Robust Bayesian Allocation

Attilio Meucci
attilio_meucci@symmys.com

Abstract

Using the Bayesian posterior distribution of the market parameters we define self-adjusting uncertainty regions, which take the investor's prior into account, for the robust mean-variance problem. Under the standard normal-inverse-Wishart conjugate assumption for the market, the ensuing robust Bayesian mean-variance efficient frontier simplifies to a parsimonious set. This set is parametrized by the exposure to overall risk, which includes market risk, estimation risk for the expected values and estimation risk for the covariances.

The classical approach to asset allocation is a two-step process: first the market distribution is estimated, then an optimization is performed, as if the estimated distribution were the true market distribution. Since this is not the case, the classical "optimal" allocation is not truly optimal. More importantly, since the optimization process is extremely sensitive to the input parameters, the sub-optimality due to estimation risk can be dramatic, see Jobson and Korkie (1980), Best and Grauer (1991), Chopra and Ziemba (1993).

Bayesian theory provides a way to limit the sensitivity of the final allocation to the input parameters by shrinking the estimate of the market parameters toward the investor's prior, see Bawa, Brown, and Klein (1979), Jorion (1986), Pastor and Stambaugh (2002). Similarly, the approach of Black and Litterman (1990) uses Bayes' rule to shrink the general equilibrium distribution of the market toward the investor's views.

The theory of robust optimization provides a different approach to dealing with estimation risk: the investor chooses the best allocation in the worst market within a given uncertainty range, see Goldfarb and Iyengar (2003), Halldorsson and Tutuncu (2003), Ceria and Stubbs (2004).

Robust allocations are guaranteed to perform adequately for all the markets within the given uncertainty range. Nevertheless, the choice of this range is quite arbitrary. Furthermore, the investor's prior knowledge, a key ingredient in any allocation decision, is not taken in consideration.

Using the Bayesian approach to estimation we can naturally identify a suitable uncertainty range for the market parameters, namely the location-dispersion ellipsoid of their posterior distribution. Robust Bayesian allocations are the solutions to a robust optimization problem that uses as uncertainty range the Bayesian location-dispersion ellipsoid. Similarly to robust allocations, these allocations account for estimation risk over a whole range of market parameters.

Similarly to Bayesian decisions, these allocations include the investor's prior knowledge in the optimization process within a sound and self-adjusting statistical framework.

Robust Bayesian allocations are discussed in De Santis and Foresi (2002), where the Bayesian setting is provided by the Black-Litterman posterior distribution. Here we consider robust Bayesian decisions that also account for the estimation error in the covariances and that explicitly process the information from the market, namely the observed time series of the past returns. As it turns out, the multi-parameter, non-conically constrained mean-variance optimization simplifies to a parsimonious efficient frontier that resembles the classical frontier, except that the classical parameterization in terms of the exposure to market risk becomes in this context a parameterization in terms of the exposure to both market risk and estimation risk.

In Section 1 we introduce the general robust Bayesian mean-variance approach to asset allocation. In Section 2 we compute explicitly the Bayesian location-dispersion ellipsoids for the market parameters under the standard normal-inverse-Wishart conjugate assumption for the market. In Section 3 we solve the ensuing mean-variance problem, computing the robust Bayesian frontier of efficient portfolios. Section 4 concludes.

1 General robust Bayesian setting

The classical mean-variance problem reads:

$$\mathbf{w}^{(i)} = \underset{\mathbf{w}}{\operatorname{argmax}}\, \mathbf{w}'\boldsymbol{\mu} \tag{1}$$
$$\text{subject to } \begin{cases} \mathbf{w} \in \mathcal{C} \\ \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \leq v^{(i)}. \end{cases}$$

