Single period portfolio choice with learning in the framework of relative regret*

Andrew E.B. Lim, J. George Shanthikumar & Gah-Yi Vahn
Department of Industrial Engineering and Operations Research
University of California
Berkeley, CA 94720
Email: {lim, jgshant}@ieor.berkeley.edu, gyvahn@gmail.com

September 9, 2009

The question as to how the form of the weight (i.e. loss) function $W(\theta, \omega)$ should be determined, is not a mathematical or statistical one. The statistician who wants to test certain hypotheses must first determine the relative importance of all possible errors, which will entirely depend on the special purpose of his investigation.

Abraham Wald (1939).

Abstract

In this paper, we formulate a single period portfolio choice problem with parameter uncertainty in the framework of relative regret. Relative regret evaluates a portfolio by comparing its return to a family of benchmarks, where the benchmarks are the wealths of fictitious investors who invest optimally given knowledge of the model parameters, and the optimal portfolio is the one that performs well in relation to all the benchmarks over the family of possible parameter values. We analyze this problem using convex duality, and show that it is equivalent to a Bayesian problem, where the Lagrange multipliers play the role of the prior distribution and the learning model involves Bayesian updating of these Lagrange multipliers/prior. This

*Preliminary and incomplete draft

Bayesian problem is unusual in that the prior distribution is endogenously chosen (by solving the dual optimization problem for the Lagrange multipliers), and the objective function involves the family of benchmarks from the relative regret problem.

Key words– parameter uncertainty, ambiguity, learning, model uncertainty, regret, relative regret, competitive analysis, portfolio selection, convex duality, Bayesian methods, objective based loss functions.

1 Introduction

Consider an investor living in a world where there are risky assets and a risk-free money market account. Log-returns for the risky assets are iid normal but the mean and variance are not known to the investor (though they are constant over the data window). We consider a single period portfolio choice problem where the agent is endowed with a finite sample of historical asset returns and makes a one shot allocation decision after observing the last data point.

Portfolio selection with parameter uncertainty is typically formulated in a Bayesian framework. In this paper, we take a different approach by adopting a relative regret objective. Positive justification for relative regret is given in Terlizzese [14] (see also Hayashi [6]), and we provide normative justification below. The essential feature of relative regret is that an investor's allocation is assessed by comparing his/her wealth to a family of benchmarks, where the benchmarks are wealths of fictitious investors who behave optimally given knowledge of the parameter value. The optimal portfolio maximizes the worst case *relative* performance with respect to this family of benchmarks.

Since returns are iid, historical data is relevant for decision making as it enables the investor to learn the unknown parameter values. The form of the learning model, however, depends on the optimization and decision making framework. In a Bayesian setting, the investor specifies a prior and Bayes' rule comes out endogenously as the learning model. With relative regret, however, there is no prior specification and the optimal learning model is less clear. One contribution of this paper is the characterization of the optimal learning model for the relative regret problem.

Some readers might question the need for considering alternatives to the Bayesian framework. One reason, which has substantial normative implications, is that the solution of the classical Bayesian model can be very sensitive to the specification of the prior. In particular, "relatively small" changes in the mean of the prior (for the expected return of a log-normal distribution) can translate into a large deterioration in performance if the prior variance is small, while a uniform prior (the "obvious" choice when the investor has no prior information) gives the solution of the classical Merton/Markowitz problem with the sample mean of realized returns plugged-in for the true expected return. It is well known, however, that this estimator converges slowly and that Merton/Markowitz portfolios with sample mean plug-in estimates of the mean return perform badly [1, 13]. More generally, it is important to recognize that the posterior distribution for the mean return converges at the same rate as the sample mean, so a poor choice of prior takes a long time to correct itself. In summary, the problem of specifying a prior in the context of a real world portfolio choice problem, may be nontrivial[1], and at a minimum, alternative learning models or new methods for choosing priors might be of interest.

Our methods and results can be summarized as follows. We analyze the single period regret problem using convex duality, and show that the dual function is a non-standard Bayesian problem with the Lagrange multipliers playing the role of the prior. The Lagrange multipliers/prior associated with the optimal solution are obtained by solving the associated dual problem, and the learning model in the optimal regret solution involves Bayesian updating of this endogenously chosen prior/multiplier. The Bayesian problem is also unusual in that the objective involves the family of benchmarks associated with the relative regret problem. In particular, the investor's decision is evaluated by comparing his/her wealth to that of each of the benchmark investors and averaging over the prior. Roughly speaking, investors are rewarded for performing well relative to benchmarks that look plausible given the posterior; if the posterior is relatively flat (so all models are still plausible) then the investor seeks to do well relative to all the benchmarks. On the other hand, the investor will narrow his/her attention to a smaller set of benchmarks/models as the posterior becomes more concentrated.

Another interpretation of our regret objective is from the perspective of loss functions in statistical inference. The statement by Abraham Wald [15] at the start of the paper points to the idea that inference and decision making should go hand in hand, and that statistical error has most meaning when it is evaluated in the context of the application in which the estimator is to be used (see also McCloskey [12]). In this regard, by embedding learning in our optimization problem and evaluating performance with respect to benchmarks, we have combined inference and decision making in a way that explicitly accounts for the "relative importance of all possible errors" to our investment performance. The learning model associated with the optimal solution gives the best estimator for this loss function.

The outline of our paper is as follows. We formulate the single period market model in Section 2, and introduce two relative regret objectives in Section 3. the first of these is more

[1]Of course, the opportunity to specify a prior is valuable when the investor does possess accurate information. On the other hand, one could also argue that parameter uncertainty is not really a problem if this were the case.

standard [6, 14] while the second (which is the major focus of this paper) is original and can be interpreted as an objective based loss function. We establish connections between the single period relative regret models and Bayesian problems in Section 4 using convex duality. In particular, we show that Lagrange multipliers in this duality relationship play the role of the prior in the Bayesian problem, and that the solution of the regret problem involves Bayesian updating of the prior/Lagrange multiplier characterized as the solution of a certain (dual) optimization problem. Comparisons between relative regret and standard worst case models are discussed in Section 5. Computational studies are presented in Section 6.

The single period model in this paper can be extended to dynamic trading, which will be discussed elsewhere. An interesting feature of this extension is that the learning model involves a posterior that is tilted using the family of benchmarks. The reader is directed to [8] for more details.

2 Model

2.1 Market, model and investor

Financial market

Assume that there is a risk free asset $S_0(t)$ and n risky assets $S_1(t), \cdots, S_n(t)$. The risk free asset has a continuously compounded interest rate of r and its price is given by $S_0(t) = e^{rt}$. We assume that the interest rate is known to the investor. The prices of the n risky assets evolve in continuous time according to

$$S_i(t) = S_i(0) \exp\Big\{\mu_i t + \sigma_i W(t)\Big\}, \quad i = 1, 2, \cdots, n, \tag{1}$$

where $W(t)$ is an n-dimensional standard Brownian motion, the scalar μ_i is the rate of return for stock i, and the n-dimensional row vector $\sigma_i = [\sigma_{i1}, \sigma_{i2}, \cdots, \sigma_{in}] \in \mathbb{R}^{1\times n}$ is the volatility of this stock. We assume throughout this paper that μ_i and σ_i are constant. The column vector

$$\mu = [\mu_1, \mu_2, \cdots, \mu_n]' \in \mathbb{R}^{n\times 1}$$

is the vector of returns for all the risky assets and the $n \times n$ matrix

$$\sigma = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_n \end{bmatrix} \in \mathbb{R}^{n\times n}$$

is the volatility. We assume, as is standard, that the non-degeneracy assumption holds: $Q = \sigma\sigma' \geq \delta I$ for some constant $\delta > 0$. The column vector of stock prices is denoted by

$$S(t) = \begin{bmatrix} S_1(t), & S_2(t), & \cdots, & S_n(t) \end{bmatrix}'.$$

Investor's observations/data

We assume in this paper that the parameters $H = (\sigma\sigma', \mu)$ for the stock price model (1) are constant. We assume however that the investor does not know the parameter values beyond the fact that they belong to some uncertainty set $\mathcal{H}$. The only assumption about the uncertainty set $\mathcal{H}$ is that it is compact. (For example, $\mathcal{H}$ might be a "confidence interval/region" associated with statistical point estimates of the parameters, subjective uncertainty regions specified by the investor around forecasted means, or a finite set of models that the investor wishes to consider, etc). We also assume that the investor does not observe the stock prices continuously but samples the process at discrete times $t\delta$ for $t = 0, 1, 2, \cdots, T$ (i.e. t indexes the number of sample points that have been seem by the investor, while $\delta > 0$ is the time interval between each observation). Equivalently, the investor is seeing a sequence of log returns $\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T)$ where

$$\mathcal{R}(t+1) = \begin{bmatrix} \mathcal{R}_1(t+1) \\ \vdots \\ \mathcal{R}_n(t+1) \end{bmatrix}, \quad t = 0, 1, 2, \cdots, T,$$

is an n-dimensional random vector with entries being the log-returns for the individual stocks over time period $[t\delta, (t+1)\delta)$:

$$\begin{aligned} \mathcal{R}_j(t+1) \quad &\triangleq \quad \ln \frac{S_j((t+1)\delta)}{S_j(t\delta)} \\ &= \quad \mu_i\delta + \sigma_i[W((t+1)\delta) - W(t\delta)] \\ &=^D \quad \mu_i\delta + \sqrt{\delta}\sigma_i Z(t+1). \end{aligned}$$

