Predicting future returns with investor views

A comparison of
Black-Litterman and Copula Opinion Pooling techniques

Supervisor: Henrik Hult

Fang Li, $\mathfrak{F}$06, <fangl@kth.se>, 870707-3030

June 13, 2011

Abstract

The mean-variance framework developed by *Markowitz (1959)* for portfolio analysis is probably one of the most elegant and widely accepted concepts in financial theory. A major problem in portfolio optimization is the estimation of the input parameters in the distributions of the underlying assets. *Jobson and Korkie (1980)* conclude that classical estimation techniques, where unknown parameters as means and covariances are replaced by their sample values perform very poorly. Small variations in sample means often lead to extreme portfolio allocation readjustments. To compliment historical data with new information that helps better predict future market outcome *Black and Litterman (1992)* found a way in which the investors could combine their market views with historical information. *Meucci (2008)* extended the idea into *Copula-opinion pooling* which is the numerical equivalence to the *Black-Litterman* framework. Also an effort has been made by *Stein (1956)* on the frontier of statistics to construct a mean estimator better than the observed average of the data sample, the shrinkage estimator. In this thesis a thorough comparison between*Black-Litterman*, *Copula-opinion pooling*, and *Black-Litterman with shrinkage estimators* will be made to find out their respective strengths and shortcomings. Certain simplifications and assumptions are made to make comparisons more convenient. New methods are suggested to increase stability of the models and cope with some of the problems encountered. The result indicates that *Black Litterman* combined with the shrinkage estimator does give consistently better results compared to the other methods.

Acknowledgements

I would like to thank my supervisor at KTH Mathematical Statistics, Henrik Hult, for his great feedback and guidance as well as my supervisors at Swedbank, Jonas Bostrom and Johan Kilander for their knowledge in the field of portfolio optimization and allocation, guidance, patience and ideas.

Stockholm, June 2011

Fang, Li

Contents

Introduction

Black and Litterman (1992) proposed a mean of estimating expected returns to achieve better-behaving portfolio models. The model sets the idealized market equilibrium as a point of reference, allows the investor to specify a chosen number of market views in the form of absolute or relative expected returns and a level of confidence for each view. The views are combined with the market equilibrium returns to give more (less) weight on assets where investors have positive (negative) opinions on.

With the increasing popularity of *Monte-Carlo* simulations normality assumptions in the original *Black Litterman* become more and more questionable. In an effort to make *Black Litterman* independent on distribution assumptions the *Copula opinion pooling* is developed by *Meucci* in 2008 where empirical distribution is used. The choice of the assumption of distributions is handed over to the user of the model. The user can either choose to fit to a certain distribution and simulate using parametric bootstrap or simply use historical observations as it is in which case fewer data points are available but no assumptions are made.

After Stein first derived in *Stein (1955)* the inadmissibility of the classical Maximum-Likelihood mean estimators and introduced the shrinkage estimator which will always dominate the ML-estimators in the sense of mean-squared error, the shrinkage estimator was initially met by frequent debate and resistance for a long period of time due to the fact that it contradicted the common sense in statistics that observed average of the data is the best estimator of the expected mean. The shrinkage estimator was later improved and developed into the *James-Stein Estimators* where the historical means are adjusted towards the grand mean of the entire sample, and the *Bayes-Stein Estimators* where the historical means are adjusted towards the expected return of the Minumum-Variance portfolio under the Markowitz framework.

Purpose

The purpose of this thesis is to find a stable model that integrates investor views into historical market information and produces more reliable future market returns. It is also in our interests to understand under which circumstances the model do produce better results compared to averaging the historical sample and the robustness of the models with respect to different input parameters. To do this a through understanding of the models is required and more focus is directed to the theoretical backgrounds of the models used in the thesis.

The analysis of the models is carried out through series of tests. All tests are individually designed to test some aspects of the model among which the mean-error, the volatility predicting power are two key components to be tested. Test scores measure performances of the models and results are analyzed and connected back to the theories behind.

The thesis only focuses on inputs and outputs of models that combines investors views with historical samples and does not evaluate the effect of portfolio optimization itself. The models used in the thesis will produce input parameters into portfolio optimization models. Performance of the portfolio optimization models on the other hand often depend on the portfolio constraints used. This dependency sometimes leads to portfolio allocations so extreme that the qualities of the original input parameters become obscure.

Outline

Chapter I, II and *III* covers the basics of the models and sketches of derivations are provided along with some discussions of how to set certain parameters. *Chapter IV* outlines the setup of the test models and gives a discussion of solutions of several problems encountered during the implementation.

Chapter V gives a thorough walk-through of the results. While several reflections and recommendations of parameter calibrations are made in the thesis, *Chapter VI* makes an attempt to construct an alternative way of determining one of the parameters used in *Black-Litterman*, the uncertainty matrix Ω.

1 The Black-Litterman model

While the standard portfolio optimization model can only take expectations and covariances of the assets intended to be included in the portfolio, *Black and Litterman (1992)* found a new way in which the investors could combine their market opinions with the expectations and covariances of the assets estimated within the Markowitz framework. The Black-Litterman framework requires a larger set of inputs: the view portfolios, the expected return on the assets in the portfolio, the confidence level of the view portfolios and the uncertainty on the reference model.

This chapter begins with an introduction of the original Black-Litterman model with its assumptions. A list of drawbacks of the Black-Litterman framework will be presented. Modifications inspired by *Stein (1955)*, *Satchell and Scowcroft (2000)*, *Idzorek (2004)*, *Meucci (2005)* and *Meucci (2010)* will be briefly discussed. The purpose of this chapter is to provide the basic background to the BL framework while highlighting useful modifications to the original framework. Some of these modifications will be implemented in Chapter IV.

1.1 The Market

Let's consider a market with N securities. The log-returns of the assets are modeled as an N-dimensional multivariate normally distributed variable $\mathbf{X} \sim (X_1, X_2, X_3, ...)$.

$$\mathbf{X} \sim N(\mu, \Sigma)$$

where μ is an N-dimensional vector and represents the expected outcome of the returns. The covariance Σ can be estimated from the observed past outcome of the returns. However when it comes to specifying μ the original BL takes into consideration the estimation error and models μ as

$$\mu \sim N(\pi, \tau\Sigma)$$

Normality is a very important assumption in the BL-framework as it allows the model to give explicit formulas for expectation and covariances of the log-return distribution, blended with expert views. The π refers to the equilibrium state of all underlying assets in the portfolio. Litterman himself argues in *Litterman (2003)* that this is a concept where supply meets demand for the underlying assets. However this idealized state never appears on real market and its returns are not readily observable.

Many BL practitioners approximates the equilibrium state π by solving the mean-variance portfolio optimization problem to obtain the market implied returns from the readily observable market allocations of the assets, say the weight vector ω^m of the assets in an index of reasonable choice. The mean-variance optimization maximizes the numerical difference between the return and the variance of the portfolio. The allocation that has the best balance between return and risk is chosen to be the optimal portfolio allocation, provided that satisfactory estimates of the mean and covariances are inserted into the model.

$$\begin{aligned} \omega^m &\equiv \mathbf{argmax} \quad \omega^m \pi - \delta \omega^m \Sigma \omega^m \\ \pi &= 2\delta \Sigma \omega^m \end{aligned} \tag{1.1}$$

Here δ is the risk-aversion factor. According to *He and Litterman (1999)* it is explicitly set to 2.5 to represent the average risk appetite of investors worldwide. The w^m is the market allocation that is again, not observable. Litterman suggested that the equilibrium portfolio could be approximated from a benchmark index, for instance the OMX30 index. The expected return derived this way is to be interpreted as the expected return by the market if all actors on the market consistently act to maximize return while minimizing risk.

Mean-variance optimization is not the only way to obtain the equilibrium estimates and the reference point for the BL-framework. Assuming that the market is in its equilibrium state in the long term and that the equilibrium is only distorted by short-term noises, we can average the historical daily returns to obtain the market equilibrium returns. In fact as *Mazzoni (2010)* pointed out in his studies, the historical sample turned out to be one of the best predictors in mean-squared error sense for predicting short term stock asset returns, compared to the GARCH(1,1) model among others.

In later sections we will introduce the shrinkage estimator for the purpose of estimating π while using the traditional historical average estimator as benchmark.

There is another parameter that caused some confusion among BL-practitioners, τ, the weight-on-views. *Black and Litterman (1992)* proposed that the constant should be set close to zero because the uncertainty in the mean is much smaller than the uncertainty in the return itself. *Satchell and Scowcroft (2000)* proposed a ingenious model where τ is stochastic, but the model also came with new parameters to calibrated. In the GS fixed income research by *Bevan and Winkelmann(1998)* they proposed a value of τ between 0.5 and 0.7.

τ is the coefficient relating the variance of the expected market returns to the estimate of the mean of the expected market returns. All estimation methods give rise to estimation errors. The standard historical average estimator $\hat{\mu}$ gives rise to variances that are proportional to the sample size.

$$\hat{\mu} \sim N(\bar{\mu}, \frac{1}{T}\Sigma)$$

It is thus sensible to set the τ to $\frac{1}{T}$ in this case where T is the sample size. In general I argue it should be set accordingly with respect to the estimation method used, the formula for τ can be derived by expressing the estimated covariances as a function of the prior distribution's covariances. An example of this will be provided in connection with introduction of the shrinkage estimation.

Meucci (2010) suggested a modification where the original BL-framework is modified to be based directly on market returns assuming no estimation error of μ, which consequently erased the parameter τ from the modified formula for the BL-expected returns and BL-covariances. The formula for this modification will be presented in the end of this section.

1.2 The views

The magic of the BL-framework lies in the fact that it does not only take the market equilibrium return and covariances and tries to compute the optimal portfolio right away. BL is not about portfolio optimization, it is about blending the investors view with the market implied asset returns. A view is a statement by the investors on how future market will perform. The view can be based on technical analysis, macro-economic analysis and so on. It can be expressed in both absolute and relative terms. The views do not necessarily have to coincide with the market. For instance, the investor might think that the second asset in the portfolio will under-perform the first asset in which case the equivalent view becomes $X_2 - X_1 > 0$. Within the frame of the thesis, the views are chosen to be expressed in absolute terms, in the form of absolute target returns. For instance we will have the view *asset 1 will have a daily target return of x percent.* Each view is coupled with a confidence level $c \in (0, \infty)$ which can be interpreted as a coefficient inversely proportional to the variances of the view distributions. With confident views c should be small with less confident views we expected the variances of the view distributions to grow, as it does for c.

The BL assumes the views to be normally distributed and only focuses on linear views. K views are represented by a $K \times N$ matrix, with its k :th row describing the k :th view. For the following views

- *asset 1 will have a target return of* q_1 *percent*
- *asset 2 will have a target return of* q_2 *percent*
- *asset 3 will have a target return of* q_3 *percent*
- ...

We have the following relationship, with pick matrix P and view vector q.

$$P = \begin{pmatrix} 1 & 0 & \cdots \\ 0 & 1 & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \quad q = \begin{pmatrix} q_1 \\ q_2 \\ \vdots \end{pmatrix}$$

with $\mu \sim N(\pi, \tau\Sigma)$ and $q = P\mu$ the distribution of $q \sim N(P\mu, P\Sigma P')$, which leaves no room for the uncertainty in the views. However the BL fixes this by introducing the conditional distribution of q given μ.

$$q \sim N(P\mu, \Omega)$$

where the uncertainty matrix Ω depends on the confidence coefficient c. We can now focus on how to systematically determine $c \in (0, \infty)$ and Ω.

Meucci (2010) and Idzorek (2004) suggested two systematic ways of determining Ω. Meucci's idea was probably inspired by how the original BL treated the parameter τ, where the author simply sets

$$\Omega = \frac{1}{c^2} P\Sigma P' \tag{1.2}$$

which yields a non-diagonal matrix proportional to original market covariance matrix. The Ω is now related to $P\Sigma P'$, which can be explained by the fact that if an asset in its equilibrium state has very high volatility, it will be harder for investors to predict its expected future return as well what investors means by *certain views*, in this case a large c, is that they are certain of the expected future returns given the historical performance of the underlying assets. Unless the investors are 100 percent sure, in which case Ω should be set to zero, it makes sense to expressed Ω in relative terms.

However to be able to express confidence on each of the views separately, assuming that the distribution of q_i has no correlation with q_j for $i \neq j$, (1.2) can be modified to

$$\Omega = \frac{1}{c^2} P_k \Sigma P_k'$$

where P_k is the k :th row in P.

Idozrek provided a step-by-step method to generate Ω. In this approach the certainty coefficients are not explicit expressed on the views themselves but rather expressed as a measure of how far the investors are prepared to go towards the extreme mean-variance portfolio allocation where the views are 100 percent certain. Idozrek first computed the portfolio weights from BL using mean-variance portfolio optimization. Comparing the result to the equilibrium market weights Idozrek deducted some portion of the difference between these two weights using a confidence argument. The less confidence the investors are views the less willingly investors would want to go towards the full-confidence weights. Finally the results is inserted into BL backwards to obtain Ω. This approach requires a choice of the portfolio optimization model to calculate the portfolio allocation. The diagonal elements in Ω become dependent on the choice of the portfolio optimization method.