In this expression $\mathbf{w}$ is the N-dimensional vector of relative portfolio weights; $\mathcal{C}$ is a set of investment constraints; the set $\left\{v^{(1)}, \ldots, v^{(I)}\right\}$ is a significative grid of target variances of the return on the portfolio, where the return at time t for a horizon τ of an asset that at the generic time t trades at the price P_t is defined as $R_{t,\tau} \equiv P_t/P_{t-\tau} - 1$; and $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ represent respectively the expected values and the covariances of the returns on the N securities in the market relative to the investment horizon:

$$\boldsymbol{\mu} \equiv \mathrm{E}\left\{\mathbf{R}_{T+\tau,\tau}\right\}, \quad \boldsymbol{\Sigma} \equiv \mathrm{Cov}\left\{\mathbf{R}_{T+\tau,\tau}\right\}. \tag{2}$$

The robust version of the mean-variance problem (1) reads:

$$\mathbf{w}^{(i)} = \underset{\mathbf{w}}{\operatorname{argmax}} \left\{ \min_{\boldsymbol{\mu} \in \widehat{\boldsymbol{\Theta}}_{\mu}} \left\{\mathbf{w}'\boldsymbol{\mu}\right\} \right\} \tag{3}$$
$$\text{subject to } \begin{cases} \mathbf{w} \in \mathcal{C} \\ \max_{\boldsymbol{\Sigma} \in \widehat{\boldsymbol{\Theta}}_{\boldsymbol{\Sigma}}} \left\{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}\right\} \leq v^{(i)}, \end{cases}$$

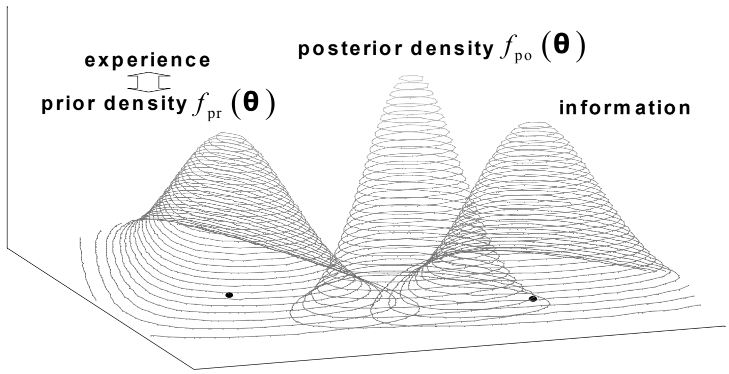

Figure 1: Bayesian approach to parameter estimation

where $\widehat{\boldsymbol{\Theta}}_{\boldsymbol{\mu}}$ and $\widehat{\boldsymbol{\Theta}}_{\boldsymbol{\Sigma}}$ are suitable uncertainty regions for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ respectively.

The Bayesian framework defines uncertainty sets in a natural way. Indeed, in the Bayesian framework the unknown generic market parameters $\boldsymbol{\theta}$ (e.g. $\boldsymbol{\mu}$ or $\boldsymbol{\Sigma}$) are random variables. The likelihood that the parameters assume given values is described by the posterior probability density function $f_{\text{po}}(\boldsymbol{\theta})$, which is determined by the information available at the time the allocation takes place and by the investor's experience and respective confidence, modeled in terms of a prior probability density function $f_{\text{pr}}(\boldsymbol{\theta})$, see Figure 1.

The region where the posterior distribution displays a higher concentration deserves more attention than the tails of the distribution: this region is a natural choice for the uncertainty set.

Consider first $\boldsymbol{\mu}$. The region where the posterior distribution displays a higher concentration is represented by the location-dispersion ellipsoid of the marginal posterior distribution of $\boldsymbol{\mu}$, see Figure 2 for the case $N \equiv 2$:

$$\widehat{\boldsymbol{\Theta}}_{\boldsymbol{\mu}} \equiv \left\{ \boldsymbol{\mu} : \ (\boldsymbol{\mu} - \widehat{\boldsymbol{\mu}}_{\text{ce}})' \, \mathbf{S}_{\boldsymbol{\mu}}^{-1} (\boldsymbol{\mu} - \widehat{\boldsymbol{\mu}}_{\text{ce}}) \leq q_{\boldsymbol{\mu}}^2 \right\}. \tag{4}$$

In this expression $q_{\boldsymbol{\mu}}$ is the radius factor for the ellipsoid; $\widehat{\boldsymbol{\mu}}_{\text{ce}}$ is a classical-equivalent estimator such as the expected value or the mode of the marginal posterior distribution of $\boldsymbol{\mu}$; and $\mathbf{S}_{\boldsymbol{\mu}}$ is a scatter matrix such as the covariance matrix or the modal dispersion of the marginal posterior distribution of $\boldsymbol{\mu}$.