In this equation, $Z(1), Z(2), \cdots, Z(T)$ is a sequence of n-dimensional iid standard normal random variables. Clearly, $\mathcal{R}(t+1)$ is multivariate normal

$$\mathcal{R}(t+1) \sim N(\delta\mu, \delta Q) \tag{2}$$

with mean $\mu\delta$ and covariance matrix δQ.

Investment decision

Consider an investor with wealth x. The investor (correctly) assumes that prices evolve in continuous time according to a model of the form (1), but does not know that parameter values (σ, μ). Instead, the investor has observed T historical returns over time periods of size δ, $\mathcal{R}(1)$, $\mathcal{R}(2), \cdots, \mathcal{R}(T)$ (or equivalently, has seen stock prices $S(0)$, $S(\delta), \cdots, S(T\delta)$) and wishes to make an investment decision over the time interval $[T\delta, (T+1)\delta)$ following the realization of the last observation. The investor can use knowledge of the T historical returns to make his/her decision but not the actual parameter values themselves. More formally, the investor's decision $\pi = [\pi_1, \pi_2, \cdots, \pi_n]'$ for the interval $[T\delta, (T+1)\delta)$ is a $\mathcal{G}_T$-measurable random vector. We assume that π_i is the proportion of wealth invested in stock i while $1-\pi'\mathbf{1}$ is the proportion invested in the bond. Under this assumption it follows that the investor's wealth at time $(T+1)\delta$ (after the realization of return $\mathcal{R}(T+1)$) is given by

$$\begin{aligned} x_\pi^H(T+1) &= \sum_{i=1}^{m} \frac{x\pi_i}{S_i(T)} e^{\mathcal{R}_i(T+1)} + x\Big(1 - \sum_{i=1}^{n} \pi_i\Big) e^{r\delta} \\ &= x\Big\{\sum_{i=1}^{n} \pi_i e^{\delta\mu_i + \sqrt{\delta}\sigma_i Z(T+1)} + \Big(1 - \sum_{i=1}^{n} \pi_i\Big) e^{r\delta}\Big\}. \end{aligned} \tag{3}$$

For the sake of mathematical convenience, let us assume that δ is small. It can then be shown that

$$x_\pi^H(T+1) \simeq x \exp\Big\{\delta\Big[r + b'\pi - \frac{1}{2}\pi' Q\pi\Big] + \sqrt{\delta}\pi'\sigma Z(T+1)\Big\}$$

where $b = [b_1, b_2, \cdots, b_n]'$ is an n-dimensional vector of real numbers $b_i \triangleq \mu_i + (1/2)|\sigma_i|^2 - r$ and $Z(T+1)$ is a standard n-dimensional normal random variable which is independent of returns $\mathcal{R}(1)$, $\mathcal{R}(2), \cdots, \mathcal{R}(T)$ (or equivalently, of realized stock prices $S(0)$, $S(\delta), \cdots, S(T\delta)$)[2]. With this in mind, we shall assume that the investor's wealth is defined by

$$x_\pi^H(T+1) = x \exp\Big\{\delta\Big[r + b'\pi - \frac{1}{2}\pi' Q\pi\Big] + \sqrt{\delta}\pi'\sigma Z(T+1)\Big\} \tag{4}$$

for the remainder of the paper. Alternatively, (4) is the wealth at $T+1$ if there is continuous rebalancing (by a computer, say) between times T and $T+1$ so as to maintain the proportions π of wealth in each stock, with the understanding that the investor does not change π between T and $T+1$ once it has been specified at T.

Finally, since we will be dealing directly with the wealth equation (4) rather than the log-returns model (2), it is more convenient for us to talk about uncertainty in $H = (Q, b)$

[2] This follows from the so-called log-linear approximation of the wealth equation (3), which becomes exact when $\delta \downarrow 0$ (see [2] for more details).

	Investor 1	Investor 2	Investor 3
prior mean m	0.15	0.2	0.25
prior precision τ	25	25	25

Figure 1: Summary of priors for Bayesian investors

instead of uncertainty in (Q, μ). In particular, we assume that the investor does not know (Q, b) beyond the fact that it lies in some compact uncertainty set $\mathcal{H}$.

2.2 Prior distributions in Bayesian models

When there is parameter uncertainty, it is common to adopt a Bayesian framework. In this section, we present an example which shows that the solution of the Bayesian problem is sensitive to the prior distribution. Sensitivity to the prior can be of concern if specification of the prior distribution is difficult (e.g. it is often difficult to translate a particular qualitative prior view into a joint distribution, particularly if there are many uncertain variables).

Suppose there is a single stock with iid log-normal returns described by (2) and parameters $\mu = 20\%$ and $\sigma = 20\%$. We assume that σ is known to all investors but μ is not. Consider three Bayesian investors. We assume that each investor knows $\sigma = 20\%$, but has a different (normal) prior on the unknown mean μ. These are summarized in Figure 1. Observe that each of the priors has a different mean m but the same precision τ (recall that precision $\tau = (\text{variance})^{-1}$). A precision of 25 is the same as a standard deviation of 20%. Observe that the mean of Investor 2's prior $m = 0.2$ equals the mean μ of the distribution generating the returns.

We simulated data consisting of $n = 10$ years of annual returns $\mathcal{R}(1), \cdots, \mathcal{R}(10)$ using the "true model" $(\mu, \sigma) = (0.2, 0.2)$, and updated the priors of each of the investors using Bayes' rule. It is well known that posteriors are normal with mean and precision

$$m' = \frac{\tau m + (n/\sigma^2)\bar{\mu}_n}{\tau + n/\sigma^2}, \quad \tau' = \tau + \frac{n}{\sigma^2}$$

where $\bar{\mu}_n = [\mathcal{R}(1) + \cdots + \mathcal{R}(n)]/n$ is the sample mean of the historical returns. Our historical sample mean took the value $\bar{\mu}_n = 0.2126$ (which is relative close to the actual value $\mu = 0.2$).

Each Bayesian investor then solved the following single period asset allocation problem

$$\begin{cases} \max_\pi \frac{1}{\gamma}\mathbb{E}\left[x(1)^\gamma\right] \\ \text{Subject to:} \\ x(1) = x(0)\exp\left\{\left[r + (\mu - r)\pi - \frac{1}{2}\pi^2\sigma^2\right] + \pi\sigma Z(T+1)\right\} \\ x(0) = 1, \\ \text{prior on } \mu \sim N(m',\, \tau') \end{cases}$$

using their updated parameters. It can be shown that

$$\pi = \left[1 - \gamma\left(1 + \frac{1}{\sigma^2\tau'}\right)\right]^{-1}\frac{m' - r}{\sigma^2}$$

is the optimal portfolio for the Bayesian investors (with $(m',\, \tau')$ as above). For the historical returns we generated, we obtained

$$\pi_1 = 22.954, \quad \pi_2 = 23.62, \quad \pi_3 = 24.3$$

for Investors 1, 2 and 3. It is interesting to compare this to the optimal portfolio $\psi = \frac{1}{1-\gamma}(\mu - r)/\sigma^2$ of a fictitious investor who knows the model parameters $(\mu,\, \sigma)$ and who solves

$$\begin{cases} \max_\psi \frac{1}{\gamma}\mathbb{E}\left[y(1)^\gamma\right] \\ \text{Subject to:} \\ y(1) = y(0)\exp\left\{\left[r + (\mu - r)\psi - \frac{1}{2}\psi^2\sigma^2\right] + \psi\sigma Z(T+1)\right\} \\ y(0) = 1. \end{cases}$$

For our model parameters $\psi = 23.4$.

