

MPRA

Munich Personal RePEc Archive

A Simple Model of Robust Portfolio Selection

Taboga, Marco

1 June 2004

Online at <https://mpra.ub.uni-muenchen.de/16472/>
MPRA Paper No. 16472, posted 29 Jul 2009 20:13 UTC

A Simple Model of Robust Portfolio Selection

Marco Taboga*
(mtaboga@fastwebnet.it)
Università di Torino

January 2004 (last revision June 2004)

Abstract

We propose a single-period portfolio selection model which allows the decision maker to easily deal with uncertainty about the distribution of asset returns. The model is preference-based and relies upon a separate parametrization of risk aversion and ambiguity aversion. A particular specification of preferences allows us to solve the portfolio selection problem and obtain a simple closed-form expression for the portfolio weights, which lends itself to a straightforward economic interpretation.

*I thank E. Castagnoli, P. Klibanoff, E. Luciano, F. Maccheroni, M. Marinacci, S. Mukerji for helpful comments and suggestions.

Traditional static portfolio selection models, such as Markowitz's (1952), assume that the distribution of asset returns is objectively known to the decision maker. Of course, this is hardly a realistic assumption, since at best investors can estimate the distribution of asset returns with limited precision. Typically, one assumes that the distribution of asset returns belongs to a given parametric family and then estimates the relevant parameters using historical data, hence being left with some uncertainty about the true parameters. Although the lack of realism of an assumption is not sufficient to disqualify a model, many researchers have argued that uncertainty about the distribution of asset returns cannot be neglected, since the optimal portfolios derived within traditional models seem to be highly sensitive to small changes in the parameters characterizing the distribution.

We adopt a Bayesian framework to model parameter uncertainty, basing the portfolio selection process on a specification of preferences which allows a simple and separate parametrization of the investor's attitudes towards parameter uncertainty and investment risk.

We assume that the investor maximizes the following objective function, recently axiomatized by Klibanoff, Marinacci and Mukerji (2003):

$$E_{\pi} \left[\varphi \left(E_{\mu} \left[u \left(\widetilde{W} \right) \right] \right) \right]$$

where \widetilde{W} is stochastic future wealth, φ and u are increasing and concave functions and μ and π are probability measures. The inner expectation (w.r.t. the measure μ) is akin to a von-Neumann Morgenstern expected utility. The investor realizes that to different distributions of asset returns correspond different values of the expected utility of wealth. The measure π assigns second-order probabilities to the distributions which are deemed plausible by the investor and hence to the different values taken by the expected utility. Instead of simply averaging over expected utilities, the decision maker applies a concave transform before taking the outer expectation, hence displaying aversion to those situations in which parameter uncertainty leads to a great

dispersion of expected utility values.

Since the Seventies a wealth of studies documenting and discussing the importance of parameter uncertainty in portfolio selection models began to appear: among them Frankfurter, Phillips and Seagle (1971), Barry (1974), Bawa and Klein (1976) and Jobson and Korkie (1980). Also, a widespread opinion came about that portfolios constructed using sample moments of returns often involve very extreme positions, which are far from being truly optimal: Green and Hollifield (1992) provide a rigorous analysis of this claim.

As a follow up to these critiques some proposals were made to improve upon Markowitz's (1952) model: for example, Frost and Savarino (1988) and Black and Litterman (1992) propose Bayesian procedures to improve the performance of estimated optimal portfolios. A very recent work by Jagannathan and Ma (2003) also gives a contribution to the debate: they discuss how imposing constraints on portfolio weights can sometimes improve the performance of estimated optimal portfolios.

The aforementioned studies are aimed at devising methods of correcting Markowitz's (1952) optimal allocation rule, in order to improve its performance in the presence of parameter uncertainty. Our paper is aimed at a different direction: adopting a Bayesian perspective, we propose a model in which parameter uncertainty is easily and directly quantifiable and becomes a further variable of the decision problem, besides risk and return.

Our paper belongs to a strand of the literature which tackles the problem of parameter uncertainty in portfolio selection by specifying the decision-maker preferences both towards investment risk and towards parameter uncertainty and by finding allocation rules consistent with these preferences. The main idea underlying most studies is that portfolio selection problems lend themselves to be analyzed with decision-theoretic tools which deviate from the orthodoxy of the subjective expected utility framework. In these models a behavioral distinction is made between risk and uncertainty. The distinction goes back to Knight (1921), who identified two kinds of uncer-

tainty: the first one is risk, arising in situations where the decision maker is able to unambiguously assign probabilities to the events which are relevant to her decision; the second one is ambiguity or Knightian uncertainty, arising when the probabilities of some events are not unanimously and uniquely agreed upon. A portfolio selection problem whose parameters are not objectively known would seem to fall into the latter of these two categories. Such a view is undoubtedly prone to many criticisms, since the standard model of decision-making, whose foundations were laid out by Savage (1954), does not allow for a meaningful distinction between risk and ambiguity: according to that model, the fact that an objectively given distribution of asset returns is not available to the investor does not prevent her from forming a unique subjective probability distribution, upon which she can base her decisions. A consequence of Savage's postulates is that information is never too imprecise to be summarized by probabilities. However, many theorists have argued that in some situations Savage's model does not allow to distinguish differences in the quality and the abundance of information used to form subjective probability distributions: Ellsberg's (1961) paradox is one of the pieces of evidence usually cited in support of this view.

Several modelling devices have been used to incorporate ambiguity into financial decision problems. A taxonomy has been proposed by Uppal and Wang (2003), who identify two main classes of models.

A first class of models is inspired by Gilboa and Schmeidler's (1989) model of multiple priors preferences. Among them, Epstein and Wang (1994) extend the Lucas (1978) general equilibrium model with multiple priors preferences and derive asset pricing and equilibrium implications; Dow and Werlang (1992) set up a static portfolio choice problem with a single risky (and ambiguous) asset and a single agent with Schmeidler and Gilboa's preferences; Chen and Epstein (2002) develop a continuous-time model and the considerable simplifications afforded by continuous-time modelling allow them to derive closed formulas for the optimal portfolio allocation.

A second class of models makes use of some tools borrowed from the robust control literature: Maenhout (2004) adapts a framework developed by Anderson, Hansen and Sargent (2003) to derive a portfolio selection model set in continuous time, where the decision maker has got a preference for robustness; Maenhout's (2004) model, which extends the classical Merton's (1990) model, assumes that the decision maker has got a reference probability measure over asset returns, but she considers also alternative probability measures, equivalent to the reference measure (in the probabilistic sense of equivalence), and she chooses among these measures according to a penalty function based on the relative entropy between the probability measures. Uppal and Wang (2003) extend Maenhout's model to take into account multiple sources of uncertainty and shed some light on real world phenomena such as underdiversification and the home bias. Both Maenhout and Uppal and Wang provide closed formulas for the optimal portfolios.

Among the models we have cited, none provides an asset allocation rule for the simplest case of a one-period horizon and a single investor. There are, however, some studies in this direction: Krasker (1982), for example, analyzes the implications of minimax behavior for portfolio choice; in his model the agent minimizes over a set of probability measures obtained by ε -contamination of a reference measure; Krasker shows how some commonly held portfolios (portfolios without short positions or portfolios replicating the market portfolio) can be rationalized by individual minimax behavior, rather than on equilibrium considerations. Also Becker, Marty and Rustem (2000) analyze minimax portfolio strategies: the focus of their research is more on computational aspects and they limit their attention to the case where the investor is able to identify a finite number of scenarios, where by scenario they mean a possible combination of means and variances for the assets to be included in the portfolio.

The model we propose is a static portfolio selection model and does not belong to either of the two classes identified by Uppal and Wang (2003), since

we use still another specification of preferences. Our investor maximizes the objective function¹:

$$E_{\pi} \left[\varphi \left(E_{\mu} \left[u \left(\widetilde{W} \right) \right] \right) \right]$$

where \widetilde{W} is stochastic future wealth, φ and u are increasing and concave functions and μ and π are probability measures. The investor does not know which probability measure μ truly describes the world, but there is a whole set of probability measures μ which she thinks to be plausible descriptions of the world; to each of these probability measures μ she subjectively assigns a degree of likelihood, according to another probability measure π . The idea behind the above specification of preferences is very simple: when the investor selects a measure μ among all the possible measures, she is able to integrate the utility of wealth with respect to that measure and hence to calculate the expected utility of future wealth $E_{\mu} \left[u \left(\widetilde{W} \right) \right]$; in general, the expected utility value thus calculated depends on the particular measure μ she selects and she subjectively assigns to these expected utility values different degrees of likelihood, according to the measure π . Instead of simply averaging over the possible expected utility values, by setting $U \left(\widetilde{W} \right) = E_{\pi} \left[E_{\mu} \left[u \left(\widetilde{W} \right) \right] \right]$, in order to rank stochastic future wealth, the investor applies a concave transformation φ to expected utility values, with the consequence, akin to what happens in standard risk theory, that she dislikes mean-preserving spreads in expected utility values.

We show that ranking stochastic future wealth according to this two-stage valuation procedure leads to a very simple necessary condition for the optimality of a portfolio of assets. We also derive a connection between this first order condition and the first order condition which arises when a standard subjective expected utility function is maximized.

