Bayesian Diagnostics for Portfolio Allocation

Francesco Corielli
Dept. of Quantitative Methods
Bocconi University
viale Isonzo 23, 20135 Milan, Italy
francesco.corielli@uni-bocconi.it

Attilio Meucci
Relative Value International
Greenwich, CT
06830, USA
ameucci@rvifunds.com

Abstract

We suggest a Bayesian approach to portfolio selection to study asset allocation robustness. A fund manager faces two major problems in real-world applications of portfolio theory. First: high correlation across different securities in a portfolio often implies that the investor can draw almost the same utility from completely different asset allocations. Second: estimation problems can severely degrade the investor's ability in evaluating the risk return properties of selected portfolios. The combined effects of these two problems makes it often safer to keep a slightly sub-optimal, more robust portfolio instead of investing in the "optimal" solution. This phenomenon is of the utmost practical importance for a portfolio manager who should decide whether or not to re-allocate her holdings. We propose a set of simple utility-based diagnostics to assist the manager's decisions and we implement these diagnostics in two empirical examples: the allocation of a global mutual fund and the computation of the optimal hedge ratio in a relative value portfolio.

Mathematics Subject Classification (2000): 91b28, 91b82
Journal of Economic Literature Classification: g11, c52
Keywords: robustness, utility diagnostics, hedge factor

1 Introduction

Fund managers that allocate their portfolios a-la-Markowitz face the following problem: is it worth incurring high transaction costs to optimize a sub-optimal allocation? The answer depends on the sensitivity of the allocation to the input parameters, but even more on the sensitivity of the objective function, i.e., the final utility, to those parameters.

There are three possible scenarios. In a first case, the final utility is very sensitive to reasonable variations of the input data: in this case no sensible suggestion can be given on the basis of available data (Michaud, 1998). In a second scenario, the maximal utility is quite insensitive to reasonable variations in the input data but even small deviations from the optimal allocation severely degrades the achieved utility optimum: in this case the optimal asset allocation should be pursued. In a third scenario, the maximal utility is quite insensitive to reasonable variations in the input data and largely different allocations yield similar utilities: in this case the current asset allocation might be apparently quite different from the optimal one but, if the respective utilities differ little, it is better not to rebalance the portfolio and to avoid useless transactions costs. This happens for example in highly correlated markets, where assets are good substitutes for one another.

In this paper we use a Bayesian framework to distinguish among the above cases and come up with an allocation decision. The current, suboptimal allocation defines implicitly the prior distribution, which is the Bayesian predictive distribution when the confidence in the prior is infinite; the Markowitz model that yield the optimal allocation corresponds to a predictive distribution with non-informative prior, i.e., zero confidence in the prior; these two extreme cases are spanned by a continuum of intermediate predictive distributions where the confidence level decreases from infinity to zero: along this continuum it is possible to test the sensitivity of the allocation and of the objective function to the input parameters, and the cost of choosing one allocation instead of another.

There are several other applications of Bayesian theory to portfolio selection and risk valuation in the financial literature. Among others, Jorion (1986) advocates Bayes methods to estimate and model risk; Black and Litterman (1990) and Pastor (2000) make use of Bayesian techniques to pool the investor's prior with an equilibrium model; Anderson et al. (2000) examine the robustness of classical results in a multiple prior setting. This literature focuses mainly on the differences between classical and Bayesian optimal allocations, without much emphasis on the objective function of the optimization. Nevertheless, in the spirit of Bayesian theory the choice of parameter values is as relevant as the values of the maximized expected utility. Indeed, recent Bayesian literature on model comparison and model robustness (see e.g. Gustafson et al.(1999), Carota et al. (1996)) proposes diagnostics based on objective functions, rather than estimates.

In Section 2 we show how to compare the current suboptimal allocation and the Markowitz optimal allocation by means of an analytical Bayesian framework. In Section 4 we apply this framework to the analysis of two different investment decisions: the allocation of wealth in a mutual fund of bond and stock indices and the choice of the optimal hedge ratio for a relative value strategy based on the spread between two stocks. In Section 5 we conclude.

2 The theoretical setting

To work in a Markowitz setting, we assume that the investor's utility be the standard quadratic function of portfolio returns:

$$u\left(r^P\right) = \frac{1}{k}\left(r^P - k\left(r^P\right)^2\right), \tag{1}$$

where k is a risk-aversion parameter. We denote the decision date by T. The investor can choose among a set of d assets whose future returns we denote by $\mathbf{r}_{T+1}$. Since $r^P_{T+1} = \boldsymbol{\omega}'\mathbf{r}_{T+1}$, where ω_i is asset i-th relative weight in the portfolio, the investor's decision amounts to picking $\boldsymbol{\omega}$ in such a way to maximize expected utility:

$$\max_{\boldsymbol{\omega}} E\left\{u\left(\boldsymbol{\omega}'\mathbf{r}_{T+1}\right)\right\}. \tag{2}$$

In order to make this decision, the investor needs the probability density $D\left(\mathbf{r}_{T+1}\right)$ of the predictive distribution In particular, given the assumptions on the utility function (1), she only needs to determine the first two moments of that distribution, i.e., expected value and covariance matrix of the returns $(\boldsymbol{\nu}, \Phi)$ and then solve:

$$\max_{\boldsymbol{\omega}} \frac{1}{k}\left\{\boldsymbol{\omega}'\boldsymbol{\nu} - k\left(\boldsymbol{\omega}'\boldsymbol{\nu}\right)^2 - k\boldsymbol{\omega}'\Phi\boldsymbol{\omega}\right\}. \tag{3}$$

In the Bayesian approach, the investor has some prior, subjective information on the distribution of returns, which she updates with more objective information to yield the predictive distribution. The predictive distribution is therefore a function of the prior and of the confidence that the investor has on this prior. If the confidence is infinite, the predictive distribution is very "subjective", i.e., simply the prior. If the confidence is zero, the prior is non-informative and the predictive distribution is very "objective".