It is interesting to observe that if prior precision τ is set to 0, the commonly accepted default when the investor has no information about μ, that

$$\pi = \left[1 - \gamma\left(1 + \frac{1}{n}\right)\right]^{-1}\frac{\bar{\mu}_n - r}{\sigma^2}$$

which is essentially the portfolio ψ with the sample mean $\bar{\mu}_n$ substituted in place of the unknown mean μ and a small correction to the risk-aversion parameter. It is well known, however, that this plug-in approach does not perform well out of sample [1, 13].

Consider now the following experiment. We generated $1,000,000$ samples of annual returns using the model $(\mu,\, \sigma) = (20\%,\, 20\%)$. For each sample we recorded the end-of-year wealth $x_i(1)$ of each of the Bayesian investors π_i $(i = 1,\, 2,\, 3)$ as well as the wealth $y(1)$ of the "knowledgeable" investor who invests according to ψ. Figures 2–4 are histograms of log relative wealth, i.e. $\log[x_i(1)/y(1)]$, for each of the Bayesian investor's. The most striking observation is the large difference between these three histograms given that the difference in prior specification is relatively small.

In the remainder of the paper, we consider relative regret as a framework for formulating portfolio selection problems with parameter uncertainty. One feature of our model is that the learning model is endogenously determined and the prior is not specified, which is appealing given the robustness concerns with prior specification. Preliminary experiments on the performance of relative performance investors are presented in Section 6.

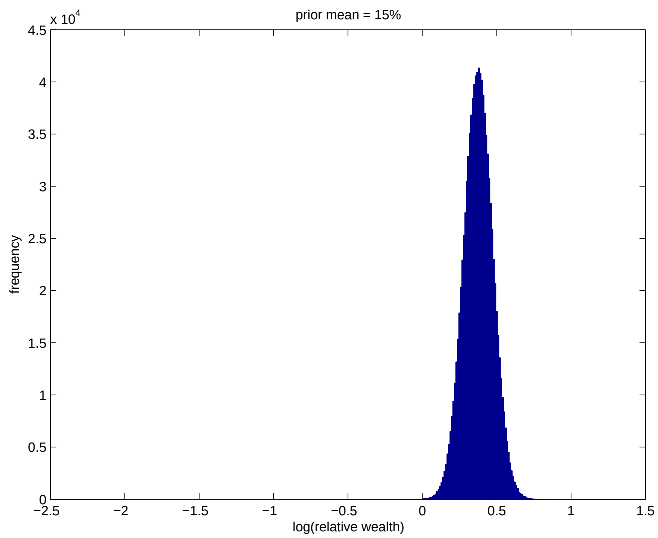

Figure 2: Histogram of log relative wealth of investor 1 relative to the knowledgeable investor.

3 Relative regret

In this section, we formulate two portfolio optimization problems within the setup described in Section 2 and analyze the solution of these problems. Both problems involve relative regret objectives. The first of these is the classical relative regret (see for example Terlizzese [14]) while the second is our own. A key feature on both problems is that the investor, though ignorant of the model parameters, has the opportunity to learn. As such a major focus of our work in subsequent sections (particularly Section 4) is the characterization of the learning model associated with the optimal solution of the problems formulated in this section.

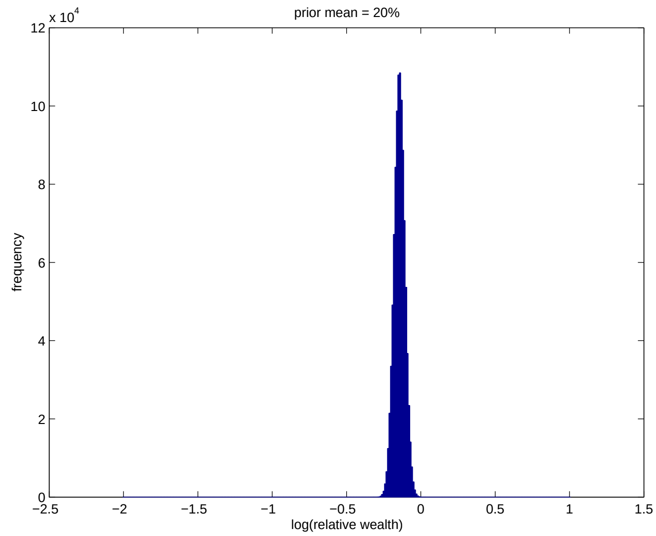

Figure 3: Histogram of log relative wealth of investor 2 relative to the knowledgeable investor.

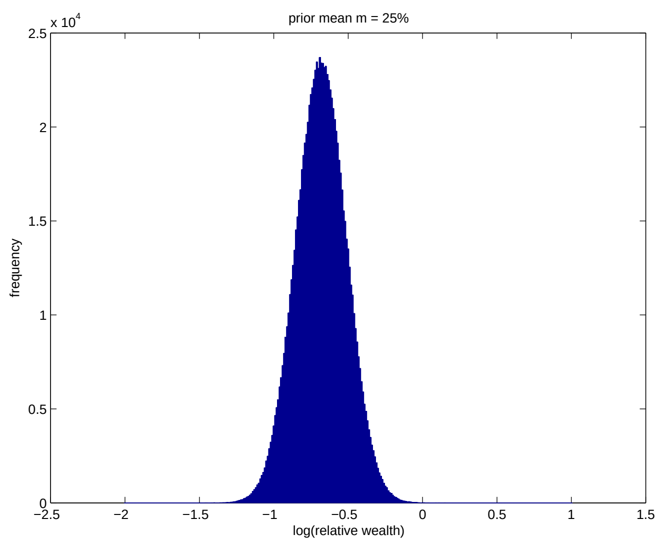

Figure 4: Histogram of log relative wealth for investor 3 relative to the knowledgeable investor.

3.1 Relative regret I: standard model

Consider the problem

$$\max_{\pi \in \mathcal{G}_T} \min_{H \in \mathcal{H}} \frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\max_\psi \mathbb{E}_H U(y_\psi^H(\delta))}. \tag{5}$$

This objective can be understood as follows:

- The investor begins by proposing a decision rule $\pi \in \mathcal{G}_T$, or equivalently, a measurable function $f : \mathbb{R}^{T+1} \to \mathbb{R}$ which maps the T observed returns $\mathcal{R}(1), \cdots, \mathcal{R}(T)$ to an investment position $\pi \equiv f(\mathcal{R}(1), \cdots, \mathcal{R}(T))$. Note that every $\mathcal{G}_T$-measurable r.v. can be represented as a measurable function of $\mathcal{R}(1), \cdots, \mathcal{R}(T)$. We emphasize that π (equivalently, f) can not depend explicitly on the parameters (Q, b) since they are not known to the investor.

- Once this decision rule $\pi \in \mathcal{G}_T$ has been revealed, nature chooses a parameter $H = (Q, b)$ from the set $\mathcal{H}$.

- For the chosen policy $\pi \in \mathcal{G}_T$ and model $H \in \mathcal{H}$, the investor's wealth at time $(T + 1)\delta$ is given by (4), and the expected utility $\mathbb{E}_H U(x_\pi^H(\delta))$ in the numerator of (5) can be computed. In addition, the denominator is calculated by optimizing over ψ given knowledge of the model H chosen by nature:

$$\begin{cases} \quad \mathbb{E}_H U(y^H(\delta)) = \max_\psi \mathbb{E}_H U(y_\psi^H(\delta)) \\ \text{subject to:} \\ \quad y_\psi^H(\delta) = y(0) \exp\Big\{ \delta \Big[r + b'\pi - \tfrac{1}{2}\psi' Q \psi \Big] + \sqrt{\delta}\psi' \sigma Z(T+1) \Big\}. \end{cases} \tag{6}$$

  The ratio of these two quantities is precisely the relative regret objective (5).

