University Of Bergamo

Doctoral Thesis

Optimal Asset Allocation, a Data-Driven Feedback Control Approach

---

Author:
Gabriele Maroni

Supervisor:
Prof. Fabio Previdi

Co-supervisor:
Prof. Simone Formentin

Abstract

Gabriele MARONI

Optimal Asset Allocation, a Data-Driven Feedback Control Approach

Typical financial problems have recently attracted the attention of the system control community, in fact many of these can be formulated as closed-loop control problems. In this thesis we deal, in particular, with two research lines: *reactive trading* and *multi-period portfolio optimization.*

Reactive trading is an innovative technical analysis approach related to the trading of individual financial assets via feedback control. The main innovation of reactive trading is to be a *model-free* approach in which stock prices are considered as external stochastic disturbances and the control system must act against them to provide certain performance guarantees. The theory comes with a trading scheme called Simultaneous-Long-Short (SLS) controller, which guarantees, under certain market hypotheses, positive expected value of the gain function regardless of market direction. However, SLS presents two problems: first of all, the original trading scheme use a static gain as the controller whose value is calibrated on in-sample data and validated by backtesting. However, stock prices are known to have dynamics that change, even suddenly, due to socio-economic events, to which a time-invariant controller may not react properly. Secondly, performance guarantees are proved by making strong assumptions on markets where price dynamics are governed by simple processes such as Geometric Brownian Motion (GBM), but real prices have non-ideal characteristics that make them more complex.

Motivated by these problems in this thesis the author proposes two solutions. As solution to the first problem a new control scheme with time-varying controllers based on the logic of an adaptive control approach called Extremum Seeking (ES). The ES is particularly suitable for reactive trading because it maximizes the output of a system without requiring knowledge of the process and adapting the value of the controller through an online estimate of the output gradient with respect to the control action.

As solution to the second problem reactive trading is reformulated under the sole assumption that the returns of the prices of the considered stock are uncertain but bounded between a range of values. Thanks to this assumption it is possible to treat the uncertainty of returns as an uncertain parameter of the system and reformulate the problem of stock trading as a robust control problem and to synthesize controllers with the H_∞ approach which guarantee robust performance in the range of variation of returns.

Multi-period portfolio optimization is an extension of the classic single-period portfolio optimization approach dating back to Markowitz. In the financial literature these problems are treated as stochastic control problems and solved by dynamic programming. However, if many assets, long investment horizons and constraints are considered, these methods fail to find a solution due to the 'curse of dimensioanily'. For this reason, recently methods based on predictive control have been very successful in solving multi-period problems. In fact, although they provide an approximate solution, they are able to naturally constrain the composition and can often be resolved efficiently even in the case of many assets and long time horizons, although they themselves are subject to the 'curse of dimensionality', in particular in their non-linear form. Of particular interest are the approaches that introduce closed-loop investment policies capable of fully exploiting the dynamic nature of the problem by employing observations on past market behavior. In the literature, until now, only linearly parameterized control policies have been implemented, although the complex dynamics of the markets. For this reason, in this thesis, statistical learning techniques based on kernel methods are used to extend multi-period investment policies to classes of non-linear functions that are better able to exploit the information hidden in the observations of past markets.

Acknowledgements

My sincere thanks are directed to the members of my research group, the Control System and Automation Laboratory of the University of Bergamo with whom I shared these doctoral years. First of all my gratitude goes to my advisor, Prof. Fabio Previdi, who not only made me grow as a scientist giving me the opportunity to face challenging projects but also made me grow as a man, thanks to his teachings and his advice. I warmly thank Prof. Simone Formentin, who has always made himself available in these years to answer my questions and help me with my research, and from which I learned a lot. I thank my colleagues and friends, Prof. Mirko Mazzoleni, Prof. Matteo Scandella and Prof. Michele Ermidoro, with them I had daily exchanges of ideas and opinions and I faced the labors of doctoral activities.

Finally I thank all my family, and in particular my mother, who has always supported me from the first years of school education up to now, they have never made me miss their support, their comfort and their help every time I needed it.

Contents

Chapter 1

Introduction

The discipline of finance and control engineering seems to be unrelated to each other. However, by analyzing the definitions of these two disciplines the connecting points are obvious [55]. Control engineering, in its essence, is the discipline that deals with controlling a process in such a way that it behaves as desired, although the presence of unforeseeable and uncontrollable external disturbances acting on it. Finance is the study of how people and organizations allocate scarce resources over time through investments, subject to uncertainty [19]. Therefore, from a controlling point of view, an investor aims to control the state of his resources through control actions, investments, aimed at dealing with uncertain economic and political events.

The generic term investment refers to the commitment of resources with the ultimate aim of obtaining a future benefit or increasing personal satisfaction. With regard to financial markets, an investment is a commitment of resources, typically in the form of money, which aims to generate a monetary profit for the investor, be it an individual, a group of people, an organization [18]. In financial markets, money investments are made through the trading of assets.

The term “asset” is a rather generic term that refers to anything that can be turned into money by the owner [22]. Assets can be divided into three main categories: real, financial or intangible [18]. Investing in a real asset means investing in physical entities such as substances or properties. Examples of real assets include precious metals, real estate, land and commodities such as sugar, oil, electricity. Intangible assets, on the other hand, do not have a physical nature, examples are patents, trademarks, intellectual property. Finally, financial assets are claims of ownership of an entity like a company, or contractual right to payments. Examples are cash, equity securities (e.g., common stocks), debt securities (e.g., bonds), derivatives (e.g., futures, options), mutual funds and bank deposits.

In this dissertation, we focus on financial assets, since they are among the three categories, the most liquid one, where liquidity is the ability to convert assets into cash quickly. Among financial assets, we mainly focus on risky financial assets tradable on the market such as stock shares. The risk component is due to the fact that the future price of these financial instruments is not known with certainty, unlike a bank deposit which guarantees a safe interest rate.

Investment strategies in financial markets can be roughly divided into three macro-categories: *fundamental analysis*, *technical analysis* and *quantitative analysis* [51]. Fundamental analysis quantifies the intrinsic value of company's securities based on economic factors such as earnings, dividends, and other indicators about the health of that company [67]. On the other hand, technical analysis looks for trading opportunities based on statistical analysis of the signals generated during trading activities such as past price patterns and traded volumes [69]. Quantitative analysis is the use of mathematical models and large datasets to analyze financial markets and securities. Common examples include the pricing of derivative securities such as options, and risk management, especially as it relates to portfolio management applications [112]. Technical analysis and quantitative analysis have some degree of overlap. In this thesis we distinguish the fact that typically technical analysis focuses on single assets or pairs of assets (eg, pairs trading), while quantitative analysis has as a basic principle the reduction of the investment risk through the management of a portfolio of diversified assets.

This thesis describes the author's contributions to two lines of research. The first falls under the umbrella of technical analysis and describes a new paradigm for stock trading via a model-free feedback controller and is called *reactive trading* [10]. The second is, instead, a quantitative analysis research line that extends the classical single-period approach of portfolio optimization dating back to Harry Markowitz [73] to a multi-period framework, which models the portfolio as a dynamic system that evolves over time, subject to the uncertainties deriving from changes in market prices and on which the investor can act through an appropriate allocation strategy [29].

1.1 Reactive trading

Among the most widely used Technical Analysis methods, we could cite the exploitation of certain indicators computable from past data, such as the Moving Average (MA), the Relative Strenght Index (RSI) and the Bollinger bands [69]. Other methodologies have been proposed by time series analysts, e.g., Auto-Regressive Moving Average (ARMA) models are sometimes used to compute the expected value of an asset price, while Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) models are employed to describe and predict the asset volatility [107].