There are two main characters in our play: the current, possibly suboptimal allocation, and the supposedly optimal allocation obtained a-la-Markowitz. To fit these two characters in a Bayesian setting, we make the current portfolio to be the optimal allocation for a Bayesian investor with infinite confidence in her prior, i.e., for a very "subjective" Bayesian investor. On the other extreme, we make the completely non-informative predictive distribution coincide with the "objective" distribution that yields the supposedly optimal allocation. In this paper we consider the objective distribution to be the normal distribution defined by the sample mean and covariance; other approaches are possible, such as the equilibrium model by Black and Litterman (1990). Setting the confidence in the prior between the above two extremes we obtain a continuum of investors and of investment decisions which allows us to evaluate one portfolio with respect to the other and choose the best allocation.

Therefore, we assume that next period's returns are normally distributed with mean $\boldsymbol{\mu}$ and covariance Σ. The same holds for all past returns up to the decision date. Therefore the probability density of the statistical model is

$$D\left(\mathbf{r}_t|\boldsymbol{\mu},\Sigma^{-1}\right)=\frac{\left|\Sigma^{-1}\right|^{\frac{1}{2}}}{(2\pi)^{\frac{d}{2}}}\exp\left\{-\frac{1}{2}\left(\mathbf{r}_t-\boldsymbol{\mu}\right)'\Sigma^{-1}\left(\mathbf{r}_t-\boldsymbol{\mu}\right)\right\}, \tag{4}$$

for all $t=1,2,\ldots,T+1$. Furthermore, returns are independent across time. The model parameters $(\boldsymbol{\mu},\Sigma)$ are not known with certainty. Nonetheless, the investor has some prior approximate knowledge of their values, which we denote by $(\boldsymbol{\mu}_0,\Sigma_0)$: the distribution of $(\boldsymbol{\mu},\Sigma)$ is Normal-Wishart centered on $(\boldsymbol{\mu}_0,\Sigma_0)$ with uncertainty parameters (c,g).

Explicitly, this means that the density of $\boldsymbol{\mu}|\Sigma^{-1}$ is normal:

$$D\left(\boldsymbol{\mu}|\Sigma^{-1}\right)=\frac{\left|c\Sigma^{-1}\right|^{\frac{1}{2}}}{(2\pi)^{\frac{d}{2}}}\exp\left\{-\frac{1}{2}\left(\boldsymbol{\mu}-\boldsymbol{\mu}_0\right)'c\Sigma^{-1}\left(\boldsymbol{\mu}-\boldsymbol{\mu}_0\right)\right\}.$$

This density is centered on the prior:

$$E\left(\boldsymbol{\mu}|\Sigma^{-1}\right)=\boldsymbol{\mu}_0,$$

and the parameter $c>0$ indicates the confidence level in the prior expect value[1]: the larger c, the more confident we are that the true model parameter $\boldsymbol{\mu}$ is close to our input $\boldsymbol{\mu}_0$.

As for the covariance matrix, Σ^{-1} is distributed as a Wishart:

$$D\left(\Sigma^{-1}\right)=\frac{\left|\frac{1}{2}g\Sigma_0\right|^{\frac{g}{2}}\left|\Sigma^{-1}\right|^{\frac{g-d-1}{2}}\exp\left\{-\frac{1}{2}Tr\left(g\Sigma_0\Sigma^{-1}\right)\right\}}{\Gamma_d\left(\frac{1}{2}g\right)}.$$

This density is also centered on the prior:

$$E\left(\Sigma^{-1}\right)=\Sigma_0^{-1},$$

and the parameter $g>d-1$ indicates the confidence level in the prior covariance[2]: the larger g, the more confident we are that the true model parameter Σ is close to our input Σ_0.

Starting from the above assumptions, the predictive distribution for next period returns $D\left(\mathbf{r}_{T+1}\right)$ is a d-dimensional Student. In particular, after some

