Title
New approaches to robustness and learning in data-driven portfolio optimization

Author
Vahn, Gah-Yi

Publication Date
2012

Peer reviewed|Thesis/dissertation

New approaches to robustness and learning in data-driven portfolio optimization

by

Gah-Yi Vahn

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering – Industrial Engineering & Operations Research

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Associate Professor Andrew E.B. Lim, Chair
Laurent El Ghaoui
Noureddine El Karoui
George Shanthikumar

Spring 2012

New approaches to robustness and learning in data-driven portfolio optimization

Abstract

New approaches to robustness and learning in data-driven portfolio optimization

by

Gah-Yi Vahn
Doctor of Philosophy in Engineering – Industrial Engineering & Operations Research

University of California, Berkeley

Associate Professor Andrew E.B. Lim, Chair

We develop two new approaches to robustness and learning in data-driven portfolio optimization, a problem that is well-known for sensitivity to model assumptions and data variability.

First, we consider the data-driven mean-CVaR problem. For this problem, we introduce and investigate *performance-based regularization* (PBR), a generalization of standard regularization techniques used in statistics and machine learning, and an alternative to worst-case approaches to improving solution robustness. We assume the available log-return data is iid, and detail the approach for two cases: nonparametric and parametric (the log-return distribution belongs in the elliptical family). We derive the asymptotic behavior of the nonparametric PBR solution, which leads to insight into the effect of penalization, and justification of the parametric PBR method. We also show via simulations that the PBR methods produce efficient frontiers that are, on average, closer to the population efficient frontier than the empirical approach to the mean-CVaR problem, with less variability.

Next, we consider portfolio optimization under parameter uncertainty, and propose optimizing a *relative regret* objective. 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 is a natural objective when there is concern about parameter uncertainty or model ambiguity. 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 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. These results show that regret is a natural means by which robust decision making and learning can be combined.

To all those who have taught me in some way.

Contents

List of Figures

Acknowledgments

I have been blessed with lots of guidance, support and friendship during my time at Berkeley. I would like to use this space to thank those who have played an important part in making my PhD a wonderful experience.

I am particularly grateful to have been advised by Andrew Lim, George Shanthikumar and Noureddine El Karoui. I have learned much from each one of you, not just in research skills but also in humility and good humor. Extra thank you to Andrew, who, as my main advisor, has always encouraged me with enthusiasm and supported me financially through multiple RA's. I also thank Laurent El Ghaoui, who has shown much interest in my research.

My family and classmates have provided lots of love and support throughout graduate school. I especially thank Gemma (it was really nice to go through this with you), Anca, Sangi, Cagla, Nguyen, Soma, Pedro and Albert (for all the fun times), Caroline and Ankit (for your words of wisdom), Hua, Wayne, Oleksa and Hang (for all those memorable sessions in Evans Hall), and of course, Gerald and my parents.

I am also grateful for my teachers and friends at the Empty Gate Zen Center. ZM Bon Seong, ZM Dae Kwang and Jason Quinn: thank you for your teaching and guidance. Ancho, Kendra, Dennis, Caren, Eunice, Young Mi, Julie: thank you for practicing with me.

Finally, I would like to thank my special furry friends of Berkeley, in particular Lemmy, Lefty and Kukui. I will truly miss you.

Chapter 1

Introduction

Portfolio optimization is one approach to allocating limited resources to a given number of alternatives. In his classic paper, Markowitz [48] introduced mean-variance portfolio optimization in a financial context. The paper proposes investing in securities in an *efficient* manner, that is, choosing a portfolio that maximizes expected return subject to an upper bound on the variance, or equivalently, a portfolio that minimizes the future variance subject to a lower bound on the expected return.

Many variations and extensions of the Markowitz problem have since been proposed. One important line of enquiry is in relaxing the assumption in Markowitz's original work that the portfolio's future expected return and variance are known. This has introduced the use of data to estimate the unknown quantities in portfolio optimization, and many empirical studies followed. A common consensus, however, is that solutions to portfolio optimization with estimated return and risk have poor out-of-sample performance because it is difficult to accurately forecast future return or variance from historical data (see [51], [40] and [13], to name a few).

This brings about the central question of this thesis — how may we incorporate data in a meaningful way in portfolio optimization? In particular, we seek methods that are robust to data and model uncertainty that nevertheless incorporate learning.

In Chapter 2, we investigate data-driven mean-CVaR portfolio optimization. CVaR (Conditional Value-at-Risk) has been of current interest as it is a risk measure that accounts for large losses, in contrast to variance, which is a symmetric measure of risk. However, as CVaR is a risk measure based on rare observations, mean-CVaR portfolio optimization is also plagued by poor out-of-sample performance. To address this, we propose *performance-based regularization* (PBR) to improve solution performance in the return-risk space. This method is a generalization of standard regularization techniques used in statistics and machine learning, and is an alternative to worst-case approaches to improving solution robustness.

Specifically, we assume the available log-return data is iid, and detail the approach for two cases: nonparametric and parametric (the log-return distribution belongs in the elliptical family). The nonparametric PBR method penalizes portfolios with

large variability in mean and CVaR estimations. The parametric PBR method solves the empirical Markowitz problem instead of the empirical mean-CVaR problem, as the solutions of the Markowitz and mean-CVaR problems are equivalent when the log-return distribution is elliptical. We derive the asymptotic behavior of the non-parametric PBR solution, which leads to insight into the effect of penalization, and justification of the parametric PBR method. We also show via simulations that the PBR methods produce efficient frontiers that are, on average, closer to the population efficient frontier than the empirical approach to the mean-CVaR problem, with less variability.

In Chapter 3, we consider portfolio optimization under parameter uncertainty. We assume the asset log-returns are iid Gaussian, but with unknown mean and covariance. We propose dealing with parameter uncertainty by solving a *relative regret* objective. Relative regret evaluates a decision by comparing its return to a family of benchmarks, where the benchmarks are the wealth of fictitious investors who invest optimally given knowledge of model parameters. The optimal relative regret portfolio maximizes the worst-case relative performance to all the benchmarks, over the family of possible parameters. Analysis of the problem using convex duality shows that it is equivalent to a nonstandard Bayesian problem, where the prior distribution is endogenously chosen by solving the dual optimization problem, with the resulting Lagrange multipliers playing the role of the prior. The learning process involves Bayesian updating of these Lagrange multipliers/prior. This approach is a genuine alternative to standard Bayesian methods to model uncertainty, which can be sensitive to the choice of the prior. It is also an alternative to maximizing the worst-case absolute performance, which may lead to conservative solutions. Our investigation shows that regret is a natural approach by which robust decision making and learning can be combined.

Chapter 2

Performance-based regularization in mean-CVaR portfolio optimization

2.1 Motivation

In recent years, there has been a growing interest in Conditional Value-at-Risk (CVaR) as a financial risk measure. This interest is based on two key advantages of CVaR over Value-at-Risk (VaR), the risk measure of choice in the financial industry over the last twenty years. Firstly, $CVaR(\beta)$, the conditional expectation of losses in the top $100(1-\beta)\%$ ($\beta = 0.95, 0.99$ are typical values used in industry), is more informative about the tail end of the loss distribution than $VaR(\beta)$, which is only the *threshold* for losses in the top $100(1-\beta)\%$. Secondly, [1] showed that CVaR satisfies the four coherence axioms of [6], whereas VaR fails the subadditivity requirement.

Portfolio optimization with CVaR as a risk measure is first studied by [55], who show that empirical CVaR minimization can be formulated as a linear program. Subsequent works include CVaR optimization for a portfolio of credit instruments [4] and derivatives [2], and portfolio optimization based on extensions of CVaR [47]. However, most discussions of CVaR in portfolio optimization to date are concerned with formulation and tractability of the problem, and assume full knowledge of the distribution of the portfolio loss. In practice, one cannot ignore the fact that the loss distribution is not known and must be estimated from historical data, constructed from expert knowledge, or a combination of both. Naive estimation of the loss distribution can pose serious problems — [42] demonstrates how fragile the solution to the empirical mean-CVaR problem is, even in the ideal situation of having iid Gaussian log-return data.

The issue of estimation errors in portfolio optimization is not, however, new knowledge. The estimation issue for the classical Markowitz (mean-variance) problem has

been raised as early as 1980 [39]. There have since been many suggestions for mitigating this issue for the Markowitz problem; two main approaches are robust optimization [27] and what we call "standard regularization" ([18], [25], [38], [21]). The robust optimization approach is to take the source of uncertainty (e.g. the asset log-returns, or its distribution), specify an uncertainty set about the source, and minimize the worst-case return-risk problem over this uncertainty set. The standard regularization approach is to solve the empirical mean-variance problem, but with a constraint on the size of the solution, as measured by L_2 or a more generalized norm. The term "regularization" is adopted from statistics and machine learning, where it refers to controlling for the size of the decision variable for better out-of-sample performance [35]. Both robust optimization and standard regularization approaches have been studied for the mean-CVaR problem; [29] and [63] show implementations of the robust optimization approach when the source of uncertainty is, respectively, the log-return vector and the log-return distribution, and [30] demonstrates implementation of standard regularization.

In this chapter, we propose performance-based regularization (PBR), a new approach to addressing estimation risk in data-driven optimization, and illustrate this method for the mean-CVaR portfolio optimization problem. We demonstrate PBR for two situations: the investor has nonparametric or parametric (specifically, the elliptical family of distributions describe the log-returns) information on the log-returns.

The nonparametric PBR method penalizes portfolios with large variability in mean and CVaR estimations. Specifically, we penalize the sample variances of the mean and CVaR estimators. The resulting problem is a combinatorial optimization problem, however we show that its convex relaxation, a quadratically-constrained quadratic program, is tight. The problem can be interpreted as a chance-constrained program that picks portfolios for which approximate probabilities of deviations of the mean and CVaR estimations from their true values are constrained.

The parametric PBR method solves the empirical Markowitz problem instead of the empirical mean-CVaR problem if the underlying log-return distribution is in the elliptical family (which includes Gaussian and t distributions). This is based on the observation that CVaR of a portfolio is a weighted sum of the portfolio mean and the portfolio variance if the log-return distribution is in the elliptical family, resulting in the equivalence of the population efficient frontiers[1] of the Markowitz and mean-CVaR problems. As we are striving to reach the population frontier with greater stability, it makes intuitive sense to use the empirical Markowitz solution in lieu of the empirical mean-CVaR solution for this model.

The PBR methods are anticipated to enhance the performance by yielding solutions that are, on average, closer to achieving the original objective (minimize the *true* CVaR subject to *true* return equal to some level). As such, the PBR approach is fundamentally different from robust optimization, in that robust optimization deals

[1]By "population" we mean having a perfect market knowledge.

with the source of uncertainty to minimize the worst-case performance, whereas PBR deals with the performance uncertainty to increase the average performance. Comparing to the statistics/machine learning literature, PBR for the nonparametric case can be seen as an extension of standard regularization, in that nonparametric PBR also constrains the decision variable, however does so indirectly through penalizing the variability of mean and CVaR estimations.

Details of the nonparametric PBR method can be found in Sec. 2.3.1 and the parametric PBR method in Sec. 2.3.2. In Sec. 2.4, we provide theoretical results for the PBR methods after deriving the Central Limit Theorem for the nonparametric PBR solution. In Sec. 2.5, we evaluate the PBR methods against the straight-forward approach of solving the empirical mean-CVaR problem for three different log-return models via simulation experiments. We find that on average, the sample efficient frontiers of the PBR solutions are closer to the population efficient frontier than those of the straight-forward approach.

2.2 Mean-CVaR portfolio optimization

Notations. Throughout the chapter, we denote convergence in probability by $\xrightarrow{P}$ and in distribution by $\Rightarrow$. The notation $X \overset{d}{=} Y$ for two random variables X and Y means they have the same distribution, and the symbol $X \sim \mathcal{D}$ is used to indicate that the random variable X follows some standard distribution $\mathcal{D}$.

2.2.1 Setup

An investor is to choose a portfolio $w \in \mathbb{R}^p$ on p different assets. Her wealth is normalized to 1, so $w^\top 1_p = 1$, where 1_p denotes $p \times 1$ vector of ones. The log-returns of the p assets is denoted by X, a $p \times 1$ random vector, which follows some absolutely continuous distribution F with twice continuously differentiable pdf and finite mean μ and covariance Σ. The investor wants to pick a portfolio that minimizes the CVaR of the portfolio loss at level $100(1-\beta)\%$, for some $\beta \in (0.5, 1)$, while reaching an expected return R. That is, she wants to solve the following problem:

$$\begin{aligned} w_0 = \underset{w}{\operatorname{argmin}} \quad & CVaR(-w^\top X; \beta) \\ s.t. \quad & w^\top \mu = R \\ & w^\top 1_p = 1, \end{aligned} \tag{CVaR-pop}$$

where

$$CVaR(-w^\top X; \beta) := \min_{\alpha} \ \alpha + \frac{1}{1-\beta}\mathbb{E}(-w^\top X - \alpha)^+, \tag{2.1}$$

as in [55].

In reality, the investor does not know the distribution F. We assume the investor observes n iid realizations of asset returns, $\mathsf{X} = [X_1, \ldots, X_n] \in \mathbb{R}^{p \times n}$. Then the most straight-forward thing is to solve the following problem, where plugged-in estimators replace the true CVaR and return values:

$$\begin{aligned} \hat{w}_n = \underset{w}{\text{argmin}} \quad & \widehat{CVaR}_n(-w^\top \mathsf{X}; \beta) \\ s.t. \quad & w^\top \hat{\mu}_n = R \\ & w^\top 1_p = 1 \end{aligned} \tag{CVaR-emp}$$

where

$$\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta) := \min_{\alpha \in \mathbb{R}} \ \alpha + \frac{1}{n(1-\beta)} \sum_{i=1}^{n} (-w^\top X_i - \alpha)^+, \tag{2.2}$$

is a sample average estimator for $CVaR(-w^\top X; \beta)$ and $\hat{\mu}_n = n^{-1} \sum_{i=1}^{n} X_i$ is the sample mean of the observed asset log-returns.

2.2.2 Estimation risk of the empirical solution

Asymptotically, as the number of observations n goes to infinity (with p constant), $\hat{w}_n$ converges in probability to w_0 [see Sec. 2.4.2 for details]. In practice, however, the investor has a limited number of relevant observations. If, for example, there are $n = 250$ iid daily observations, and the investor wishes to control the top 5% of the losses, then there are only $250 \times 0.05 = 12.5$ points to estimate the portfolio CVaR at level $\beta = 0.95$. For stock log-returns, $n = 250$ iid daily observations is rather optimistic; there is ample empirical evidence that suggests daily log-returns are non-stationary over this period of time [50]. Even for time scales with more evidence for stationarity (e.g. bi-weekly/montly), the stationarity tends to last for no more than 5 years [50].

As a result, solving (CVaR-emp) using real data results in highly unreliable solutions. Let us illustrate this point, assuming an ideal market scenario. There are $p = 10$ stocks, with daily returns following a Gaussian distribution[2]: $X \sim \mathcal{N}(\mu_{sim}, \Sigma_{sim})$, and the investor has $n = 250$ iid observations of X. In the following, we conduct an experiment similar to those found in [42], to evaluate the performance and reliability of solving (CVaR-emp) under this ideal scenario. Briefly, the experimental procedure is as follows:

- Simulate 250 historical observations from $\mathcal{N}(\mu_{sim}, \Sigma_{sim})$.
- Solve (CVaR-emp) with $\beta = 0.95$ and some return level R to find an instance of $\hat{w}_n$.

[2]the parameters are the sample mean and covariance matrix of data from 500 daily returns of 10 different US stocks from Jan 2009– Jan 2011

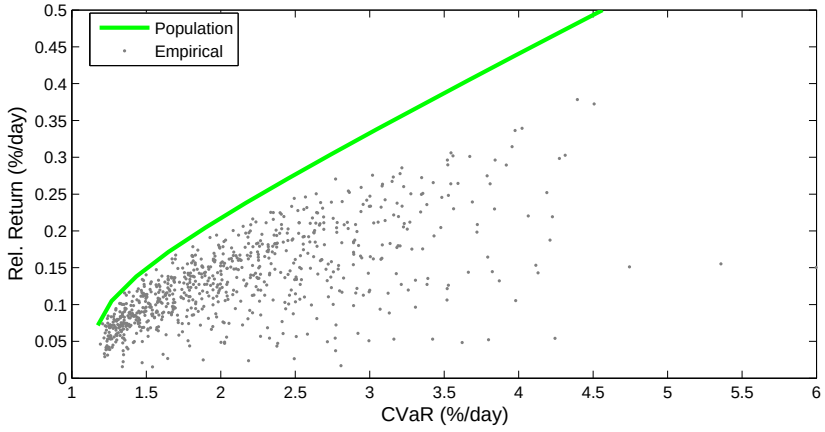

Figure 2.1 Distribution of realized daily return (%) vs. daily risk (%) of empirical solution $\hat{w}$. Green line represent the population frontier, i.e. the efficient frontier corresponding to solving (CVaR-pop).

- Plot the realized return $\hat{w}_n^\top \mu$ versus realized risk $CVaR(-\hat{w}_n^\top X; \beta)$; this corresponds to one grey point in Fig. (2.1).
- Repeat for different values of R to obtain a sample efficient frontier.
- Repeat many times to get a distribution of the sample efficient frontier.

The result of the experiment is summarized in Fig. (2.1). The green curve corresponds to the population efficient frontier. Each of the grey dots corresponds to a solution instance of (CVaR-emp). There are two noteworthy observations: the solutions $\hat{w}_n$ are sub-optimal, and they are highly variable. For instance, for a daily return of 0.1%, the CVaR ranges from 1.3% to 4%.

In the following section, we introduce *performance-based regularization* (PBR) as an approach to improve upon (CVaR-emp). The PBR approach is so-called because its goal is to improve upon $\hat{w}_n$ in terms of its performance, i.e. closeness to the population efficient frontier, ideally with less variability. We describe PBR for two cases: the investor has nonparametric or parametric knowledge of the market.

2.3 Performance-based regularization

2.3.1 Nonparametric case

In the nonparametric case, we assume the asset log-returns X follows some distribution P with finite mean μ and covariance Σ, and the investor has n iid observations: $\mathsf{X} = [X_1, \ldots, X_n] \in \mathbb{R}^{p \times n}$. The nonparametric PBR approach to (CVaR-pop) is to

solve the following problem:

$$\begin{array}{ll} \min\limits_{w} & \widehat{CVaR}_n(-w^\top \mathsf{X}; \beta) \\ s.t. & w^\top \hat{\mu}_n = R \\ & w^\top 1_p = 1 \\ & P_1(w) \leq U_1 \\ & P_2(w) \leq U_2 \end{array} \tag{2.3}$$

where P_1 and P_2 are penalty functions that characterize the uncertainty associated with $w^\top \hat{\mu}_n$ and $\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta)$ respectively. The idea is to penalize decisions w for which the uncertainty about the true values $w^\top \mu$ and $CVaR(-w^\top X; \beta)$ is large.

What, then, are appropriate penalty functions? Recall that we are trying to find solutions that yield efficient frontiers that are closer to the population efficient frontier, ideally with smaller variability. Thus the variances of $w^\top \hat{\mu}_n$ and $\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta)$ make appropriate penalty functions, as they characterize the deviation from the respective population values. The variance of $w^\top \hat{\mu}_n$ is given by

$$Var(w^\top \hat{\mu}_n) = \frac{1}{n^2} \sum_{i=1}^{n} Var(w^\top X_i) = \frac{1}{n} w^\top \Sigma w,$$

and the variance of $\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta)$ is approximately equal to $\gamma_0^2/n(1-\beta)^2 = Var[\max(-w^\top X - \alpha_\beta)]$, where

$$\alpha_\beta = \inf\{\alpha : P(-w^\top X \geq \alpha) \leq 1 - \beta\}$$

is the Value-at-Risk (VaR) of the portfolio w at level β, due to the following lemma.

Lemma 1. *Suppose* $\mathsf{X} = [X_1, \ldots, X_n] \overset{iid}{\sim} F$*, where* F *is absolutely continuous with twice continuously differentiable pdf. Then*

$$\frac{\sqrt{n}(1-\beta)}{\gamma_0} (\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta) - CVaR(-w^\top X; \beta)) \Rightarrow \mathcal{N}(0, 1). \tag{2.4}$$

Proof. See Appendix 2.6.1. □

Of course, we do not know the true variances, so we contend with *sample* variances of the estimators $w^\top \hat{\mu}_n$ and $\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta)$. That is, we consider the following penalty functions:

$$\begin{aligned} P_1(w) &= \frac{1}{n} w^\top \hat{\Sigma}_n w, \text{ where } \hat{\Sigma}_n = Cov(\mathsf{X}), \\ P_2(w) &= \frac{1}{n(1-\beta)^2} z^\top \Omega_n z, \text{ where} \\ \Omega_n &= \frac{1}{n-1}[I_n - n^{-1} 1_n 1_n^\top], \ I_n = n \times n \text{ identity matrix, and} \\ z_i &= \max(0, -w^\top X_i - \alpha) \text{ for } i = 1, \ldots, n. \end{aligned}$$

For the rest of this chapter, we investigate the nonparametric PBR method with sample variance penalty functions. Of course, this is just one particular choice, and it opens up the question of how different penalty functions affect the solution performance, and whether there are such things as "optimal" penalty functions. These are difficult questions worthy of further research, and we do not investigate them in this chapter. Nevertheless, we derive the asymptotic behavior of the solution of nonparametric PBR method in Sec. 2.4, which gives us some insight into i) how one could compare the effects of different penalty functions and ii) the first-order effect of many typical penalty functions.

