Portfolio Selection with Estimation Risk: a Test Based Approach*

Bertille Antoine†

October 9, 2008

Preliminary - Comments welcome

Abstract

An important challenge of portfolio allocation arises when the (true) characteristics of returns distribution are replaced by some estimates. This introduces estimation risk, which is crucial for portfolio management just like traditional financial risk. This paper contrasts with existing literature by focusing on a different measure of performance. We borrow from practitioners and evaluate different funds allocations through their likelihood of beating a benchmark. Then, the optimal portfolio which accounts for estimation risk is known in closed-form and does not depend on any nuisance parameter. This investment rule corresponds to a mean-variance investor with a corrected, sample-dependent and counter-cyclical risk aversion parameter.

JEL Classification: C4, D8, G0.

Key words: Portfolio theory; Estimation risk; Benchmark performance; Mean-variance efficiency; Test.

*I would like to thank Eric Renault for his constant encouragements and advises. I also thank Bas Werker for helpful discussions.

†*Simon Fraser University. Email: bertille_antoine@sfu.ca.*

1 Introduction

An optimal portfolio is the *best* allocation of funds across available assets[1]. Of course, what best means depends on the performance measure we use. Markowitz (1959) offers the classic definition of portfolio efficiency: a portfolio is efficient if it has the largest expected return for a given target of risk as measured by the variance. For a given level of risk-aversion, this mean-variance efficiency provides a convenient single-period framework and remains among the most important benchmark models used by practitioners nowadays (see Meucci (2005)). In practice, however, its associated optimal investment rule depends on unknown parameters, the mean and the variance of returns distribution. To get a feasible version of this optimal rule, Markowitz (1959) simply replaces the unknown parameters by some estimates. Applying such a plug-in method gives rise to several issues. First the estimation risk is overlooked: in practice samples are finite, hence estimates are different from their respective true values. This new source of risk even appears in well-specified parametric models and adds to the traditional financial risk[2]. Second, is this feasible rule optimal? A (suboptimal) two-step approach can only be motivated when one believes that the estimated rule is *not too far* from the true optimal one.

In sharp contrast with existing literature, we focus on a different (more conservative) measure of performance. We borrow from practitioners and evaluate different funds allocations through their likelihood of beating a benchmark. Several industries are actually interested in such a goal: for instance, institutional money managers, and among others the defined benefits pension plans and the endowment plans are devoted to guarantee the (chosen) minimal performance. For a given benchmark, we deduce a closed-form and workable optimal investment rule which naturally incorporates the estimation risk of the mean, and does not depend on any nuisance parameter. Hence it is directly applicable without requiring any additional (suboptimal) plug-in step.

More precisely, our portfolio selection method is based on a one-sided test ensuring that

[1] Brandt (2004) provides a broad survey on general issues related to portfolio choice.

[2] Kan and Zhou (2006) provide an extensive study of the financial consequences of ignoring estimation risk.

the portfolio performance is above a given threshold; then we obtain the optimal allocation from the maximization of the associated p-value. The specific design of the p-value selection method has three advantages. First, testing is a natural and valid statistical tool to compare random quantities (here the estimated portfolio performance). Hence the uncertainty of the problem is directly accounted for: we will see that this is crucial to get an exact optimal investment rule. Second, maximizing the p-value increases the likelihood of our objective of interest (here to beat the chosen benchmark). Finally the optimal investment rule belongs to the class of two-fund investment rules, similar to the (feasible) optimal mean-variance rule: investing in the (sample) tangency portfolio and in the riskless asset[3]. This investment rule corresponds to a mean-variance investor with a corrected, sample-dependent risk-aversion parameter. While existing literature recommends to increase the risk-aversion parameter to account for estimation risk, we advocate more flexibility: we may indeed decrease the risk-aversion parameter depending on the realized sample.

Our work relates to the literature as follows. First, estimation risk in portfolio allocation has been known for a while. One of the earliest and maybe most natural solution appears to be Bayesian. The Bayesian approach is based on the predictive distribution introduced by Zellner and Chetty (1965) under which expectations are now considered. It provides a general framework where estimation risk is naturally accounted for when considering the parameters as random variables: the posterior distribution captures their possible outcomes and is combined to a prior model to derive the predictive distribution. The study by Bawa, Brown and Klein (1979) surveys the early literature, and is then followed by many others. It is not always clear how the prior model can be chosen, even though it is based on the investor's knowledge and experience: different priors may lead to contrastive investment strategies. We only consider non-informative prior models.
More recently, some authors decided to directly focus on the expected financial loss when the optimal investment rule is replaced by some feasible one. Due to the com-

[3]This specific class of investment rules has already been considered in the literature: see ter Horst, de Roon and Werker (2006) and Kan and Zhou (2006). However, here it directly follows from our portfolio selection method (just like the mean-variance procedure) and not from a simplifying assumption.

plexity of the problem, ter Horst, de Roon and Werker (2006) and Kan and Zhou (2006) restrict their attention to the class of two-fund investment rules (similar to the feasible optimal mean-variance rule). While ter Horst *et al.* (2006) ignore the estimation risk of the variance, Kan and Khou (2006) (under the normality assumption of the returns) provide a closed-form solution to the simplified problem. However, the optimal rule depends on nuisance parameters. So, in order to implement this optimal strategy, one needs to add a suboptimal plug-in step[4]. The mean-variance framework seems to be limited. As shown by Kan and Zhou (2006), who were able to exhibit the optimal two-fund rule while accounting for estimation risk, the outcome is not completely satisfactory as the optimal rule is unfeasible. This more general issue arises when one maximizes some expected quantity. This motivates our approach: we take some distance with the traditional minimization of an expected financial loss function and maximize the likelihood of some desirable event[5].
Finally, previous studies have already focused on defeating a benchmark: see for instance Stutzer (2003) and references therein. However, to our knowledge, this has not yet been related to estimation risk. Moreover, these studies work in a continuous time framework: this is definitely not our interest here.

To conclude, a simple Monte-Carlo study involving five risky assets (calibrated from monthly unhedged returns of stock indices for the G5 countries) is used to compare eleven investment strategies. These are compared with respect to their out-of-sample expected performances as well as with respect to their maintenance costs and stability over some investment horizon. The p-value selection method performs surprisingly well considering it is not specifically designed to maximize the mean-variance performance. Moreover, it avoids extreme positions in the assets and remains relatively stable over

[4]Of course, by construction, Kan and Zhou (2006) theoretical two-fund rule outperforms any p-value investment rule. However nothing is guaranteed when one considers its feasible version as shown in our simulation exercise.

[5]Others have also departed from the classical mean-variance approach: Garlappi, Uppal and Wang (2006) propose a sequential max-min method where the worst performance (when the unknown parameters fall into a confidence interval) is maximized with respect to the portfolio weights; Harvey, Liechty, Liechty and Muller (2004) adopt a Bayesian setting under the assumption that the returns follow a skew-normal distribution.

time.

The remainder of the paper is organized as follows. Section 2 solves the classical mean-variance problem. The p-value selection method is introduced in section 3. Section 4 reviews some competing investment strategies. Section 5 presents the results of a simple simulation study calibrated from real data. Section 6 concludes.
The details of the calculations are gathered in the appendix.

