Outperformance and Tracking:

Dynamic Asset Allocation for Active and Passive Portfolio Management

Ali Al-Aradi, Sebastian Jaimungal

Department of Statistical Sciences, University of Toronto

Abstract

Portfolio management problems are often divided into two types: active and passive, where the objective is to outperform and track a preselected benchmark, respectively. Here, we formulate and solve a dynamic asset allocation problem that combines these two objectives in a unified framework. We look to maximize the expected growth rate differential between the wealth of the investor's portfolio and that of a *performance* benchmark while penalizing risk-weighted deviations from a given *tracking* portfolio. Using stochastic control techniques, we provide explicit closed-form expressions for the optimal allocation and we show how the optimal strategy can be related to the growth optimal portfolio. The admissible benchmarks encompass the class of functionally generated portfolios (FGPs), which include the market portfolio, as the only requirement is that they depend only on the prevailing asset values. Finally, some numerical experiments are presented to illustrate the risk-reward profile of the optimal allocation.

Keywords: Active portfolio management; Stochastic Portfolio Theory; stochastic control; portfolio selection; growth optimal portfolio; functionally generated portfolios.

1. Introduction

Active portfolio management aims at constructing portfolios that generate superior returns. This is in contrast to *passive portfolio management*, where the goal is to track a given index. The goals of active management can be further divided into those of an *absolute* nature and those of

a *relative* nature. Absolute goals do not involve any external processes, e.g. maximizing expected growth, minimizing the variance of terminal wealth or the probability of ruin. Relative goals are those that involve an external benchmark often given in the form of a portfolio of assets, e.g. maximizing the probability of outperforming the market portfolio. In some cases, the performance of an active manager may also incorporate their deviation from the benchmark, as they are penalized for tracking error or taking on excessive active risk. It may also be desirable to hold portfolios with minimal levels of absolute risk, measured in terms of the volatility of the portfolio's wealth process. The main goal of this paper is to construct portfolios that achieve optimal relative returns against a *performance benchmark* given by a portfolio that depends on prevailing asset values. We also allow for an investor to anchor their portfolio to a *tracking benchmark* by penalizing deviations from the latter. It is often the case for active managers that these two portfolios coincide. Care must be taken in setting up and solving the control problem considering that the weights of the benchmark process are themselves stochastic processes, and must be incorporated as additional state variables.

The impetus for our work is Oderda (2015), where the author investigates optimizing relative returns using stochastic portfolio theory (SPT), but only partially solves the problem via static optimization. Here, we significantly improve on those results by using optimal stochastic control techniques. We succeed in providing explicit closed form expressions for the optimal allocation. In addition, we demonstrate how the optimal portfolio we derive relates to the growth optimal portfolio, which maximizes expected growth over any horizon and plays an important role in the financial literature.

One major issue that is faced by investors that seek outperformance is the need to robustly estimate the growth rates of individual assets. This can be a very difficult goal to achieve. However, allowing the investor to anchor to a portfolio of their choice embeds in our stochastic control problem a portfolio tracking problem which can be used to overcome the issues of estimation. In particular, an investor can choose to track a portfolio known to have certain outperformance properties, such as certain functionally generated portfolios or universal portfolios, and attempt to improve their risk-adjusted performance via the stochastic control problem that we pose.

2. Literature Review

There is a great deal of literature dedicated to solving various portfolio selection problems via stochastic control theory. The seminal work of Merton (1969) introduced the dynamic asset allocation and consumption problem, utilizing stochastic control techniques to derive optimal investment and consumption policies. Extensions can be found in Merton (1971), Magill and Constantinides (1976), Davis and Norman (1990), Browne (1997) and more recently Blanchet-Scalliet et al. (2008), Liu and Muhle-Karbe (2013) and Ang et al. (2014) to name a few. The focus in these papers is generally on maximizing the utility of discounted consumption and terminal wealth or minimizing shortfall probability, or other related absolute performance measures that are independent of any external benchmark or relative goal.

The question of optimal *active management* was introduced in Browne (1999a) and later refined in Browne (1999b). In these papers, an investor is able to trade in a number of assets modeled by geometric Brownian motions (GBMs) and a risk-free asset. The investor aims to maximize the performance of their portfolio *relative* to a stochastic benchmark which is also modeled as an exogenously given GBM. They assume markets to be "incomplete" by modeling the benchmark with a Brownian motion that is independent of the risky securities. This modeling approach allows for a wide range of benchmarks, including inflation, exchange rates and other portfolios. The paper provides a general result and is applied to find the optimal portfolio strategy for a number of active management problems, namely: (i) maximizing the probability that the investor's wealth achieves a certain performance goal relative to the benchmark, before falling below it to a predetermined shortfall; (ii) minimizing the expected time to reach the performance goal; (iii) maximizing the expected reward obtained upon reaching the goal; (iv) minimizing the expected penalty paid upon falling to the shortfall level; (v) maximizing the utility of relative wealth. Browne (2000) extends the work to include a risk constraint in the optimization problem, although the setup there is restricted to a complete market where the benchmark portfolio consists only of assets from the asset universe. Browne (2000) has the same objectives as Browne (1999a) and Browne (1999b), with the additional constraint that the probability of success is bounded below by a given level.

One important shortcoming in these papers, particularly regarding the problem of beating a benchmark portfolio, is that the solutions only apply to benchmark portfolios with constant weights and can be extended at most to the case of deterministic weight processes. This is because

the case of stochastic weights would require including the benchmark weights as state variables, hence significantly increasing the dimensionality of the control problem. It can be argued that the case of stochastic weights is more relevant, considering that one of the simplest and arguably most widespread benchmark is the market portfolio, whose weights evolve stochastically through time.

Much of the model setup in this paper is based on stochastic portfolio theory (see e.g., Fernholz (2002) or Karatzas and Fernholz (2009) for a more recent overview), which is a flexible framework for analyzing market structure and portfolio behavior. SPT is a descriptive, rather than a normative, approach to addressing these issues and relies on a minimal set of assumptions that are readily satisfied in real equity markets. One of the seminal papers in the development of SPT is Fernholz and Shay (1982) who introduce the central tools used in SPT, primarily the notion of excess growth, along with a characterization of long-term behavior of stocks and equity markets and the conditions for achieving market equilibrium. Fernholz (1999a) expands on the work of Fernholz and Shay (1982) by building on the notion of excess growth and introducing the concept of market diversity and entropy as a measure of diversity, and their role in describing stock market equilibrium. Fernholz (1999b) further develops the SPT framework by introducing portfolio generating functions, a tool that can be used to construct dynamic portfolios that depend only on the market weights of the asset universe. Fernholz (2001) extends the concept of functionally generated portfolios to functions of ranked market weights, which allows for the construction of portfolios based on company size.

A main focus in SPT is determining conditions under which relative arbitrage opportunities - portfolios that are guaranteed to outperform the market portfolio - exist over various time horizons, and methods for constructing such portfolios. This is discussed in works such as Fernholz and Karatzas (2005), Banner and Fernholz (2008), Pal and Wong (2013), Wong (2015), Pal and Wong (2016) and Fernholz et al. (2016), among others. It should be noted that, unlike the present work, SPT is typically concerned with outperformance in the a.s. sense. An exception to this is the work of Samo and Vervuurt (2016), in which machine learning techniques are utilized to achieve outperformance in expectation by maximizing the investor's Sharpe ratio.

An application of SPT in the context of active management is found in Oderda (2015), where the author seeks to explain the superior risk-adjusted performance of rule-based, non cap-weighted investment strategies relative to the market portfolio. The dynamics of relative wealth

are described using results from SPT and a static optimization is applied to solve for the investor's optimal portfolio. The solution to this utility maximization problem yields an optimal portfolio that holds a proportion in five passive rule-based portfolios: the market portfolio, the equal-weight portfolio, the global minimum variance portfolio, the risk parity portfolio and the high-cash flow portfolio.

As mentioned earlier, here, we significantly improve on the results in Oderda (2015) by using optimal stochastic control techniques. We succeed in providing explicit closed form expressions for the optimal allocation. In addition, we demonstrate how the optimal portfolio relates to the growth optimal portfolio and how to apply our techniques to leverage the relative arbitrage properties of certain functionally generated portfolios.

3. Model Setup

3.1. Market Model

We begin with the typical SPT market model as described in, for example, Fernholz (2002). Let $\boldsymbol{W} = \{\boldsymbol{W}(t) = (W_1(t), ..., W_k(t))'\}_{t \geq 0}$ be a standard k-dimensional Wiener process defined on the filtered probability space $(\Omega, \mathcal{F}, \mathfrak{F}, \mathbb{P})$, where $\mathfrak{F} = \{\mathcal{F}_t\}_{t \geq 0}$ is the $\mathbb{P}$-augmentation of the natural filtration generated by $\boldsymbol{W}$, $\mathcal{F}_t = \sigma(\{\boldsymbol{W}(s)\}_{s \in [0,t]})$. The Wiener processes are the drivers of the assets returns in the economy. Further, we will assume that the market consists of n stocks, where $n \geq k$.

Definition 1. *The **stock price process** for asset i, $X_i = \{X_i(t)\}_{t \geq 0}$, is a continuous semimartingale of the form*

$$X_i(t) = X_i(0) \cdot \exp\left(\int_0^t (\gamma_i(s) + \delta_i(s))\, ds + \int_0^t \sum_{\nu=1}^k \xi_{i\nu}(s)\, dW_\nu(t)\right) \tag{3.1}$$

*for $i = 1, ..., n$, where γ_i, δ_i and $\xi_{i\nu}$ are deterministic functions corresponding to the asset's **growth rate**, **dividend rate** and **volatility** with respect to the ν-th source of randomness.*

We require the following assumption on the model parameters:

Assumption 1. *The growth, dividend and volatility functions, γ_i, δ_i and $\xi_{i\nu}$, are bounded and differentiable for $i = 1, ..., n$ and $v = 1, ..., k$.*

5

One model which satisfies this assumption is one where the growth, dividend and volatility parameters are constant; a multivariate extension of the typical Black-Scholes model. This model will be used in the implementation section to obtain some numerical results. Note that our market model does not include a risk-free asset, however it would not be difficult to extend the control problem to incorporate such an asset. Additionally, Assumption 1 is stronger than the usual integrability assumptions made in SPT contexts, but simplifies the process of solving the stochastic control problem via a dynamic programming approach. While it is possible to allow for the model parameters to be stochastic, this will require additional assumptions and it will make the proofs of optimality more cumbersome without adding a great deal to the main results of this paper.

Without loss of generality, we assume each asset has a single outstanding share, so that X_i represents the total market capitalization for the corresponding security. Also, it is more convenient to work with log prices, $\log X_i$, which satisfy the stochastic differential equation (SDE):

$$d \log X_i(t) = (\gamma_i(t) + \delta_i(t)) \, dt + \boldsymbol{\xi}_i(t)' \, d\boldsymbol{W}(t), \tag{3.2}$$

where $\boldsymbol{\xi}_i(t) = (\xi_{i1}(t), ..., \xi_{ik}(t))'$ is a $k \times 1$ vector of volatilities. This can also be expressed in vector notation as follows:

$$d \log \boldsymbol{X}(t) = (\boldsymbol{\gamma}(t) + \boldsymbol{\delta}(t)) \, dt + \boldsymbol{\xi}(t) \, d\boldsymbol{W}(t), \tag{3.3}$$

where

$$\log \boldsymbol{X}(t) \underset{(n \times 1)}{=} (\log X_1(t), ..., \log X_n(t))', \qquad \boldsymbol{\gamma}(t) \underset{(n \times 1)}{=} (\gamma_1(t), ..., \gamma_n(t))',$$

$$\boldsymbol{\delta}(t) \underset{(n \times 1)}{=} (\delta_1(t), ..., \delta_n(t))', \qquad \boldsymbol{\xi}(t) \underset{(n \times k)}{=} (\boldsymbol{\xi}_1(t), ..., \boldsymbol{\xi}_n(t))'.$$

Next, we introduce the covariance process:

Definition 2. *The **covariance process** is a matrix-valued function given by:*

$$\underset{(n \times n)}{\boldsymbol{\Sigma}(t)} = \boldsymbol{\xi}(t) \cdot \boldsymbol{\xi}(t)'. \tag{3.4}$$

The covariance of asset i with asset j can then be expressed as $\boldsymbol{\Sigma}_{ij}(t) = \boldsymbol{\xi}_i(t)' \boldsymbol{\xi}_j(t)$. We assume next that the market satisfies the usual **nondegeneracy condition** that ensures that the covariance matrix is invertible:

6

Assumption 2. *The covariance process $\boldsymbol{\Sigma}(t)$ is nonsingular for all $t \geq 0$, and there exists $\varepsilon_1 > 0$ such that for all $\mathbf{x} \in \mathbb{R}^n$ and $t \geq 0$*

$$\mathbf{x}'\boldsymbol{\Sigma}(t)\mathbf{x} \;\geq\; \varepsilon_1\|\mathbf{x}\|^2 \qquad \mathbb{P}\text{-}a.s. \tag{3.5}$$

Thus, $\boldsymbol{\Sigma}(t)$ is positive definite for all $t \geq 0$, and $\boldsymbol{\Sigma}^{-1}(t)$ exists for all $t \geq 0$. We also assume the market has **bounded variance**:

Assumption 3. *There exists $M_1 > 0$ such that for all $\mathbf{x} \in \mathbb{R}^n$ and $t \geq 0$,*

$$\mathbf{x}'\boldsymbol{\Sigma}(t)\mathbf{x} \;\leq\; M_1\|\mathbf{x}\|^2 \qquad \mathbb{P}\text{-}a.s. \tag{3.6}$$

3.2. Portfolio Dynamics

Definition 3. *A **portfolio** is a measurable, $\mathfrak{F}$-adapted, vector-valued process $\boldsymbol{\pi} = \{\boldsymbol{\pi}(t)\}_{t\geq 0}$, where $\boldsymbol{\pi}(t) = (\pi_1(t), ..., \pi_n(t))$ such that, for all $t \geq 0$, $\boldsymbol{\pi}(t)$ is bounded $\mathbb{P}$-almost surely and satisfies:*

$$\pi_1(t) + \cdots + \pi_n(t) = 1 \quad \mathbb{P}\text{-}a.s. \tag{3.7}$$

Each component of $\boldsymbol{\pi}$ represents the proportion of wealth invested in the corresponding stock; negative values of $\pi_i(t)$ indicate a short position in stock i. From a well-known result in SPT, the logarithm of the **portfolio value process**, $Z_{\boldsymbol{\pi}} = \{Z_{\boldsymbol{\pi}}(t)\}_{t\geq 0}$, satisfies the SDE

$$d\log Z_{\boldsymbol{\pi}}(t) = (\gamma_{\boldsymbol{\pi}}(t) + \delta_{\boldsymbol{\pi}}(t))\ dt + \boldsymbol{\xi}_{\boldsymbol{\pi}}(t)'\ d\boldsymbol{W}(t)\,, \tag{3.8}$$

where
$$\underset{(1\times 1)}{\gamma_{\boldsymbol{\pi}}(t)} = \boldsymbol{\pi}(t)'\boldsymbol{\gamma}(t) + \gamma_{\boldsymbol{\pi}}^*(t)\,, \qquad \underset{(1\times 1)}{\gamma_{\boldsymbol{\pi}}^*(t)} = \tfrac{1}{2}\left[\boldsymbol{\pi}(t)'\mathrm{diag}(\boldsymbol{\Sigma}(t)) - \boldsymbol{\pi}(t)'\boldsymbol{\Sigma}(t)\boldsymbol{\pi}(t)\right],$$

$$\underset{(1\times 1)}{\delta_{\boldsymbol{\pi}}(t)} = \boldsymbol{\pi}(t)'\boldsymbol{\delta}(t)\,, \qquad \underset{(k\times 1)}{\boldsymbol{\xi}_{\boldsymbol{\pi}}(t)} = \boldsymbol{\xi}(t)'\boldsymbol{\pi}(t)\,.$$

Here, $\gamma_{\boldsymbol{\pi}}$, $\delta_{\boldsymbol{\pi}}$ and $\boldsymbol{\xi}_{\boldsymbol{\pi}}$ are the **portfolio growth rate**, **portfolio dividend** and **portfolio volatility** processes, respectively, and $\gamma_{\boldsymbol{\pi}}^*$ is the **excess growth rate** process of the portfolio $\boldsymbol{\pi}$. The excess growth rate is equal to half of the difference between the portfolio-weighted average of asset return variances and the portfolio variance. This is a pivotal quantity which plays an important role in SPT, as discussed in more detail in Fernholz (2002).

7

One portfolio of particular interest to us is the **market portfolio**. The market portfolio plays a central role in SPT and, in our context, is one of the most common benchmarks assigned to active managers.

Definition 4. *The **market portfolio** (process), $\boldsymbol{\mu} = \{\boldsymbol{\mu}(t)\}_{t\geq 0}$, is the portfolio with weights given by:*

$$\mu_i(t) = \frac{X_i(t)}{X_1(t) + \cdots + X_n(t)} = \frac{X_i(t)}{Z_{\boldsymbol{\mu}}(t)}, \qquad for\ i = 1, ..., n. \tag{3.9}$$

In other words, the market portfolio holds each stock according to its proportional market capitalization.

Remark 1. *Note that the market portfolio is a vector-valued stochastic process. It is important to highlight the fact that it is a stochastic process and its dynamics must be taken into account when setting up and solving our optimization problem.*

3.3. Relative Return Dynamics

The main state variable in our optimization problem will be the ratio of the wealth of an arbitrary portfolio $\boldsymbol{\pi}$ relative to that of a preselected performance benchmark $\boldsymbol{\rho}$. The following proposition specifies the dynamics of the logarithm of this process.

Proposition 1. *Let $Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(t) = \log\left(\frac{Z_{\boldsymbol{\pi}}(t)}{Z_{\boldsymbol{\rho}}(t)}\right)$ denote the logarithm of relative portfolio wealth for the portfolios $\boldsymbol{\pi}$ and $\boldsymbol{\rho}$. Then this process satisfies the SDE:*

$$dY^{\boldsymbol{\pi},\boldsymbol{\rho}}(t) = a(t, \boldsymbol{\rho}(t), \boldsymbol{\pi}(t))\ dt + \left(\boldsymbol{\xi}_{\boldsymbol{\pi}}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t)\right)'\ d\boldsymbol{W}(t), \tag{3.10}$$

where

$$\begin{aligned}
a(t, \boldsymbol{\rho}(t), \boldsymbol{\pi}(t)) &= (\gamma_{\boldsymbol{\pi}}(t) + \delta_{\boldsymbol{\pi}}(t)) - (\gamma_{\boldsymbol{\rho}}(t) + \delta_{\boldsymbol{\rho}}(t)) \\
&= \boldsymbol{\pi}(t)'\boldsymbol{\alpha}(t) - \tfrac{1}{2}\boldsymbol{\pi}(t)'\boldsymbol{\Sigma}(t)\boldsymbol{\pi}(t) - (\gamma_{\boldsymbol{\rho}}(t) + \delta_{\boldsymbol{\rho}}(t)) \\
\boldsymbol{\alpha}(t) &= \boldsymbol{\gamma}(t) + \boldsymbol{\delta}(t) + \tfrac{1}{2}diag\left(\boldsymbol{\Sigma}(t)\right)
\end{aligned}$$

Proof. The proof follows from noticing that $Y^{\boldsymbol{\pi}}(t) = \log Z_{\boldsymbol{\pi}}(t) - \log Z_{\boldsymbol{\rho}}(t)$, and then applying the portfolio dynamics given in Equation (3.8). ∎

Note that in the proposition above $\boldsymbol{\alpha}(t) = (\alpha_1(t), ..., \alpha_n(t))'$ is a vector of the instantaneous **rate of return** for each asset in the economy.

Corollary 1. *The quadratic variation of* $Y^{\boldsymbol{\pi},\boldsymbol{\rho}}$ *is given by:*

$$d \langle Y^{\boldsymbol{\pi},\boldsymbol{\rho}}, Y^{\boldsymbol{\pi},\boldsymbol{\rho}} \rangle_t = (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t))' \, \boldsymbol{\Sigma}(t) \, (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t)) \ dt \ . \tag{3.11}$$

Proof. Using the volatility of $Y^{\boldsymbol{\pi},\boldsymbol{\rho}}$ given in Proposition 1 and the definition of $\boldsymbol{\Sigma}$ and $\boldsymbol{\xi}_{\boldsymbol{\pi}}$ in (3.4) and (3.8), respectively, we have:

$$
\begin{aligned}
d \langle Y^{\boldsymbol{\pi},\boldsymbol{\rho}}, Y^{\boldsymbol{\pi},\boldsymbol{\rho}} \rangle_t &= \left(\boldsymbol{\xi}_{\boldsymbol{\pi}}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t) \right)' \left(\boldsymbol{\xi}_{\boldsymbol{\pi}}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t) \right) \ dt \\
&= (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t))' \, \boldsymbol{\xi}(t) \cdot \boldsymbol{\xi}(t)' \, (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t)) \ dt \\
&= (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t))' \, \boldsymbol{\Sigma}(t) \, (\boldsymbol{\pi}(t) - \boldsymbol{\rho}(t)) \ dt
\end{aligned}
$$

which is the required result. ∎

In the context of active management this quantity is typically referred to as (instantaneous) **active risk** or **tracking error**. An active manager is typically penalized for incurring excessive active risk over their investment horizon.

Remark 2. *Fernholz (2002) refers to this quantity as the relative variance process of* $\boldsymbol{\pi}$ *versus* $\boldsymbol{\rho}$, *and denotes by* $\tau^{\boldsymbol{\rho}}_{\boldsymbol{\pi}\boldsymbol{\pi}}(t)$. *Note that this quantity is equal to zero if and only if* $\boldsymbol{\pi}$ *and* $\boldsymbol{\rho}$ *are equal.*

4. Stochastic Control Problem

In this section we formulate the main stochastic control problem. First, we fix two portfolios against which we measure our outperformance and our active risk, respectively. That is, we choose a **performance benchmark** $\boldsymbol{\rho}$, which the investor wishes to outperform, and a **tracking portfolio** $\boldsymbol{\eta}$, which the investor will penalize deviations from.

Remark 3. *The common setup in active portfolio management is for the performance benchmark to be the same as the tracking portfolio. However, this general setup allows for added flexibility and certainly accommodates the case where* $\boldsymbol{\rho} = \boldsymbol{\eta}$. *This separation is also useful for the situation where we remove the performance benchmark and track functionally generated portfolios to achieve our outperformance goal - this will be discussed in more detail later in the paper.*

The primary objective is to determine the portfolio process $\boldsymbol{\pi}$ that maximizes the expected growth rate of relative wealth to $\boldsymbol{\rho}$ over the investment horizon T. This is equivalent to maximizing the expected utility of relative return assuming a log-utility function. Additionally, the investor is penalized for taking on excessive levels of active risk (measured against $\boldsymbol{\eta}$). We also incorporate a general quadratic penalty term that is independent of the two benchmarks. Putting this together, the **performance criteria** of a portfolio $\boldsymbol{\pi}$ is given by:

$$H^{\boldsymbol{\pi}}(t, y, \boldsymbol{x}) = \mathbb{E}_{t,y,\boldsymbol{x}}\left[\zeta_0 \cdot Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(T) - \frac{\zeta_1}{2}\int_t^T (\boldsymbol{\pi}(s) - \boldsymbol{\eta}(s))'\boldsymbol{\Sigma}(s)(\boldsymbol{\pi}(s) - \boldsymbol{\eta}(s))\ ds - \frac{\zeta_2}{2}\int_t^T \boldsymbol{\pi}(s)'\mathbf{Q}(s)\boldsymbol{\pi}(s)\ ds\right],$$

(4.1)

where $\zeta_0, \zeta_1, \zeta_2 \geq 0$ and $\mathbb{E}_{t,y,\boldsymbol{x}}[\ \cdot\ ]$ is shorthand for denoting $\mathbb{E}[\ \cdot\ |Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(t) = y, \boldsymbol{X}(t) = \boldsymbol{x}]$. We also need the following assumptions:

Assumption 4. *The matrix-valued function $\mathbf{Q}$ is symmetric, positive-definite, (element-wise) differentiable and satisfies*

$$\varepsilon_2\|\mathbf{x}\|^2 \ \leq \ \mathbf{x}'\mathbf{Q}(t)\mathbf{x} \ \leq \ M_2\|\mathbf{x}\|^2$$

for some $\varepsilon_2 > 0$ and $M_2 < \infty$ for all $\mathbf{x} \in \mathbb{R}^n$ and $t \geq 0$.