1.3 Combining the views with market equilibrium

The problem of combining the view with market equilibrium returns is to combine the distribution of μ with the conditional distribution $v|\mu$, a direct application of Bayes'

Theorem. Below is a sketch of the proof from *Meucci (2010)*.

$$f_\mu(\mu) = \frac{|(\tau\Sigma)|^{-\frac{1}{2}}}{(2\pi)^{\frac{N}{2}}} e^{-\frac{1}{2}(\mu-\pi)'(\tau\Sigma)^{-1}(\mu-\pi)}$$
$$f_{q|\mu}(q) = \frac{|(\Omega)|^{-\frac{1}{2}}}{(2\pi)^{\frac{K}{2}}} e^{-\frac{1}{2}(q-P\mu)'(\Omega)^{-1}(q-P\mu)}$$

Applying Bayes theorem we write

$$f_{\mu|q}(\mu) = \frac{f_{\mu,q}(\mu,q)}{f_q(q)} = \frac{f_{q|\mu}(q)f_\mu(\mu)}{\int f_{q|\mu}(q)f_\mu(\mu)d\mu} \tag{1.3}$$

The numerator in (1.3) is after some matrix calculations

$$\begin{aligned}
f_{\mu,q}(\mu,q) &= f_{q|\mu}(q)f_\mu(\mu) \\
&\propto |(\tau\Sigma)|^{-\frac{1}{2}}\,|\Omega|^{-\frac{1}{2}}\,e^{-\frac{1}{2}[(\mu-\pi)'(\tau\Sigma)^{-1}(\mu-\pi)+(q-P\mu)'\Omega^{-1}(q-P\mu)]} \\
&= |(\tau\Sigma)^{-1}+P'\Omega P|^{\frac{1}{2}} e^{-\frac{1}{2}(\mu-\mu_{BL})'((\tau\Sigma)^{-1}+P'\Omega^{-1}P)(\mu-\mu_{BL})} \\
&\times\ |\Omega+P\Sigma P'|^{\frac{1}{2}} e^{-\frac{1}{2}(q-\bar{q})'(\Omega+P\Sigma P')^{-1}(q-\bar{q})} \\
&= h(\mu,q)\times m(q)
\end{aligned}$$

where

$$\mu_{BL} = ((\tau\Sigma)^{-1}+P'\Omega^{-1}P)^{-1}((\tau\Sigma)^{-1}\pi+P'\Omega^{-1}q)$$

and $\bar{q}$ is some constant that does not depend on μ or q. This means that the entire second part of the expression will disappear dividing with the denominator.We have

$$\begin{aligned}
f_{\mu|q}(\mu) &= \frac{f_{\mu,q}(\mu,q)}{f_q(q)} \\
&= \frac{h(\mu,q)m(q)}{f_q(q)} \\
&\propto h(\mu,q)
\end{aligned}$$

which gives us the covariance matrix of the blended distribution as well

$$\Sigma^{\mu}_{BL} = ((\tau\Sigma)^{-1}+P'\Omega^{-1}P)^{-1}$$

Note that we have now computed the mean and covariances of the distribution of $\mu|q$, however we want the distribution of $\mathbf{X}|q$.However we have $\boldsymbol{(}X\boldsymbol{)} = \mu + Z$ where Z is independent of q and $Z \sim N(\mathbf{0},\Sigma)$ which leads to

$$\begin{aligned}
\mathbf{X}|q &= \mu|q+Z \\
\Sigma_{BL} &= \Sigma+\Sigma^{\mu}_{BL}
\end{aligned}$$

By re-shuffling the terms we obtain the BL-formulas in its usual form

$$\begin{aligned} \mu_{BL} &= \pi + \tau\Sigma P'(\tau P\Sigma P' + \Omega)^{-1}(q - P\pi) \\ \Sigma_{BL} &= (1+\tau)\Sigma - \tau^2\Sigma P'(\tau P\Sigma P' + \Omega)^{-1}P\Sigma \end{aligned} \tag{1.4}$$

By highlighting contradictions to intuitions in the extreme null- and full-confidence limits of the BL posterior distribution, *Meucci (2005)* suggested to set $\mu \equiv \pi$ and eliminating the estimation risk and the use of τ in the model. Note that is only a special case of the original model with $(X) \sim N(\pi, \Sigma)$.Derivation is exactly the same as in the original model, which leads to the results:

$$\begin{aligned} \mu^m_{BL} &= \pi + \Sigma P'(P\Sigma P' + \Omega)^{-1}(q - P\pi) \\ \Sigma^m_{BL} &= \Sigma - \Sigma P'(P\Sigma P' + \Omega)^{-1}P\Sigma \end{aligned} \tag{1.5}$$

For testing purposes we will use the original BL-model with different values for τ depending on the estimation method used.

2 The Copula-opinion pooling model

The BL-framework is a breakthrough in portfolio optimization, opening up another territory worth years of research. Many studies has been done aiming at improving the existing framework and tackling parameter estimation problems within the original framework, which remains to be one of the biggest drawbacks for BL. Another problem with BL is its assumptions for normality on both market equilibrium returns and views.This assumption gives rise to easily editable formulas but not sufficient to describe the highly non-normal market return conditions.

With the increasing popularity of Monte-Carlo simulations, the use of empirical distributions becomes all the more popular. Given large samples, which could be achieved by bootstrapping or other methods of data-mining, a fairly accurate picture of the market returns could be constructed. This method is already widely used when it comes to risk management calculations; however it is a fairly new method with regard to blending market equilibrium returns with investor views.

This chapter begins with an introduction of the copula-opinion pooling framework introduced by *Meucci (2006).* Certain modifications will be presented in order for this framework to be more comparable to the original BL-framework. Further testing results and comparisons between the two methods will be provided in CHAPTER IV.

2.1 The market

According to *Meucci (2006)* the COP model is to address the following issues of the BL framework.

> *In reality the markets are in general highly non-normal. Furthermore, practitioners might wish to input their views in less informative ways than the alpha plus normal noise prescription of BL,for instance by means of uniform distributions on ranges. Finally, the markets are not necessarily represented by a set of returns, for instance for a trader of Asian options the market is the set of prices of one single underlying at specific monitoring times.*

Like BL, the model starts off with a multivariate market represented by an N-dimensional market returns M. Unlike BL, M is a matrix containing all the available historical returns that are appropriate to take into consideration. For instance, if there are five assets and 200 days of data the M matrix would be a 200×5 matrix.

We can argue that 200 data points are hardly enough to motivate the use of empirical distribution. *Meucci (2006)* suggested parametric bootstrap where the data matrix is first fitted into a multivariate skew-t distribution, which is later used to simulate another market matrix large enough to motivate the use of empirical distribution. However this approach involves assumption of a distribution unknown to the investors. To be able to blend market views with investor views without involving any distributional assumption on the market, we might suggest to bootstrap more data points through randomly drawing from the old sample. However this method does not introduce any new information into the existing dataset and thus will not improve the accuracy significantly. Since we already have parametric method for blending market returns with expert views (BL), we choose to implement this method without any additional assumption on the market distribution.

For later conveniences the original empirical market cumulative distribution function (cdf) and the marginal cdf for i :th asset is assumed to be:

$$\begin{aligned} \text{cdf:} \quad & F^m(\mathbf{X^m}) = F^m(X_1^m, X_2^m, \cdots, X_N^m) \\ \text{marginal cdf:} \quad & F_j^m(x_j) = F^m(\cdots, \infty, x_j, \infty, \cdots) \\ & j = 1, \cdots, N \end{aligned}$$

2.2 The views

As in the BL framework, the views are a set of K linear statements on market asset returns. *Meucci (2006)* adopted author's own modification on the original BL and expresses the views directly on market returns. The views results as usual in a $K \times N$ pick matrix, with its k :th row representing the k :th view of the investors'. Furthermore, P does not necessarily have to be a square matrix which leads to problems if P needs to be invertible(as shown later in the section). The author suggested adding rows to P where no views are explicitly expressed. It has the same effect of adding $(N-K) \times N$ rows of zeros to P. Rather than adding rows of zeros, a $(N-K) \times N$ matrix, $\vec{P}$, that spans up the complimentary space of P is added, resulting in the invertible $N \times N$ matrix $\bar{P}$.

With D days of market log-return data points, N assets and K views the model becomes:

$$M = \begin{pmatrix} X_{i,j}^m & \cdots & X_{i,N}^m \\ X_{i+1,j}^m & \ddots & X_{i+1,N}^m \\ \vdots & \ddots & \vdots \\ X_{D,j}^m & \cdots & X_{D,N}^m \end{pmatrix} \quad P = \begin{pmatrix} 1 & 0 & \cdots \\ 0 & \ddots & \vdots \\ \vdots & \vdots & 1 \end{pmatrix} \quad q = \begin{pmatrix} q_1 \\ \vdots \\ q_K \end{pmatrix}$$

$$\bar{P} = \begin{pmatrix} P \\ \vec{P} \end{pmatrix} \quad V = \vec{P}M \quad j = 1, \cdots, N \quad i = 1, \cdots, D$$

where $X_{i,j}^m$ means the i :th day log-return observation of the j :th asset X_j^m. The above yields a $D \times N$ matrix V, which is the view-adjusted market returns. V_{cj} is the i :th column of matrix, representing the observations of stochastic variable V_j during a D day period. Only the first K columns in V are relevant and subject to further computations.

$$V = \begin{pmatrix} V_{i,j} & \cdots & V_{i,N} \\ V_{i+1,j} & \ddots & V_{i+1,N} \\ \vdots & \ddots & \vdots \\ V_{D,j} & \cdots & V_{D,N} \end{pmatrix} = [V_{c1}, V_{c2}, \cdots, V_{cN}]$$

$$V_{cj} = \begin{pmatrix} V_{i,j} \\ V_{i+1,j} \\ \cdots \\ V_{D,j} \end{pmatrix}$$

Apart from computing the view-adjusted matrix V, a marginal cumulative distribution function for the i :th view, $F_i^v(v_j) = F^v(\cdots, \infty, v_j, \infty, \cdots)$ needs to assigned to each of the views. While market returns are difficult to estimate with distributions the views are far easier and perhaps more intuitive as well to model with an assumed distribution. Any distribution that suits the investor views could be applied. A couple of suggestions are:

Multivariate Normal Distribution This assumption is used in the original BL-framework. An application of this assumption in the COP model will make the two models comparable in the sense that both models will have the same inputs. However the results from the two models will still differ. This result will be shown in later sections.

$$F_j^v(v_j) = \Phi_{q,\Omega}(v_j) = P(V_j \leq v_j)$$

where q and Ω is the mean and covariances of the K -dimensional multivariate normal distribution, just as in BL.

Uniform Distribution This class of distribution is not used in the BL-framework. It is suggested by *Meucci (2006)*.The distribution is sensible in the way that when an investor thinks of a view, target returns in particular, she normally thinks of a range in which the future return is likely to occur rather than a normal distributed stochastic variable. The uniform distribution provides them with such a tool and spares the problem of estimating the covariances in the normal distribution case.

$$F_j^v(v_j) = U_{a,b}(v_j) = P(V_j \leq v_j)$$

where U is the cdf for an uniform distribution between (a, b). For instance, if the view is *asset 1 return will be between two and five percent* the distribution will be uniform with $a = 0.02$ and $b = 0.05$.

2.3 Combining the views with market returns

For those assets where the investor has views upon, a numerical view-blending technique is be used to make the views count. $Meucci$(2006) suggests the opinion pooling technique to obtain the posterior marginal distribution F_j^{post}.

$$F_j^{post}(x_j) = F^{post}(\cdots, \infty, x_j, \infty, \cdots) = (1 - p_j)F_j^m(x_j) + p_j F_j^v(x_j)$$

$p_j \in (0, 1)$ is the confidence coefficient as in the the BL-framework. The less certain the view is, the less becomes p_j and the more weight will be put on the original marginal market distributions. However even though p_j has the same interpretation as c_j:s in the BL, they are defined for different intervals. A transformation function will be introduced in Chapter IV to solve the problem.

The second step in the COP approach is to determine a joint distribution which is consistent with the posterior marginal distribution and also inherits the dependency structure of the original market returns. This calls for the use of Copula.

Copula is a distribution function C of a random vector $U = (U_1, U_2, \cdots, U_k)$ whose components U_j are uniformly distributed on $(0, 1)$. It is defined as

$$C(u_1, \cdots, u_k) = P(U_1 \leq u_1, \cdots, U_k \leq u_k)$$

Note that both F_j^v and F_j^{post} are continuous marginal distribution functions and that

$$\begin{pmatrix} F_1^v(V_1) \\ \vdots \\ F_k^v(V_k) \end{pmatrix} \equiv \begin{pmatrix} U_1 \\ \vdots \\ U_k \end{pmatrix}$$

where the distribution of U_j can be specified empirically through observations of V_j, V_cj. Now we can use the marginal distribution of F_j^{post} to transform the uniform U_j:s back to the posterior empirical distribution of V, V^{post}.

$$V^{post} = \begin{pmatrix} V_1^{post} \\ \vdots \\ V_k^{post} \end{pmatrix} = \begin{pmatrix} F_1^{post-1}(U_1) \\ \vdots \\ F_k^{post-1}(U_k) \end{pmatrix}$$

When inverse transforming from U_j to V^{post} interpolation technique is used in MATLAB.Note that one drawback of the COP approach is that we have to explicit give the marginal cdf:s for each of the view distributions so they have to be independent of each other.Note also that if p_j for view j is very small the term $p_j F_j^v(x_j)$ will be close to zero and we will not change the original empirical view distribution V implied by the market by a noticeable amount. But if p_j is large or that V_j has very narrow variance around the mean,the term $p_j F_j^v(x_j)$ will dominate in the posterior marginal distribution.

Finally, to determine the posterior joint distribution of the market F^{mpost} the inverse of the pick matrix is used.

$$M^{post} = \bar{P}^{-1}V^{post}$$

Note that the extra column added to P is used here to ensure that the inverse of P exists. The first column of M^{post} is only a function of the first column in V^{post}, which means that the posterior return distribution of asset j is only dependent on the view distribution of the same asset. If the view distribution for asset j did not change as there were no view imposed on the asset, the i :th column in V^{post} will be the same as in V, which multiplied by $\bar{P}^{-1}$ gives back the i :th column in M.

The COP method will be compared head-to-head with the BL-framework in CHAPTER IV. But already here we have a feeling that the method can be more sensitive to the choice of the views compare to BL.

3 The shrinkage estimator

The best guess about future is usually obtained by computing the average of past events. Stein's paradox defines circumstances in which there are estimators better than the arithmetic average. - Bradley Efron & Carl Morris

When Charles Stein of Stanford University first published the result in statistics in 1955, it was initially met by frequent and sometimes angry debates and resistance for long period of time. In traditional statistical theory it was for long believed that no other estimation rule is uniformly better than the observed average. Stein's paradox lies in the statement that if there are three or more groups of means within a sample and if we are interested in predicting future means of the groups there is a procedure that always gives the better estimation in the sense of mean-squared error than simply extrapolating from the observed averages of the groups.

Applying this to the context of market asset returns estimation, Stein's theory says that if we are interested in predicting future market asset returns we can do better than simply averaging the past returns.

This chapter will start by introducing the James-Stein estimator, giving a sketch of the proof provided by Stein in the original paper in 1955 and *James & Stein (1961)*. We will move on to several modifications of the original model suggested by *Stanley Sclove (1960)*, *Jobson & Korkie (1980, 1981)* and *Jorion (1986)*. The chapter will conclude with a model which will be implemented in Chapter IV.