The nonparametric PBR method with sample variance of return and CVaR estimators as penalties is:

$$
\begin{aligned}
(\hat{\alpha}_n^v, \hat{w}_n^v, \hat{z}_n^v) \quad = \quad & \underset{\alpha,w,z}{\operatorname{argmin}} \quad \alpha + \frac{1}{n(1-\beta)} \sum_{i=1}^{n} z_i \\
& s.t. \quad w^\top \hat{\mu}_n \quad = \quad R \\
& \qquad\quad w^\top 1_p \quad = \quad 1 \\
& \qquad \frac{1}{n} w^\top \hat{\Sigma}_n w \quad \leq \quad U_1 \\
& \qquad \frac{1}{n(1-\beta)^2} z^\top \Omega_n z \quad \leq \quad U_2 \\
& z_i = \max(0, -w^\top X_i - \alpha), \quad i = \quad 1, \ldots, n.
\end{aligned}
\tag{CVaR-pen}
$$

At first glance, (CVaR-pen) is a combinatorial optimization problem due to the cutoff variables z_i, $i = 1, \ldots, n$. However, it turns out that the convex relaxation of (CVaR-pen), a quadratically-constrained quadratic program (QCQP), is tight, thus we can solve (CVaR-pen) efficiently. Before stating the result, let us first introduce the convex relaxation of (CVaR-pen):

$$
\begin{aligned}
\min_{\alpha,w,z} \quad & \alpha + \frac{1}{n(1-\beta)} \sum_{i=1}^{n} z_i \\
& s.t. \quad w^\top \hat{\mu}_n \quad = R \qquad (\nu_1) \\
& \qquad\quad w^\top 1_p \quad = 1 \qquad (\nu_2) \\
& \qquad \frac{1}{n} w^\top \hat{\Sigma}_n w \quad \leq U_1 \qquad (\lambda_1) \\
& \qquad \frac{1}{n(1-\beta)^2} z^\top \Omega_n z \quad \leq U_2 \qquad (\lambda_2) \\
& \qquad\quad z_i \geq 0 \quad i = 1, \ldots, n \quad (\eta_1) \\
& z_i \geq -w^\top X_i - \alpha, \quad i = 1, \ldots, n \quad (\eta_2)
\end{aligned}
\tag{CVaR-relax}
$$

and its dual (where the dual variables correspond to the primal constraints as indi-

cated above):

$$\begin{array}{ll} \max\limits_{\nu_1,\nu_2,\lambda_1,\lambda_2,\eta_1,\eta_2} & g(\nu_1,\nu_2,\eta_1,\eta_2,\lambda_1,\lambda_2) \\ s.t. & \eta_2^\top 1_n = 1 \\ & \lambda_1 \geq 0, \lambda_2 \geq 0 \\ & \eta_1 \geq 0, \eta_2 \geq 0 \end{array} \qquad \text{(CVaR-relax-d)}$$

where

$$\begin{aligned} & g(\nu_1,\nu_2,\lambda_1,\lambda_2,\eta_1,\eta_2) \\ & = -\frac{n}{2\lambda_1}(\nu_1\hat{\mu}_n + \nu_2 1_p - \mathsf{X}\eta_2)^\top \hat{\Sigma}_n^{-1}(\nu_1\hat{\mu}_n + \nu_2 1_p - \mathsf{X}\eta_2) \\ & \quad - \frac{n(1-\beta)^2}{2\lambda_2}(\eta_1+\eta_2)^\top \Omega_n^\dagger(\eta_1+\eta_2) + R\nu_1 + \nu_2 - U_1\lambda_1 - U_2\lambda_2, \end{aligned}$$

and $\Omega_n^\dagger$ is the Moore-Penrose pseudo inverse of the singular matrix Ω_n.

We now show (CVaR-pen) can be solved efficiently by its convex relaxation:

Theorem 1. *Let $(\alpha^*, w^*, z^*, \lambda_1^*, \lambda_2^*, \eta_1^*, \eta_2^*)$ be the primal-dual optimal point of (CVaR-relax) and (CVaR-relax-d). If $\eta_2^* \neq 1_n/n$, then (α^*, w^*, z^*) is an optimal point of (CVaR-pen). Otherwise, if $\eta_2^* = 1_n/n$, we can find the optimal solution to (CVaR-relax) by solving (CVaR-relax-d) with an additional constraint $\eta_1^\top 1_n \geq \delta$, where δ is a constant $0 < \delta \ll 1$.*

Proof. See Appendix 2.6.2. □

Remark 1 – Bias introduced by penalty functions.
Note that if the penalties induce active constraints (i.e. U_1, U_2 are small enough), $\hat{w}_n^v$ does not converge to w_0 as $n \to \infty$, i.e. the penalty constraints introduce bias. This is not a problem, however, because we are concerned with finite sample performance, not asymptotic consistency. In Sec. 2.5, we see that the bias introduced by the penalized solution is actually in the direction that improves performance in the return-risk space.

Remark 2 – Interpretation as chance-programming.
Both $w^\top\hat{\mu}_n$ and $\widehat{CVaR}_n(-w^\top\mathsf{X};\beta)$ are asymptotically normally distributed, so constraining their variances results in the reduction of the corresponding confidence intervals at some fixed level ϵ. Hence penalizing their variances can be interpreted as chance-programming [16]. Analytically, the chance constraint on $|w^\top\hat{\mu}_n - w^\top\mu|$ can be transformed to a penalty constraint in the following manner:

$$\begin{aligned} P\left(|w^\top\hat{\mu}_n - w^\top\mu| \leq t\right) &\geq 1-\epsilon \\ \approx 2\Phi\left(\frac{t}{\sqrt{w^\top\Sigma w/n}}\right) - 1 &\geq 1-\epsilon \text{ for large } n \\ \iff \qquad \frac{1}{n}w^\top\Sigma w &\leq \left(\frac{t}{\Phi^{-1}(1-\epsilon/2)}\right)^2. \end{aligned}$$

That is, for a fixed level ϵ, there is a one-to-one mapping between the parameter U_1 of the penalty constraint $w^\top \hat{\Sigma}_n w/n \leq U_1$ and the parameter t of the chance constraint. The (asymptotic) variance penalty on $\widehat{CVaR}_n(-w^\top \mathsf{X}; \beta)$ has a similar interpretation as a chance constraint.

However, the penalty method can be interpreted as chance programming only if we choose the variance of the respective estimators as the penalty functions. Although we focus on the sample variance penalty function in this chapter, we assert that the penalty method need not be restricted to this particular choice.

2.3.2 Parametric case

In the parametric case, we assume the asset log-returns follow an elliptical distribution; i.e. the level sets of the distribution density function form ellipsoids. An elliptical distribution has a stochastic representation as follows (see [3] or [52]):

$$X \stackrel{d}{=} \mu + Y\Sigma^{1/2}U \tag{2.5}$$

where μ is the mean vector, U is a $p \times 1$ random vector uniformly distributed on the p-dimensional sphere of radius 1 (i.e. $U \stackrel{d}{=} Z_p/||Z_p||_2$, $Z_p \sim \mathcal{N}(0, I_p)$), and Y is a non-negative random variable independent of U. A special case is the Gaussian model: choosing $Y = \chi_p$, we get $X \sim \mathcal{N}(\mu, \Sigma)$. The elliptical family of distributions can thus be thought of as a generalization of the Gaussian family, and may be more reasonable for financial modeling because the non-random mixing of covariances can capture non-trivial tail dependence and heavier tails [50]. In particular, t-distributions also belong in the elliptical family.

The parametric PBR method is to solve the empirical Markowitz problem instead of (CVaR-emp) if X belongs in the elliptical family:

$$\begin{aligned} \hat{w}_n^M = \underset{w}{\text{argmin}} \quad & w^\top \hat{\Sigma}_n w \\ s.t. \quad & w^\top \hat{\mu}_n = R \\ & w^\top 1_p = 1. \end{aligned} \tag{Mark-emp}$$

The method is based on Lemma 2, which shows that the solutions of (CVaR-pop) and the population Markowitz problem [which is the same as (Mark-emp) except with (Σ, μ) replacing $(\hat{\Sigma}_n, \hat{\mu}_n)$] are equivalent if X is elliptically distributed. Lemma 2 is an extension of results mentioned elsewhere ([55], [20]) that show the equivalence of the solutions of (CVaR-pop) and the population Markowitz problem when X is Gaussian. However, to our knowledge, the implication that we can solve (Mark-emp) in lieu of (CVaR-emp) to obtain a better-performing solution has not been asserted.

Lemma 2. *Suppose $X \sim Ellip(\mu, \Sigma, Y)$ as in (2.5) and $Y > 0$. Then the solution of the population mean-CVaR problem (CVaR-pop) and the population Markowitz problem are equivalent.*

Proof. The proof is straightforward: we show $CVaR(-w^\top X;\beta)$ is a weighted sum of the portfolio mean $w^\top\mu$ and portfolio std $\sqrt{w^\top\Sigma w}$.

First, the portfolio loss is:

$$L(w) := -w^\top X \stackrel{d}{=} -w^\top\mu + Yv^\top U\sqrt{w^\top\Sigma w},$$

where $v^\top = w^\top\Sigma^{1/2}/\sqrt{w^\top\Sigma w}$, with $||v||_2 = 1$. Before we compute $CVaR(-w^\top X;\beta) = CVaR(L(w);\beta)$, we need to compute α_β, the VaR of $L(w)$ at level β [equivalently, the $(1-\beta)$-quantile of $L(w)$]. Since $L(w)$ is a continuous random variable, $\alpha_\beta = F^{-1}_{L(w)}(1-\beta)$, where $F^{-1}_{L(w)}$ is the inverse cdf of $L(w)$. Now

$$F_{L(w)}(x) = P(L(w)\le x) = P\left(Yv^\top U \ge \frac{-x-w^\top\mu}{\sqrt{w^\top\Sigma w}}\right),$$

so to compute α_β, we need the distribution of $Yv^\top U$. Since v has norm 1, $v^\top Z_p \stackrel{d}{=} Z_1$, where $Z_1\sim\mathcal{N}(0,1)$, and since $U\stackrel{d}{=} Z_p/||Z_p||_2$,

$$v^\top U \stackrel{d}{=} \frac{Z_1}{\sqrt{Z_1^2+\chi^2_{p-1}}},$$

where χ^2_{p-1} is independent of Z_1. Thus $(v^\top U)^2 \sim Beta(1/2,(p-1)/2)$, and by the symmetry of the normal, we have

$$P(Yv^\top U\ge x) = P(YI(1/2)\sqrt{B}\ge x),$$

where $B\sim Beta(1/2,(p-1)/2)$ and $I(1/2)\sim Bernoulli(1/2)$, independent of the rest. This quantity clearly does not depend on our choice of w, hence the solution to the equation

$$F_{L(w)}(x) = 1-\beta$$

is given by

$$\alpha_\beta = -w^\top\mu + q(1-\beta;YI(1/2)\sqrt{B})\sqrt{w^\top\Sigma w},$$

where q is a function that does not depend on w, and is unique since $L(w)$ is a continuous random variable.

Thus CVaR at level β is given by

$$\begin{aligned} CVaR(L(w);\beta) &= \frac{1}{1-\beta}\mathbb{E}[L(w)I(L(w)\ge\alpha_\beta)] \\ &= -w^\top\mu + G(1-\beta;YI(1/2)\sqrt{B})\sqrt{w^\top\Sigma w}, \end{aligned} \tag{2.6}$$

where G does not depend on w. Hence minimizing $CVaR(L(w);\beta)$ subject to $w^\top\mu = R$ and $w^\top 1_p = 1$ is equivalent to minimizing $w^\top\Sigma w$ subject to the same constraints, which is precisely the population Markowitz problem. □

2.4 Theory

We have thus far introduced nonparametric and parametric PBR methods to improve upon the empirical mean-CVaR problem (CVaR-emp). While we evaluate these methods in Sec. 2.5 via simulation experiments, it is still desirable to obtain some theoretical understanding of $\hat{w}_n$, $\hat{w}_n^v$ and $\hat{w}_n^M$.

The solution to the empirical Markowitz problem $\hat{w}_n^M$ has an explicit form and its asymptotic behavior has been studied elsewhere (for $X \sim \mathcal{N}(\mu, \Sigma)$, see [39], and for $X \sim \textit{Elliptical}$, see [23]). So we focus on deriving the asymptotic behavior of $\hat{w}_n$ and $\hat{w}_n^v$ — specifically, we show that they follow the Central Limit Theorem (CLT). Application of the delta method from classical statistics (see for example, Chapter 3 of [61]) then allows us to conclude that the corresponding sample efficient frontiers also follow the CLT. From these results, we can get some insight into the effect of the penalty functions in the nonparametric PBR method, and (indirectly) justify the parametric PBR method when the log-returns are Gaussian.

Notations. In this section, we make use of stochastic little-o and big-O notations: for a given sequence of random variables R_n, $X_n = o_P(R_n)$ means $X_n = Y_n R_n$ where $Y_n \xrightarrow{P} 0$, and $X_n = O_P(R_n)$ means $X_n = Y_n R_n$ where $Y_n = O_P(1)$, i.e. for every $\varepsilon > 0$ there exists a constant M such that $\sup_n P(|Y_n| > M) < \varepsilon$.

Measurability Issues. We also encounter quantities that may not be measurable (e.g. supremum over uncountable families of measurable functions). We note that whenever the "probability" of such quantities are written down, we actually mean the outer probability. For further details, see Appendix C of [54].

2.4.1 Preliminaries

The quantities $\hat{w}_n$ and $\hat{w}_n^v$ are solutions to non-trivial optimization problems so they cannot be written down analytically, and it seems characterizing their asymptotic distributions would be difficult. However, we are not at a complete loss. In statistics, an M-estimator[3] is an estimator that minimizes an empirical function of the type

$$\theta \mapsto M_n(\theta) := \frac{1}{n} \sum_{i=1}^{n} m_\theta(X_i), \tag{2.7}$$

where $X_1, \ldots, X_n$ are iid observations, over some parameter space Θ. The solution $\hat{\theta}_n$ is then a reasonable estimator of the minimizer θ_0 of the true mean $M(\theta) = \mathbb{E}[m_\theta(X_1)]$. It is well-known that $\hat{\theta}_n$ obeys the Central Limit Theorem (i.e. is asymptotically normally distributed) under some regularity conditions. Intuitively, assuming θ is

[3] "M" stands for Minimization (or Maximization). For readers unfamiliar with M-estimation, maximum likelihood estimation falls in this category.

one-dimensional and M_n is sufficiently smooth, the CLT result is based on Taylor expansion of the first-order condition $dM_n(\hat{\theta}_n)/d\theta = 0$ about θ_0:

$$0 \;=\; \frac{dM_n(\hat{\theta}_n)}{d\theta} = \frac{dM_n(\theta_0)}{d\theta} + (\hat{\theta}_n - \theta_0)\frac{d^2M_n(\theta_0)}{d\theta^2} + O_P(|\hat{\theta}_n - \theta_0|^2).$$

Under reasonable assumptions that $d^2M_n(\theta_0)/d\theta^2$ obeys the Weak Law of Large Numbers and $\hat{\theta}_n$ is a consistent estimator of θ_0 (i.e. $|\hat{\theta}_n - \theta_0| \xrightarrow{P} 0$), we have

$$\sqrt{n}(\hat{\theta}_n - \theta_0) = -\frac{1}{\mathbb{E}[\frac{d^2M_n(\theta_0)}{d\theta^2}]}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\frac{dm_{\theta_0}(X_i)}{d\theta} + o_P(1),$$

with the latter expression obeying the standard CLT as it is a normalized sum of iid random variables.

So we ask, can we transform (CVaR-emp) and (CVaR-pen) to a problem for which we can use the M-estimation results?

The first step towards transforming (CVaR-emp) and (CVaR-pen) is to make them into constraint-free optimization problems. This is achievable, albeit with some thoughts, and we defer the details to Sec. 2.4.2. Next, we need to show $\hat{w}_n$ and $\hat{w}_n^v$ are consistent, i.e. they converge in probability to the corresponding population solutions. The proof of consistency is also provided in Sec. 2.4.2.

Once (CVaR-emp) is transformed to a global optimization problem, it is equivalent to an M-estimation problem in that the objective is a sample average of iid random variables of the form Eq. (2.7). Thus we conclude $\hat{w}_n$ is asymptotically normally distributed with mean w_0 and covariance matrix Σ_{w_0}, which we can compute.

However, (CVaR-pen) after transformation into a global problem is not quite an M-estimation problem, because, after some algebra, the objective is of the form (see Sec. 2.4.2 for details):

$$\theta \mapsto M_n(\theta) = \frac{1}{n(n-1)}\sum_{i\neq j} m_\theta^U(X_i, X_j), \tag{2.8}$$

where $m^U(\cdot,\cdot)$ is a permutation-symmetric function, and the sum is over all possible pairs (i,j) for $1 \le i,j \le n$, resulting in a sample average of identically distributed but non-independent terms.

For fixed θ, statistics of the form Eq. (2.8) are known as U-statistics, and we believe the solution $\hat{w}_n^v$ is still well-behaved because U-statistics can be decomposed into a term of the form $M_n^1(\theta) = \sum_{i=1}^n m_\theta^1(X_i)$ (known as its Hajék projection or first term in its Hoeffding decomposition; see [37]) and a remainder which converges to zero in probability at rate $\sqrt{n}$. Thus we intuit that the asymptotic behavior of $\hat{w}_n^v$ is equivalent to the minimizer of $M_n^1(\theta)$, the latter for which we can apply the standard M-estimation result. We make this intuition rigorous in Sec. 2.4.3. In Sec. 2.4.4, we provide details of the asymptotic distributions of $\hat{w}_n$ and $\hat{w}_n^v$ when $X \sim \mathcal{N}(\mu, \Sigma)$, and provide a justification of the parametric PBR method.

2.4.2 Consistency of $\hat{w}_n$ and $\hat{w}_n^v$

In this subsection we show consistency of $\hat{w}_n^v = \hat{w}_n^v(\lambda_1, \lambda_2)$. The result goes through for $\hat{w}_n$ by setting $\lambda_1 = \lambda_2 = 0$.

Transformation into global optimization

The penalized CVaR portfolio optimization problem with dualized mean and sample/asymptotic variance penalty constraints is

$$\begin{array}{ll} \min\limits_{(\alpha,w)\in\mathbb{R}\times\mathbb{R}^p} & M_n(\alpha, w; \lambda_1, \lambda_2) \\ s.t. & w^\top 1_p = 1, \end{array} \tag{CVaR-dual}$$

where

$$M_n(\theta; \lambda_1, \lambda_2) = \frac{1}{n}\sum_{i=1}^n m_\theta(X_i) + \frac{\lambda_1}{n} w^\top \hat{\Sigma}_n w + \frac{\lambda_2}{n-1}\sum_{i=1}^n \left(z_\theta(X_i) - \frac{1}{n}\sum_{j=1}^n z_\theta(X_j) \right)^2, \tag{2.9}$$

$$m_\theta(x) = \alpha + \frac{1}{1-\beta} z_\theta(x) - \lambda_0 w^\top x, \tag{2.10}$$

and $\lambda_0 > 0$, $\lambda_1, \lambda_2 \geq 0$ are pre-determined constants.

We dualize the mean constraint $w'\hat{\mu}_n = R$ because it makes the analysis of the corresponding solution much easier. While dualizing the mean constraint adds a sample average of iid terms to the objective, leaving it as a constraint results in a solution that has a non-trivial dependence on the underlying randomness.

Now eliminating the non-random constraint $w^\top 1_p = 1$ is straight-forward; one possible way is to re-parameterize w as $w = w_1 + Lv$, where $L = [0_{(p-1)\times 1}, I_{(p-1)\times(p-1)}]^\top$, $v = [w_2, \ldots, w_p]^\top$ and $w_1 = [1 - v^\top 1_{(p-1)}, 0_{1\times(p-1)}]^\top$. The transformed problem is thus

$$\min_{\theta\in\mathbb{R}^p} M_n(\theta; \lambda_1, \lambda_2), \tag{2.11}$$

where $\theta = (\alpha, v) \in \mathbb{R}\times\mathbb{R}^{p-1}$ is free of constraints, and the corresponding population problem is

$$\min_{\theta\in\mathbb{R}^p} M(\theta; \lambda_1, \lambda_2) = \mathbb{E}[M_n(\theta; \lambda_1, \lambda_2)]. \tag{2.12}$$

In what follows, we assume $M(\theta; \lambda_1, \lambda_2)$ has a unique minimizer $\theta_0(\lambda_1, \lambda_2)$. We also let $\hat{\theta}_n(\lambda_1, \lambda_2)$ be a near-minimizer of $M_n(\theta; \lambda_1, \lambda_2)$, i.e.

$$M_n(\hat{\theta}_n; \lambda_1, \lambda_2) < \inf_{\theta\in\mathbb{R}^p} M_n(\theta; \lambda_1, \lambda_2) + o_P(1). \tag{2.13}$$

Transformation of the objective to a U-statistic

Let $\theta = (\alpha, v) \in \mathbb{R} \times \mathbb{R}^{p-1}$ and $z_\theta(x) := (-x^\top(w_1 + Lv) - \alpha)^+$. With simple algebra, we can re-write the objective Eq. (2.9) as a U-statistic:

$$M_n(\theta; \lambda_1, \lambda_2) = \frac{1}{\binom{n}{2}} \sum_{\substack{1 \le i,j \le n \\ i \ne j}} m^U_{(\theta;\lambda_1,\lambda_2)}(X_i, X_j), \tag{2.14}$$

where

$$\begin{aligned} m^U_{(\theta;\lambda_1,\lambda_2)}(x_i, x_j) \;\; := \;\; & \frac{1}{2}\left[m_\theta(x_i) + m_\theta(x_j)\right] + \frac{\lambda_1}{2}[(w_1 + Lv)^\top (x_i - x_j)]^2 \\ & + \frac{\lambda_2}{2}(z_\theta(x_i) - z_\theta(x_j))^2. \end{aligned} \tag{2.15}$$

Consistency of $\hat{\theta}_n(\lambda_1, \lambda_2)$

Let us now prove consistency of $\hat{\theta}_n(\lambda_1, \lambda_2)$ for fixed $\lambda_1, \lambda_2 \ge 0$. The intuition behind the proof is as follows: if $M(\theta; \lambda_1, \lambda_2)$ is well-behaved such that for every $\varepsilon > 0$ there exists $\eta > 0$ such that $||\hat{\theta}_n(\lambda_1, \lambda_2) - \theta_0(\lambda_1, \lambda_2)||_2 > \varepsilon \implies M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) > \eta$, then consistency follows from showing that the probability of the event $\{M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) > \eta\}$ goes to zero for all $\varepsilon > 0$. In the proof, we show that $0 \le M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) \le -(M_n(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\hat{\theta}_n; \lambda_1, \lambda_2)) + o_P(1)$, hence the result follows by proving Uniform Law of Large Numbers (ULLN) for $M_n(\theta; \lambda_1, \lambda_2)$:

$$\sup_{\theta \in \mathbb{R}^p} |M_n(\theta; \lambda_1, \lambda_2) - M(\theta; \lambda_1, \lambda_2)| \xrightarrow{P} 0. \tag{2.16}$$

ULLN has been extensively studied in the statistics and empirical processes literature and one of the standard approaches to showing ULLN is through bracketing numbers. Given two functions l, u, the bracket $[l, u]$ is the set of all functions g with $l \le g \le u$. An ε-bracket in $L_r(P)$ is a bracket $[l, u]$ with $\mathbb{E}_P(u - l)^r < \varepsilon^r$, and the bracketing number $N_{[\,]}(\varepsilon, \mathcal{F}, L_r(P))$ is the minimum number of ε-brackets needed to cover $\mathcal{F}$. Having a finite bracketing number $N_{[\,]}(\varepsilon, \mathcal{F}, L_r(P)) < \infty$ for every $\varepsilon > 0$ means one can find a finite approximation to $\mathcal{F}$ with ε-accuracy for all $\varepsilon > 0$, and ULLN holds for such $\mathcal{F}$ (Theorem 19.4 [61]).