More recently, machine learning algorithms have been also widely employed, e.g., to provide more complex nonlinear models of the stock prices [111], to predict the sign of the returns of the price [63] and to combine information from past market data and textual information deriving from financial news [40].

In the control systems community, the problem of single stock trading has been sometimes reformulated as a control problem, where the price dynamics is typically described as a stochastic differential equation like the Ito process, see, e.g., [59], [28], or as a regime switching model [36].

A common feature of all the above trading approaches is that they are *model-based*, in that they rely on a "model" of the stock price dynamics. However, the accurate estimation of the price model from past data is all but a trivial task, and modeling errors may lead to detrimental performances.

These observations led Barmish, in [8] and later in [10], to introduce a new revolutionary approach to stock trading, the *reactive trading* approach, aimed at studying *model-free* strategies based on feedback control theory, where buy and sell signals are generated on the basis of gain-loss performance rather than a parametrized model of the prices. The innovative feature of this approach lies in the fact that the price of the considered stock is seen as an exogenous disturbance to reject, rather than a stochastic process to be modelled (this approach is typical of the control literature in which we try to guarantee robustness properties of the closed-loop system against external disturbances). Barmish develops a trading scheme called Simultaneous Long Short controller (SLS), equipped with two linear feedbacks running in parallel and capable of implementing a long and a short investment strategy at the same time. A key feature of such a trading scheme is that, even within such a poor information environment, under some market assumptions and the assumption that stock price follows simple dynamics (e.g., Geometric Brownian Motion (GBM)), the profit is guaranteed to have positive expected value regardless of the direction of the market. This result is formalized by the so called Robust Positive Expectation (RPE) Theorem [10].

However, the reactive trading approach presents two major open problems that will be addressed in this thesis. The first problem lies in the way of tuning the controller parameters and is explained as follows. The controllers in the feedback loops of the original SLS scheme are static gains. Later, in [72], the authors generalize the static controllers to dynamic controllers with Proportional Integral (PI) action. However, both in the case of static and dynamic controllers, the controller parameters must be tuned on the basis of backtesting simulations and then kept fixed. The problem with this strategy is that stock prices, as documented in the literature [64, 98, 84], are strongly non-stationary and subject to sudden changes in their dynamics due to socio-economic events. This could make the current controller parameter values inadequate.

A first solution could be to frequently recalibrate the parameters of the controller. Another solution, which is the one proposed in this thesis represents the author's first innovative contribution to the research line of reactive trading and consists of introducing time-varying controllers based on the rationale of an adaptive control technique called Extremum Seeking (ES) [4]. Such an approach appears to be very suitable for the problem at hand, for the following reasons: (i) its aim is to maximize the output of a system whose dynamics is unknown, like the excess return; (ii) it is intrinsically model-free; (iii) it provides a time-varying feedback gain, so it may adapt to market time-varying conditions; (iv) unlike many other adaptive methods, it is theoretically guaranteed to converge to a local optimum. This approach will be described in detail in Chapter 5 of this thesis.

The second problem is the lack of robustness guarantees with real market data and is explained as follows. The RPE Theorem is based on strong assumptions about the processes that govern stock prices. In its original version, the RPE Theorem guarantees positive expected value of the gain-loss function in the case of prices governed by GBM dynamics. In later works, this theorem is extended to the case of prices governed by GBM dynamics with tyme varying mean and volatility [94] and to the case of prices governed by diffusion processes with stochastic jumps [11]. However, real market prices at the same time present complex dynamics such as non-stationarity, non-gaussianity, stochastic jumps and therefore the hypotheses of the RPE theorem do not hold.

As a solution to the second problem and second innovative contribution of this thesis to the reactive trading research line, the author propose a reformulation of the reactive trading scheme, in which the return (i.e., the normalized price trend) is not treated as an unknown disturbance, but rather as an uncertain parameter within a limited range of values. The key observation motivating this analysis is that, in high frequency trading, the return can be well approximated as an Uncertain But Bounded (UBB) parameter, as we will also show via an extensive empirical study. It follows that, by assuming a very mild knowledge of the process dynamics (i.e., the return bounds), a robust controller can be designed, which not only provides good average performance, but also more robust behavior to sudden changes in price dynamics. Specifically, to formulate the design problem as a H_∞ problem, a trend following scheme will be employed, where the desired gain is reasonably selected by considering the current situation. This approach will be described in detail in Chapter 5.

1.2 Multi-period portfolio optimization

Optimizing the composition of a portfolio is one of the key problems of modern quantitative finance. The main objective of portfolio optimization is to guide the investor in the optimal allocation based on some utility function of his monetary resources among different financial assets. The first to deal with this problem with a quantitative approach was the Nobel prize Harry Markowitz who in 1952 ([73]) proposed a single-period allocation strategy based on a trade off between return (measured by the expected return of the portfolio) and risk (measured by portfolio's variance) that could be efficiently resolved through quadratic programming.

The single-period formulation of the portfolio optimization proposed by Markowitz was also soon extended to a dynamic formulation as a stochastic control problem in which the fluctuations in share prices are modeled through stochastic processes and the investor tries to optimally vary over time the number of shares of each asset so as to maximize some utility function of his future wealth.

In the seminal works of Samuelson [96] and Merton [77] the problem of portfolio optimization is formulated as a stochastic optimal control problem and solved through dynamic programming. However multi-period formulations like the one of Merton, in particular, considered only the maximization of the final wealth without considering the risk to obtain it and, moreover, they were based on the assumption of trading in an idealized frictionless market. Therefore they did not take into account the costs of transactions or any restrictions on the composition of the portfolio. For this reason, over the years, much of the work in the multi-period portfolio optimization setting has been aimed at relaxing many of the hypotheses of the idealized market [43, 30]. For this reason, recently, control enginnering approaches such as the predictive control frameworks have received growing interest from the financial community. [93]. Predictive control easily allows to incorporate realistic features of the financial market by modeling them as constraints imposed on the model of the optimization problem.

Of particular interest for this thesis are the approaches described in [29, 30] which propose closed-loop control actions which are parametric functions of past realizations of market returns. This makes it possible to exploit the sequential nature of the multi-stage problem. To maintain reasonable calculation times for practical applications, the author opts for affine control actions.

However, the financial series are known for their complex dynamics and for having a very low signal to noise ratio. It follows that linear control actions may not be optimal for capturing complex price dynamics.

As a solution to this problem and third innovative contribution of this thesis, the author propose to extend the linear parameterization of control actions to nonlinear functions. This is done using a statistical learning non-parametric technique based on kernel methods ([97]) for finding the best nonlinear control policy. The advantage of this approach is that the investment strategy is chosen within a much wider class of functions than the class of linear functions, moreover the use of kernel methods allows to preserve a convex formulation of the problem and therefore guarantee the solution efficiently. This approach will be described in detail in Chapter 6.

1.3 Outline

Chapter 2 provides an overview of reactive trading. The market assumptions are first introduced and then the SLS trading scheme is presented together with the main theoretical result, the RPE Theorem. Finally, recent extensions to the theory are described and the open problems are reported.

Chapter 3 provides an overview of the theory of portfolio optimization. Begin by describing the classic single-period mean-variance approach and its drawbacks. Recent developments are then described, such as portfolio approaches with regularization of parameters and portfolio allocation based on robust optimization. In the second part the multi-period approach is described with particular emphasis on approaches based on predictive control.

Chapter 4 describes the first innovation of the thesis, the integration, in the SLS trading scheme, of time-variant controllers based on the ES rationale.

Chapter 5 describes the second innovative contribution of this thesis, the reformulation of the stock trading problem as a robust control problem of a system characterized by parametric uncertainties represented by price returns.

Chapter 6 describes the last innovative contribution of the thesis, the use of kernel methods to select non-linear investment policies for multi-period portfolio optimization.