Assumption 5. *The performance benchmark process $\boldsymbol{\rho} = \{\boldsymbol{\rho}(t)\}_{t\geq 0}$ and the tracking portfolio process $\boldsymbol{\eta} = \{\boldsymbol{\eta}(t)\}_{t\geq 0}$ are Markovian in $\boldsymbol{X}$; i.e. there exist bounded functions $\rho : [0, T] \times \mathbb{R}_+^n \to \mathbb{R}^n$ and $\eta : [0, T] \times \mathbb{R}_+^n \to \mathbb{R}^n$ such that $\boldsymbol{\rho}(t) = \rho(t, \boldsymbol{X}(t))$ and $\boldsymbol{\eta}(t) = \eta(t, \boldsymbol{X}(t))$. Additionally, we assume that the functions ρ and η are once differentiable with respect to all of their variables.*

Note that the assumption above is satisfied by the class of (time-dependent) functionally generated portfolios, which includes the ubiquitous market portfolio, as they depend only on the prevailing market weights which in turn depend only on the asset values.

Let us now discuss the components of the performance criteria (4.1):

1. The first term, $\zeta_0 Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(T)$, is a terminal reward term which corresponds to the investor wishing to **maximize the expected growth rate differential between their portfolio and the performance benchmark $\boldsymbol{\rho}$**. It is also equivalent to **maximizing the expected utility of relative wealth assuming a log-utility function**. The constant ζ_0 is a subjective parameter determining the amount of emphasis the investor wishes to place on benchmark outperformance. Setting ζ_0 equal to 0 changes our portfolio problem to a tracking problem

where the goal is to track the benchmark while minimizing some running quadratic penalty term. This will become particularly useful when we choose to track a functionally generated portfolio to achieve outperformance as it allows us to avoid estimating asset growth rates which can be extremely difficult.

2. The second term is a running penalty term which **penalizes tracking error/active risk**. That is, deviations from the tracking portfolio are penalized on a running basis, with deviations in riskier assets being penalized more heavily. Alternatively, this penalty term can be viewed as an attempt to **minimize the quadratic variation of relative wealth with respect to $\boldsymbol{\eta}$, $Y^{\boldsymbol{\pi},\boldsymbol{\eta}}(t)$**. The parameter ζ_1 represents the investor's tolerance for active risk.

3. The final term is a **general quadratic running penalty term** - that does not involve either benchmark - with an associated tolerance parameter ζ_2. One possible choice for $\mathbf{Q}(t)$ is the covariance matrix $\boldsymbol{\Sigma}(t)$, which can be adopted to minimize the *absolute* risk of the portfolio measured in terms of the quadratic variation of the portfolio wealth process, $Z_{\boldsymbol{\pi}}(t)$. Another option is to let $\mathbf{Q}(t)$ be a constant diagonal matrix, which has the effect of penalizing allocation in each asset according to the magnitude of the corresponding diagonal entry. The investor can use this choice of $\mathbf{Q}$ as a way of imposing a set of "soft" constraints on allocation to each asset. We will demonstrate later on how this penalty term can be interpreted as tilting away from certain assets or, more generally, "shrinking" towards some desired portfolio.

Remark 4. *While it may be possible to extend this control problem by considering alternative utility functions, particularly to observe the effect of the performance benchmark on the optimal portfolio, this extension does not appear to be trivial. Using general utility functions would likely require other approaches and lies outside the scope of this work.*

Now, the performance criteria (4.1) can be simplified by noticing that:

$$Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(T) = Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(t) + \int_t^T a(s, \boldsymbol{\rho}(s), \boldsymbol{\pi}(s)) \, ds + \int_t^T \left(\boldsymbol{\xi}_{\boldsymbol{\pi}}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t) \right)' \, d\boldsymbol{W}(s)$$

and that, since $\boldsymbol{\xi}(t)$ is square-integrable, the conditional expectation of the second integral above is equal to zero due to the martingale property of the stochastic integral. This allows us to rewrite the performance criteria as:

11

$$H^{\boldsymbol{\pi}}(t,y,\boldsymbol{x}) = \zeta_0 y + \mathbb{E}_{t,\boldsymbol{x}}\left[\int_t^T \zeta_0 \cdot a(s,\boldsymbol{\rho}(s),\boldsymbol{\pi}(s))\right.$$

$$\left. - \tfrac{\zeta_1}{2}\cdot(\boldsymbol{\pi}(s)-\boldsymbol{\eta}(s))'\boldsymbol{\Sigma}(s)(\boldsymbol{\pi}(s)-\boldsymbol{\eta}(s)) - \tfrac{\zeta_2}{2}\cdot\boldsymbol{\pi}(s)'\mathbf{Q}(s)\boldsymbol{\pi}(s)\ ds\right]$$

$$= \zeta_0 y + h^{\boldsymbol{\pi}}(t,\boldsymbol{x}) \tag{4.2}$$

where $h^{\boldsymbol{\pi}}(t,\boldsymbol{x})$ can be interpreted as the performance criteria for the related stochastic control problem consisting only of the running reward/penalty term appearing in the integral of (4.2) above. The **value function** for this auxiliary stochastic control problem is given by:

$$h(t,\boldsymbol{x}) = \sup_{\boldsymbol{\pi}\in\mathcal{A}} h^{\boldsymbol{\pi}}(t,\boldsymbol{x})\ , \tag{4.3}$$

where $\mathcal{A}$ is the set of admissible strategies consisting of all portfolios, which, in our context is defined to be the collection of almost surely bounded vector-valued processes that sum up to 1, and hence all portfolios have finite $\mathbb{L}^2(\Omega\times[0,T])$-norm. Notice that, due to the assumptions made in Section 3, the performance criteria (4.2) is finite for any portfolio in the admissible set. Assuming the appropriate differentiability conditions, the dynamic programming principle suggests that the value function in (4.3) satisfies the Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{cases}\partial_t h + \sup_{\boldsymbol{\pi}\in\mathcal{A}}\left\{\mathcal{L}^{\boldsymbol{X}}h + \zeta_0\cdot a(t,\rho(t,\boldsymbol{x}),\boldsymbol{\pi}) - \tfrac{\zeta_1}{2}\cdot(\boldsymbol{\pi}-\eta(t,\boldsymbol{x}))'\boldsymbol{\Sigma}(t)(\boldsymbol{\pi}-\eta(t,\boldsymbol{x})) - \tfrac{\zeta_2}{2}\cdot\boldsymbol{\pi}'\mathbf{Q}(t)\boldsymbol{\pi}\right\} = 0\ ,\\[2mm] \hspace{10cm} h(T,\boldsymbol{x}) = 0\ .\end{cases}$$
$$\tag{4.4}$$

where $\mathcal{L}^{\boldsymbol{X}}$ is the infinitesimal generator of the asset value processes. Notice that Assumption 5 ensures that the asset values are the only state variables in this PDE.

Proposition 2. *The solution to the HJB equation* (4.4) *is given by:*

$$h(t,\boldsymbol{x}) = \mathbb{E}_{t,\boldsymbol{x}}\left[-\int_t^T G(s,\boldsymbol{X}(s))\ ds\right]\ , \tag{4.5}$$

where the expectation is taken under the physical measure $\mathbb{P}$, *and the function* G *is given by:*

$$G(t,\boldsymbol{x}) = C(t,\boldsymbol{x}) + \frac{1}{2}\left[\frac{1-\mathbf{1}'\mathbf{A}^{-1}(t)\boldsymbol{B}(t,\boldsymbol{x})}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}}\cdot\mathbf{1} + \boldsymbol{B}(t,\boldsymbol{x})\right]'\mathbf{A}^{-1}(t)\left[\frac{1-\mathbf{1}'\mathbf{A}^{-1}(t)\boldsymbol{B}(t,\boldsymbol{x})}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}}\cdot\mathbf{1} - \boldsymbol{B}(t,\boldsymbol{x})\right]\ .$$

with

$$\mathbf{A}(t) = (\zeta_0 + \zeta_1)\boldsymbol{\Sigma}(t) + \zeta_2\mathbf{Q}(t) \,,$$
$$\underset{(n\times n)}{}$$

$$\boldsymbol{B}(t, \boldsymbol{x}) = \zeta_0\boldsymbol{\alpha}(t) + \zeta_1\boldsymbol{\Sigma}(t)\eta(t, \boldsymbol{x}) \,, \quad and$$
$$\underset{(n\times 1)}{}$$

$$C(t, \boldsymbol{x}) = \zeta_0 \cdot \left(\boldsymbol{\gamma}_{\rho(t,\boldsymbol{x})}(t) + \boldsymbol{\delta}_{\rho(t,\boldsymbol{x})}(t)\right) + \tfrac{\zeta_1}{2} \cdot \eta(t, \boldsymbol{x})'\boldsymbol{\Sigma}(t)\eta(t, \boldsymbol{x}) \,.$$
$$\underset{(1\times 1)}{}$$

Proof. See Appendix A.1. ∎

Next, we present one of the key results of the paper: a verification theorem showing that the candidate solution provided in Proposition 2 coincides with the value function. As well, we present the explicit form of the optimal weights.

Theorem 1. *The candidate value function given in Proposition 2 is indeed the solution to the stochastic control problem* (4.3). *Moreover, the optimal portfolio is given by:*

$$\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*(t) = \mathbf{A}^{-1}(t) \cdot \left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}(t) \cdot \boldsymbol{B}(t, \boldsymbol{X}(t))}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}} \cdot \mathbf{1} + \boldsymbol{B}(t, \boldsymbol{X}(t))\right] \,, \tag{4.6a}$$

where

$$\mathbf{A}(t) = (\zeta_0 + \zeta_1)\boldsymbol{\Sigma}(t) + \zeta_2\mathbf{Q}(t) \,, \quad and \quad \boldsymbol{B}(t, \boldsymbol{x}) = \zeta_0\boldsymbol{\alpha}(t) + \zeta_1\boldsymbol{\Sigma}(t)\eta(t, \boldsymbol{x}) \,. \tag{4.6b}$$
$$\underset{(n\times n)}{} \qquad\qquad\qquad \underset{(n\times 1)}{}$$

Proof. See Appendix A.2. ∎

If an investor is faced with the market portfolio as their performance and tracking benchmark, as is the case with the vast majority of active management mandates, then the optimal strategy above can be applied by replacing η with the function $\mu : \mathbb{R}_+^n \to \Delta^n$:

$$\mu(\boldsymbol{X}) = \left(\frac{X_1}{\sum_{i=1}^n X_i}, ..., \frac{X_n}{\sum_{i=1}^n X_i}\right) \,, \tag{4.7}$$

where $\Delta^n = \{\boldsymbol{\mu} \in \mathbb{R}^n : \mu_i \in (0, 1), i = 1, ..., n; \ \boldsymbol{\mu}'\mathbf{1} = 1\}$ is the standard n-dimensional simplex.

Before discussing the properties of the optimal portfolio in (4.6), we first present two portfolios that are of particular importance: the growth optimal portfolio (GOP) and the minimum quadratic variation portfolio (MQP). As the names suggest, the GOP is the portfolio with maximal expected growth over any time horizon while the MQP is the portfolio with the smallest quadratic variation

of all portfolios over the investment horizon. Alternatively, the GOP corresponds to the optimal portfolio for an investor that wishes to maximize the expected utility of terminal wealth when their utility function is logarithmic. The GOP plays an important role in financial theory for a variety of reasons; see e.g., Christensen (2012) for a comprehensive literature review. In certain limiting cases, the solution to our stochastic control problem can be written in terms of the GOP and MQP. Formally, the GOP and MQP are the solutions to the following optimization problems:

$$\boldsymbol{\pi}_{GOP} = \arg\sup_{\pi \in \mathcal{A}} \mathbb{E}\left[\log\left(\frac{Z_{\boldsymbol{\pi}}(T)}{Z_{\boldsymbol{\pi}}(0)}\right)\right] , \tag{4.8}$$

$$\boldsymbol{\pi}_{MQP} = \arg\inf_{\pi \in \mathcal{A}} \mathbb{E}\left[\int_0^T \boldsymbol{\pi}'\boldsymbol{\Sigma}(t)\boldsymbol{\pi} \ dt\right] . \tag{4.9}$$

Corollary 2. *The growth optimal portfolio is given by:*

$$\boldsymbol{\pi}_{GOP}(t) = \left(1 - \mathbf{1}'\boldsymbol{\Sigma}^{-1}(t)\boldsymbol{\alpha}(t)\right) \cdot \boldsymbol{\pi}_{MQP}(t) + \boldsymbol{\Sigma}^{-1}(t)\boldsymbol{\alpha}(t) . \tag{4.10}$$

where $\boldsymbol{\pi}_{MQP}$ is the minimum quadratic variation portfolio given by:

$$\boldsymbol{\pi}_{MQP}(t) = \frac{1}{\mathbf{1}'\boldsymbol{\Sigma}^{-1}(t)\mathbf{1}} \cdot \boldsymbol{\Sigma}^{-1}(t)\mathbf{1} . \tag{4.11}$$

Proof. Use the optimal control given in Theorem 1 with $\zeta_0 = \zeta_1 = 0$ and $\mathbf{Q} = \boldsymbol{\Sigma}$ to obtain the MQP, and with $\zeta_1 = \zeta_2 = 0$ to obtain the GOP. ∎

Remark 5. *Oderda (2015) refers to the MQP as the global minimum variance (GMV) portfolio. However, this is a misnomer as the latter term is typically reserved for the portfolio that minimizes the variance of terminal wealth, $\mathbb{V}[Z_{\boldsymbol{\pi}}(T)]$, for a given investment horizon. Though it should be noted that in a discrete-time, single-period model where the covariance matrix of returns is given by $\boldsymbol{\Sigma}$, the GMV is in fact equal to $\frac{1}{\mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1}} \cdot \boldsymbol{\Sigma}^{-1}\mathbf{1}$.*

Armed with the definitions of the GOP and MQP, we present the following corollary which summarizes some of the properties of the optimal strategy given in (4.6), as well as its relation to the aforementioned portfolios.