- The investor chooses the policy $\pi \equiv f(\cdot)$ and nature the model H to satisfy the equilibrium condition (5).

An axiomatic justification for this objective is given in Terlizzese [14]; see also Hayashi [6].

CRRA utility function: $U(x) = \frac{1}{\eta} x^\eta$, $\eta < 1$

A more explicit computation can be done for the model (5) if the utility function is assumed to be CRRA. More specifically, observing that $y(\delta)$ is a log-normal random variable given by (6), it follows that

$$\begin{aligned} & \frac{1}{\eta}\mathbb{E}_H \left(y_\pi^H(\delta) \right)^\eta \\ &= \frac{1}{\eta} y(0)^\eta \mathbb{E}_H \exp\Big\{ \eta\delta \Big[r + b'\pi - \frac{1}{2}\psi' Q\psi \Big] + \eta\sqrt{\delta}\psi'\sigma Z(T+1) \Big\} \\ &= \frac{1}{\eta} y(0)^\eta \exp\Big\{ \eta\delta \Big[r + b'\psi - \frac{1-\eta}{2}\psi' Q \psi \Big] \Big\} \\ &= \frac{1}{\eta} y(0)^\eta \exp\Big\{ \delta\eta \Big[r - \frac{1-\eta}{2} \Big(\psi - \frac{Q^{-1}b}{1-\eta} \Big)' Q \Big(\psi - \frac{Q^{-1}b}{1-\eta} \Big) + \frac{1}{2(1-\eta)} b' Q^{-1} b \Big] \Big\}, \end{aligned} \tag{7}$$

(where the first equality is just the moment generating function of a normal r.v.). It now follows that the benchmark investor's optimal portfolio (the solution of (6) using (7)) is given by

$$\psi^H = \arg\max_{\psi} = \frac{1}{\eta}\mathbb{E}_H U(y_\pi^H(\delta)) = \frac{1}{1-\eta}Q^{-1}b \tag{8}$$

when the model is $H = (Q, b)$. Substituting ψ^H into the wealth equation in (6) it follows that the benchmark investor's optimal wealth is given by

$$y^H(\delta) = y(0)\exp\left\{\delta\Big[r + \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] + \frac{\sqrt{\delta}}{1-\eta}b'Q^{-1}\sigma Z(T+1)\right\} \tag{9}$$

and the denominator of (5) (from substituting (8) into (7)) is

$$\mathbb{E}_H U(y^H(\delta)) = \frac{1}{\eta}y(0)^\eta \exp\left\{\delta\eta\Big[r + \frac{1}{2(1-\eta)}b'Q^{-1}b\Big]\right\}. \tag{10}$$

On the other hand, for the portfolio $\pi \in \mathcal{G}_T$ and model $H = (Q, b) \in \mathcal{H}$, the investor's utility function (the numerator of (5)) satisfies

$$\mathbb{E}_H U(x_\pi^H(\delta)) = \mathbb{E}_H\left[\mathbb{E}_H\left\{U(x_\pi^H(\delta))|\mathcal{G}_T\right\}\right].$$

Observing that (conditional on $\mathcal{G}_T$) the exponent of

$$U(x_\pi^H(\delta)) = \frac{x(0)^\eta}{\eta}e^{\delta\eta[r+b'\pi-\frac{1}{2}\pi'Q\pi]+\eta\sqrt{\delta}\pi'\sigma Z(T+1)}$$

is a standard normal r.v. with mean $\delta\eta(r + b'\pi - \frac{1}{2}\pi'Q\pi)$ and variance $\delta\eta^2\pi'Q\pi$, it follows from the formula for the moment generating function of a normal r.v. that

$$\mathbb{E}_H[U(x_\pi^H(\delta))|\mathcal{G}_T] = \frac{1}{\eta}x(0)^\eta \exp\left\{\delta\eta\Big[r + b'\pi - \frac{1-\eta}{2}\pi'Q\pi\Big]\right\}$$

and hence

$$\mathbb{E}_H U(x_\pi^H(\delta)) = \frac{1}{\eta}x(0)^\eta \mathbb{E}_H \exp\left\{\delta\eta\Big[r + b'\pi - \frac{1-\eta}{2}\pi'Q\pi\Big]\right\}. \tag{11}$$

Substituting (10) and (11) into the relative regret objective in (5), we obtain

$$\frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\max_\psi \mathbb{E}_H U(y_\psi^H(\delta))} = \frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\mathbb{E}_H U(y^H(\delta))} = \mathbb{E}_H \exp\left\{\delta\eta\Big[b'\pi - \frac{1-\eta}{2}\pi'Q\pi - \frac{1}{2}\frac{b'Q^{-1}b'}{1-\eta}\Big]\right\}$$

and it follows that (5) is equivalent to

$$\max_{\pi\in\mathcal{G}_T}\min_{H\in\mathcal{H}} \frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\max_\psi \mathbb{E}_H U(y_\psi^H(\delta))} = \max_{\pi\in\mathcal{G}_T}\min_{H\in\mathcal{H}} \mathbb{E}_H \exp\left\{\delta\eta\Big[b'\pi - \frac{1-\eta}{2}\pi'Q\pi - \frac{1}{2}\frac{b'Q^{-1}b'}{1-\eta}\Big]\right\}. \tag{12}$$

It does not appear that an explicit expression for the equilibrium solution of (12) is possible. However, we show in Section 4 that an approximate solution, which becomes exact when $\delta \downarrow 0$, can be obtained.

3.2 Relative regret II: objective based loss function

In this section we introduce an alternative relative regret problem. We adopt the same family of models $\mathcal{H}$ as in Section 3.2, the same (approximate) wealth equation (4) for the investor, and the same definition (6) of the benchmark $y^H(\delta)$. The essential difference comes in the way that the investor's wealth $x_\pi^H(\delta)$ is compared to that of the benchmark investor $y^H(\delta)$.

Benchmark investor

As in (6), the benchmark investor solves a portfolio selection problem with full knowledge of the model parameters $H = (Q, b)$. More specifically, suppose that the benchmark investor has utility function $U^B(y)$ and that he/she solves the portfolio selection problem

$$\left\{\begin{array}{l} \quad \mathbb{E}_H U^B(y^H(\delta)) \equiv \max_\psi \mathbb{E}_H U^B(y_\psi^H(\delta)) \\ \text{subject to:} \\ \quad y_\psi^H(\delta) = y(0)\exp\left\{\delta\Big[r + b'\psi - \frac{1}{2}\psi' Q\psi\Big] + \sqrt{\delta}\psi'\sigma Z(T+1)\right\}, \end{array}\right. \tag{13}$$

where as before, ψ^H denotes the optimal solution of this problem and $y^H(\delta)$ (a random variable) is the associated optimal wealth.

Relative regret problem

Consider the following relative regret problem

$$\left\{\begin{array}{l} \quad \max_{\pi\in\mathcal{G}_T} \min_{H\in\mathcal{H}} \mathbb{E}_H\left[U\left(\frac{x_\pi^H(\delta)}{y^H(\delta)}\right)\right] \\ \text{subject to:} \\ \quad x_\pi^H(\delta) \text{ is given by (4)} \\ \quad y^H(\delta) \text{ is defined via (13).} \end{array}\right. \tag{14}$$

The key difference between (14) and the relative regret problem (5) is the way that $x_\pi^H(\delta)$ and $y^H(\delta)$ are compared. In this regard, it is worth noting that the "comparison function" $U(z)$ in the objective and the utility function $U^B(y)$ need not be the same.

The model (14) can be described as follows.

- The investor begins by declaring a policy $\pi \in \mathcal{G}_T$ (or equivalently, by specifying some measurable function $f : \mathbb{R}^{T+1} \to \mathbb{R}$). Nature, endowed with knowledge of this policy π (equivalently, function f), follows up by choosing a model $H \in \mathcal{H}$. Asset returns $\mathcal{R}(1), \cdots, \mathcal{R}(T)$ are then generated under nature's model H.

  At time T, the investor adopts the position $\pi = f(\mathcal{R}(1), \cdots, \mathcal{R}(T))$ (after seeing all the returns) while nature invests according to ψ^H (the optimal solution of (13) corresponding to H).