Similarly, consider $\boldsymbol{\Sigma}$. The region where the posterior distribution displays a higher concentration is represented by the location-dispersion ellipsoid of the marginal posterior distribution of $\boldsymbol{\Sigma}$:

$$\widehat{\boldsymbol{\Theta}}_{\boldsymbol{\Sigma}} \equiv \left\{ \boldsymbol{\Sigma} : \ \text{vech}\left[\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}_{\text{ce}}\right]' \mathbf{S}_{\boldsymbol{\Sigma}}^{-1} \, \text{vech}\left[\boldsymbol{\Sigma} - \widehat{\boldsymbol{\Sigma}}_{\text{ce}}\right] \leq q_{\boldsymbol{\Sigma}}^2 \right\}. \tag{5}$$

In this expression vech is the operator that stacks the columns of a matrix

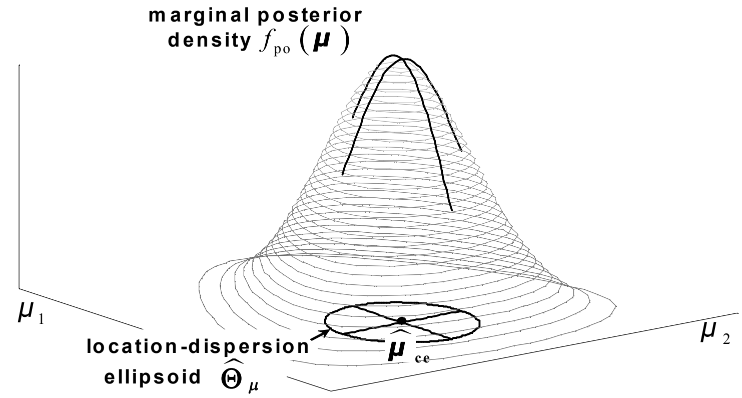

Figure 2: Bayesian posterior distribution and uncertainty set

skipping the redundant entries above the diagonal; $q_{\boldsymbol{\Sigma}}$ is the radius factor for the ellipsoid; $\widehat{\boldsymbol{\Sigma}}_{\mathrm{ce}}$ is a classical-equivalent estimator such as the expected value or the mode of the marginal posterior distribution of $\boldsymbol{\Sigma}$; and $\mathbf{S}_{\boldsymbol{\mu}}$ is a scatter matrix such as the covariance matrix or the modal dispersion of the marginal posterior distribution of $\mathrm{vech}\,[\boldsymbol{\Sigma}]$. The matrices $\boldsymbol{\Sigma}$ in this ellipsoid are always symmetric, because the vech operator only spans the non-redundant elements of a matrix. When the radius factor $q_{\boldsymbol{\Sigma}}$ is not too large the matrices $\boldsymbol{\Sigma}$ in this ellipsoid are also positive definite: indeed, positivity is a continuous property and $\widehat{\boldsymbol{\Sigma}}_{\mathrm{ce}}$ is positive definite, see Figure 3 for the case $N \equiv 2$, which is completely determined by three entries.

The robust Bayesian mean-variance approach to allocation consists in using the Bayesian elliptical uncertainty sets (4) and (5) in the robust mean-variance allocation problem (3). Notice that this problem is parametrized by the radius factor $q_{\boldsymbol{\mu}}$, the radius factor $q_{\boldsymbol{\Sigma}}$ and the target variance $v^{(i)}$. The term $q_{\boldsymbol{\mu}}$ represents aversion to estimation risk for the expected values $\boldsymbol{\mu}$: considering a larger ellipsoid increases the chances of including the true, unknown value of $\boldsymbol{\mu}$ within the ellipsoid. Similarly, the term $q_{\boldsymbol{\Sigma}}$ represents aversion to estimation risk for the covariances $\boldsymbol{\Sigma}$: considering a larger ellipsoid increases the chances of including the true, unknown value of $\boldsymbol{\Sigma}$ within the ellipsoid. On the other hand, the term $v^{(i)}$ represents exposure to market risk. As a result, in principle, the robust Bayesian mean-variance efficient frontier should constitute a three-dimensional surface in the N-dimensional space of the allocations.