Chapter 2

Trading single assets

2.1 Introduction

The objective of this chapter is to describe in greater detail the approach to stock trading via feedback control, called reactive trading and introduced in Chapter 1.

In Section 2.2 the basic scheme of reactive trading is introduced, furthermore the motivation behind the "reactive" nature of the approach is given. Section 2.3 describes the market assumptions on which the theory is based, namely that of *idealized frictionless market.* Then, in Section 2.4, the main reactive trading control scheme is introduced, the SLS, and the fundamental theoretical result, the RPE theorem, is described. Furthermore, in the same section, a practical example of SLS functioning is presented. Section 2.5 introduces the typical evaluation metrics of a stock trading strategy. Finally, in sections 2.6 and 2.7 respectively, recent developments and open problems of the reactive trading approach are discussed.

2.2 Stock trading via feedback control

The objective of this section is to describe Barmish reactive approach to stock trading, and the motivation behind the term reactive will be clear soon.

As stated in Chapter 1, the stock trading problem can be formulated as a model-free feedback control problem. The main innovation of this approach lies in the fact that no parametric price model of the traded stock is assumed to be used to regulate the investment level. Instead, the stock price $p(t)$ is treated as an external uncontrolled disturbance affecting the control system.

As the main example of this model-free strategy we consider trading strategies obtained through a linear feedback rule, which scheme is depicted in Figure 2.1.

The investment level $I(t)$, i.e. the amount of money invested in the stock, represents the adjustable control variable. The cumultive gain-loss function $g(t)$ over $[0, t]$

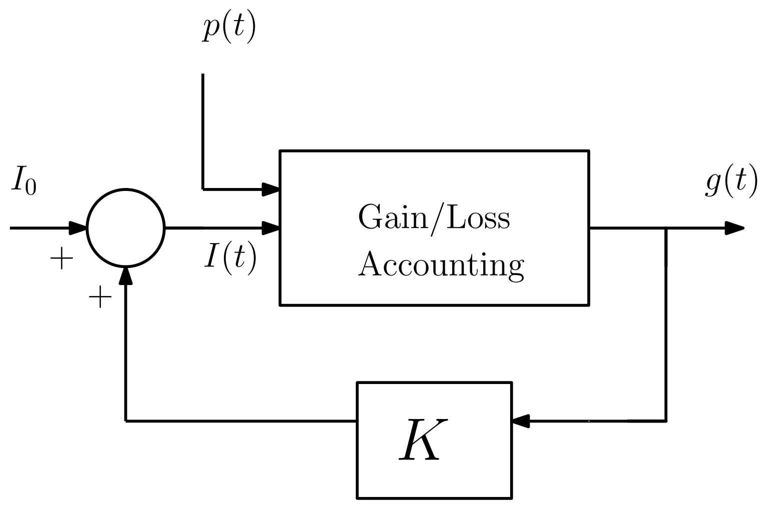

FIGURE 2.1: Basic reactive trading scheme.

represents the output to be controlled. K is the static controller gain and I_0 is a constant value denoting the initial investment. The linear control trading rule is thus given by:

$$I(t) = I_0 + Kg(t) \tag{2.1}$$

The functioning of the control scheme is simply to adjust the investment level to variations in the gain-loss function. Since the system reacts to changes in the gain-loss function, and therefore to price changes, rather than making predictions about future price's trend, this strategy is called reactive trading.

With the trading scheme of Figure 2.1, it is possible to implement two different investment positions called long-selling and short-selling. A long investment is made by the investor whenever he/she trades shares he/she owns, on the contrary a short investment involves the buying and selling of shares that he/she does not own directly but must be borrowed from a bank or a financial intermediary (broker).

A long type position is modeled by choosing a positive initial investment $I_0 > 0$ and a positive proportional controller value $K > 0$. With this choice of parameters, the trading process begins with a positive investment level of $I(t) > 0$. An increase in the price incurred in the trading period considered (which can range from the fraction of a second in the case of high frequency trading to 1 hour, a day, a week, ...) generates an increase in the gain function $g(t)$ that forces the linear control feedback to an increase in the level of investment $I(t)$ in the stock. On the contrary, a decrease in price would generate a loss that will force the scheme to weaken the investment. If the losses are such to bring the level of investment to zero $I(t) = 0$ the long position is automatically closed as the trader no longer owns shares.

A short position works in the opposite way, the model parameters are chosen negative, $I_0 < 0$, $K < 0$ and a negative investment level $I(t) < 0$ stands for short sales. In this case, if the price increases, the trader experiments a loss that drives the scheme to decrease the absolute value of the investment position. On the contrary, if the price decreases the trader benefits from a gain that drives the scheme to an increasing in the investment level.

A strategy of this type where the level of investment adapts to the price trend is well known in the financial literature and takes the name of *trend-following strategy* [35].

2.3 Idealized frictionless markets

Another innovation introduced by Barmish with respect to the classic technical analysis is in the evaluation of the reliability of a trading strategy. In fact, in the financial literature about technical analysis an investment strategy is evaluated only on the basis of statistical simulations such as backtesting with historical price [91]. It is however known in finance that those who rely on backtesting only as a criterion for the validation of a strategy have a high probability of incurring in the "data snooping" phenomenon, known also as backtesting overfitting in the finance context [5]. Data snooping occurs when a dataset is used multiple times for the purpose of inference or model selection. In these cases, any successful claimed results are due to the chances rather than to the goodness of the strategy. For these reasons Barmish aims to develop a theoretical framework based on formal theorems of performance certifications accompanied by proper backtesting with past market data to provide reliable indications on the future performance of a strategy.

Stock trading via the feedback control theoretical framework is based on the assumptions of *idealized frictionless markets*. The assumptions of this type of market are common to relevant results of classical finance including the Black-Scholes model [80] and are described below:

- *continuous trading assumption*: the trader can continuously adjust his current investment level $I(t)$. This assumption is actually consistent with high frequency trading in which many transactions can take place every second.

- *Perfect liquidity and trader as price-taker assumption*: the trader can buy and sell all the shares he/she wants at the current market price, which means that the volumes of the traded shares are not large enough to cause large changes in the price of the stock itself (difference between bid and ask prices). This assumption is consistent for small and medium investors who typically do not move large volumes of shares, while this is not applicable for larger traders such as hedge or mutual funds.

- *Cost-less trading assumption*: the trader does not incur transaction costs. This assumption typically does not discourage large investors who, although they move large volumes of orders, favor customized commission structures that make transaction costs often negligible. On the contrary, small and medium-sized traders must carefully evaluate the volume and frequency of transactions in order to avoid prohibitive transaction costs, although they have experienced a decrease in these costs over the years.

- *Adequate resources*: It is assumed that the trading strategy meets the collateral requirements of the broker so that all trades are admissible and no transactions are stopped by the broker. This requirements can be satisfied assuming that the account of the trader has a suitable large cash balance to provide adequate collateral. It should be pointed out that in real settings, trades will be permitted if $|I(t)| < \gamma w(t)$ with typical value of such that $2 \leq \gamma \leq 4$, for the stock market so that the investor must take account of these restrictions and possibly use techniques to limit the maximum amount of investment, such as for example the one introduced in chapter (2.4.2).