Choosing a particular combination of the two functions φ and u and of

¹In the paper by Klibanoff, Marinacci and Mukerji (2003) the role of π and μ is reversed. We have chosen to use the symbol μ to denote the inner measure to keep consistency with the finance literature on portfolio choice, since in a later part of the paper there will be an identification of μ with the vector of expected returns.

the two distributions μ and π , we analyze thoroughly the issue of uncertainty about the expected returns on the assets to be included in the portfolio. We perform some comparative statics and some experiments to understand how changing the degree of parameter uncertainty and individual attitudes towards ambiguity affects the optimal portfolio. Some of the comparative statics are carried out analytically, resorting to some recent advances in matrix perturbation theory, while some are carried out numerically. The results we obtain from our comparative statics exercises are qualitatively similar to those obtained by Maenhout (2004) and Uppal and Wang (2003), although concentrating on a single-period problem we obtain richer characterizations. The main implications of our model is that an increase in ambiguity about expected returns or in ambiguity aversion induces the investor to form a less aggressive portfolio of risky assets, with the consequence that both the expected return and the standard deviation of the portfolio are reduced. While in Uppal and Wang (2003) the reduction in portfolio weights is proportional, in our model the vector of portfolio weights is altered in a nonlinear fashion. Furthermore, we are able to understand how introducing ambiguity aversion influences the efficiency of selected portfolios.

The paper is organized as follows: Section 1 describes the portfolio selection model and some of its general properties. Section 2 specializes the model to a tractable case, which allows explicit derivation of the portfolio weights. Section 3 describes some empirical results obtained with the model described in section 2. Section 4 concludes the paper.

1 The portfolio selection model

We consider the one-period allocation problem of an agent who has to decide how to invest her wealth W_0 at time 0, dividing it among $n + 1$ assets. The gross return on the i -th asset after one period is a random variable denoted

by R_i . The $(n \times 1)$ vector of the returns on the first n assets is denoted by R and the $(n \times 1)$ vector of portfolio weights, indicating the fraction of wealth invested in each of the first n assets is denoted by w . The end-of-period wealth is denoted by \widetilde{W} and is equal to:

$$\widetilde{W} = W_0 \left[R_{n+1} + w' \left(R - \vec{1} R_{n+1} \right) \right]$$

where $\vec{1}$ is a column vector of ones of dimension n . The above definition of \widetilde{W} implicitly accomodates the requirement that the portfolio weights sum up to unity.

We assume that there are no frictions of any kind: securities are perfectly divisible; there are no transaction costs or taxes; the agent is a price-taker, in that she believes that her choices do not affect the distribution of asset returns; there are no institutional restrictions, so that the agent is allowed to buy, sell or short sell any desired amount of any security (this assumption can be weakened, by simply requiring that at an optimum institutional restrictions are not binding).

Let $(\Omega, \sigma(\Omega), \pi)$ be a measure space and assume that for each $\omega \in \Omega$ we are given a measure $\mu(\omega; \cdot)$ on $\mathcal{B}(\mathbb{R}^{n+1})$. Assume also that, for each $B \in \mathcal{B}(\mathbb{R}^{n+1})$, $\mu(\omega; B)$ is $\sigma(\Omega)$ -measurable. Both π and $\mu(\omega; \cdot)$ are assumed to be probability measures. As a consequence, there exists a probability measure m defined on $\sigma(\Omega) \times \mathcal{B}(\mathbb{R}^{n+1})$ such that:

$$m(A \times B) = \int_A \mu(\omega; B) d\pi(\omega) \quad \forall A \in \sigma(\Omega), B \in \mathcal{B}(\mathbb{R}^{n+1})$$

Furthermore, if $f \in L^1(\Omega \times \mathbb{R}^{n+1})$, then the function

$$\omega \rightarrow \int_{\mathbb{R}^n} f(\omega, r) d\mu(\omega; r)$$

is π -a.s. well-defined, belongs to $L^1(\Omega)$ and the conditional version of Fubini's

theorem (see Ash and Doléans Dade (1999)) ensures that:

$$\int_{A \times B} f(\omega, r) dm(\omega, r) = \int_A \int_B f(\omega, r) d\mu(\omega; r) d\pi(\omega)$$

for any $A \in \sigma(\Omega)$ and $B \in \mathcal{B}(\mathbb{R}^{n+1})$.

We identify the return on the i -th asset R_i with the i -th component of the element extracted from the sample space $(\mathbb{R}^{n+1}, \mathcal{B}(\mathbb{R}^{n+1}))$, so that each conditional probability measure $\mu(\omega; \cdot)$ on $\mathcal{B}(\mathbb{R}^{n+1})$ can be interpreted as a distribution of asset returns. We assume that $R_i \in L^2(\Omega \times \mathbb{R}^{n+1})$ and $R_i \in L^2(\mathbb{R}^{n+1})$ for each i and each $\mu(\omega; \cdot)$ in a set of π -measure 1.

The investor chooses the portfolio weights w in order to solve the following maximization problem²:

$$\max_w \mathbb{E}_\pi \left[\varphi \left(\mathbb{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right) \right] \quad (1)$$

which is a shorthand for:

$$\max_w \int_\Omega \varphi \left(\int_{\mathbb{R}^{n+1}} u \left(\widetilde{W}(r) \right) d\mu(\omega; r) \right) d\pi(\omega)$$

As we have anticipated in the introduction, the above objective function, first axiomatized as a specification of preferences by Klibanoff, Marinacci and Mukerji (2003), is conceptually very simple. Uncertainty about future asset returns cannot be described by a unique probability distribution, but the investor is able to identify a set of probability measures which could be plausible descriptions of the randomness inherent in the asset allocation problem. Each probability measure $\mu(\omega; \cdot)$ yields a von-Neumann Morgenstern expected utility value $\mathbb{E}_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right]$, where u is concave, strictly increasing and finite-valued. Since u is concave and strictly increasing and \widetilde{W} is linear in the asset returns, then $u \left(\widetilde{W} \right) \in L^1(\Omega \times \mathbb{R}^{n+1})$, $u \left(\widetilde{W} \right) \in L^1(\mathbb{R}^{n+1})$ and

²See note (1)

$E_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right]$ is a $\sigma(\Omega)$ -measurable function belonging to $L^1(\Omega)$. By considering all the probability measures $\mu(\omega; \cdot)$ for ω in Ω , we obtain a whole range of expected utility values; formally, we have a mapping $U : \Omega \rightarrow \mathbb{R}$ defined by:

$$U(\omega) = E_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right]$$

The investor subjectively assigns a degree of likelihood to the probability measures $\mu(\omega; \cdot)$, that is she forms a subjective measure π , assigning a probability to a σ -algebra $\sigma(\Omega)$ of subsets of Ω . This is compatible, for example, with a Bayesian framework in which the distribution of asset returns is parametrized by a vector θ and the investor assigns a non-degenerate probability distribution to the parameter vector θ .

Since $U(\omega)$ is a measurable function and belongs to $L^1(\Omega)$, the investor is able to evaluate the integral:

$$E_{\pi} [\varphi(U)]$$

where φ is again taken to be a concave, strictly increasing and finite-valued function. The above expectation is an "expected utility of expected utilities": if π is non-degenerate, the expected utility U is a random variable; the investor, instead of simply maximizing the expected value $E_{\pi}[U]$ of the possible values of U , maximizes the expectation $E_{\pi}[\varphi(U)]$ of a concave transform of U . By doing so, she incorporates ambiguity aversion into her objective function: this is a consequence of the fact that by Jensen's inequality:

$$E_{\pi} [\varphi(U)] \leq \varphi(E_{\pi}[U])$$

so that a setting in which the expected utility of wealth is equal to $E_{\pi}[U]$ without ambiguity is preferred to one in which the expectation of the expected utility of wealth has got the same value, but some ambiguity is present.

If the maximization problem (1) has got an interior solution and if both

u and φ are differentiable, the following first order necessary condition must be satisfied:

$$\mathbb{E}_\pi \left[\varphi' \left(\mathbb{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right) \mathbb{E}_\mu \left[u' \left(\widetilde{W} \right) \left(R - \overline{1} R_{n+1} \right) \right] \right] = 0 \quad (2)$$

where we have assumed that differentiation under the integral sign is legitimate and the dominated convergence theorem applies: for example it suffices to assume that $u' \left(\widetilde{W} \right)$ is bounded $\mu(\omega; \cdot)$ -a.s. for any ω in a set of π -measure 1 and $\varphi' \left(\mathbb{E}_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right] \right)$ is bounded π -a.s.

It is worth commenting briefly on this necessary condition: when there is no ambiguity it simplifies to

$$\mathbb{E}_\mu \left[u' \left(\widetilde{W} \right) \left(R - \overline{1} R_{n+1} \right) \right] = 0 \quad (3)$$

which has a simple economic interpretation: if the portfolio allocation is optimal, the marginal increase in utility obtained by selling one dollar worth of asset $n + 1$ and investing the proceeds into any one of the other assets, must have zero expected value. When ambiguity is present, i.e. π is non-degenerate, condition (3) does not necessarily hold for every $\mu(\omega; \cdot)$; there might be some probability measures $\mu(\omega; \cdot)$ under which a reallocation of the portfolio increases the expected utility value $\mathbb{E}_\mu \left[u \left(\widetilde{W} \right) \right]$ calculated under those measures; however, once one accounts for the effect of the reallocation on the whole range of expected utility values, the overall marginal benefit of the reallocation must be zero.

Imposing very mild conditions on the structure of the problem (1), the optimality condition (2) can be written in a form which is very similar to (3). We state the result in the following proposition, which we prove in the Appendix:

Proposition 1 *Let φ and u be of class C^1 . Assume at an optimum $u' \left(\widetilde{W} \right)$ is bounded $\mu(\omega; \cdot)$ -a.s. for any ω in a set of π -measure 1 and $\varphi' \left(\mathbb{E}_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right] \right)$*

is bounded π -a.s. Then, there exists a probability measure π^* , equivalent to π , such that condition (2) can be written as:

$$E_{\pi^*} \left[E_{\mu} \left[u' \left(\widetilde{W} \right) \left(R - \overline{1} R_{n+1} \right) \right] \right] = 0$$

or as:

$$E_P \left[u' \left(\widetilde{W} \right) \left(R - \overline{1} R_{n+1} \right) \right] = 0$$

where $P = \pi^* \times \mu$.