[1] Indeed,

$$Cov\left(\boldsymbol{\mu}|\Sigma^{-1}\right)=\frac{1}{c}\Sigma$$

[2] as in Aitchison and Dunsmore (1975) we have

$$Cov\left[Vec\left(\Sigma^{-1}\right)\right]=\frac{1}{g}\left(I_{d^2}+K\right)\left(\Sigma_0^{-1}\otimes\Sigma_0^{-1}\right),$$

where I_{d^2} and K are the $d^2\times d^2$ identity and commutation matrices respectively computations available upon request in a technical appendix, we obtain the expected value $\boldsymbol{\nu}$ and the covariance matrix Φ of this distribution:

$$
\begin{aligned}
\boldsymbol{\nu} &= \frac{c\boldsymbol{\mu}_0 + T\mathbf{m}}{c+T} \qquad (5)\\
\Phi &= \frac{1+1/(c+T)}{g-(d-1)+T}\left(g\Sigma_0 + TS + \frac{c}{1+c/T}(\mathbf{m}-\boldsymbol{\mu}_0)(\mathbf{m}-\boldsymbol{\mu}_0)'\right),
\end{aligned}
$$

where $(\mathbf{m}, S)$ are the sample mean and covariance respectively:

$$
\begin{aligned}
\mathbf{m} &= \frac{1}{T}\sum_{t=1}^{T}\mathbf{r}_t \qquad (6)\\
S &= \frac{1}{T}\sum_{t=1}^{T}(\mathbf{r}_t-\mathbf{m})(\mathbf{r}_t-\mathbf{m})'.
\end{aligned}
$$

The continuum of Bayesian investors (5) is parameterized by the confidence levels c and g. To better understand the order of magnitude of these parameters in (5), we notice the symmetric role that (c, g) and the number of observations T play in this formula. The confidence levels can be interpreted as the number of "pseudo observations" that give rise to the prior. If the confidence is very large, i.e., $c, g \gg T$, the predictive distribution is simply the prior

$$\boldsymbol{\nu} = \boldsymbol{\mu}_0, \qquad \Phi = \Sigma_0. \qquad (7)$$

We choose the prior parameters $(\boldsymbol{\mu}_0, \Sigma_0)$ to be such that, if set in (2), they yield the current, possibly suboptimal portfolio. If the confidence is very low, i.e., $c, g \ll T$, the predictive distribution is equivalent to the normal distribution defined by the sample mean and covariance

$$\boldsymbol{\nu} = \mathbf{m}, \qquad \Phi = S \qquad (8)$$

Therefore the non-informative distribution gives rise to the "objective" optimal portfolio

3 The recipe

First of all we want to make sure that whatever we might state about an allocation makes sense from a statistical point of view. If the allocation is extremely sensitive to reasonable variations of the input data no sensible suggestion can be given on the basis of available data (Michaud, 1998). Therefore the recipe starts with

Step 1. Check the sensitivity of the current, potentially suboptimal portfolio to variations in the input parameters $(\boldsymbol{\nu}, \Phi)$.

Due to the high number of parameters, it is advisable to take a "comparative statics???" approach and see the effects of changing in turn each of a) expected returns and b) eigenvalues of the covariance matrix. If this sensitivity is high, there is little we can say about the optimality of a portfolio. Otherwise, we can move on to

Step 2. Check how abrupt the change is between the current, possibly suboptimal portfolio, and the supposedly optimal portfolio.

It is easy to perform this step using the Bayesian framework introduced in the previous section: by definition, solving (3) with the choice (7), which corresponds to a high confidence in the prior, yields the current portfolio; solving (3) with the choice (8), which corresponds to a low confidence in the prior, yields the supposedly optimal portfolio. Decreasing smoothly the confidence level in the prior in (5) it is possible to change smoothly from one allocation to the other and detect at which point the optimal allocation diverges from the current one. Now we are ready for

Step 3. Check how abrupt the change is between the *utility* of the current, possibly suboptimal portfolio, and the *utility* of the supposedly optimal portfolio.

4 Empirical analysis

When an investor decides to rebalance her asset allocation according to the results of an optimization tool, she always faces the following question: "Is the portfolio really that different? Is it worth facing the ensuing transaction costs?".

In this section we suggest some Bayesian diagnostics to distinguish these scenarios and assess the optimality an allocation. We consider two practical examples: a standard optimal asset allocation problem typical in mutual fund management and the computation of a "market neutral" position typical in hedge fund management. To develop the diagnostics we consider a spectrum of Bayesian investors that includes the classical decision maker as a special case. We also analyze the dependence of the optimal classical allocation on the input data. Finally, we study the allocations of the classical investor and of neighboring Bayesian investors in terms of their objective function.

4.1 Mutual funds: optimal portfolio

In this first example investor I is a fund manager. Therefore she maximizes the expected value of her utility (1) under the constraint that the relative weights of each asset in portfolio are positive and sum to one. Her allocation is the solution $\boldsymbol{\omega}_I$ to

$$U_I \equiv \max_{\substack{\boldsymbol{\omega} \geq \mathbf{0} \\ \boldsymbol{\omega}'\mathbf{i}=1}} \left\{ \boldsymbol{\omega}'\boldsymbol{\nu}_I - k\left(\boldsymbol{\omega}'\boldsymbol{\nu}_I\right)^2 - k\boldsymbol{\omega}'\Phi_I\boldsymbol{\omega} \right\}, \tag{9}$$

where $(\boldsymbol{\nu}_I, \Phi_I)$ are the expected value and covariance matrix as perceived by the investor of next period's returns on the d assets she can trade.