Corollary 3. *The optimal portfolio given in Theorem 1 satisfies:*

(i) The limiting portfolios as the preference parameters tend to infinity are given by:

$$\lim_{\zeta_0 \to \infty} \boldsymbol{\pi}_\zeta^*(t) = \boldsymbol{\pi}_{GOP}(t)$$

$$= \boldsymbol{\pi}_\zeta^*(t) \text{ with } \boldsymbol{\zeta} = (\zeta_0, 0, 0) \qquad \text{(no running penalties)}$$

$$\lim_{\zeta_1 \to \infty} \boldsymbol{\pi}_\zeta^*(t) = \boldsymbol{\eta}(t)$$

$$= \boldsymbol{\pi}_\zeta^*(t) \text{ with } \boldsymbol{\zeta} = (0, \zeta_1, 0) \qquad \text{(tracking penalty only)}$$

$$\lim_{\zeta_2 \to \infty} \boldsymbol{\pi}_\zeta^*(t) = \frac{1}{\mathbf{1}'\mathbf{Q}^{-1}(t)\mathbf{1}} \cdot \mathbf{Q}^{-1}(t)\mathbf{1}$$

$$= \boldsymbol{\pi}_\zeta^*(t) \text{ with } \boldsymbol{\zeta} = (0, 0, \zeta_2) \qquad \text{(absolute penalty only)}$$

(ii) When $\zeta_0, \zeta_1 > 0$ and $\zeta_2 = 0$ (no absolute penalty):

$$\boldsymbol{\pi}_\zeta^*(t) = c \cdot \boldsymbol{\pi}_{GOP}(t) + (1 - c) \cdot \boldsymbol{\eta}(t)$$

where $c = \frac{\zeta_0}{\zeta_0 + \zeta_1} \in (0, 1)$ is a fixed constant.

(iii) When $\mathbf{Q}(t) = \boldsymbol{\Sigma}(t)$ (minimize relative and absolute risk):

$$\boldsymbol{\pi}_\zeta^*(t) = c_1 \cdot \boldsymbol{\pi}_{GOP}(t) + c_2 \cdot \boldsymbol{\eta}(t) + c_3 \cdot \boldsymbol{\pi}_{MQP}(t)$$

where $c_i = \frac{\zeta_i}{\zeta_0 + \zeta_1 + \zeta_2} \in (0, 1)$ for $i = 1, 2, 3$ are fixed constants summing to 1.

Proof. See Appendix A.3. ∎

A few points are worth mentioning at this juncture:

1. The limiting portfolios in the corollary above show that as the emphasis on outperformance increases, the optimal portfolio tends to the GOP. This is a consequence of the GOP having the highest expected growth rate, and hence delivering the highest level of (expected) outperformance. Similarly, as the aversion to tracking error increases, the optimal portfolio tends to the tracking benchmark. This reflects the fact that the investor is heavily penalizing any deviation from the tracking benchmark and decides to hold the latter to avoid any running penalties. Finally, when the third aversion parameter tends to infinity, the investor holds the portfolio that minimizes the absolute running penalty at each point in time. These three

15

limiting cases are equivalent to investors that are *only* interested in outperformance, tracking or reducing an absolute running penalty, which lead to the expected result of holding the GOP, the benchmark itself or the absolute penalty minimizer, respectively.

2. In the absence of the absolute running penalty, the optimal strategy consists of holding a constant proportion in the tracking benchmark and the remainder in the GOP. The idea is to benefit from the expected growth rate of the GOP while modulating its high levels of active risk by investing in the tracking benchmark. The proportions are determined by the relative importance the investor places on outperformance versus tracking. The fact that a *constant* proportion of wealth is invested in each portfolio is related to the fact that log-utility falls under the class of constant relative risk aversion (CRRA) utility functions, wherein the fraction of wealth invested in risky assets is independent of the level of wealth. Since we are concerned with *active* risk relative to the tracking benchmark in this case, the tracking benchmark plays the role of the risk-free asset.

3. When the investor wishes to minimize both relative and absolute risk, their optimal strategy is to hold constant proportions in the GOP, MQP and tracking benchmark. These holdings contribute to increasing growth, decreasing absolute risk and decreasing active risk, respectively.

4. The optimal solution does not depend on the performance benchmark ρ. The reason for this is simple: the investor wishes to maximize the difference between the expected growth rates of their portfolio and the performance benchmark. Regardless of the choice of ρ, the way this difference can be maximized is by choosing the portfolio with the highest possible growth rate, i.e. the GOP. More fundamentally, ρ does not appear since we have that $\log\left(\frac{Z_{\pi}(T)}{Z_{\rho}(T)}\right) = \log(Z_{\pi}(T)) - \log(Z_{\rho}(T))$ and the second term does not involve the control and hence can be left out of the performance criteria (4.2) altogether. This is unlikely to be the case if alternative utility functions are adopted.

5. The optimal portfolio is myopic, i.e. independent of the investment horizon T. This is expected as the GOP can be shown to maximize expected growth at *any* time horizon. This property is also a typical of logarithmic utility functions.

6. When the investor is not interested in outperformance, i.e. $\zeta_0 = 0$, the optimal strategy does not require knowledge/estimation of the assets' growth rates $\boldsymbol{\gamma}$, which can be quite difficult to estimate robustly. Estimation of the covariance matrix, $\boldsymbol{\Sigma}$, is still needed. ζ_0 can also be used to reflect the investor's confidence in their growth rate estimates: a small value of ζ_0 (relative to ζ_1 and ζ_2) is akin to a small allocation in the growth optimal portfolio (relative to the other subportfolios). Furthermore, depending on the nature of the performance benchmark, outperformance can be achieved in an indirect manner. When that benchmark is the market portfolio, as is often the case for active managers, the tools of SPT provide a wide range of functionally generated portfolios that constitute relative arbitrages with respect to the market portfolio and do not require the estimation of asset growth rates; see Fernholz (2002) for more details. The investor can simply track one of these portfolios setting $\zeta_0 = 0$ while imposing their own absolute penalty term via the matrix $\mathbf{Q}(t)$.

7. The absolute penalty term forces the optimal strategy towards the "shrinkage" portfolio given by $\frac{1}{\mathbf{1}'\mathbf{Q}^{-1}(t)\mathbf{1}} \cdot \mathbf{Q}^{-1}(t)\mathbf{1}$. When $\mathbf{Q}$ is a diagonal matrix $\mathbf{Q}(t) = \mathrm{diag}\left(w_1(t), ..., w_n(t)\right)$, then $\mathbf{Q}^{-1}(t) = \mathrm{diag}\left(\frac{1}{w_1(t)}, ..., \frac{1}{w_n(t)}\right)$. From this, we see that the absolute penalty term forces us to shrink to a portfolio proportional to $\left(\frac{1}{w_1(t)}, ..., \frac{1}{w_n(t)}\right)$; when w_i is large $1/w_i$ is small and the shrinkage portfolio allocates less capital to asset i. This can be used to tilt away from undesired assets; taking $w_i \to \infty$ forces the allocation in asset i to zero. Additionally, taking $\mathbf{Q} = \mathbf{I}$ penalizes large positions in any asset by shrinking to the equal-weight portfolio, while setting $w_i(t) = \boldsymbol{\Sigma}_{ii}(t)$ forces shrinkage to the risk parity portfolio.

The points made above highlight the motivation for including the two running penalties. In principle, when an investor's goal is simply to outperform a performance benchmark, they would hold the GOP to maximize their expected growth. However, it is well-documented that the GOP is associated with very large levels of risk (in terms of portfolio variance) as well as potentially large short positions in a number of assets. Adding the relative and absolute penalty terms mitigates some of this risk. It is also worth noting that the decompositions in Corollary 3 (ii) and (iii) can help guide the choice of the subjective parameters ζ_i which can be used to express the proportion of wealth the investor wishes to place in the GOP, tracking benchmark and MQP. The simulation results provided later can also give a rough idea of the relative and absolute risk-return characteristics of optimal portfolios for different choices of $\boldsymbol{\zeta}$ to further refine the investor's decision.

Finally, we can interpret the optimal portfolio as being the growth optimal portfolio for an alternative market:

Corollary 4. *The optimal portfolio $\boldsymbol{\pi}_{\zeta}^{*}$ is the growth optimal portfolio for a market with a modified rate of return process $\boldsymbol{\alpha}^{*}$ and a modified covariance process $\boldsymbol{\Sigma}^{*}$ given by:*

$$\boldsymbol{\alpha}^{*}(t) = \zeta_0 \boldsymbol{\alpha}(t) + \zeta_1 \boldsymbol{\Sigma}(t)\eta(t, \boldsymbol{X}(t)) \tag{4.12a}$$

$$\boldsymbol{\Sigma}^{*}(t) = (\zeta_0 + \zeta_1)\boldsymbol{\Sigma}(t) + \zeta_2 \mathbf{Q}(t) \tag{4.12b}$$

Proof. The result follows from direct comparison of the optimal portfolio in Theorem 1 with the form of the growth optimal portfolio in Corollary 2. ∎

The last corollary has an interesting interpretation. Consider first the adjustment term added to the rate of return process, $\zeta_1 \boldsymbol{\Sigma}(t)\eta(t, \boldsymbol{X}(t))$. It is straightforward to show that the elements of this vector constitute the quadratic covariation between between each asset and the wealth process associated with the tracking benchmark $\boldsymbol{\eta}$, namely $\langle \log X_i, \log Z_{\boldsymbol{\eta}} \rangle$; see Section 1.2 of Fernholz (2002). This implies that the investor is modifying the assets' rates of return to reward those assets that are more closely correlated with the portfolio they are trying to track, and these artificial rewards are proportional to the investor's desire to track $\boldsymbol{\eta}$ represented by ζ_1. Moreover, if we consider the case where $\mathbf{Q}(t)$ is a diagonal matrix, the modification to the covariance matrix amounts to increasing the variance of each asset according to the corresponding diagonal entry of $\mathbf{Q}(t)$. This in turn makes certain assets less desirable and is tied to the notion of tilting away from those assets as discussed earlier in this section.

5. Implementation

In this section we use Fama-French industry data[1] to simulate the optimal portfolio (OP) given by (4.6) for different values of $\boldsymbol{\zeta}$. We use the simulations to gain an understanding of the portfolios' features and risk-return profiles compared to the other portfolios, namely: the growth optimal portfolio (GOP) given by (4.10) and the minimum quadratic variance portfolio (MQP) given by (4.11), as well as the optimal solution derived in Oderda (2015) which we will refer to as the maximal drift portfolio (MDP).

---

[1] Source: `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

5.1. Data

The data consists of monthly returns (with and without dividends) for five industries (Consumer, Manufacturing, Hi-Tech, Health, and Other) for the period January 2005 to July 2017. Industry portfolios are constructed by assigning each stock to one of the five groupings listed above according to their SIC code at the time. We then treat the resulting industries as the constituents of our market, i.e. our market consists of five assets that are the industries themselves. The data also contains the number of firms and the average market capitalization of firms for each industry, which allows us to compute a time series of market weights for the five industries.

We work with industry returns rather than individual securities to avoid difficulties that arise from varying investment sets caused by individual companies entering and exiting the market and we restrict ourselves to a market of five industries for computational simplicity and ease of displaying and interpreting results. In our experiments, we found that simulation results are qualitatively similar for various granularity levels, though it should be noted that parameter estimation, particularly of the growth rates, becomes less reliable and robust as the number of assets increases.

5.2. Parameter Estimation

For simplicity, we assume that the parameters γ, δ and ξ are constants and we use the entirety of the dataset for the estimation. Other estimation methods that allow for time-varying parameters are likely more appropriate, however this is beyond the scope of the current work. The parameter estimation procedure involves estimating:

1. the covariance matrix, Σ, estimated by the sample covariance of changes in asset log-values

2. the matrix of volatilities, ξ, which is given as the Cholesky decomposition of Σ

3. the growth rate, γ, estimated as the sample mean of changes in asset log-values (ex-dividends)

4. the dividend growth rate, δ, estimated as the sample mean of changes in asset log-values (with dividends) minus γ.