- Once the positions π and ψ^H have been taken, one more return realization $\mathcal{R}(T+1)$ is generated under nature's model H and the wealths of the investor $x_\pi^H(\delta)$ (given by (4)) and the benchmark $y^H(\delta)$ (given by (13)) are realized. Conditional on $\pi = f(\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T))$ and H, the distribution of the ratio $x_\pi^H(\delta)/y^H(\delta)$ is fully characterized, and we can calculate

$$\mathbb{E}_H U\Big(\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big).$$

  We use this objective to compare $x_\pi^H(\delta)$ and $y^H(\delta)$.

- The investor and nature choose π and H to satisfy the equilibrium condition associated with (14).

We reiterate that the essential difference between the models (5) and (14) is the way that $x_\pi^H(\delta)$ and $y^H(\delta)$ are compared. In (5) they are compared by evaluating the ratio of their expected utilities while in (14) we compute the expectation of the comparison function $U(z)$ applied to the ratio $\frac{x_\pi^H(\delta)}{y^H(\delta)}$.

Several additional comments are worth making. Firstly, unlike the objective (5), we are not aware of an axiomatic foundation for (14) though we believe that this is an issue worth pursuing. On the other hand, it will be shown that the solution of (5) is a limiting case of (5) when the utility/comparison functions are CRRA/power type. Another advantage of (14) is that it gives us some degree of control over the "risk aversion" of the benchmark investor (through the choice of $U^B(y)$) as well as the distance measure between $x_\pi^H(\delta)$ and $y^H(\delta)$ through the choice of $U(z)$. Finally, there is a natural extension of (14) to multi-period problems that is relatively easy to analyze. (This will be done elsewhere). The same can not be said about (5).

Power utility and comparison functions: $U^B(y) = \frac{1}{\eta}y^\eta$ and $U(z) = \frac{1}{\gamma}z^\gamma$

In this section we consider the relative regret problem (14) under the assumption that the benchmark investor's utility function as well as the comparison function are power-type: $U^B(y) = \frac{1}{\eta}y^\eta$ ($\eta < 1$) and $U(z) = \frac{1}{\gamma}z^\gamma$ ($\gamma < 1$). As shown in (8)-(9), the benchmark investor's problem (13) with a CRRA utility has an explicit solution

$$\psi^H = \frac{1}{1-\eta}Q^{-1}b \tag{15}$$

and

$$y^H(\delta) = y(0)\exp\Big\{\delta\Big[r + \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] + \frac{\sqrt{\delta}}{1-\eta}b'Q^{-1}\sigma Z(T+1)\Big\} \tag{16}$$

is the associated benchmark investor's wealth. It now follows that the normalized wealth process $z^H_\pi(\delta) = \frac{x^H_\pi(\delta)}{y^H(\delta)}$ satisfies

$$\frac{x^H_\pi(\delta)}{y^H(\delta)} = \frac{x(0)}{y(0)} \exp\Big\{\delta\Big[b'\pi - \frac{1}{2}\pi' Q\pi - \frac{1-2\eta}{2(1-\eta)^2} b' Q^{-1} b\Big] + \sqrt{\delta}\Big[\pi - \frac{1}{1-\eta} Q^{-1} b\Big]' \sigma Z(T+1)\Big\}$$

so the relative regret problem (14) becomes

$$\begin{cases} \quad \max_{\pi\in\mathcal{G}_T} \min_{H\in\mathcal{H}} \mathbb{E}_H[U(z^H_\pi(\delta))] \\ \text{subject to:} \\ \quad z^H_\pi(\delta) = z(0)\exp\Big\{\delta\Big[b'\pi - \frac{1}{2}\pi'Q\pi - \frac{1-2\eta}{2(1-\eta)^2} b'Q^{-1}b\Big] + \sqrt{\delta}\Big[\pi - \frac{1}{1-\eta}Q^{-1}b\Big]'\sigma Z(T+1)\Big\} \end{cases} \tag{17}$$

where we have substituted $z^H_\pi(\delta)$ for $\frac{x^H_\pi(\delta)}{y^H(\delta)}$. Choosing $U(z) = \frac{1}{\gamma} z^\gamma$ as the comparison function in (17), it follows from the log-normality of $z^H_\pi(\delta)$ (conditional on $\mathcal{G}_T$ and H) that the conditional expectation

$$\begin{aligned} &\mathbb{E}_H\left[U(z^H_\pi(\delta)) \,\middle|\, \mathcal{G}_T\right] \\ &\quad = \frac{1}{\gamma}\mathbb{E}_H\left[z^H_\pi(\delta)^\gamma \,\middle|\, \mathcal{G}_T\right] \\ &\quad = \frac{z(0)^\gamma}{\gamma}\mathbb{E}_H\Big[\exp\Big\{\delta\gamma\Big[b'\pi - \frac{1}{2}\pi'Q\pi - \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] + \gamma\sqrt{\delta}\Big[\pi - \frac{1}{1-\eta}Q^{-1}b\Big]'\sigma Z(T+1)\Big\}\,\Big|\,\mathcal{G}_T\Big] \\ &\quad = \frac{z(0)^\gamma}{\gamma}\exp\Big\{\delta\gamma\Big[\frac{1-\eta-\gamma}{1-\eta}b'\pi - \frac{1-\gamma}{2}\pi'Q\pi - \frac{1-\gamma-\eta}{2(1-\eta)^2}b'Q^{-1}b\Big]\Big\}. \end{aligned}$$

It now follows that (17) is equivalent to

$$\max_{\pi\in\mathcal{G}_T}\min_{H\in\mathcal{H}} \frac{z(0)^\gamma}{\gamma}\mathbb{E}_H\Big[\exp\Big\{\delta\gamma\Big[\frac{1-\eta-\gamma}{1-\eta}b'\pi - \frac{1-\gamma}{2}\pi'Q\pi - \frac{1-\gamma-\eta}{2(1-\eta)^2}b'Q^{-1}b\Big]\Big\}\Big]. \tag{18}$$

It is interesting to note the similarity between (12) and (18). We shall expand further on this in later sections.

4 Optimal portfolio and learning model

We characterize the learning model and optimal portfolio for the relative regret problem with power-type utility and comparison functions. We also derive an approximate solution which is exact in the limit that $\delta \downarrow 0$ that shows the explicit dependence of the optimal portfolio on data.

4.1 Convex duality and robust learning

We begin by characterizing the learning model and optimal portfolio using results from convex analysis. For the purposes of accessibility, we derive our results for the case that the uncertainty set $\mathcal{H} = \{H_1, \cdots, H_m\} = \{(b_1, Q_1), \cdots, (b_m, Q_m)\}$ is finite, and for simplicity, assume that the comparison and utility functions are power-type (i.e. $U(z) = \frac{1}{\gamma} z^\gamma$ and $U^B(y) = \frac{1}{\eta} y^\eta$). The case of compact and uncountable uncertainty set $\mathcal{H}$ is handled using similar ideas but requires some heavier machinery and is proven in the Appendix.

Observe firstly that when $\mathcal{H} = \{H_1, \cdots, H_m\} = \{(b_1, Q_1), \cdots, (b_m, Q_m)\}$, that (14) is equivalent to

$$\begin{cases} \nu^* = \max_{\pi,\, \kappa} \kappa \\ \text{subject to:} \\ \mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] \geq \kappa, \quad i = 1, \cdots, m. \end{cases} \tag{19}$$

When the utility and comparison functions are power-type, (18) gives

$$\begin{aligned} &\mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] \\ &= \frac{z(0)^\gamma}{\gamma}\mathbb{E}_{H_i}\Big[\exp\Big\{\delta\gamma\Big[\frac{1-\eta-\gamma}{1-\eta} b_i'\pi - \frac{1-\gamma}{2}\pi' Q_i \pi - \frac{1-\gamma-\eta}{2(1-\eta)^2} b_i' Q_i^{-1} b_i\Big]\Big\}\Big]. \end{aligned} \tag{20}$$

Observe that $\mathbb{E}_{H_i}[\frac{1}{\gamma}(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)})^\gamma]$ is concave in π whenever $\gamma < 0$, and (19) is a convex optimization problem in (π, κ) when this condition holds. We shall enforce this assumption unless otherwise stated.