Accordingly, the classical investor finds an optimal allocation $\boldsymbol{\omega}_C$ solving (9) with the choice (**??**). We assume the covariance matrix $\widehat{\Sigma}$ is estimated as the non-central second sample moment of exponentially smoothed past returns (a standard methodology applied, e.g., by RiskMetrics); the expected returns are assumed defined by a risk-premium model of the form:

$$\widehat{\boldsymbol{\mu}} = \mathbf{r}_f + \lambda \times \widehat{\boldsymbol{\sigma}}, \tag{10}$$

where $\mathbf{r}_f$ is the risk-free rate, λ is the risk-premium and $\widehat{\boldsymbol{\sigma}}$ is the vector of volatilities, i.e., the square root of the diagonal of $\widehat{\Sigma}$.

On the other hand, the Bayesian investor finds an optimal allocation $\boldsymbol{\omega}_B$ solving(9) with the choice (5), where the prior is "centered" on the classical estimate, as in (**??**). The Bayesian solution $\boldsymbol{\omega}_B$ depends on the confidence levels c and g.

To illustrate, we choose a risk aversion parameter $k = 25$ in (1) and define a market of $d = 4$ assets: a mid-duration US government bond price index (3-5yrs G3O2), a long-duration US government bond price index (7-10yrs G9O2), a low volatility stock price index (S&P 500) and a high volatility stock price index (Nasdaq). All assets are denominated in the same currency (the US dollar). The database of daily returns spans the period January 4th 1999 through the decision date July 25th 2000, for a total of $T = 408$ observations. The classical covariance matrix is estimated by smoothing the returns with a decay factor equal to 0.95. The classical expected returns are obtained setting $\mathbf{r}_f = \mathbf{0}$ and $\lambda = 0.25$ in (10). Instead of considering the whole spectrum of Bayesian frameworks for arbitrary values of the confidence levels c and g we consider the case $c = g$, that is, we assume equal confidence level on both the expected value and the covariance matrix of returns. We let the confidence decrease from $c = g = 10T$ to $c = g = T/10$. In the first case, which corresponds to about 4000 "pseudo observations" in the classical estimate, we obtain an approximation of the classical framework. As we approach the lower extreme (around 40 "pseudo observations") the investor becomes more and more Bayesian.

4.1.1 Optimal allocation

We solve (9) for the above grid of confidence parameters and we first turn our attention to the portfolio allocation. In Figure 1 we plot the results. As we pointed out in (**??**) the Bayesian investor is more "cautious", therefore she allocates more wealth in low-volatility assets: we see in the figure that the overall exposure to bonds increases as the investor turns more and more Bayesian. Nevertheless, this change is not dramatic. Instead, we notice huge changes in the balance between mid-duration and long-duration bonds. As the investor turns more and more Bayesian, the former increases and the latter decreases by several orders of magnitude, changing completely the allocation.

4.1.2 Allocation sensitivity to the input data

Is the phenomenon in Figure 1 statistically significant (Michaud, 1998)? We focus on the classical investor ($c = g \gg T$) and check whether a little change in her estimated parameters severely affects the optimal allocation. If this is the case, the manager should not be too surprised to see the dramatic changes in allocation displayed in the figure.

First we perform a comparative static analysis on the expected values. In Figure 2, the first column to the left represents the baseline case, i.e., the composition of the classical investor that also appears as the leftmost case in 1. The other columns represent respectively the optimal allocation if the expected returns (10) of the mid-duration bonds, the long-duration bonds, the S&P and the Nasdaq are increased in turn by 10%. As expected (Best and Grauer, 1991), the effect of these changes on the optimal allocation are significant, but not as much as to justify the changes in allocation displayed in Figure 1.

Now we focus on the sensitivity to the covariance matrix $\widehat{\Sigma}$. Instead of varying in turn all the four variances and the six covariances, we act on the eigenvalues and keep the eigenvector structure constant. Implicitly we are assuming that the factor structure underlying the correlation across securities is constant while the weight of each factor is subject to change. This simplifies our analysis and is consistent with common influence-based statistical procedures (Huber, 1986). In Figure 3 we display the results. As expected (see Best and Grauer, 1991), the optimal portfolio is more robust to changes in the covariance structure than it is to modifications in the expected values.

4.1.3 Objective function sensitivity to allocation

In this section we consider to what extent the allocations in 1 are really different. In order to do this, we compute the objective function (9) that corresponds to those allocations, In Figure 4 we display the result: the more Bayesian the investor the less happy she is with her portfolio. This is not trivial, since her portfolio is optimal. The information we draw from 4 is only qualitative: the units of utility on the vertical axis are not well defined. To gain a more quantitative insight in the utility of an allocation, we notice in (9) that higher expected returns increase utility. Therefore a decrease in the value of the objective function can be offset by an increase in the risk premium λ in (10). In Figure 5 we plot the locus of risk premia that the Bayesian investor would require to attain the utility of the classical decision maker. We see how for almost all of the classical to Bayesian spectrum a very small increase in the risk premium is all the Bayesian investor asks to compensate her for the utility loss. For this phenomenon to happen, the utility of the Bayesian and of the classical investors have to be quite similar: the portfolio compositions in Figure 1 are only apparently different. From a statistical point of view we could say that the objective function is robust, whereas portfolio compositions are not.