3.1 The original proposal

In *James & Stein (1961)* the author proposed a new class of estimator for the mean of a multivariate normally distributed dataset. On the surface this new class of estimators seems counter-intuitive. Each independent experiment should not affect other experiment, for instance given that students takes tests independently one student's performance should only be judged on the student's personal test scores of the different tests. Yet according to Stein the performance of the entire class should be taken into consideration. What Stein proved in his paper is that by making use of test scores of the entire class we can construct better estimates for individual performances.

Let X be a random n-vector whose expected value is the completely unknown vector ζ and whose components are normally distributed with known covariance matrix Σ. We want to solve for ζ while minimizing the loss function L given by

$$L(\zeta, d) = (\zeta - d)'\Sigma^{-1}(\zeta - d)$$

where d is the vector of estimates. In the case of independent multivariate normal distribution with covariance matrix $\sigma^2 I$ the loss function becomes:

$$\begin{aligned}
L(\zeta, d) &= (\zeta - d)'\Sigma^{-1}(\zeta - d) \\
&= \frac{(\zeta - d)'I(\zeta - d)}{\sigma^2} \\
&= \frac{\sum(\zeta_i - d_i)^2}{\sigma^2} \\
&\propto E[(\zeta - d)^2]
\end{aligned}$$

with I as the diagonal unity matrix. Let us now introduce two independent stochastic variable Y and S where Y is p-dimensional and S represents the known variance of the sample of Y.

$$\begin{aligned}
Y &\sim N(\zeta, \sigma^2 I) \\
S &\sim \sigma^2 \chi^2(n)
\end{aligned}$$

Also we introduce two estimators of the form

$$\begin{aligned}
d_1 &= Y \\
d_2 &= (1 - \frac{aS}{\|Y\|^2})Y
\end{aligned}$$

We can show that d_2 is a more efficient estimator than d_1 by showing that

$$\begin{aligned}
E(d_2 - \zeta)^2 &< E(d_1 - \zeta)^2 \\
&= E(\underbrace{Y - \zeta}_{N(0,\sigma^2 I)})^2 \\
&= \sigma^2 p
\end{aligned}$$

we extend the error term for d_2 and obtain

$$\begin{aligned}
E(d_2 - \zeta)^2 &= E((1 - \frac{aS}{\|Y\|^2})Y - \zeta)^2 \\
&= E(Y - \zeta)^2 - 2aE(S\frac{S(Y - \zeta)'Y}{\|Y\|^2}) + a^2 E(\frac{S^2}{\|Y\|^2}) \\
&= \sigma^2\{p - 2aE(\frac{S}{\sigma^2})E(\frac{(\frac{Y}{\sigma} - \frac{\zeta}{\sigma})'\frac{Y}{\sigma}}{\|\frac{Y}{\sigma}\|^2}) + a^2 E(\frac{S}{\sigma^2})^2 E(\frac{1}{\|\frac{X}{\sigma^2}\|^2})\} \\
&= \sigma^2\{p - 2an(p-2)E\frac{1}{p - 2 + 2K} + a^2 n(n+2)E\frac{1}{p - 2 + 2K}\}
\end{aligned}$$

using the following relationship (without proof)

$$\begin{aligned}
E\frac{1}{\|Y\|^2} &= E\frac{1}{\chi^2_{p+2K}} = E\frac{1}{p - 2 + 2K} > 0 \\
E\frac{(X - \zeta)'X}{\|X\|^2} &= (p-2)E\frac{1}{p - 2 + 2K}
\end{aligned}$$

Now taking the derivative with respect to a yields that $a = \frac{p-2}{n+2}$ minimize the loss function.

$$\begin{aligned} E(d_2 - \zeta)^2 &= \sigma^2\{p - \frac{n(p-2)^2}{n+2}E\frac{1}{p-2+2K}\} \\ &< \sigma^2 p \end{aligned}$$

As *James & Stein (1961)* suggested the result could be further generalized to the case where the covariance matrix is unknown but an estimate based on a Wishart matrix with n degrees of freedom is available.

$$\begin{aligned} d_3 &= (1 - \frac{c}{Y'SY})Y \\ c &= \frac{p-2}{n-p+3} \end{aligned}$$

We want to show now that an estimator of this form is superior, regardless of the vector it shrink towards. We start by rewriting:

$$\begin{aligned} d_3 &= (1 - \frac{c}{Y'SY})Y \\ &= (1 - \frac{c}{(\bar{Y}-\bar{0})'S(\bar{Y}-\bar{0})})(\bar{Y}-\bar{0}) \end{aligned}$$

now we have replaced Y with its mean $\bar{Y}$ and according to the theorem d_3 should be superior $\bar{Y}$ if $p \geq 3$.But since a normal distributed variable plus a constant is another normal distributed variable we can without loss of generosity replace $\bar{0}$ with some shrinkage vector K.

$$\begin{aligned} d_3 &= (1 - \frac{c}{(\bar{Y}-K)'S(\bar{Y}-K)})(\bar{Y}-K) + K \\ &= (1-\omega)\bar{Y} + \omega\bar{Y} \\ \omega &= \frac{c}{(\bar{Y}-K)'S(\bar{Y}-K)} \end{aligned}$$

Looking at the formula so far we realize that there is no guarantee that $1-\omega$ will be positive. In the late 1960's, Stanley Sclove worked out a modification of the original James-Stein shrinkage estimator where the term $(\bar{Y}-K)'S(\bar{Y}-K)$ is interpreted as a F-test result against the null hypothesis that $\bar{Y} = K$. If test is passed with confidence level c, all weight will be put on K, with ω taking the value 1, otherwise the estimator will shrink towards K depending on how strongly the F-test came out.

3.2 The extensions

Jobson, Korkie and Ratti (1979) are the first to use James-Stein estimators in context of portfolio optimization. They used simulation analysis to show that James-Stein estimators does lead to more reasonable results compared to the historical average estimators. The idea of shrinkage has further been developed in a Bayesian framework by *Jorion (1986)* in which he proposed the following Baye-Stein estimator:

$$\begin{aligned}
E(\hat{Y}) &= d_4 = (1-\omega)\bar{Y} + \omega Y_0 I \\
\omega &= \frac{N+2}{(N+2) + (\bar{Y} - Y_0 I)' T \Sigma^{-1} (\bar{Y} - Y_0 I)} \\
\lambda &= \frac{T\omega}{1-\omega} \\
V(\hat{Y}) &= \Sigma(1 + \frac{1}{T+\lambda} + \frac{\lambda}{T(T+1+\lambda)} \frac{I'I}{I'\Sigma^{-}1I}) \qquad (3.1)
\end{aligned}$$

Inspecting closely this is a slightly modified James-Stein estimator with shrinkage vector K equals to $Y_0 \times I$ which happens to be the average return of the minimum variance portfolio optimization. Note that I here is the unity vector. Also the introduction of the shrinkage covariance matrix is new and not included in the original James-Stein shrinkage estimator. Most importantly the Baye-Stein shrinkage estimators eliminates the need of finding a suitable shrinkage vector that gives noticeable improvement compared to shrinking towards a zero vector.

The shrinkage vectors is under constant improvement. While they all achieve less mean-squared error compared to the observing-average, their ability to predict future log-return remains questionable. We will implement Baye-Stein estimator into the original BL-framework to see if it results in more stable market equilibrium return estimation.

4 Implementation

While combining expert views with market return seems a very nice idea for predicting future returns on a portfolio of assets, it involves more estimates as inputs to the models. We have already discussed some approaches of setting the extra parameters in previous chapters. The goal of this chapter is to put some of these approaches to the test and eventually find out whether expert views do help in the prediction of future returns. We setup three methods and one benchmark. The three methods are:

- **COP:** *Copula-opinion pooling*
- **BL:** *Original Black-Litterman*
- **BLS:** *Black-Litterman with shrinkage estimators*

All three methods are compared to the benchmark **BM**, the observed historical average estimator. A neutral view and a good view are used with varying degrees of certainty. Three measures are designed to test the quality of the expected return estimates, the variance estimates and the overall degree of matching with market-realized log-return distributions.

4.1 The data

800-day time series of OMX30, STOXX50, NIKKEI 225, SPX500, OMRX and the three- month STIBOR rates are used for testing. The data series starts on 2008/01/03, undergoes the financial crisis with increasing volatility and decreasing interest rates on the markets and stabilizes on post-crisis financial markets in 2009. Certain impacts of the European debt crisis is reflected in the data as well. The data ends on 2011/04/27. As could be seen in Figure 1, for all indices market volatility is constantly varying over the test period.

In line with market practice a history of 200 trading days will be considered for predicting future quarterly returns. Shorter historical periods may result in lack of empirical data for the copula opinion-pooling method, as it may not capture enough variations of the market itself and thus cannot be used to represent the equilibrium state of the market. Longer historical periods may use data that is too old and does not have any impact on the concurrent market. A future predicting period of between one and three months is preferred to represent realistic holding period of the portfolio. Shorter future predicting period may need far more sophisticated models to reflect short-term market movements than only expert target returns. This means that a longer prediction period is preferred to include enough mean-reverting in the data. This motivates the choice of a quarterly predicting window.

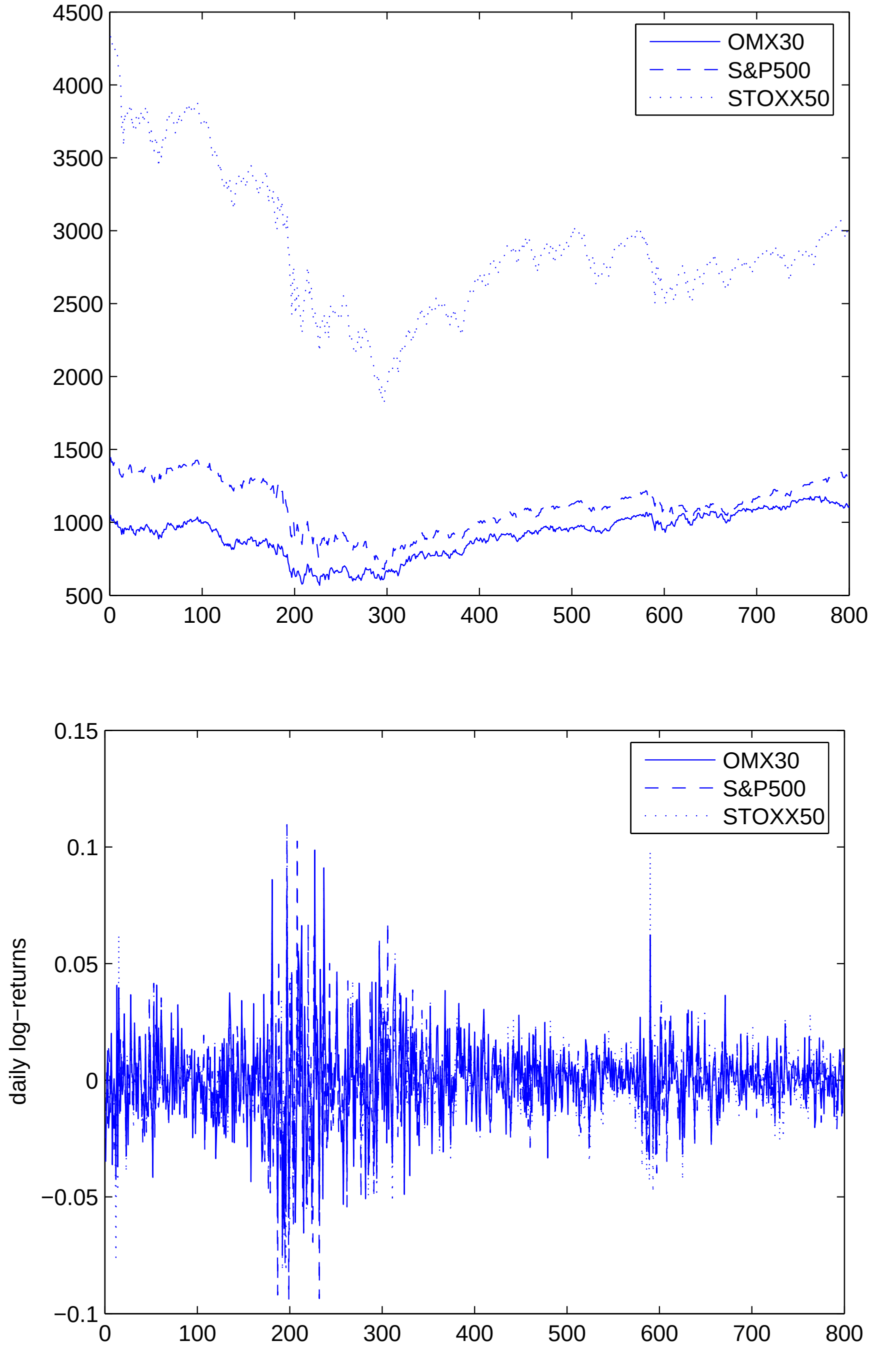

Figure 1: absolute and log-return plot for 3 of the under-lyings through out the testing period. The two volatility peaks come from the financial crisis and the European debt crisis.

4.2 The models

We start here with a list of variables.

- M_{Prior}: the 200 days of historical data describing the log-return time series of the five indices
- M_{Post}: the 60 day forecast of the log-return time series of the five indices
- P: pick matrix describing the views
- q: target return vector
- ω_s: shrinkage coefficient
- Y_0: minimum variance portfolio return, shrinkage vector
- $\hat{\mu_{BL}}$ $\hat{\Sigma_{BL}}$: shrinkage estimator of market prior mean and covariances
- μ_{real} Σ_{real} r_{real}: realized average log-return, covariances and market total return of the future 60 days
- μ_{COP} μ_{BL} μ_{BLS} μ_{BM}: estimated expected daily log-returns from COP, BL, BLS and BM
- Σ_{COP} Σ_{BL} Σ_{BLS} Σ_{BM}: estimated covariances for daily log-returns from COP, BL, BLS and BM

4.3 Copula-opinion pooling

The general theory has been outlaid in Chapter II. A 200×5 M_{Prior} is used and re-sampled to a 2000×5 matrix with non-parametric bootstrap. An outcome of the model is the 60×5 view-adjusted posterior matrix M_{Post}. If no views are expressed M_{Post} will be the same as M_{Prior} apart from the length. Depending on the confidence levels of the views M_{Prior} will be more or less adjusted towards the view distributions. The only distribution assumption comes from the normal distribution assumption of the views. While it is not necessary to assume normality it gives COP the same sets of inputs as in other models and make them easier to compare. However we should bear in mind at all time the COP is different from BL. In the extreme case where we do not have any confidence in the views, the model simply estimate future expected returns directly from the 200 historical data points without any modification.