There are certainly known sufficient conditions for finite bracketing numbers. For our problem, if we can replace $\mathbb{R}^p$ with a compact set, we can show F is a Lipschitz class of functions (defined in the next paragraph), which is known to have finite $N_{[\,]}(\varepsilon, \mathcal{F}, L_r(P))$ for every $\varepsilon > 0$. Now for all practical purposes, we need only consider a compact subset of Θ, $[-K, K]^p$ where K is appropriately large enough, because the elements of $\theta = (\alpha, v)$ are only meaningful if bounded in size (α is the Value-at-Risk of the portfolio $w = w_1 + Lv$). Hence for the rest of this section we assume a K exists such that $\hat{\theta}_n \in [-K, K]^p$ for all n and $\theta_0 \in [-K, K]^p$.

Definition 1 (**Lipschitz class**). *Consider a class of measurable functions $\mathcal{F} = \{f_\theta : \theta \in \Theta\}$, $f_\theta : \mathcal{X} \to \mathbb{R}$, under some probability measure P. We say $\mathcal{F}$ is a* Lipschitz class *about $\theta_0 \in \Theta$ if $\theta \mapsto f_\theta(x)$ is differentiable at θ_0 for P-almost every x with derivative $\dot{f}_{\theta_0}(x)$ and such that, for every θ_1 and θ_2 in a neighborhood of θ_0, there exists a measurable function $\dot{f}$ with $\mathbb{E}[\dot{f}^2(X_1)] < \infty$ such that*

$$|f_{\theta_1}(x) - f_{\theta_2}(x)| \leq \dot{f}(x)||\theta_1 - \theta_2||_2.$$

Example 19.7 of [61] shows that if $\mathcal{F} = \{f_\theta : \theta \in \Theta\}$ is a class of measurable functions with bounded $\Theta \subset \mathbb{R}^d$ and $\mathcal{F}$ is Lipschitz about $\theta_0 \in \Theta$ then for every $0 < \varepsilon < diam(\Theta)$, there exists C such that

$$N_{[\,]}(\varepsilon\sqrt{\mathbb{E}(|\dot{f}(X)|^2)}, \mathcal{F}, L_2(P)) \leq C\left(\frac{diam(\Theta)}{\varepsilon}\right)^d, \tag{2.17}$$

i.e. has a finite bracketing number for all $\varepsilon > 0$. This result is needed in proving consistency in the following.

Theorem 2. *For fixed $\lambda_1, \lambda_2 \geq 0$, let $\hat{\theta}_n(\lambda_1, \lambda_2)$ be a near-minimizer of $M_n(\theta; \lambda_1, \lambda_2)$ as in Eq. (2.13), and let $\theta_0(\lambda_1, \lambda_2)$ be the unique minimizer of $M(\theta; \lambda_1, \lambda_2)$. Also let*

$$\mathcal{F}_1 = \{m_\theta : \theta \in [-K, K]^p\}, \quad \mathcal{F}_2 = \{m^U_{(\theta;\lambda_1,\lambda_2)} : \theta \in [-K, K]^p\},$$

where m_θ and $m^U_{(\theta;\lambda_1,\lambda_2)}$ are defined in Eqs. (2.10) and (2.15). Suppose the following:

Assumption 1. $\theta \mapsto M(\theta; \lambda_1, \lambda_2)$ is continuous and $\liminf_{|\theta|\to\pm\infty} M(\theta; \lambda_1, \lambda_2) > M(\theta_0; \lambda_1, \lambda_2)$.

Assumption 2. $X_1, \ldots, X_n$ are iid continuous random vectors with finite fourth moment.

Then

$$||\hat{\theta}_n(\lambda_1, \lambda_2) - \theta_0(\lambda_1, \lambda_2)||_2 \xrightarrow{P} 0.$$

Proof. See Appendix 2.4.5. □

2.4.3 Central Limit Theorem for $\hat{\theta}_n(\lambda_1, \lambda_2)$

We are now ready to show the CLT for $\hat{\theta}_n(\lambda_1, \lambda_2)$. The CLT for $\hat{\theta}_n(0, 0)$ is a straightforward application of known M-estimation results for Lipschitz class of objective functions (e.g. Theorem 5.23 of [61]).

The CLT for $\hat{\theta}_n(\lambda_1, \lambda_2)$ when λ_1, λ_2 are not both zero does not follow straightforwardly from M-estimation results because $M_n(\theta; \lambda_1, \lambda_2)$ is a sample average of identically distributed but non-independent terms. However, statistics of the form $M_n(\theta; \lambda_1, \lambda_2)$ are known as U-statistics, and we can decompose them into a sum of iid random variables and a component which is $o_P(1/\sqrt{n})$ [37]:

$$M_n(\theta; \lambda_1, \lambda_2) = \frac{1}{n}\sum_{i=1}^n m^1_{(\theta;\lambda_1,\lambda_2)}(X_i) + E_n(\theta; \lambda_1, \lambda_2), \tag{2.18}$$

where

$$m^1_{(\theta;\lambda_1,\lambda_2)}(X_i) = 2\mathbb{E}_{X_j}[m^U_{(\theta;\lambda_1,\lambda_2)}(X_i, X_j)] - \mathbb{E}_{X_1,X_2}[m^U_{(\theta;\lambda_1,\lambda_2)}(X_1, X_2)]$$

and $E_n(\theta;\lambda_1,\lambda_2)$ is $o_P(1/\sqrt{n})$. Hence we suspect $|R^U_n(\hat{\theta}_n;\lambda_1,\lambda_2)| \xrightarrow{P} 0$, where

$$R^U_n(\theta;\lambda_1,\lambda_2) = \sqrt{n}(\theta-\theta_0) - [\nabla^2_{\theta_0}\mathbb{E}m^1_{(\theta;\lambda_1,\lambda_2)}(X_i)]^{-1}\frac{1}{\sqrt{n}}\sum_{i=1}^n m^1_{(\theta;\lambda_1,\lambda_2)}(X_i).$$

Now $\hat{\theta}_n$ changes with every n so we need uniform probabilistic convergence of $R^U_n(\theta;\lambda_1,\lambda_2)$, and implicitly of $E_n(\theta;\lambda_1,\lambda_2)$. For this we need to show convergence of particular stochastic processes; an empirical process and a U-process.

Definition 2. *Let $X_1,\ldots,X_n$ be iid random vectors from $\mathcal{X}$. For a measurable function $f:\mathcal{X}\to\mathbb{R}$, the* empirical process at f *is*

$$\mathbb{G}_n f := \frac{1}{\sqrt{n}}\sum_{i=1}^n [f(X_i) - \mathbb{E}f(X_1)],$$

and for a measurable function $g:\mathcal{X}\times\mathcal{X}\to\mathbb{R}$, the U-process at g *is*

$$\mathbb{U}_n g := \frac{\sqrt{n}}{\binom{n}{2}}\sum_{i\neq j}[g(X_i,X_j) - \mathbb{E}_{X_1,X_2}g(X_1,X_2)].$$

To show convergence of quantities such as $\sup_{t\in T}|X_n(t)|$ for some stochastic process $\{X_n(t): t\in T\}$, we need to introduce the notion of weak convergence of stochastic processes. If $X_n(\cdot,\omega)$ is a bounded function for every $\omega\in\Omega$, then we can consider $X_n(\cdot,\omega)$ to be a point in the function space $\ell^\infty(T)$, the space of bounded functions on T which is equipped with the supremum norm. Hence, showing the convergence of $\sup_{t\in T}|X_n(t)|$ is equivalent to showing weak convergence of X_n in this function space.

Definition 3 (**Weak convergence of a stochastic process**). *A sequence of $X_n:\Omega_n\mapsto\ell^\infty(T)$ converges weakly to a tight random element* [4] *X iff both of the following conditions hold:*

1. *Finite approximation: the sequence $(X_n(t_1),\ldots,X_n(t_k))$ converges in distribution in $\mathbb{R}^k$ for every finite set of points $t_1,\ldots,t_k$ in T.*
2. *Maximal inequality: for every $\varepsilon,\eta>0$ there exists a partition of T into finitely many sets $T_1,\ldots,T_k$ such that*

$$\limsup_{n\to\infty} P\left[\sup_i \sup_{s,t\in T_i}|X_n(s)-X_n(t)|\geq\varepsilon\right]\leq\eta.$$

[4] A random element is a generalization of a random variable. Let $(\Omega,\mathcal{G},P)$ be a probability space and $\mathbb{D}$ a metric space. Then the $\mathcal{G}$-measurable map $X:\Omega\mapsto\mathbb{D}$ is called a random element.

The point at the end of this is, as taking the supremum is a continuous map in the topology of $\ell^\infty(T)$, weak convergence of $X_n(\cdot)$ to $X(\cdot)$ would allow us to conclude $\sup_{t\in T}|X_n(t)| \to \sup_{t\in T}|X(t)|$.

Regarding empirical processes, we say a class of measurable functions $\mathcal{F}$ is *P-Donsker* if $\{\mathbb{G}_n f : f \in \mathcal{F}\}$ converges weakly to a tight random element in $\ell^\infty(\mathcal{F})$. This property is related to the bracketing numbers introduced in Sec. 2.4.2: a class $\mathcal{F}$ is P-Donsker if $\varepsilon \log[N_{[\,]}(\varepsilon, \mathcal{F}, L_2(P))] \to 0$ as $\varepsilon \to 0$ (due to Donsker; see Theorem 19.5 of [61]). Many sufficient conditions for the weak convergence of $\{\mathbb{U}_n f : f \in \mathcal{F}\}$ are provided in [5], and we make use of one in our proof of CLT for $\hat{\theta}_n(\lambda_1, \lambda_2)$ below.

Theorem 3. *Fix $\lambda_1, \lambda_2 \geq 0$, λ_1, λ_2 not both zero and assume the same setting as Theorem 2. Also let*

$$\dot{m}^U_{(\theta_0;\lambda_1,\lambda_2)}(x) = \nabla_\theta m^U_{(\theta_0;\lambda_1,\lambda_2)}(x)|_{\theta=\theta_0(\lambda_1,\lambda_2)}, \text{ for } x \in \mathbb{R}^p,$$

and further assume

Assumption 3. $\mathbb{E}_{X_1,X_2}[m^U_{(\theta_0;\lambda_1,\lambda_2)}(X_1,X_2)^2] < \infty$.

Assumption 4. $\theta \mapsto M(\theta;\lambda_1,\lambda_2)$ *admits a second-order Taylor expansion at its point of minimum* $\theta_0(\lambda_1,\lambda_2)$ *with nonsingular symmetric second derivative matrix* $V_{\theta_0(\lambda_1,\lambda_2)}$.

Then

$$\sqrt{n}(\hat{\theta}_n(\lambda_1,\lambda_2) - \theta_0(\lambda_1,\lambda_2)) \;=\; -V^{-1}_{\theta_0(\lambda_1,\lambda_2)} \frac{1}{\sqrt{n}} \sum_{i=1}^n \dot{m}^1_{(\theta_0;\lambda_1,\lambda_2)}(X_i) + o_p(1)$$

where

$$\dot{m}^1_{(\theta;\lambda_1,\lambda_2)}(X_i) = 2\mathbb{E}_{X_2}[\dot{m}^U_{(\theta;\lambda_1,\lambda_2)}(X_1,X_2)] - \mathbb{E}_{X_1,X_2}[\dot{m}^U_{(\theta;;\lambda_1,\lambda_2)}(X_1,X_2)]$$

is the first-order term in the Hoeffding decomposition of $M_n(\theta;\lambda_1,\lambda_2)$.

Remark – Implication on the choice of penalty functions.
We have just shown that asymptotically, the sample variance penalty functions affect the solution performance only through its Hajék projection. This observation can generalize to many typical penalty functions (e.g. different statistics of mean and CVaR estimators), and as such, the implication is that of all possible penalty functions to consider, one may focus on a subclass of functions that can be expressed as a sample average of iid terms.

Corollary 1. *Assume the same setting as Theorem 3. Then*

$$\sqrt{n}(\hat{\theta}_n(\lambda_1,\lambda_2) - \theta_0(\lambda_1,\lambda_2)) \Rightarrow \mathcal{N}\left(0, \Sigma_{\theta_0}(\lambda_1,\lambda_2)\right), \tag{2.19}$$

where $\Sigma_{\theta_0}(\lambda_1,\lambda_2) = A^{-1}BA^{-1}$,

$$\begin{aligned} A = A_{\theta_0}(\lambda_1,\lambda_2) &= \nabla_\theta^2\ \mathbb{E}[m^1_{(\theta;\lambda_1,\lambda_2)}(X_1)]\Big|_{\theta=\theta_0} \\ &= \nabla_\theta^2\ [\frac{1}{1-\beta}\mathbb{E}z_\theta(X_1) + \lambda_1 w^\top \Sigma w + \lambda_2 Var(z_\theta(X_1))]\Big|_{\theta=\theta_0} \\ B = B_{\theta_0}(\lambda_1,\lambda_2) &= \mathbb{E}[\nabla_{\theta_0} m^1_{(\theta;\lambda_1,\lambda_2)}(X_1)\nabla_{\theta_0} m^1_{(\theta;\lambda_1,\lambda_2)}(X_1)^\top] \end{aligned}$$

where

$$\begin{aligned} &\nabla_\theta\ m^1_{(\theta;\lambda_1,\lambda_2)}(x) \\ &= \begin{bmatrix} 1-\frac{1}{1-\beta}\mathbb{I} \\ -\frac{1}{1-\beta}L^\top X\mathbb{I} - \lambda_0 L^\top X \end{bmatrix} \\ &+ \begin{bmatrix} 2\lambda_2\mathbb{E}[(z_\theta(X)-\mathbb{E}z_\theta(X))(-\mathbb{I}+\mathbb{E}\mathbb{I})] \\ 2\lambda_1 L^\top(X-\mu)(X-\mu)^\top w + 2\lambda_2\mathbb{E}[(z_\theta(X)-\mathbb{E}z_\theta(X))(-L^\top X\mathbb{I}+\mathbb{E}L^\top X\mathbb{I})] \end{bmatrix}, \end{aligned}$$

and $\mathbb{I} = \mathbb{I}(z_\theta(X)\geq 0)$.

Remarks.

1. For asymptotics of $\hat{w}_n(\lambda_1,\lambda_2)$ we have
$$\sqrt{n}(\hat{w}_n(\lambda_1,\lambda_2) - w_0(\lambda_1,\lambda_2)) \Rightarrow \mathcal{N}\left(0, \Sigma_{w_0}(\lambda_1,\lambda_2)\right), \tag{2.20}$$
where $\Sigma_{w_0}(\lambda_1,\lambda_2) = (0_p\ L)\Sigma_{\theta_0}(\lambda_1,\lambda_2)(0_p\ L)^\top$.

2. Setting $\lambda_1,\lambda_2 = 0$, we get back the unpenalized mean-CVaR problem.

3. **Asymptotic distribution of the efficient frontier.**
With Eq. (2.20), we can state the distribution of the true efficient frontier — that is, the distribution of $\hat{w}_n(\lambda_1,\lambda_2)^\top\mu$ and
$$g(\hat{w}_n(\lambda_1,\lambda_2)) := CVaR(-\hat{w}_n(\lambda_1,\lambda_2)^\top X_{n+1};\beta),$$
where $X_{n+1}\sim F$, independent of $X_1\ldots,X_n$. For the portfolio mean, we have
$$\sqrt{n}(\hat{w}_n(\lambda_1,\lambda_2)^\top\mu - w_0(\lambda_1,\lambda_2)^\top\mu) \Rightarrow \mathcal{N}(0,\mu^\top\Sigma_{w_0}((\lambda_1,\lambda_2))\mu)$$
and for the true CVaR, by the delta Method
$$\begin{aligned} &\sqrt{n}(g(\hat{w}_n(\lambda_1,\lambda_2)) - g(w_0(\lambda_1,\lambda_2))) \\ &\qquad \Rightarrow \mathcal{N}\left\{0, g'(w_0(\lambda_1,\lambda_2))^\top\Sigma_{w_0}(\lambda_1,\lambda_2)g'(w_0(\lambda_1,\lambda_2))\right\}. \end{aligned} \tag{2.21}$$
The asymptotic distribution of $g(\hat{w}_n(\lambda_1,\lambda_2))$ clearly depends on the distribution of the assets X. In the case when $X\sim\ Ellip(\mu,\Sigma,Y)$, $g(w) = -w^\top\mu + G\sqrt{w^\top\Sigma w}$ according to our previous calculations in Eq. (2.6). Hence
$$\sqrt{n}(g(\hat{w}_n) - g(w_0)) \Rightarrow \mathcal{N}\left\{0, \left(-\mu + G\frac{\Sigma w_0}{\sqrt{w_0\Sigma w_0}}\right)^\top \Sigma_{w_0}\left(-\mu + G\frac{\Sigma w_0}{\sqrt{w_0\Sigma w_0}}\right)\right\}. \tag{2.22}$$

2.4.4 Example. Asymptotic analysis when $X \sim \mathcal{N}(\mu, \Sigma)$

In the following, we provide the detailed computation of $\Sigma_{\theta_0}(0,0)$ for the unpenalized solution $\hat{\theta}_n(0,0)$ when $X \sim \mathcal{N}(\mu, \Sigma)$.

Lemma 3. *Suppose $X \sim \mathcal{N}(\mu, \Sigma)$. Then*

$$\begin{aligned} z_{\theta_0}(X) &= -w_0^\top X - \alpha_0 \sim \sigma_0 \mathcal{N}(-\Phi^{-1}(\beta), 1), \text{ and} \\ p_0 &= f_{-w_0^\top X}(0) = \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left\{ -\frac{1}{2\sigma_0^2} (\Phi^{-1}(\beta))^2 \right\}, \end{aligned}$$

where $\sigma_0 = \sqrt{w_0^\top \Sigma w_0}$. Then $\Sigma_{\theta_0}(0,0) = A_0^{-1} B_0 A_0^{-1}$, where A_0, B_0 are symmetric matrices with

$$\begin{aligned} A_0(1,1) &= \frac{p_0}{1-\beta} \\ A_0(j,l) &= \frac{p_0}{(1-\beta)} \mathbb{E}[L_j^\top X L_l^\top X | z_{\theta_0}(X) = 0] \quad \text{for } 2 \le j, l \le p \\ A_0(1,j) &= -\frac{p_0}{(1-\beta)} \mathbb{E}[L_j^\top X | z_{\theta_0}(X) = 0] \qquad \text{for } 2 \le j \le p, \end{aligned}$$

where L_j is the j-th column of L, and

$$\begin{aligned} B_0(1,1) &= \frac{\beta}{1-\beta} \\ B_0(j,l) &= \frac{1}{(1-\beta)} \left(\frac{1}{1-\beta} + 2\lambda_0 \right) \mathbb{E}[L_j^\top X L_l^\top X \mathbb{I}(z_{\theta_0}(X) \ge 0)] \\ &\quad + \lambda_0^2 (L_j^\top \Sigma L_l + L_j^\top \mu L_l^\top \mu) \quad \text{for } 2 \le j, l \le p \\ B_0(1,j) &= 0 \qquad \text{for } 2 \le j \le p. \end{aligned}$$

Proof. This is a straight-forward application of Corollary 1 for the case $X \sim \mathcal{N}(\mu, \Sigma)$. □

Let us now compare the asymptotic results derived above with simulations with finite number of observations. Consider 5 assets, a range of observations ($n = 250, 500, 1000, 2000$) and $X \sim \mathcal{N}(\mu_{sim}, \Sigma_{sim})$, where the model parameters are the same as the model parameters of the first five assets used in Sec. 2.2.2. For simulations, we solve the mean-CVaR problem with dualized mean constraint:

$$\begin{aligned} \min_{w} \quad & \widehat{CVaR}_n(-w^\top \mathsf{X}; \beta) - \lambda_0 w^\top \hat{\mu}_n \\ s.t. \quad & w^\top 1_p = 1, \end{aligned}$$

and follow steps similar to Sec. 2.2.2.

In Fig. 2.2, we summarize the empirical frontiers by plotting their averages and indicating $\pm 1/2$ standard deviation error bars, in both true mean (vertical) and true risk estimations (horizontal) in grey. The population frontier is also plotted, and

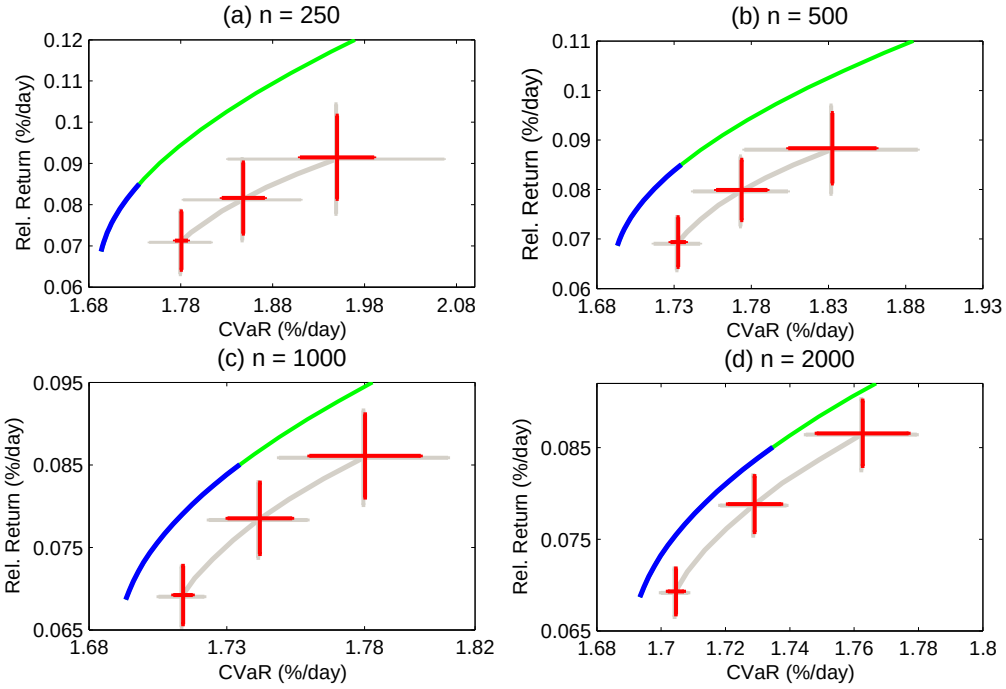

Figure 2.2 Comparison of theoretical (red) and simulated (grey) distributions of the empirical efficient frontier when $X \sim \mathcal{N}(\mu_{sim}, \Sigma_{sim})$ for increasing number of observations $n = [250, 500, 1000, 2000]$. The error bars indicate $\pm 1/2$ std variabilities in the mean and CVaR. Green is the population efficient frontier, and blue indicates the portion that corresponds to the return range considered for the simulations. Observe that the asymptotic variance calculated theoretically (red bars) approach the simulated variance (grey bars) with increasing n.

is shown in green, and the theoretical $\pm 1/2$ standard deviations of mean and risk estimations are juxtaposed with the empirical error bars in red. We make a couple of observations:

1. With increasing n, the theoretical error bars approach the simulated ones, as expected.

2. The theory seems to better predict the mean estimation error (vertical) better than the risk estimation error (horizontal). With finite n, the mean estimation error, which is computed using Eq. (2.21), depends only on one approximate quantity $\Sigma_{w_0}(0,0)$, whereas the risk estimation error, computed using Eq. (2.22), depends on $\Sigma_{w_0}(0,0)$ and w_0. Although $\hat{w}_n$ is a consistent estimator of w_0 asymptotically, with finite n the difference does play a role, as shown by the relative inaccuracy of the horizontal error bars compared to the vertical ones. The finite sample bias also explains the gap in the positions of the population and simulated efficient frontiers.