Another way to look at this problem is to compute the cost, i.e., the utility loss of an investor that accepts someone else's allocation. Accordingly, we define the cost of choosing the Bayesian allocation for a classical investor as follows:

$$C_{B4C} \equiv U_C - \frac{1}{k}\left(\boldsymbol{\omega}_B'\boldsymbol{\nu}_C - k\left(\boldsymbol{\omega}_B'\boldsymbol{\nu}_C\right)^2 - k\boldsymbol{\omega}_B'\Phi_C\boldsymbol{\omega}_B\right).$$

Similarly we define the cost of being classical for a Bayesian investor as:

$$C_{C4B} \equiv U_B - \frac{1}{k}\left(\boldsymbol{\omega}_C'\boldsymbol{\nu}_B - k\left(\boldsymbol{\omega}_C'\boldsymbol{\nu}_B\right)^2 - k\boldsymbol{\omega}_C'\Phi_B\boldsymbol{\omega}_C\right).$$

We compute these quantities for the usual spectrum of classical to Bayesian investor. In Figure 6 we display the results. Again, the mutual costs for the Bayesian investor and the classical decision maker to choose the other's portfolio is negligible as long as the Bayesian investor is sufficiently confident in her prior: the two investors see each other's portfolios as good substitutes and the differences in portfolio composition between them are more apparent than real.

4.2 Hedge funds: spread positions

In the second example we tackle a simple optimal hedge problem. This is a stylized version of problems met by hedge fund managers: the investor takes spread positions on a set of d assets. In this case the expected returns (which dominate the asset allocation problem) become negligible and the investor focuses only on the risk of her portfolio. This setting is the limiting case of the utility (1) for an infinite risk-aversion parameter k.

Accordingly, the classical investor finds an optimal allocation $\boldsymbol{\omega}_C$ solving (9) with the choice (**??**). We set $\widehat{\boldsymbol{\mu}} = \mathbf{0}$ and the covariance matrix $\widehat{\Sigma}$ is estimated as in the previous example as the non-central second sample moment of exponentially smoothed past returns with a factor 0.95. On the other hand, the Bayesian investor finds an optimal allocation $\boldsymbol{\omega}_B$ solving(9) with the choice (5), where the prior is centered on the classical estimate, as in (**??**). Again, the Bayesian solution $\boldsymbol{\omega}_B$ depends on the confidence levels c and g.

To illustrate, we consider only two stocks in the same currency and sector: BASF and BAYER. The respective daily time series span the period November 24, 1998 through November 21, 2000, for a total of $T = 507$ observations. Again, instead of considering the whole spectrum of Bayesian frameworks for arbitrary values of the confidence levels c and g we consider the case $c = g$, that is, we assume equal confidence level on both the expected value and the covariance matrix of returns. We let the confidence decrease from $c = g = 10T$ to $c = g = T/10$. In the first case, which corresponds to about 5000 "pseudo observations" in the classical estimate, we obtain an approximation of the classical framework. As we approach the lower extreme (around 50 "pseudo observations") the investor becomes more and more Bayesian.

4.2.1 Optimal allocation

We compute the optimal portfolios assuming we are short the BASF shares. This amounts to computing the optimal number of BAYER shares to buy to hedge the market risk, In Figure 7 we plot the optimal hedge ratio (i.e., how much to invest in BAYER if we are short one unit of currency of BASF). For the sake of clarity in Table 1 we report the values of the hedge ratio for tree significant values of the confidence parameters. As in the mutual fund example, we notice a significative change in the hedge ratio as the investor turns Bayesian.

4.2.2 Allocation sensitivity to the input data

Empirical experience with optimal hedge ratios advocates for instability in the hedge result. In Table 2 we report the sensitivity of the hedge ratio of the classical investor to changes in the covariance matrix eigenvalues. A 10% change in each of the two eigenvalues implies approximately 6% changes in the optimal hedge ratio. This is equivalent to change from 5000 to 1000 of the confidence parameters ($c = g$). Unlike the example of asset allocation,. in this case a reasonable change in the input data has a relevant effect in the value of the optimal hedge ratio.

4.2.3 Objective function sensitivity to allocation

We can complement the analysis of the optimal position in Figure 7 with a study of the objective function, which in this case is the risk. In Figure 8 we plot the risk of the Bayesian hedge ratio for the Bayesian investor, which by definition corresponds to the least possible risk; the risk of the classical hedge ratio for the Bayesian investor; and the risk of the Bayesian hedge ratio for the classical investor. Again, we see the same pattern encountered in the case of mutual funds: the risk is perceived as almost the same for a wide range of the spectrum: in this range the optimal hedge ratio is only apparently different. As the confidence in the prior turns less than 1000 "pseudo observation", this perception diverges. We see that the sensitivity to input data stressed in the previous section is only borderline relevant if evaluated in utility scale.

5 Conclusions

This paper suggests a set of simple diagnostics to support asset allocation decisions: by simply looking at a few charts, the investor is able to assess if and how to rebalance her portfolio. These diagnostics focus on the objective function of the investor, not on the actual allocation.

We model a continuum of approaches: we start from a classical framework, where no uncertainty on the estimated parameters is assumed, and we move on toward a more and more Bayesian setting. This way we create a "neighborhood" of the classical asset allocation framework. In order to do this we adopt the Bayesian interpretation of Black and Litterman (1990).