2 Classical Mean-Variance problem

This section discusses the mean-variance problem and introduces estimation risk. Consider an investor who chooses a portfolio among N financial risky assets and the riskless asset. At time t, denote respectively by $R_t \equiv (r_{1t} \ \cdots \ r_{Nt})'$ and R_{ft} the rates of returns on the N risky assets and the riskless asset. The vector of excess returns is defined as $\tilde{R}_t \equiv R_t - R_{ft}\iota$ where ι is the conformable vector of ones. The following standard assumption is maintained on the probability distribution of excess returns $\tilde{R}_t$:

ASSUMPTION 1 *The vector of excess returns $\tilde{R}_t$ is independent and identically distributed over time. In addition, R_t is normally distributed with mean $\tilde{\mu}_0$ and variance Σ_0.*

At time t, the portfolio is built after investing a vector θ into the risky assets and $(1 - \theta'\iota)$ in the riskless asset. The portfolio excess return is $r_t^P(\theta) \equiv \theta'\tilde{R}_t$, and its associated mean and variance are then respectively,

$$\mu_P = \theta'\tilde{\mu}_0 \quad \text{and} \quad \sigma_P^2 = \theta'\Sigma_0\theta$$

The vector of weights θ defines the investment rule which maximizes the following mean-variance objective function:

$$\max_{\theta} \left\{ E\left[r_t^P(\theta)\right] - \frac{\eta}{2} Var\left[r_t^P(\theta)\right] \right\} \iff \max_{\theta} \left\{ \theta'\tilde{\mu}_0 - \frac{\eta}{2}\theta'\Sigma_0\theta \right\}$$

where η is the coefficient of relative risk-aversion. This leads to the following optimal vector of weights:

$$\theta_{MV} = \frac{1}{\eta}\Sigma_0^{-1}\tilde{\mu}_0 \tag{2.1}$$

In practice, the parameters $\tilde{\mu}_0$ and Σ_0 are unknown: therefore the optimal mean-variance investment rule θ_{MV} is *unfeasible* and cannot be calculated in practice. Markowitz (1959) simply replaces the unknown parameters by some estimates. This easily provides a *feasible* version of the above optimal rule. More precisely, for some estimates $\hat{\tilde{\mu}}$ and $\hat{\Sigma}$ of the unknown parameters $\tilde{\mu}_0$ and Σ_0, one defines the feasible (random) investment rule and its associated (random) performance as:

$$\hat{\theta}_{MV} = \frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\tilde{\mu}} \quad \text{and} \quad Q(\hat{\theta}_{MV}) = \frac{1}{2\eta}\hat{\tilde{\mu}}'\hat{\Sigma}^{-1}\hat{\tilde{\mu}} \tag{2.2}$$

where $\hat{\tilde{\mu}}$ and $\hat{\Sigma}$ are, for instance, the maximum likelihood estimators,

$$\hat{\tilde{\mu}} = \frac{1}{T}\sum_{t=1}^{T}\tilde{R}_t \quad \text{and} \quad \hat{\Sigma} = \frac{1}{T}\sum_{t=1}^{T}(\tilde{R}_t - \hat{\tilde{\mu}})(\tilde{R}_t - \hat{\tilde{\mu}})' \tag{2.3}$$

Applying this *plug-in* method comes at a price. First, estimation risk is overlooked. In practice, the sample size is only T (finite), hence $\hat{\tilde{\mu}}$ and $\hat{\Sigma}$ are different from their respective true values. Second, precisely because the feasible rule $\hat{\theta}_{MV}$ is numerically different from the true optimal one, its optimality cannot be guaranteed. In the next section, we propose a portfolio selection method that incorporates estimation risk and does not require any additional (suboptimal) step.

3 Maximization of the p-value

This section introduces the p-value selection method and derives the associated optimal investment rule for a given benchmark c. In a second step, the question of the existence of an optimal benchmark is raised.

3.1 Definition and Optimal investment rule

As emphasized earlier, this paper takes some distance with the classical mean-variance framework and the common idea of minimizing some (expected) financial risk function. More precisely, in sharp contrast with existing literature, we do not maximize any usual measure of portfolio performance. We rather compare available funds allocations through their likelihood of beating the chosen benchmark. Of course, our portfolio selection method crucially depends on this benchmark. *Reasonable* benchmark choices yield to more *conservative* objective functions than the classic maximization of the (mean-variance) performance. Our investor is more conservative in the sense that she is not interested in achieving the maximal performance at every period; she rather selects the investment rule that maximizes the likelihood of defeating the benchmark. This selection method directly accounts for the random nature of the problem while being of primary concern for several industries, like institutional money managers.

Our portfolio selection method is based on a one-sided test that the chosen measure of portfolio performance is above the given threshold. Obviously, two unknowns remain here: first the choice of the performance measure and second the threshold. As pointed out earlier, Markowitz's mean-variance efficiency is a convenient framework privileged by practitioners. Accordingly, we consider the following measure of portfolio performance:

$$Q(\mu_P, \sigma_P^2) = \mu_P - \frac{\eta}{2}\sigma_P^2 \tag{3.1}$$

where (μ_P, σ_P^2) are respectively the first two moments of the probability distribution of the portfolio. This measure of performance has mainly been chosen for comparison purposes: our p-value selection method works with any other measure $Q(.)$[6]. Not only the test is the natural statistical tool to compare random quantities and incorporate estimation risk, but also it directly focuses on the well-defined objective for a portfolio

[6] Any performance measure works under regularity assumption like assumption 2. In particular, we could think of incorporating higher moments to account for effects of skewness, kurtosis... This only affects the tractability of the optimal investment rule when one wants to account for the associated estimation risk. This is indeed related to the (asymptotic) distribution of the estimated portfolio performance (that may need to be simulated) .

manager, to beat the performance of a benchmark index.

Formally, the null hypothesis of interest is stated as follows:

$$H_0 : \; Q(\mu_P, \sigma_P^2) > c \tag{3.2}$$

where c is the (deterministic) performance of the (chosen) benchmark index. To construct the associated test statistic, some assumptions are needed on the probability distribution of the returns. Consider an investor at time T who has observed the N risky asset returns from time $t = 1$ to T.

ASSUMPTION **2** *The vectors of the N financial excess returns of interest at time t, $\tilde{R}_t = [\tilde{r}_{1t} \cdots \tilde{r}_{Nt}]'$ for t=1 to T, are identically distributed and serially independent. More formally,*

1) $\tilde{R}_t \sim \mathcal{F}(\tilde{\mu}_0, \Sigma_0)$ *for any* $t = 1, \cdots, T$ *where $\mathcal{F}$ is some smooth distribution function whose first two moments exist*
2) $\tilde{R}_t$ *and* $\tilde{R}_{t'}$ *are independent for any t and* $t' = 1, \cdots, T$ *such that* $t \neq t'$

We consider from now on the portfolio excess return $\tilde{r}_t^P(\theta) = \theta' \tilde{R}_t$. Note that this only shifts the deterministic benchmark c: only strictly positive benchmarks c are now considered. A null benchmark corresponds to the minimal acceptable performance, guaranteed when always investing in the riskless asset. The measure of portfolio performance is then written as:

$$Q_P(\theta) = E\tilde{r}_t^P(\theta) - \frac{\eta}{2} Var(\tilde{r}_t^P(\theta))$$

and is estimated by[7]:

$$\hat{Q}_P(\theta) = \theta' \hat{\tilde{\mu}} - \frac{\eta}{2} \theta' \hat{\Sigma} \theta \tag{3.3}$$

$$\text{with} \;\; \hat{\tilde{\mu}} = \frac{1}{T} \sum_{t=1}^{T} \tilde{R}_t \;\; \text{and} \;\; \hat{\Sigma} = \frac{1}{T} \sum_{t=1}^{T} (\tilde{R}_t - \hat{\tilde{\mu}})(\tilde{R}_t - \hat{\tilde{\mu}})' \tag{3.4}$$

[7] The procedure remains similar for any other set of consistent estimates. We could even think of the selection problem as starting right here, with a set of estimates given by a practitioner.

The application of the vectorial central limit theorem yields the asymptotic distribution of the estimated performance: $\sqrt{T}\left[\hat{Q}_P(\theta) - Q_P(\theta)\right]$ is asymptotically normally distributed with mean 0 and variance $Var(\hat{Q}_P(\theta))$. Then, for an estimator $\hat{S}$ of its standard deviation, the test statistic and associated p-value are defined as follows:

$$St(\theta) = \frac{\hat{Q}_P(\theta) - c}{\hat{S}/\sqrt{T}} \quad \text{and} \quad \text{p-value}(\theta) = \int_{-\infty}^{St(\theta)} f_{\mathcal{T}}(u)du$$

with $f_{\mathcal{T}}$ the density function of a student random variable with $(T-1)$ degrees of freedom. Hence the maximization problem is finally stated as:

$$\max_{\theta} \left[\text{p-value}(\theta)\right] \iff \max_{\theta} \left[St(\theta)\right]$$

The p-value selection method can be linked to the well-known financial risk measure, the Value-at-Risk (VaR hereafter). Briefly the VaR at level α represents an estimate of the level of loss on a portfolio which is expected to be equaled or exceeded with the given, small probability α: risk regulations usually dictates the choice of this level of confidence. Our selection method rather guarantees the chosen minimal level of performance with the highest level of confidence. We think that choosing the benchmark is more inline with institutional money managers concerns.