4.4 Original Black-Litterman

An original BL-framework is set up to obtain μ_{BL} and Σ_{BL} from M_{Prior}, P and q. We assume that market is in equilibrium and we are thus able to obtain the equilibirum return vector $\pi = \mu_{BM}$ and its covariances directly from M_{Prior}. Using μ_{BL} and Σ_{BL} we use parametric bootstrap to obtain the future 60 day forecast of log-returns. The choice of τ plays a very important part in the performance of BL.

4.5 BL with shrinkage estimator

Following the concept of shrinkage estimators in Chapter III, it is possible to construct a multivariate mean estimator that is superior to the benchmark. The Bayes-stein shrinkage estimator is used to replace μ_{BM} with $\hat{\mu_{BM}}$ as inputs in the original BL. Following this change we need to modify τ in the BL model as well. Recall that the shrinkage estimator is

$$\begin{aligned}
\hat{\mu_{BM}} &= (1-\omega_s)\mu_{BM} + \omega_s Y_0 \\
\mathrm{Var}(\hat{\mu_{BM}}) &= (1-\omega_s)^2 \mathrm{Var}(\mu_{BM}) = \frac{(1-\omega_s)^2}{N}\Sigma_{BM} \\
&\Rightarrow \quad \tau = \frac{(1-\omega_s)^2}{N} \qquad (4.1)
\end{aligned}$$

where Y_0 is a scalar and the minimum variance portfolio return containing all five indices. While everything else stays the same as in the original BL, a ω_s of $(0.4, 0.8)$ implies quite a change from the observed historical average. The shrinkage estimator gives more stable inputs into the BL framework, which should at least lead to more stable results from the model.

4.6 Views and confidence levels

The views are expressed as independent normally distributed variables $q \sim N(P\mu, \Omega)$. Ω is defined as a diagonal matrix with each element representing the uncertainty in each view.

$$\Omega = \begin{pmatrix} \omega_1 & 0 & \cdots \\ 0 & \ddots & \cdots \\ \cdots & \cdots & \omega_5 \end{pmatrix} \qquad \omega_i = \frac{1}{c^2} P_i' \Sigma P_i \quad c \in (0, \infty)$$

when pooling together the marginal prior- and view distributions in COP we use the pooling factor $p \in (0, 1)$.

$$F_j^{post}(x_j) = (1-p_j)F_j^m(x_j) + p_j F_j^v(x_j) \qquad (4.2)$$

To make COP relatively comparable to BL we already choose to give them exactly the same sets of inputs. However we also need to find a way to transform c into p, since the p and c have the same interpretation.

We assume now that both distribution cdfs F^m and F^v are normal distributions so that any linear combination of them is still normally distributed. We might start thinking that $F^{post} \sim N(\mu_{BL}, \Sigma_{BL})$ if $\mu_{COP} = \mu_{BL}$ and $\Sigma_{COP} = \Sigma_{BL}$. However we can see that this is not true by reviewing the following:

$$\begin{aligned}
\Sigma_{BL} &= \underbrace{\Sigma}_{marketpriorpart} + \underbrace{\Sigma^{\mu}_{BL}}_{viewcorrectionpart} \\
\lim_{\Omega\to 0}\Sigma_{BL} &= \lim_{\Omega\to 0}(1+\tau)\Sigma - \tau^2\Sigma P'(\tau P\Sigma P' + \Omega)^{-1}P\Sigma \\
&= (1+\tau)\Sigma - \tau^2\Sigma P'(\tau P\Sigma P')^{-1}P\Sigma \\
&= (1+\tau)\Sigma - \tau\Sigma P'(\Sigma P')^{-1}P^{-1}P\Sigma \\
&= (1+\tau)\Sigma - \tau\Sigma \\
&= \Sigma
\end{aligned}$$

In the extreme case $|\Sigma^{\mu}_{BL}|$ goes to zero but Σ_{BL} does not. This means that in the original BL the posterior covariances will never be less than $|\Sigma|$. However COP does not involve the concept of market equilibrium distribution $\pi \sim N(\pi, \tau\Sigma)$, it expresses the posterior distribution directly on the market log return distributions. We can see this by letting p go to 1 in Equation (4.2).

$$\lim_{p\to 1} F^{post}_j(x_j) = F^v_j(x_j)$$

where the variances of $F^v_j(x_j)$ will also go to zero as views become more confident. Thus the original COP is only comparable with Meucci's modified version of the BL where the posterior distribution is calculated directly upon the market returns X, not the equilibrium return μ. We will see more consequences of this in the result section. For now this means that we can use Meucci's modified BL and solve the following problem to obtain the $p \in (0,1)$ that minimize the sum of mean and variances for all $c \in (0, \infty)$. We use period $(350, 550)$ as reference since it is a normal period with relatively constant volatility shown in Figure 1.

$$\min_p |\mu^m_{BL} - \mu_{COP}|^2 + \sum |\Sigma^m_{BL}(i,i) - \Sigma_{COP}(i,i)| \tag{4.3}$$

The square on the means is to bring the mean-error to the same level as variances are quadratic. We try here to find out c and p that gives the most similar single-asset performances with same view and same prior market information. μ^m_{BL} and Σ^m_{BL} is the expectation and covariances of the modified BL, consistent with Chapter I. The solution is plotted below with several values highlighted.

$$\begin{aligned}
c = 0.5 &\Leftrightarrow p = 0.1910 \\
c = 1 &\Leftrightarrow p = 0.4725 \\
c = 5 &\Leftrightarrow p = 0.9283
\end{aligned}$$

Even after the above transformation we should still bear in mind that COP and BL are not directly comparable due to the fact that COP in our case is purely implemented with empirical distributions and both BL methods uses parametric bootstrap to obtain future forecast.

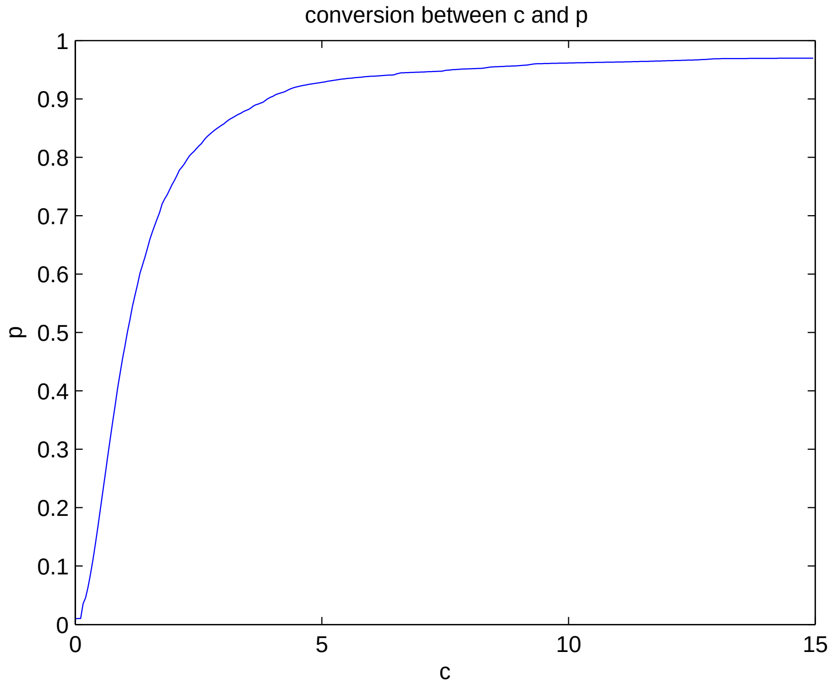

Figure 2: c-p plot shows clearly that the relation between c and p is not linear but the function is strictly increasing which confirms the similarity in interpretation.

4.7 Test methods

Three tests have been implemented. The results are generated using the data described above and may vary under other circumstances. The tests are designed specifically to test the estimation quality of the mean, the variances and the distribution predicting power as a whole.

4.7.1 Mean tests

An appropriate method of measuring the matching degree of the mean vector is the mean-squared error of the estimated mean vector compared to its realized equivalence. Since M_{Prior} represents daily returns we obtain from the models the expectation and covariances forecast of daily log-returns for the prediction period. If we assume time independency we can calculate the future d days return by simply scaling up μ_{BL} and μ_{BLS}. For instance in the BL case with n underlying assets the MSE becomes

$$MSE_{BL} = \frac{|d\mu_{BL} - r_{real}|^2}{n}$$

where r_{real} is the realized future 60-day market returns. In COP we obtain the entire posterior distribution M_{post} as an empirical distribution and we need to estimate μ_{COP} and Σ_{COP} from M_{Post}. MSE gives the overall performance of the models in the sense of mean estimation but it does not tell us anything about the error in individual asset mean estimations. The overall MSE may still be low if the model returns good result for some of the assets while it gives totally wrong estimations for other assets. We can calculate the mean-error for the assets by taking

$$ME_{BL} = \frac{\sqrt{nMSE_{BL}}}{n} = \frac{1}{\sqrt{n}}\sqrt{MSE_{BL}}$$

4.7.2 Variance tests

Assuming that the daily log returns are normally distributed χ^2-tests could be conducted to test the degree of variance matching. If Y_i is normally distributed with $N(0,1)$ then $Y_1^2 + Y_2^2 + \cdots + Y_n^2$ is χ^2 distributed with degree of freedom n.

In our model if we only concentrate on the i:th asset log-returns in BL they are normally distributed with $N(\mu_{BL}(i), \Sigma_{BL}(i,i))$. Since we want to test variances we can minimize the noises by replacing μ_{BL} with the sample mean mu_{real}, the average log-return of the future 60-day markets. We thus lose one degree of freedom. For 60 days forecast we will have a χ^2 distribution with 59 degrees of freedom, and the following 95 percent confidence interval for the test scores, for the i:th asset:

$$\chi^2_{0.025}(59) = 39.6619 \leq \chi_i^2 = \frac{\sum (M_{post}(i) - \mu_{real})^2}{\Sigma_{BL}(i,i)} \leq \chi^2_{0.975}(59) = 82.1174$$

If the test score lies outside the interval we can reject the H0 hypothesis that the variance of the i:th asset is $\Sigma_{BL}(i,i)$, in the BL case. Here we could also test the correlations between the assets by adding asset returns along a given direction. But this effect will not be considered here.

4.7.3 Distribution test

Apart from conducting tests on means and variances it would be interesting to see how well the realized and modeled distributions match as a whole. For the purpose of this we will use the empirical data available from COP while we bootstrap parametrically using multivariate normal distribution from BL, BLS and BM.

QQ-plot QQ-plot could be obtained by sorting and plotting the actual outcome of the market log-return distribution for the prediction period against any of the model results from COP, BL, BLS or BM. If they line up in a 45 degree line they match perfectly, otherwise we will have to analyze each of the plots separately to see how well they match. Due to the number of plots we have to generate (5 assets, 4 methods, 20 plots per test) the QQ-plot becomes unpractical.

The KS test While QQ-plot is a very powerful tool for examining the relationship between two samples, it does not give a score of how similar or different two samples are. For this purpose we introduce the Kolmogorov-Smirnov test.

Suppose we have i.i.d sample $X_1, \cdots, X_n$ with some unknown distribution $F(x) = P_f(X < x)$ and we would like to test if the sample comes from a known distribution $G(x) = P_g(X < x)$ with the following hypothesis:

$$\begin{aligned} H_0 : &\quad F = G \\ H_1 : &\quad F \neq G \end{aligned}$$

The idea of the KS test is to measure the maximum distance between the cumulative distribution functions of the two distributions. If the sample distribution (empirical distribution) matches the assumed distribution perfectly we have

$$\sup |F(x) - G(x)| \to 0$$

If they do not match perfectly we have

$$P_h(\sqrt{n} \sup |F(x) - G(x)| < t) \to H(t) = 1 - 2 \sum (-1)^{i-1} e^{2i^2 t}$$

where $H(t)$ is the Kolmogorov-Smirnov distribution. It is worth noting that for large n

$$\begin{aligned} D_n = \sqrt{n} \sup |F(x) - G(x)| &\quad > \quad \sqrt{n}\delta \quad \delta \neq 0 \\ \lim_{n \to \infty} D_n &\quad \to \quad \infty \end{aligned}$$

The result is very intuitive since by law of large numbers $F(x)$ should converge to $G(x)$ for large n. If it does not converge, that is if $\delta \neq 0$ the test fails. Like other statistics the test will fail on the predefined level of significance α.

The KS-test could be extended to test for similarity between two samples as well. Let $X_1, \cdots, X_m$ have the unknown cdf $A(x)$ and empirical distribution $A_e(x)$ and $Y_1, \cdots, Y_n$ have the unknown *cdf* $B(x)$ and empirical distribution $B_e(x)$

$$\begin{aligned} H_0 &\quad A(x) = B(x) \\ H_1 &\quad A(x) \neq B(x) \\ D_{mn} = (\frac{mn}{m+n})^0.5 \sup |A_e(x) - B_e(x)| \\ \alpha = P_h(D_n \geq c_\alpha) \approx 1 - H(c_\alpha) \\ \text{confirm } H_0 : &\quad D_n \leq c_\alpha \\ \text{reject } H_0 : &\quad D_n \geq c_\alpha \end{aligned}$$

An example For instance it might be interesting to see if parametric bootstrapping is any better at approximating future log-return distribution than using pure historical data. For this purpose we choose to plot the 300-500 day data against 500-700 day data and compute the 500-average of the KS p-value for all assets. We also compare the KS p-value of all 5 asset pairs. Using 5 percent confidence level, the below matrix shows the p-values of the tests. Each value can be interpreted as a measure of how its equivalent QQ-plot performs. Taking the average we see that purely historical data

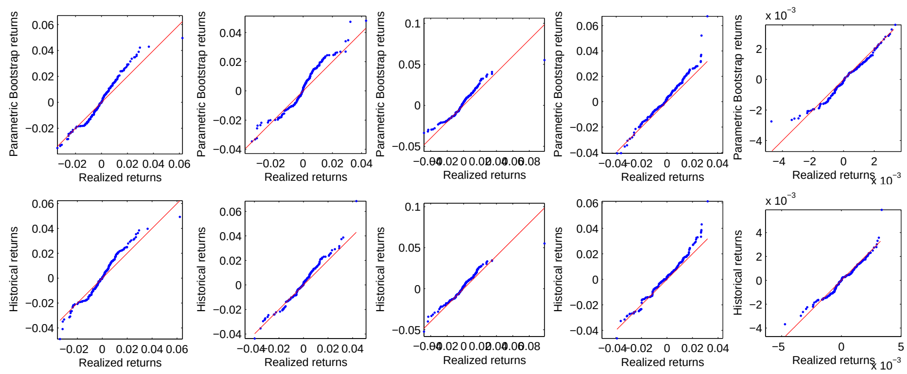

Figure 3: QQ-plots for parametric bootstrapping vs using pure historical data. We can identify any method is noticeably superior to the other.