Thus, it is possible to find a probability measure P (an unambiguous distribution of asset returns) such that solving the problem (1) is equivalent to solving the classical portfolio optimization problem:

$$\max_w E_P \left[u \left(\widetilde{W} \right) \right] \quad (4)$$

Although solving (4) is equivalent to solving (1), the formulation (1) is not superfluous since it allows the decision maker to recover the measure P constructively. Furthermore, inspection of the proof shows that the measure π^* is problem dependent, since its Radon-Nikodym derivative with respect to the measure π is a function of the return to the optimal portfolio:

$$\frac{d\pi^*}{d\pi} = \frac{\varphi' \left(E_{\mu} \left[u \left(\widetilde{W} \right) \right] \right)}{E_{\pi} \left[\varphi' \left(E_{\mu} \left[u \left(\widetilde{W} \right) \right] \right) \right]}$$

A qualitative analysis of the Radon-Nykodym derivative also provides some insights on the optimal behavior of the agent: when φ is concave φ' is decreasing; this implies that the original measure is distorted by subtracting weight from those measures μ which yield high expected utility values and adding more weight to those measures which yield low expected utility values.

In the next section we will analyse a special case in which, besides being able to find closed formulas for the portfolio weights, we are able to derive P explicitly.

2 A tractable case

To implement the model described in the previous section, one needs to specify the functional forms of u and φ and the distributions μ and π . We propose a specification which allows to derive the optimal portfolio weights explicitly and gain some insight into the model. We first make the distributional assumptions: for convenience, we specialize to the case in which the $(n + 1)$ -th asset is a risk-free asset, so that R_{n+1} is equal to a constant R_f ; we further assume that the vector R of returns is jointly normally distributed with $E[R] = \mu$ and $\text{Var}[R] = \Sigma$ (it will soon become clear why we have chosen to assign the symbol μ both to the expected return and to the inner probability measure). However, the vector μ of expected returns is not known with certainty by the investor, who assigns the following distribution to the parameter μ :

$$\mu \sim N(m, S)$$

In the language of Bayesian statistics, this can be thought of as the prior distribution of the parameter. Adopting the Bayesian perspective, the joint distributional hypothesis on the parameter μ and on the vector R can be written as:

$$\begin{aligned} f(R|\mu) &= (2\pi)^{-\frac{n}{2}} |\Sigma|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(R - \mu)' \Sigma^{-1} (R - \mu)\right) \\ f(\mu) &= (2\pi)^{-\frac{n}{2}} |S|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mu - m)' S^{-1} (\mu - m)\right) \end{aligned} \quad (5)$$

As to the covariance matrix Σ , the fact that we assume that it is known with certainty might seem unrealistic; however, this is not obviously so, since it is often possible to sample the price processes generating the returns R at

frequencies higher than the frequency considered for calculating the returns. When sampling at higher frequencies is possible and the price increments are serially independent, we can obtain more precise estimates of Σ , without adding any precision to our estimate of μ ; for a discussion of this point see, for example, Gouriéroux and Jasiak (2001), Campbell and Viceira (2002) or Merton (1990). Another reason for considering Σ known, besides keeping the model tractable, is that, when the matrix Σ is known, $f(\mu)$ is conjugate for $f(R|\mu)$, which means that, after observing some sample data and updating the prior distribution of μ , its posterior distribution remains normal.

The choice of the functions u and φ characterizing the investor's preferences turns out to be quite delicate. We make for u a choice which is standard in one-period portfolio choice, i.e. we choose a CARA specification:

$$u(\widetilde{W}) = -e^{-\tau\widetilde{W}}$$

where τ is the coefficient of absolute risk aversion.

Since $u(\widetilde{W})$ takes on only negative values, so does the inner expectation $E_\mu[u(\widetilde{W})]$ in (1). As a consequence, the function φ , taking $E_\mu[u(\widetilde{W})]$ as an argument, must be defined on the negative part of the real line. We propose the following specification:

$$\varphi(U) = -\frac{(-U)^{\zeta+1}}{\zeta+1} \tag{6}$$

where ζ is a positive parameter and φ is defined on the negative part of the real axis. φ is both strictly increasing and strictly concave when $\zeta > 0$, since $\varphi'(U) > 0$ and $\varphi''(U) < 0$. Note that

$$\zeta = \frac{\varphi''(U)}{\varphi'(U)}U \tag{7}$$

so that ζ could be interpreted as a coefficient of relative ambiguity aver-

sion, akin to an Arrow and Pratt's coefficient of relative risk aversion; note, however, that, when compared to an Arrow and Pratt's coefficient, a minus sign is missing in front of the right hand side of (7): the reason is that U here is negative, while in the standard definition it is positive.

To gain a better understanding of the specification of φ , (6) can be written as:

$$\begin{aligned}\varphi(U) &= \frac{(\xi(U))^{1-(\zeta+2)}}{1-(\zeta+2)} \\ \xi(U) &= -\frac{1}{U}\end{aligned}$$

ξ is a strictly monotone transformation and when U ranges in the interval $(-\infty, 0)$, $\xi(U)$ takes values in the interval $(0, \infty)$. Once this monotone transformation is accomplished, the ξ values are retransformed with a standard power function with coefficient $\zeta + 2$. Roughly speaking, the choice of $\zeta + 2$ as an exponent instead of ζ is just a concavity adjustment done to keep the coefficient of relative ambiguity aversion equal to ζ . Mapping utility values from the interval $(-\infty, 0)$ to the interval $(0, \infty)$ via the ξ function also allows us to switch from a scale of measurement $((-\infty, 0))$ which does not possess an absolute zero, to a scale $((0, \infty))$ which does: this is extremely important since defining relative ambiguity aversion calls for a relative scale and a relative scale is well-defined only when an absolute zero can be identified.

Another feature of the proposed specification of φ about which it is worth commenting is that when $\zeta = 0$ the objective function reduces to:

$$\mathbb{E}_\pi \left[\mathbb{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right]$$

i.e. to the case in which the investor simply averages over expected utility values, so that this specification also encompasses the case of no aversion to ambiguity. This is equivalent to a Von-Neumann Morgenstern expected utility where integration is with respect to the predictive distribution of the

returns, as in Barberis (2000).

From now on, since the measure μ in (1) is completely identified by the vector of expected return, we will use the symbol μ to denote both of them and we will write the objective function, consistently with (5), as:

$$\mathbf{E}_\pi \left[\varphi \left(\mathbf{E} \left[u \left(\widetilde{W} \right) \mid \mu \right] \right) \right]$$

Now, denote by $R_w = R_f + w' \left(R - \vec{1} R_f \right)$ the return on the portfolio with weights w . We have:

$$\begin{aligned} \mathbf{E}[R_w \mid \mu] &= R_f + w' \left(\mu - \vec{1} R_f \right) \\ \text{Var}[R_w \mid \mu] &= w' \Sigma w \end{aligned}$$

The maximization problem is:

$$\max_w \mathbf{E}_\pi \left[\varphi \left(-\mathbf{E} \left[e^{-\tau R_w} \mid \mu \right] \right) \right] \quad (8)$$

where the initial wealth W_0 has been standardized to unity without loss of generality.

Since

$$-\mathbf{E} \left[e^{-\tau R_w} \mid \mu \right] = -e^{-\tau \mathbf{E}[R_w \mid \mu] + \frac{1}{2} \tau^2 \text{Var}[R_w \mid \mu]}$$

the objective function becomes:

$$\mathbf{E}_\pi \left[-\frac{1}{\zeta + 1} e^{-(\zeta+1)\tau [R_f + w'(\mu - \vec{1} R_f)] + \frac{1}{2}(\zeta+1)\tau^2 w' \Sigma w} \right]$$

and, calculating the expected value:

$$-\frac{1}{1 + \zeta} e^{-(\zeta+1)\tau [R_f + w(m - \vec{1} R_f)] + \frac{1}{2}(\zeta+1)\tau^2 w' \Sigma w + \frac{1}{2}(\zeta+1)^2 \tau^2 w' S w}$$

After taking monotone transformations, the maximization problem can

be rewritten as:

$$\max_w R_f + w' \left(m - \vec{1} R_f \right) - \frac{1}{2} \tau w' \Sigma w - \frac{1}{2} (\zeta + 1) \tau w' S w$$

The first order condition is:

$$m - \vec{1} R_f - \tau \Sigma w - (\zeta + 1) \tau S w = 0$$

which yields the optimal portfolio weights:

$$w^* = \frac{1}{\tau} [\Sigma + (1 + \zeta) S]^{-1} \left(m - \vec{1} R_f \right) \quad (9)$$

The above formula for the optimal portfolio weights is better understood when compared to the formula that arises in the case where parameter uncertainty is not taken into account by the decision maker. Suppose that the investor solves the standard optimization problem

$$\max_w -\mathbb{E} \left[e^{-\tau R_w} \right] \quad (10)$$

disregarding the fact that the distribution of μ is non-degenerate and setting $\mu = m$: the optimal portfolio weights would be:

$$w^{**} = \frac{1}{\tau} \Sigma^{-1} \left(m - \vec{1} R_f \right) \quad (11)$$

The optimal portfolio (9) differs from (11) only for the fact that the matrix $(1 + \zeta) S$ has been added to the covariance matrix Σ . According to Proposition 1, we have been able to find an unambiguous distribution of asset returns (a measure P) such that solving (4) is equivalent to solving (1); such a distribution is:

$$R \sim N(m, \Sigma + (1 + \zeta) S)$$

To gain some intuition about the correction thus made to the variance

covariance matrix and to the portfolio weights, consider the case in which Σ and S are diagonal (there is no covariance between asset returns and no covariance between the parameters). If parameter uncertainty is not taken into account the optimal weight given to a generic asset i is:

$$w_i^{**} = \frac{1}{\tau} \frac{m_i - R_f}{\sigma_i^2} \quad (12)$$

where w_i is the weight of asset i in the optimal portfolio and m_i and σ_i^2 are the mean and the variance of the gross return on asset i . The weight is proportional to the expected return in excess of the risk-free rate and inversely proportional to the variance of the return; the coefficient of proportionality is the risk tolerance parameter $\frac{1}{\tau}$.