Obviously, estimation risk is related to the estimation of both the mean and the variance of the portfolio. If it is commonly accepted that the estimation error on the sample mean is much larger than on the sample variance, recent studies suggest that it might not always be the case: see Cho (2007) and Kan and Zhou (2007). The latter authors conclude that this is only acceptable when the ratio number of assets over sample size (that is N/T) is *small*: in particular there is an interactive effect between both estimation errors. Here, to simplify the problem (and get an interpretable closed-form investment rule), we ignore the estimation risk of the variance[8]. The simplified maximization problem is now:

$$\theta_p(c) = \arg\max_{\theta} \left[\frac{\theta'\hat{\mu} - \eta/2\theta'\hat{\Sigma}\theta - c}{(\theta'\hat{\Sigma}\theta)^{1/2}/\sqrt{T}}\right]$$

[8]Note that, in our simulation study in section 5, the ratio N/T is kept small. In this sense, our methodology applies more to pension funds than mutual funds.

where $\hat{\hat{\mu}}$ and $\hat{\Sigma}$ have been defined in equation (3.4).

DEFINITION 1 *Let $\hat{\hat{\mu}}$ and $\hat{\Sigma}$ respectively be estimators of the first two moments of the distribution of the excess returns as in (3.4). Then, for a given (deterministic) benchmark c, the optimal p-value investment rule is defined as:*

$$\theta_p(c) = \sqrt{\frac{2\eta c}{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}}}}\frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\hat{\mu}} \tag{3.5}$$

Several comments are worth mentioning. First, the optimal p-value rule $\theta_p(c)$ is random and depends on the (chosen) estimates of the mean and variance of the excess returns distribution. However, this random rule (3.5) is the genuine rule that solves our optimization problem. In other words, our workable rule does not come from an additional (suboptimal) plug-in step (see also section 4). The deep reason for this exactness lies in the definition of our p-value selection method: the randomness of the problem precisely defines our selection procedure. Without uncertainty, there would not be any purpose to run a test and therefore no p-value maximization. Second, the rule (3.5) is a two-fund investment rule, just like the (feasible) mean-variance optimization problem $\hat{\theta}_{MV}$ (see equation (2.2)): both rules yield to the same repartition of wealth among the different financial risky assets. This allows us to reinterpret the p-value investor in terms of mean-variance behavior with a corrected risk-aversion parameter in section 4. Finally, note that the optimal p-value investment rule works for a given c. The next section naturally asks whether there exists an optimal benchmark or not.

3.2 An optimal choice for the benchmark?

The above selection method depends on the choice of the benchmark c: it represents the minimal level of portfolio performance the investor wants to guarantee with the highest possible level of confidence. In some sense, this benchmark is not a choice variable and we cannot genuinely define its optimal value. Recall that the choice of the benchmark reflects the degree of conservatism of the financial institution. However, it is still helpful to exhibit the *optimal* benchmark for comparison purposes.

We maximize the expected performance of the portfolio associated with the optimal p-value rule for a given benchmark:

$$\max_{c \geq 0} E\left[Q_P(\theta_p(c))\right]$$

The *optimal* benchmark c^* reads:

$$c^* = \frac{1}{2\eta} \times \frac{\left[E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\tilde{\mu}_0}{\sqrt{\hat{\gamma}^2}}\right)\right]^2}{\left[E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\Sigma_0\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{\hat{\gamma}^2}\right)\right]^2} \quad \text{where} \quad \hat{\gamma}^2 \equiv \hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}} \tag{3.6}$$

The optimal benchmark c^* is clearly unfeasible since it depends on the unknown parameters $\tilde{\mu}_0$ and Σ_0[9]. Interestingly enough, without estimation risk (or assuming $\tilde{\mu}_0$ and Σ_0 are known), we can check that the associated investment rule is numerically equal to the true mean-variance rule, which is also the optimal rule in absence of estimation risk. See also section 4.3.

4 Theoretical comparison with existing literature

This section is dedicated to the comparison of competing investment strategies after introducing the useful concept of *corrected* risk-aversion parameter, already considered in ter Horst *et al.* (2006).

4.1 Overview of some competing selection methods

This subsection briefly introduces some of the existing investment rules. In particular, we emphasize the different methodologies to account for estimation risk[10].

- Mean-variance (Markowitz (1959)) (see section 2): this rule selects the portfolio with the maximal mean-variance performance. The optimal allocation is unfeasible: it

[9] This is not really surprising since we maximize the expected performance for a given c.
[10] See also Kan and Zhou (2006).

depends on the first two unknown moments of the excess returns distribution. When some estimates (see equation (2.3)) of the unknowns are plugged into the formula, it becomes feasible. The estimation risk is then ignored. This rule is given by:

$$\hat{\theta}_{MV} = \frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

• Bayesian (Bawa, Brown and Klein (1979)): the Bayesian approach maximizes the expected performance of the portfolio where the expectation is computed according to the predictive distribution of the market. In turn, this predictive distribution is built from a combination of historical observations and the prior. Estimation risk is made explicit by considering the unknown parameters as random variables, described by the posterior distribution. However, it is not always clear how the prior can be chosen. Under the standard assumption of diffuse priors on both the mean and the variance of the excess returns, it can be shown that the Bayesian optimal portfolio weights are:

$$\theta_B = \frac{1}{\eta}\left(\frac{T-N-2}{T+1}\right)\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

• ter Horst, de Roon and Werker (2006): the portfolio weights are chosen to minimize the risk function based on the loss of replacing the true (unknown) mean of the portfolio by its sample estimate. They restrict their attention to the class of two-fund rules and ignore the estimation risk of the variance:

$$\theta_{HRW} = \frac{1}{\eta}\left(\frac{\gamma^2}{\gamma^2 + N/T}\right)\hat{\Sigma}^{-1}\hat{\hat{\mu}} \quad \text{with} \quad \gamma^2 = \tilde{\mu}_0'\Sigma_0^{-1}\tilde{\mu}_0$$

The resulting optimal rule θ_{HRW} is unfeasible: γ^2 is then replaced by its sample counterpart $\hat{\gamma}^2 = \hat{\hat{\mu}}\hat{\Sigma}^{-1}\hat{\hat{\mu}}$. Again, optimality is not guaranteed.

• Kan and Zhou (2006): they extend the previous selection method to incorporate the estimation risk of the variance:

$$\theta_{KZ2} = \frac{1}{\eta}\left(\frac{(T-N-4)(T-N-1)}{T(T-2)} \times \frac{\gamma^2}{\gamma^2 + N/T}\right)\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

Just like θ_{HRW}, the resulting optimal rule θ_{KZ2} is unfeasible: see appendix B for its feasible version. They also explore the class of three-fund investment rules by considering the sample global mean-variance portfolio. The associated optimal rule θ_{KZ3} is unfeasible as well: see also appendix B for additional details.

- Garlappi, Uppal and Wang (2006): they consider a model that allows for multi priors and where the investor is averse to ambiguity. The standard mean-variance framework is modified by adding a preliminary minimization step. A constraint restricts the expected return to fall into a confidence interval around its estimated value and recognizes the existence of estimation risk. The minimization over the possible expected returns subject to this constraint reflects the investor's aversion to ambiguity. While this approach has a solid axiomatic foundation, its sequentiality cannot be directly liked to an optimality criterion. The optimal rule θ_{GUW} is defined in appendix B.

The following theoretical rankings have been derived by Kan and Zhou (2006):

$$\theta_{KZ3} >> \theta_{KZ2} >> \theta_B >> \hat{\theta}_{MV} \ , \quad \theta_{KZ2} >> \theta_{HRW} \quad \text{and} \quad \theta_{KZ2} >> \theta_{GUW}$$

where $>>$ stands for "outperforms in terms of mean-variance performance". We argue that this ranking might not be guaranteed in practice (even in simple simulation frameworks where the returns are normally distributed) when θ_{HRW}, θ_{KZ2} and θ_{KZ3} are replaced by their feasible counterparts. Kan and Zhou (2006) already mentioned this issue when comparing their (feasible) optimal two-fund rule to the one of Garlappi *et al.*. See also section 5.

4.2 Corrected risk-aversion parameter

Despite their differences, most of the selection methods described above yield an optimal rule within the class of two-fund rules, just like the (feasible) Markowitz's mean-variance approach[11].