KS p-value

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
parametric	0.0351	0.1025	0.1024	0.2054	0.3735
historical	0.0545	0.0506	0.1551	0.1947	0.4773

performs slightly better with an average p-value at 0.1864 compared to 0.1638. The two methods thus come quite close to each other on normal markets. This tells us that it is reasonable comparing COP's empirical distribution to the bootstrapped distribution in BL, BLS and BM. Also KS tests performs sufficiently well to replace QQ-plots in our tests.

5 Results

Reviewing Figure 1 we identify the following two periods for single period test purposes.

- Normal market conditions: day 300-500 (2009-03-12 to 2009-12-23) to predict day 500-560 (2009-12-23 to 2010-03-24)
- Abnormal market conditions: day 0-200 (2008-01-03 to 2008-10-15) to predict day 200-260 (2008-10-14 to 2009-01-14)

Also we will use two sets of views. Since all indices could be decomposed into a drift part, a noise part and a jump part according to *Merton 1975a*, the indices return under sufficiently long period should thus be comparable to the drift part (excluding jump effects). Since $e^x \approx 1 + x$ for very small x and e^x is the index pricing formula with drift part only we could argue to set the view to risk-free rates. Since our investment takes place quarterly and we use the 3 months STIBOR rates on the first day of the prediction period as our first view, $q = r_f$. The view is usually not good enough for a 3-month investment.

One advantage with back-testing is that we already know the actual outcome of the markets. We will thus use the actual market return as our second view, $q = r_{real}$. This view is as good as we can get and gives a sense of best-case-scenario performance of the models.

5.1 Single period test

5.1.1 Mean-error test

COP is very view sensitive. For the same reason it performs bad in abnormal markets with high confidence on r_f where the means get pulled in the wrong direction. COP performance is on-par with Meucci's modified BL as in Table 1, partially due to the consistent p and c used in these two models, partially because of the lack of τ in both methods and larger part of view correction is taken into consideration.

The choice of τ has an huge impact on the BL and BLS results. This is the reason why Meucci's BL almost always outperforms the original BL, apart from high confidence r_f view in abnormal markets. Meucci's BL is more sensitive because

$$\lim_{\tau \to 0} \tau \Sigma P' (\tau P \Sigma P' + \Omega)^{-1} = 0$$

and the term grows larger with larger τ given non-zero Ω. By setting τ closer to zero the view correction part will be taken less seriously and in the limit we just get back the market prior part which equals μ_{BM}. This issue can be solved by either using Meucci's modification or by setting a larger τ.

Mean-error in day 300-500				
60-day ME	$r_f, c = 0.5$	$r_f, c = 5$	$r_{real}, c = 0.5$	$r_{real}, c = 5$
COP	0.0304	0.0039	0.0285	0.0122
BL	0.0390	0.0313	0.0390	0.0308
Meucci BL	0.0266	0.0026	0.0260	0.0152
BLS	0.0178	0.0173	0.0178	0.0172
BM	0.0391	0.0391	0.0391	0.0391

Table 1: Mean-error in normal markets. BLS achieves very stable results while COP and Meucci's BL performs best with correct views

Mean-error in day 0-200				
60-day ME	$r_f, c = 0.5$	$r_f, c = 5$	$r_{real}, c = 0.5$	$r_{real}, c = 5$
COP	0.0294	0.0043	0.0242	0.0455
BL	0.0279	0.0275	0.0277	0.0221
Meucci BL	0.0266	0.0049	0.0233	0.0489
BLS	0.0310	0.0293	0.0310	0.0318
BM	0.0279	0.0279	0.0279	0.0279

Table 2: Mean-error in abnormal markets. None of the methods managed to beat BM with $q = r_f$

BLS is all-around better in the ME sense on normal markets, regardless of good or bad views. The stability comes from two sources. The first is the construction of $\tau = \frac{(1-\omega_s)^2}{N}$ that yields even smaller τ than BL with $\omega \in (0.5, 0.9)$. This means that the BLS mean goes towards full confidence view solutions slower compared to BL. The second edge comes from the addition of OMRX with is an interest rate index. This managed to reduce the overall risk of the portfolio. For instance during day $300 - 500$ with $q = r_f$ and $c = 0.5$ the shrinkage estimator was able to scale down the means heavily towards the OMRX return and thus lowering the overall MSE.

r_real	μ_{COP}	μ_{BL}	μ_{BLS}	μ_{BM}
0.0011	0.0017	0.0018	0.0009	0.0018
0.0006	0.0019	0.0020	0.0010	0.0020
-0.0002	0.0019	0.0020	0.0010	0.0021
0.0003	0.0018	0.0019	0.0010	0.0020
0.0002	0.0002	0.0001	0.0002	0.0001

Table 3: daily log-return vector for all models, note that BLS shrinks towards the middle much more than other models due to the shrinkage estimator Y_0

If we throw away the OMRX index from the portfolio we have Table 4 during normal markets. The COP and BL perform the same while the BLS performs worse. With more asset classes and different return levels the shrinkage estimator works best. For single-assets portfolios we need to go for larger τ if we want to use BLS to enable more view-sensitivity.

Mean-error in day 300-500				
60-day ME	$r_f, c = 0.5$	$r_f, c = 5$	$r_{real}, c = 0.5$	$r_{real}, c = 5$
COP	0.0362	0.0034	0.0270	0.0170
BL	0.0488	0.0205	0.0488	0.0179
Meucci BL	0.0338	0.0005	0.0303	0.0195
BLS	0.0473	0.0472	0.0473	0.0472
BM	0.0489	0.0489	0.0489	0.0489

Table 4: Mean-error without OMRX in normal markets. BLS barely changes due close-to-one ω_s which decreases the value of τ

A good guess generally improve the ability to predict regardless of markets while risk-neutral views seems to work on normal markets. All view combining methods give rise to better ME-score than the benchmark during normal market conditions.

5.1.2 Variance test

We carry out χ^2 test with H_0 stating that realized variances during the prediction period of the i:th assets can be described by $\Sigma_{COP}(i,i)$, $\Sigma_{BL}(i,i)$, $\Sigma_{BLS}(i,i)$ and $\Sigma_{BM}(i,i)$. Larger test scores imply that variances from models are too small and vice versa. We only carry out tests on marginal distributions and disregard the effects of correlations between the assets. To capture the correlations between them we can for instance carry out the same χ^2 test with certain linear combinations of the assets. All distributions are assumed to be normal during the test.

daily χ^2 tests for day 300-500 with r_f and $c = 0.5$					
	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	22.392	25.348	44.161	39.126	39.57
BL	18.21	23.036	35.91	32.339	34.19
BLS	18.216	23.043	35.921	32.35	34.135
BM	18.301	23.151	36.089	32.501	34.36
daily χ^2 tests for day 300-500 with r_f and $c = 5$					
	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	278.41	381.88	573.76	39.126	39.54
BL	18.228	23.057	35.949	32.359	34.214
BLS	18.216	23.044	35.923	32.351	34.136
BM	18.301	23.151	36.089	32.501	34.36

Table 5: daily χ^2 tests in normal markets, BL, BLS and BM are not particularly confidence-sensitive while COP is highly sensitive

Variance estimates are consistently too large for BL, BLS and BM and H_0 is rejected for almost all tests. This is a period which started off with large volatilities in the reference period. The prediction period realized volatilities are much less than expected, in line with test results.

Something happened to COP and leads us back to the assumptions of the models. Reviewing the formula for Σ_{BL} we have

$$\Sigma_{BL} = \underbrace{(1+\tau)\Sigma}_{marketpriorpart} \underbrace{-\tau^2\Sigma P'(\tau P\Sigma P' + \Omega)^{-1}P\Sigma}_{viewcorrectionpart}$$

The view correction part is scaled down significantly by τ and is in general less than market prior part unless the views are extremely un-confident. For confidence views the view correction part goes to zero. Indeed it does so much faster than for μ_{BL} due to τ^2. Thus neither BL or BLS should come too far from the benchmark with reasonable levels of view confidence. However with COP we have shown before that the variances goes to zero, just like in Meucci's BL. This is the reason we could not really compare COP with the rest of the models in variance sense. This also gives us an initiative to modify the original Σ_{COP} as follows

$$\tilde{\Sigma_{COP}} = \Sigma_{COP} + \Sigma_{BM}$$

daily χ^2 tests for day 0-200 with r_f and $c = 0.5$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	81.131	93.98	73.474	111.09	19.312
BL	149.5	188.42	137.31	208.53	90.165
BLS	149.63	188.59	137.43	208.72	90.048
BM	150.24	189.36	138	209.58	90.617

daily χ^2 tests for day 0-200 with r_f and $c = 5$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	133.19	166.05	119.81	191.15	39.837
BL	149.67	188.58	137.47	208.68	90.231
BLS	149.63	188.59	137.44	208.72	90.05
BM	150.24	189.36	138	209.58	90.617

Table 6: χ^2 tests in abnormal markets, BL, BLS and BM are not particularly confidence-sensitive while the modified COP is now catching up with the rests but a certain degree of double-counting leads to less test score values

It remains interesting to discuss whether the covariance matrix should become 0 or Σ_{BM} with extreme certain views. It makes sense to set this to 0 as extreme certain views mean that we are sure that asset returns will turn out the way we expected. But on the other side the true covariances of the assets will not drop to zero just because we have certain views. Luckily for portfolio optimizations only the relative differences in variances and correlations matter so Meucci's modifications will still work like a charm for its designed purposes. However we should pay attention to this problem when we calculate the portfolio risk from the models. With the modified the COP we obtain the χ^2 results for abnormal markets in Table 6.

In Table 6 variance estimates are consistently too small, resulting in large χ^2 test scores. Reference period has too low volatility compared to the prediction period. It is worth noting that the prediction period is right after the financial crisis in 2008. In general it is hard to predict future volatilities with historical volatilities, particularly if reference period is as long as 200 days.

COP variances are now in line with the rest. The modification leads to certain double counting effect if the views are not strong enough. If p goes towards 0 COP will just return $M_{post} = M_{prior}$. If we now add Σ_{BM} to it we double-count the covariances while the views stay the same. This double-counting effect decreases quite drastically with increasing p. Just as in Meucci's BL, the market posterior distributions are directly expressed on market prior log-return distributions, adding Σ_{BM} is totally unnecessary and even violates the basic assumptions of $X \equiv N(\pi, \Sigma)$. Because of this we continue using the original Σ_{COP}, but we will bear in mind the problems that it causes.

5.1.3 Distribution test

In this section we will keep COP unmodified so we can get a better understanding of how the original Σ_{COP} behaves. By using the original model we do not have to assume any distribution. By adding Σ_{BM} certain distribution assumptions have to be done in order to obtain a bootstrapped sample for QQ-plots and KS tests.

QQ-plots In the daily log-return distribution plot in Figure 4 and 5 we see a clear anti-clockwise twist from the 45-degree line. With realized quantiles on the x-axis and model quantiles on y-axis this anti-clockwise twist demonstrates that model variances are larger than realized variances. Already at $c = 0.5$ we see traces of COP distributions being different than BL, BLS, BM and also the realized market distributions. It starts to shrink towards the views and highlights even smaller variances than benchmark.

KS test To get a better understanding of the degree of matching of overall distributions we use the KS test to quantify it to a number. To achieve stability in test scores we take the average score of 500 tests. The KS tests are carried out with 5 % confidence with H_0 stating that the model distributions of COP, BL, BLS and BM match the real outcome of the market log distribution.

COP variances goes to zero with higher confidences in the views, as predicted in the χ^2 results. COP scores in normal markets happen to be better because the views are qualified and variances estimates from the reference period are too large but COP tends to make the variances smaller compared to BL.

Overall distribution matching decreases when views are bad and certain, which is shown with Table 9 and Table 10. It increases when views are qualified and certain which is shown in Table 7 and Table 8.