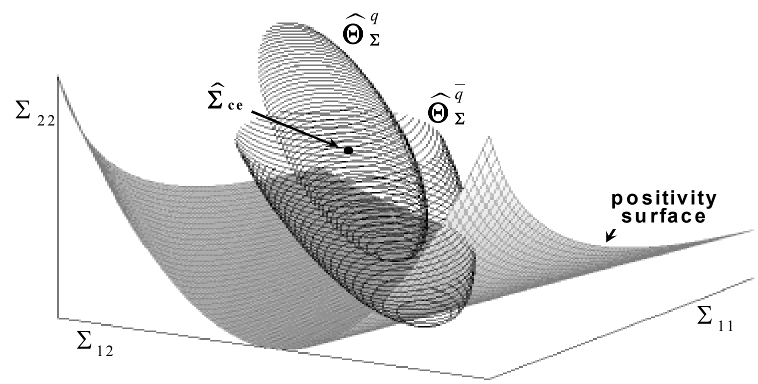

Figure 3: Bayesian location-dispersion ellipsoid for covariance

2 Specification of a market model

In order to determine the Bayesian elliptical uncertainty sets (4) and (5) we need to compute the posterior distributions of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. To this purpose, we make the following assumptions: first, the market consists of equity-like securities for which the returns are independently and identically distributed across time; second, the estimation interval is the same as the investment horizon; third, the returns are normally distributed:

$$\mathbf{R}_{t,\tau}|\boldsymbol{\mu},\boldsymbol{\Sigma} \sim \mathrm{N}\left(\boldsymbol{\mu},\boldsymbol{\Sigma}\right). \tag{6}$$

Furthermore, we model the investor's prior experience as a normal-inverse-Wishart distribution:

$$\boldsymbol{\mu}|\boldsymbol{\Sigma} \sim \mathrm{N}\left(\boldsymbol{\mu}_0, \frac{\boldsymbol{\Sigma}}{T_0}\right), \quad \boldsymbol{\Sigma}^{-1} \sim \mathrm{W}\left(\nu_0, \frac{\boldsymbol{\Sigma}_0^{-1}}{\nu_0}\right). \tag{7}$$

In this expression $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ represent the investor's experience on the parameters, whereas (T_0, ν_0) represent the respective confidence.

Under the above hypotheses it is possible to compute the posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ analytically, see Aitchison and Dunsmore (1975). First of all, the information from the market is summarized in the sample mean and the sample covariance of the past realizations of the returns:

$$\widehat{\boldsymbol{\mu}} \equiv \frac{1}{T}\sum_{t=1}^{T} \mathbf{r}_{t,\tau}, \quad \widehat{\boldsymbol{\Sigma}} \equiv \frac{1}{T}\sum_{t=1}^{T} \left(\mathbf{r}_{t,\tau} - \widehat{\boldsymbol{\mu}}\right)\left(\mathbf{r}_{t,\tau} - \widehat{\boldsymbol{\mu}}\right). \tag{8}$$

The posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, like the prior distribution (7), is also normal-inverse-Wishart, where the respective parameters read:

$$T_1 \equiv T_0 + T \tag{9}$$

$$\boldsymbol{\mu}_1 \equiv \frac{1}{T_1}\left[T_0\boldsymbol{\mu}_0 + T\widehat{\boldsymbol{\mu}}\right] \tag{10}$$

$$\nu_1 \equiv \nu_0 + T \tag{11}$$

$$\boldsymbol{\Sigma}_1 \equiv \frac{1}{\nu_1}\left[\nu_0\boldsymbol{\Sigma}_0 + T\widehat{\boldsymbol{\Sigma}} + \frac{(\boldsymbol{\mu}_0 - \widehat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \widehat{\boldsymbol{\mu}})'}{\frac{1}{T} + \frac{1}{T_0}}\right]. \tag{12}$$