Let $\mu_i \geq 0$ denote the Lagrange multiplier for the i^{th} constraint in (19). Clearly, if (π, κ) is feasible for (19) and $\mu_i \geq 0$ for $i = 1, \cdots, m$, then

$$\begin{aligned} L(\pi, \kappa, \mu) &= \kappa + \sum_{i=1}^m \mu_i \Big\{\mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] - \kappa\Big\} \\ &= \kappa\Big(1 - \sum_{i=1}^m \mu_i\Big) + \sum_{i=1}^m \mu_i \mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] \\ &\geq \nu^* \end{aligned}$$

Define the dual objective function

$$\begin{aligned} \psi(\mu) &= \max_{\pi \in \mathcal{G}_T,\, \kappa} L(\pi, \kappa, \mu) \\ &= \max_{\pi \in \mathcal{G}_T,\, \kappa} \Big\{\kappa\Big(1 - \sum_{i=1}^m \mu_i\Big) + \sum_{i=1}^m \mu_i \mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big]\Big\}. \end{aligned}$$

Clearly, $\psi(\mu)$ is finite if and only if $\sum_{i=1}^m \mu_i = 1$, under which it follows that

$$\psi(\mu) = \max_{\pi \in \mathcal{G}_T} \sum_{i=1}^m \mu_i \mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big].$$

Observing that the Lagrange multipliers $\mu = [\mu_1, \cdots, \mu_m]'$ are all non-negative and sum to 1, it follows that μ can be interpreted as a probability measure on the class of models $\mathcal{H} = \{H_1, \cdots H_m\}$, and that the summation in the dual function is nothing but an expectation where the Lagrange multiplier μ plays the role of a prior distribution. Particularly

$$\psi(\mu) = \max_{\pi \in \mathcal{G}_T} \mathbb{E}_\mu \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big], \tag{21}$$

which is nothing but a Bayesian problem with prior μ.

All that remains is to relate the dual function $\psi(\mu)$ to the original optimization problem (19). The following result is an immediate consequence of Lagrangian duality (see Theorem 1 on pp. 224 of [11]). The general proof of this result (which extends the analysis to the case when $\mathcal{H}$ is possibly uncountable though compact) can be found in the Appendix.

Proposition 4.1 *Suppose that $\mathcal{H}$ is compact, that $\gamma < 0$ and $\eta < 1$. Let ν^* denote the optimal value of the relative regret problem (14) (or (19)), and $\psi(\mu)$ be the value function (dual function) for the Bayesian problem (21) when the prior (Lagrange multiplier) is μ. Then dual optimization problem*

$$\psi(\mu^*) = \min_{\mu \geq 0,\, \mu(\mathcal{H})=1} \psi(\mu) \tag{22}$$

has a solution μ^* and the optimal regret objective value satisfies $\nu^* = \psi(\mu^*)$. The optimal portfolio for the relative regret problem is the maximizer in (21) under μ^*, namely

$$\pi^* = \arg\max_{\pi \in \mathcal{G}_T} \mathbb{E}_{\mu^*} \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big] \tag{23}$$

Proposition 4.1 tells us that the solution of the relative regret problem is also the solution of a Bayesian problem (23) where the optimal Lagrange multiplier μ^* characterized by (22) also plays the role of a prior distribution on the set of models $\mathcal{H}$.

This result allows us to see the dependence of the optimal portfolio π^* on the observations $\underline{\mathcal{R}}_T = \{\mathcal{R}(1), \cdots, \mathcal{R}(T)\}$, and hence to characterize the optimal learning model. Specifically, if μ^* is the optimal prior/Lagrange multiplier from (22), and $\mu_T^* = [\mu_T^*(1), \cdots, \mu_T^*(m)]$ is the posterior obtained from Bayesian updating conditional on data $\underline{\mathcal{R}}_T$, then the objective function can be written as

$$\mathbb{E}_{\mu^*} \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big] = \mathbb{E}_{\mu^*} \Big[\mathbb{E}_{\mu^*} \Big\{ \Big(\frac{1}{\gamma} \frac{x(\delta)}{y(\delta)} \Big)^\gamma \Bigm| \underline{\mathcal{R}}_T \Big\} \Big] = \mathbb{E}_{\mu^*} \Big\{ \mathbb{E}_{\mu_T^*} \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big] \Big\},$$

where

$$\mathbb{E}_{\mu_T^*} \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big] = \begin{cases} \displaystyle\sum_{i=1}^m \mu_T^*(i) \mathbb{E}_{H_i} \Big[\frac{1}{\gamma} \Big(\frac{x^{H_i}(\delta)}{y^{H_i}(\delta)} \Big)^\gamma \Big], & \mathcal{H} \text{ is finite,} \\ \displaystyle\int_{\mathcal{H}} \mathbb{E}_H \Big[\frac{1}{\gamma} \Big(\frac{x_\pi^H(\delta)}{y^H(\delta)} \Big)^\gamma \Big] \mu^*(dH), & \mathcal{H} \text{ is uncountable and compact.} \end{cases}$$

This means that

$$\psi(\mu^*) = \max_{\pi\in\mathcal{G}_T} \mathbb{E}_{\mu^*}\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^{\gamma}\Big] = \mathbb{E}_{\mu^*}\Big\{\max_{\pi\in\mathcal{G}_T} \mathbb{E}_{\mu_T^*}\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^{\gamma}\Big]\Big\},$$

where the second equality follows from the $\mathcal{G}_T$-measurability of π. The dependence of the optimal portfolio π^* on the data can now be summarized as follows.

Proposition 4.2 *Suppose that $\mathcal{H}$ is compact, that $\gamma < 0$ and $\eta < 1$, and that the Lagrange multiplier/prior distribution μ^* on $\mathcal{H}$ is the solution of the dual problem (22). Let $\mu_T^* = [\mu_T^*(1), \cdots, \mu_T^*(m)]$ denote the posterior distribution on $\mathcal{H}$ obtained from Bayesian updating of μ^* given the observations $\underline{\mathcal{R}}_T = \{\mathcal{R}(1), \cdots, \mathcal{R}(T)\}$. Then*

$$\pi^* = \arg\max_{\pi\in\mathcal{G}_T} \mathbb{E}_{\mu_T^*}\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^{\gamma}\Big] \tag{24}$$

is optimal for the relative regret problem (14).

4.2 Optimal portfolio: approximate solution

Although we have characterized the optimal portfolio (24) and optimal dual variable/prior (22), both appear difficult to compute. In this section we derive an approximate characterization of the optimal portfolio π^* and the associated Lagrange multiplier/prior μ^*, which becomes exact as $\delta \to 0$. The basic idea is to approximate the RHS of (24) so that the maximization over π can be solved explicitly and a closed form expression for the dual function $\psi(\mu)$ can be obtained. The resulting dual problem is easier to solve than (22) while the expression for π^* shows us the explicit dependence of the (approximately) optimal solution on the posterior μ_T^*.

As a start, recall that

$$\mathbb{E}_H\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^{\gamma}\Big] = \frac{1}{\gamma}\mathbb{E}_H\Big[\exp\Big\{\delta\gamma\Big[\frac{1-\eta-\gamma}{1-\eta}b'\pi - \frac{1-\gamma}{2}\pi' Q\pi - \frac{1-\gamma-\eta}{2(1-\eta)^2}b'Q^{-1}b\Big]\Big\}\Big].$$

A Taylor expansion of $\exp\{\cdots\}$ in orders of δ gives

$$\begin{aligned}
&\frac{1}{\gamma}\mathbb{E}_H\exp\{\cdots\} \\
&= \frac{1}{\gamma}\mathbb{E}_H\Big[1 + \delta\gamma\Big\{\frac{1-\eta-\gamma}{1-\eta}b'\pi - \frac{1-\gamma}{2}\pi' Q\pi - \frac{1-\eta-\gamma}{2(1-\eta)^2}b'Q^{-1}b\Big\} + o(\delta)\Big] \\
&= \mathbb{E}_H\Big[\frac{1}{\gamma} + \delta\frac{1-\gamma}{2}\Big\{2\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}b'\pi - \pi' Q\pi - \frac{1-\eta-\gamma}{(1-\eta)^2(1-\gamma)}b'Q^{-1}b\Big\} + o(\delta)\Big].
\end{aligned}$$