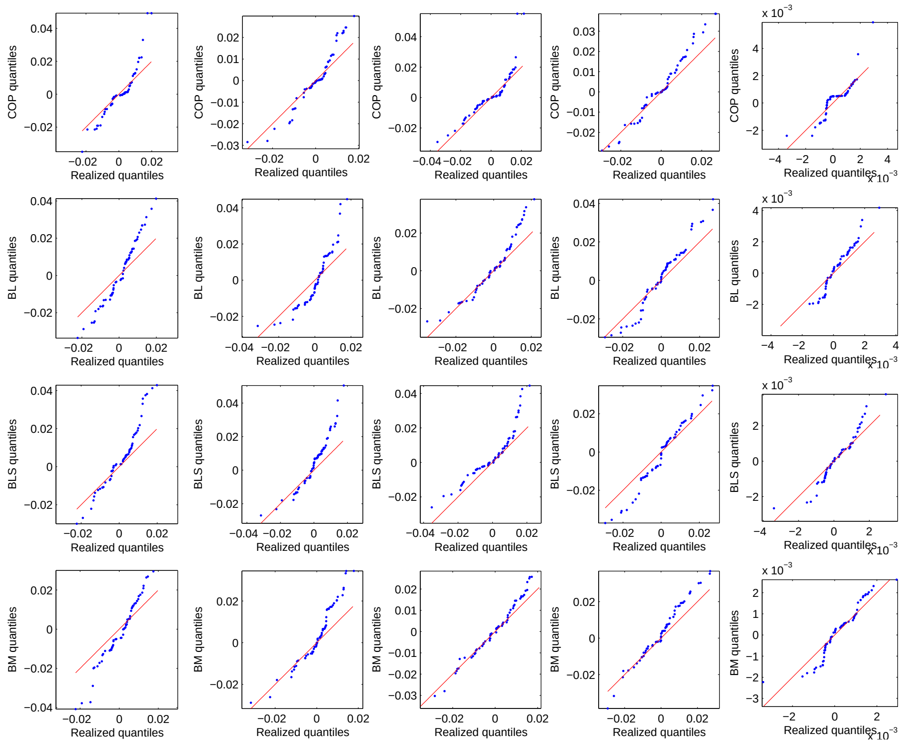

Figure 4: QQ-plots in normal markets with r_f and $c = 0.5$, note that models often return too large variances compared to the actual outcome

daily KS tests for day 300-500 with r_f and $c = 0.5$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	0.232	0.479	0.439	0.515	0.240
BL	0.134	0.113	0.379	0.325	0.223
BLS	0.135	0.131	0.421	0.340	0.272
BM	0.135	0.120	0.379	0.319	0.241

Table 7: p-value for KS tests on normal markets with neutral views and low confidence level

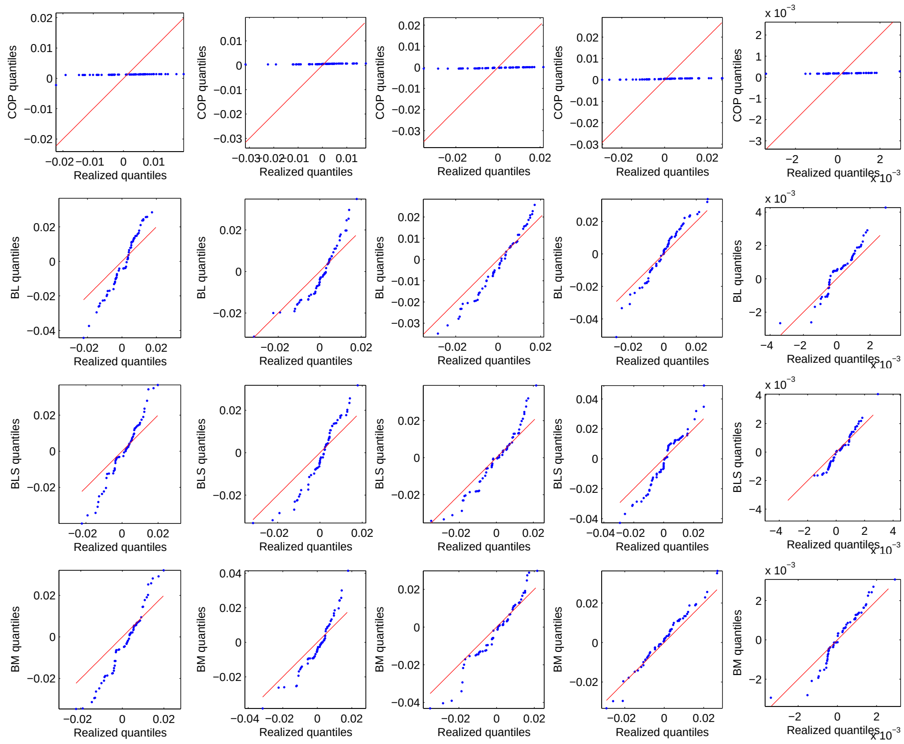

Figure 5: QQ-plots in normal markets with r_f and $c = 100$, note that while BL, BLS stayed roughly the same while COP heavily concentrates its distribution towards the views

daily KS tests for day 300-500 with r_f and $c = 100$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	5.8×10^{-9}	5.6×10^{-8}	1.9×10^{-8}	4.4×10^{-7}	5.3×10^{-8}
BL	0.157	0.124	0.458	0.367	0.273
BLS	0.148	0.124	0.448	0.353	0.272
BM	0.135	0.120	0.379	0.319	0.241

Table 8: p-value for KS tests on normal markets with neutral views and high confidence level. A good guess(r_f is quite good as shown in Table 1) results in slightly better distribution matching. BL and BLS have both reasonable outputs on the variances(compared to benchmark) while their mean estimates in this case come closer to the true values.

daily KS tests for day 0-200 with r_f and $c = 0.5$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	0.237	0.018	0.163	0.082	0.240
BL	0.471	0.163	0.502	0.191	0.183
BLS	0.463	0.158	0.516	0.201	0.124
BM	0.484	0.162	0.528	0.201	0.165

Table 9: KS tests on abnormal markets with neutral views and low confidence level

daily KS tests for day 0-200 with r_f and $c = 100$

	OMX30	STOXX50	NIKKEI225	SPX500	OMRX
COP	2.3×10^{-13}	5.9×10^{-11}	1.5×10^{-12}	6.5×10^{-11}	1.2×10^{-25}
BL	0.062	0.019	0.075	0.090	1.04×10^{-25}
BLS	0.197	0.019	0.161	0.125	2.4×10^{-24}
BM	0.479	0.164	0.512	0.188	0.163

Table 10: p-value for KS tests with excellent views and high confidence level. Bad guesses (see Table 2) result in worse distribution matching than benchmark while variances stay approximately the same.

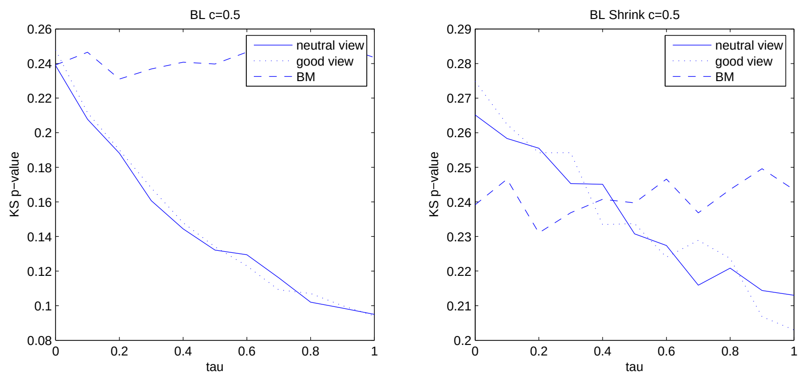

Figure 6: τ amplifies the weights on the view correction part

BL is more affected by view certainty than BLS due to the fact the BLS in our case has a smaller τ, which puts less weight on the view correction part. Reviewing BL formulas, we see that smaller τ yields smaller view correction and the prior market distribution dominates more, which should be having the same effect as varying the confidence coefficient c. Let us investigate this closer.

Figure 6 and 7 consists of a τ-KS score plot and a c-KS score plot where we average KS score for all five assets between day 300-500 with $q = r$. For larger τ BL and BLS scores both decreased which means that the distributions are becoming more unlikely to match the realized market distributions. If the view correction part generated with $c = 0.5$ is not satisfying increasing τ will only amplify the un-satisfying effect and lead to worse matching. Note also that by construction BLS uses the same τ as in BL but scaled with $(1 - \omega_s)^2$, and thus yields more stable KS scores.

Increasing confidence generally yields better degrees of distributional matching provided that the views are sophisticated. Since covariances grow closer to the market prior covariances more correct expectations intuitively leads to better distributional matching. Here we see that τ and c together controls the weight to be put on the views. It thus makes sense to do as in *Meucci (2006)* to eliminate the use of τ and re-calibrate the confidence levels c. Or we can consider $\tau\Sigma^{-1}$ as a whole since it always appears as a whole in BL formulas.

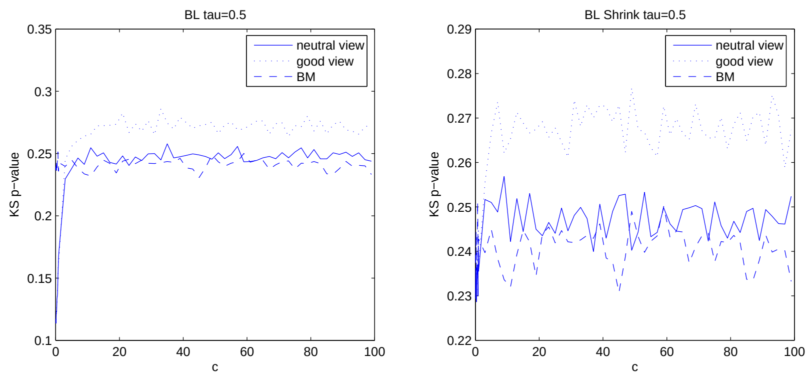

Figure 7: c calibrates the view correction part to make it more or less satisfactory

5.1.4 Single-period summaries

Blending market returns with investor views do help on normal market conditions to achieve better prediction for future returns. If views are certain and sophisticated, COP or Meucci-BL outperforms the others. If views are not sophisticated enough but we have cross-assets porfolio BLS is the all-round winner, with good distribution match and stable mean-error. View-sensitivities could be customized by increasing values of

τ, the weight-on-view coefficient, while view-correctness could be calibrated by varying the confidence levels of the views.

The original COP covariances is only comparable to Meucci's BL since they both express posterior distribution on market log-return distributions directly.

5.2 Multi-period test

The section's result will give insights on whether single period results could be generalized. We will use the methods used in the single period tests. In general Meucci's BL performs quite similar to that of COP and we will not show Meucci's BL separately. To compensate for small τ values we now set τ=0.5 for BL to allow more visible difference between BL and BM.

We consider the entire test period. For every test a reference period of 200 days and a prediction window of 60 days will be used. We start from using day $1-200$ data to predict day $200-260$ returns, and $60-260$ to predict $260-320$ and vice versa.

5.2.1 Mean-error test

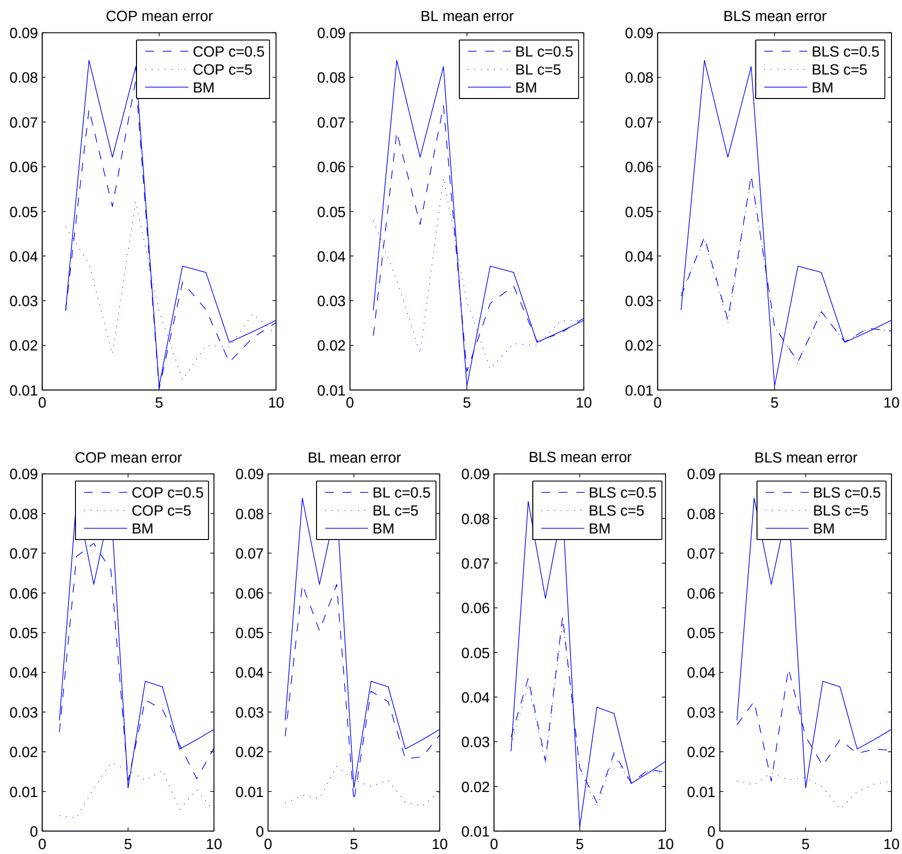

Figure 8: Multi-period 60-day mean-errors using $q = r_f$ (above) and $q = r$ (below). The last two BLS plots are generated with $\tau = \frac{(1-\omega_s)^2}{N}$ and $\tau = 0.5$ respectively. τ is set to 0.5 for BL.

The mean-error fluctuates with market conditions. BL now performs similar to COP. r_f is not very good apart from the first four periods when the interest rate is high (falling) and equity returns are negative or close to zero in the aftershock of the financial crisis. After period four interest rates lingers close to zero and equity market boomed in 2009. As expected during 2009 and 2010 the model performance results with $q = r_f$ is quite disappointing.

BLS is not confidence-sensitive in our tests, but yields stable results. The stability comes from the shrinkage estimator which is designed for MSE-minimization and the confidence stability comes from the small weigh-on-view coefficient τ. In Figure 8 an increased τ made a difference and is able to decrease the mean-errors a notch.

The results seem to be consistent with single-period tests with COP and BL leading in situations with high-confident and excellent views while BLS is best with neutral views under normal market conditions.

5.2.2 Variance test

The plots are generated by averaging the χ^2 scores for all assets and roll the test over all periods. We also keep using the original Σ_{COP}, which leads to the expectation that variances will tend to 0 with higher confidence levels and χ^2 test scores will be on the upper extremes.

In Figure 9 COP variances peak for high confidence levels indicates that Σ_{COP} grows close to 0, exactly as expected. Both BL and BLS comes close to benchmark, particularly BLS as shown in Figure 9. The view does not have an huge effect on the variance estimates. According to the tests none of the methods are particularly good at estimating variances since the χ^2 scores fail the test most of the time. By looking at the log-return plot we realize that we need less than 200 historical days to achieve more accurate variance estimates.

Figure 10 uses historical 60 days to estimate future 60 days. The results is blended. We have both higher peaks and more test scores within the confidence interval. In times where market volatilities suddenly increase or decrease the variance tests fail and we get the two tops and one dip in the plots. It seems to be generating slightly more stable test scores compared to Figure 9 as the tops are now around 110 compared to 120 before and the dip is a bit above 20 as opposed to well below 20. This confirms the rule of the thumb that we should use shorter historical days to make predictions on future volatilities.