According to these rules, the same repartition of wealth among the different risky financial assets is recommended: their differences lie in the share of wealth invested in risky assets relative to the riskless asset. The (feasible) mean-variance rule can be reinterpreted as a function of the risk-aversion parameter η:

$$\hat{\theta}_{MV}(\eta) = \frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

Note that $\left[\hat{\Sigma}^{-1}\hat{\hat{\mu}}\right]$ defines how wealth is allocated among risky assets while η weights the share of wealth assigned to the risky assets: the greater η, the lower the (global) share to the risky assets.

We can then write each two-fund rule as a mean-variance rule with a *corrected* risk aversion parameter. In fact, any two-fund rule vector of weights θ_r can be rewritten as follows:

$$\theta_r = \hat{\theta}_{MV}(\tilde{\eta}) \qquad \text{for some} \quad \tilde{\eta} > 0 \tag{4.1}$$

Therefore, the behavior of any two-fund investor can be characterized in terms of a mean-variance associated to a new (corrected) risk-aversion parameter $\tilde{\eta}$. The following corrected risk-aversion parameters can be deduced for the two-fund rules discussed above[12]:

$$\begin{aligned}
\tilde{\eta}_{HRW} &= \eta\frac{\gamma^2 + N/T}{\gamma^2} \\
\tilde{\eta}_{KZ2} &= \eta\frac{\gamma^2 + N/T}{\gamma^2} \times \frac{T(T-2)}{(T-N-4)(T-N-1)} \\
\tilde{\eta}_{B} &= \eta\frac{T+1}{T-N-2} \\
\tilde{\eta}_{p}(c) &= \eta\sqrt{\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{2\eta c}}
\end{aligned}$$

[11]This is especially surprising for our p-value selection method since it does not come from any simplifying assumption (as for θ_{KZ2} and θ_{HRW}).

[12]We could also consider θ_{GUW} as a two-fund rule with a corrected risk-aversion parameter that can be infinite with non-zero probability.

4.3 Comparison of the reinterpreted investment rules

Our original mean-variance investor always becomes more risk-averse when applying any of the competing rules we consider here. However, this is not true when she applies the p-value rules: her risk-aversion parameter does not always increase.

On one hand, the investors respectively associated with the three competing rules θ_B, θ_{HRW} and θ_{KZ2} are always more risk-averse than the mean-variance investor. Moreover the following ranking can even be observed:

$$\tilde{\eta}_{KZ2} > \tilde{\eta}_{HRW} > \eta \quad \text{and} \quad \tilde{\eta}_B > \eta$$

Recall that θ_{KZ2} is nothing but θ_{HRW} where the additional estimation risk coming from the variance is accounted for. So one could be tempted to conclude that increasing the risk-aversion parameter is a sensible way to account for estimation risk.

On the other hand, the p-value corrected risk-aversion linearly depends on the original risk-aversion parameter: hence, the p-value investor might be characterized as a mean-variance investor either by increasing or decreasing the risk-aversion η. The corrected risk-aversion parameter can actually be rewritten as follows:

$$\tilde{\eta}_p(c) = \eta\sqrt{\frac{Q(\hat{\theta}_{MV})}{c}}$$

where $Q(\hat{\theta}_{MV})$ is the performance associated to the feasible mean-variance investment rule (see equation (2.2)). Depending on the choice of the benchmark c, one falls into one of the following cases:

(i) if $c = Q(\hat{\theta}_{MV})$ then $\tilde{\eta}_p = \eta$
(ii) if $c > Q(\hat{\theta}_{MV})$ then $\tilde{\eta}_p < \eta$
(iii) if $c < Q(\hat{\theta}_{MV})$ then $\tilde{\eta}_p > \eta$

Intuitively, this additional flexibility might be profitable, especially because it can be linked to the actual sample realizations. Consider an investor who chooses a moderate benchmark c. Assume now that, *by chance*, she faces a profitable financial environment (or a sample associated to a relatively high performance): likely $c < Q(\hat{\theta}_{MV})$ and so

$\tilde{\eta}_p > \eta$. Overall, the part invested in the risky assets is going to be lower. The profitable financial conditions offer additional safety to the p-value investor: it is more likely to beat the target. On the contrary, with a *not so good* financial environment, one may expect the investor to become less risk-averse, still hoping to defeat the benchmark. Hence, we found our risk-aversion parameter is counter-cyclical. In the literature, there is a lot of evidence of this feature: see Aydemir (2008), and Chavas and Holt (1996).

Intuitively, it makes sense to incorporate the sample-information into the decision process. The p-value selection method might also overcome the well-known problem of the mean-variance investment rule which takes extreme positions. The next section further investigates this.

5 Monte-Carlo study

This section presents the results of a simple Monte-Carlo study. The simulation exercise involves five risky assets and the riskless asset. The risky returns follow a multivariate normal distribution and the true model parameters are calibrated from monthly unhedged returns of stock indices for the G5 countries over the period January 1974 to December 1998. The G5 stock indices are the MSCI indices for France, Germany, Japan, the UK and the US as done, for instance, in ter Horst *et al.* (2006). Table 1 contains the summary statistics.

A financial strategy is considered over an investment horizon T_h. More precisely, at time $t = 1$ investors have access to T (past) historical observations of the financial returns. These are used to estimate the unknown parameters (typically $\tilde{\mu}_0$ and Σ_0) required to evaluate their investment strategy. The induced portfolio is hold for one period until $t = 2$, whereas the investment strategy is reevaluated using again the T most recent observations to build the estimators. A new portfolio is constructed, and so on until T_h.

We compare eleven investment strategies around two objects of interest for portfolio managers. First, we compare their respective performance over the investment horizon.

The performance is evaluated through the expected (one-period) mean-variance performance. Second, we compare the stability of the investment rules as measured by the transaction costs incurred to reallocate the portfolio at each period.

We consider the following investment rules: (1) the mean-variance optimal rule in absence of uncertainty θ_{MV}; (1f) the feasible counterpart of (1) $\hat{\theta}_{MV}$; (2) the optimal two-fund rule θ_{KZ2}; (2f) the feasible counterpart of (2) $\hat{\theta}_{KZ2}$; (3) the optimal two-fund rule when the variance is known θ_{HRW}; (3f) the feasible counterpart of (3) $\hat{\theta}_{HRW}$; (4) the Bayesian rule with diffuse prior θ_B; (5) the sequential min-max rule θ_{GUW}; (6) the optimal three-fund rule θ_{KZ3}; (6f) the feasible counterpart of (6) $\hat{\theta}_{KZ3}$; (7) the p-value rule for four different benchmarks. In this convenient Monte-Carlo setup, the benchmarks can be evaluated directly with respect to the maximal performance $Q(\theta_{MV})$. Typically, we consider here $c_1 = .1Q(\theta_{MV})$; $c_2 = .5Q(\theta_{MV})$; $c_3 = .9Q(\theta_{MV})$; and c^* the *optimal* benchmark (according to section 3.2) which is evaluated by simulation for each size of the rolling window (see table 2). In practice, one can think of at least two ways to get a convenient benchmark: c might be a numerical target that has been decided by the board of directors; c can also be based on the historical performance of some benchmark index.

We choose to set the risk-aversion parameter η equal to 5. For each portfolio rule r, defined by the vector of weights θ_r, the associated (one-period) expected performance is evaluated as follows[13]:

$$E\left[Q(\theta_r)\right] = E\left(\theta_r'\tilde{\mu}_0\right) - \frac{\eta}{2}E\left(\theta_r'\Sigma_0\theta_r\right) \tag{5.1}$$

where the true moments $\tilde{\mu}_0$ and Σ_0 are known (but only at this stage!) in our convenient Monte-Carlo framework: this helps isolate the effects of estimation risk.

Most of the above rules lead to a random vector of weights θ_r. Hence, this is not always possible to obtain a closed-form solution for the expected performance. If so, the performance is evaluated by averaging over many replications of the experiment. This is the case for the rules (2f), (3f), (5), (6f) and (7). For the remaining rules,

[13] Note that to simplify the notations we do not make explicit its dependence to the date of the investment

expected performances are formally provided in appendix B.