When parameter uncertainty is taken into account, the optimal weight given to asset i is:

$$w_i^* = \frac{1}{\tau} \frac{m_i - R_f}{\sigma_i^2 + (1 + \zeta) s_i^2} \quad (13)$$

where s_i^2 is the variance of the parameter μ_i . Compared to the previous case, the portfolio weight is reduced when there is parameter uncertainty (s_i^2 is positive) and the higher the ambiguity aversion coefficient ζ , the greater this reduction is. This result is in accordance with Uppal and Wang's (2003) findings: although their portfolio selection model is dynamic and set in continuous time, it is comparable to ours because they also assume that the only ambiguity is about expected returns (drifts); in their model all the probability measures can be recovered from each other via a Girsanov change of measure, hence leaving the volatility unchanged. What they find is that in the presence of ambiguity and with a preference for robustness, the weight given to an ambiguous asset is smaller than the weight that would be given to the same asset in a traditional intertemporal Merton model; the role played in our model by the coefficient of ambiguity aversion is akin to the role played in their model by the degree of confidence in the reference measure.

Resorting to results from matrix perturbation theory, one can also do

some comparative statics on the unconditional expected return and the unconditional variance of the optimal portfolio.

Note that the first two unconditional moments of the return on the optimal portfolio are:

$$\begin{aligned} E[R_{w^*}] &= R_f + w^{*'} \left(m - \vec{1} R_f \right) \\ \text{Var}[R_{w^*}] &= E[\text{Var}[R_{w^*} | \mu]] + \text{Var}[E[R_{w^*} | \mu]] = w^{*'} (\Sigma + S) w^* \end{aligned}$$

The sensitivity of the two above moments to changes in the ambiguity aversion coefficient ζ is characterized by the following proposition:

Proposition 2 *Let R_{w^*} be the return on the optimal portfolio, selected according to the rule (9). Then the unconditional expected return $E[R_{w^*}]$ on the optimal portfolio is a decreasing function of the coefficient of ambiguity aversion ζ . Furthermore, its derivative is:*

$$\frac{\partial E[R_{w^*}]}{\partial \zeta} = -\frac{1}{\tau} \sum_{i=1}^n \sum_{j=1}^n \frac{1}{\lambda_i \lambda_j} u'_i S u_j u'_j \left(m - \vec{1} R_f \right) \left(m - \vec{1} R_f \right)' u_i$$

where λ_i and u_i are the eigenvalues and the eigenvectors of the matrix $A = \Sigma + (1 + \zeta) S$. In the special case in which $\zeta = 0$ also the unconditional variance of the return on the optimal portfolio is a locally decreasing function of ζ and:

$$\frac{\partial \text{Var}[R_{w^*}]}{\partial \zeta} = \frac{1}{\tau} \frac{\partial E[R_{w^*}]}{\partial \zeta}$$

We have not been able to find the derivative of $\text{Var}[R_{w^*}]$ for the more general case in which $\zeta \neq 0$. Furthermore, we would like to be able to say something about the impact of introducing (or increasing) ambiguity aversion on the efficiency of the portfolios selected by the investor, for example by examining the behavior of the Sharpe ratio:

$$\text{Sh}(R_w) = \frac{E[R_w] - R_f}{\sqrt{\text{Var}[R_w]}}$$

For this reason we have conducted some simulations (whose results are commented in the next section) to understand more about the comparative statics of the optimal portfolio.

3 Empirical results

In this section we conduct an empirical analysis of the portfolio allocation rule proposed in the previous section.

We assume that returns are generated according to (5). R is taken to be a vector of monthly returns on 30 stocks. The covariance matrix Σ is set equal to the sample covariance matrix of the monthly returns on the 30 stocks included in the Dow Jones Industrial Average Index, calculated from a sample of 120 monthly returns (January 1994 to December 2003).

We impose the following structure on m and S :

$$\begin{aligned} m &= \bar{m} \vec{1} \\ S &= \bar{s}^2 I + \bar{c} \left(\vec{1} \vec{1}' - I \right) \end{aligned} \tag{14}$$

where \bar{m} , \bar{s} and \bar{c} are scalars and \bar{c} is constrained to lie in the open interval $\left(-\frac{\bar{s}^2}{n-1}, \bar{s}^2 \right)$ to ensure that S is positive definite (we include in the Appendix a proof of the fact that this condition is both necessary and sufficient). Thus, we require that both the expectation and the variance of the parameters μ_i are equal across assets; we also require that all the covariances between the parameters are equal. This is the same structure assumed by Frost and Savarino (1986), although in a slightly different setting. It has the advantage that by varying a single parameter (\bar{c}) one is able to analyse the impact of a change in the covariance structure on the portfolio statistics. We will later deal with a more general structure.

As a first step, a reference model is defined by choosing a particular set of values for the parameters. We set $\tau = 10$, $\zeta = 2$, $\bar{m} = 1 + 0.5\%$, $\bar{s} = 0.3\%$,

$\sqrt{\bar{c}} = 0.1\%$ and $R_f = 1 + 0.2\%$ (note that the returns are monthly returns and the optimization horizon is the month). Comparative statics are performed by varying one parameter at a time over a range of values. The figures in the Appendix report the results of this exercise, displaying the impact of changing the parameters on the unconditional expected return $E[R_{w^*}]$, on the unconditional standard deviation $Sd[R_{w^*}] = \sqrt{\text{Var}[R_{w^*}]}$, on the Sharpe ratio $\text{Sh}[R_{w^*}]$ and on the sum of absolute portfolio weights:

$$a = \sum_{i=1}^{30} |w_i|$$

Increasing the coefficient of ambiguity aversion ζ decreases the expected return (as predicted by Proposition 2) and the standard deviation. Also the Sharpe ratio slightly declines, while the sum of absolute portfolio weights decreases substantially. Increasing \bar{s} or \bar{c} seems to produce the same effect as an increase in ζ . Hence, an increase in ambiguity or in ambiguity aversion induces the investor to form a less aggressive portfolio of risky assets (as reflected in the fact that the sum of absolute portfolio weights diminishes), with the consequence that both the expected return and the standard deviation of the portfolio are reduced.

Since the results presented so far could depend on the fact that we have assumed a very particular structure for m and S , we perform some simulations in which m and S assume a different structure. Roughly speaking, the idea behind the simulations is to take a prior as in (14), generate consistently a parameter μ at each simulation, then generate a path of asset returns and update the prior based on the observed sample, thus losing the structure (14).

We perform four sets of 100,000 simulations: a first set is performed with the reference parameters detailed above and three more sets are performed choosing different values for \bar{m} , \bar{s} and \bar{c} , as reported in the tables in the Appendix.

We start each simulation by randomly generating a value for μ , drawn from the distribution with density

$$f(\mu) = (2\pi)^{-\frac{n}{2}} |S|^{-\frac{1}{2}} \exp((\mu - m)' S^{-1} (\mu - m))$$

Conditioning on this value of μ , we simulate a sample of 24 monthly returns $\{R_t, t = 1, \dots, 24\}$, generated independently from the distribution with density

$$f(R_t|\mu) = (2\pi)^{-\frac{n}{2}} |\Sigma|^{-\frac{1}{2}} \exp((R_t - \mu)' \Sigma^{-1} (R_t - \mu))$$

and we use the returns thus simulated to update the prior distribution according to the Bayesian updating formulas:

$$f(\mu|R_1, \dots, R_{24}) = (2\pi)^{-\frac{n}{2}} |S_1|^{-\frac{1}{2}} \exp((\mu - m_1)' S_1^{-1} (\mu - m_1))$$

where

$$\begin{aligned} m_1 &= (S^{-1} + 24\Sigma^{-1})^{-1} (S^{-1}m + 24 \cdot \Sigma^{-1}\bar{R}) \\ S_1 &= (S^{-1} + 24 \cdot \Sigma^{-1})^{-1} \end{aligned}$$

and \bar{R} is the sample average of the returns R_t .

After updating the prior, we use it to calculate the optimal portfolio w^* which solves the maximization problem (8):

$$w^* = w^*(\bar{R}) = \frac{1}{\tau} [\Sigma + (1 + \zeta) S_1]^{-1} (m_1 - \vec{1} R_f)$$

where the notation $w(\bar{R})$ has been chosen to emphasize the dependence of the optimal portfolio on the observed average returns.

We calculate the conditional expected return on the portfolio:

$$\mathbb{E}[R_{w^*}|\mu] = w^{*'} (\mu - \vec{1} R_f) + R_f$$

its conditional variance:

$$\text{Var} [R_{w^*} | \mu] = w^{*\prime} \Sigma w^*$$

and the sum of absolute portfolio weights. All the above quantities are calculated conditioning on the true value of μ , which is not known to the investor and they are calculated repeatedly, changing the value of the ambiguity aversion coefficient ζ . We report in the Appendix the results obtained when ζ takes values in the set $\{0, 2, 4, 8, 16, 32\}$. We report also the results obtained following the standard asset allocation rule:

$$w^{**} = \Sigma^{-1} \left(m_1 - \vec{1} R_f \right)$$

This is, for example, the rule followed by Frost and Savarino (1986), with the difference that in their simulation exercise Σ is unknown and they use a Normal-Wishart prior on (μ, Σ) .