A summary of the estimated parameters is given in Table 1 below.

Industry	Growth	Dividends	Std. Dev.		Cnsmr	Manuf	HiTec	Hlth	Other
Cnsmr	7.4 %	2.1 %	12 %	Cnsmr	1				
Manuf	5.9 %	2.3 %	16 %	Manuf	0.76	1			
HiTec	8.6 %	1.5 %	16 %	HiTec	0.87	0.81	1		
Hlth	8.3 %	1.9 %	13 %	Hlth	0.73	0.61	0.68	1	
Other	3.2 %	1.8 %	19 %	Other	0.87	0.76	0.82	0.69	1

Table 1: Summary of estimated model parameters: growth rate γ, dividend rate δ and standard deviations $\sqrt{\Sigma_{ii}}$ (left); correlation matrix (right).

5.3. Simulation

Next, we use the estimated parameters to simulate industry value paths and compute the strategy's performance in the simulated market model. For the simulation, we consider an investment horizon of $T = 5$ years with daily rebalancing time-steps ($\Delta t = 1/252$; 252 days per year so that the total number of steps is $N = \Delta t \cdot T = 1260$ days) and 1,000 simulations. The asset values are initialized using the market weights at the end of the dataset (market weights on July 2017). Furthermore, we will take the market portfolio to be the performance benchmark and the tracking benchmark, i.e. $\mu = \rho = \eta$, as is commonly the case in active portfolio management. Also, we take the absolute penalty matrix to be the covariance matrix, i.e. $\mathbf{Q} = \mathbf{\Sigma}$.

For each simulation, we compute weight vectors for the optimal portfolio at each time step using the estimated model parameters and using various ζ values, namely: $\zeta_0 \in \{0.1, 0.5, 5\}$ and $\zeta_1, \zeta_2 \in \{0.001, 0.05, 0.1, ..., 1\}$. We also compute the GOP and MQP; note that since the model parameters are assumed to be constant, the weight vectors for the GOP and MQP are constant.

Remark 6. *The procedure outlined above implicitly assumes that the investor knows the true parameters underlying the data generating process. Clearly, this is not the case in practice and the procedure is more akin to an in-sample form of backtesting. Moreover, in our experiments we find that model misspecification can lead to adverse results as one would expect and it is well-known that parameter estimation, particularly of the asset growth rates, is difficult to perform robustly. However, our point of interest is in studying the features of the optimal portfolio with varying ζ values rather than demonstrating its efficacy in actual trading situations. In other words, we wish to separate the parameter estimation problem (which we do not concern ourselves with in this paper)*

from the stochastic control problem. Additionally, one can heuristically adjust the ζ_0 parameter to reflect the level of "confidence" they place in their growth rate estimates.

In terms of performance metrics, we will consider the following quantities for each portfolio:

1. The performance criteria given in (4.1), where the running penalties given by the Riemann integrals are approximated using the discretization:

$$\text{Terminal Reward} \ = \ \zeta_0 \log \left(\frac{Z_{\boldsymbol{\pi}}(T)}{Z_{\boldsymbol{\mu}}(T)} \right)$$

$$\text{Relative Running Penalty} \ = \ \zeta_1 \Delta t \cdot \sum_{t=1}^{N} (\boldsymbol{\pi}(t) - \boldsymbol{\mu}(t))' \ \boldsymbol{\Sigma}(t) \ (\boldsymbol{\pi}(t) - \boldsymbol{\mu}(t))$$

$$\text{Absolute Running Penalty} \ = \ \zeta_2 \Delta t \cdot \sum_{t=1}^{N} \boldsymbol{\pi}(t)' \ \mathbf{Q}(t) \ \boldsymbol{\pi}(t)$$

2. The (average) absolute return and active return of each portfolio, defined as the time-series sample mean of portfolio returns and excess returns over the market (for a single path), respectively:

$$\text{Absolute Return} \ = \ \frac{1}{\Delta t} \frac{1}{N} \sum_{t=1}^{N} \boldsymbol{\pi}(t)' \ \boldsymbol{R}(t, t+1)$$

$$\text{Active Return} \ = \ \frac{1}{\Delta t} \frac{1}{N} \sum_{t=1}^{N} (\boldsymbol{\pi}(t) - \boldsymbol{\mu}(t))' \ \boldsymbol{R}(t, t+1)$$

where $\boldsymbol{R}(t, t+1)$ is the vector of asset returns between t and $t+1$. Note that the factor $\frac{1}{\Delta t}$ is an annualization factor.

3. The absolute and active risk of each strategy, defined as the time-series sample standard deviation of portfolio returns and excess returns over the market, respectively:

$$\text{Absolute Risk} \ = \ \sqrt{\frac{1}{\Delta t} \frac{1}{N} \sum_{t=1}^{N} (\boldsymbol{\pi}(t)' \ \boldsymbol{R}(t, t+1) - \text{Absolute Return})^2}$$

$$\text{Active Risk} \ = \ \sqrt{\frac{1}{\Delta t} \frac{1}{N} \sum_{t=1}^{N} ((\boldsymbol{\pi}(t) - \boldsymbol{\mu}(t))' \ \boldsymbol{R}(t, t+1) - \text{Active Return})^2}$$

4. The Sharpe ratio and information ratio defined as the ratio of absolute return to absolute risk and active return to active risk, respectively:

$$\text{Sharpe Ratio} \ = \ \frac{\text{Absolute Return}}{\text{Absolute Risk}}$$

$$\text{Information Ratio} \ = \ \frac{\text{Active Return}}{\text{Active Risk}}$$

21

5.4. Results

Figure 1 shows the average absolute and active risk and return values across all simulations for the optimal strategy with different $\boldsymbol{\zeta}$ values. We find that as ζ_1 increases, the performance metrics converge to the metrics of the market portfolio; as ζ_2 increases, they converge to the metrics of the MQP (recall that we set $\mathbf{Q} = \boldsymbol{\Sigma}$). The rate of change is higher for lower values of ζ_0; the optimal strategy is less sensitive to changes in ζ_1 and ζ_2 for high values of ζ_0, in which case it is more closely tied to the GOP.

The convergence features discussed above can be seem more apparently in the scatterplots shown in Figure 2, which show the absolute and relative risk-return characteristics of each of the portfolios we analyze. It is apparent that the GOP outperforms all other portfolios in terms of expected return, however, this comes at a high level of absolute risk. On the other hand, we find that the optimal strategies with $\zeta_0 = 0.1$ (shown in red) have a higher return per unit of risk. Additionally, the optimal portfolios dominate the MDP and markt portfolio and trace out an "efficient frontier" tying the GOP and MQP. The alignment of the optimal portfolios in the second plot in Figure 2 is a manifestation of the observation that the information ratio is insensitive to ζ_1. Also, the dots converge to the origin, since the optimal solutions converge to the market portfolio as ζ_1 tends to infinity.

The right panel in Figure 3 shows the effect of varying ζ_2 on the information ratio of the optimal strategy (note that the information ratio of the optimal portfolio is not sensitive to ζ_1). As $\zeta_2 \rightarrow 0$, the information ratio of the optimal strategy converges to that of the GOP; as $\zeta_2 \rightarrow \infty$, the information ratio of the optimal strategy decreases to that of the MQP. Note that if the information ratio of the GOP was less than that of the MQP, the information ratio of the optimal strategy would *increase* up to the MQP metric as ζ_2 increases. The left panel in Figure 3 shows the Sharpe ratio of the optimal portfolio for different $\boldsymbol{\zeta}$ values. The effect of varying these parameters on Sharpe ratio is more complex.

Finally, Figure 4 shows historgrams of the Sharpe ratio and information ratio of the optimal portfolio for $\boldsymbol{\zeta} = (0.5, 0.5, 0.5)$ compared to the MDP, GOP, MQP and the market portfolio. This gives a sense of the level of dispersion of these metrics for all portfolios.
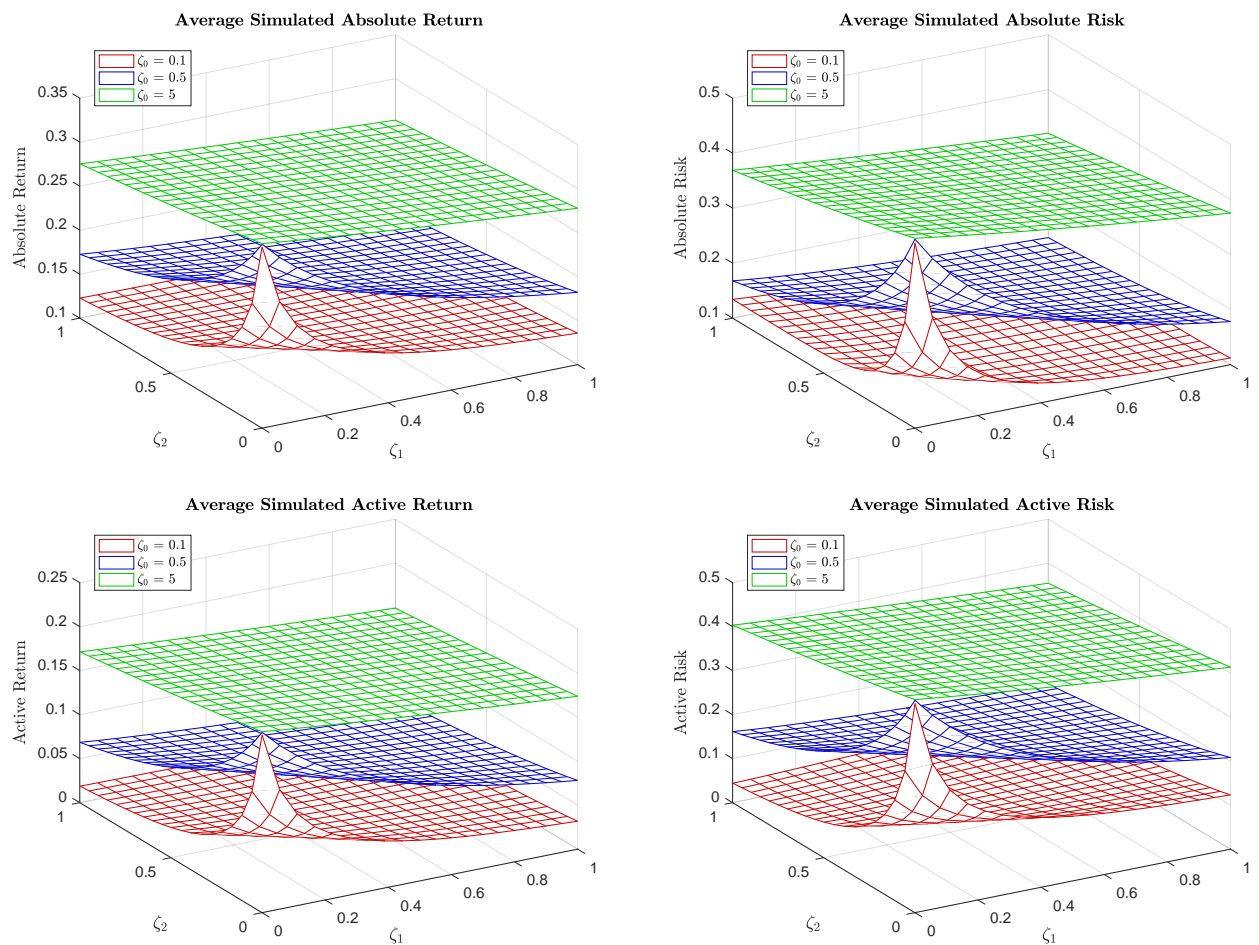
Figure 1: Surfaces of average simulated absolute return and risk (top panels) and active return and risk (bottom panels) for all ζ values; each point is the average metric value across all simulations for the optimal portfolio with the corresponding ζ.
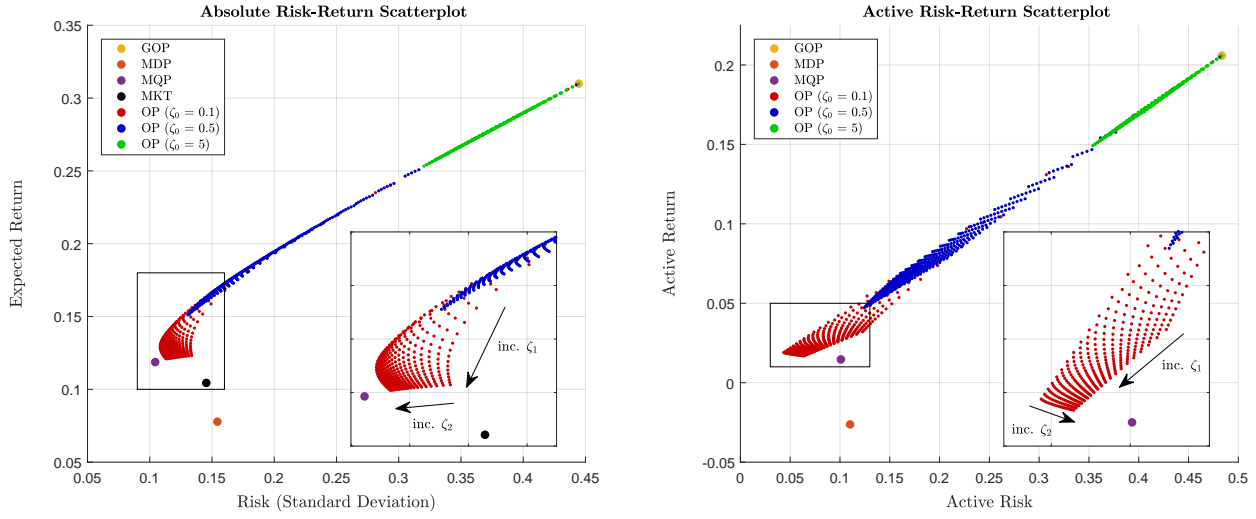
Figure 2: Scatterplots of absolute risk vs. absolute return (left panel) and active risk vs. active return (right panel) for the optimal portfolios with $\zeta_0 = 0.1$ (red), $\zeta_0 = 0.5$ (blue), $\zeta_0 = 5$ (green) and varying ζ_1, ζ_2 and GOP, MDP, MQP and market portfolio. Each dot is the average active risk/return for the corresponding portfolio across all simulations.
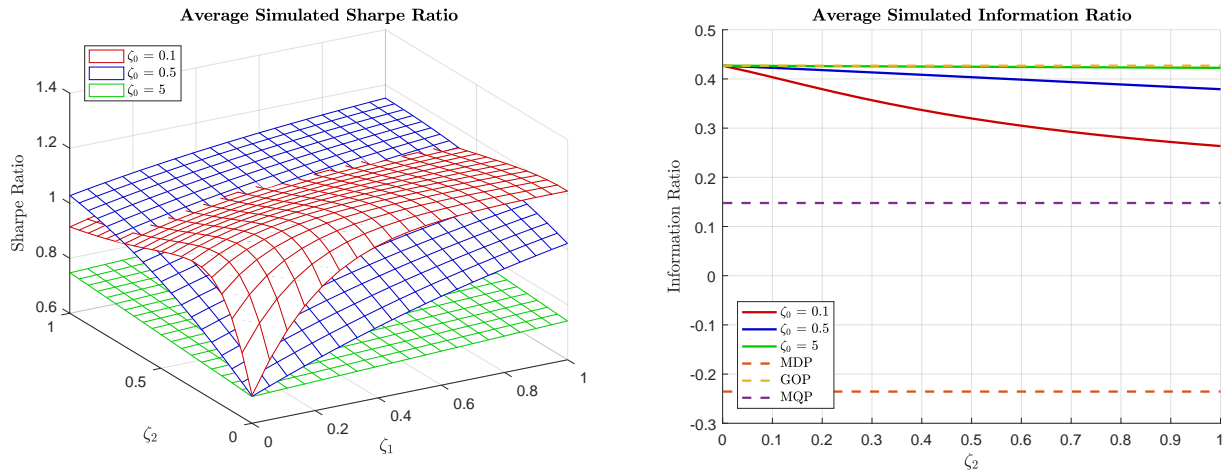