- *Interest and Margin Assumptions*: In practice, denoting with $w(t)$ the wealth function of the trader, when $w(t) > I(t)$ the investor has free cash in his account which he could invest at the risk-free rate of return $r_f \geq 0$. Conversely, when $w(t) < I(t)$, margin interest is owed to the broker with an interest rate m. Assuming for simplicity that $m = r$ (typically $m > r$), previous considerations can be captured by the equation

$$w(t) = w_0 + g(t) + r \int_0^t (w(\tau) - I(\tau)) d\tau \tag{2.2}$$

  where w_0 is the initial wealth of the trader. For simplicity of presentation, in the sequel, an interest rate of $r = 0$ is assumed, so that in this idealized setting the equation updating the trader account simplifies to

$$w(t) = w_0 + g(t) \tag{2.3}$$

2.4 Simultaneous Long Short (SLS) controller

As mentioned in Section 2.2, the scheme of Figure 2.1 allows one to implement either a long position or a short position but not both simultaneously. A long position makes profits if, on average, prices tend to have a positive trend (*bull markets*), on the contrary a short position leads to gains if, on average, the share price has a negative trend (*bear markets*). This implies that the trader must make a choice, a priori, on the investment strategy to follow and consequently must make a priori bets regarding the future price trend. Although this seems an obvious and obligatory choice, it is precisely what one want to avoid.

Ideally a trader would like to hedge against the direction of the market. Motivated by this desire, in [10], a new feedback control scheme called *Simultaneous Long Short (SLS) controller* has been developed that can perform well in both bull and bear markets. This is accomplished by implementing two parallel feedback loops running in parallel, one that executes a long strategy and another that executes a short strategy. The SLS block diagram is depicted in Figure 2.2.

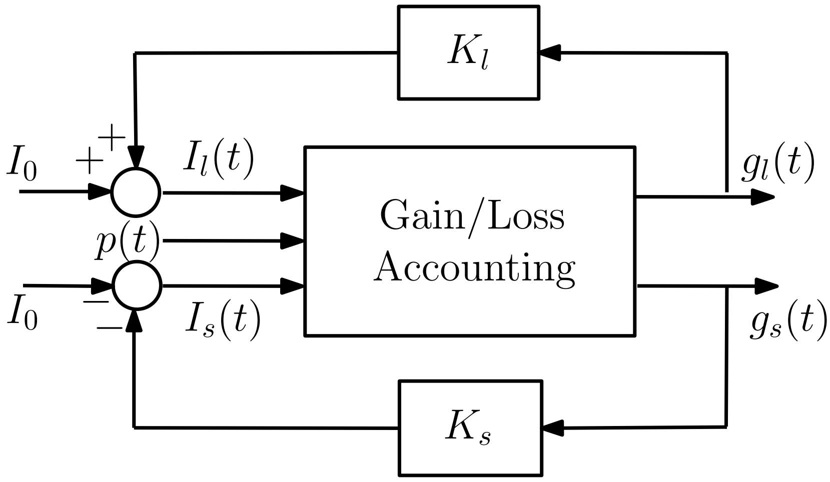

FIGURE 2.2: Simultaneous Long Short (SLS) control scheme.

Supposing $I_0 > 0$ and denoting with K_l the proportional controller of the feedback loops that implements the long position and with $g_l(t)$ the gain-loss function of this position, the amount to invest in the long trade is given by:

$$I_l(t) \doteq I_0 + K_l g_l(t) \tag{2.4}$$

Denoting instead with $K_s > 0$ and with $g_s(t)$ respectively the gain controller of the short feedback loop and the gain loss function of the short position, the short trade is given by:

$$I_s(t) \doteq -I_0 - K_s g_s(t) \tag{2.5}$$

Assuming for simplicity $K_l = K_s = K$, the total investment is given by the sum of the long investment level $I_l(t)$ and the short investment level $I_s(t)$, i.e.:

$$I(t) = I_l(t) + I_s(t) = K(g_l(t) - g_s(t)) \tag{2.6}$$

and the total gain-loss function is given by:

$$g(t) = g_l(t) + g_s(t) \tag{2.7}$$

Finally, it is useful to define the wealth function of the trader. Assuming that the trader has an initial wealth of w_0, the function that keeps track of the investor's wealth is given by:

$$w(t) = w_0 + g(t) \tag{2.8}$$

The operating logic of the SLS scheme is as follows: if the price of the stock is maintaining a positive trend, the gain of the long feedback loop $g_l(t)$ increases accordingly, and the control law progressively forces an increase in the long investment level $I_l(t)$ and a progressive decrease, in module, of the short investment level $I_s(t)$. In the case, instead, of a progressive decrease in the price, the control scheme would strengthen the short position by increasing, in module, the short investment level and, at the same time, weakening the long position by reducing the long investment level. If the price trend remains constant in one direction for enough time the scheme will automatically turn off the losing position.

The fundamental theoretical result of this line of research, the RPE theorem, is described below.

2.4.1 Robust Positive Expectation (RPE) theorem

Theorem's fundamental assumption is that prices are driven by Geometric Brownian Motion (GBM), one of the most well-known price process in the financial literature [88]. This assumption does not reflect reality: price dynamics are more complex, are not stationary and are subject to stochastic jumps. However, each new theory starts with assumptions that are gradually relaxed. The GBM process is described by the stochastic equation:

$$\frac{dp}{p} = \mu dt + \sigma dZ_t \tag{2.9}$$

where $\frac{dp}{p}$ is the percentage change in price value over the time interval $[t, t + dt]$, μ is called the *drift* of the process and $\sigma > 0$ is called the *volatility* of the process and Z_t is a standard Wiener process such that $dZ_t \sim \mathcal{N}(0,\ dt)$.

Theorem 2.1 (Robust Positive Expectation (RPE) theorem [10]). *In an idealized frictionless market with GBM prices with drift μ and volatility σ, for $t \geq 0$, the expectation, variance and worst-case loss resulting from the SLS feedback control are given, respectively, by:*

$$\mathbb{E}[g(t)] = \frac{I_0}{K}[e^{\mu K t} + e^{-\mu K t} - 2] \tag{2.10}$$

$$Var[g(t)] = \frac{I_0^2}{K^2}(e^{\sigma^2 K^2 t} - 1)(e^{2\mu K t} + e^{-2\mu K t} + e^{\sigma^2 K^2 t}) \tag{2.11}$$

$$g^*(t) = \frac{2I_0}{K}[e^{-\frac{1}{2}\sigma^2 K^2 t} - 1] \tag{2.12}$$

Moreover, except for the trivial break-even case when $\mu = 0$:

$$\mathbb{E}[g(t)] > 0 \tag{2.13}$$

The theorem states that, regardless the sign of the drift μ of the price, the control strategy guarantess that the gain loss function will be strictly positive in expected value.

2.4.2 Practical implementation of the SLS controller

The original scheme is derived in continuous time, but in practical implementations the scheme is reformulated in discrete time. Hence, we now assume trading occurs at discrete time, the inter-sample time can be either small such as fraction of seconds for the high-frequency trader or large such days for the mutual fund. Indeed, we let $p(k)$, $I(k)$, $w(k)$ and $g(k)$ denote the discrete-time version of $p(t)$, $I(t)$, $w(t)$ and $g(t)$ respectively.