We also include in the tables a measure of the distance of the selected portfolio from the frontier of efficient portfolios. Note that ex-post (when μ is perfectly known) all minimum variance portfolios have a stock component which is a scalar multiple of the portfolio:

$$w^{***} = \Sigma^{-1} \left(\mu - \vec{1} R_f \right)$$

We propose to measure the distance of a generic portfolio w from the frontier of minimum variance portfolios as

$$\delta(w) = \min_{\lambda} \|w - \lambda w^{***}\|_2 \tag{15}$$

where λ is a scalar and $\|\cdot\|_2$ is the Euclidean norm. The efficient portfolio $\lambda^* w^{***}$ which solves the above minimum norm problem is the orthogonal

projection of w on the minimum variance frontier, which corresponds to

$$\lambda^* = \frac{w'w^{***}}{w^{***}w^{***}}$$

Since it can happen that a portfolio is projected on the inefficient part of the minimum variance frontier ($\lambda^* < 0$), we choose in those cases to set $\lambda^* = 0$ (in the simulations we have performed this happened in no more than 0.12% of all cases), so that we are rather measuring the distance of the portfolios from the efficient frontier, which is not a linear space.

After performing the simulations, the following sample moments are computed to estimate the unconditional moments:

$$\begin{aligned}\widehat{\mathbb{E}}[R_{w^*}] &= \frac{1}{N} \sum_{k=1}^N \mathbb{E}[R_{w^*} | \mu_k] \\ \widehat{\text{Var}}[R_{w^*}] &= \frac{1}{N} \sum_{k=1}^N \text{Var}[R_{w^*} | \mu_k] + \frac{1}{N} \sum_{i=1}^N \left(\mathbb{E}[R_{w^*} | \mu_k] - \widehat{\mathbb{E}}[R_{w^*}] \right)^2 \\ \widehat{\mathbb{E}}[a] &= \frac{1}{N} \sum_{k=1}^N a_k \\ \widehat{\mathbb{E}}[\delta(w^*)] &= \frac{1}{N} \sum_{k=1}^N \delta(w_k^*)\end{aligned}$$

where N is the number of simulations and the subscript k has been added to indicate the value obtained in simulation k .

These sample moments are displayed in Tables I, II, III and IV, along with an estimate of the Sharpe ratio:

$$\widehat{\text{Sh}}(R_{w^*}) = \frac{\widehat{\mathbb{E}}[R_{w^*}] - R_f}{\sqrt{\widehat{\text{Var}}[R_{w^*}]}}$$

The results emerging from the simulations are not qualitatively different

from those that emerged from the previous simpler comparative statics exercises: the more the ambiguity aversion coefficient ζ is increased, the less aggressive the portfolio is; increasing ζ leads to a reduction of the expected return, of the standard deviation of returns and of the sum of absolute portfolio weights. A novel feature is that, when ζ is increased, the Sharpe ratio remains virtually unchanged, but the average distance from the efficient frontier diminishes, so that these two measures provide conflicting evidence on the effect of increased ambiguity aversion on portfolio efficiency: while the behavior of the Sharpe ratio would suggest that more ambiguity aversion does not lead to improvements in efficiency, the diminishing distance from the efficient frontier seems to be a sign of improved robustness. The impact on portfolio statistics of switching from a traditional portfolio rule to the portfolio rule we propose is similar to the impact just described of an increase in ambiguity aversion.

Table V reports some features of the empirical distribution of the conditional expected return on the optimal portfolio: the expected return $E[R_{w^*}|\mu]$ is itself a random variable, being a function of the parameter μ and of the portfolio selected, and hence it takes a different value at each simulation, giving rise to an empirical distribution of conditional expected returns, which approximates the true distribution. The standard deviation of this distribution decreases as the coefficient of ambiguity aversion increases, so that more ambiguity aversion leads to a less volatile conditional expected return across scenarios. It is also interesting that the first percentile of the distribution increases when the coefficient of ambiguity aversion increases: this means that, although the unconditional expected return is lower when ambiguity aversion is increased, so that the distribution of conditional expected returns is shifted to the left, its left tail (comprising those cases in which the conditional expected return of the portfolio is disappointingly low), is considerably shortened and moved to the right.

4 Conclusions

We have addressed the problem of optimally selecting a portfolio of assets when the probabilistic distribution of asset returns is not known with precision. Many researchers argue that standard portfolio allocation rules are not robust, in the sense that even slight departures from the true distribution of asset returns may lead to select portfolios that are far from being optimal. We have proposed a model which allows to incorporate individual attitudes towards ambiguity (uncertainty about the true distribution of asset returns) into the portfolio selection process: we have done this by using a specification of preferences which to our knowledge has never been used in asset allocation models and which allows to easily parametrize both risk aversion and aversion to ambiguity.

We have fully analysed a special case of our model where we are able to explicitly derive the optimal portfolio weights when expected returns are uncertain. This is in a sense complementary to a recent work by Jagannathan and Ma (2003) which deals with uncertainty about the covariance matrix of returns.

Our model has both normative and positive implications: besides giving a portfolio allocation rule, it yields sharp predictions about the behavior of an ambiguity averse individual; an ambiguity averse individual reduces her position in those assets whose expected returns are uncertain and the greater the uncertainty and the aversion to uncertainty, the greater the reduction is.

We have conducted some simulations to gain a better understanding of the implications of our portfolio rule. Each time an investor is confronted with a different data set to estimate the same distribution of asset returns, he chooses a different portfolio, with a different expected return: with our allocation rule the variability of portfolio expected returns seems to be lower than with a standard allocation rule.

Although we have thoroughly investigated only a tractable case, the framework set up in the paper is general enough to leave space for further

research: a possible direction of future efforts could be to find an allocation rule for the more general case in which both the means and the variances of asset returns are uncertain and to understand the interaction between these two kinds of ambiguity.

References

- [1] Anderson, E. W., L. P. Hansen and T. J. Sargent, 2003, "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk and Model Detection", *Journal of the European Economic Association*, 1, 68-123.
- [2] Ash, R. and C. Doléans-Dade, 1999, *Probability and Measure Theory*, Harcourt Academic Press, San Diego.
- [3] Barberis, N., 2000, "Investing for the Long Run when Returns are Predictable", *Journal of Finance*, 55, 225-264.
- [4] Barry, C. B., 1974, "Portfolio Analysis under Uncertain Mean, Variances and Covariances", *Journal of Finance*, 29, 515-522.
- [5] Bawa, V. S. and R. W. Klein, 1976, "The Effect of Estimation Risk on Optimal Portfolio Choice", *Journal of Financial Economics*, 3, 215-231.
- [6] Becker B., W. Marty and B. Rustem, 2000, "Robust Min-Max Portfolio Strategies for Rival Forecast and Risk Scenarios", *Journal of Economic Dynamics and Control*, 24, 1591-1621.
- [7] Black, F. and R. Litterman, 1992, "Global Portfolio Optimization", *Financial Analyst Journal*, 48, 28-43.
- [8] van Bossum, M. A., 2002, *Semidefinite Optimization, a Spectral Approach*, unpublished PhD Dissertation, University of Utrecht.
- [9] Campbell, J. Y. and L. M. Viceira, 2002, *Strategic asset allocation*, Oxford University Press, New York.
- [10] Chen, Z. and L. Epstein, 2002, "Ambiguity, Risk and Asset Returns in Continuous Time", *Econometrica*, 70, 1403-1443.

- [11] Dow, J. and S. R. da Costa Werlang, 1992, "Uncertainty Aversion, Risk Aversion and the Optimal Choice of Portfolio", *Econometrica*, 60, 197-204.
- [12] Ellsberg, D., 1961, "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics*, 75, 643-669.
- [13] Epstein, L. and T. Wang, 1994, "Intertemporal Asset Pricing under Knightian Uncertainty", *Econometrica*, 62, 283-322.
- [14] Frankfurter, G. M., H. E. Phillips and J. P. Seagle, 1971, "Portfolio Selection: The Effect of Uncertain Mean, Variances and Covariances", *Journal of Financial and Quantitative Analysis*, 6, 1251-1262.
- [15] Frost, P. A. and J. Savarino, 1988, "An Empirical Bayes Approach to Efficient Portfolio Selection", *Journal of Financial and Quantitative Analysis*, 21, 293-305.
- [16] Gouriéroux, C. and J. Jasiak, 2001, *Financial econometrics: problems, models and methods*, Princeton University Press, Princeton.
- [17] Green, R. C. and B. Hollifield, 1992, "When will Mean-Variance Efficient Portfolios be Well Diversified?", *Journal of Finance*, 47, 1785-1809.
- [18] Jagannathan, R. and T. Ma, 2003, "Risk Reduction in Large Portfolios - Why Imposing the Wrong Constraints Helps", *Journal of Finance*, 58, 1651-1684.
- [19] Jobson, J. D. and B. Korkie, 1980, "Estimation for Markowitz Efficient Portfolios", *Journal of the American Statistical Association*, 75, 544-554.
- [20] Klibanoff, P., M. Marinacci and S. Mukerji, 2003, "A smooth model of decision making under ambiguity", ICER Working Paper no. 11/2003.
- [21] Knight, F. H., 1921, *Risk, Uncertainty and Profit*, Houghton Mifflin, Boston.