The marginal posterior distribution of $\boldsymbol{\mu}$ is multivariate Student t. From its expression it is possible to compute explicitly the classical-equivalent estimator and the scatter matrix that appear in the location-dispersion ellipsoid (4), see Meucci (2005):

$$\widehat{\boldsymbol{\mu}}_{\mathrm{ce}} = \boldsymbol{\mu}_1 \tag{13}$$

$$\mathbf{S}_{\boldsymbol{\mu}} = \frac{1}{T_1}\frac{\nu_1}{\nu_1 - 2}\boldsymbol{\Sigma}_1. \tag{14}$$

It is also possible to compute explicitly the classical-equivalent estimator and the scatter matrix of the inverse-Wishart marginal posterior distribution of $\boldsymbol{\Sigma}$ that appear in the location-dispersion ellipsoid (5), see Meucci (2005):

$$\widehat{\boldsymbol{\Sigma}}_{\mathrm{ce}} = \frac{\nu_1}{\nu_1 + N + 1}\boldsymbol{\Sigma}_1 \tag{15}$$

$$\mathbf{S}_{\boldsymbol{\Sigma}} = \frac{2\nu_1^2}{(\nu_1 + N + 1)^3}\left(\mathbf{D}_N'\left(\boldsymbol{\Sigma}_1^{-1} \otimes \boldsymbol{\Sigma}_1^{-1}\right)\mathbf{D}_N\right)^{-1}. \tag{16}$$

In this expression $\mathbf{D}_N$ is the duplication matrix that reinstates the redundant entries above the diagonal of a symmetric matrix, see Magnus and Neudecker (1999), and $\otimes$ is the Kronecker product.

3 Computation of robust Bayesian efficient allocations

The robust Bayesian mean-variance problem, i.e. the robust problem (3) where the Bayesian elliptical uncertainty sets (4) and (5) are specified by (13)-(16), simplifies as follows:

$$\mathbf{w}_{\mathrm{rB}}^{(i)} = \underset{\mathbf{w}}{\operatorname{argmax}}\left\{\mathbf{w}'\boldsymbol{\mu}_1 - \gamma_{\mu}\sqrt{\mathbf{w}'\boldsymbol{\Sigma}_1\mathbf{w}}\right\} \tag{17}$$

$$\text{subject to} \left\{\begin{array}{l} \mathbf{w} \in \mathcal{C} \\ \mathbf{w}'\boldsymbol{\Sigma}_1\mathbf{w} \leq \gamma_{\Sigma}^{(i)}, \end{array}\right.$$

where:

$$\gamma_\mu \equiv \sqrt{\frac{q_\mu^2}{T_1}\frac{\nu_1}{\nu_1 - 2}} \tag{18}$$

$$\gamma_\Sigma^{(i)} \equiv \frac{v^{(i)}}{\frac{\nu_1}{\nu_1+N+1} + \sqrt{\frac{2\nu_1^2 q_\Sigma^2}{(\nu_1+N+1)^3}}}. \tag{19}$$

We prove the above result in a technical appendix available upon request. As we show in that appendix, under standard regularity assumption for the investment constraints $\mathcal{C}$ the maximization (17) can be can be cast in the form of a second-order cone programming problem. Therefore the robust Bayesian frontier can be computed numerically, see Ben-Tal and Nemirovski (2001).