Now we introduce the one-period percentage change in stock price, called return of the price, as:

$$\rho(k+1) \doteq \frac{p(k+1) - p(k)}{p(k)} \tag{2.14}$$

The dynamic update equations for the SLS are thus given by

$$g_l(k+1) = g_l(k) + \rho(k+1)I_l(k) \tag{2.15}$$

$$g_s(k+1) = g_s(k) + \rho(k+1)I_s(k) \tag{2.16}$$

$$g(k+1) = g_l(k+1) + g_s(k+1) \tag{2.17}$$

$$w(k+1) = w(0) + g(k+1) \tag{2.18}$$

$$I_l(k+1) = I_0 + K_l g_l(k+1) \tag{2.19}$$

$$I_s(k+1) = -I_0 - K_s g_s(k+1) \tag{2.20}$$

$$I(k+1) = I_l(k+1) + I_s(k+1) \tag{2.21}$$

To handle the sign restriction condition on $I_l(k)$ and $I_s(k)$ we modify their update equation:

$$I_l(k+1) = \max\{(I_0 + K_l g_l(k+1)), 0\} \tag{2.22}$$

$$I_s(k+1) = \min\{(-I_0 - K_s g_s(k+1)), 0\} \tag{2.23}$$

Furthermore, in practice, it is reasonable to suppose that a trader has constraints on the maximum amount of money I_{max} he/she can invest. To model this fact it is possible to add two further saturation constraints, one for the long side and the other for the short side of the trade

$$I_l(k+1) = \min\{(I_0 + K_l g_l(k+1)), I_{max}\} \tag{2.24}$$

$$I_s(k+1) = \max\{(-I_0 - K_s g_s(k+1)), I_{max}\} \tag{2.25}$$

To demonstrate how the discrete-time implementation of the SLS works, we perform two numerical simulation experiments. In particular we simulate price realizations assuming the Black-Scholes-Merton setup for option pricing [15]. In this setup, the price of a stock index $p(T)$ at a future date T given the price $p(0)$ as of today is given according to

$$p(T) = p(0)\exp\left((r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{(T)}z\right) \tag{2.26}$$

where r is the constant risk-free annual rate of return, σ is the volatility of returns of the underlying asset assumed constant and z is a standard normally distributed random variable. We consider now the Black-Scholes-Merton model in its dynamic form, as described by the stochastic differential equation 2.9 that represents a GBM process. Equation 2.9 can be discretized exactly by an Euler scheme, resulting in

$$p(k+1) = p(k)\exp\left((r - \frac{1}{2}\sigma^2)\Delta_t + \sigma\sqrt{\Delta_t}z\right) \tag{2.27}$$

Setting $\Delta_t = \frac{1}{252}$ we are able to simulate price realizations made of 252 trading day samples and governed by GBM dynamics so that the assumptions of Theorem 2.1 are satisfied. The first price realization was generated with $r = 0.20$ and $\sigma = 0.25$ to simulate a bull market (Figure 2.3), the second price realization was generated with with $r = -0.20$ and $\sigma = 0.25$ to simulate a bear market (Figure 2.5). In both cases the SLS was initialized with $I(0) = I_0 = 100$, $w(0) = 50$ and the feedback gain was set to $K_l = K_s = 2$.We expect the investment strategy to be profitable in both scenarios, as stated by Theorem 2.1.

It is necessary to point out that the simulations reported in this section serve only for illustrative purposes to visually describe the behavior of the SLS scheme in two ideal situations in which the bull and bear behaviors are well marked.

Example 1 - Bull market

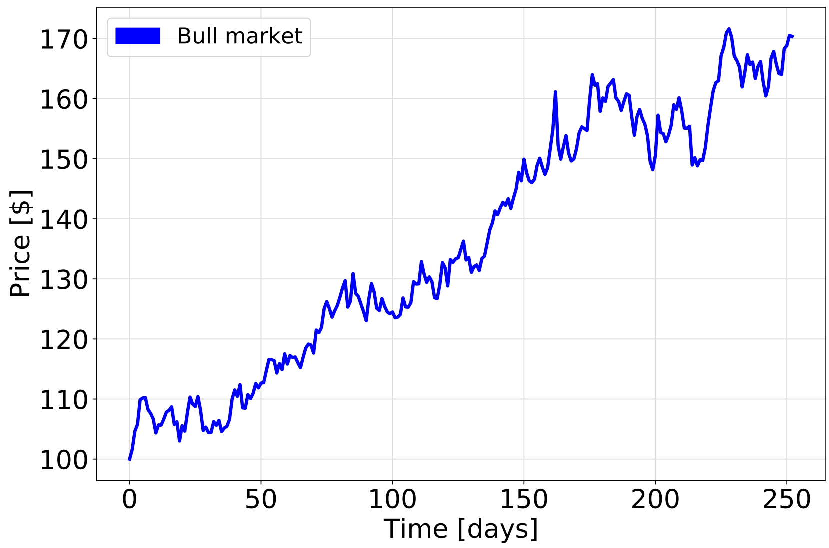

FIGURE 2.3: Bull market.

Figure 2.3 shows an increasing simulated price trend, while Figure 2.4 depicts the behavior of the SLS investment strategy during the investment horizon, and in particular shows the trends of long and short investment levels (top picture) and the performances of the gain-loss function of the long position and the short position (bottom picture). Since the price increases, the control scheme make sure that the short position tends to close up so that the long investment remains the only active position. The short position suffers a loss of money as evidenced by the sign of its own gain-loss function, while the long position incurs a high profit. The total return of the strategy at the end of the investment horizon is positive.

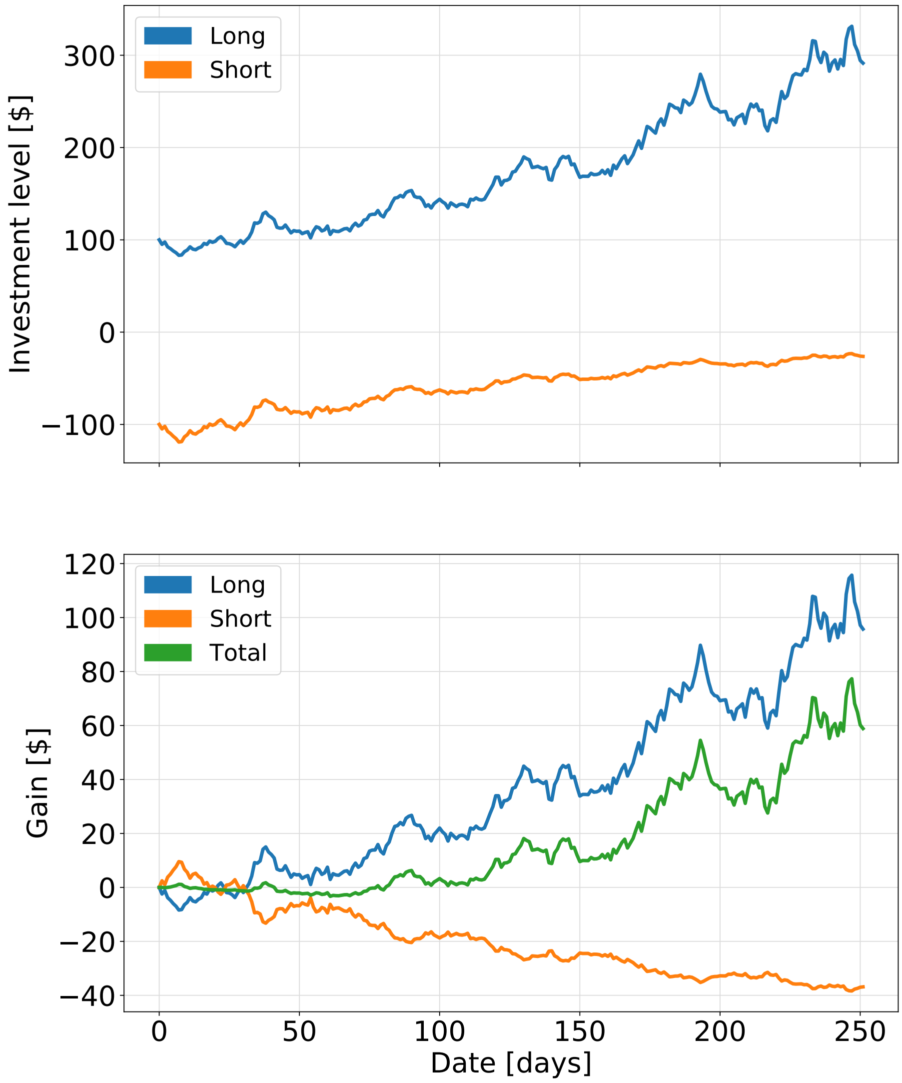

FIGURE 2.4: (Top) Investment levels of Long and short positions. (Bottom) Long, short and total gain-loss functions.