- [22] Krasker, W. S., 1982, "Minimax Behavior in Portfolio Selection", *Journal of Finance*, 37, 609-614.
- [23] Lucas, R. E., 1978, "Asset Prices in an Exchange Economy", *Econometrica*, 46, 1429-1445.
- [24] Maenhout, P. J., 2004, "Robust Portfolio Rules and Asset Pricing", forthcoming in *Review of Financial Studies*.
- [25] Markowitz, H. M., 1952, "Portfolio Selection", *Journal of Finance*, 7, 77-91.
- [26] Merton, R. C., 1990, *Continuous time finance*, Basil Blackwell, Cambridge.
- [27] Savage, L. J., 1954, *The Foundations of Statistics*, John Wiley and Sons, New York.
- [28] Toriki, M., 2001, "Second-order Directional Derivatives of all Eigenvalues of a Symmetric Matrix", *Nonlinear Analysis, Theory, Methods and Applications*, 46, 1133-1150.
- [29] Uppal, R. and T. Wang, 2003, "Model Misspecification and Underdiversification", *Journal of Finance*, 58, 2465-2486.

Appendix

Proof of Proposition 1. The optimality condition is:

$$\mathbf{E}_\pi \left[\varphi' \left(\mathbf{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right) \mathbf{E}_\mu \left[u' \left(\widetilde{W} \right) \left(R - \vec{1} R_{n+1} \right) \right] \right] = 0$$

where $u' \left(\widetilde{W} \right)$ is bounded $\mu(\omega; \cdot)$ -a.s. for any ω in a set of π -measure 1 and $\varphi' \left(\mathbf{E}_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right] \right)$ is bounded π -a.s.

Under the assumptions stated in Section 1, $\mathbf{E}_\mu \left[u \left(\widetilde{W} \right) \right]$ is $\sigma(\Omega)$ -measurable; furthermore, φ' is continuous since $\varphi \in C^1$, hence $\varphi' \left(\mathbf{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right)$ is $\sigma(\Omega)$ -measurable; it is also strictly positive. Define:

$$\xi = \frac{\varphi' \left(\mathbf{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right)}{\mathbf{E}_\pi \left[\varphi' \left(\mathbf{E}_\mu \left[u \left(\widetilde{W} \right) \right] \right) \right]}$$

where the denominator is finite given the boundedness assumption on $\varphi' \left(\mathbf{E}_{\mu(\omega)} \left[u \left(\widetilde{W} \right) \right] \right)$.

The optimality condition can be rewritten as:

$$\mathbf{E}_\pi \left[\xi \mathbf{E}_\mu \left[u' \left(\widetilde{W} \right) \left(R - \vec{1} R_{n+1} \right) \right] \right] = 0$$

Since $\xi > 0$ and $\mathbf{E}_\pi[\xi] = 1$, ξ can be used to define a change of measure and write the optimality condition as:

$$\mathbf{E}_{\pi^*} \left[\mathbf{E}_\mu \left[u' \left(\widetilde{W} \right) \left(R - \vec{1} R_{n+1} \right) \right] \right] = 0$$

where π^* is another probability measure, absolutely continuous with respect to π , with Radon-Nikodym derivative

$$\frac{d\pi^*}{d\pi} = \xi$$

The above double expectation is just a double integral:

$$\int_{\Omega} \int_{\mathbb{R}^n} u'(\widetilde{W}) \left(R - \vec{1} R_{n+1} \right) d\mu(\omega; r) d\pi^*(\omega)$$

Since $u'(\widetilde{W})$ is bounded $\mu(\omega; \cdot)$ -a.s. for any ω in a set of π -measure 1 and $R_i \in L^2(\mathbb{R}^n)$ for any i , the inner integral belongs to $L^1(\Omega)$, Tonelli's theorem guarantees that the integral

$$\int_{\Omega \times \mathbb{R}^n} u'(\widetilde{W}) \left(R - \vec{1} R_{n+1} \right) d\pi^* \times \mu(\omega, r)$$

is well-defined and equals the above double integral.

Thus, the optimality condition becomes:

$$\mathbb{E}_P \left[u'(\widetilde{W}) \left(R - \vec{1} R_{n+1} \right) \right] = 0$$

where $P = \mu \times \pi^*$. ■

Proof of Proposition 2. To prove this proposition we will need the following lemma, which can be found in van Bossum (2002) and descends from a more general theorem in Torki (2001).

Lemma 3 : *let A be a symmetric and positive definite ($n \times n$) matrix which can be written as:*

$$A = y_1 B_1 + y_2 B_2 + \dots + y_k B_k$$

where B_1, B_2, \dots, B_k are positive semidefinite matrices and y_1, y_2, \dots, y_k are positive scalars. Let A have n distinct eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ (they are real and strictly positive) and denote by u_1, u_2, \dots, u_n their respective eigenvectors (chosen so as to be orthonormal). Then:

$$\begin{aligned} \frac{\partial \lambda_i}{\partial y_j} &= u_i' B_j u_i \\ \frac{\partial^2 \lambda_i}{\partial y_i \partial y_j} &= 2 \sum_{r \neq i} \frac{u_i' B_1 u_r u_i' B_j u_r}{\lambda_i - \lambda_r} \end{aligned}$$

Define

$$A(\zeta, t) = \Sigma + (1 + \zeta) S + t z z'$$

where $z = m - \vec{1} R_f$.

We can diagonalize $A(\zeta, 0)$ as follows:

$$A(\zeta, 0) = U \Lambda U'$$

where U is the matrix whose columns u_1, u_2, \dots, u_n are the orthonormal eigenvectors of $A(\zeta, 0)$ and Λ is the diagonal matrix of eigenvalues.

The expected return on the optimal portfolio can be written as:

$$E[R_{w^*}] = \frac{1}{\tau} z' U \Lambda^{-1} U' z + R_f$$

Expanding the product:

$$\tau E [R_{w^*}] = \frac{z'u_1u_1'z}{\lambda_1} + \frac{z'u_2u_2'z}{\lambda_2} + \dots + \frac{z'u_nu_n'z}{\lambda_n} + \tau R_f$$

or:

$$\tau E [R_{w^*}] = \frac{u_1'zz'u_1}{\lambda_1} + \frac{u_2'zz'u_2}{\lambda_2} + \dots + \frac{u_n'zz'u_n}{\lambda_n} + \tau R_f$$

which, applying the above lemma, becomes:

$$\tau E [R_{w^*}] = \sum_{i=1}^n \frac{1}{\lambda_i} \frac{\partial \lambda_i}{\partial t} + \tau R_f$$

Taking the derivative with respect to ζ , we get:

$$\begin{aligned} \tau \frac{\partial E [R_{w^*}]}{\partial \zeta} &= \sum_{i=1}^n \frac{1}{\lambda_i} \frac{\partial^2 \lambda_i}{\partial \zeta \partial t} - \sum_{i=1}^n \frac{1}{\lambda_i^2} \frac{\partial \lambda_i}{\partial t} \frac{\partial \lambda_i}{\partial \zeta} = \\ &= \sum_{i=1}^n \frac{1}{\lambda_i} \left[2 \sum_{j \neq i} \frac{u_i' S u_j u_i' z z' u_j}{\lambda_i - \lambda_j} - \frac{u_i' S u_i u_i' z z' u_i}{\lambda_i} \right] \end{aligned}$$

Setting $Q_{ij} = u_i' S u_j u_i' z z' u_i$ and noting that $Q_{ij} = Q_{ji}$, we can write:

$$\begin{aligned} \tau \frac{\partial E [R_{w^*}]}{\partial \zeta} &= \sum_{i=1}^n \left[2 \sum_{j=1}^{i-1} \left(\frac{Q_{ij}}{\lambda_i (\lambda_i - \lambda_j)} + \frac{Q_{ji}}{\lambda_j (\lambda_j - \lambda_i)} \right) - \frac{Q_{ii}}{\lambda_i^2} \right] = \\ &= \sum_{i=1}^n \left[2 \sum_{j=1}^{i-1} \left(\frac{(\lambda_j - \lambda_i) Q_{ij}}{\lambda_i \lambda_j (\lambda_i - \lambda_j)} \right) - \frac{Q_{ii}}{\lambda_i^2} \right] = \\ &= - \sum_{i=1}^n \sum_{j=1}^n \frac{Q_{ij}}{\lambda_i \lambda_j} \end{aligned}$$

Since the eigenvalues are strictly positive and the Q_{ij} are positive,

$$\frac{\partial E [R_{w^*}]}{\partial \zeta} \leq 0$$

The second part of the proposition is easily proved by noting that:

$$\text{Var} [R_{w^*}] = w^{*'} (\Sigma + S) w^* = \frac{1}{\tau^2} z' [\Sigma + (1 + \zeta) S]^{-1} (\Sigma + S) [\Sigma + (1 + \zeta) S]^{-1} z$$

When $\zeta = 0$ it simplifies to:

$$\text{Var} [R_w] = w^{*'} (\Sigma + S) w^* = \frac{1}{\tau^2} z' [\Sigma + S]^{-1} z = \frac{1}{\tau} (\text{E} [R_w] - R_f)$$

so that:

$$\frac{\partial \text{Var} [R_{w^*}]}{\partial \zeta} = \frac{1}{\tau} \frac{\partial \text{E} [R_{w^*}]}{\partial \zeta}$$

■

Upper and lower bounds on \bar{c} . Deriving upper and lower bounds for \bar{c} is an easy task, once one makes use of the following:

Lemma 4 *Lemma: if A is a $(n \times n)$ matrix with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ and μ is a scalar, then the eigenvalues of the matrix $A + \mu I$ are $\lambda_1 + \mu, \lambda_2 + \mu, \dots, \lambda_n + \mu$.*

The eigenvalues of the $(n \times n)$ matrix $\vec{1} \vec{1}'$ are $\lambda_1 = n, \lambda_2 = 0, \dots, \lambda_n = 0$. This is easily seen by noting that the rank of $\vec{1} \vec{1}'$ is 1, so that $n - 1$ of its eigenvalues must be zero, and n is an eigenvalue associated to the eigenvector $\vec{1}$.