Figure 3: Left panel: Surfaces of average simulated Sharpe ratios for all $\boldsymbol{\zeta}$ values; each point on the surface is the average Sharpe ratio across all simulations for the optimal portfolio with the corresponding $\boldsymbol{\zeta}$. Right panel: average information ratio across simulations as a function of ζ_2 compared to MDP, GOP, MQP.
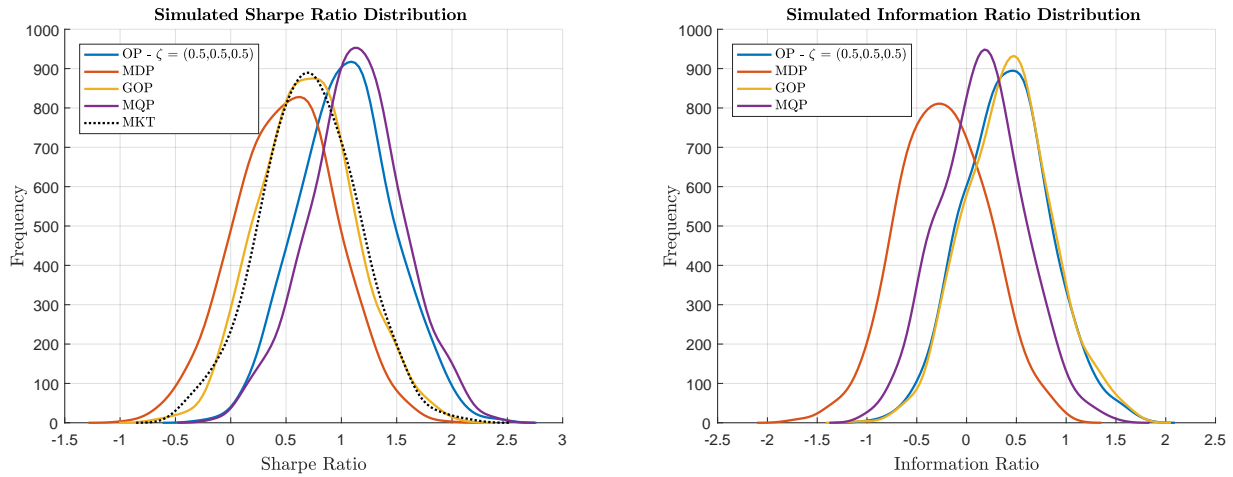
Figure 4: Distribution of simulated Sharpe ratios (left panel) and information ratios (right panel) for the optimal portfolio with $\boldsymbol{\zeta} = (0.5, 0.5, 0.5)$, MDP, GOP, MQP and the market portfolio.

6. Conclusions

In this article, we addressed the problem of how to dynamically allocate funds to beat a performance benchmark while remaining tethered to a tracking benchmark. An auxiliary penalty term is included to achieve additional goals, including controlling the quadratic variation of the wealth path and fine-tuning the allocation to individual securities. We related the solution to the growth optimal portfolio and minimum quadratic variance portfolio and discussed some of its limiting properties with respect to these portfolios. We furthermore provide a verification theorem demonstrating that our solution solves the optimal investment problem. Based on simulations of the market model, we show that the investor is able to control their risk-reward profile by tuning various tolerance parameters. Moreover, the optimal portfolios we find appear to outperform the static optimization of Oderda (2015) for all values of the penalty parameters in our simulation setting.

There are several future directions left open for research. One of these directions for which we already have preliminary results is related to incorporating latent information in the return and dividend process. Returns are notoriously difficult to estimate, and having a model which allows them to be stochastic, but also driven by latent factors, is essential to making the strategy robust to differing market environments. Extending the results to rank-based market models (e.g. Atlas models) that are prevalent in SPT literature is another interesting direction.

Appendix A   Proofs

A.1   Proof of Proposition 2

The PDE in (4.4) can be written as:

$$\begin{cases} \left(\partial_t + \mathcal{L}^{\boldsymbol{X}}\right) h + \sup_{\boldsymbol{\pi} \in \mathcal{A}} \{F(t, \boldsymbol{x}, \boldsymbol{\pi})\} = 0 \,, \\ \\ \qquad\qquad\qquad\qquad h(T, \boldsymbol{x}) = 0 \,. \end{cases} \tag{A.1}$$

where, using the definition of $a(t, \boldsymbol{\rho}, \boldsymbol{\pi})$ given in Proposition 1 along with Assumption 5, we have:

$$\begin{aligned} F(t, \boldsymbol{x}, \boldsymbol{\pi}) &= \zeta_0 \cdot a(t, \rho(t, \boldsymbol{x}), \boldsymbol{\pi}) - \tfrac{\zeta_1}{2} \cdot (\boldsymbol{\pi} - \eta(t, \boldsymbol{x}))' \boldsymbol{\Sigma}(t)(\boldsymbol{\pi} - \eta(t, \boldsymbol{x})) - \tfrac{\zeta_2}{2} \cdot \boldsymbol{\pi}' \mathbf{Q}(t) \boldsymbol{\pi} \\ &= \zeta_0 \cdot \left[\boldsymbol{\pi}' \boldsymbol{\alpha}(t) - \tfrac{1}{2} \boldsymbol{\pi}' \boldsymbol{\Sigma}(t) \boldsymbol{\pi} - \left(\boldsymbol{\gamma}_{\rho(t, \boldsymbol{x})}(t) + \boldsymbol{\delta}_{\rho(t, \boldsymbol{x})}(t) \right) \right] \\ &\quad - \tfrac{\zeta_1}{2} \cdot \left[\boldsymbol{\pi}' \boldsymbol{\Sigma}(t) \boldsymbol{\pi} - 2 \boldsymbol{\pi}' \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) + \eta(t, \boldsymbol{x})' \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) \right] - \tfrac{\zeta_2}{2} \cdot \boldsymbol{\pi}' \mathbf{Q}(t) \boldsymbol{\pi} \\ &= -\tfrac{1}{2} \boldsymbol{\pi}' \left[(\zeta_0 + \zeta_1) \boldsymbol{\Sigma}(t) + \zeta_2 \mathbf{Q}(t) \right] \boldsymbol{\pi} \\ &\quad + \boldsymbol{\pi}' \left[\zeta_0 \boldsymbol{\alpha}(t) + \zeta_1 \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) \right] - \zeta_0 \cdot \left(\boldsymbol{\gamma}_{\rho(t, \boldsymbol{x})}(t) + \boldsymbol{\delta}_{\rho(t, \boldsymbol{x})}(t) \right) - \tfrac{\zeta_1}{2} \cdot \eta(t, \boldsymbol{x})' \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) \\ &= -\tfrac{1}{2} \boldsymbol{\pi}' \mathbf{A}(t) \boldsymbol{\pi} + \boldsymbol{\pi}' \boldsymbol{B}(t, \boldsymbol{x}) - C(t, \boldsymbol{x}) \end{aligned}$$

which allows us to then rewrite the PDE above in the following concise form:

$$\begin{cases} \left(\partial_t + \mathcal{L}^{\boldsymbol{X}}\right) h - C(t, \boldsymbol{x}) + \sup_{\boldsymbol{\pi} \in \mathcal{A}} \left\{ -\tfrac{1}{2} \boldsymbol{\pi}' \mathbf{A}(t) \boldsymbol{\pi} + \boldsymbol{\pi}' \boldsymbol{B}(t, \boldsymbol{x}) \right\} = 0 \,, \\ \\ \qquad\qquad\qquad\qquad\qquad\qquad h(T, \boldsymbol{x}) = 0 \,. \end{cases} \tag{A.2}$$

where

$$\mathbf{A}(t) = (\zeta_0 + \zeta_1) \boldsymbol{\Sigma}(t) + \zeta_2 \mathbf{Q}(t) \,,$$
$$\text{\textcolor{red}{(n \times n)}}$$

$$\boldsymbol{B}(t, \boldsymbol{x}) = \zeta_0 \boldsymbol{\alpha}(t) + \zeta_1 \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) \,, \quad \text{and}$$
$$\text{\textcolor{red}{(n \times 1)}}$$

$$C(t, \boldsymbol{x}) = \zeta_0 \cdot \left(\boldsymbol{\gamma}_{\rho(t, \boldsymbol{x})}(t) + \boldsymbol{\delta}_{\rho(t, \boldsymbol{x})}(t) \right) + \tfrac{\zeta_1}{2} \cdot \eta(t, \boldsymbol{x})' \boldsymbol{\Sigma}(t) \eta(t, \boldsymbol{x}) \,.$$
$$\text{\textcolor{red}{(1 \times 1)}}$$

The next step is to solve for the optimal control by determining the first order conditions. The optimization problem of interest in this case is:

$$\max_{\boldsymbol{\pi}} \left\{ -\tfrac{1}{2} \boldsymbol{\pi}' \mathbf{A}(t) \boldsymbol{\pi} + \boldsymbol{\pi}' \boldsymbol{B}(t, \boldsymbol{x}) \right\} \quad \text{subject to} \quad \mathbf{1}' \boldsymbol{\pi} = 1 \,.$$