5.2.3 Distribution test

As in the single-period case when it comes to neutral views we should not have a too much weight-on-views, using the original setup we obtain the following plot which is

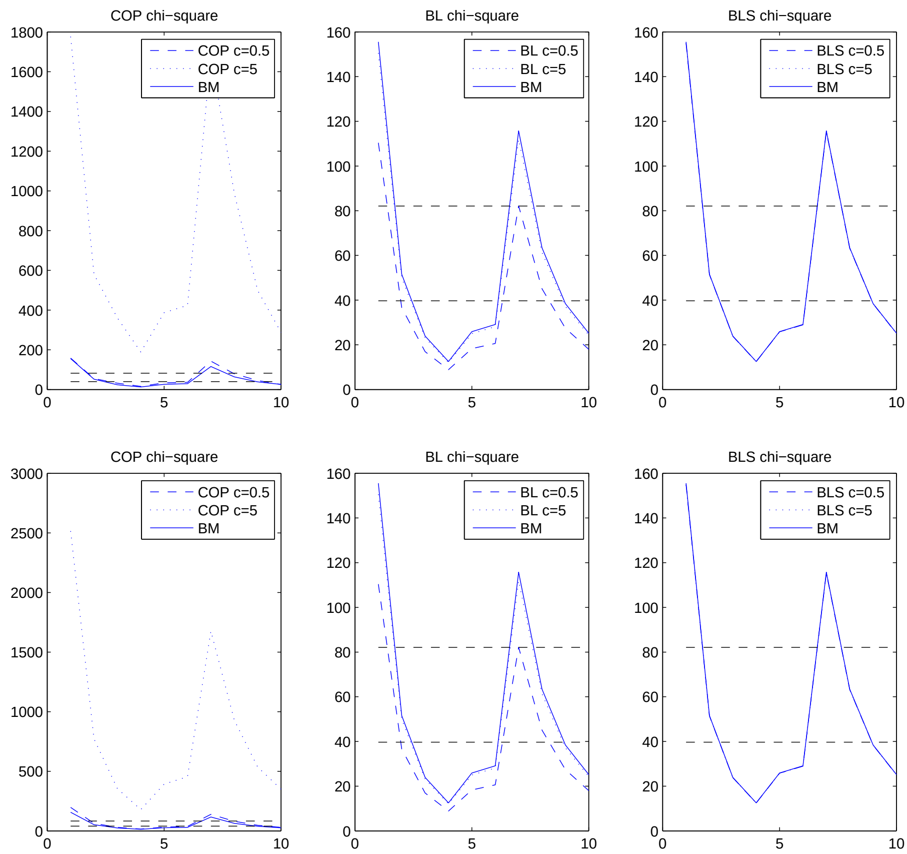

Figure 9: Multi-period 500-average χ^2 test scores using $q = r_f$ (above) and $q = r$ (below).

consistent with single period results, that is, COP variances is not comparable with the rest as it is so close to zero that it is barely visible. BL and BLS comes close to market prior distribution due to its small values of τ.

With good view we can afford to have a larger τ. With $\tau = 0.5$ we can in general achieve better results when views are qualified and certain. Figure 12 is generated with $c = 0.5$ and $c = 100$ respectively. The gain in KS score is around 3 percent(from 24 to 27 percent) in the single-period case according to Figure 6 and 7. In the multi-period case the trend mainly follows the market prior distributions with slight gains, particularly in the first 3 periods where the gains is clearly visible. In general it is hard to be confident of the views to such an extent($c = 100$) to motivate the slight gains in distribution matching. Another interesting point here is that COP actually managed to beat BM by an huge amount between period three to seven. During these 240 days the reference period volatility is very high and the prediction period volatility is low. The original

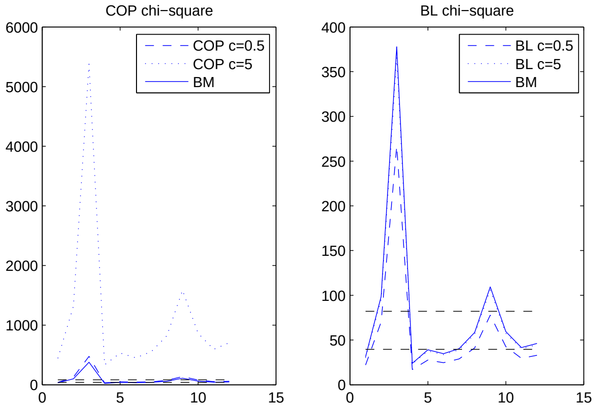

Figure 10: Multi-period 500-average χ^2 with 60-day reference period and 60-day prediction period using $q = r$

COP decreases the volatilities drastically and it helped in this case to minimize the distribution differences.

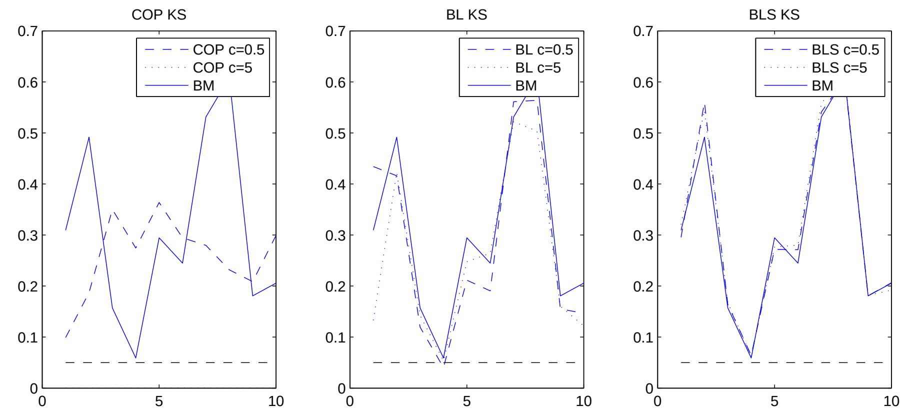

Figure 11: Multi-period 500-average KS-test using $q = r_f$ and the original τ

5.2.4 Multi-period summaries

Functionality of τ is to amplify the view correction part in the BL expressions. Both τ and c controls the weight put on the views and it makes sense to consider these two parameters together.

COP has very similar characteristics compared to BL with τ close to 1. This is why COP

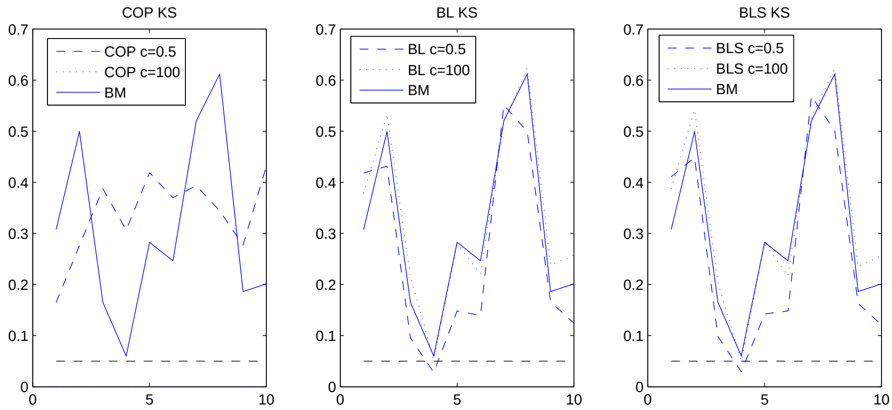

Figure 12: Multi-period 500-average KS-test using $q = r$ and $\tau = 0.5$

is very sensitive to confidence level changes, just as Meucci's BL. Also the covariances of COP goes to zero with extreme confident views which can be corrected by adding the market prior covariances to it. However addition could not be motivated by theories as it double-counts the covariance with low confidence levels.

BLS is still overall the most stable thanks to the cross-assets characteristics of the portfolio. It minimized the ME very quickly even under neutral views with neutral confidences. It is also more stable reacting to confidence level and weight-on-view parameter changes due to the smaller τ used.

Blending market log-return distributions with sophisticated views consistently improves the mean prediction throughout all periods. The cost we pay is the estimation of the extra parameters τ and Ω. While we have given an interpretation of τ in this section we will try to find a systematic way of determining Ω in the following chapter.

6 Determining the uncertainty matrix Ω

One difficult parameter to set in view-combining models is view confidences. One might know what views one have but it is always hard to know to what degree one is ready to take the bet. For instance if one is certain of her views she might want to set $c >= 1$ in our case, but c is defined for $(0, \infty)$ and what should she do with an this handful of choices of c:s?

This section will provide a concept to calibrate and estimate future confidences with help from historical view certainties. It starts with a description of the idea and concludes with a series of mean-error tests. We will choose to focus on mean-error tests since we have shown before that distributions or variances are not sensitive to view confidence variations unless they become extreme.

6.1 The model

We continue to assume Ω to be a diagonal matrix and views to be independent normally distributed with $N(q, \Omega)$. Following the same convention used in previous sections we have

$$\Omega = \begin{pmatrix} \omega_1 & 0 & \cdots \\ 0 & \ddots & \cdots \\ \cdots & \cdots & \omega_5 \end{pmatrix} \quad \omega_i = \frac{1}{c^2} P_i' \Sigma P_i \quad c \in (0, \infty)$$

For all past allocations we can estimate market realized expectation μ_{real} and covariances Σ_{real} directly from the market outcome. In best case scenarios the output from the model should match μ_{real} and Σ_{real} perfectly. In reality they could never match and we do best by minimizing the quadratic differences. For instance in the case of BL We have three problems ready to be solved depending on our priorities, with μ_{BL} and Σ_{BL} as functions of c.

Mean-error prioritized optimization We minimize the 60-day mean-error using r_{real}, the 60-day market realized return.

$$\min_c \frac{|60\mu_{BL} - r_{real}|}{dim(\mu_B L)}$$
$$r_{real} \approx 60\mu_{real}$$

Mean-Covariance prioritized optimization We minimize the sum of daily mean-squared- and covariance error.

$$\min_c |\mu_{BL} - \mu_{real}|^2 + |\Sigma_{BL} - \Sigma_{real}|$$

single-period during day 1-200 with $q = r_f$

asset	q	benchmark return	μ_{real}	optimized confidence
1	9.4×10^{-5}	-0.0024	-0.0007	
2	9.4×10^{-5}	-0.0023	-0.0012	
3	9.4×10^{-5}	-0.0026	-0.0022	$c = 0.5342$, $me = 0.0221$
4	9.4×10^{-5}	-0.0022	-0.0029	
5	9.4×10^{-5}	-0.0002	-0.0009	

single-period during day 300-500 with $q = r_f$

asset	q	benchmark return	μ_{real}	optimized confidence
1	5×10^{-6}	0.0018	0.0011	
2	5×10^{-6}	0.0020	0.0006	
3	5×10^{-6}	0.0021	-0.0002	$c = 2.4710$, $me = 0.0136$
4	5×10^{-6}	0.0020	0.0003	
5	5×10^{-6}	0.0001	0.0002	

Table 11: optimized confidence using BL with $\tau = 0.5$ and neutral guess on abnormal(above) and normal(below) markets

In this section we will focus on mean-error optimizations and assume an overall confidence level $c_i = c$, $i = 1, 2, 3, \cdots$. This reduces the optimization problem to one dimension and achieves much better stability in the solutions. It is possible to construct an optimizer that optimizes confidence levels on multi-dimensional basis since the construction of Ω allows different confidences for different assets. However the results turned out to be largely affected by numerical truncation errors when the derivatives of the expressions become very flat. Much smaller tolerance levels are required. Also the results are harder to interpret than the one-dimensional case but the two are in general consistent with each other.

6.2 The results

We focus now on how much better mean-error scores we can achieve by using the new method of calibrating Ω. The tests consist of two parts, one for single-period and another for multi-periods. A reference period of 200 days and prediction period of 60 will be used just as the tests in the Chapter V. To allow more sensitivities to the views we set $\tau = 0.5$ for both BL and BLS. All optimizations are based on the BL model but it is easy to modify based on COP or BLS as well. Also we only consider $c \in (0, 10)$ where 10 represents very high confidence in the views.

single-period during day 1-200 with $q = r_f$

asset	q	benchmark return	μ_{real}	optimized confidence
1	-0.0007	-0.0024	-0.0007	
2	-0.0023	-0.0023	-0.0012	
3	-0.0022	-0.0026	-0.0022	$c = 10$, $me = 0.0029$
4	-0.0029	-0.0022	-0.0029	
5	-0.0009	-0.0002	-0.0009	

single-period during day 300-500 with $q = r_f$

asset	q	benchmark return	μ_{real}	optimized confidence
1	0.0011	0.0018	0.0011	
2	0.0006	0.0020	0.0006	
3	-0.0002	0.0021	-0.0002	$c = 10$, $me = 0.0014$
4	0.0003	0.0020	0.0003	
5	0.0002	0.0001	0.0002	

Table 12: optimized confidence using BL with $\tau = 0.5$ and good guess on abnormal(above) and normal(below) markets

6.2.1 Single-period analysis

As we can see in Table 11 and 12 that r_f generally yields much lower confidence than r. This is an intuitive result since we will minimize the error if we maximize the confidence when views are correct. This is also the reason we have to set a boundary for the optimization since we always achieve minimum error when $c \to \infty$. All mean-errors obtained using optimized confidence level are lower than the best of the scores we obtained in single-period tests with $c = 0.5$ and $c = 5$. The method thus gives a very good numerical indication of how certain the views are. The larger the difference between the μ_{real} and q, the lower becomes the confidence level. In the $q = r_f$ case the interest rate is close to 0 at the same time as market index log-returns are very low(close to 0) between day 300 and 500, why we obtained higher confidence in the views.

6.2.2 Multi-period analysis

To see if the single-period results could be generalized we apply the same test on all periods.

Now we see clearly from Figure 13 that r_f is not a very sophisticated view. Its optimized confidence levels varies largely in the definition space of c. The confidence level first increased with lower interest rate and uncertain equity market(lower return) after the crisis. It plummeted during the equity boom(high return) in 2009 while interest rate stayed low, but only to bounce back up again during the European debt crisis with yet another uncertain equity market. The confidence level fluctuates with markets and there is not much consistency here.

Mean-error fluctuates as much as the confidence levels for $q = r_f$. For $q = r$ the implied confidence levels are maximized at all times and thus yields very low mean-error. This indicates that if we have a consistent view we can calibrate future view-confidences from a historical track record. If the record is consistently good we can afford to use higher certainty in the views to achieve better prediction for the future.

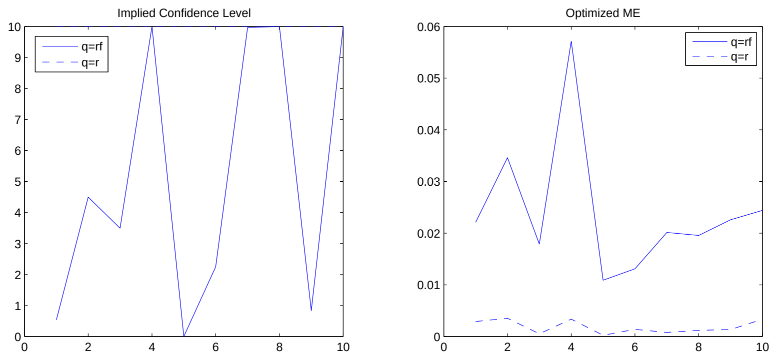

Figure 13: BL optimization with $\tau = 0.5$

6.3 Strategy back-testing

In this test we want to see if using historical confidence track-record has an effect on minimizing future returns. We will use $q = r_f$, $q = r$ and $q = r + noise$ to simulate expert views in reality.