Table 3 provides the expected performances (in percentages per month) associated with every rule for several sample sizes of the rolling window used to calculate the estimators. Generally speaking, things get better when the sample size increases: i) the performance of each investment rule gets closer to the true optimal one $Q(\theta_{MV})$; ii) the feasible rules get closer to their theoretical counterparts - see also table 4; iii) finally, the estimation risk coming from the variance matters less when T increases. There is an additional loss of 15% per month when using $\hat{\theta}_{HRW}$ instead of $\hat{\theta}_{KZ2}$ for $T = 120$ and it drops to less than 1% when $T = 240$.
Figure 1 provides a visual comparison of the performances of all the feasible rules, as a function of the rolling window size. The dominance of the feasible three-fund rule is obvious. Hence, diversification appears to matter quite a bit when accounting for estimation risk. The p-value with a medium benchmark performs fairly well.
Figure 2 provides the same information for the p-value rules compared to their associated target. Note first that the rank of the expected performances is preserved: a low target is associated to a lower expected performance. Then, except for the highest target (chosen as 90% of the maximal theoretical performance), the minimal target is always ensured and indeed outperformed.
The performance of the p-value investment rule is positively surprising. Recall that compared to most of its competitors, it does not maximize the mean-variance performance. Of course, its performance crucially depends on the benchmark[14]. However, for quite a wide range of potential benchmarks ($c_1 = .035$ to $c_3 = .315$ percentages per month) the p-value performs quite well: the medium benchmark even outperforms $\hat{\theta}_{KZ3}$ when $T = 60$. It clearly outperforms the mean-variance, the Bayesian and the min-max sequential investment rules.

We now compare the stability of the portfolio rules via the transaction costs incurred to reallocate the portfolio at each period. This cost is the averaged amount (in arbitrary units) payed by the investor to modify her positions. The arbitrary cost is the same for

[14] Simulations with *unreasonable* targets, both very low and high, confirm this, and are not reported here.

each risky asset. More precisely, for each rule r defined by the vector of weights θ_t, at date t, the maintenance cost is defined as follows:

$$C_r = E\left[\sum_{t=1}^{T_h-1} |\theta_{t+1} - \theta_t|\,\iota\right] \sim \frac{1}{M}\sum_{m=1}^{M}\left[\sum_{t=1}^{T_h-1} \left|\theta_{t+1}^m - \theta_t^m\right|\iota\right] \tag{5.2}$$

where $\iota = [1\ 1\ \cdots\ 1]'$ and M is the number of replications.

Table 5 collects the average transaction costs for each rule, over an investment horizon $T_h = 60$ with several rolling window sizes and $M = 50000$ replications. Even though the transaction costs are calculated in (5.2) in a relatively basic and crude way, several comments are worth mentioning.

First, generally speaking, the transaction costs decrease when the size of the rolling window increases: the estimators naturally become more accurate and stable when the sample size increases, so do the financial positions. Only the most economical rule (θ_{GUW}) does not satisfy this. The reason comes from its definition. Contrary to any other investment rule, θ_{GUW} has a nonzero probability of entirely investing in the riskfree asset: this mechanically lowers its maintenance cost. Note that when the sample increases (and hence the estimators can be trusted more), the GUW-investor has higher transaction costs, meaning that she invests more in the risky assets.

The true (unfeasible) rules (θ_{KZ2}, θ_{HRW}, and θ_{KZ3}) tend to be more economical than their associated feasible counterparts. When the sample size increases, these rules, as well as most of the remaining ones, get closer to each others; this has already been noticed with the expected performance.

The p-value rules are naturally ranked as a function of their associated benchmark. More precisely, the lowest target c_1 yields a more economical rule $\theta_p(c_1)$. In order to fulfill her objective, this investor does not need to invest as much in the risky assets.

Finally the feasible rules can be ranked from the most economical as follows (the ranking does not depend on the size of the sample used to produce the estimators):

$$\theta_{GUW} >> \theta_p(c_1) >> \theta_p(c_2) \sim \hat{\theta}_{KZ3} >> \hat{\theta}_{KZ2} >> \theta_p(c_3) \sim \hat{\theta}_{HRW} >> \theta_B >> \hat{\theta}_{MV}$$

where $>>$ stands for "more economical" and $\sim$ for "economically equivalent".

6 Conclusion

In this paper we propose a new way to account for estimation risk when selecting the optimal portfolio. In sharp contrast with existing literature, the optimal portfolio is not defined as the one maximizing some mean-variance performance: we consider here a more conservative definition of optimality which focuses on guaranteeing some minimal performance. More precisely, our portfolio selection method is based on a one-sided test ensuring that the portfolio performance is above a given threshold. The optimal weights are then obtained from the maximization of the p-value associated to the above test. The test provides an integrated method to account for estimation risk. Moreover, after neglecting the estimation risk of the sample variance, it leads to a closed-form investment rule which can be used without requiring any additional (suboptimal) step.

Of course the performance of the p-value investment rule (which is not designed or meant to achieve the maximal performance) depends on the chosen benchmark c. However, as illustrated in our simple Monte-Carlo simulation study where we consider a wide range of benchmarks, the overall performance is quite satisfactory. In particular, it performs pretty well for relatively small samples (we believe mainly because it does not require an additional suboptimal plug-in step) and outperforms reasonable choices of targets. We find these preliminary results really encouraging.

The great advantage of the simple framework we consider here consists in providing closed-form optimal investment rules, interpretable in terms of mean-variance behavior. Compared to competing two-fund rules (e.g. Kan and Zhou (2007) and ter Horst, de Roon and Werker (2006)), we have shown that this is not always an increase of the original risk-aversion parameter that works to account for estimation risk.

For future research, several directions might be worth examining. First, one could extend our selection method to random targets. This would permit to track the performance of benchmark indices, rather than deterministic targets that may not always be inline with the financial environment. Second, considering that we generally do better than the feasible optimal two-fund rule and the very good results of the feasible three-fund rule, we may wonder how the p-value selection method, adapted to consider three-fund investment rules, would perform: as suggested by Kan and Zhou (2006), even more than three assets may help. Finally, recent papers have considered the related issue of model uncertainty. In particular, Cavadini, Sbuelz and Trojani (2001) extend the study of ter Horst, de Roon and Werker (2006) to incorporate model risk: they use robust inference methods *à la Huber*, or local deviations to the chosen initial distribution. Of course, the interpretability of the investment rules is likely the price to pay to consider such extensions.

Appendix

A Analogy with the Value-at-Risk

The Value-at-Risk (VaR) is a well-known financial risk measure summarizing the worst expected loss the investor is ready to accept. More precisely, the choice of a level of confidence $(1-\alpha)$ is associated to an α-quantile or $VaR(\alpha)$. When X represents the financial return of interest assumed normally distributed with mean μ and variance σ^2, we have:

$$\begin{aligned} P(X < -VaR_\alpha) = \alpha \quad &\Leftrightarrow \quad P\left(\frac{X-\mu}{\sigma} < \frac{-VaR_\alpha - \mu}{\sigma}\right) = \alpha \\ &\Leftrightarrow \quad \Phi\left(\frac{-VaR_\alpha - \mu}{\sigma}\right) = \alpha \end{aligned}$$

where $\Phi(.)$ is the cumulative distribution function of a standard gaussian random variable. So,

$$\begin{aligned} P(X < -VaR_\alpha) = \alpha \quad &\Leftrightarrow \quad \frac{-VaR_\alpha - \mu}{\sigma} = \Phi^{-1}(\alpha) \\ &\Leftrightarrow \quad -VaR_\alpha = \mu + \sigma\Phi^{-1}(\alpha) \end{aligned}$$

Reasonable values of α are small (and for sure < 0.5), so $\Phi^{-1}(\alpha) < 0$.

to be completed

B Additional results on other investment rules

These calculations were derived in Kan and Zhou (2006).