Let us now derive asymptotic properties of the penalized solution $\hat{\theta}_n(\lambda_1, \lambda_2)$, $\lambda_1, \lambda_2 \geq 0$, when $X \sim \mathcal{N}(\mu, \Sigma)$. First, we show that when $X \sim \mathcal{N}(\mu, \Sigma)$, penalizing variance of CVaR estimation is redundant if one penalizes the sample variance of the mean.

Lemma 4. *Suppose $X \sim \mathcal{N}(\mu, \Sigma)$ and let $z_\theta(X) = -\alpha - w^\top X$. Then $z_\theta(X) \sim \mathcal{N}(\mu_1, \sigma_1^2)$ where $\mu_1 = -\sigma_1 \Phi^{-1}(\beta)$, $\sigma_1^2 = w^\top \Sigma w$, and*

$$Var[\max(z_\theta(X), 0)] = C(\beta)\sigma_1^2,$$

where $C(\beta)$ is a constant that only depends on β. Thus penalizing the sample variance of CVaR via $P_2(w) = \widehat{Var}_n[z_{\theta(w)}(X), 0)] \leq U_2$ is redundant if one penalizes the sample variance of the mean via $P_1(w) = w^\top \hat{\Sigma}_n w = \hat{\sigma}_{1,n}^2 \leq U_1$.

Proof. Straight-forward calculations show

$$Var[\max(z_\theta(X), 0)] = \left\{([\Phi^{-1}(\beta)]^2 + 1)(1-\beta) - 3\Phi^{-1}(\beta) f_{Z_0}[\Phi^{-1}(\beta)]\right\} \sigma_1^2,$$

where f_{Z_0} is the pdf of the standard normal random variable Z_0. □

The implication now is that when $X \sim \mathcal{N}(\mu, \Sigma)$, we need only consider $\lambda_1 \geq 0, \lambda_2 = 0$ to characterize the asymptotic properties of the penalized solution, which we describe below.

Lemma 5. *Suppose $X \sim \mathcal{N}(\mu, \Sigma)$. Then*

$$\Sigma_{\theta_0}(\lambda_1, 0) = A_1^{-1} B_1 A_1^{-1},$$

where A_1, B_1 are symmetric matrices with

$$\begin{array}{lll} A_1(1,1) & = A_0(1,1) & \\ A_1(j,l) & = A_0(j,l) + \lambda_1 L_j^\top \Sigma L_l & \textit{for } 2 \le j,l \le p \\ A_1(1,j) & = A_0(1,j) & \textit{for } 2 \le j \le p \end{array}$$

where L_j is the j-th column of L, and

$$\begin{array}{lll} B_1(1,1) & = B_0(1,1) & \\ B_1(j,l) & = B_0(j,l) + \lambda_1 \mathbb{E}[b_{0,j}b_{1,l} + b_{0,l}b_{1,j} + \lambda_1 b_{1,j}b_{1,l}] & \textit{for } 2 \le j,l \le p \\ B_1(1,j) & = B_0(1,j) + \lambda_1 \mathbb{E}[b_{0,1}b_{1,j}] & \textit{for } 2 \le j \le p \end{array}$$

where for $2 \le j,l \le p$,

$$\begin{aligned} \mathbb{E}[b_{0,j}b_{1,l}] &= -\frac{2}{1-\beta}\mathbb{E}[L_j^\top X L_l^\top (X-\mu) w^\top (X-\mu) \mathbb{I}(z_{\theta_0}(X) \ge 0)] - 2\lambda_0 L_j^\top \mu L_l^\top \Sigma w \\ \mathbb{E}[b_{1,j}b_{1,l}] &= 4\mathbb{E}[L_j^\top (X-\mu)(X-\mu)^\top L_l w^\top (X-\mu)(X-\mu)^\top w] \\ \mathbb{E}[b_{0,1}b_{1,l}] &= 2L_l \Sigma w - \frac{2}{1-\beta}\mathbb{E}[L_l^\top (X-\mu) w^\top (X-\mu) \mathbb{I}(z_{\theta_0}(X) \ge 0)]. \end{aligned}$$

Proof. This is a straight-forward application of Corollary 1 for the case $X \sim \mathcal{N}(\mu, \Sigma)$. □

Remark – Justification of the parametric PBR method.
The nonparametric PBR method with only a penalty on the mean estimation is a linear combination of the empirical mean-CVaR problem (CVaR-emp) and the empirical Markowitz problem (Mark-emp) because the penalty is precisely the portfolio variance estimate $w^\top \hat{\Sigma}_n w$. In particular, this single-penalty problem approaches (Mark-emp) with increasing λ_1. In Figure 2.3, we plot 1 std of $w_n^v(\lambda_1, 0)^\top \mu$ and $CVaR(-w_n^v(\lambda_1, 0)^\top X; \beta)$ for the single-penalty problem as λ_1 is increased, for different values of λ_0, computed using Lemma 5. Observe that the asymptotic standard deviations for both portfolio mean and CVaR decrease with increasing λ_1, uniformly in λ_0. Given that both solutions to (CVaR-emp) and (Mark-emp) converge to the population solution w_0, the asymptotic theory deems the empirical Markowitz solution superior.

2.4.5 Proof of Theorem 2

Proof. By uniqueness of $\theta_0(\lambda_1, \lambda_2)$ and Assumption 1 (and compactness arguments), for every $\varepsilon > 0$, there exists $\eta > 0$ such that

$$||\hat{\theta}_n(\lambda_1, \lambda_2) - \theta_0(\lambda_1, \lambda_2)||_2 > \varepsilon \implies M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) > \eta.$$

Thus if we can show the probability of the event $\{M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) > \eta\}$ goes to zero for every $\varepsilon > 0$, then we have consistency.

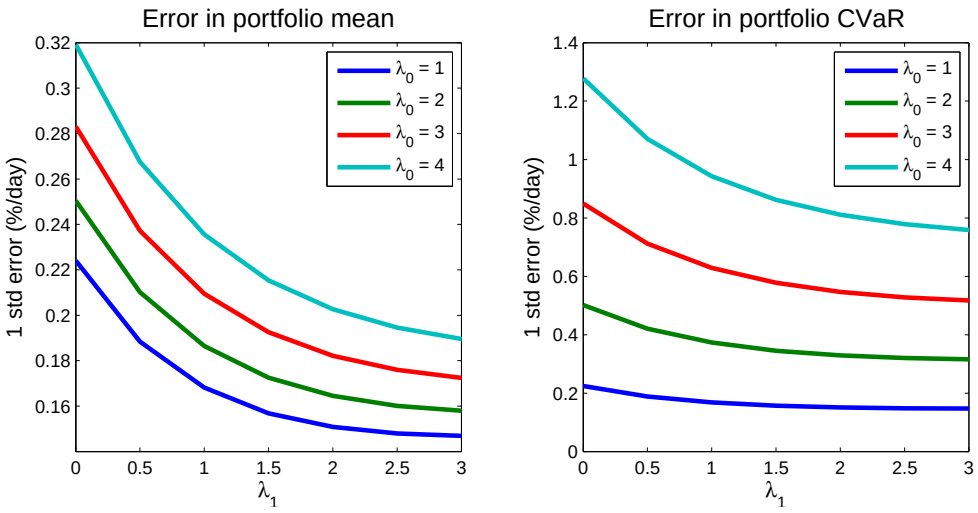

Figure 2.3 1 asymptotic std of the portfolio mean and CVaR for the single-penalty problem as λ_1 is increased when $X \sim \mathcal{N}(\mu, \Sigma)$, for different values of λ_0.

We also have

$$M_n(\hat{\theta}_n; \lambda_1, \lambda_2) \leq M_n(\theta_0; \lambda_1, \lambda_2) + o_P(1) = M(\theta_0; \lambda_1, \lambda_2) + o_P(1), \qquad (\star)$$

the first inequality because $\hat{\theta}_n(\lambda_1, \lambda_2)$ is a near-minimizer of M_n, and the second equality by the Weak Law of Large Numbers (WLLN) on $M_n(\theta_0; \lambda_1, \lambda_2)$.

Thus

$$\begin{aligned} 0 &\leq M(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2) \\ &= [M(\hat{\theta}_n; \lambda_1, \lambda_2) - M_n(\hat{\theta}_n; \lambda_1, \lambda_2)] + [M_n(\hat{\theta}_n; \lambda_1, \lambda_2) - M_n(\theta_0; \lambda_1, \lambda_2)] \\ &\quad + [M_n(\theta_0; \lambda_1, \lambda_2) - M(\theta_0; \lambda_1, \lambda_2)] \\ &\leq M(\hat{\theta}_n; \lambda_1, \lambda_2) - M_n(\hat{\theta}_n; \lambda_1, \lambda_2) + o_P(1), \end{aligned}$$

because the second term in [ ] is $o_P(1)$ by $(\star)$, and the last term in [ ] is $o_P(1)$ by WLLN. We are left to prove $|M_n(\hat{\theta}_n; \lambda_1, \lambda_2) - M(\hat{\theta}_n; \lambda_1, \lambda_2)| \xrightarrow{P} 0$. At first glance, one may consider invoking the WLLN again. However, as $\hat{\theta}_n(\lambda_1, \lambda_2)$ is a random sequence of vectors that changes for every n, we cannot apply the WLLN which is a pointwise result (i.e. for each fixed $\theta \in \Theta$), and we need to appeal to the stronger ULLN.

Case I: $\lambda_1 = \lambda_2 = 0$. To show ULLN for the original objective, we show that $\mathcal{F}_1$ is a Lipschitz class of functions, hence $N_{[\,]}(\varepsilon, \mathcal{F}_1, L_r(P))$ for every $\varepsilon > 0$. Now $\theta \mapsto m_\theta(x) = \alpha + (1-\beta)^{-1}(-\alpha - w_0^\top x - v^\top L^\top x)^+$ is clearly differentiable at θ_0 for all $x \in \mathbb{R}^p$. Furthermore,

$$\nabla_\theta m_\theta(x) = \begin{bmatrix} -1 \\ -L^\top x \end{bmatrix} I(x),$$

where $I(x) := \mathbb{I}(-\alpha - w_0^\top x - v^\top L^\top x \geq 0)$, hence

$$\dot{m}(x) := \max(1, ||L^\top x||_\infty) \tag{2.23}$$

is an upper bound on $||\nabla_\theta m_\theta(x)||_\infty$ and is independent of θ. Thus $|m_{\theta_1}(x) - m_{\theta_2}(x)| \leq \dot{m}(x)||\theta_1 - \theta_2||_2$ for all $\theta_1, \theta_2 \in [-K, K]^{1+p}$, and together with Assumption 2 (here a weaker assumption that X has finite second moment suffices), $\mathcal{F}_1$ is a Lipschitz class.

Case II: $\lambda_1 \geq 0, \lambda_2 \geq 0$, λ_1, λ_2 not both zero. Corollary 3.5 in [5] says that ULLN also holds for the penalized objective if $N_{[\,]}(\varepsilon, \mathcal{F}_2, L_2(P \times P)) < \infty$ for every $\varepsilon > 0$. Let us now show that $\mathcal{F}_2$ is also a Lipschitz class of functions. Again, it is clear that

$$\begin{aligned}
&\theta \mapsto m^U_{(\theta;\lambda_1,\lambda_2)}(x_1, x_2) \\
&\quad = \frac{1}{2}\left[m_\theta(x_1) + m_\theta(x_2)\right] + \frac{\lambda_1}{2}[(w_1 + Lv)^\top (x_1 - x_2)]^2 + \frac{\lambda_2}{2}(z_\theta(x_1) - z_\theta(x_2))^2
\end{aligned}$$

is differentiable at θ_0 for all $(x_1, x_2) \in \mathbb{R}^p \times \mathbb{R}^p$. Also for all $\theta \in [-K, K]^{1+p}$,

$$\begin{aligned}
\nabla_\theta \frac{\lambda_1}{2}[(w_1 + Lv)^\top (x_1 - x_2)]^2 &= \lambda_1 (x_1 - x_2)(x_1 - x_2)^\top (w_1 + Lv) \\
\implies ||\nabla_\theta \frac{\lambda_1}{2}[(w_1 + Lv)^\top (x_1 - x_2)]^2||_\infty &\leq \lambda_1 ||x_1 - x_2||^2_\infty ||w_1 + Lv||_\infty \\
&\leq \lambda_1 C(K) ||x_1 - x_2||^2_\infty,
\end{aligned}$$

for some constant $C(K)$ dependent on K. Also,

$$\nabla_\theta \frac{\lambda_2}{2}(z_\theta(x_1) - z_\theta(x_2))^2 = \lambda_2 (z_\theta(x_1) - z_\theta(x_2)) \begin{bmatrix} -I(x_1) + I(x_2) \\ -L^\top x_1 I(x_1) + L^\top x_2 I(x_2) \end{bmatrix}$$

and

$$\begin{aligned}
|z_\theta(x_1)| &= |-(\alpha - w_0^\top x_1 - v^\top L^\top x_1)^+| \leq |\alpha - w_0^\top x_1 - v^\top L^\top x_1| \\
&\leq K + |w_0^\top x_1| + K|e^\top x_1|
\end{aligned}$$

implies

$$\begin{aligned}
||\nabla_\theta \frac{\lambda_2}{2}(z_\theta(x_1) - z_\theta(x_2))^2||_\infty &\leq \lambda_2 |z_\theta(x_1) - z_\theta(x_2)|(\dot{m}(x_1) + \dot{m}(x_2)) \\
&\quad \dot{m} \text{ as defined in Eq. (2.23)} \\
&\leq \lambda_2 C'(K)(||x_1||_\infty + ||x_2||_\infty)(\dot{m}(x_1) + \dot{m}(x_2)),
\end{aligned}$$

for some constant $C'(K)$ dependent on K. Hence

$$\begin{aligned}
\dot{m}^U_{(\lambda_1,\lambda_2)}(x_1, x_2) &:= \frac{1}{2}[\dot{m}(x_1) + \dot{m}(x_2)] + \lambda_1 C(K)||x_1 - x_2||^2_\infty \\
&\quad + \lambda_2 C'(K)(||x_1||_\infty + ||x_2||_\infty)(\dot{m}(x_1) + \dot{m}(x_2))
\end{aligned} \tag{2.24}$$

is an upper bound on $||\nabla_\theta m^U_{(\theta;\lambda_1,\lambda_2)}(x_1, x_2)||_\infty$ that is independent of θ. Thus

$$|m^U_{(\theta_1;\lambda_1,\lambda_2)}(x_1, x_2) - m^U_{(\theta_2;\lambda_1,\lambda_2)}(x_1, x_2)| \leq \dot{m}^U_{(\lambda_1,\lambda_2)}(x_1, x_2)||\theta_1 - \theta_2||_2,$$

and together with Assumption 2, $\mathcal{F}_2$ is a Lipschitz class. □

2.4.6 Proof of Theorem 3

In what follows, we suppress the dependence on λ_1, λ_2 for notational convenience.

Proof. The proof parallels the proof of Theorem 5.23 of [61]. Let us assume for now that

1. For every given random sequence h_n that is bounded in probability,

$$\mathbb{U}_n[\sqrt{n}(m^U_{\theta_0+h_n/\sqrt{n}} - m^U_{\theta_0}) - h_n^\top \dot{m}^U_{\theta_0}] \xrightarrow{P} 0, \qquad (*)$$

   and

2. $\sqrt{n}(\hat{\theta}_n - \theta_0) = O_P(1)$.

Since $\theta \mapsto M(\theta)$ is twice-differentiable, and $\nabla_\theta M(\theta)|_{\theta=\theta_0} = 0$ by first-order condition, we can rewrite Eq. (*) to get

$$\begin{aligned} n\binom{n}{2}^{-1} \sum_{i\neq j} [m^U_{\theta_0+h_n/\sqrt{n}}(X_i, X_j) - m^U_{\theta_0}(X_i, X_j)] &= \frac{1}{2} h_n^\top V_{\theta_0} h_n + h_n^\top \mathbb{U}_n[\dot{m}^U_{\theta_0}] + o_p(1) \\ &= \frac{1}{2} h_n^\top V_{\theta_0} h_n + h_n^\top \mathbb{G}_n[\dot{m}^1_{\theta_0}] + o_p(1), \end{aligned}$$

where we use the fact, from Hoeffding decomposition,

$$\begin{aligned} \mathbb{U}_n[\dot{m}^U_{\theta_0}] &= \frac{\sqrt{n}}{\binom{n}{2}} \sum_{i\neq j} \left[\dot{m}^U_{\theta_0}(X_i, X_j) - \mathbb{E}_{X_1,X_2}[\dot{m}^U_{\theta_0}(X_1, X_2)]\right] \\ &= \frac{1}{\sqrt{n}} \sum_{i=1}^n [\dot{m}^1_{\theta_0}(X_i) - \mathbb{E}\dot{m}^1_{\theta_0}(X_1)] + o_p(1) = \mathbb{G}_n[\dot{m}^1_{\theta_0}] + o_p(1), \end{aligned}$$

with $\dot{m}^1_\theta$ as in the statement of the theorem.

The above statement is valid for both $\hat{h}_n = \sqrt{n}(\hat{\theta}_n - \theta_0)$ and for $\tilde{h}_n = -V_{\theta_0}^{-1}\mathbb{G}_n \dot{m}^1_{\theta_0}$. Upon substitution, we obtain

$$n\binom{n}{2}^{-1} \sum_{i\neq j} [m^U_{\theta_0+\hat{h}_n/\sqrt{n}}(X_i, X_j) - m^U_{\theta_0}(X_i, X_j)] = \frac{1}{2} \hat{h}_n^\top V_{\theta_0} \hat{h}_n + \hat{h}_n^\top \mathbb{G}_n[\dot{m}^1_{\theta_0}] + o_p(1)$$

$$\leq n\binom{n}{2}^{-1} \sum_{i\neq j} [m^U_{\theta_0+\tilde{h}_n/\sqrt{n}}(X_i, X_j) - m^U_{\theta_0}(X_i, X_j)] = -\frac{1}{2} \mathbb{G}_n[\dot{m}^1_{\theta_0}]^\top V_{\theta_0}^{-1} \mathbb{G}_n[\dot{m}^1_{\theta_0}] + o_p(1)$$

where the inequality is from the definition of $\hat{\theta}_n = \theta_0 + \hat{h}_n/\sqrt{n}$ as a near-minimizer.

Taking the difference and completing the square, we get

$$\frac{1}{2}(\hat{h}_n + V_{\theta_0}^{-1}\mathbb{G}_n \dot{m}^1_{\theta_0})^\top V_{\theta_0} (\hat{h}_n + V_{\theta_0}^{-1}\mathbb{G}_n \dot{m}^1_{\theta_0}) + o_p(1) \leq 0,$$

and because V_{θ_0} is nonsingular, the quadratic form on the left must converge to zero in probability. The same must be true for $||\hat{h}_n + V_{\theta_0}^{-1}\mathbb{G}_n\dot{m}^1_{\theta_0}||_2$.

To complete the proof, we need to show (*) and $\sqrt{n}(\hat{\theta}_n - \theta_0) = O_P(1)$ hold.

Proof of (*).
Let $f_h := \sqrt{n}(m^U_{\theta_0+h/\sqrt{n}} - m^U_{\theta_0}) - h^\top \dot{m}^U_{\theta_0}$. As we are considering only sequences h_n that are bounded in probability, it suffices to show $\sup_{h:||h||_2\leq 1}|\mathbb{U}_n[f_h]|$ goes to zero in probability. Again by Hoeffding decomposition, for any given random sequence h_n that is bounded in probability, $\mathbb{U}_n[f_{h_n}] = \mathbb{G}_n[f^1_{h_n}] + E_n(h_n)$, where f^1_h is the first term in the Hoeffding decomposition of $\mathbb{U}_n[f_h]$ given by

$$\begin{aligned} f^1_h &= \sqrt{n}(m^1_{\theta_0+h/\sqrt{n}} - m^1_{\theta_0}) - h^\top \dot{m}^1_{\theta_0}, \\ m^1_\theta(x_1) &= 2\mathbb{E}_{X_2}[m^U_\theta(x_1, X_2)] - \mathbb{E}_{X_1,X_2}[m^U_\theta(X_1, X_2)], \end{aligned}$$

and $\dot{m}^1_\theta$ as defined in the statement of the theorem. According to Lemma 19.31 in [61], if $\mathcal{F}'_2 := \{m^1_\theta : \theta \in [-K, K]^{1+p}\}$ is a Lipschitz class of functions,

$$\sup_{h:||h||_2\leq 1} |\mathbb{G}_n[f^1_{h_n}]| \xrightarrow{P} 0.$$

Now by Assumption 2 that X_i's are iid continuous random vectors with finite fourth moment, $\theta \mapsto m^1_\theta(x)$ is differentiable at θ_0 for all $x \in \mathbb{R}$. Further, by triangle inequality,

$$\begin{aligned} &|m^1_{\theta_1}(x) - m^1_{\theta_2}(x)| \\ &\leq 2\mathbb{E}_{X_2}|m^U_{\theta_1}(x, X_2) - m^U_{\theta_2}(x, X_2)| + \mathbb{E}_{X_1,X_2}|m^U_\theta(X_1, X_2) - m^U_\theta(X_1, X_2)| \\ &\leq m^1(x)||\theta_1 - \theta_2||_2, \end{aligned}$$

where $m^1(x) = (2\mathbb{E}_{X_2}|\dot{m}^U(x, X_2)| + \mathbb{E}_{X_1,X_2}|\dot{m}^U(X_1, X_2)|)$, $\dot{m}^U$ as in Eq. (2.24). Since X_i's have finite fourth moment, $\mathbb{E}[m^1(X_1)^2] < \infty$ and thus $\mathcal{F}'_2$ is a Lipschitz class.

Now we are left to show $\sup_{h:||h||_2\leq 1}|E_n(h)| \xrightarrow{P} 0$. Let $\mathcal{F}_h := \{f_h : ||h||_2 \leq 1\}$. According to Theorem 4.6 of [5], $\sup_{h:||h||_2\leq 1}|E_n(h)| \xrightarrow{P} 0$ if $\mathcal{F}_h$ has a finite, integrable envelope function and both $\mathcal{F}_h$ and $\mathcal{F}'_h := \{f^1_h : ||h||_2 \leq 1\}$ are Lipschitz classes about $h = 0$. $\mathcal{F}_h$ has a finite, integrable envelope function $F(x_1, x_2) = \dot{m}^U(x_1, x_2) + ||\dot{m}_{\theta_0}(x_1, x_2)||_2 < \infty$ due to Assumption 2 and the Lipschitz property of m^U_θ:

$$\begin{aligned} |f_h| &\leq |\sqrt{n}(m^U_{\theta_0+h/\sqrt{n}} - m^U_{\theta_0}) - h^\top \dot{m}_{\theta_0}| \\ &\leq (\dot{m}^U + ||\dot{m}_{\theta_0}||_2)||h||_2. \end{aligned}$$

It is now straight-forward to check that $\mathcal{F}_h$ is a Lipschitz class about $h = 0$, and $\mathcal{F}'_h$ also, because it inherits the key properties from $\mathcal{F}_h$.