Example 2 - Bear market

Unlike the previous case, figure 2.5 shows a decreasing price trend and figure 2.6 depicts the behaviour of the SLS in this bear market. In this case, since the price decreases, the control scheme makes sure that the long position tends to close up so that the short investment remains the only active position. In this case, the gain loss function of the short position is the profitable one, while the gain loss function of the long position is negative. However, as in the previous case, the overall gain of the strategy at the end of the investment horizon is positive.

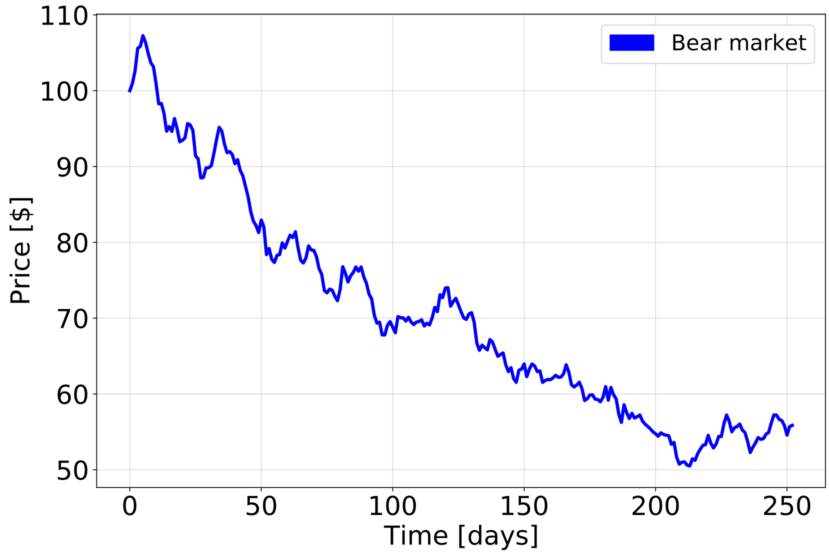

FIGURE 2.5: Bear market.

2.5 Performance evaluation

In this section we describe additional performance measures, other than profit, to asses the performance of a trading strategy given the outputs of simulations and backtests, where backtesting is the process of applying a trading strategy to historical data to see how the strategy would have performed. This is necessary as a trader is not only interested in the final profit of a strategy but also in the risk it takes to obtain it. We assume an investment horizon of T time periods with $k = 0, 1, \dots T$.

Final return

Final return measures the return of the wealth of the trader at the end of the investment horizon, and is given by:

$$r(T) \doteq \frac{w(T) - w(0)}{w(0)} \tag{2.28}$$

Average return and variance

Defining the returns as the natural log difference:

$$r(k) = \ln w(k) - \ln w(k-1) \tag{2.29}$$

the average return of a strategy over the period of time $k = 0, 1, \dots T$ is computed as

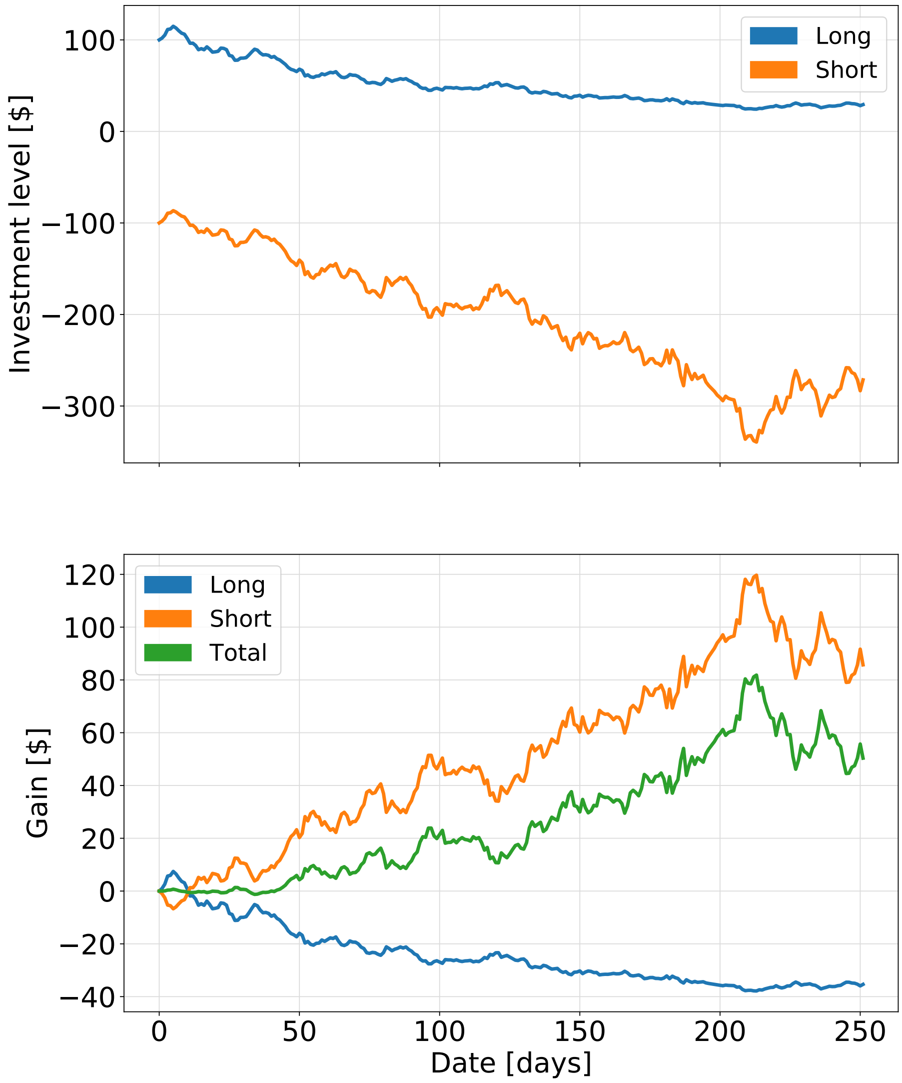

FIGURE 2.6: (Top) Investment levels of Long and short positions. (Bottom) Long, short and total gain-loss functions.

$$\bar{r} \doteq \frac{1}{T}\sum_{k=1}^{T} r(k) \tag{2.30}$$

The variance of a trading strategy measures the degree of variation of the series of returns and is given by:

$$\mathrm{Var}[r(k)] \doteq \frac{1}{T-1}\sum_{k=1}^{T} (\bar{r} - r(k))^2 \tag{2.31}$$

Sharpe ratio (SR)

To define the Sharpe Ratio (SR) we have to first define the return r_f of a risk-free asset. Risk-free assets has certain future return, for example Treasury bills are considered risk-free because they are backed by the U.S. government and are so sure that their returns are very close to the current interest rates.