- Two-fund rule of Kan and Zhou:

$$\theta_{KZ2} = \frac{1}{\eta}\left[\left(\frac{(T-N-1)(T-N-4)}{T(T-2)}\right)\left(\frac{\gamma^2}{\gamma^2 + N/T}\right)\right]\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

with $\gamma^2 = \tilde{\mu}'\Sigma^{-1}\tilde{\mu}$. Kan and Zhou (2006) recommend the following feasible rule $\hat{\theta}_{KZ2}$ where γ^2 is replaced by

$$\hat{\gamma}_a^2 = \frac{(T-N-2)\hat{\gamma}^2 - N}{T} + \frac{2(\hat{\gamma}^2)^{N/2}(1+\hat{\gamma}^2)^{-(T-2)/2}}{TB_{\hat{\gamma}^2/(1+\hat{\gamma}^2)}(N/2,(T-N)/2)}$$

with $\hat{\gamma}^2 = \hat{\tilde{\mu}}'\hat{\Sigma}^{-1}\hat{\tilde{\mu}}$ and $B_x(a,b)$ is the incomplete beta function

$$B_x(a,b) = \int_0^x y^{a-1}(1-y)^{b-1}dy$$

- Three-fund rule of Kan and Zhou:

$$\theta_{KZ3} = \frac{c_3}{\eta}\left[\left(\frac{\psi^2}{\psi^2+N/T}\right)\hat{\Sigma}^{-1}\hat{\tilde{\mu}} + \left(\frac{N/T}{\psi^2+N/T}\right)\mu_g\hat{\Sigma}^{-1}\iota\right]$$

with

$$\begin{aligned}
\mu_g &= \frac{\iota'\Sigma^{-1}\tilde{\mu}}{\iota'\Sigma^{-1}\iota} \\
c_3 &= \left(\frac{T-N-4}{T}\right)\left(\frac{T-N-1}{T-2}\right) \\
\psi^2 &= (\mu-\mu_g\iota)'\,\Sigma^{-1}\,(\mu-\mu_g\iota)
\end{aligned}$$

Kan and Zhou (2006) recommend the following feasible rule $\hat{\theta}_{KZ3}$ where μ_g and ψ^2 are respectively replaced by:

$$\begin{aligned}
\hat{\mu}_g &= \frac{\hat{\tilde{\mu}}'\hat{\Sigma}^{-1}\iota}{\iota'\hat{\Sigma}^{-1}\iota} \\
\hat{\psi}_a^2 &= \frac{(T-N-1)\hat{\psi}^2-(N-1)}{T} + \frac{2(\hat{\psi}^2)^{(N-1)/2}(1+\hat{\psi}^2)^{-(T-2)/2}}{TB_{\hat{\psi}^2/(1+\hat{\psi}^2)}((N-1)/2,(T-N+1)/2)}
\end{aligned}$$

- Sequential min-max of Garlappi, Uppal and Wang:

$$\theta_{GUW} = \frac{1}{\eta}d\frac{T-1}{T}\hat{\Sigma}^{-1}\hat{\tilde{\mu}}$$

$$\text{where}\quad d = \begin{cases} 1-(\epsilon/\hat{\gamma}^2)^{1/2} & if \quad \hat{\gamma}^2 > \epsilon \\ 0 & if \quad \hat{\gamma}^2 \le \epsilon \end{cases} \qquad \text{with} \quad \epsilon = N\mathcal{F}^{-1}_{N,T-N}(p)/(T-N)$$

where $\mathcal{F}^{-1}_{N,T-N}$ is the inverse of the cumulative distribution function of a central F-distribution with $(N, T-N)$ degrees of freedom and p is a probability. We use $p = .99$ as suggested in Garlappi *et al.* (2006).

• Expected performances of the investment rules considered in section 5. See also Kan and Zhou (2006) for additional details.

(1) Parameter certainty optimal: $E_1 = \gamma^2/(2\eta)$.

(1f) Feasible counterpart of (1):

$$E_{1f} = k_1 \frac{\gamma^2}{2\eta} - \frac{NT(T-2)}{2\eta(T-N-1)(T-N-2)(T-N-4)}$$
$$\text{with } \ k_1 = \left(\frac{T}{T-N-2}\right)\left[2 - \frac{T(T-2)}{(T-N-1)(T-N-4)}\right]$$

(2) Optimal 2-fund rule:

$$E_2 = \frac{\gamma^2}{2\eta} \times \frac{(T-N-4)(T-N-1)}{(T-2)(T-N-2)} \times \frac{\gamma^2}{\gamma^2 + N/T}$$

(2f) Feasible counterpart of (2): E_{2f} must be evaluated by simulation.

(3) Optimal 2-fund rule with known variance:

$$E_3 = \frac{\gamma^2}{2\eta} \times \frac{\gamma^2}{\gamma^2 + N/T} \times \frac{T}{T-N-2}\left[2 - \frac{(T+N)(T-2)}{(T-N-1)(T-N-4)}\right]$$

(3f) Feasible counterpart of (3): E_{3f} must be evaluated by simulation.

(4) Bayesian with diffuse priors:

$$E_4 = k_2 \frac{\gamma^2}{2\eta} - \frac{NT(T-2)(T-N-2)}{2\eta(T+1)^2(T-N-1)(T-N-4)}$$
$$\text{with } \ k_2 = \frac{T}{T+1}\left[2 - \frac{T(T-2)(T-N-2)}{(T+1)(T-N-1)(T-N-4)}\right]$$

(5) Uncertainty aversion rule: E_5 must be evaluated by simulation.

(6) Optimal 3-fund rule:

$$E_6 = \frac{\gamma^2}{2\eta} \frac{(T-N-1)(T-N-4)}{(T-2)(T-N-2)}\left[1 - \frac{N/T}{\gamma^2 + \left(\frac{\gamma^2}{\psi^2}\right)\left(\frac{N}{T}\right)}\right]$$

(6f) Feasible counterpart of (6): E_{6f} must be evaluated by simulation.

(7) P-value maximization given c: $E_7(c)$ must be evaluated by simulation.

C Proofs of the main results

• Proof of equation (3.5) $\theta_p(c)$:

The first order conditions can be reinterpreted as a function of the (feasible) vector of the mean-variance weights $\hat{\theta}_{MV}$ defined in equation (2.2) as follows:

$$(\hat{\tilde{\mu}} - \eta\hat{\Sigma}\theta_p)\sqrt{\theta_p'\hat{\Sigma}\theta_p} - \frac{\theta_p'\hat{\tilde{\mu}} - \eta/2\theta_p'\hat{\Sigma}\theta_p - c}{\sqrt{\theta_p'\hat{\Sigma}\theta_p}}\hat{\Sigma}\theta_p = 0$$

$$\Leftrightarrow \quad \hat{\tilde{\mu}} - \eta\hat{\Sigma}\theta_p - \frac{\theta_p'\hat{\tilde{\mu}} - c}{\theta_p'\hat{\Sigma}\theta_p}\hat{\Sigma}\theta_p + \frac{\eta}{2}\hat{\Sigma}\theta_p = 0$$

$$\Leftrightarrow \quad \hat{\tilde{\mu}} - \frac{\eta}{2}\hat{\Sigma}\theta_p - \frac{\theta_p'\hat{\tilde{\mu}} - c}{\theta_p'\hat{\Sigma}\theta_p}\hat{\Sigma}\theta_p = 0$$

$$\Leftrightarrow \quad \hat{\Sigma}^{-1}\hat{\tilde{\mu}} - \frac{\eta}{2}\theta_p - \frac{\theta_p'\hat{\tilde{\mu}} - c}{\theta_p'\hat{\Sigma}\theta_p}\theta_p = 0$$

$$\Leftrightarrow \quad \eta \times \hat{\theta}_{MV} = \left[\frac{\theta_p'\hat{\tilde{\mu}} - c}{\theta_p'\hat{\Sigma}\theta_p} + \frac{\eta}{2}\right]\theta \tag{C.1}$$

Now for a given threshold c, we can always define a constant real number k_c such that:

$$k_c \times \eta = \frac{\theta_p'\hat{\tilde{\mu}} - c}{\theta_p'\hat{\Sigma}\theta_p} \tag{C.2}$$

Then, substituting (C.2) into (C.1) yields to:

$$\eta \times \hat{\theta}_{MV} = \left(k_c + \frac{1}{2}\right) \times \eta \times \theta_p \quad \Leftrightarrow \quad \theta_p = \frac{1}{k_c + 1/2}\frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\tilde{\mu}} \tag{C.3}$$

If I consider $\hat{\theta}_{MV}$ as a function of η such that $\hat{\theta}_{MV} = (\hat{\Sigma}^{-1}\tilde{R})/\eta$ and θ_p as a function of η (where θ_p is the weighting vector maximizing the p-value of the test with a parameter η of risk-aversion), then I get:

$$\theta_p(\eta) = \hat{\theta}_{MV}(\tilde{\eta}) \quad \text{with} \quad \tilde{\eta} = \eta \times (k_c + \frac{1}{2}) \tag{C.4}$$

The interpretation of $\tilde{\eta}$ as a *corrected* parameter of risk aversion is valid if and only if $k_c + 1/2 > 0$. This result may appear a bit *ad hoc* at first because k_c depends on the unknown vector of weights θ. But from equation (C.2), we are actually able to deduce its explicit expression as a function of known quantities only:

$$(C.2) \iff \theta_p'\hat{\hat{\mu}} - c = \theta_p'\hat{\Sigma}\theta_p \times k_c \times \eta$$