Proof of $\sqrt{n}(\hat{\theta}_n - \theta_0) = O_P(1)$.
The proof of $\sqrt{n}(\hat{\theta}_n(0,0) - \theta_0(0,0)) = O_p(1)$ can be found in Theorem 5.52 and Corollary 5.53 of [61], and is a standard M-estimation result. In essence, Theorem 5.52 shows that, under some regularity conditions, $P(\sqrt{n}||\hat{\theta}_n(0,0) - \theta_0(0,0)||_2 > \alpha)$ can be bounded by $P(|\mathbb{G}_n[m_\theta]| > \alpha') = P(\sqrt{n}|M_n(\theta) - M(\theta)| > \alpha')$, which is shown to go to zero via some maximal inequalities. Corollary 5.53 shows that the Lipschitz condition on $\{m_\theta : \theta \in [-K, K]^{1+p}\}$ is sufficient to satisfy the regularity conditions of the theorem.

We can extend Theorem 5.52 to show $\sqrt{n}(\hat{\theta}_n(\lambda_1, \lambda_2) - \theta_0(\lambda_1, \lambda_2))$, $\lambda_1, \lambda_2 \geq 0$ not both zero, by bounding $P(\sqrt{n}||\hat{\theta}_n(\lambda_1, \lambda_2) - \theta_0(\lambda_1, \lambda_2)||_2 > \alpha)$ by

$$P(|\mathbb{U}_n[m_\theta^U]| > \alpha') \leq P(|\mathbb{G}_n[m_\theta^1]| + |E_n'(\theta)| > \alpha'),$$

where E_n' is the remainder term after first-order projection of the U-process $\mathbb{U}_n[m_\theta^U]$. It remains to show that for every sufficiently small $\delta > 0$,

$$\sup_{\theta : ||\theta - \theta_0||_2 < \delta} |E_n'(\theta)| \xrightarrow{P} 0, \tag{2.25}$$

which can be proven using the same reasoning for $\sup_{h : ||h||_2 \leq 1} |E_n(h)| \xrightarrow{P} 0$ in the proof of (*). □

2.4.7 Computation of key statistics

Given the distribution for X, both $A_0 = A_{\theta_0}(0,0)$ and $B_0 = B_{\theta_0}(0,0)$ are computable. The lemma below computes the key quantities that constitute A_0 and B_0 when $X \sim \mathcal{N}(\mu, \Sigma)$.

Lemma 6. *Suppose $X \sim \mathcal{N}(\mu, \Sigma)$, and $z_\theta(X) = -\alpha - w^\top X \sim \mathcal{N}(\mu_1, \sigma_1^2)$, where $\mu_1 = -\sigma_1 \Phi^{-1}(\beta)$ and $\sigma_1^2 = w^\top \Sigma w$. Then*

$$p_0 = P(z_\theta(X) = 0) = \frac{1}{\sqrt{2\pi}\sigma_1} \exp\left(-\frac{1}{2}\Phi^{-1}(\beta)^2\right) \tag{2.26}$$

$$\mathbb{E}[\max(z_\theta(X), 0)] = \frac{\sigma_1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\Phi^{-1}(\beta)^2\right) - \sigma_1(1-\beta)\Phi^{-1}(\beta) \tag{2.27}$$

$$\mathbb{E}[L_j^\top X \mathbb{I}(Z_1 \geq 0)] = (1-\beta)(L_j^\top \mu - \Phi^{-1}(\beta)\frac{L_j^\top \Sigma w}{\sigma_1}) - \frac{L_j^\top \Sigma w}{\sigma_1^2}\mathbb{E}[\max(z_\theta(X), 0)] \tag{2.28}$$

$$\mathbb{E}[L_j^\top X | Z_1 = 0] = L_j^\top \mu - \Phi^{-1}(\beta)\frac{L_j^\top \Sigma w}{\sigma_1} \tag{2.29}$$

$$\begin{aligned}\mathbb{E}[L_j^\top X L_l^\top X \mathbb{I}(Z_1 \geq 0)] &= \frac{1}{4}(g(\mu_1, (L_j+L_l)^\top \mu, \sigma_1, \sigma_2, -(L_j+L_l)^\top \Sigma w_1) \\ &\quad -g(\mu_1, (L_j-L_l)^\top \mu, \sigma_1, \sigma_2, -(L_j-L_l)^\top \Sigma w_1))\end{aligned} \tag{2.30}$$

$$\begin{aligned}\mathbb{E}[L_j^\top X L_l^\top X | Z_1 = 0] &= \frac{1}{4}(h(\mu_1, (L_j+L_l)^\top \mu, \sigma_1, \sigma_2, -(L_j+L_l)^\top \Sigma w_1) \\ &\quad -h(\mu_1, (L_j-L_l)^\top \mu, \sigma_1, \sigma_2, -(L_j-L_l)^\top \Sigma w_1))\end{aligned} \tag{2.31}$$

where

$$\begin{aligned} g(\mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_{12}) &= (1-\beta)\left[\mu_2^2 + \sigma_2^2\right] + p_0 \sigma_{12}\left[-\Phi^{-1}(\beta)\frac{\sigma_{12}}{\sigma_1} + 2\mu_2\right] \\ h(\mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_{12}) &= = (\mu_2 + \frac{\sigma_{12}}{\sigma_1}\Phi^{-1}(\beta))^2 + \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}\ . \end{aligned}$$

Proof. Equations (2.26) and (2.27) are straight-forward, for Eqs. (2.28)–(2.31) we use the fact that if $Z_1 \sim \mathcal{N}(\mu_1, \sigma_1)$ and $Z_2 \sim \mathcal{N}(\mu_2, \sigma_2)$,

$$Z_2 | Z_1 = \mathcal{N}(\mu_2 + \sigma_{12}/\sigma_1^2(Z_1 - \mu_1), \sigma_2^2 - \sigma_{12}^2/\sigma_1^2), \tag{2.32}$$

where $\sigma_{12} = Cov(Z_1, Z_2)$.

• **Terms involving only $L_j^\top X$.**
Note that from (2.32), $\mathbb{E}[Z_2 | Z_1 = 0] = \mu_2 - \frac{\sigma_{12}}{\sigma_1^2}\mu_1$. Let $Z_2 = L_j^\top X$, and recall that $\mathbb{E}(L_j^\top X) = L_j^\top \mu$ and $\mathbb{E}(Z_1) = -\sigma_1 \Phi^{-1}(\beta)$. Also, note that $\sigma_{12} = -L_j^\top \Sigma w$. After some algebra, we get (2.29).

Since we know the distribution of $Z_2 | Z_1$, we have

$$\begin{aligned}\mathbb{E}[Z_2 \mathbb{I}(Z_1 \geq 0)] &= \mathbb{E}[\mathbb{I}(Z_1 \geq 0)(\mu_2 + \frac{\sigma_{12}}{\sigma_1^2}(Z_1 - \mu_1))] \\ &= (1-\beta)(\mu_2 - \frac{\sigma_{12}}{\sigma_1^2}\mu_1) + \frac{\sigma_{12}}{\sigma_1^2}\mathbb{E}[Z_1 \mathbb{I}(Z_1 \geq 0)] \\ &= (1-\beta)(L_j^\top \mu - \Phi^{-1}(\beta)\frac{L_j^\top \Sigma w}{\sigma_1}) - \frac{L_j^\top \Sigma w}{\sigma_1^2}\mathbb{E}[\max(Z_1, 0)]\end{aligned}$$

• **Terms involving $L_j^\top X L_l^\top X$.**
To compute $\mathbb{E}[L_j^\top X L_l^\top X \mathbb{I}(Z_1 \geq 0)]$ and $\mathbb{E}[L_j^\top X L_l^\top X | Z_1 = 0]$, first note that

$$\mathbb{E}[L_j^\top X L_l^\top X \mathbb{I}(Z_1 \geq 0)] = \frac{1}{4}\mathbb{E}\left[[(L_j^\top X + L_l^\top X)^2 - (L_j^\top X - L_l^\top X)^2]\mathbb{I}(Z_1 \geq 0)\right]\ .$$

and similarly

$$\mathbb{E}[L_j^\top X L_l^\top X | Z_1 = 0] = \frac{1}{4}\mathbb{E}[\left[(L_j^\top X + L_l^\top X)^2 - (L_j^\top X - L_l^\top X)^2\right] | Z_1 = 0] \, .$$

Hence it is sufficient to first find expressions for $\mathbb{E}[Z_2^2 \mathbb{I}(Z_1 \geq 0)]$ and $\mathbb{E}[Z_2^2 | Z_1 = 0]$ for some normal Z_2, then apply the resulting formulae to $Z_2 = (L_j \pm L_l)^\top X$. This results in $\mu_2 = (L_j \pm L_l)^\top \mu$, $\sigma_{12} = -(L_j \pm L_l)^\top \Sigma w$ and $\sigma_2^2 = (L_j \pm L_l)^\top \Sigma (L_j \pm L_l)$.

From tower property and the conditional distribution of $Z_2 | Z_1$,

$$\mathbb{E}[Z_2^2 \mathbb{I}(Z_1 \geq 0)] = \mathbb{E}[\mathbb{I}(Z_1 \geq 0) \left[(\mu_2 + \frac{\sigma_{12}}{\sigma_1^2}(Z_1 - \mu_1))^2 + \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}\right]] \, .$$

By simple computations,

$$\begin{aligned}
\mathbb{E}[(Z_1 - \mu_1)\mathbb{I}_{Z_1 \geq 0}] &= \frac{\sigma_1}{\sqrt{2\pi}} \exp(-\mu_1^2/(2\sigma_1^2)) = \sigma_1^2 f_{Z_1}(0) = \sigma_1^2 p_0 \, , \text{ and} \\
\mathbb{E}[(Z_1 - \mu_1)^2 \mathbb{I}_{Z_1 \geq 0}] &= \sigma_1^2(\mu_1 p_0 + (1 - \beta)) \, .
\end{aligned}$$

Now $\mu_1/\sigma_1 = -\Phi^{-1}(\beta)$, and

$$\begin{aligned}
\mathbb{E}[Z_2^2 \mathbb{I}_{Z_1 \geq 0}] &= (1 - \beta)\left[\mu_2^2 + \sigma_2^2\right] + p_0 \left[\mu_1 \frac{\sigma_{12}^2}{\sigma_1^2} + 2\sigma_{12}\mu_2\right] \\
&= (1 - \beta)\left[\mu_2^2 + \sigma_2^2\right] + p_0 \sigma_{12} \left[-\Phi^{-1}(\beta)\frac{\sigma_{12}}{\sigma_1} + 2\mu_2\right] \\
&:= g(\mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_{12})
\end{aligned}$$

Similarly,

$$\begin{aligned}
\mathbb{E}[Z_2^2 | Z_1 = 0] &= (\mu_2 - \frac{\sigma_{12}}{\sigma_1^2}\mu_1)^2 + \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2} \\
&= (\mu_2 + \frac{\sigma_{12}}{\sigma_1}\Phi^{-1}(\beta))^2 + \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2} \\
&:= h(\mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_{12})
\end{aligned}$$

□

2.5 Numerical results

In this section, we present simulation results to evaluate the nonparametric and parametric PBR methods presented in Sec. 2.3 against the straight-forward approach (CVaR-emp). We consider $p = 10$ assets and three distributional models for the asset log-returns: X is multivariate Gaussian, elliptical and mixture of multivariate

Gaussian and negative exponential. For each model, we follow the procedure outlined in Section 2 to construct sample efficient frontiers corresponding to (CVaR-emp), (CVaR-pen) and (Mark-emp).

One question that arises while solving (CVaR-pen) is how one chooses the penalty terms U_1 and U_2 in the constraints

$$\begin{aligned} \frac{1}{n} w^\top \hat{\Sigma}_n w &\leq U_1 \\ \frac{1}{n(1-\beta)^2} z^\top \Omega_n z &\leq U_2. \end{aligned}$$

If U_1, U_2 are too small, the problem becomes infeasible, whereas if they are too large, the penalization does not have any effect. It is sensible to choose U_1, U_2 as a proportion of $\hat{w}_n^\top \hat{\Sigma}_n \hat{w}_n / n$ and $\hat{z}_n^\top \Omega_n \hat{z}_n / (n(1-\beta)^2)$ respectively, where $(\hat{w}_n, \hat{z}_n)$ is the solution to the unpenalized problem (CVaR-emp). We denote the proportions r_1 and r_2 respectively. In practice, one would perform cross-validation to find values of $(r_1, r_2) \in [0,1] \times [0,1]$ that maximize out-of-sample performance.

2.5.1 Gaussian/elliptical models

Here we consider

$$X \sim \mu_{sim} + \lambda \mathcal{N}(0, \Sigma_{sim})$$

where λ is as in (2.5), with $\lambda = 1$ for a Gaussian model and $\lambda \sim \Gamma(3, 0.5)$ for an elliptical model. The parameters μ_{sim} and Σ_{sim} are the same as those used in Sec. 2.2.2. We plot the histograms for $100,000$ sample returns for an equally-weighted portfolio $w = 1_p/p$ under the Gaussian and elliptical models in Fig. (2.4).

We summarize the simulation results in Fig. (2.5), where $(r_1, r_2) = (0.92, 1)$ for both the Gaussian and elliptical models (recall that the second penalty is redundant due to Lemma 2). Notice that for both models, the empirical Markowitz efficient frontier dominates the penalized efficient frontier which in turn dominates the empirical mean-CVaR efficient frontier, in both *position* of the average of the simulated frontiers and *variability*, as indicated by the vertical and horizontal error bars.

For the Gaussian case, $r_1 = 0.92$ was just feasible in that further reduction in this value led to most instances of the problem being infeasible. From Fig. (2.5b), we can see that this is because the penalized solutions are approaching the empirical Markowitz solutions with this choice of r_1 as the average simulated efficient frontiers of penalized (grey) and empirical Markowitz (blue) solutions are close. For the elliptical model, $r_1 = 0.92$ could be further reduced with the resulting penalized efficient frontier approaching the empirical Markowitz efficient frontier. In summary, the empirical Markowitz solutions perform uniformly better than both the original and penalized mean-CVaR solutions, with the penalized efficient frontier nearing the empirical Markowitz efficient frontier with decreasing r_1.

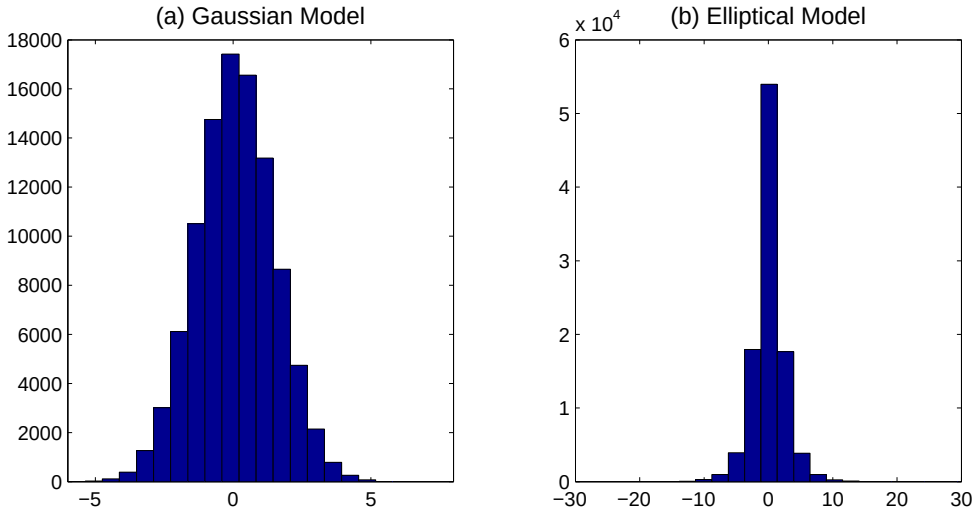

Figure 2.4 Distribution of equally weighted portfolio under (a) Gaussian and (b) elliptical model

2.5.2 Mixture model

Let us now consider returns being driven by a mixture of multivariate normal and negative exponential distributions, such that with a small probability, all assets undergo a perfectly correlated exponential-tail loss. Formally,

$$X \sim (1 - I(q))N(\mu_{sim}, \Sigma_{sim}) + I(q)(Y 1_p + f), \tag{2.33}$$

where $(\mu_{sim}, \Sigma_{sim})$ are parameters with the same value as in the Gaussian/elliptical models, $I(q) \sim Bernoulli(q)$, and $f = [f_1, \ldots, f_p]^\top$ is a $p \times 1$ vector of constants, and Y is a negative exponential random variable with density

$$P(Y = y) = \begin{cases} \lambda e^{\lambda y}, & \text{if } y \leq 0 \\ 0 & \text{otherwise.} \end{cases}$$

In our simulations, we consider $q = 0.05$, $f_i = \mu_i - \sqrt{\Sigma_{ii}}$ for $i = 1, \ldots, p$ and $\lambda = 1$. The histogram for $100,000$ sample returns of an equally-weighted portfolio under this mixture model is shown in Fig. (4a).

We summarize the simulation results in Fig. (4b), where $(r_1, r_2) = (0.5, 0.5)$. In this case, the penalized efficient frontiers perform better on average than the efficient frontiers generated by the other two methods. The empirical Markowitz efficient frontiers do not seem to perform any better than the original efficient frontiers on average, which is not surprising because the empirical Markowitz solution is only intended for X having an elliptical distribution.

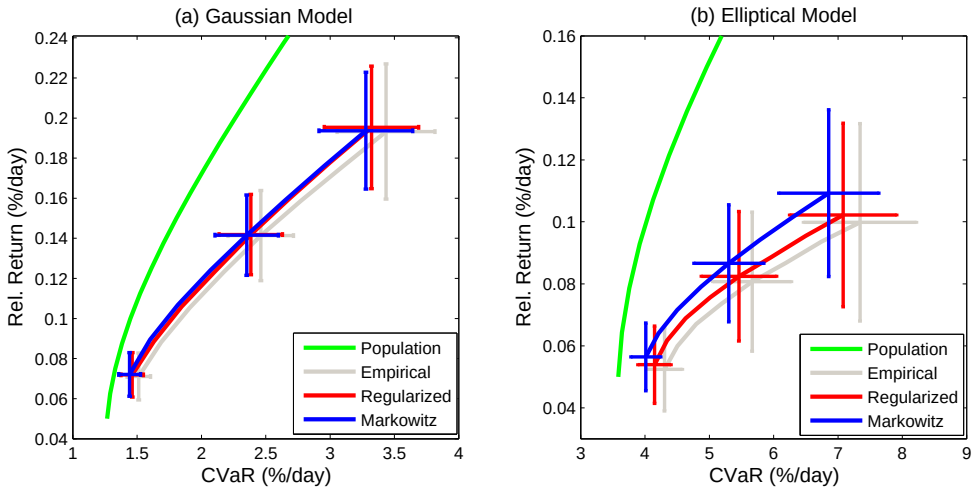

Figure 2.5 Average of population risk vs return for solutions to (CVaR-emp) in grey, (CVaR-pen) in red and (Mark-emp) in blue under (a) Gaussian model and (b) elliptical model. Green curve denotes the population efficient frontier. Horizontal and vertical lines show $\pm 1/2$ std error.

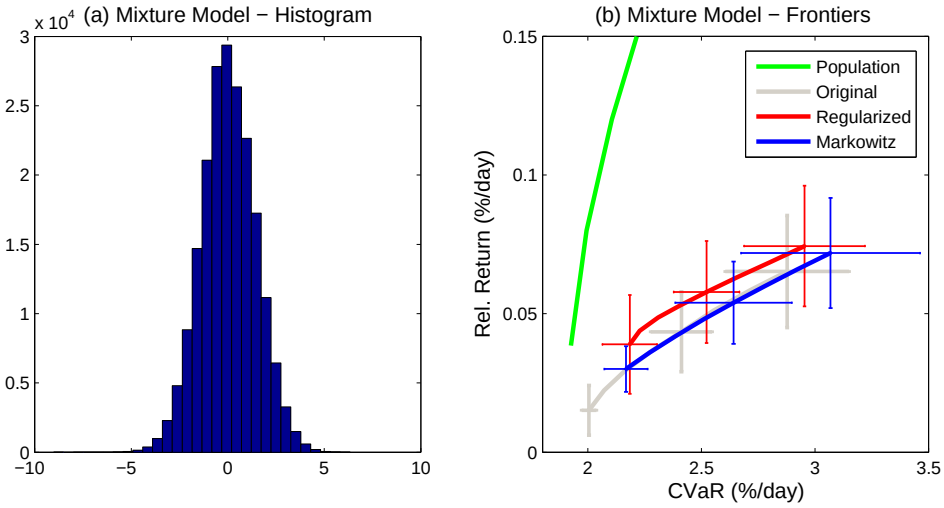

Figure 2.6 (a) Distribution of returns for an equally weighted portfolio under the mixture model. (b) Average of population risk vs return for solutions to (CVaR-emp) in grey, (CVaR-pen) in red and (Mark-emp) in blue under the mixture model. Green curve denotes the population efficient frontier. Horizontal and vertical lines show $\pm 1/2$ std error.

2.6 Appendix

2.6.1 Asymptotics of the CVaR estimator

Setting. Let $\mathsf{L} = [L_1, \ldots, L_n]$ be n iid observations (of portfolio losses) from a distribution F which is absolutely continuous, has a twice continuously differentiable pdf and a finite second moment.

In this section, we prove the asymptotic distribution of the estimator $\widehat{CVaR}_n(L;\beta)$ introduced in Eq. (2.2) of Sec. 2.1. First, we define a closely related CVaR estimator:

Definition 4 (**Type 1 CVaR estimator**.)**.** *For $\beta \in (0.5, 1)$, we define Type 1 CVaR estimator to be*

$$\widehat{CV1}_n(\mathsf{L};\beta) := \min_{\alpha\in\mathbb{R}} \ (1-\varepsilon_n)\alpha + \frac{1}{n - \lceil n\beta\rceil + 1}\sum_{i=1}^{n}(L_i-\alpha)^+,$$

where ε_n is some constant satisfying $0 < \varepsilon_n < (n - \lceil n\beta\rceil + 1)^{-1}$, $\sqrt{n}\varepsilon_n \xrightarrow{P} 0$.