SR quantifies the excess return (the average return of a strategy in excess of the risk-free rate, i.e., $\bar{r} - r_f$) normalized by the risk (the original definition consider the standard deviation of the returns $\sqrt{\mathrm{Var}[r(k)]}$ as the measure of risk), the higher the Sharpe Ratio of a strategy, the greater the reward of the investment for the risk taken. SR is then given by:

$$SR \doteq \frac{\bar{r} - r_f}{\sqrt{\mathrm{Var}[r(k)]}} \tag{2.32}$$

In the particular case in which we consider $r_f = 0$ we speak of Information Ratio (IR):

$$IR \doteq \frac{\bar{r}}{\sqrt{\mathrm{Var}[r(k)]}} \tag{2.33}$$

Percentage drawdown

Percentage drawdown at time k is defined as the maximum percentual decline from a historical peak of the wealth $w(j)$, such that:

$$d(k) \doteq \max_{j=0,1,\ldots,k} \frac{w(j) - w(k)}{w(j)} \tag{2.34}$$

Since Percentage drawdown is a function of time it is usual to report its maximum value, called maximum percentage drawdown:

$$d_{max} \doteq \max_{k} d(k) \tag{2.35}$$

2.6 Numerical experiment

In this section we describe a numerical experiment on real financial data. The objective of the experiment is to show the functioning of the SLS in the attempt of "beating the market", that is, earning an investment return that exceeds the performance of the Standard & Poor’s 500 index. Commonly called the S&P 500, it’s one of the most popular benchmarks of the overall U.S. stock market performance. The performance of the SLS algorithm will be compared with the naive strategy called Buy and Hold (B& H). B& H is a passive investment strategy in which an investor buys stocks (or other types of securities such as ETFs) and holds them for a long period regardless of fluctuations in the market. In particular, in this experiment we consider the trading of SPDR S&P 500 trusts, that is an exchange-traded fund which trades on the NYSE under the symbol SPY and is designed to track the S&P 500 stock market index.

In Figure 2.7 it is possible to observe the historical data of the SPY price in the period 2011-2014. These data have been divided into two parts, a first part of in-sample data ranging from 2011-2012 and used to tune the value of the static gain $K = K_l = K_s$. The second part of data, which covers the period 2013-2014 it is used to evaluate the out-of-sample performance of the investment strategies.

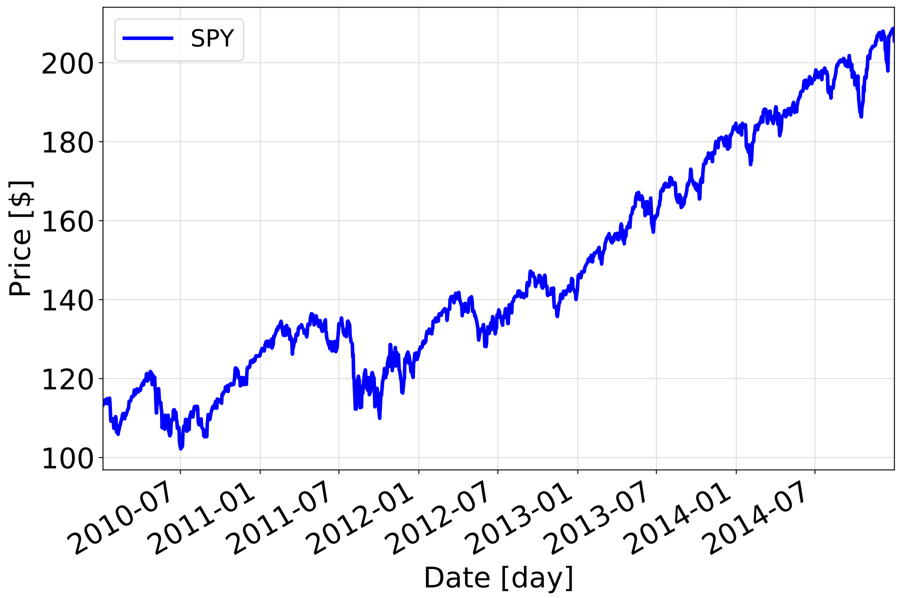

FIGURE 2.7: Time evolution of SPY ETF.

2.6.1 Out-of-sample results

To make the evaluation of the performances more reliable, in addition to the actual in-sample and out-of-sample realizations, we generate a number $N_{in} = 1000$ of scenarios from the in-sample data and a number $N_{out} = 1000$ from the out-of-sample data with same length as the original realizations. The scenarios are obtained by sampling with replacement the original series, technique called bootstrap sampling [44]. That is, in order not to rely on a single realization of the data (the real one) to evaluate the performance of the algorithm, starting from the vector of the in-sample data we generate other N_{in} vectors of the same length by sampling with replacement, creating N_{in} possible alternative realizations / scenarios with the same probability distribution as the original data. These scenarios generated from the in-sample data are used to determine the optimal value of the parameters of the algorithm, the gain K. This procedure is repeated to generate N_{out} scenarios from the out-of-sample data, which will instead be used to evaluate the performance of the algorithm. The values used for the experiment are an initial investment $I_0 = 5000$ (in the case of SLS this value changes during the time horizon while in the case of B&H it remains unchanged) and an initial wealth of $w(0) = 2500$. Moreover, an optimal value of $K = 6$ was selected from a grid of 10 values separated linearly in the interval $[1, 10]$ as the value that maximizes the average in-sample IR of the SLS strategy, see Figure 2.8. The Figure 2.9 shows the empirical distribution of the IRs obtained by testing the SLS strategy and the B&H strategy on the 1000 out-of-sample scenarios. The two distributions have, respectively, mean 0.078 and standard deviation 0.058 for the SLS strategy and mean 0.10 and standard deviation 0.046 for the B& H strategy. This results tells us that, on average, for this particular experiment setting, B& H allows to obtain higher IR values with a lower variance in the results than those obtainable with SLS.

In Figure 2.10, it is possible to observe the empirical distribution of the values of maximum percentage drawdowns obtained by testing the two strategies on the out-of-sample realizations. The two distributions have mean 0.26, standard deviation 0.08 and a positive skew of 0.76 in the case of the SLS and mean 0.1, standard deviation 0.03 and positive skew of 1.15 in the case of B&H. From the point of view of this second evaluation metric, B&H allows to obtain on average better performances than the SLS.

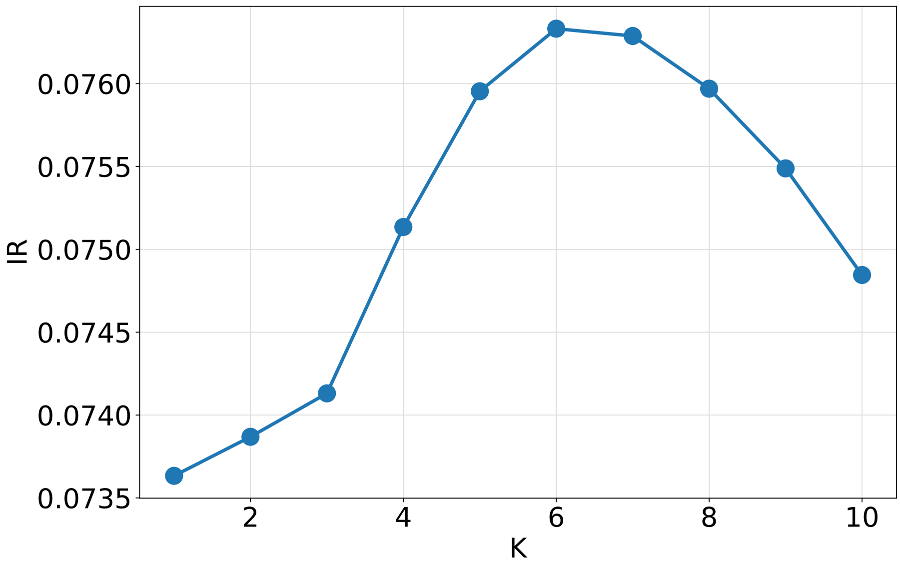

FIGURE 2.8: Average IR for different values of K.