The efficient allocations (17) can be parametrized equivalently in terms of one single positive multiplier λ as follows:

$$\mathbf{w}(\lambda) = \underset{\mathbf{w}\in\mathcal{C}}{\operatorname{argmax}} \left\{ \mathbf{w}'\boldsymbol{\mu}_1 - \lambda\sqrt{\mathbf{w}'\boldsymbol{\Sigma}_1\mathbf{w}} \right\}. \tag{20}$$

The multiplier λ is determined by the scalars (18) and (19). It is easy to check that the value of λ is directly related to the aversion to estimation risk for $\boldsymbol{\mu}$, namely $q_{\boldsymbol{\mu}}$, and to the aversion to estimation risk for $\boldsymbol{\Sigma}$, namely $q_{\boldsymbol{\Sigma}}$, and inversely related to the exposure to market risk $v^{(i)}$. Accordingly, the term under the square root in (20) represents both estimation risk and market risk and the coefficient λ represents aversion to both types of risk.

In other words, the a-priori three-dimensional robust Bayesian efficient frontier collapses to a line. Hence the robust Bayesian mean-variance efficient frontier is conceptually similar to, and just as parsimonious as, the classical mean-variance efficient frontier (1). Nevertheless, in the classical setting "risk" only refers to market risk, whereas in the robust Bayesian setting "risk" blends both market risk and estimation risk.

Also notice the self-adjusting nature of the Bayesian setting. From (9)-(12) the expected values $\boldsymbol{\mu}_1$ and the covariance matrix $\boldsymbol{\Sigma}_1$ that determine the efficient allocations (20) are mixtures of the investor's prior knowledge $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and of information from the market $\left(\widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\Sigma}}\right)$: the balance between these components is steered by the relative weight of the confidence in the investor's prior parameters, represented respectively by (T_0, ν_0), and the amount of information, represented by number of observations T in the time series.

In particular, when the number of observations T is large with respect to the confidence levels T_0 and ν_0 in the investor's prior, the expected values $\boldsymbol{\mu}_1$ tend to the sample mean $\widehat{\boldsymbol{\mu}}$ and the covariance matrix $\boldsymbol{\Sigma}_1$ tends to the sample covariance $\widehat{\boldsymbol{\Sigma}}$. Therefore we obtain a sample-based efficient frontier:

$$\mathbf{w}(\lambda) = \underset{\mathbf{w}\in\mathcal{C}}{\operatorname{argmax}} \left\{ \mathbf{w}'\widehat{\boldsymbol{\mu}} - \lambda\sqrt{\mathbf{w}'\widehat{\boldsymbol{\Sigma}}\mathbf{w}} \right\}. \tag{21}$$

Similarly, when the confidence levels T_0 and ν_0 in the investor's prior are large with respect to the number of observations T, the expected values $\boldsymbol{\mu}_1$ tend to the prior $\boldsymbol{\mu}_0$ and the covariance matrix $\boldsymbol{\Sigma}_1$ tends to the prior $\boldsymbol{\Sigma}_0$. Therefore we obtain a prior efficient frontier that disregards any information from the market:

$$\mathbf{w}(\lambda) = \underset{\mathbf{w}\in\mathcal{C}}{\operatorname{argmax}} \left\{ \mathbf{w}'\boldsymbol{\mu}_0 - \lambda\sqrt{\mathbf{w}'\boldsymbol{\Sigma}_0\mathbf{w}} \right\}. \tag{22}$$

