadratic constraint.

References

[1] Ang A., Hodrick R., Xing Y. and Zhang X., *The Cross-Section of Volatility and*

Expected Returns, The Journal of Finance, Vol. 61, No. 1, 2006, pp. 259-299.

[2] Bouchaud J.P. and Potters M., 2003, *Theory of Financial Risk and Derivative Pricing*, Cambridge University Press.

[3] Bouchaud J.P. and Potters M., 2005, *Financial Applications of Random Matrix Theory: Old Laces and New Pieces*, Acta Physica Polonica B 36 2767-2784.

[4] Bouchaud J.P. and Potters M., 2009, *Financial Applications of Random Matrix Theory: a short review*, arXiv-qfinn 0910.1205.

[5] DeMiguel, V., Garlappi L., Nogales J., Uppal R., 2009,*A generalized approach to portfolio optimization: improving performance by constraining portfolio norms.* Management Science, 55, pp. 798-812.

[6] Haugen R.. and Baker N., *The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios*, Journal of Portfolio Management, Vol. 17, No. 3, 1991, pp. 35-40

[7] Haugen R. and Baker N., *Case Closed*, Handbook of Portfolio Construction: Contemporary Applications of Markowitz techniques, Springer, 2010.

[8] Jagannathan J. and Ma T. (2003), *Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps*, Journal of Finance, 58(4), pp. 1651-1683.

[9] DeMiguel V., Garlappi L. , Nogales F.J., and Uppal R., *A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms*, Management Science, vol.55, pp.798-812, 2009.

[10] Ledoit, O. and Wolf, M., 2003, *Improved estimation of the covariance matrix of stock returns with an application to portfolio selection.* J. Empirical Finance10 603-621.

[11] Ledoit, O. and Wolf, M., 2004, *A well-conditioned estimator for large-dimensional covariance matrices.* J. Multivariate Anal. 88 365-411.

[12] T. Roncalli, *Understanding the Impact of Weights Constraints in Portfolio Theory*, working paper at SSRN: http://ssrn.com/abstract=1761625, 2011.

About Ossiam

Ossiam is a research-driven French asset management firm (authorized by the Autorité des Marchés Financiers) and specializes in delivering smart beta* solutions. Efficient indexing is at the core of Ossiam's business model. The firm was founded in response to a post-subprime crisis demand from investors for simplicity, liquidity and transparency. Given the environment, there was a growing need among investors for enhanced beta exposure and risk hedging. Ossiam is focused on the development of innovative investment solutions for investors via a new generation of indices.

**'Smart beta' refers to systematically managed, non-market-cap-weighted strategies covering any asset class.*

ossiam

80, avenue de la Grande Armée

75017 Paris – France

info@ossiam.com

+ 33 1 78 40 56 90