This can be solved using a Lagrange multiplier to find that the optimal control is

$$\boldsymbol{\pi}_{\zeta}^*(t) = \mathbf{A}^{-1}(t) \cdot \left[\frac{1 - \mathbf{1}' \mathbf{A}^{-1}(t) \boldsymbol{B}(t, \boldsymbol{x})}{\mathbf{1}' \mathbf{A}^{-1}(t) \mathbf{1}} \cdot \mathbf{1} + \boldsymbol{B}(t, \boldsymbol{x}) \right] = \mathbf{A}^{-1}(t) \cdot \boldsymbol{d}(t, \boldsymbol{x}) \,, \tag{A.3}$$

where $d(t, x) = \frac{1 - \mathbf{1}'\mathbf{A}^{-1}(t)B(t,x)}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}} \cdot \mathbf{1} + B(t, x)$. Note that $\mathbf{A}$ is positive definite, so the optimal control is

$\underset{(n\times 1)}{}$

indeed a maximizer. Next, notice that we can write the quadratic term in the PDE above (suppressing the dependence on t and x for notational simplicity) as follows:

$$-\tfrac{1}{2}\left(\boldsymbol{\pi}_\zeta^*\right)' \mathbf{A}\boldsymbol{\pi}_\zeta^* + \left(\boldsymbol{\pi}_\zeta^*\right)' \boldsymbol{B} = -\tfrac{1}{2}\boldsymbol{d}'\mathbf{A}^{-1}\cdot\mathbf{A}\cdot\mathbf{A}^{-1}\cdot\boldsymbol{d} + \boldsymbol{d}'\mathbf{A}^{-1}\boldsymbol{B} \qquad \text{(note that } \mathbf{A} \text{ and } \mathbf{A}^{-1} \text{ are symmetric)}$$

$$= \boldsymbol{d}'\mathbf{A}^{-1}\left[-\tfrac{1}{2}\boldsymbol{d} + \boldsymbol{B}\right]$$

$$= -\tfrac{1}{2}\left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B}}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}} \cdot \mathbf{1} + \boldsymbol{B}\right]' \mathbf{A}^{-1}\left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B}}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}} \cdot \mathbf{1} - \boldsymbol{B}\right] \qquad (A.4)$$

Therefore, the PDE in (4.4) becomes:

$$\begin{cases} \left(\partial_t + \mathcal{L}^{\boldsymbol{X}}\right) h - G(t, \boldsymbol{x}) = 0 \,, \\[2mm] h(T, \boldsymbol{x}) = 0 \,, \end{cases} \qquad (A.5)$$

where

$$G(t, \boldsymbol{x}) = C(t, \boldsymbol{x}) + \frac{1}{2}\left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}(t)B(t,\boldsymbol{x})}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}} \cdot \mathbf{1} + B(t,\boldsymbol{x})\right]' \mathbf{A}^{-1}(t)\left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}(t)B(t,\boldsymbol{x})}{\mathbf{1}'\mathbf{A}^{-1}(t)\mathbf{1}} \cdot \mathbf{1} - B(t,\boldsymbol{x})\right] \,.$$

The remainder of the proof relies on finding a Feynman-Kac representation for the solution to the PDE above. It is slightly more convenient at this point to perform a change of variables and use $\log \boldsymbol{X}$ as the state variables. Note that this mainly affects the infinitesimal generator in the PDE above. Namely, we have:

$$\mathcal{L}^{\boldsymbol{X}} = \tfrac{1}{2}\sum_{i,j=1}^{n} \underbrace{\boldsymbol{\Sigma}_{ij}(t)}_{=\,a_{ij}} \cdot \partial_{ij} + \sum_{i=1}^{n} \underbrace{(\gamma_i(t) + \delta_i(t))}_{=\,b_i} \cdot \partial_i$$

First, we establish the existence of a sufficiently smooth solution to the PDE (A.5). For this, we rely on the sufficient conditions given in Remark 7.8 in Chapter 5 of Karatzas and Shreve (1998), namely: (i) uniform ellipticity of the PDE on $[0, \infty) \times \mathbb{R}^n$; (ii) boundedness of a_{ij}, b_i in $[0, T] \times \mathbb{R}^n$; (iii) uniform Hölder continuity of a_{ij}, b_i and G in $[0, T] \times \mathbb{R}^n$; (iv) G satisfies $|G(t, \boldsymbol{x})| \leq K(1 + \|\boldsymbol{x}\|^{2\lambda})$ for some $K > 0, \lambda \geq 1$.

The first condition holds by the nondegeneracy assumption we made on the covariance matrix $\boldsymbol{\Sigma}$. For the second condition, we have assumed that $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ are bounded above and below for all t and we can use the nondegeneracy and bounded variance assumptions made in Section 3 to obtain upper and lower bounds on the elements of $\boldsymbol{\Sigma}(t)$. In particular, since we have that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and any $t \geq 0$:

$$\varepsilon_1\|\mathbf{x} + \mathbf{y}\|^2 \leq (\mathbf{x} + \mathbf{y})'\boldsymbol{\Sigma}(t)(\mathbf{x} + \mathbf{y}) \leq M_1\|\mathbf{x} + \mathbf{y}\|^2$$

$$\implies \varepsilon_1\|\mathbf{x} + \mathbf{y}\|^2 \leq \mathbf{x}'\boldsymbol{\Sigma}(t)\mathbf{x} + \mathbf{y}'\boldsymbol{\Sigma}(t)\mathbf{y} + 2\mathbf{x}'\boldsymbol{\Sigma}(t)\mathbf{y} \leq M_1\|\mathbf{x} + \mathbf{y}\|^2$$

$$\implies \tfrac{1}{2}\left(\varepsilon_1\|\mathbf{x} + \mathbf{y}\|^2 - M_1\left(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2\right)\right) \leq \mathbf{x}'\boldsymbol{\Sigma}(t)\mathbf{y} \leq M_1\|\mathbf{x} + \mathbf{y}\|^2$$

27

Taking $\mathbf{x} = \mathbf{e}_i$ and $\mathbf{y} = \mathbf{e}_j$, where $\mathbf{e}_i$ is vector of zeros with a 1 in the i^{th} position, we obtain the following bound on the covariance elements:

$$\varepsilon_1 - M_1 \;\leq\; \mathbf{\Sigma}_{ij}(t) \;\leq\; 4M_1 \;.$$

This shows that condition (ii) is satisfied. Next, we show that the function G is bounded and hence satisfies the polynomial growth condition (iv), which will also be needed to prove that condition (iii) holds. For this, we note that since the elements of $\mathbf{\Sigma}$ along with $\rho(t, \boldsymbol{x})$ and $\eta(t, \boldsymbol{x})$ are bounded, we conclude that the functions $\boldsymbol{B}(t, \boldsymbol{x})$ and $C(t, \boldsymbol{x})$ are bounded. Now, it suffices to show that quadratic forms $\mathbf{x}'\mathbf{A}^{-1}(t)\mathbf{y}$ for $\mathbf{x}, \mathbf{y}$ in some bounded subset of $\mathbb{R}^n$ are also bounded. To this end, recall that $\mathbf{A}(t) = (\zeta_0 + \zeta_1)\mathbf{\Sigma}(t) + \zeta_2\mathbf{Q}(t)$ and that for all $\mathbf{x} \in \mathbb{R}^n$:

$$\varepsilon_1\|\mathbf{x}\|^2 \;\leq\; \mathbf{x}'\mathbf{\Sigma}(t)\mathbf{x} \;\leq\; M_1\|\mathbf{x}\|^2$$
$$\varepsilon_2\|\mathbf{x}\|^2 \;\leq\; \mathbf{x}'\mathbf{Q}(t)\mathbf{x} \;\leq\; M_2\|\mathbf{x}\|^2$$

which implies that there exists $\epsilon, M > 0$ such that for all $\mathbf{x} \in \mathbb{R}^n$:

$$\epsilon\|\mathbf{x}\|^2 \;\leq\; \mathbf{x}'\mathbf{A}(t)\mathbf{x} \;\leq\; M\|\mathbf{x}\|^2$$

Denoting the eignevalues of $\mathbf{A}(t)$ by $\lambda_1(t) \geq ... \geq \lambda_n(t)$, we note that the bounds above are tight when we replace M and ε with λ_1 and λ_n, respectively. Thus, the eigenvalues satisfy:

$$\varepsilon \leq \lambda_n(t) \leq \lambda_1(t) \leq M \qquad \Longrightarrow \qquad \frac{1}{M} \leq \frac{1}{\lambda_1(t)} \leq \frac{1}{\lambda_n(t)} \leq \frac{1}{\varepsilon}$$

from which we can conclude there exist positive constants $\tilde{\varepsilon}, \tilde{M}$ such that for any $\mathbf{x} \in \mathbb{R}^n$:

$$\tilde{\varepsilon}\|\mathbf{x}\|^2 \;\leq\; \mathbf{x}'\mathbf{A}^{-1}(t)\mathbf{x} \;\leq\; \tilde{M}\|\mathbf{x}\|^2$$

and by similar reasoning as above we can obtain:

$$\tfrac{1}{2}\left(\tilde{\varepsilon}\|\mathbf{x} + \mathbf{y}\|^2 - \tilde{M}\left(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2\right)\right) \;\leq\; \mathbf{x}'\mathbf{A}^{-1}(t)\mathbf{y} \;\leq\; \tilde{M}\|\mathbf{x} + \mathbf{y}\|^2$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. When $\mathbf{x}, \mathbf{y}$ are in a bounded subset of $\mathbb{R}^n$, the norms on either side of the inequality above are also bounded. Combining this with the fact that $\boldsymbol{B}$ and C are bounded we can conclude that G is in fact bounded (since we are concerned with quadratic forms involving $\mathbf{1}$ and $\boldsymbol{B}$), and hence satisfies the polynomial growth condition (iv).

To prove condition (iii) we note that the assumption of differentiability on γ_i, δ_i and $\xi_{i\nu}$ implies that $a_{ij} = \mathbf{\Sigma}_{ij} = \sum_{\nu=1}^k \xi_{i\nu}\xi_{j\nu}$ and $b_i = \gamma_i + \delta_i$ are differentiable and hence uniformly Hölder continuous. Similarly, G can be shown to be differentiable and hence uniformly Hölder continuous on $[0, T] \times \mathbb{R}^n$ by

noting that $\mathbf{A}$, $\boldsymbol{B}$ and C are all bounded and differentiable which implies the same for G. Since conditions (i)-(iv) are satisfied, a solution exists to the PDE (A.5) and, by Theorem 7.6 in Chapter 5 of Karatzas and Shreve (1998), the unique solution is given by the following Feynman-Kac representation:

$$h(t, \boldsymbol{x}) = \mathbb{E}_{t,\boldsymbol{x}} \left[-\int_t^T G(s, \boldsymbol{X}(s)) \, ds \right], \tag{A.6}$$

where the expectation is taken under the physical measure $\mathbb{P}$. ∎

A.2 Proof of Theorem 1

Along the lines of Theorem 5.1 of Touzi (2012), we provide a verification argument to demonstrate that the candidate solution given in Proposition 2 is in fact the value function, and that the corresponding control given in (4.6) is the optimal control. Since $h(t, \boldsymbol{x})$ in (4.5) is a classical solution to the HJB equation, and $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*(t)$ is the maximizer of

$$\boldsymbol{\pi} \longmapsto F(t, \boldsymbol{x}, \boldsymbol{\pi}),$$

it suffices to check that the process $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^* = \left(\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*(t) \right)_{t \geq 0}$ is a well-defined admissible control process and that the controlled stochastic differential equation

$$dY^{\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*, \boldsymbol{\rho}}(t) = a\left(t, \boldsymbol{\rho}(t), \boldsymbol{\pi}_{\boldsymbol{\zeta}}^*(t)\right) \, dt + \left(\boldsymbol{\xi}_{\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t) \right)' \, d\boldsymbol{W}(t),$$

defines a unique solution for each given initial data.