Now consider the following CVaR estimator, expressed without the minimization:

Definition 5 (**Type 2 CVaR estimator**.)**.** *For $\beta \in (0.5, 1)$, we define Type 2 CVaR estimator to be*

$$\widehat{CV2}_n(\mathsf{L};\beta) := \frac{1}{n - \lceil n\beta\rceil + 1}\sum_{i=1}^{n} L_i 1(L_i \geq \hat{\alpha}_n(\beta)),$$

where $\hat{\alpha}_n(\beta) := L_{(\lceil n\beta\rceil)}$, the $\lceil n\beta\rceil$-th order statistic of the sample $L_1, \ldots, L_n$.

Type 2 CVaR estimator is asymptotically normally distributed [17]. In the remainder of this section, we show that $\widehat{CV2}_n(\mathsf{L};\beta)$ is asymptotically equivalent to $\widehat{CV1}_n(\mathsf{L};\beta)$, which is in turn asymptotically equivalent to $\widehat{CVaR}_n(\mathsf{L};\beta)$. We then conclude $\widehat{CVaR}_n(\mathsf{L};\beta)$ is also asymptotically normal, converging to the same asymptotic distribution as $\widehat{CV2}_n(\mathsf{L};\beta)$.

Proposition 1. *The solution $\alpha^* = L_{(\lceil n\beta\rceil)}$ is unique to the one-dimensional optimization problem*

$$\min_{\alpha\in\mathbb{R}} \ \left\{ G_n(\alpha) := (1-\varepsilon_n)\alpha + \frac{1}{n - \lceil n\beta\rceil + 1}\sum_{i=1}^{n}(L_i-\alpha)^+ \right\},$$

where ε_n is some constant satisfying $0 < \varepsilon_n < (n - \lceil n\beta\rceil + 1)^{-1}$.

Proof. The expression to be minimized is a piecewise linear convex function with nodes at $L_1, \ldots, L_n$. We show that $G_n(\alpha)$ has gradients of opposite signs about a single point, $L_{(\lceil n\beta \rceil)}$, hence this point must be the unique optimal solution. Now consider, for $m \in \{-\lceil n\beta \rceil + 1, \ldots, n - \lceil n\beta \rceil\}$:

$$\begin{aligned} \Delta(m) &= G_n(L_{(\lceil n\beta \rceil + m + 1)}) - G_n(L_{(\lceil n\beta \rceil + m)}) \\ &= (1 - \varepsilon_n)(L_{(\lceil n\beta \rceil + m + 1)} - L_{(\lceil n\beta \rceil + m)}) - \frac{1}{n - \lceil n\beta \rceil + 1} A, \end{aligned}$$

where

$$\begin{aligned} A &= \sum_{i=1}^{n} \left[(L_i - L_{(\lceil n\beta \rceil + m + 1)})^+ - (L_i - L_{(\lceil n\beta \rceil + m)})^+ \right] \\ &= (n - \lceil n\beta \rceil - m)(L_{(\lceil n\beta \rceil + m + 1)} - L_{(\lceil n\beta \rceil + m)}). \end{aligned}$$

Thus

$$\Delta(m) = \left(L_{(\lceil n\beta \rceil + m + 1)} - L_{(\lceil n\beta \rceil + m)} \right) \left((1 - \varepsilon_n) - \frac{n - \lceil n\beta \rceil - m}{n - \lceil n\beta \rceil + 1} \right).$$

Now $\Delta(0) > 0$ since $(L_{(\lceil n\beta \rceil + 1)} - L_{(\lceil n\beta \rceil)}) > 0$ and $(1 - \varepsilon_n) > (n - \lceil n\beta \rceil)(n - \lceil n\beta \rceil + 1)^{-1}$ by the restriction on ε_n, and $\Delta(-1) < 0$ since $(L_{(\lceil n\beta \rceil)} - L_{(\lceil n\beta \rceil - 1)}) > 0$ and $(1 - \varepsilon_n) < 1$ again by the choice of ε_n. Thus $G_n(\alpha)$ has a unique minimum at $\alpha^* = L_{(\lceil n\beta \rceil)}$. □

Remark. Note if $\varepsilon_n = 0$, then multiple solutions occur because $\Delta(-1) = 0$.

Corollary 2. *Type 1 and Type 2 CVaR estimators are related by*

$$\widehat{CV2}_n(\mathsf{L}; \beta) = \widehat{CV1}_n(\mathsf{L}; \beta) + \varepsilon_n L_{(\lceil n\beta \rceil)}.$$

Proof. Rewriting Type 2 CVaR estimator:

$$\begin{aligned} \widehat{CV2}_n(\mathsf{L}; \beta) &= \frac{1}{n - \lceil n\beta \rceil + 1} \sum_{i=1}^{n} L_i 1(L_i \geq L_{(\lceil n\beta \rceil)}) \\ &= L_{(\lceil n\beta \rceil)} + \frac{1}{n - \lceil n\beta \rceil + 1} \sum_{i=1}^{n} (L_i - L_{(\lceil n\beta \rceil)}) 1(L_i \geq L_{(\lceil n\beta \rceil)}) \\ &= \widehat{CV1}_n(\mathsf{L}; \beta) + \varepsilon_n L_{(\lceil n\beta \rceil)}, \end{aligned}$$

where the final equality is due to Proposition 1. □

We now show asymptotic normality of $\widehat{CV1}_n(\mathsf{L}; \beta)$.

Lemma 7. *Type 1 CVaR estimator is asymptotically normal as follows:*

$$\frac{\sqrt{n}(1 - \beta)}{\gamma_0} \left(\widehat{CV1}_n(\mathsf{L}; \beta) - CVaR(L_1; \beta) \right) \Rightarrow \mathcal{N}(0, 1), \tag{2.34}$$

where $\gamma_0^2 = Variance[(L_1 - \alpha_\beta) 1(L_1 \geq \alpha_\beta)]$, *and* $\alpha_\beta = \inf\{\alpha : P(L_1 \geq \alpha) \leq 1 - \beta\}$, *Value-at-Risk of the random loss* L_1 *at level* β.

Proof. Asymptotic normality for Type 2 CVaR estimator is proven in [17], and the result is immediate from invoking Slutsky's lemma on Corollary 2 and the assumption $\sqrt{n}\varepsilon_n \to 0$. □

Proof of Lemma 1

Proof. The asymptotic distribution of $\widehat{CVaR}_n(\mathsf{L};\beta)$ is the same as $\widehat{CV1}_n(\mathsf{L};\beta)$ because

$$\sqrt{n}|\widehat{CVaR}_n(\mathsf{L};\beta) - \widehat{CV1}_n(\mathsf{L};\beta)| = o_P(1).$$

□

2.6.2 Proof of Theorem 1

Lemma 8. *Consider the optimization problem*

$$\begin{array}{lrll} \min\limits_{z\in\mathbb{R}^n} & z^\top 1_n & & \\ s.t. & z_i \geq & 0 & \forall\ i \\ & z_i \geq & c_i & \forall\ i \\ & z^\top \Omega_n z \leq & f & \end{array} \tag{2.35}$$

where $c_i > 0\ \forall\ i$, $f > 0$, $\Omega_n = (n-1)^{-1}(I_n - n^{-1}1_n 1_n^\top)$, the sample covariance operator. Suppose (2.35) is feasible with an optimal solution (x^*, z^*). Let $S_1(z) := \{1 \leq i \leq n :\ z_i = 0\}$, $S_2(z) := \{1 \leq i \leq n :\ z_i = c_i\}$ and $V(z) := S_1^c \cap S_2^c$ (i.e. $V(z)$ is the set of indices for which $z_i > \max(0, c_i)$). Then the optimal solution z^* falls into one of two cases: either $S_1(z^*) \neq \emptyset$ and $V(z^*) = \emptyset$, or $S_1(z^*) = \emptyset$ and $V(z^*) \neq \emptyset$.

Proof. The problem (2.35) is a convex optimization problem because Ω_n is a positive semidefinite matrix. The problem is also strictly feasible, since $z_0 = 2\max_i\{c_i\}1_n$ is a strictly feasible point: clearly, $z_{0,i} > \max\{0, c_i\}\ \forall\ i$ and $z_0^\top \Omega_n z_0 = 0 < f$ as 1_n is orthogonal to Ω_n. Thus Slater's condition for strong duality holds, and we can derive properties of the optimal solution by examining KKT conditions.

The Lagrangian is

$$\mathcal{L}(z, \eta_1, \eta_2, \lambda) = \lambda z^\top \Omega_n z + (1_n - \eta_1 - \eta_2)^\top z + \eta_2^\top c - \lambda f$$

The KKT conditions are

- Primal feasibility
- Dual feasibility: $\eta_1^*, \eta_2^* \geq 0$ component-wise and $\lambda^* \geq 0$
- Complementary slackness:
  $z_i^* \eta_{1,i}^* = 0\ \forall\ i$, $(z_i^* - c_i)\eta_{2,i}^* = 0\ \forall\ i$ and $\lambda^*[(z^*)^\top \Omega_n z^* - f] = 0$

- First Order Condition:

$$\nabla_{z^*}\mathcal{L} \;=\; 2\lambda\Omega_n z^* + (1_n - \eta_1^* - \eta_2^*) = 0 \tag{2.36a}$$

By substituting for Ω_n, (2.36a) can be written as

$$\frac{2\lambda}{n-1}\left(z^* - \frac{1}{n}(1_n^\top z^*)1_n\right) \;=\; -1_n + \eta_1^* + \eta_2^*. \tag{2.37}$$

Suppose $S_1(z^*) \neq \emptyset$ at the optimal primal-dual point $(z^*, \eta_1^*, \eta_2^*, \lambda^*)$. Then $\exists\ i_0 \in S_1(z^*)$ such that $z_{i_0}^* = 0$. The i_0-th component of (2.37) gives

$$-\frac{2\lambda^*}{n(n-1)}(1_n^\top z^*) \;=\; -1 + \eta_{1,i_0}^* + \eta_{2,i_0}^*. \tag{2.38}$$

Now suppose $V(z^*) \neq \emptyset$ at the optimal primal-dual point $(z^*, \eta_1^*, \eta_2^*, \lambda^*)$. Then $\exists\ j_0 \in V(z^*)$ such that $z_{j_0}^* > \max(0, c_i)$, $\eta_{1,j_0}^* = 0$ and $\eta_{2,j_0}^* = 0$. The j_0-th component of (2.37) gives

$$\frac{2\lambda^*}{n-1}\left(z_{j_0}^* - \frac{1}{n}(1_n^\top z^*)\right) \;=\; -1, \tag{2.39}$$

which also implies $\lambda^* > 0$.

Now suppose $S_1(z^*)$ and $V(z^*)$ are both nonempty. Combining (2.38) and (2.39), we arrive at the necessary condition

$$\frac{2\lambda^*}{n-1} z_{j_0}^* = -\eta_{1,i_0}^* - \eta_{2,i_0}^*.$$

which is clearly a contradiction since $LHS > 0$ whereas $RHS \leq 0$. Hence $S_1(z^*)$ and $V(z^*)$ cannot both be nonempty. □

Proof of Theorem 1

Proof. Clearly, (CVaR-relax) is a relaxation of (CVaR-pen): the components of the variable z in (CVaR-relax) are relaxations of $\max(0, -w^\top X_i - \alpha)$. Thus the two problem formulations are equivalent if at optimum, $z_i = \max(0, -w^\top X_i - \alpha)\ \forall\ i = 1, \ldots, n$ for (CVaR-relax).

Let $(\alpha^*, w^*, z^*, \nu_1^*, \nu_2^*, \eta_1^*, \eta_2^*, \lambda_1^*, \lambda_2^*)$ be the primal-dual optimal point for (CVaR-relax) and (CVaR-relax-d). Our aim is to show that $V(z^*)$, the set of indices for which $z_i^* > \max(0, -w^\top X_i - \alpha)$, is empty. Suppose the contrary. Then by Lemma 8, $S_1(z^*)$, the set of indices for which $z_i^* = 0$, is empty. This means $z_i^* > 0\ \forall\ i$ and $\eta_{1,i}^* = 0\ \forall\ i$ by complementary slackness.

Now consider the sub-problem for a fixed η_2 in the dual problem (CVaR-relax-d):

$$\max_{\eta_1:\eta_1\geq 0} -(\eta_1+\eta_2)\Omega_n^\dagger(\eta_1+\eta_2). \tag{2.40}$$

As 1_n is orthogonal to $\Omega_n^\dagger$, and $\Omega_n^\dagger$ is positive semidefinite, the optimal solution is of the form $\eta_1 = a1_n - \eta_2$, where a is any constant such that $a \geq \max_i(\eta_{2,i})$, with a corresponding optimal objective 0. Hence, bearing in mind the constraints $\eta_2 \geq 0$ and $\eta_2^\top 1_n = 1$ in (CVaR-relax-d), $\eta_1 = 0$ is one of the optimal solutions iff $\eta_2^* = 1_n/n$. Thus if $\eta_2^* \neq 1_n/n$, we get a contradiction. Otherwise, we can force the dual problem to find a solution with $\eta_1 \neq 0$ by introducing an additional constraint $\eta_1^\top 1_n \geq \delta$ for some constant $0 < \delta \ll 1$. □

Chapter 3

Relative regret optimization in portfolio choice under parameter uncertainty

3.1 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 which may be small). We consider a single period portfolio choice problem where the agent is endowed with a finite sample of historical returns data and makes a one shot allocation decision after observing the last data point. Parameter estimation is possible since data is generated iid from a known class of models with constant but unknown parameters, while concern for robustness is legitimate since the data window may be small and convergence of the learning model may be slow. We would like to understand how parameter estimation/learning and robust decision making can be combined in this setting.

Portfolio selection with parameter uncertainty is commonly formulated in a Bayesian framework where the investor makes an exogenous specification of the prior and maximizes expected utility of terminal wealth. One challenge with the Bayesian approach, which has substantial normative implications, is that the solution can be very sensitive to the specification of the prior[1]. 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 or even moderate, while a uniform prior (the "obvious" choice when the investor has no prior information) gives the solution of the classical Merton/Markowitz problem

[1]Of course, this is a desirable feature if the investor has strong prior views that he wishes to incorporate in the optimization problem.

with the sample mean of realized returns plugged-in for the true expected return, and it is well known that this strategy performs badly [12, 56]. More generally, the posterior distribution for the mean return converges to the truth at the same rate as the sample mean, so a poor choice takes a long time to correct itself.

In this chapter, we address the problem of portfolio choice with parameter uncertainty by adopting a relative regret objective. 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 benchmarks are wealth of fictitious investors who behave optimally given knowledge of the parameter value, and the goal is to maximize the worst case *relative* performance with respect to this family of benchmarks. Early results on this notion can be found in [11, 32, 57] while more recent analysis of problems with this objective include [7, 43, 44, 53]. Also related is the work on universal portfolios ([19]), though its focus is the asymptotic regret of certain online policies, as opposed to finite horizon results which is the focus of this chapter (see also [9, 15, 24]). More generally, though portfolio choice with parameter uncertainty and model ambiguity has attracted substantial attention in recent years (see for example [26, 28, 31, 46, 58, 60]), most of this literature adopts an absolute worst case, as opposed to a relative performance/regret, objective. One exception is [43] which we discuss in more detail below.

A second feature of our model is that a finite sample of historical returns (assumed iid normal but with unknown mean and variance) is available to the investor when he/she makes a decision. Historical data has value because it can be used to learn model parameters, and our regret optimizing investor being endowed with a moderate amount of data lives somewhere between the highly non-stationary data-absent world in which decisions need to be made with robustness being a primary concern, at the one extreme, and the stationary data abundant world in which parameters can be learned with little statistical uncertainty, at the other. The analysis of his/her problem is interesting to us because it gives insight into the tradeoff between robust decision making, data driven optimization and learning for the world "in between". Our model puts no restrictions on how the investor uses the data, and one contribution of this chapter is to characterize both the investment decision and update rule for learning parameters that are optimal in this setting.

We analyze our regret problem using convex duality, and show that the optimal learning model involves Bayesian updating of the Lagrange multiplier that solves the dual problem, which plays the role of the prior. We also show that the optimal portfolio is the solution of an interesting (though non-standard) Bayesian portfolio choice problem where 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 posterior. 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. We mention here that probabilistic interpretations of Lagrange multipliers in the context of regret are also made in [53] and [43], though neither considers learning or updating because the decision maker has no data or has already incorporated it in the uncertainty set ([53]) or lives in a highly non-stationary world in which even recent data has no value ([43]). The problem of combining robustness with learning is also discussed in [33, 34] and [41] though from a worst case perspective.

Our relative regret objective can be interpreted as a loss function that evaluates an estimator in the context of the application and the associated decision making goals. The idea that inference and decision making should go hand in hand goes back at least to [62] who states

> *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*

(see also [49]). Crudely stated, a poor estimator (according to some loss function) may be perfectly acceptable if it can be combined with an investment decision that consistently delivers large profits, while an "optimal" estimator is of little value if decisions using the resulting estimates consistently perform poorly.

The outline of our chapter is as follows. We formulate the single period market model in Section 3.2, and introduce two relative regret objectives in Section 3.3. The first of these is more standard while the second (which is the major focus of this chapter) is original and can be interpreted as an objective-based loss function. We establish connections between our relative regret model and Bayesian problems in Section 3.4 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 Lagrange multiplier/prior characterized as the solution of the associated dual problem. Computational studies are provided in Section 3.5.

The single period model in this chapter 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.

3.2 Model

3.2.1 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 as a geometric Brownian motion with constant drift and volatility, as in the classical Black-Scholes model of stock prices [10]. Specifically,

$$S_i(t) = S_i(0) \exp\Big\{\mu_i t + \sigma_i' W(t)\Big\}, \quad i = 1, 2, \cdots, n, \tag{3.1}$$

where $W(t)$ is an n-dimensional standard Brownian motion, the scalar μ_i is the (constant) 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 (constant) volatility of this stock. 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

$$Q = \sigma\sigma' \geq \delta I$$

holds for some constant $\delta > 0$. The column vector of stock prices is denoted by

$$S(t) = \big[\ S_1(t), \quad S_2(t), \quad \cdots, \quad S_n(t)\ \big]'.$$

Investor's observations/data

We assume in this chapter that the parameters $H = (\sigma\sigma', \mu)$ for the stock price model (3.1) are constant, that the investor knows the model family (3.1) but 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 seen 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 $[t\delta,\ (t+1)\delta)$:

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

where

$$Z(t+1) \triangleq \frac{W((t+1)\delta) - W(t\delta)}{\sqrt{\delta}}.$$

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

$$\mathcal{R}(t+1) \sim N(\delta\mu,\ \delta Q) \tag{3.2}$$

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 (3.1), but does not know the parameter values $(\sigma,\ \mu)$. 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. More formally, let

$$\mathcal{G}_T \triangleq \sigma\{\mathcal{R}(1),\ \mathcal{R}(2), \cdots,\ \mathcal{R}(T)\} = \sigma\{S(0),\ S(\delta), \cdots,\ S(T\delta)\} = \sigma\{\underline{S}(t\delta)\}$$

denote the σ-algebra generated by relative returns $\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T)$ (equivalently, stock prices $S(0), S(\delta), S(2\delta), \cdots, S(T\delta)$). 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 use π_i to denote the proportion of wealth invested in stock i and assume the rest $1 - \pi'\mathbf{1}$ is invested in the risk free asset. 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

$$\begin{aligned} x_\pi^H(T+1) &= \sum_{i=1}^n x\pi_i e^{\mathcal{R}_i(T+1)} + x\Big(1 - \sum_{i=1}^n \pi_i\Big)e^{r\delta} \\ &= x\Big\{\sum_{i=1}^n \pi_i e^{\delta\mu_i + \sqrt{\delta}\sigma_i Z(T+1)} + \Big(1 - \sum_{i=1}^n \pi_i\Big)e^{r\delta}\Big\}. \end{aligned} \tag{3.3}$$

Assuming δ is small relative to $r, ||\mu||_1$ and $||Q||_1$, one can show via Taylor expansions

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

where $b = [b_1, b_2, \cdots, b_n]'$ is an n-dimensional vector of real numbers $b_i \triangleq \mu_i + (1/2)\sigma_i\sigma_i' - r$ 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

$$x_\pi^H(T+1) = x \exp\Big\{\delta\Big[r + b'\pi - \frac{1}{2}\pi' Q\pi\Big] + \sqrt{\delta}\pi'\sigma Z(T+1)\Big\} \tag{3.4}$$

for the remainder of the chapter. Alternatively, (3.4) 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 $T\delta$ and $(T+1)\delta$ once it has been specified at $T\delta$.

Finally, since we will be dealing directly with the wealth equation (3.4) rather than the log-returns model (3.2), it is more convenient for us to talk about uncertainty in $H = (Q, b)$ 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 $\mathcal{H}$.

3.2.2 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

[2]This follows from the so-called log-linear approximation of the wealth equation (3.3), which becomes exact when $\delta \downarrow 0$ (see [14] 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 3.1 Summary of priors for Bayesian investors

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 (3.2) 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 3.1. 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 (see, for example, [8]). Each Bayesian investor then solved the following single period asset allocation problem

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

using their updated parameters. It can be shown that

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

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

$$\pi_{B,1}^* = 22.95, \quad \pi_{B,2}^* = 23.62, \quad \pi_{B,3}^* = 24.30$$

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 with Constant Relative Risk Aversion (CRRA) utility who knows the model parameters (μ, σ) and who solves

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

For our model parameters, $\psi^* = 23.4$.

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

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

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 [12, 56].

Consider now the following experiment. We generated $1,000,000$ samples of annual returns using the model (μ, σ) = (20%, 20%). For each sample we recorded the end-of-year wealth $x_i(1)$ of each of the Bayesian investors $\pi^*_{B,i}$ ($i = 1, 2, 3$) as well as the wealth $y(1)$ of the "knowledgeable" investor who invests according to ψ^*. Figures 3.2–3.4 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[3].

[3]Figures 3.2-3.4 may be puzzling to some in that investor 1 outperforms investors 2 and 3 relative to the benchmark, even though the mean of the prior chosen by investor 2 coincides with the mean of the data generating process (i.e. $\mu = 0.2$). This is because Figures 3.2-3.4 show relative performance *conditional on a particular realization of 10 years of data* and by chance, this realization was such that investor 1's portfolio outperformed those of the other two investors. For other data samples, the ordering will differ. The main point of this example is not that one particular investor outperforms the others, but rather, that relatively small differences in the prior can substantially affect performance.

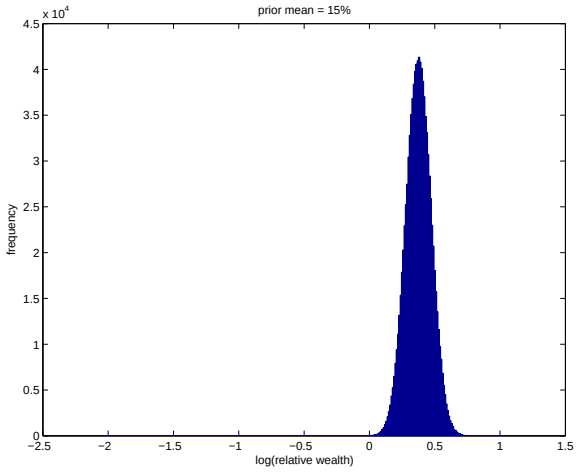

Figure 3.2 Histogram of log relative wealth of investor 1 relative to the knowledgeable investor.