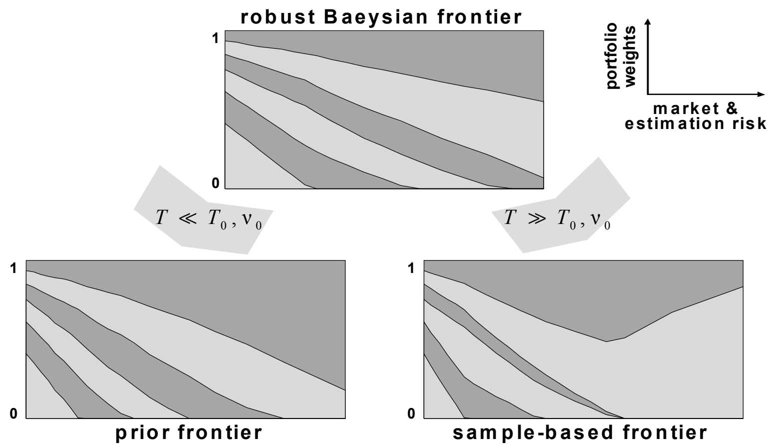

Figure 4: Robust Bayesian mean-variance efficient allocations

To illustrate, we consider a market of $N \equiv 6$ stocks from the utilities sector of the S&P 500. We estimate the sample mean and the sample covariance from a database of weekly returns. To specify the prior, we use the estimated sample standard deviations $\widehat{\boldsymbol{\sigma}}$: to determine the prior expected values we assume an equilibrium argument $\boldsymbol{\mu}_0 \equiv \rho\widehat{\boldsymbol{\sigma}}$, where we set the risk premium $\rho \approx 0.3$; to determine the prior covariances we set all the cross-correlations equal to 0.5.

Assume that the investor is bound by the standard budget constraint $\mathbf{w}'\mathbf{1} = 1$ and no-short-sale constraint $\mathbf{w} \geq \mathbf{0}$: in Figure 4 we plot the general robust Bayesian efficient frontier (20) and the limit cases (21) and (22).

4 Conclusions

Conceptually, the robust Bayesian approach to allocation displays the optimality features of robust optimization as well as the self-adjusting Bayesian mechanism that accounts for the investor's prior knowledge within a sound statistical framework.

Under fairly standard assumptions for the market, the robust Bayesian mean-variance efficient frontier reduces to a parsimonious set, namely a line, as in the classical mean-variance approach. Unlike in the classical setting, where the efficient frontier is parametrized by the exposure to market-risk, in the robust Bayesian setting the efficient frontier is parameterized by the exposure to overall risk, which blends market risk, estimation risk for the expected values and estimation risk for the covariances.

As customary in the Bayesian framework, the robust Bayesian frontier is a mix of a purely sample-based frontier, which completely disregards the investor's prior knowledge, and a purely prior frontier, which completely disregards the information from the market. The interplay between these two components is steered by the relative weight of the confidence in the prior with respect to the amount of information from the market.

References

Aitchison, J., and I. R. Dunsmore, 1975, *Statistical Prediction Analysis* (Cambridge University Press).

Bawa, V. S., S. J. Brown, and R. W. Klein, 1979, *Estimation Risk and Optimal Porfolio Choice* (North Holland).

Ben-Tal, A., and A. Nemirovski, 2001, *Lectures on modern convex optimization: analysis, algorithms, and engineering applications* (Society for Industrial and Applied Mathematics).

Best, M. J., and R. R. Grauer, 1991, On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results, *Review of Financial Studies* 4, 315–342.

Black, F., and R. Litterman, 1990, Asset allocation: combining investor views with market equilibrium, *Goldman Sachs Fixed Income Research*.

Ceria, S., and R. A. Stubbs, 2004, Incorporating estimation errors into portfolio selection: Robust efficient frontiers, *Axioma Inc. Technical Report*.

Chopra, V., and W. T. Ziemba, 1993, The effects of errors in means, variances, and covariances on optimal portfolio choice, *Journal of Portfolio Management* pp. 6–11.

De Santis, G., and S. Foresi, 2002, Robust optimization, *Goldman Sachs Technical Report*.

Goldfarb, D., and G. Iyengar, 2003, Robust portfolio selection problems, *Mathematics of Operations Research* 28, 1–38.

Halldorsson, B. V., and R. H. Tutuncu, 2003, An interior-point method for a class of saddle-point problems, *Journal of Optimization Theory and Applications* 116, 559–590.

Jobson, J. D., and B. Korkie, 1980, Estimation for Markowitz efficient portfolios, *Journal of the American Statistical Association* 75, 544–554.

Jorion, P., 1986, Bayes-Stein estimation for portfolio analysis, *Journal of Financial and Quantitative Analysis* 21, 279–291.

Magnus, J. R., and H. Neudecker, 1999, *Matrix Differential Calculus with Applications in Statistics and Econometrics, Revised Edition* (Wiley).

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

Pastor, L., and R. F. Stambaugh, 2002, Investing in equity mutual funds, *Journal of Financial Economics* 63, 351–380.