We find out that the allocation changes considerably in the neighborhood of the classical approach. Therefore we test for the robustness of the allocation with respect to the data: the swings in allocation due to changes in the input data are not enough to justify such different allocations. This is in line with the literature (Michaud, 1998) which warns against a direct application of maximization tools to compute optimal portfolios. Then we turn our attention to the objective function and we find out that these are much more robust than the allocations.

Much work can be done in this direction. For instance, one can explore the classical neighborhood along different choices of the parameters c and g. Secondly, it would be interesting to apply the methodology of this paper to a non-mean-variance context. Thirdly, we make use here of only some principles drawn from the literature on the robustness of Bayesian procedures (Gustafson et al. 1999). We believe most of this literature could be usefully applied in the context of portfolio optimization and statistical hedging.

References

Aitchison J., Dunsmore,I.R. (1975). *Statistical prediction analysis.* Cambridge University Press

Anderson E. W., Hansen L. P., Sargent T. J. (2000) *Robustness, Detection and the Price of Risk.* Working Paper.

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

Berger J. O. (1985) *Statistical decision theory and Bayesian analysis*, Springer.

Best M.J., Grauer R.R. (1991) *On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results.* Review of Financial Studies, 4 pp.315-342.

Black, F. and Litterman R. (1990) *Asset allocation: combining investor views with market equilibrium.* Goldman Sachs and Co., Fixed income research.

Campbell, J. (2000) *Asset pricing and the millenium.* The Journal of Finance LV,4.

Carota, C., Parmigiani, G., Polson, N. G. (1996) *Diagnostic measures for model criticism.* Journal of the American Statistical Association, 91, 753-762.

Gustafson P., Srinivasan C. and Wasserman L. (1999). *Local sensitivity analysis.* Working Paper, Carnegie Mellon University.

Jorion, P. (1986) *Bayes-Stein estimation for portfolio analysis*, Journal of Financial and Quantitative Analysis, 21-3, pp. 279-292.

Michaud, R. (1998). *Efficient Asset Management*, Boston: Harvard Business School Press

Morris, N. (1983) *Parametric empirical based inference: theory and applications* Journal of the American Statistical Association 78 pp. 47-65.

Huber, P. J. (1981) *Robust statistics* Wiley and sons, New York.

Theil, H., Goldberger A.S. (1961) *On pure and mixed estimation in economics.* International Economic Review 2 pp. 65-78.

Table 1: HEDGE RATIOS FOR DIFFERENT CONFIDENCE VALUES

Confidence $c = g$	Position
5070 ($= 10T$)	0.70
522 ($\sim T$)	0.64
50.7 ($= T/10$)	0.58

Table 2: SENSITIVITY TO CHANGES IN THE COVARIANCE MATRIX

no changes	change eig1	change eig2
0.7023	0.7416	0.6635

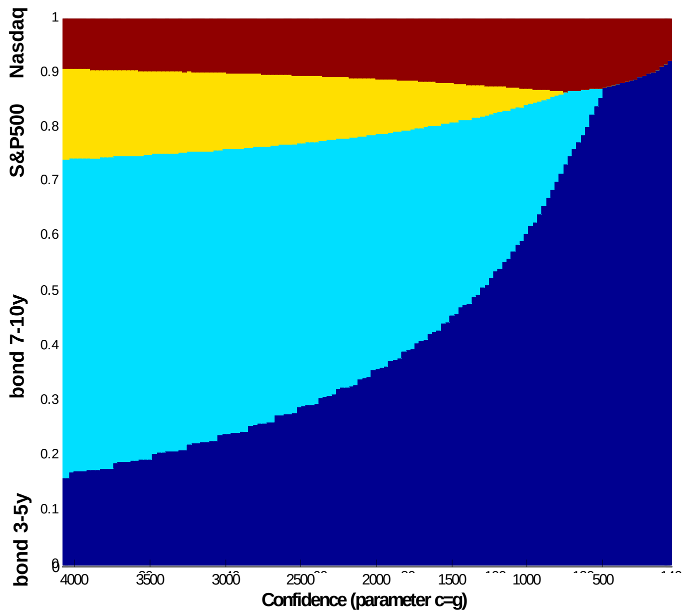

Figure 1: optimal allocation for a spectrum of Bayesian investors (parametrized by their confidence $c = g$ in the prior, which is a classical estimate). The leftmost allocation is an approximation to the portfolio of a classical decision maker. The further to the right we move in the picture, the more Bayesian the investor.

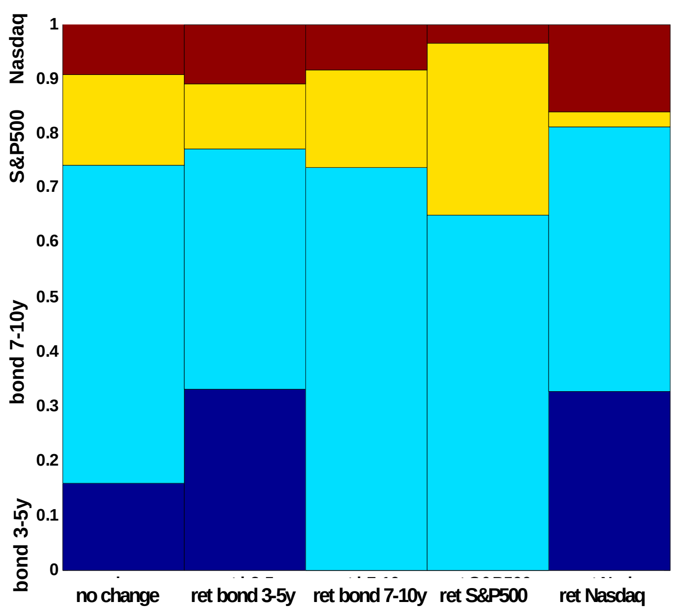

Figure 2: sensitivity of the optimal portfolio to changes in the expected returns. The leftmost bar correspond to the optimal allocation for the classical investor. In the subsequent bars we increase in turn the values of the classical estimate of each expected return by 10%.