From Figure 14 we see that one-period lag is much more significant than for the rest. It is thus argue-able to set

$$c_i = c_{i-1} \quad i = 2, 3, \cdots \tag{6.1}$$

As expected r_f is not sophisticated enough to allow consistency in the views. The views are relatively independent and historical view confidences does not have an huge impact on future view confidence levels. On the other hand historically tracked view confidences generally produce better results than both $c = 0.5$ and $c = 5$ with $q = r$. But this is quite expected since we have correct views and we obtain large c:s which yield smaller mean-error. This scenario is best-case and never happens in real-life.

Perhaps we should add a noise to $q = r$ to reflect real-world scenarios. To set the variance of the noise we consider the following: the VIX-index reflects annualized volatility on the *SPX500* and is usually around 20 to 30% (crisis in 2008 made it peak at around

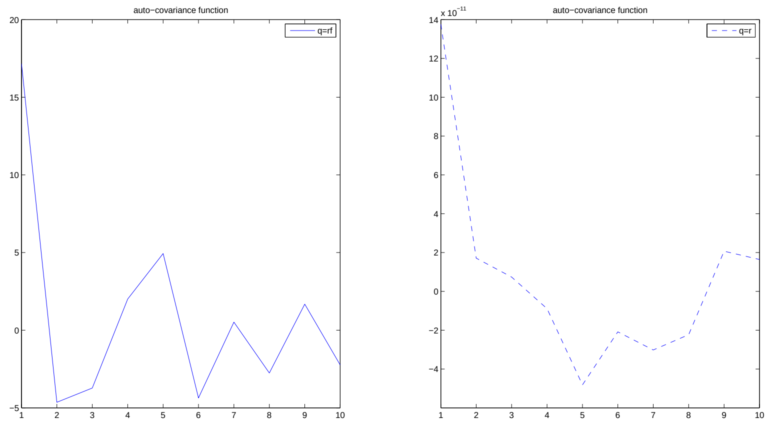

Figure 14: Auto-covariance plot for $q = r_f$ (left) and $q = r$ (right)

80%) which means an daily variance at $\frac{0.3\times0.3}{250} \approx 0.04$ percent. It thus makes sense to add a normal distributed noise to r with $N(0, 0.0004)$.

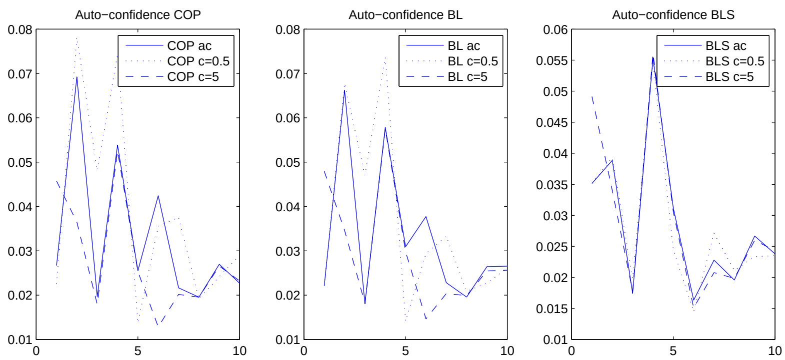

Figure 15: 60-day mean-error with $q = r_f$

We can also move towards the more extreme values, a daily noise with $N(0, 0.001)$ is equivalent to an annual volatility around 50%, which VIX achieved during the European debt crisis.

At 30% annualized volatility, by using historically tracked confidences the performance quickly become superior to $c = 0.5$ even if they shared the same start value. It also performs better than $c = 10$ in the sense that the mean-errors tend to be less volatile. At 50 percent annualized volatility, the mean-error become much more volatile and no clear winner can be named. The high correlation with $c = 10$ implies that we tend to

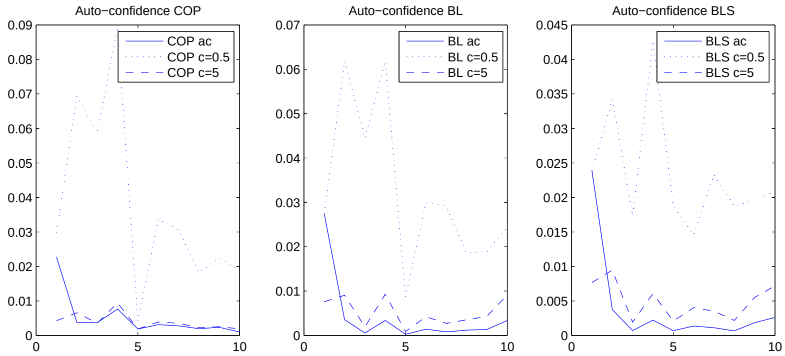

Figure 16: 60-day mean-error with $q = r$

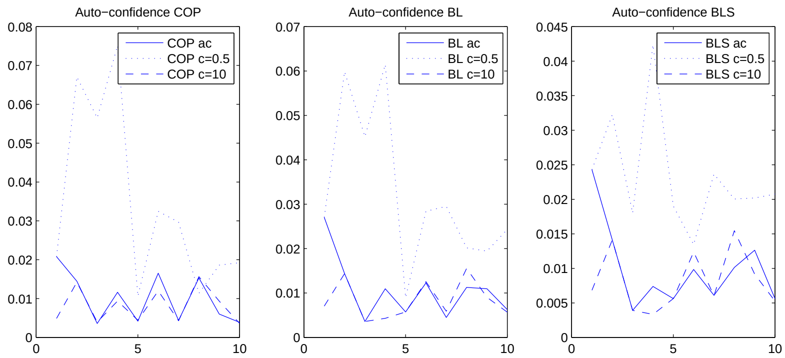

Figure 17: 60-day mean-error with $q = r$ plus 30 percent annualized volatility on all assets

be overconfident with historically tracked confidences.

While we could not use historical returns to predict futures, given sufficient consistency in the views we could use historical confidences on the views to improve future performances.

We should also note that it is possible to use time-series models with one or two periods of lag to achieve more sophisticated prediction for future confidence levels if we have access to vast set of historical data and historical views. But we should implement boundaries in the model to see to that c stay positive at all times.

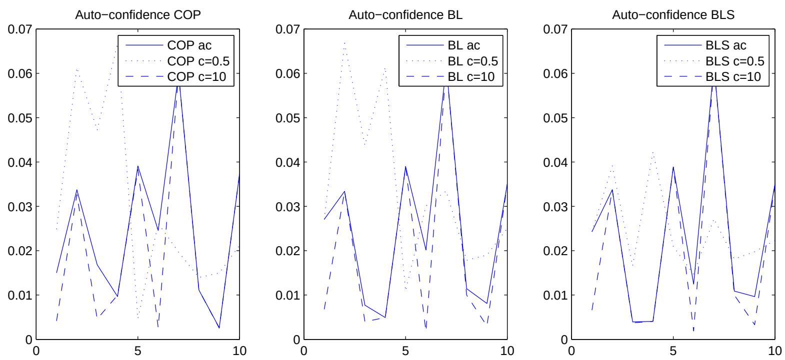

Figure 18: 60-day mean-error with $q = r$ plus 50 percent annualized volatility on all assets

7 Conclusions

There are two ways of improving the prediction quality of future market returns. One can either choose to use models that take advantage of the momentum in the data series and derive future values by extrapolation, or use additional information apart from historical data. Blending market returns with investor views belongs to the later category. It does help to improve the prediction quality of future market returns given sophisticated views. However with more information there are more parameters to be calibrated. The thesis has chosen several popular methods in the research frontier in this area and makes a head-to-head performance comparison between the methods in order to find out whether any or any one of the methods can be preferred over others.

Copula opinion pooling (COP), *Black-Litterman* (BL), *Black-Litterman with shrinkage estimators* (BLS) are chosen in the thesis. While BL opened our eyes by first introducing the concept, COP aims to do what BL does without the normality assumption on market return distributions. BLS is a refined version of the original BL with same sets of inputs adding features of the shrinkage estimator to obtain less mean-squared error in mean prediction. Back-testing are conducted where model results are compared to real market outcome in the sense of mean-error, volatility and return distribution predicting power as a whole.

During the back testing we notice that assigning appropriate values to parameters such as τ and Ω is difficult during application of BL and BLS. COP (and the modified version BL by *Meucci*) would be facing the same problems if not τ is explicitly set to 1 in the model by construction. τ controls the weight put on the views, Ω controls confidence levels in the views but they together gives the sensitivities of the model to the views. While calibrating them separately can be difficult we only need $\tau\Omega^{-1}$ as a whole. This suggests that from a practitioner's point of view we could fix τ and calibrate Ω for a given τ. The last chapter provides a way of calibrating the confidence levels on the views by extrapolating historical confidence levels. We show that it is appropriate to set a one or two period lag in the view confidence levels and compute future confidence levels from historical ones, given that the views are consistent. We simulate the views by combining correct market outcomes with normal distributed noises. With an annulized volatility less than 50% the simulations show that we do get better mean prediction if we calibrate the view confidences against its historical values.

Returning to the comparisons between COP, BL and BLS we benchmark all testing scores against the benchmark (BM) where means and covariances comes directly from the historical sample. We find out that COP is giving similar results as the modified BL by *Meucci* since they both share similar assumptions when it comes prior market return distributions. Both models react more to view corrections, resulting in sensitive mean-errors and variance predictions that are not quite comparable to the real market outcome of return distributions. BL and BLS both yield good predictions when it comes to volatility and return distribution prediction as a whole where the effect of τ or Ω does not come into play unless the views become extreme. BLS clearly dominates the BL if we have a multi-asset portfolio. The shrinkage estimator compares the historical returns with each other and scales down higher asset returns towards asset returns with lower volatilities. This slightly neutralizes the market and makes BLS the most stable method overall. All three methods demonstrate that blending market log-return distributions with sophisticated views consistently improves the mean prediction throughout all test periods.

All three methods produce interesting results while each has its own edge compared to others. The sensitivity on views makes COP combined with certain views maximize the benefits of the confident views. BLS and BLS are best to use when confidence on views are neutral under normal market conditions. However we only focus on mean and covariance in this thesis. The differences between portfolio performance after portfolio optimization with model results is left out. Also further research on systematic ways of determining τ and Ω would be interesting from a practitioner's point of view.

References

[1] Bevan, A., K. Winkelmann (1998), *Using the Black-Litterman Global Asset Allocation Model: Three Years of Practical Experience*, GS Fixed Income Research (September 1998)

[2] Black, F., R. Litterman (1990), *Asset Allocation: Combining Investors Views with Market*, GS Fixed Income Research (September 1990)

[3] Black, F., R. Litterman (1992), *Global Portfolio Optimization*, Financial Analysts Journal, September-October, pp. 28-43

[4] Efron, B., Morris, C. (1973), *Steins estimation rule and its competitors - an empirical Bayes approach*, Journal of the American Statistical Association, March 1973, Volume 68, Number 341

[5] Efron, B., Morris, C. (1975), *Data analysis using Steins estimator and its generalizations*, Journal of the American Statistical Association, March 1975, Volume 70, Number 350

[6] Hull, J.C. (2002). Options, Futures, and Other, Derivatives (5th ed.). New Jersey: Prentice Hall

[7] Hult, H., Lindskog, F. (2009). Lecture notes in Risk management, Royal Institute of Technology

[8] Idzorek, Tom, (2004), *A Step-by-Step Guide to the Black-Litterman Model*, Available on http://www.blacklitterman.org/papers.html

[9] James, W., Stein, C. (1961), *Estimation with quadratic loss*, Proceedings of the 4rd Berkeley Symposium on Probability and Statistics I (pp. 361-379). Berkeley: University of California Press.

[10] Jobson, J. D., Korkie, B., Ratti, V. (1979), *Improved estimation for Markowitz portfolios using James-Stein type estimators*, Proceedings of the American Statistical Association Business and Economics Statistics Section, 41, 279-284.

[11] Jobson, J. D., Korkie, B. (1980), *Estimation for Markowitz efficient portfolios*, Journal of the American Statistical Association, Volume 75, 544-554.

[12] Jobson, J. D., Korkie, B. (1981), *Putting Markowitz theory to work*, Journal of Portfolio Management, 7, 70-74

[13] Jorion, P. (1986), *Bayes-Stein estimation for portfolio analysis*, Journal of Financial and Quantitative Analysis, 21,279-292.

[14] Litterman, Robert (2003), *Beyond equilibirium, the Black-Litterman Approach*, Modern Investment Management- an equilibrium approach, New Jersey: Wiley pp 76-88

[15] Mankert.C, (2006), *The Black-Litterman Model - mathematical and behavioral finance approaches towards its use in practice*, Thesis at Royal Institute of Technology 2006

[16] Mazzoni, T. (2010), *Are short term stock asset returns predictable?*, Working paper: Discussion Material at FernUniversitat in Hagen

[17] Meucci, A. (2005), *Risk and Asset Allocation*, Springer Finance

[18] Meucci, A. (2006), *Beyond Black-Litterman: views on non-normal markets*, Risk Magazine 19, 87-92

[19] Meucci, A. (2006), *Beyond Black-Litterman in Practice: a five-step recipe to input views on non-normal markets*, available on http://www.symmys.com/

[20] Meucci, A. (2008), *The Black-Litterman Approach: Original Model and Extensions*, Working paper: The Encyclopedia of Quantitative Finance, Wiley, 2010

[21] Satchell, S., A. Scowcroft (2000) *A Demystification of the B-L model: Managing Quantitative and Tradtional Portfolio Construction* Journal of Asset Management 1(2), pp.138- 150.

[22] Stein, C. (1956), *Inadmissibility of the usual estimator for the mean of a multivariate normal distribution*, Proceedings of the 3rd Berkeley Symposium on Probability and Statistics I (pp. 197-206). Berkeley: University of California Press.

[23] Walters, J. (2010), *The factor tau in the Black-Litterman Model*, available on http://www.blacklitterman.org/

[24] Zellner, A. and V.K. Chetty (1965), *Prediction and decision problems in regression models from the Bayesian point of view*, Journal of the American Statistical Association, 60, 608-615.