Then after replacing θ_p by its expression (C.3), we get:

$$\frac{1}{k_c+1/2}\frac{1}{\eta}\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}} - c = \frac{k_c}{(k_c+1/2)^2}\frac{1}{\eta}\hat{\hat{\mu}}\hat{\Sigma}^{-1}\hat{\hat{\mu}} \;\Rightarrow\; c = \frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{2\eta(k_c+1/2)^2}$$

$$\Rightarrow\; (k_c+1/2) = \sqrt{\frac{\hat{\hat{\mu}}'\hat{\Sigma}\hat{\hat{\mu}}}{2\eta}} \times \frac{1}{\sqrt{c}}$$

- Proof of equation (3.6) c^*:

$$Q(\theta_p(c)) \;=\; \theta_p'(c)\tilde{\mu}_0 - \frac{\eta}{2}\theta_p'(c)\Sigma_0\theta_p(c) = \frac{\sqrt{2c}}{\sqrt{\eta\hat{\gamma}^2}}\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\tilde{\mu}_0 - \frac{c}{\hat{\gamma}^2}\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\Sigma_0\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

where $\hat{\gamma}^2 = \hat{\hat{\mu}}'\hat{\Sigma}^{-1}\hat{\hat{\mu}}$. We then maximize it with respect to c:

$$\max_{c\geq 0} E\left[Q_P(\theta_p(c))\right]$$

The associated first order conditions are:

$$\frac{1}{\sqrt{2c^*\eta}}E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\tilde{\mu}_0}{\sqrt{\hat{\gamma}^2}}\right) = E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\Sigma_0\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{\hat{\gamma}^2}\right) \implies c^* = \frac{1}{2\eta} \times \frac{\left[E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\tilde{\mu}_0}{\sqrt{\hat{\gamma}^2}}\right)\right]^2}{\left[E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\Sigma_0\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{\hat{\gamma}^2}\right)\right]^2}$$

The associated optimal vector of weights is the following:

$$\theta_p(c^*) = \frac{\left|E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\tilde{\mu}_0}{\sqrt{\hat{\gamma}^2}}\right)\right|}{\left|E\left(\frac{\hat{\hat{\mu}}'\hat{\Sigma}^{-1}\Sigma_0\hat{\Sigma}^{-1}\hat{\hat{\mu}}}{\hat{\gamma}^2}\right)\right|}\frac{1}{\sqrt{\hat{\gamma}^2}}\frac{1}{\eta}\hat{\Sigma}^{-1}\hat{\hat{\mu}}$$

Note in particular that if $\tilde{\mu}_0$ and Σ_0 were known, we would get $\theta_p(c^*) = \theta_{MV}$, which corresponds to the best portfolio rule in absence of estimation risk.

References

[1] V. Bawa, S. Brown, and R. Klein, *Estimation Risk and Optimal Portfolio Choice*, North Holland, Amsterdam, 1979.

[2] M. Brandt, *Portfolio Choice Problems*, in Handbook of Financial Econometrics (2004), Y. Aït-Sahalia and L.P. Hansen, eds.

[3] F. Cavadini, A. Sbuelz, and F. Trojani, *A simplified way of incorporating Model Risk, Estimation Risk and Robustness in Mean Variance Portfolio Management*, Working paper (2001).

[4] D. Cho, *Uncertainty in Second Moments: Implications for Portfolio Allocation*, Working paper (2003), SUNY at Buffalo.

[5] L. Garlappi, R. Uppal, and T. Wang, *Portfolio selection with Parameter and Model Uncertainty: A Multi-Prior Approach*, The Review of Financial Studies (2006), forthcoming.

[6] C.R. Harvey, J. Liechty, M. Liechty, and P. Müller, *Portfolio Selection with Higher Moments*, Working paper (2004), Duke university.

[7] J. Horst, F. de Roon, and B. Werker, *An Alternative Approach to Estimation Risk*, in Advances in Corporate Finance and Asset Pricing, Elsevier Amsterdam (2006), Luc Renneboog, eds.

[8] P. Jorion, *Value-at-Risk: the new Benchmark for controlling Market Risk*, second ed., McGraw-Hill, Chicago, Irwin, 1996.

[9] R. Kan and G. Zhou, *Optimal Portfolio Choice with Parameter Uncertainty*, Journal of Financial and Quantitative Analysis (2006), forthcoming.

[10] H.M. Markowitz, *Portfolio Selection: Efficient Diversification of Investments*, Wiley, New-York, 1959.

[11] A. Meucci, *Risk and Asset Allocation*, Springer-Verlag, New-York, 2005.

[12] M. Stutzer, *Portfolio Choice with Endogenous Utility: a Large Deviations Approach*, Journal of Econometrics **116** (2003), 365–386.

[13] A. Zellner and V.K. Chetty, *Prediction and Decision Problems in Regression Models from the Bayesian Point of View*, Journal of the American Statistical Association **60** (1065), 608–616.

	Mean	Standard deviation
France	0.014	0.069
Germany	0.013	0.059
Japan	0.011	0.067
UK	0.015	0.073
USA	0.012	0.044

and $\rho_0 = \begin{pmatrix} 1 & .590 & .390 & .541 & .456 \\ & 1 & .338 & .424 & .347 \\ & & 1 & .342 & .221 \\ & & & 1 & .506 \\ & & & & 1 \end{pmatrix}$

Table 1: Summary statistics and matrix of correlations for the MSCI of G5 countries over the period January 1974 to December 1998.

	Size of the rolling window T				
	60	120	180	240	300
Optimal benchmark c^*	0.0765	0.1415	0.1840	0.2124	0.2325

Table 2: Optimal benchmark c^* (in percentages per month) for several sample sizes of the rolling window used to calculate the estimators of the first two moments of the portfolio distribution. c^* has been evaluated by simulation with $M = 50000$ replications.

Rule		Size of the rolling window T				
		60	120	180	240	300
(1)	θ_{MV}	0.3503	0.3503	0.3503	0.3503	0.3503
(1f)	$\hat{\theta}_{MV}$	-0.8977	-0.1667	0.0256	0.1146	0.1648
(2)	θ_{KZ2}	0.0929	0.1518	0.1888	0.2141	0.2326
(2f)*	$\hat{\theta}_{KZ2}$	-0.0046	0.1033	0.1510	0.1832	0.2067
(3)	θ_{HRW}	0.0741	0.1418	0.1816	0.2084	0.2279
(3f)*	$\hat{\theta}_{HRW}$	-0.2577	0.0518	0.1371	0.1813	0.2090
(4)	θ_B	0.0028	0.0014	0.0010	0.0007	0.0006
(5)*	θ_{GUW}	0.0036	0.0121	0.0223	0.0356	0.0511
(6)	θ_{KZ3}	0.2827	0.3007	0.3074	0.3113	0.3140
(6f)*	$\hat{\theta}_{KZ3}$	0.0266	0.1770	0.2274	0.2530	0.2683
(7)*	$\theta_p(c_1)$	0.0835	0.1167	0.1333	0.1439	0.1509
(7)*	$\theta_p(c_2)$	0.0690	0.1545	0.1949	0.2204	0.2374
(7)*	$\theta_p(c_3)$	-0.0050	0.1190	0.1761	0.2117	0.2352
(7)*	$\theta_p(c^*)$	0.0933	0.1564	0.1950	0.2223	0.2417

Table 3: Expected performances (in percentages per month) for several sample sizes of the rolling window used to calculate the required estimators of the first two moments of the portfolio distribution. A star (*) identifies a rule whose expected performance has been evaluated through a simulation with $M = 50000$ replications. For $\hat{\theta}_{KZ2}$ and $\hat{\theta}_{KZ3}$ we follow the recommendations of Kan and Zhou (2006); for θ_{GUW} we follow Garlappi *et al.* (2006). The benchmarks are chosen as $c_1 = 0.3503$, $c_2 = .1751$ and $c_3 = .3153$; for the optimal c^*, see table 2.