The dual function (21) can now be written as

$\psi(\mu)$

$$
\begin{aligned}
&= \mathbb{E}_\mu\Big\{\max_{\pi\in\mathcal{G}_T}\mathbb{E}_{\mu_T}\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^\gamma\Big]\Big\}, \\
&= \mathbb{E}_\mu\Big\{\max_{\pi\in\mathcal{G}_T}\int_{\mathcal{H}}\mathbb{E}_H\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma\Big]\mu_T(dH)\Big\} \\
&= \frac{1}{\gamma}+\delta\frac{1-\gamma}{2}\mathbb{E}_\mu\Big\{\max_{\pi\in\mathcal{G}_T}\mathbb{E}_{\mu_T}\Big[2\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}b'\pi-\pi'Q\pi-\frac{1-\eta-\gamma}{(1-\eta)^2(1-\gamma)}b'Q^{-1}b\Big]\Big\}+o(\delta).
\end{aligned}
$$

With this in mind, define

$$
\begin{aligned}
\bar{\psi}(\mu) &= \mathbb{E}_\mu\Big[\max_{\pi\in\mathcal{G}_T}\mathbb{E}_{\mu_T}\Big\{2\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}b'\pi-\pi'Q\pi-\frac{1-\eta-\gamma}{(1-\eta)^2(1-\gamma)}b'Q^{-1}b\Big\}\Big] \\
&= \mathbb{E}_\mu\Big\{\max_{\pi\in\mathcal{G}_T}\Big[2\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}\mathbb{E}_{\mu_T}(b)'\pi-\pi'\mathbb{E}_{\mu_T}(Q)\pi\Big]\Big\} \\
&\quad -\frac{1-\eta-\gamma}{(1-\eta)^2(1-\gamma)}\mathbb{E}_\mu(b'Q^{-1}b)
\end{aligned}
$$

where

$$
\mathbb{E}_{\mu_T}(Q) \equiv \int_{\mathcal{H}} Q\mu_T(dH), \quad \mathbb{E}_{\mu_T}(b)\equiv\int_{\mathcal{H}} b\mu_T(dH)
$$

and

$$
\mathbb{E}_\mu(b'Q^{-1}b)=\mathbb{E}_\mu[\mathbb{E}_{\mu_T}(b'Q^{-1}b)].
$$

Several lines of algebra show that

$$
\begin{aligned}
\bar{\psi}(\mu) &= -\mathbb{E}_\mu\Big\{\min_{\pi\in\mathcal{G}_T}\Big[\pi-\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}\left[\mathbb{E}_{\mu_T}(Q)\right]^{-1}\mathbb{E}_{\mu_T}(b)\Big]'\mathbb{E}_{\mu_T}(Q) \\
&\quad\times\Big[\pi-\frac{1-\eta-\gamma}{(1-\eta)(1-\gamma)}\left[\mathbb{E}_{\mu_T}(Q)\right]^{-1}\mathbb{E}_{\mu_T}(b)\Big]\Big\}+\Big[\frac{\eta}{(1-\gamma)(1-\eta)}\Big]^2\mathbb{E}_\mu(b'Q^{-1}b) \\
&\quad+\Big[\frac{1-\gamma-\eta}{(1-\gamma)(1-\eta)}\Big]^2\Big[\mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\}-\mathbb{E}_\mu(b'Q^{-1}b)\Big].
\end{aligned}
$$

Clearly

$$
\bar{\pi}^\mu=\frac{1-\gamma-\eta}{(1-\eta)(1-\gamma)}[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b) \tag{25}
$$

is the optimal portfolio from which it follows that

$$
\begin{aligned}
\bar{\psi}(\mu) &\triangleq \Big[\frac{\eta}{(1-\gamma)(1-\eta)}\Big]^2\mathbb{E}_\mu(b'Q^{-1}b) \\
&\quad+\Big[\frac{1-\gamma-\eta}{(1-\gamma)(1-\eta)}\Big]^2\Big[\mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\}-\mathbb{E}_\mu(b'Q^{-1}b)\Big]. \qquad (26)
\end{aligned}
$$

Recalling that

$$
\psi(\mu)=\frac{1}{\gamma}+\delta\frac{1-\gamma}{2}\bar{\psi}(\mu)+o(\delta)
$$

(and noting that the coefficient of $\bar{\psi}(\mu)$ is positive) it follows that an approximate solution $\bar{\mu}^*$ of the dual problem (22) can be obtained by solving

$$\bar{\mu}^* = \arg \min_{\mu \geq 0,\, \mu(\mathcal{H})=1} .\bar{\psi}(\mu) \tag{27}$$

while

$$\bar{\pi}^* = \frac{1-\gamma-\eta}{(1-\eta)(1-\gamma)}[\mathbb{E}_{\bar{\mu}_T^*}(Q)]^{-1}\mathbb{E}_{\bar{\mu}_T^*}(b) \tag{28}$$

is an approximate solution for the regret problem.

A similar calculation/approximation can be carried out for (12). In this case, the (approximate) dual problem is again given by (27) where

$$\bar{\psi}(\mu) \quad = \quad \mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\} - \mathbb{E}_\mu(b'Q^{-1}b).$$

The (approximate) optimal policy is given by

$$\bar{\pi}^* \quad = \quad \frac{1}{1-\eta}[\mathbb{E}_{\mu_T^*}(Q)]^{-1}\mathbb{E}_{\mu_T^*}(b).$$

Interestingly, this is the extreme case of $\gamma \to -\infty$ in (26) and (28), which coincides with a large aversion to missing the benchmark.

5 Relative regret v's worst case

In this section we compare the (approximate) solution (27)-(28) of the relative regret problem ((14) or (18)) to that of the typical "worst case" robust portfolio selection problem

$$\begin{cases} \max_{\pi \in \mathcal{G}_T} \min_{H \in \mathcal{H}} \frac{1}{\gamma}\mathbb{E}^H[x_\pi^H(\delta)^\gamma] \\ \text{subject to:} \\ x_\pi^H((\delta)) = x \exp\Big\{\delta\Big[r + b'\pi - \frac{1}{2}\pi' Q\pi\Big] + \sqrt{\delta}\pi'\sigma Z(T+1)\Big\}. \end{cases} \tag{29}$$

which is given by

$$\pi^* \quad = \quad \frac{1}{1-\gamma}[Q^*]^{-1}b^* \tag{30}$$

$$H^* \quad = \quad (Q^*,\, b^*) = \arg \min_{\mu \geq 0,\, \mu(\mathcal{H})=1} \mathbb{E}_\mu[b'Q^{-1}b] = \arg \min_{(Q,\, b) \in \mathcal{H}} b'Q^{-1}b. \tag{31}$$

The solution π^* of the worst case problem (29) can be described as follows: (i) Find the model $(Q^*,\, b^*)$ in $\mathcal{H}$ with the smallest Sharpe ratio (i.e. equation (31)), and (ii) solve a standard Bayesian problem with a prior that puts all its mass on $(Q^*,\, b^*)$ (i.e. equation (30)).

The solution (30)-(31) of the worst case problem (29) is problematic on several grounds. Firstly, investing according to a prior that puts all its mass on the model with the smallest Sharpe ratio seems overly pessimistic and is sensitive to the choice of uncertainty set. This recommendation is a consequence of the worst case objective which is only concerned about performance for the worst-case model (Q^*, b^*), but is unconcerned about under-performing when "better" models (e.g. the one with the largest Sharpe ratio) apply. In contrast, the relative regret objective (14) favors decisions that perform well for both pessimistic (small Sharpe ratio) as well as optimistic (large Sharpe ratio) models. This manifests itself in a prior (defined as the solution of (27)) that puts mass on pessimistic as well as optimistic models. Secondly, since the worst case prior (31) is degenerate, then so too is the posterior. That is, the "worst case" equilibrium portfolio (30) resolutely sticks to the "worst case" model (Q^*, b^*) and ignores the data $\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T)$, even if it strongly suggests that returns are not being generated by (Q^*, b^*) but by some other model. In other words, learning does not occur, even when it is possible. In contrast, non-degeneracy of the prior in the relative regret solution implies that the posterior will not be static but updated in response to historical returns $\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T)$.

6 Experiments

Coming soon.