By the above lemma, the eigenvalues of the matrix

$$S = (\bar{s}^2 - \bar{c}) I + \bar{c} \vec{1} \vec{1}'$$

are $\lambda_1 = \bar{c}n + \bar{s}^2 - \bar{c}, \lambda_2 = \bar{s}^2 - \bar{c}, \dots, \lambda_n = \bar{s}^2 - \bar{c}$.

The requirement that all the eigenvalues be strictly greater than zero yields the two conditions:

$$\begin{aligned} \bar{c} &< \bar{s}^2 \\ \bar{c} &> -\frac{\bar{s}^2}{n-1} \end{aligned}$$

■

Table 1
Summary statistics of the portfolios

This table presents some summary statistics of the portfolios selected by an ambiguity averse investor, according to the optimal rule described in the paper. $\widehat{\mathbb{E}}[R_{w^*}]$ is the Monte Carlo estimate of the unconditional monthly expected return on the portfolio, $\widehat{\text{Sd}}[R_{w^*}|\mu]$ estimates the standard deviation, $\widehat{\text{Sh}}[R_{w^*}|\mu]$ the Sharpe ratio, $\widehat{\mathbb{E}}[\delta(w^*)]$ the expected distance of the portfolio from the truly efficient frontier and $\widehat{\mathbb{E}}[a]$ the expected value of the sum of absolute portfolio weights. All the values displayed in the table are averages, calculated from a sample of 100,000 simulations and conditioning on the true value of the vector of expected returns μ , which is not known to the investor (the ambiguity in the model is about μ). The table contains the statistics relative to six different portfolios (corresponding to different values of the ambiguity aversion coefficient ζ) plus a seventh portfolio, marked with a (*), chosen by an investor who disregards ambiguity and adopts a standard portfolio selection rule. The vector of monthly expected returns μ is drawn at each simulation from a multivariate normal distribution, with $\mathbb{E}[\mu] = \bar{m} \vec{1}$ and $\text{Var}[\mu] = \bar{s}^2 I + \bar{c} (\vec{1} \vec{1}' - I)$. In the present table $\bar{m} - 1 = 0.5\%$, $\bar{s} = 0.3\%$ and $\sqrt{\bar{c}} = 0.1\%$. The net riskfree rate of return is set equal to 0.2% and the coefficient of absolute risk aversion is $\tau = 10$.

	$\bar{m} - 1 = 0.5\% \quad \bar{s} = 0.3\% \quad \sqrt{\bar{c}} = 0.1\%$				
ζ	$\widehat{\mathbb{E}}[R_{w^*}]$	$\widehat{\text{Sd}}[R_{w^*}]$	$\widehat{\text{Sh}}[R_{w^*}]$	$\widehat{\mathbb{E}}[\delta(w^*)]$	$\widehat{\mathbb{E}}[a]$
*	0.3716%	1.3154%	0.1304	0.2570	1.2049
0	0.3708%	1.3097%	0.1304	0.2554	1.1977
2	0.3694%	1.2984%	0.1304	0.2524	1.1835
4	0.3679%	1.2874%	0.1304	0.2494	1.1696
8	0.3651%	1.2662%	0.1304	0.2436	1.1430
16	0.3599%	1.2270%	0.1303	0.2330	1.0937
32	0.3507%	1.1589%	0.1301	0.2145	1.0081

Table 2
Summary statistics of the portfolios

This table presents some summary statistics of the portfolios selected by an ambiguity averse investor, according to the optimal rule described in the paper. $\widehat{\mathbb{E}}[R_{w^*}]$ is the Monte Carlo estimate of the unconditional monthly expected return on the portfolio, $\widehat{\text{Sd}}[R_{w^*}|\mu]$ estimates the standard deviation, $\widehat{\text{Sh}}[R_{w^*}|\mu]$ the Sharpe ratio, $\widehat{\mathbb{E}}[\delta(w^*)]$ the expected distance of the portfolio from the truly efficient frontier and $\widehat{\mathbb{E}}[a]$ the expected value of the sum of absolute portfolio weights. All the values displayed in the table are averages, calculated from a sample of 100,000 simulations and conditioning on the true value of the vector of expected returns μ , which is not known to the investor (the ambiguity in the model is about μ). The table contains the statistics relative to six different portfolios (corresponding to different values of the ambiguity aversion coefficient ζ) plus a seventh portfolio, marked with a (*), chosen by an investor who disregards ambiguity and adopts a standard portfolio selection rule. The vector of monthly expected returns μ is drawn at each simulation from a multivariate normal distribution, with $\mathbb{E}[\mu] = \bar{m} \vec{1}$ and $\text{Var}[\mu] = \bar{s}^2 I + \bar{c} \left(\vec{1} \vec{1}' - I \right)$. In the present table $\bar{m} - 1 = 0.7\%$, $\bar{s} = 0.2\%$ and $\sqrt{\bar{c}} = 0$. The net riskfree rate of return is set equal to 0.2% and the coefficient of absolute risk aversion is $\tau = 10$.

	$\bar{m} - 1 = 0.7\%$	$\bar{s} = 0.2\%$	$\sqrt{\bar{c}} = 0$		
ζ	$\widehat{\mathbb{E}}[R_{w^*}]$	$\widehat{\text{Sd}}[R_{w^*}]$	$\widehat{\text{Sh}}[R_{w^*}]$	$\widehat{\mathbb{E}}[\delta(w^*)]$	$\widehat{\mathbb{E}}[a]$
*	0.4617%	1.6179%	0.1617	0.2513	1.2204
0	0.4613%	1.6155%	0.1617	0.2505	1.2166
2	0.4605%	1.6107%	0.1617	0.2489	1.2091
4	0.4597%	1.6060%	0.1617	0.2474	1.2017
8	0.4582%	1.5967%	0.1617	0.2443	1.1872
16	0.4553%	1.5792%	0.1616	0.2385	1.1595
32	0.4499%	1.5471%	0.1615	0.2278	1.1088

Table 3
Summary statistics of the portfolios

This table presents some summary statistics of the portfolios selected by an ambiguity averse investor, according to the optimal rule described in the paper. $\widehat{\mathbb{E}}[R_{w^*}]$ is the Monte Carlo estimate of the unconditional monthly expected return on the portfolio, $\widehat{\text{Sd}}[R_{w^*}|\mu]$ estimates the standard deviation, $\widehat{\text{Sh}}[R_{w^*}|\mu]$ the Sharpe ratio, $\widehat{\mathbb{E}}[\delta(w^*)]$ the expected distance of the portfolio from the truly efficient frontier and $\widehat{\mathbb{E}}[a]$ the expected value of the sum of absolute portfolio weights. All the values displayed in the table are averages, calculated from a sample of 100,000 simulations and conditioning on the true value of the vector of expected returns μ , which is not known to the investor (the ambiguity in the model is about μ). The table contains the statistics relative to six different portfolios (corresponding to different values of the ambiguity aversion coefficient ζ) plus a seventh portfolio, marked with a (*), chosen by an investor who disregards ambiguity and adopts a standard portfolio selection rule. The vector of monthly expected returns μ is drawn at each simulation from a multivariate normal distribution, with $\mathbb{E}[\mu] = \bar{m} \vec{1}$ and $\text{Var}[\mu] = \bar{s}^2 I + \bar{c} \left(\vec{1} \vec{1}' - I \right)$. In the present table $\bar{m} - 1 = 0.3\%$, $\bar{s} = 0.4\%$ and $\sqrt{\bar{c}} = 0.3\%$. The net riskfree rate of return is set equal to 0.2% and the coefficient of absolute risk aversion is $\tau = 10$.

	$\bar{m} - 1 = 0.3\%$	$\bar{s} = 0.4\%$	$\sqrt{\bar{c}} = 0.3\%$		
ζ	$\widehat{\mathbb{E}}[R_{w^*}]$	$\widehat{\text{Sd}}[R_{w^*}]$	$\widehat{\text{Sh}}[R_{w^*}]$	$\widehat{\mathbb{E}}[\delta(w^*)]$	$\widehat{\mathbb{E}}[a]$
*	0.2968%	0.9917%	0.0976	0.2141	0.9958
0	0.2962%	0.9854%	0.0976	0.2127	0.9895
2	0.2950%	0.9731%	0.0976	0.2100	0.9772
4	0.2938%	0.9612%	0.0976	0.2074	0.9652
8	0.2916%	0.9384%	0.0976	0.2024	0.9423
16	0.2875%	0.8968%	0.0976	0.1932	0.9002
32	0.2805%	0.8260%	0.0974	0.1774	0.8281

Table 4
Summary statistics of the portfolios

This table presents some summary statistics of the portfolios selected by an ambiguity averse investor, according to the optimal rule described in the paper. $\widehat{\mathbb{E}}[R_{w^*}]$ is the Monte Carlo estimate of the unconditional monthly expected return on the portfolio, $\widehat{\text{Sd}}[R_{w^*}|\mu]$ estimates the standard deviation, $\widehat{\text{Sh}}[R_{w^*}|\mu]$ the Sharpe ratio, $\widehat{\mathbb{E}}[\delta(w^*)]$ the expected distance of the portfolio from the truly efficient frontier and $\widehat{\mathbb{E}}[a]$ the expected value of the sum of absolute portfolio weights. All the values displayed in the table are averages, calculated from a sample of 100,000 simulations and conditioning on the true value of the vector of expected returns μ , which is not known to the investor (the ambiguity in the model is about μ). The table contains the statistics relative to six different portfolios (corresponding to different values of the ambiguity aversion coefficient ζ) plus a seventh portfolio, marked with a (*), chosen by an investor who disregards ambiguity and adopts a standard portfolio selection rule. The vector of monthly expected returns μ is drawn at each simulation from a multivariate normal distribution, with $\mathbb{E}[\mu] = \bar{m} \vec{1}$ and $\text{Var}[\mu] = \bar{s}^2 I + \bar{c} \left(\vec{1} \vec{1}' - I \right)$. In the present table $\bar{m} - 1 = 0.5\%$, $\bar{s} = 1\%$ and $\sqrt{\bar{c}} = 0.8\%$. The net riskfree rate of return is set equal to 0.2% and the coefficient of absolute risk aversion is $\tau = 10$.