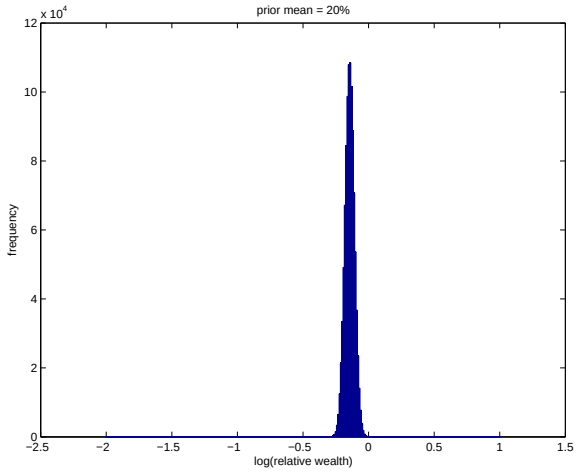

Figure 3.3 Histogram of log relative wealth of investor 2 relative to the knowledgeable investor.

3.2.3 Worst case model

Worst case models are commonly proposed when there is model uncertainty. A typical formulation of this problem is

$$\begin{cases} \max\limits_{\pi\in\mathcal{G}_T}\min\limits_{H\in\mathcal{H}}\mathbb{E}^H\left[\dfrac{1}{\gamma}x_\pi^H(\delta)^\gamma\right] \\ \text{subject to:} \\ x_\pi^H((\delta)) = x\exp\left\{\delta\left[r+b'\pi-\frac{1}{2}\pi'Q\pi\right]+\sqrt{\delta}\pi'\sigma Z(T+1)\right\}, \end{cases} \tag{3.5}$$

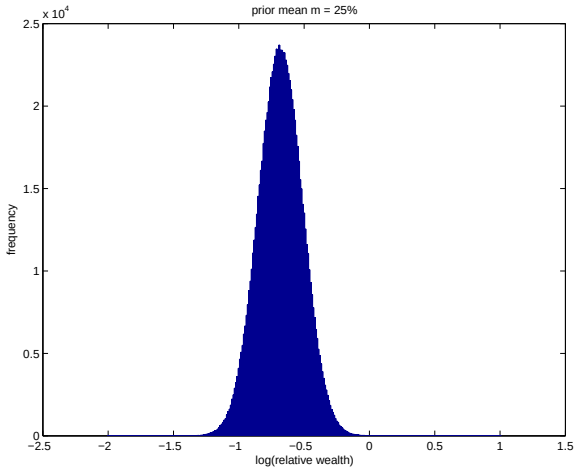

Figure 3.4 Histogram of log relative wealth for investor 3 relative to the knowledgeable investor.

for which the solution is[4]

$$\pi^* = \frac{1}{1-\gamma}[Q^*]^{-1}b^* \tag{3.6}$$

$$H^* = (Q^*,\, b^*) = \arg \min_{\mu \geq 0,\, \mu(\mathcal{H})=1} \mathbb{E}_\mu[b' Q^{-1} b] = \arg \min_{(Q,\, b) \in \mathcal{H}} b' Q^{-1} b. \tag{3.7}$$

The solution can be described as follows: (i) Find the model $(Q^*,\, b^*)$ in $\mathcal{H}$ with the smallest Sharpe ratio (i.e. equation (3.7)), and (ii) solve a standard Bayesian problem with a prior that puts all its mass on $(Q^*,\, b^*)$ (i.e. equation (3.6)).

The solution (3.6)-(3.7) 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 feature 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. Secondly, since the worst case prior (3.7) is degenerate, then so too is the posterior. That is, the "worst case" equilibrium portfolio (3.6) 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 the remainder of the chapter, we consider relative regret as a framework for formulating portfolio selection problems with parameter uncertainty. In contrast to

[4]In writing the equality in (3.7) we mean that the solution of the first optimization problem (over probability measures with support $\mathcal{H}$) is degenerate with mass 1 on the solution $(Q^*,\, b^*) \in \mathcal{H}$ of the second.

the standard worst case approach, relative regret favors decisions that perform well in both pessimistic (low Sharpe ratio) and optimistic (large Sharpe ratio) scenarios and will use (rather than ignore) data to infer the model parameters so as to increase the likelihood that it performs reasonably well in all cases. That being said, one feature of our model is that there is no prior specification and the update rule is not imposed. Rather, we are interested in understanding the learning model that comes up as part of the optimal solution.

3.3 Relative regret

In this section, we formulate two portfolio optimization problems within the setup described in Section 3.2 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 [59]) while the second is our own. A key feature of 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 3.4) is the characterization of the learning model associated with the optimal solution of the problems formulated in this section. We refer the reader to [43] for a version of this chapter where learning is not possible.

3.3.1 Relative regret I: standard model

Consider the problem

$$\max_{\pi\in\mathcal{G}_T}\min_{H\in\mathcal{H}}\frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\max\limits_{\psi}\mathbb{E}_H U(y_\psi^H(\delta))}. \tag{3.8}$$

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) cannot 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 (3.4), and the expected utility $\mathbb{E}_H U(x_\pi^H(\delta))$ in the

numerator of (3.8) can be computed. In addition, the denominator is calculated by optimizing over ψ given knowledge of the model H chosen by nature:

$$\begin{cases} \mathbb{E}_H U(y^H(\delta)) = \max\limits_{\psi} \mathbb{E}_H U(y^H_\psi(\delta)) \\ \text{subject to:} \\ y^H_\psi(\delta) = y(0)\exp\Big\{\delta\Big[r + b'\pi - \frac{1}{2}\psi' Q\psi\Big] + \sqrt{\delta}\psi'\sigma Z(T+1)\Big\}. \end{cases} \tag{3.9}$$

The ratio of these two quantities is precisely the relative regret objective (3.8).

- The investor chooses the policy $\pi \equiv f(\cdot)$ and nature the model H to satisfy the equilibrium condition (3.8).

An axiomatic justification for this objective is given in [59]; see also [36].

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

A more explicit computation can be done for the model (3.8) if the utility function is assumed to be CRRA. More specifically, observing that $y(\delta)$ is a log-normal random variable given by (3.9), it follows that

$$\begin{aligned} &\frac{1}{\eta}\mathbb{E}_H\left(y^H_\pi(\delta)\right)^\eta \\ &= \frac{1}{\eta}y(0)^\eta \mathbb{E}_H \exp\Big\{\eta\delta\Big[r + b'\pi - \frac{1}{2}\psi' Q\psi\Big] + \eta\sqrt{\delta}\psi'\sigma Z(T+1)\Big\} \\ &= \frac{1}{\eta}y(0)^\eta \exp\Big\{\eta\delta\Big[r + b'\psi - \frac{1-\eta}{2}\psi' Q\psi\Big]\Big\} \\ &= \frac{1}{\eta}y(0)^\eta \exp\Big\{\delta\eta\Big[r - \frac{1-\eta}{2}\Big(\psi - \frac{Q^{-1}b}{1-\eta}\Big)' Q\Big(\psi - \frac{Q^{-1}b}{1-\eta}\Big) + \frac{1}{2(1-\eta)}b'Q^{-1}b\Big]\Big\}, \end{aligned} \tag{3.10}$$

(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 (3.9) using (3.10)) is given by

$$\psi^H = \arg\max_{\psi} \frac{1}{\eta}\mathbb{E}_H U(y^H_\pi(\delta)) = \frac{1}{1-\eta}Q^{-1}b \tag{3.11}$$

when the model is $H = (Q, b)$. Substituting ψ^H into the wealth equation in (3.9) it follows that the benchmark investor's optimal wealth is given by

$$y^H(\delta) = y(0)\exp\Big\{\delta\Big[r + \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] + \frac{\sqrt{\delta}}{1-\eta}b'Q^{-1}\sigma Z(T+1)\Big\} \tag{3.12}$$

and the denominator of (3.8) (from substituting (3.11) into (3.10)) is

$$\mathbb{E}_H U(y^H(\delta)) = \frac{1}{\eta} y(0)^\eta \exp\left\{\delta\eta\Big[r + \frac{1}{2(1-\eta)} b' Q^{-1} b\Big]\right\}. \tag{3.13}$$

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 (3.8)) 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

$$\mathbb{E}_H U(x_\pi^H(\delta)) = \frac{1}{\eta} x(0)^\eta \mathbb{E}_H \exp\left\{\delta\eta\Big[r + b'\pi - \frac{1-\eta}{2}\pi' Q\pi\Big]\right\}. \tag{3.14}$$

Substituting (3.13) and (3.14) into the relative regret objective in (3.8), we obtain

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

and it follows that (3.8) is equivalent to

$$\begin{aligned} &\max_{\pi\in\mathcal{G}_T} \min_{H\in\mathcal{H}} \frac{\mathbb{E}_H U(x_\pi^H(\delta))}{\max\limits_\psi \mathbb{E}_H U(y_\psi^H(\delta))} \\ &= \max_{\pi\in\mathcal{G}_T} \min_{H\in\mathcal{H}} \mathbb{E}_H \exp\left\{\delta\eta\Big[b'\pi - \frac{1-\eta}{2}\pi' Q\pi - \frac{1}{2}\frac{b' Q^{-1} b}{1-\eta}\Big]\right\}. \end{aligned} \tag{3.15}$$

It does not appear that an explicit expression for the equilibrium solution of (3.15) is possible. However, we show in Section 3.4 that an approximate solution, which becomes exact when $\delta \downarrow 0$, can be obtained.

3.3.2 Relative regret II: objective based loss function

In this section we introduce an alternative relative regret problem. We adopt the same (approximate) wealth equation (3.4) for the investor and the same definition (3.9) 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 (3.9), 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

$$\begin{cases} \mathbb{E}_H U^B(y^H(\delta)) \equiv \max\limits_{\psi} \mathbb{E}_H U^B(y^H_\psi(\delta)) \\ \text{subject to:} \\ \quad y^H_\psi(\delta) = y(0) \exp\Big\{\delta\Big[r + b'\psi - \frac{1}{2}\psi' Q \psi\Big] + \sqrt{\delta}\psi'\sigma Z(T+1)\Big\}. \end{cases} \tag{3.16}$$

Relative regret problem

Consider the following relative regret problem

$$\begin{cases} \max\limits_{\pi\in\mathcal{G}_T} \min\limits_{H\in\mathcal{H}} \mathbb{E}_H\Big[U\Big(\frac{x^H_\pi(\delta)}{y^H(\delta)}\Big)\Big] \\ \text{subject to:} \\ \quad x^H_\pi(\delta) \text{ is given by (3.4)} \\ \quad y^H(\delta) \text{ is defined via (3.16).} \end{cases} \tag{3.17}$$

The key difference between (3.17) and the relative regret problem (3.8) is the way that $x^H_\pi(\delta)$ 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 (3.17) 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 (3.16) 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 wealth of the investor $x^H_\pi(\delta)$ (given by (3.4)) and the benchmark $y^H(\delta)$ (given by (3.16)) are realized. Conditional on $\pi = f(\mathcal{R}(1), \mathcal{R}(2), \cdots, \mathcal{R}(T))$ and H, the distribution of the ratio $x^H_\pi(\delta)/y^H(\delta)$ is fully characterized, and we can calculate

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

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

- The investor and nature choose π and H to satisfy the equilibrium condition associated with (3.17).

We reiterate that the essential difference between the models (3.8) and (3.17) is the way that $x_\pi^H(\delta)$ and $y^H(\delta)$ are compared. In (3.8) they are compared by evaluating the ratio of their expected utilities while in (3.17) 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 (3.8), we are not aware of an axiomatic foundation for (3.17) though we believe that this is an issue worth pursuing. On the other hand, it will be shown that the solution of (3.8) is a limiting case of (3.17) when the utility/comparison functions are CRRA/power type. Another advantage of (3.17) 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 (3.17) to multi-period problems that is relatively easy to analyze (this will be done elsewhere). The same cannot be said about (3.8).

Objective based loss function: Suppose we have a set of models

$$\mathcal{M}(\mathcal{H}) = \{M(H),\, H \in \mathcal{H}\},$$

parameterized by $H \in \mathcal{H}$ and data $(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ generated from one of the models $M(H^*)$ in this family. A classical problem in statistics is to estimate the unknown parameter H^*. The quality of an estimator $\hat{H}_T = g(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ of H^* is evaluated using a loss function $\mathcal{L}(g(\mathcal{R}_1, \cdots, \mathcal{R}_T),\, H)$ and some criterion; for example, the min-max criterion is

$$\min_g \max_{H \in \mathcal{H}} \mathbb{E}_H[\mathcal{L}(g(\mathcal{R}_1, \cdots, \mathcal{R}_T),\, H)]. \tag{3.18}$$

The objective in (3.17) can be interpreted as a min-max criterion for an *objective based loss function* $U(x_\pi^H(\delta),\, y^H(\delta)) \equiv U(\frac{x_\pi^H(\delta)}{y^H(\delta)})$ which we now describe. A fundamental difference is that our primary concern is performance of a *decision* $\pi = f(\mathcal{R}_1, \cdots, \mathcal{R}_T)$, that may involve an estimate of the parameters somewhere in its definition, rather than the quality of the estimator itself (though of course, both issues are related).

The components of our model have analogs with (3.18) as follows:

- The optimal (random) wealth $y^H(\delta)$ for every given parameter value in (3.17) is analogous to the parameter value H in the loss function $\mathcal{L}(g(\mathcal{R}_1, \cdots, \mathcal{R}_T),\, H)$. The benchmark in the parameter estimation problem (3.18) is H whereas our benchmark is the optimal wealth for the data generating parameter;

- Investment decisions $\pi = f(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ are analogous to the estimator $\hat{H} = g(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ in (3.18);

- While the objective in (3.18) compares the closeness of an estimator $g(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ to the parameter H, we compare the performance $x_\pi^H(\delta)$ of the mapping $\pi = f(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ to the optimal performance $y^H(\delta)$ for each parameter.

- We are concerned with the distribution of the investor performance $x_\pi^H(\delta)$ relative to $y^H(\delta)$, represented by the comparison function $U(x,\, y) = \frac{1}{\gamma}(x/y)^\gamma$, which is analogous to concern about the distribution of $g(\mathcal{R}_1, \cdots, \mathcal{R}_T)$ as expressed through the choice of loss function in (3.18). In this regard, classical regret (3.8) is different as it compares average performance instead of performance for each realization.

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

Let us now consider the relative regret problem (3.17) under the assumption that the benchmark investor's utility function as well as the comparison function are of 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 (3.11)-(3.12), the benchmark investor's problem (3.16) with a CRRA utility has an explicit solution

$$\psi^H = \frac{1}{1-\eta}Q^{-1}b \tag{3.19}$$

and

$$y^H(\delta) = y(0)\exp\Big\{\delta\Big[r + \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] + \frac{\sqrt{\delta}}{1-\eta}b'Q^{-1}\sigma Z(T+1)\Big\} \tag{3.20}$$

is the associated benchmark investor's wealth. It now follows that the normalized wealth process $z_\pi^H(\delta)$ satisfies

$$\begin{aligned} z_\pi^H(\delta) = \frac{x_\pi^H(\delta)}{y^H(\delta)} \;\; &= \;\; \frac{x(0)}{y(0)}\exp\Big\{\delta\Big[b'\pi - \frac{1}{2}\pi'Q\pi - \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] \\ &\qquad + \sqrt{\delta}\Big[\pi - \frac{1}{1-\eta}Q^{-1}b\Big]'\sigma Z(T+1)\Big\} \end{aligned}$$

so the relative regret problem (3.17) becomes

$$\left\{\begin{array}{l} \max\limits_{\pi\in\mathcal{G}_T}\min\limits_{H\in\mathcal{H}}\mathbb{E}_H[U(z_\pi^H(\delta))] \\ \text{subject to:} \\ \quad z_\pi^H(\delta) = z(0)\exp\Big\{\delta\Big[b'\pi - \frac{1}{2}\pi'Q\pi - \frac{1-2\eta}{2(1-\eta)^2}b'Q^{-1}b\Big] \\ \qquad\qquad + \sqrt{\delta}\Big[\pi - \frac{1}{1-\eta}Q^{-1}b\Big]'\sigma Z(T+1)\Big\} \end{array}\right. \tag{3.21}$$

where we have substituted $z_\pi^H(\delta)$ for $\frac{x_\pi^H(\delta)}{y^H(\delta)}$. Choosing $U(z) = \frac{1}{\gamma}z^\gamma$ as the comparison function in (3.21), it follows from the log-normality of $z_\pi^H(\delta)$ (conditional on $\mathcal{G}_T$ and H) that the conditional expectation

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

It now follows that (3.21) is equivalent to

$$\max_{\pi\in\mathcal{G}_T}\min_{H\in\mathcal{H}}\frac{z(0)^\gamma}{\gamma}\mathbb{E}_H\Big\{\exp\Big[\delta\gamma\Big(\frac{1-\eta-\gamma}{1-\eta}b'\pi - \frac{1-\gamma}{2}\pi' Q\pi - \frac{1-2\eta-\gamma}{2(1-\eta)^2}b'Q^{-1}b\Big)\Big]\Big\}. \tag{3.22}$$

It is interesting to note the similarity between (3.15) and (3.22). We shall expand further on this in later sections.

3.4 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 $\delta \downarrow 0$ that shows the explicit dependence of the optimal portfolio on data.

3.4.1 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 $\mathcal{H} = \{H_1, \cdots, H_m\} = \{(b_1, Q_1), \cdots, (b_m, Q_m)\}$ and utility and comparison functions are power-type, that (3.17) is equivalent to

$$\begin{cases} \nu^* = \max\limits_{\pi,\,\kappa} \kappa \\ \text{subject to:} \\ \mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] \geq \kappa, \quad i = 1, \cdots, m. \end{cases} \tag{3.23}$$

For every model, (3.22) gives

$$\begin{aligned} &\mathbb{E}_{H_i}\Big[\frac{1}{\gamma}\Big(\frac{x_\pi^{H_i}(\delta)}{y^{H_i}(\delta)}\Big)^\gamma\Big] \\ &= \frac{z(0)^\gamma}{\gamma}\mathbb{E}_{H_i}\Big\{\exp\Big[\delta\gamma\Big(\frac{1-\eta-\gamma}{1-\eta}b_i'\pi - \frac{1-\gamma}{2}\pi' Q_i \pi - \frac{1-2\eta-\gamma}{2(1-\eta)^2}b_i' Q_i^{-1} b_i\Big)\Big]\Big\}. \end{aligned} \tag{3.24}$$

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 (3.23) is a convex optimization problem in $(\pi,\, \kappa)$ when this condition holds.

Let $\mu_i \geq 0$ denote the Lagrange multiplier for the i^{th} constraint in (3.23). Clearly, if $(\pi,\, \kappa)$ is feasible for (3.23) 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

$$\psi(\mu) = \max_{\pi \in \mathcal{G}_T} \mathbb{E}_\mu \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big], \tag{3.25}$$

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 (3.23). The following result is an immediate consequence of Lagrangian duality (see Theorem 1 on pp. 224 of [45]). 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 2. *Suppose that $\mathcal{H}$ is compact, that $\gamma < 0$ and $\eta < 1$. Let ν^* denote the optimal value of the relative regret problem (3.17) (or (3.23)), and $\psi(\mu)$ be the value function (dual function) for the Bayesian problem (3.25) when the prior (Lagrange multiplier) is μ. Then dual optimization problem*

$$\psi(\mu^*) = \min_{\mu \geq 0,\, \mu(\mathcal{H})=1} \psi(\mu) \tag{3.26}$$

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 (3.25) under μ^*, namely

$$\pi^* = \arg\max_{\pi \in \mathcal{G}_T} \mathbb{E}_{\mu^*} \Big[\frac{1}{\gamma} \Big(\frac{x(\delta)}{y(\delta)} \Big)^\gamma \Big] \tag{3.27}$$

Proposition 2 tells us that the solution of the relative regret problem is also the solution of a Bayesian problem (3.27) where the optimal Lagrange multiplier μ^* characterized by (3.26) 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 (3.26), 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

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

where

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

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 3. *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 (3.26). 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*

$$\pi^* = \arg\max_{\pi\in\mathcal{G}_T} \mathbb{E}_{\mu_T^*}\Big[\frac{1}{\gamma}\Big(\frac{x(\delta)}{y(\delta)}\Big)^\gamma\Big] \tag{3.28}$$

is optimal for the relative regret problem (3.17).

We mention here that probabilistic interpretations of Lagrange multipliers in the context of regret are also made in [53] and [43], though neither considers learning or parameter estimation because the decision maker has no data or has already incorporated it in the uncertainty set ([53]) or lives in a highly non-stationary world in which even recent data has no value ([43]).

3.4.2 Optimal portfolio: approximate solution

Although we have characterized the optimal portfolio (3.28) and optimal dual variable/prior (3.26), 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 (3.28) 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 (3.26) 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

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

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

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

The dual function (3.25) can now be written as

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

With this in mind, define

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

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

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

By completing the square, the above is equivalent to

$$\begin{aligned}\bar{\psi}(\mu) &= -\mathbb{E}_\mu\Big\{\min_{\pi\in\mathcal{G}_T}\Big[\pi-C_1\left[\mathbb{E}_{\mu_T}(Q)\right]^{-1}\mathbb{E}_{\mu_T}(b)\Big]'\mathbb{E}_{\mu_T}(Q)\Big[\pi-C_1\left[\mathbb{E}_{\mu_T}(Q)\right]^{-1}\mathbb{E}_{\mu_T}(b)\Big]\Big\}\\ &\quad+C_1^2\,\mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\}-C_2\mathbb{E}_\mu(b'Q^{-1}b)\end{aligned}$$

where $C_1=\frac{1-\gamma-\eta}{(1-\gamma)(1-\eta)}$ and $C_2=\frac{1-2\eta-\gamma}{(1-\eta)^2(1-\gamma)}$. Clearly,

$$\bar{\pi}^\mu \triangleq C_1[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b) \tag{3.29}$$

is the optimal portfolio, from which it follows that

$$\bar{\psi}(\mu) = C_1^2\,\mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\}-C_2\mathbb{E}_\mu(b'Q^{-1}b). \tag{3.30}$$

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 (3.26) can be obtained by solving

$$\bar{\mu}^* = \arg \min_{\mu \geq 0,\ \mu(\mathcal{H})=1} \bar{\psi}(\mu) \tag{3.31}$$

while

$$\bar{\pi}^* = C_1[\mathbb{E}_{\bar{\mu}_T^*}(Q)]^{-1}\mathbb{E}_{\bar{\mu}_T^*}(b) \tag{3.32}$$

is an approximate solution for the regret problem.

A similar calculation/approximation can be carried out for (3.15). In this case, the (approximate) dual problem is again given by (3.31) where

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

The (approximate) optimal policy is given by

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

Interestingly, this is the extreme case of $\gamma \to -\infty$ in (3.30) and (3.32), which coincides with a large aversion to missing the benchmark.