T	60		120		180		240		300	
$\hat{\theta}_{KZ2}$	105.0	(101.3)	32.0	(70.5)	20.0	(56.9)	14.4	(47.7)	11.1	(41.0)
$\hat{\theta}_{HRW}$	447.8	(173.6)	63.5	(85.2)	24.5	(60.9)	13.0	(48.2)	8.3	(40.3)
$\hat{\theta}_{KZ3}$	90.6	(92.4)	41.1	(49.5)	26.0	(35.1)	18.7	(27.8)	14.5	(23.4)

Table 4: Expected performance losses (in percentages per month) when using the feasible rule instead of its theoretical counterpart. For convenience, we also report in parentheses the loss of using the feasible rule instead of the true optimal one θ_{MV}.

Rule		Size of the rolling window T				
		60	120	180	240	300
(1)	θ_{MV}	0	0	0	0	0
(1f)	$\hat{\theta}_{MV}$	27.9694	12.6632	8.1790	6.0430	4.7878
(2)	θ_{KZ2}	6.5508	5.1686	4.2363	3.5860	3.1047
(2f)	$\hat{\theta}_{KZ2}$	11.3447	6.5516	4.8417	3.9102	3.3040
(3)	θ_{HRW}	8.2777	5.7838	4.5618	3.7894	3.2443
(3f)	$\hat{\theta}_{HRW}$	18.5680	8.8222	5.9800	4.5882	3.7469
(4)	θ_B	24.3013	11.8256	7.8175	5.8424	4.6606
(5)	θ_{GUW}	0.5518	0.5340	0.6264	0.6942	0.7478
(6)	θ_{KZ3}	3.8310	2.0374	1.4739	1.2097	1.0583
(6f)	$\hat{\theta}_{KZ3}$	12.7672	6.4304	4.3444	3.3113	2.6818
(7)	$\theta_p(c_1)$	4.6841	2.7357	1.9583	1.5330	1.2595
(7)	$\theta_p(c_2)$	10.4739	6.1172	4.3789	3.4278	2.8162
(7)	$\theta_p(c_3)$	14.0521	8.2071	5.8749	4.5989	3.7784
(7)	$\theta_p(c^*)$	6.9220	5.4983	4.4881	3.7747	3.2447

Table 5: Average transaction costs over an investment horizon $T_h = 60$. We consider several rolling window sizes (to evaluate the estimators of the first two moment of the distributions of the returns) and $M = 50000$ replications.

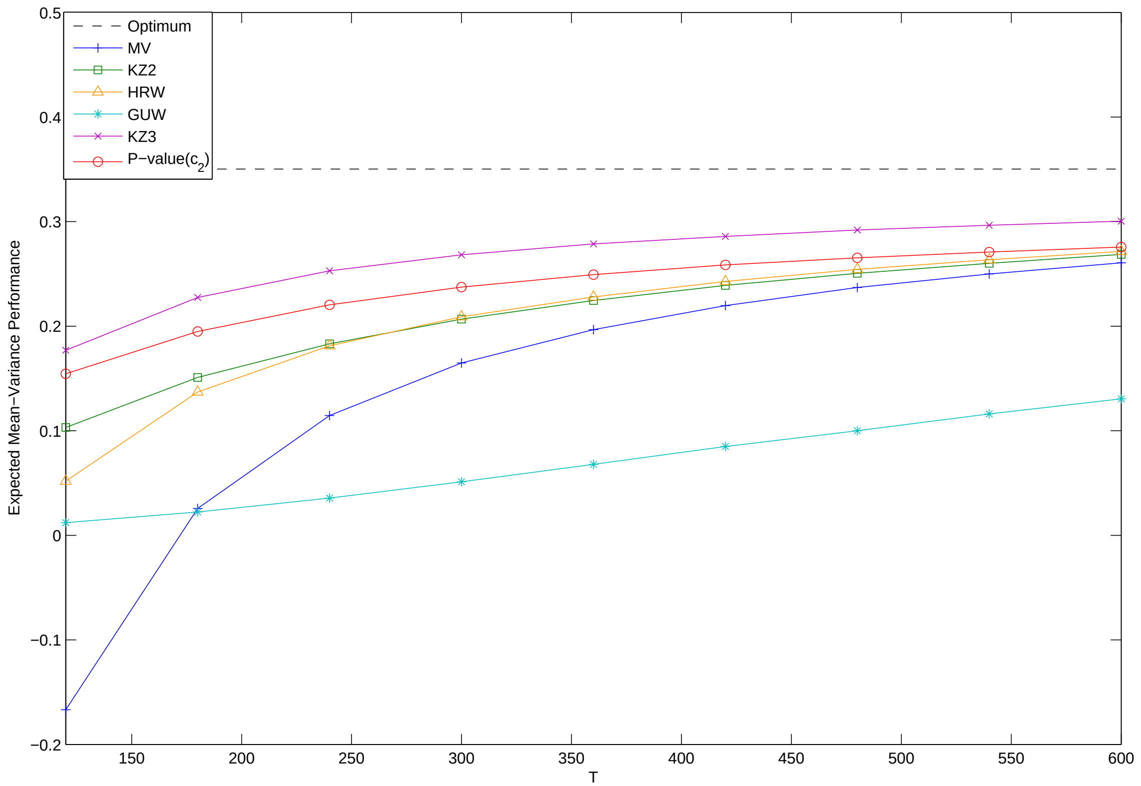

Figure 1: Expected performances (in percentages per month) for several feasible investment rules as a function of the size of the rolling window used to calculate the required estimators.

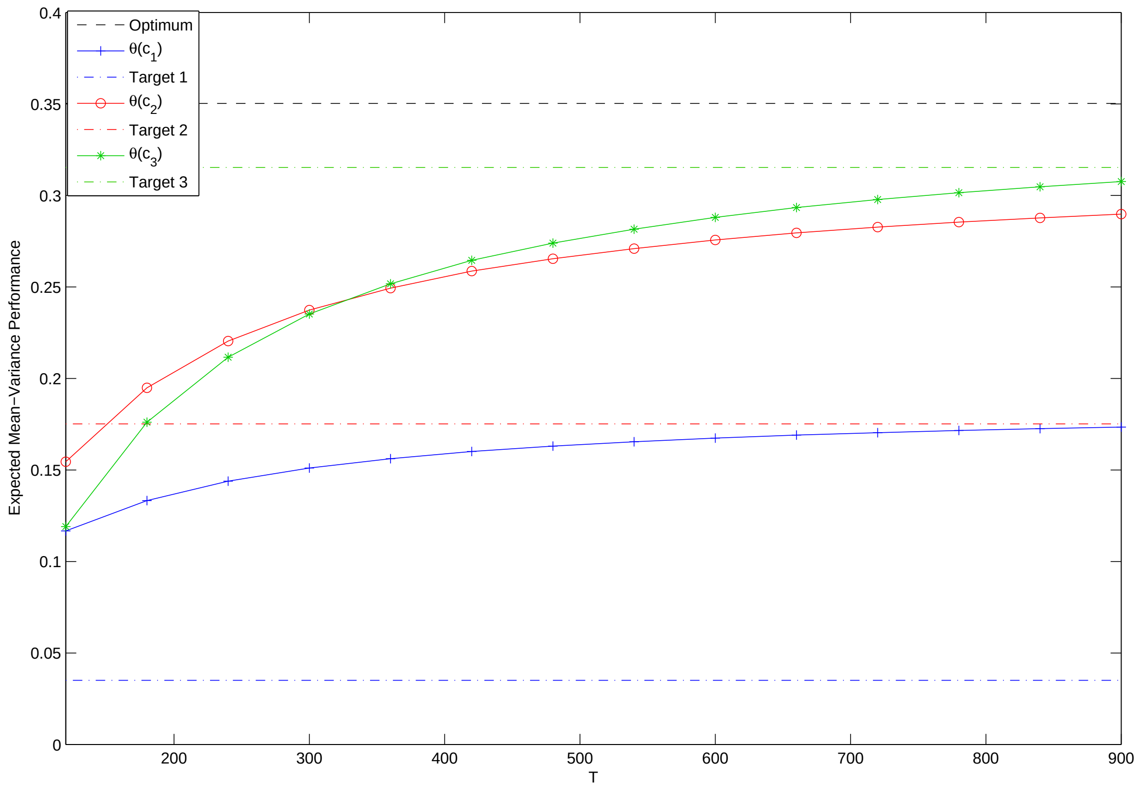

Figure 2: Expected performances (in percentages per month) for the p-value investment rules with several benchmarks as a function of the size of the rolling window used to calculate the required estimators.