	$\bar{m} - 1 = 0.5\%$	$\bar{s} = 1\%$	$\sqrt{\bar{c}} = 0.8\%$		
ζ	$\widehat{\mathbb{E}}[R_{w^*}]$	$\widehat{\text{Sd}}[R_{w^*}]$	$\widehat{\text{Sh}}[R_{w^*}]$	$\widehat{\mathbb{E}}[\delta(w^*)]$	$\widehat{\mathbb{E}}[a]$
*	1.9046%	4.2800%	0.3982	0.7183	4.1071
0	1.8724%	4.1980%	0.3983	0.7046	4.0299
2	1.8121%	4.0446%	0.3985	0.6789	3.8849
4	1.7564%	3.9038%	0.3986	0.6552	3.7511
8	1.6568%	3.6539%	0.3987	0.6131	3.5121
16	1.4945%	3.2499%	0.3983	0.5448	3.1219
32	1.2639%	2.6826%	0.3965	0.4482	2.5673

Table 5
Empirical distribution of the expected rate of return

This table displays some features of the empirical distribution of the monthly conditional expected return $E[R_{w^*}|\mu]$ on the portfolios selected by an ambiguity averse investor, according to the optimal rule described in the paper. $E[R_{w^*}|\mu]$ depends on the portfolio selected by the investor and on the realization of the parameter vector μ ; hence, we obtain a different value for each of the 100,000 simulations we perform, which gives rise to an empirical distribution of $E[R_{w^*}|\mu]$. We report both the standard deviation and some percentiles of the distribution. Each row corresponds to a different value of the ambiguity aversion coefficient ζ , while the row marked with a (*) corresponds to the case in which the investor disregards ambiguity and adopts a standard portfolio selection rule. The vector of monthly expected returns μ is drawn at each simulation from a multivariate normal distribution, with $E[\mu] = \bar{m} \vec{1}$ and $\text{Var}[\mu] = \bar{s}^2 I + \bar{c} (\vec{1} \vec{1}' - I)$. In the present table $\bar{m} = 0.5\%$, $\bar{s} = 0.3\%$ and $\bar{c} = 0.1\%$. The riskfree rate of return is set equal to 0.2%.

Percentiles of the distributions							
ζ	Std deviation	1st	10th	20th	80th	90th	99th
*	0.1058%	0.168%	0.252%	0.285%	0.450%	0.507%	0.699%
0	0.1051%	0.168%	0.252%	0.285%	0.449%	0.506%	0.686%
2	0.1038%	0.169%	0.252%	0.284%	0.447%	0.503%	0.681%
4	0.1025%	0.170%	0.252%	0.284%	0.444%	0.500%	0.676%
8	0.1001%	0.172%	0.252%	0.283%	0.440%	0.494%	0.666%
16	0.0956%	0.175%	0.251%	0.281%	0.431%	0.483%	0.646%
32	0.0878%	0.180%	0.250%	0.278%	0.417%	0.464%	0.610%

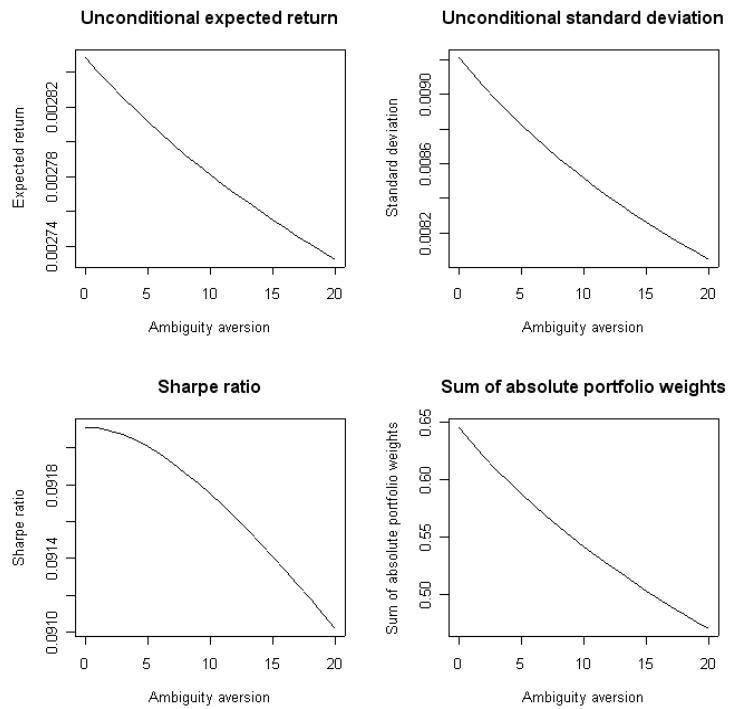


Figure 1: The impact of an increase in ambiguity aversion. These plots depict the impact of increasing the coefficient of ambiguity aversion ζ on some statistics describing the optimal portfolio (unconditional expected return, standard deviation, Sharpe ratio and average value of the sum of absolute portfolio weights).

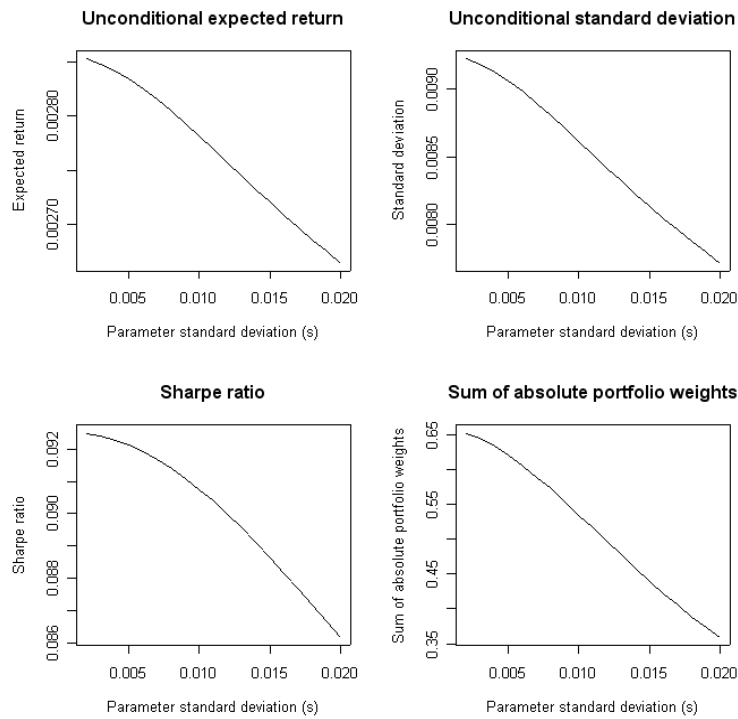


Figure 2: The impact of an increase in ambiguity. These plots depict the impact of increasing the variance \bar{s} of the unknown parameter on some statistics describing the optimal portfolio (unconditional expected return, standard deviation, Sharpe ratio and average value of the sum of absolute portfolio weights).

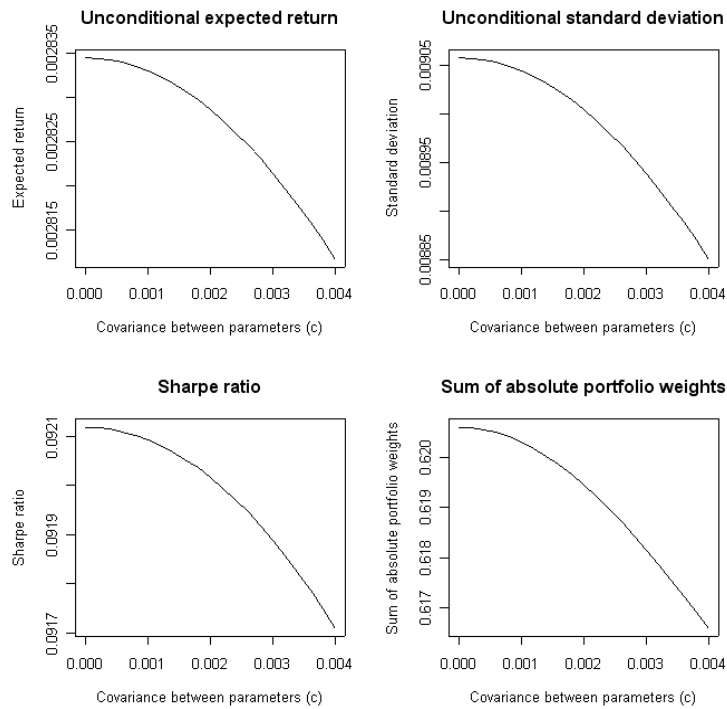


Figure 3: The impact of an increase in the covariance between parameters. These plots depict the impact of increasing the covariance \bar{c} between the unknown parameters on some statistics describing the optimal portfolio (unconditional expected return, standard deviation, Sharpe ratio and average value of the sum of absolute portfolio weights).