First, we verify that $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*$ is an admissible control. This requires us to show that $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*$ is an $\mathfrak{F}$-adapted vector-valued process with $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*(t)'\mathbf{1} = 1$ for all t and that $\pi_i^*(t)$ is bounded almost surely for all t and for $i = 1, ..., n$. Clearly, $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*$ is $\mathfrak{F}$-adapted and its elements sum up to 1 from the following observation:

$$
\begin{aligned}
\mathbf{1}'\boldsymbol{\pi}_{\boldsymbol{\zeta}}^* &= \mathbf{1}'\mathbf{A}^{-1} \cdot \left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B}}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}} \cdot \mathbf{1} + \boldsymbol{B} \right] \\
&= \frac{1 - \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B}}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}} \cdot \mathbf{1}'\mathbf{A}^{-1}\mathbf{1} + \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B} \\
&= 1 - \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B} + \mathbf{1}'\mathbf{A}^{-1}\boldsymbol{B} \\
&= 1 .
\end{aligned}
$$

The fact that $\boldsymbol{\pi}_{\boldsymbol{\zeta}}^*$ is bounded follows from the boundedness of $\boldsymbol{B}$ and the elements of $\mathbf{A}^{-1}$ which was established in the previous proof. Next, we need to show the existence and uniqueness of the solution to the SDE given in the outset of this proof. To this end, note that the drift and volatility terms for any portfolio $\boldsymbol{\pi}$ are $a(t, \boldsymbol{\rho}(t), \boldsymbol{\pi}(t)) = \boldsymbol{\pi}'\boldsymbol{\alpha}(t) - \frac{1}{2}\boldsymbol{\pi}'\boldsymbol{\Sigma}(t)\boldsymbol{\pi} - (\gamma_{\boldsymbol{\rho}}(t) + \delta_{\boldsymbol{\rho}}(t))$ and $\left(\boldsymbol{\xi}_{\boldsymbol{\pi}}(t) - \boldsymbol{\xi}_{\boldsymbol{\rho}}(t) \right)$ respectively. Since any admissible $\boldsymbol{\pi}$, $\boldsymbol{\rho}$, $\boldsymbol{\gamma}$, $\boldsymbol{\delta}$ and the elements of $\boldsymbol{\Sigma}(t)$ are all bounded, so too is the drift term $a(t, \boldsymbol{\rho}(t), \boldsymbol{\pi}(t))$.

In particular it is square integrable for any $t \geq 0$. Moreover, since $\boldsymbol{\xi}(t)$ is assumed to be square integrable, so is the volatility term. Finally, since neither the drift nor volatility depends on $Y^{\boldsymbol{\pi},\boldsymbol{\rho}}$, both satisfy the required Lipschitz condition given in Theorem 2.2 of Touzi (2012). Then, by the same theorem, the SDE admits a unique solution for any choice of initial data $Y^{\boldsymbol{\pi},\boldsymbol{\rho}}(0)$. $\blacksquare$

A.3   Proof of Corollary 3

We first prove (i). Recall that the optimal portfolio is given by:

$$\boldsymbol{\pi}_{\zeta}^{*} = \mathbf{A}^{-1}\left[\frac{1 - \mathbf{1}'\mathbf{A}^{-1}\mathbf{B}}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}} \cdot \mathbf{1} + \mathbf{B}\right] = \left(1 - \mathbf{1}'\mathbf{A}^{-1}\mathbf{B}\right) \cdot \frac{1}{\mathbf{1}'\mathbf{A}^{-1}\mathbf{1}}\mathbf{A}^{-1}\mathbf{1} + \mathbf{A}^{-1}\mathbf{B}$$

where $\mathbf{A} = (\zeta_0 + \zeta_1)\boldsymbol{\Sigma} + \zeta_2\mathbf{Q}$ and $\mathbf{B} = \zeta_0\boldsymbol{\alpha} + \zeta_1\boldsymbol{\Sigma}\boldsymbol{\eta}$ and the dependence on t and $\boldsymbol{x}$ is dropped for brevity. Now, we can write $\mathbf{A}$ in the following two ways:

$$\mathbf{A} = (\zeta_0 + \zeta_1)\left(\boldsymbol{\Sigma} + \frac{\zeta_2}{\zeta_0 + \zeta_1}\mathbf{Q}\right) \qquad\qquad \mathbf{A} = \zeta_2\left(\frac{\zeta_0 + \zeta_1}{\zeta_2}\boldsymbol{\Sigma} + \mathbf{Q}\right)$$

$$= (\zeta_0 + \zeta_1)\mathbf{A}_1 \qquad\qquad\qquad\qquad = \zeta_2\mathbf{A}_2$$

from which it follows that:

$$\lim_{\zeta_0 \to \infty} \mathbf{A}_1 = \lim_{\zeta_1 \to \infty} \mathbf{A}_1 = \boldsymbol{\Sigma} \qquad \text{and} \qquad \lim_{\zeta_2 \to \infty} \mathbf{A}_2 = \mathbf{Q}$$

Additionally, we can write the inverse of $\mathbf{A}$ as follows:

$$\mathbf{A}^{-1} = \frac{1}{\zeta_0 + \zeta_1}\mathbf{A}_1^{-1} = \frac{1}{\zeta_2}\mathbf{A}_2^{-1}$$

Writing the optimal portfolio in terms of $\mathbf{A}_1$ we have:

$$\boldsymbol{\pi}_{\zeta}^{*} = \left[1 - \frac{1}{\zeta_0 + \zeta_1}\mathbf{1}'\mathbf{A}_1^{-1}\left(\zeta_0\boldsymbol{\alpha} + \zeta_1\boldsymbol{\Sigma}\boldsymbol{\eta}\right)\right] \cdot \frac{1}{\mathbf{1}'\mathbf{A}_1^{-1}\mathbf{1}}\mathbf{A}_1^{-1}\mathbf{1} + \frac{1}{\zeta_0 + \zeta_1}\mathbf{A}_1^{-1}\left(\zeta_0\boldsymbol{\alpha} + \zeta_1\boldsymbol{\Sigma}\boldsymbol{\eta}\right)$$

$$= \left[1 - \frac{\zeta_0}{\zeta_0 + \zeta_1}\mathbf{1}'\mathbf{A}_1^{-1}\boldsymbol{\alpha} - \frac{\zeta_1}{\zeta_0 + \zeta_1}\mathbf{1}'\mathbf{A}_1^{-1}\boldsymbol{\Sigma}\boldsymbol{\eta}\right] \cdot \frac{1}{\mathbf{1}'\mathbf{A}_1^{-1}\mathbf{1}}\mathbf{A}_1^{-1}\mathbf{1} + \frac{\zeta_0}{\zeta_0 + \zeta_1}\mathbf{A}_1^{-1}\boldsymbol{\alpha} + \frac{\zeta_1}{\zeta_0 + \zeta_1}\mathbf{A}_1^{-1}\boldsymbol{\Sigma}\boldsymbol{\eta}$$

Now, noting that:

$$\lim_{\zeta_0 \to \infty} \frac{\zeta_0}{\zeta_0 + \zeta_1} = \lim_{\zeta_1 \to \infty} \frac{\zeta_1}{\zeta_0 + \zeta_1} = 1$$

$$\lim_{\zeta_0 \to \infty} \frac{\zeta_1}{\zeta_0 + \zeta_1} = \lim_{\zeta_1 \to \infty} \frac{\zeta_0}{\zeta_0 + \zeta_1} = 0$$

it follows that:

$$\lim_{\zeta_0 \to \infty} \boldsymbol{\pi}_{\zeta}^{*} = \boldsymbol{\pi}_{GOP} \qquad , \qquad \lim_{\zeta_1 \to \infty} \boldsymbol{\pi}_{\zeta}^{*} = \boldsymbol{\eta} \ .$$

Similarly, writing the optimal portfolio in terms of $\mathbf{A}_2$ we have:

$$\boldsymbol{\pi}_{\zeta}^{*} = \left[1 - \frac{\zeta_0}{\zeta_2}\mathbf{1}'\mathbf{A}_2^{-1}\boldsymbol{\alpha} - \frac{\zeta_1}{\zeta_2}\mathbf{1}'\mathbf{A}_2^{-1}\boldsymbol{\Sigma}\boldsymbol{\eta}\right] \cdot \frac{1}{\mathbf{1}'\mathbf{A}_2^{-1}\mathbf{1}}\mathbf{A}_2^{-1}\mathbf{1} + \frac{\zeta_0}{\zeta_2}\mathbf{A}_2^{-1}\boldsymbol{\alpha} + \frac{\zeta_1}{\zeta_2}\mathbf{A}_2^{-1}\boldsymbol{\Sigma}\boldsymbol{\eta}$$

from which it follows that:

$$\lim_{\zeta_2 \to \infty} \boldsymbol{\pi}_\zeta^* = \frac{1}{\mathbf{1}'\mathbf{Q}^{-1}\mathbf{1}} \mathbf{Q}^{-1}\mathbf{1} \ .$$

Statements (ii) and (iii) can be verified directly by substituting in the appropriate values of $\zeta_0, \zeta_1, \zeta_2$ and $\mathbf{Q}$ and noting that the GOP is given by $\boldsymbol{\pi}_{GOP} = (1 - \mathbf{1}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}) \cdot \boldsymbol{\pi}_{MQP} + \boldsymbol{\Sigma}^{-1}\boldsymbol{\alpha}.$ ∎

References

Ang, A., D. Papanikolaou, and M. M. Westerfield (2014). Portfolio choice with illiquid assets. *Management Science 60*(11), 2737–2761.

Banner, A. D. and D. Fernholz (2008). Short-term relative arbitrage in volatility-stabilized markets. *Annals of Finance 4*(4), 445–454.

Blanchet-Scalliet, C., N. E. Karoui, M. Jeanblanc, and L. Materllini (2008). Optimal investment decisions when time-horizon is uncertain. *Journal of Mathematical Economics 44*(11), 1100–1113.

Browne, S. (1997). Survival and growth with a liability: Optimal portfolio strategies in continuous time. *Mathematics of Operations Research 22*(2), 468–493.

Browne, S. (1999a). Beating a moving target: Optimal portfolio strategies for outperforming a stochastic benchmark. *Finance and Stochastics 3*(3), 275–294.

Browne, S. (1999b). Reaching goals by a deadline: Digital options and continuous-time active portfolio management. *Advances in Applied Probability 31*(2), 551–577.

Browne, S. (2000). Risk-constrained dynamic active portfolio management. *Management Science 46*(9), 1188–1199.

Christensen, M. M. (2012). On the history of the growth-optimal portfolio. In *Machine Learning For Financial Engineering*, pp. 1–79. World Scientific.

Cvitanić, J. and I. Karatzas (1992). Convex duality in constrained portfolio optimization. *The Annals of Applied Probability*, 767–818.

Davis, M. and A. Norman (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research 15*(4), 676–713.

Fernholz, E. R., I. Karatzas, and J. Ruf (2016). Volatility and arbitrage. *arXiv preprint arXiv:1608.06121*.

Fernholz, R. (1999a). On the diversity of equity markets. *Journal of Mathematical Economics 31*(3), 393–417.

Fernholz, R. (1999b). Portfolio generating functions. In *Quantitative Analysis in Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar*, pp. 344. World Scientific.

Fernholz, R. (2001). Equity portfolios generated by functions of ranked market weights. *Finance and Stochastics 5*(4), 469–486.

Fernholz, R. (2002). *Stochastic portfolio theory*. Springer.

Fernholz, R. and I. Karatzas (2005). Relative arbitrage in volatility-stabilized markets. *Annals of Finance 1*(2), 149–177.

Fernholz, R. and B. Shay (1982). Stochastic portfolio theory and stock market equilibrium. *The Journal of Finance 37*(2), 615–624.

Karatzas, I. and R. Fernholz (2009). Stochastic portfolio theory: an overview. *Handbook of numerical analysis 15*, 89–167.

Karatzas, I. and S. Shreve (1998). *Brownian motion and stochastic calculus* (2 ed.).

Liu, R. and J. Muhle-Karbe (2013). Portfolio choice with stochastic investment opportunities: a user's guide. *arXiv preprint arXiv:1311.1715*.

Magill, M. and G. Constantinides (1976). Portfolio selection with transaction costs. *J. Economic Theory 13*, 245–263.

Merton, R. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics LI*, 247–257.

Merton, R. (1971). Optimal consumption and portfolio rules in a continuous-time model. *J. Economic Theory 3*, 373–413.

Oderda, G. (2015). Stochastic portfolio theory optimization and the origin of rule-based investing. *Quantitative Finance 15*(8), 1259–1266.

Pal, S. and T.-K. L. Wong (2013). Energy, entropy, and arbitrage. *arXiv preprint arXiv:1308.5376*.

Pal, S. and T.-K. L. Wong (2016). The geometry of relative arbitrage. *Mathematics and Financial Economics 10*(3), 263–293.

Samo, Y.-L. K. and A. Vervuurt (2016). Stochastic portfolio theory: A machine learning perspective. *arXiv preprint arXiv:1605.02654*.

Touzi, N. (2012). *Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE*. Volume 29 of Fields Institute Monographs. Springer New York.

Wong, T.-K. L. (2015). Optimization of relative arbitrage. *Annals of Finance 11*(3-4), 345–382.