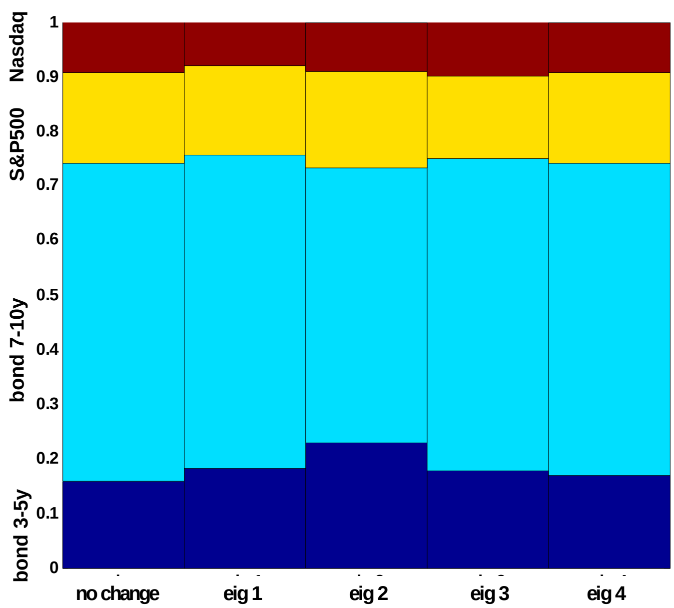

Figure 3: sensitivity of the optimal portfolio to changes in the eigenvalues of the covariance matrix. The leftmost bar correspond to the optimal allocation for the classical investor. In the subsequent bars we increase in turn the eigenvalues of the classical estimate of the covariance matrix by 10%.

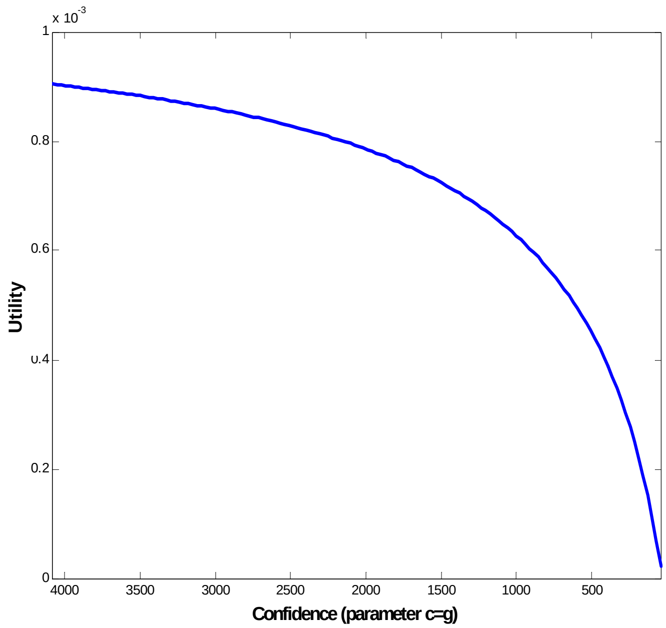

Figure 4: utility drawn on optimal allocation for a spectrum of Bayesian investors (parametrized by their confidence $c = g$ in the prior, which is a classical estimate). The leftmost value approximates the utility of a classical decision maker. The further to the right we move in the picture, the more Bayesian the investor.

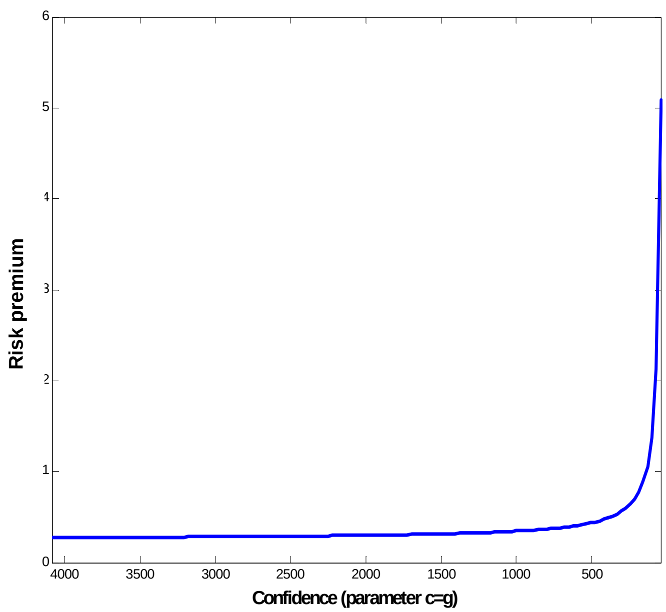

Figure 5: risk premium that makes the utility from optimal allocation of a spectrum of Bayesian investors (parametrized by their confidence $c = g$ in the prior, which is a classical estimate) equal to the utility of a classical decision maker.