7 Conclusion

We have formulated the problem of portfolio selection with parameter uncertainty in the framework of relative regret. We show that the solution of this problem is also that of another non-standard Bayesian problem where the objective involves the benchmarks associated with the relative regret model, and the prior is chosen endogenously. We solve this "Bayesian benchmarking" problem and show that the solution of the dynamic relative regret problem involves a "tilted" posterior, where the tilting is done by way of a likelihood ratio which depends on the family of benchmarks and favors the model which has performed the best on the history of realized returns. The results in this paper can be generalized to continuous time jump-diffusion models with uncertainty about the jump intensity (as well as the mean return). This is an interesting extension given the difficulties associated with estimating the intensity of the jump process. It is also possible to consider alternative benchmarks, such as those which weight recent returns more heavily than past returns. This is of interest when returns are only locally stationary. These results will be reported elsewhere.

Appendix

Proof of Proposition 4.1

The problem (14) is equivalent to

$$\left\{\begin{array}{l} \quad \max_{\pi\in\mathcal{G}_T,\,\kappa}\kappa \\ \text{subject to:} \\ \quad \mathbb{E}_H\left(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\right)^\gamma \geq \kappa, \quad \forall\, H\in\mathcal{H}, \end{array}\right. \tag{32}$$

which is a convex optimization problem in $(\pi,\,\kappa)$ for all values of $\gamma < 0$ and $\eta < 1$. We will analyze this problem using convex duality, for which the following definitions are required. For more details, the reader should consult Luenberger [11].

Let $C(\mathcal{H})$ denote the space of real-valued continuous functionals on $\mathcal{H}$ with sup-norm

$$\|g\| \triangleq \sup_{(Q,\,b)\in\mathcal{H}} |g(Q,\,b)|, \quad \forall\, g\in C(\mathcal{H}).$$

The linear space $C(\mathcal{H})$ with this norm is a Banach space [5]. Let

$$\mathcal{P} \triangleq \{g\in C(\mathcal{H})\,|\,g(Q,\,b)\geq 0,\ \forall\,(Q,\,b)\in\mathcal{H}\}$$

define the positive cone in $C(\mathcal{H})$. It is easy to see that $\mathcal{P}$ has non-empty interior[3]. We say that $f\geq g$ for $f,\,g\in C(\mathcal{H})$ if $f-g\in\mathcal{P}$ and $g\leq 0$ if $-g\in\mathcal{P}$. We write $g>0$ if $g(Q,\,b)>0$ for every $(Q,\,b)\in\mathcal{H}$. Next, let $\mathcal{B}(\mathcal{H})$ denote the set of Borel sets of $\mathcal{H}$. The dual (or conjugate) space $C^*(\mathcal{H})$ is (isomorphic to) the set of measures defined on $\mathcal{B}(\mathcal{H})$ with bounded total variation:

$$C^*(\mathcal{H}) = \left\{\mu\,\Big|\,\int_{\mathcal{H}}|\mu(dH)|<\infty\right\};$$

see for example Section IV.6.3 in Dunford and Schwartz [5]. Observe that elements of $C^*(\mathcal{H})$ are signed measures. The dual cone of $\mathcal{P}$ is defined by $\mathcal{P}^*\triangleq\{\mu\in C^*(\mathcal{H})\,|\,\int_{\mathcal{H}}f\,d\mu\geq 0,\ \forall\, f\in\mathcal{P}\}$ (see [11]) and is equal to the subset of $\mathcal{C}^*(\mathcal{H})$ consisting of positive measures:

$$\mathcal{P}^* = \{\mu\in C^*(\mathcal{H})\,|\,\mu(A)\geq 0,\ \forall\, A\in\mathcal{B}(\mathcal{H})\}. \tag{33}$$

We write $\mu\geq 0$ when $\mu\in\mathcal{P}^*$.

Let $\mu\in\mathcal{P}^*$ be arbitrary and $(\pi,\,\kappa)$ be feasible for (32). It is clear that

$$\begin{aligned} L(\pi,\,\kappa,\,\mu) &\triangleq \kappa + \int_{H\in\mathcal{H}}\left[\mathbb{E}_H\left(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\right)^\gamma - \kappa\right]\mu(dH) \\ &= \kappa(1-\mu(\mathcal{H})) + \int_{H\in\mathcal{H}}\mathbb{E}_H\left(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\right)^\gamma\mu(dH) \\ &\geq \nu^*. \end{aligned} \tag{34}$$

[3]This is needed for certain strong duality results.

Define the dual function $\psi(\mu)$ as

$$\begin{aligned}\psi(\mu) &\triangleq \max_{\kappa\in\mathbb{R},\,\pi\in\mathcal{G}_T} L(\pi,\,\kappa,\,\mu)\\ &= \max_{\kappa\in\mathbb{R},\,\pi\in\mathcal{G}_T} \kappa(1-\mu(\mathcal{H})) + \int_{H\in\mathcal{H}} \mathbb{E}_H\Big(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma \mu(dH).\end{aligned}$$

From our construction of $L(\pi,\,\kappa,\,\mu)$, $\psi(\mu)$ is an upper bound on ν^* for every $\mu \in \mathcal{P}^*$. This upper bound is finite if and only if $\mu(\mathcal{H}) = 1$ (i.e. a probability measure on $\mathcal{H}$) from which it follows that

$$\psi(\mu) = \max_{\pi\in\mathcal{G}_T} \int_{H\in\mathcal{H}} \mathbb{E}_H\Big(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma \mu(dH).$$

Observing that the integral is nothing but an expectation with respect to a probability distribution μ, we adopt the notation

$$\mathbb{E}_\mu\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^\gamma\Big] \equiv \int_{H\in\mathcal{H}} \mathbb{E}_H\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma\Big]\mu(dH)$$

and we can write

$$\psi(\mu) = \max_{\pi\in\mathcal{G}_T} \mathbb{E}_\mu\Big\{\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^\gamma\Big\}$$

which is precisely the dual function (21). It now follows from Theorem 1 on pg. 224 of [11] that the optimal relative regret objective value ν^* and the dual objective value $\psi(\mu)$ are related by

$$\nu^* = \psi(\mu^*) = \min_{\mu\geq 0,\,\mu(\mathcal{H})=1} \psi(\mu)$$

and that the optimal solution π^{μ^*} of (23) with prior μ^* is also the optimal solution of the relative regret problem (32).

References

[1] M.W. Brandt. Portfolio choice problems. *Handbook of Financial Econometrics*, Y. Air-Sahalia and L.P. Hansen (eds), forthcoming.

[2] J.Y. Campbell and L.M. Viciera. *Strategic Asset Allocation*, Oxford University Press, 2002.

[3] T.M. Cover. Universal portfolios. *Mathematical Finance*, 1, pp 1–29, 1991.

[4] P.M. DeMarzo, I. Kremer and Y. Mansour. On Hannan and Blackwell's approachability and options - a game theoretic approach fpr option pricing. Working paper, Graduate School of Business, Stanford University, 2005.

[5] N. Dunford and J.T. Schwartz. *Linear Operators, Part 1: General Theory*. Wiley, 1988.

[6] T. Hayashi. Regret aversion and opportunity dependence. Working paper, Department of Economics, University of Texas (Austin), 2006.

[7] T.A. Knox. Learning how to invest when returns are uncertain. Working paper, 2002.

[8] A.E.B. Lim and J.G. Shanthikumar. Multi-period portfolio choice with learning in the framework of relative regret. Working paper, IEOR Department, University of California, Berkeley, 2009.

[9] A.E.B. Lim and J.G. Shanthikumar. Robust portfolio selection with benchmarked objectives. Working paper, IEOR Department, University of California, Berkeley, 2007.

[10] A.E.B. Lim, J.G. Shanthikumar and Z.J. (Max) Shen. Model uncertainty, robust optimization, and learning. *TutORials in Operations Research*, 3, pp 66-94, 2006.

[11] D.G. Luenberger. *Optimization by Vector Space Methods*, John Wiley, New York, 1968.

[12] D.M. McCloskey. The loss function has been mislaid: The rhetoric of significance tests. *American Economic Review*, 75(2), pp 201–205, May 1985.

[13] L.C.G. Rogers. The relaxed investor and parameter uncertainty. *Finance and Stochastics*, 5, pp 131–154, 2001.

[14] D. Terlizzese. Relative minimax. Working paper, 2006.

[15] A. Wald. Contributions to the theory of statistical estimation and testing hypotheses. *Annals of Mathematical Statistics*, 10, 00 299-326, December 1939.