3.5 Example

In this section, we plot the approximate dual function (3.30) and solve for the approximate optimal prior and portfolio (3.31)-(3.32) for a simple example. We illustrate the effect of the number of data points on the dual function and the optimal prior.

For simplicity, we consider 2 assets and 3 models, and assume the covariance matrix is known. Thus our uncertainty set is $\mathcal{H} = \{(b_1, Q), \ldots, (b_3, Q)\}$, and we choose the following models for annualized mean returns

$$b_1 = \begin{pmatrix} 0.055 \\ 0.020 \end{pmatrix},\ b_2 = \begin{pmatrix} 0.035 \\ 0.010 \end{pmatrix},\ b_3 = \begin{pmatrix} 0.015 \\ 0.035 \end{pmatrix}$$

and the following annualized covariance matrix:

$$Q = \begin{pmatrix} 0.0087 & 0.0037 \\ 0.0037 & 0.0063 \end{pmatrix}.$$

We also set $\eta = -3,\ \gamma = -5,\ r = 0.03$ and observation frequency $\delta = 1/12$, which corresponds to monthly observations. T denotes the number of monthly return samples in the investor's data set.

To begin, observe that the approximate dual function (3.30) can be written as

$$\begin{aligned}
\bar{\psi}(\mu) &= C_1^2\ \mathbb{E}_\mu\Big\{\mathbb{E}_{\mu_T}(b)'[\mathbb{E}_{\mu_T}(Q)]^{-1}\mathbb{E}_{\mu_T}(b)\Big\} - C_2\mathbb{E}_\mu(b'Q^{-1}b) \\
&= C_1^2\ \mathbb{E}_\mu\left\{\left(\sum_{i=1}^3 \mu_T(i)b_i\right)' Q^{-1}\left(\sum_{i=1}^3 \mu_T(i)b_i\right)\right\} - C_2\sum_{i=1}^3 \mu_i(b_i'Q^{-1}b_i) \\
&= C_1^2\ \sum_{i=1}^3\sum_{j=1}^3 b_i'Q^{-1}b_j\mathbb{E}_\mu\left[\mu_T(i)\mu_T(j)\right] - C_2\sum_{i=1}^3 \mu_i(b_i'Q^{-1}b_i) \\
&= C_1^2\ \sum_{i=1}^3\sum_{j=1}^3\sum_{k=1}^3 b_i'Q^{-1}b_j\mu_k\mathbb{E}_\mu\left[\mu_T(i)\mu_T(j)|H_k \text{ is true}\right] - C_2\sum_{i=1}^3 \mu_i(b_i'Q^{-1}b_i)
\end{aligned} \tag{3.33}$$

where $\mu_T(i)$ is the posterior probability that the model H_i is correct after T observations of log-returns $\{\mathcal{R}(1),\ldots,\mathcal{R}(T)\}$:

$$\begin{aligned}
\mu_T(i) &:= p(H_i \text{ is true}|\mathcal{R}(1),\ldots,\mathcal{R}(T)) \\
&= \frac{\mu_i\ p(\mathcal{R}(1),\ldots,\mathcal{R}(T)|H_i \text{ is true})}{\sum_{l=1}^m \mu_l\ p(\mathcal{R}(1),\ldots,\mathcal{R}(T)|H_l \text{ is true})} \\
&= \frac{\mu_i \prod_{t=1}^T p(\mathcal{R}(t)|H_i \text{ is true})}{\sum_{l=1}^m \mu_l \prod_{t=1}^T p(\mathcal{R}(t)|H_l \text{ is true})} \\
&= \frac{\mu_i \exp\left\{-\frac{1}{2}\sum_{t=1}^T(\mathcal{R}(t)-\delta\nu_i)'(\delta Q)^{-1}(\mathcal{R}(t)-\delta\nu_i)\right\}}{\sum_{l=1}^m \mu_l \exp\left\{-\frac{1}{2}\sum_{t=1}^T(\mathcal{R}(t)-\delta\nu_l)'(\delta Q)^{-1}(\mathcal{R}(t)-\delta\nu_l)\right\}}
\end{aligned}$$

Now for a given μ, computing the second expression of $\bar{\psi}(\mu)$ in (3.33) is easy; for the first term we employ Monte Carlo integration — see Fig. (3.5) for the algorithm. We find the optimal prior by initially computing $\bar{\psi}(\mu)$ over the whole domain $\sum_{i=1}^3 \mu_i = 1,\ \mu_i \geq 0,\ i = 1,2,3$ coarsely discretized, and zooming in on the region of interest with higher degree of accuracy in both domain discretization and Monte Carlo integration.

We plot the approximate dual function surfaces in Fig. (3.7) for $T = 24, 36, 48$ and 60 observations, corresponding to $2-5$ years of monthly observations. Observe that the dual function surface becomes flatter with increasing number of observations. This makes sense intuitively; when the number of observations is small, the impact of learning is not significant and robustness concerns dominate. It may be worthwhile to focus on doing well relative to some models while ignoring others. However, with an increasing number of observations, learning takes over and it eventually becomes worthwhile to consider all models; for a large number of data points, any prior that weighs all the models performs reasonably well, hence the flatter surface. For this example, robustness means focusing on models 1 and 2 when there are only 24 data

Initialize $Nsim$, $Nint$, $Value = Nsim \times 1$ zero vector
for $csim = 1$ to $Nsim$ **do**
  for $cint = 1$ to $Nint$ **do**
    Generate $\mathbf{R} = \{\mathcal{R}(1), \ldots, \mathcal{R}(T)\}$ under model H_k
    $Value(csim) = Value(csim) + \mu_T(i; \mathbf{R})\mu_T(j; \mathbf{R})$
  end for
  $Value(csim) = \dfrac{1}{Nint} Value(csim)$
end for
return sample mean($Value$), sample std($Value$)

Figure 3.5 Algorithm for computing $\mathbb{E}_\mu\left[\mu_T(i)\mu_T(j)|H_k \text{ is true}\right]$ for given μ and i, j, k.

points (see Fig. 3.6). When there are ≥ 36 data points, learning becomes significant enough that the prior puts weight on all models. Hence for a substantial number of data points, the optimal prior still weights all models but the flatness of the dual function means that any prior that weights all of the models performs reasonably well.

In this regard, we compare the performance of the approximate optimal portfolio (3.32) corresponding to the approximate optimal prior $\bar{\mu}^*$ with the portfolio that corresponds to using a uniform prior $\mu_{unif} = [1/3, 1/3, 1/3]$. For the comparison, we simulate annualized portfolio log-return $12 \times \log(x_\pi^H(\delta)/x(0))$ using (3.4) under the three different models. We plot the mean log-returns with 1 std error bars in Fig. (3.8). The mean log-returns of the approximate solution performs better than or equal to the uniform prior across all three models and for all time periods. In addition, the mean log-return gap decreases with increasing learning periods T, which is consistent with the flattening of the dual function surface — that the performance of μ_{unif} increases over time.

For the final reported values, we discretize the domain by increments of 0.05 and use $Nsim = 100$ and $Nint = 5000$ for Monte Carlo integration. To examine the validity of the approximate solution, we also report objective values $\bar{\nu}^*$ for the (primal) relative regret problem (3.23) where $\pi = \bar{\pi}^*$ is the approximately optimal prior given by the solution of (3.32). Clearly, $\bar{\nu}^*$ is a lower bound to the true optimal regret value of (3.23). We note that the gap between $\bar{\nu}^*$ and the approximate dual function $\bar{\psi}(\cdot)$ evaluated at the optimal solution $\bar{\mu}^*$ of the approximate dual problem is always within 0.4% of the value of $\bar{\nu}^*$ which suggests that (3.31)-(3.32) is a good approximation of the exact problem (3.23).

T	$\bar{\mu}^*$	$\psi(\bar{\mu}^*)$ mean (std.)	$\bar{\nu}^*$ mean (std.)	Duality Gap at $\bar{\mu}^*$ (% of $\psi(\bar{\mu}^*)$)
24	$(0.421, 0.053, 0.526)'$	-0.2018 (0.0002)	-0.2024 (0.0003)	0.0006 (0.3%)
36	$(0.368, 0.158, 0.474)'$	-0.2016 (0.0002)	-0.2022 (0.0003)	0.0006 (0.3%)
48	$(0.316, 0.211, 0.474)'$	-0.2014 (0.0002)	-0.2020 (0.0002)	0.0006 (0.3%)
60	$(0.208, 0.313, 0.480)'$	-0.2012 (0.0002)	-0.2018 (0.0002)	0.0006 (0.3%)

Figure 3.6 Optimal prior computation for the approximate dual problem.

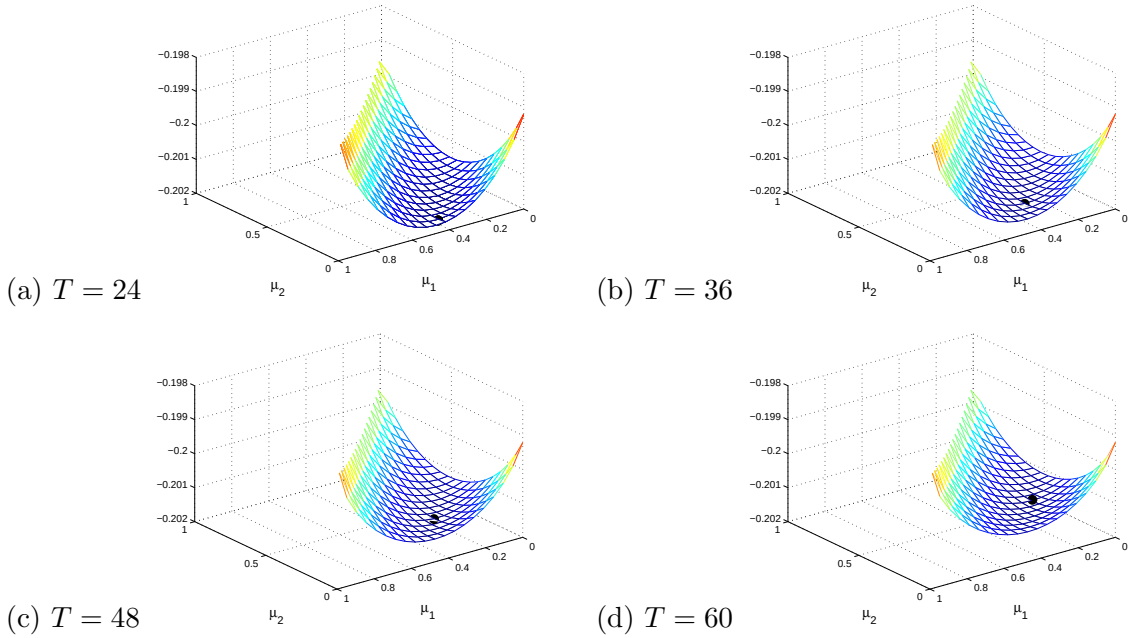

(a) $T = 24$ (b) $T = 36$

(c) $T = 48$ (d) $T = 60$

Figure 3.7 Dual function surfaces to first order δ for different values of T. The minimum value is shown by a filled black circle.

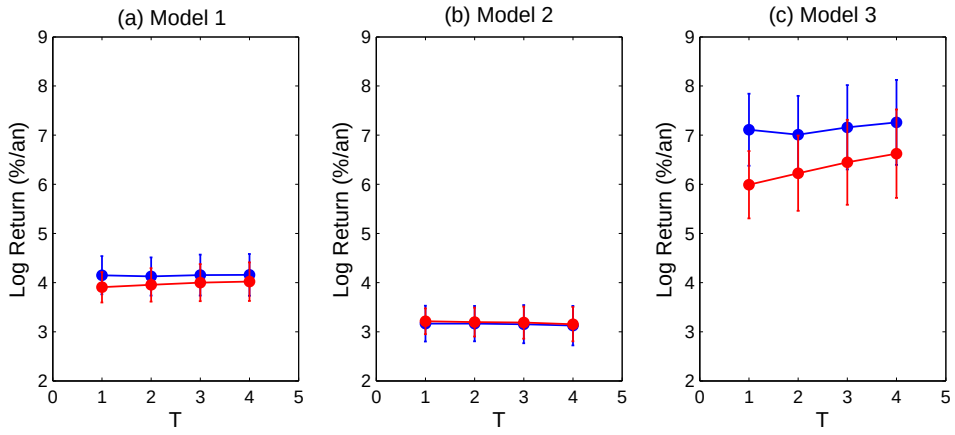

Figure 3.8 Annualized portfolio mean log-returns of the approximate optimal prior (blue) and the uniform prior (red) for the three different models. Error bars indicate 1 std error. The mean log-returns of the approximate solution performs no worse than the uniform prior for all time periods. In addition, the mean log-return gap decreases with increasing learning periods T, which is consistent with the observation of flattening dual function surface in the previous paragraph.

3.6 Appendix

Proof of Proposition 2 The problem (3.17) is equivalent to

$$\begin{cases} \max_{\pi\in\mathcal{G}_T,\,\kappa}\,\kappa \\ \text{subject to:} \\ \quad \mathbb{E}_H\left(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\right)^\gamma \geq \kappa, \quad \forall\, H\in\mathcal{H}, \end{cases} \tag{3.34}$$

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 [45].

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 [22]. 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[5]. 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 $g>0$

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

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:

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

see for example Section IV.6.3 in Dunford and Schwartz [22]. 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 [45]) and is equal to the subset of $\mathcal{C}^*(\mathcal{H})$ consisting of positive measures:

$$\mathcal{P}^* = \{\mu \in C^*(\mathcal{H}) \,|\, \mu(A) \geq 0,\ \forall\, A \in \mathcal{B}(\mathcal{H})\}. \tag{3.35}$$

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

Let $\mu \in \mathcal{P}^*$ be arbitrary and $(\pi,\, \kappa)$ be feasible for (3.34). It is clear that

$$\begin{aligned} L(\pi,\, \kappa,\, \mu) &\triangleq \kappa + \int_{H\in\mathcal{H}} \Big[\mathbb{E}_H\Big(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma - \kappa\Big]\mu(dH) \\ &= \kappa(1-\mu(\mathcal{H})) + \int_{H\in\mathcal{H}} \mathbb{E}_H\Big(\frac{1}{\gamma}\frac{x_\pi^H(\delta)}{y^H(\delta)}\Big)^\gamma \mu(dH) \\ &\geq \nu^*. \end{aligned} \tag{3.36}$$

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 (3.25). It now follows from Theorem 1 on pg. 224 of [45] 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 (3.27) with prior μ^* is also the optimal solution of the relative regret problem (3.34).

Chapter 4

Conclusion

In this thesis, we develop two new approaches to robustness and learning in data-driven portfolio optimization.

In Chapter 2, we investigate performance-based regularization (PBR) to reduce estimation risk in empirical mean-CVaR portfolio optimization. The nonparametric PBR method solves the empirical mean-CVaR problem with penalties on the uncertainties in mean and CVaR estimations. The parametric PBR method solves the empirical Markowitz problem instead if the underlying model is elliptically distributed. Both theoretical analysis and simulation experiments show the PBR methods improve upon the naive approach to data-driven mean-CVaR portfolio optimization.

In Chapter 3, we combine learning and robust portfolio allocation via solving a relative regret objective. We show using convex duality that the optimal learning model is Bayesian where the prior is endogenously specified and corresponds to the Lagrange multipliers that solve the dual problem. The optimal investment portfolio is the solution of a non-standard Bayesian problem where the posterior is obtained by Bayesian updating of the optimal dual variables/prior using Bayes' rule. The problem of minimizing relative regret can be interpreted as one of minimizing a loss function in which estimators are evaluated, not by the associated statistical errors, but by the performance of investment policies that use these estimators relative to a benchmark. Our results can be interpreted as a characterization of the optimal estimator for this loss function.

From a larger perspective, both PBR and relative regret optimization are new and promising methods of dealing with estimation risk and model uncertainty while learning from data. While this thesis has focused on the portfolio optimization problem, both methods can be adapted to other problems that involve data-driven optimization with large uncertainty. We leave this exploration for future work.

Bibliography

[1] C. Acerbi and D. Tasche. On the coherence of expected shortfall. *Journal of Banking and Finance*, 26:1487–1503, 2001.

[2] S. Alexander, T.F. Coleman, and Y. Li. Minimizing CVaR and VaR for a portfolio of derivatives. *Journal of Banking & Finance*, 30(2):583–605, 2006.

[3] T.W. Anderson. *An introduction to multivariate statistical analysis*. John Wiley & Sons, 1958.

[4] F. Andersson, H. Mausser, D. Rosen, and S. Uryasev. Credit risk optimization with conditional value-at-risk criterion. *Mathematical Programming*, 89(2):273–291, 2001.

[5] M.A. Arcones and E. Gine. Limit theorems for u-processes. *The Annals of Probability*, 21(3):1494–1542, 1993.

[6] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

[7] D. Bergemann and K. Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 2011.

[8] J.O. Berger. *Statistical decision theory and Bayesian analysis*. Springer, 1985.

[9] O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.

[10] F. Black and M. Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, pages 637–654, 1973.

[11] D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.

[12] M.W. Brandt. Portfolio choice problems. *Handbook of financial econometrics*, 1:269–336, 2009.

[13] M. Broadie. Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, 45(1):21–58, 1993.

[14] J.Y. Campbell and L.M. Viceira. *Strategic asset allocation: portfolio choice for long-term investors*. OUP Oxford, 2002.

[15] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge Univ Pr, 2006.

[16] A. Charnes, W. W. Cooper, and G. H. Symonds. Cost horizons and certainty equivalents: An approach to stochastic programming of heating oil. *Management Science*, 4(3):pp. 235–263, 1958.

[17] S.X. Chen. Nonparametric estimation of expected shortfall. *Journal of Financial Econometrics*, 6(1):87, 2008.

[18] V.K. Chopra. Improving Optimization. *The Journal of Investing*, 2(3):51–59, 1993.

[19] T.M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.

[20] E.G. De Giorgi. A Note on Portfolio Selection under Various Risk Measures. *SSRN eLibrary*, 2002.

[21] V. DeMiguel, L. Garlappi, F.J. Nogales, and R. Uppal. A Generalized Approach to Portfolio Optimization: Improving Performance by Constraining Portfolio Norms. *Management Science*, 55(5):798, 2009.

[22] N. Dunford and J.T. Schwartz. *Linear Operators: Vol.: 1: General Theory*. Interscience Publishers, Incorporated, 1967.

[23] N. El Karoui. High-dimensionality effects in the markowitz problem and other quadratic programs with linear constraints: Risk underestimation. *The Annals of Statistics*, 38(6):3487–3566, 2010.

[24] D.P. Foster and R. Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, 29(1-2):7–35, 1999.

[25] P.A. Frost and J.E. Savarino. For better performance: Constrain portfolio weights. *The Journal of Portfolio Management*, 15(1):29–34, 1988.

[26] L. Garlappi, R. Uppal, and T. Wang. Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Review of Financial Studies*, 20(1):41–81, 2007.

[27] D. Goldfarb and G. Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38, 2003.

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

[29] J. Gotoh and K. Shinozaki. A Robust Portfolio Technique for Mitigating the Fragility of CVaR Minimization. Technical report, Dept. of Industrial and Systems Engineering, Chuo University, 2010.

[30] Jun-ya Gotoh and Akiko Takeda. On the role of norm constraints in portfolio selection. *Computational Management Science*, 8(4):323–353, 2010.

[31] A. Gundel. Robust utility maximization for complete and incomplete market models. *Finance and Stochastics*, 9(2):151–176, 2005.

[32] J. Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.

[33] L.P. Hansen and T.J. Sargent. Robust estimation and control under commitment. *Journal of Economic Theory*, 124(2):258–301, 2005.

[34] L.P. Hansen and T.J. Sargent. Recursive robust estimation and control without commitment. *Journal of Economic Theory*, 136(1):1–27, 2007.

[35] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer-Verlag, 2 edition, 2009.

[36] T. Hayashi. Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1):242–268, 2008.

[37] W. Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948.

[38] R. Jagannathan and T. Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1684, 2003.

[39] J.D. Jobson and B. Korkie. Estimation for Markowitz Efficient Portfolios. *Journal of the American Statistical Association*, 75(371):544–554, 1980.

[40] P. Jorion. Portfolio optimization in practice. *Financial Analysts Journal*, 48(1):68–74, 1992.

[41] T.A. Knox. Foundations for learning how to invest when returns are uncertain. Working paper, University of Chicago Graduate School of Business, 2003.

[42] A.E.B. Lim, J.G. Shanthikumar, and G.-Y. Vahn. Conditional Value-at-Risk in portfolio optimization: coherent but fragile. *Operations Research Letters*, 39(3):163 – 171, 2011.

[43] A.E.B. Lim, J.G. Shanthikumar, and T. Watewai. Robust asset allocation with benchmarked objectives. *Mathematical Finance*, 21(4):643–679, 2011.

[44] Ruben Lobel and Georgia Perakis. Dynamic Pricing Through Sampling Based Optimization. *SSRN eLibrary*, 2010.

[45] D.G. Luenberger. *Optimization by vector space methods*. Wiley-Interscience, 1997.

[46] P.J. Maenhout. Robust portfolio rules and asset pricing. *Review of Financial Studies*, 17(4):951–983, 2004.

[47] R. Mansini, W. Ogryczak, and M.G. Speranza. Conditional value at risk and related linear programming models for portfolio optimization. *Annals of Operations Research*, 152(1):227–256, 2007.

[48] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

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

[50] A.J. McNeil, R. Frey, and P. Embrechts. *Quantitative risk management: Concepts, techniques and tools*. Princeton Univ Pr, 2005.

[51] R.O. Michaud. The Markowitz optimization enigma: Is optimized optimal? *ICFA Continuing Education Series*, 1989(4):43–54, 1989.

[52] R.J. Muirhead. *Aspects of multivariate statistical theory*. Wiley Online Library, 1982.

[53] G. Perakis and G. Roels. Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188, 2008.

[54] D. Pollard. *Convergence of stochastic processes*. Springer, 1984.

[55] R.T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.

[56] L.C.G. Rogers. The relaxed investor and parameter uncertainty. *Finance and Stochastics*, 5(2):131–154, 2001.

[57] L.J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951.

[58] A. Schied. Optimal investments for robust utility functionals in complete market models. *Mathematics of Operations Research*, 30(3):750–764, 2005.

[59] D. Terlizzese. Relative minimax. Working paper, Einaudi Institute for Economics and Finance (EIEF), 2006.

[60] R.H. Tütüncü and M. Koenig. Robust asset allocation. *Annals of Operations Research*, 132(1):157–187, 2004.

[61] A.W. Van der Vaart. *Asymptotic statistics*. Cambridge Univ Press, 2000.

[62] A. Wald. Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics*, 10(4):299–326, 1939.

[63] S. Zhu and M. Fukushima. Worst-case conditional value-at-risk with application to robust portfolio management. *Operations research*, 57(5):1155–1168, 2009.