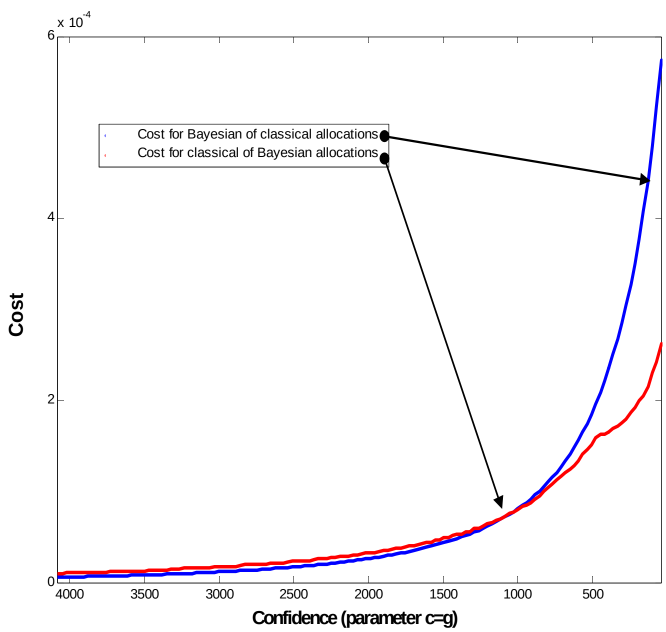

Figure 6: cost (defined as the utility loss) of a spectrum of Bayesian investors (parametrized by their confidence $c = g$ in the prior, which is a classical estimate) and of the classical decision maker when they choose the other's optimal allocation

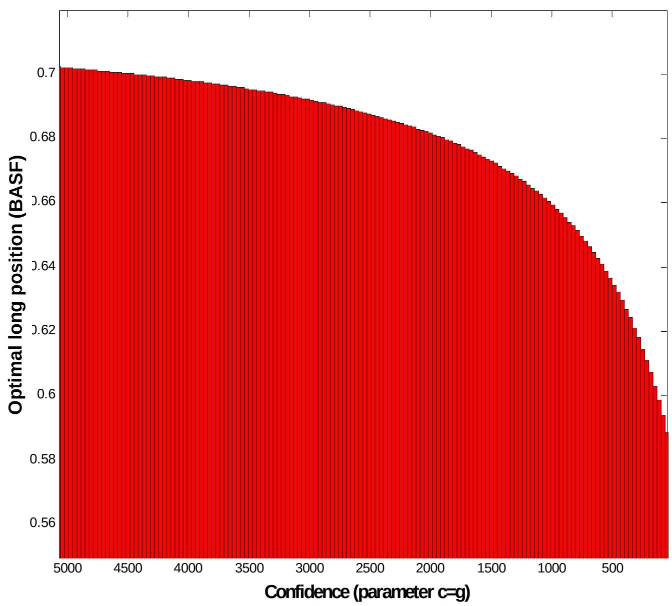

Figure 7: optimal hedge ratio for a spectrum of Bayesian investors (parametrized by their confidence $c = g$ in the prior, which is a classical estimate). The hedge ratio is defined as the amount to invest in one stock (BASF in this example) to best hedge the risk of a one unit of currency short position in another stock (BAYER in this example). The leftmost hedge ratio is an approximation to the classical decision maker's. The further to the right we move in the picture, the more Bayesian the investor.

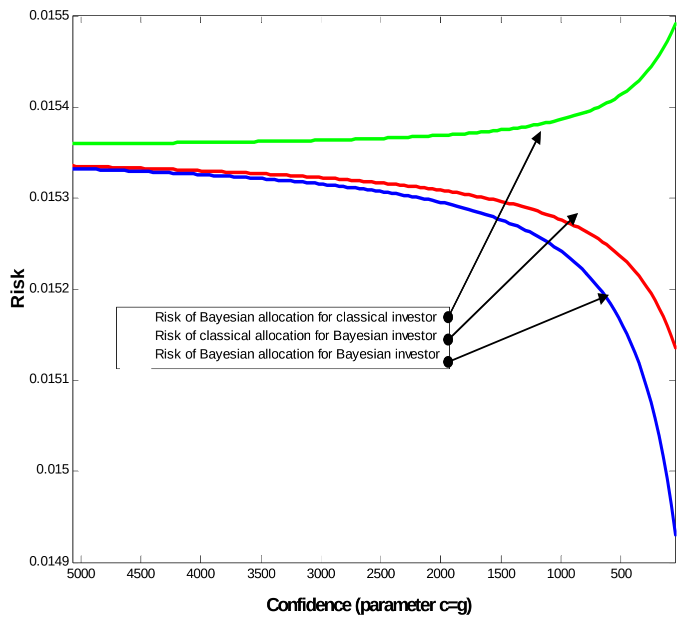

Figure 8: comparison of the risks faced by a spectrum of Bayesian investor (parametrized by their confidence level $c = g$) and by a classical decision maker when they choose different hedge ratios