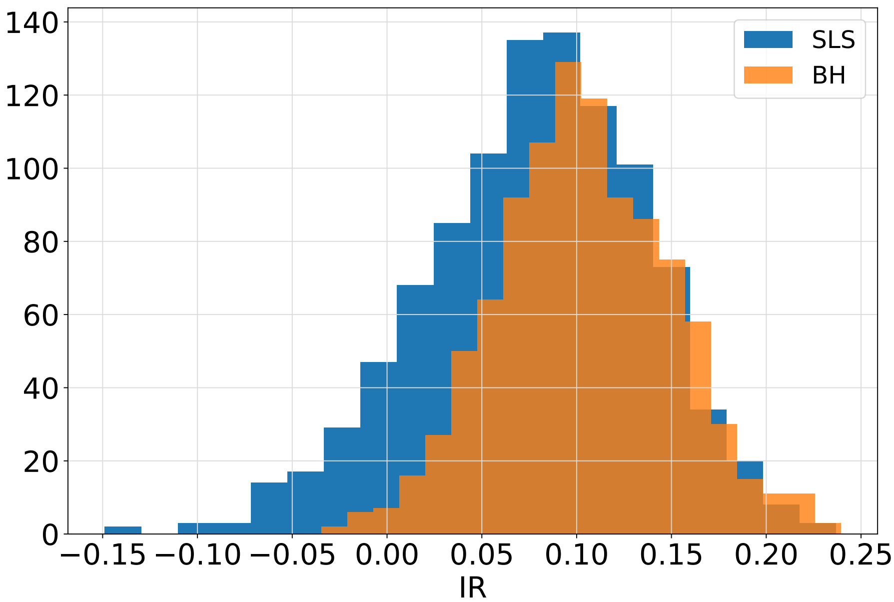

FIGURE 2.9: Emprical distribution of the Information Ratios of the SLS and B& H strategy.

2.7 Extensions and recent developments

This section summarizes briefly some recent interenting developments in this research area.

Generalization to dynamic controller

In [72] the Robust positive expectation theorem is extended to dynamic controllers, and in particular a Proportional-Integrative (PI) controller, of the form

$$I(t) = I_0 + K_P g(t) + K_I \int_0^t g(\tau) d\tau \tag{2.36}$$

is considered. Assuming that the prices of the traded stocks are Geometric Brownian Motion and still adopting a Simultaneous Long Short control scheme with investment laws given by

$$I_l(t) \doteq I_0 + K_P g_l(t) + K_I \int_0^t g_l(\tau) d\tau \tag{2.37}$$

$$I_s(t) \doteq I_0 + K_P g_s(t) + K_I \int_0^t g_s(\tau) d\tau \tag{2.38}$$

it is shown once again that, regardless of the sign, the expected value of the gain function is strictly positive

$$\mathbb{E}[g(t)] > 0 \tag{2.39}$$

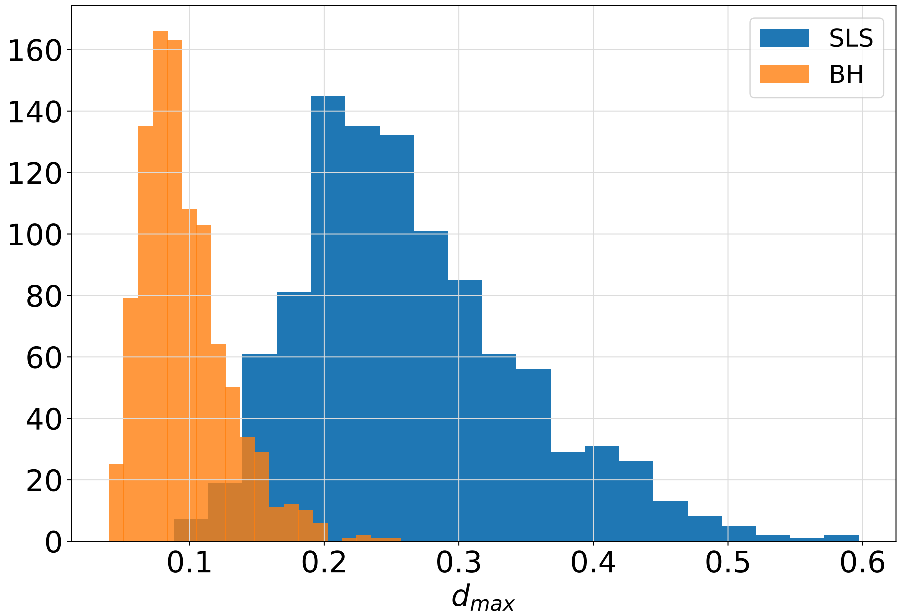

FIGURE 2.10: Emprical distribution of the maximum percentage draw-downs of the SLS and B& H strategy.

Generalization to time-varying price dynamics

In [94] is provided a generalization of the Robust positive expectation theorem to Geometric Brownian Motion price dynamics with time-varying parameters

$$\frac{dp}{p} = \mu(t)dt + \sigma(t)dZ_t \tag{2.40}$$

where the drift $\mu(t)$ and the volatility $\sigma(t)$ are assumed to be continuos functions of time. To the authors' surprise, the positivity of the expected gain-loss function is demonstrated.

Generalization to discontinuos stock returns

In [11] the results of the RPE Theorem are extended to price processes governed by Merton's jump diffusion model [78]. The relative price change in Merton's model is given by:

$$\frac{dp}{p} = (\alpha - \lambda\kappa)dt + \sigma dZ_t + dN_t \tag{2.41}$$

where N_t is a Poisson-driven process with jump intensity $\lambda > 0$, α denotes the jumpless trend and σ the volatility. Suppose in the small time interval dt the asset price jumps from p to yp , where y is the absolute price jump size. So the percentage change in the asset price caused by the jump is:

$$\frac{dp}{p} = \frac{yp - p}{p} = y - 1 \tag{2.42}$$

where Merton assume y to be a nonnegative random variable drawn from lognormal distribution, $\ln(y) \sim \mathcal{N}(\mu, \delta^2)$. Finally we define $\kappa \doteq \mathbb{E}[y - 1]$.

It is proved that the expected gain of the SLS trading strategy with a stock price following 2.41 is:

$$\mathbb{E}[g(t)] = \frac{I_0}{K}[e^{\alpha K t} + e^{-\alpha K t} - 2] \tag{2.43}$$

and therefore is in general positive and indipendent of intensity, kind (that means, the distribution of y), and size of the jumps.

Generalization to a pair of stocks

In [39] is developed a new version of the Robust positive expectation theorem for the case of trading two directionally correlated stocks with bounded non-zero momenta. They consider two stocks with stochastically varying prices in discrete time $p_1(k)$ and $p_2(k)$ with returns given by:

$$\rho_i(k) \doteq \frac{p_i(k+1) - p_i(k)}{p_i(k)} \tag{2.44}$$

with $i = 1, 2$ and expected value $\mu_1 \doteq \mathbb{E}[\rho_1(k)]$, $\mu_2 \doteq \mathbb{E}[\rho_2(k)]$. These returns are assumed to be *directionally correlated* that is, there exists a constant $\beta \neq 0$ such that $\mu_2 = \beta\mu_1$, moreover *bounded non-zero momentum* is assumed, that is there exist two positive constant μ_{min} and μ_{max} known to the trader such that $\mu_{min} \leq |\mu| \leq \mu_{max}$.

Given β, μ_{min} and μ_{max} the new version of the theorem proposed in this work provide necessary and sufficient conditions on the controller gain K under which robustly positive expected trading gain $\mathbb{E}[g(t)] > 0$ is guaranteed. If no K satisfies the conditions the stock pairs are considered not tradable. It is noted that this result fits well in the context of pairs trading.

2.8 Open problems

Of the theory of reactive trading described so far we report the two main open problems and the relative solutions proposed as innovative contributions of this thesis.

Problem 1 - tuning of the controller's parameters

In the original scheme the controller is static, ie it is a pure gain $\mathrm{R}(k) = K$. Motivated by the idea that a dynamic controller could better adapt to the dynamics of prices, in [72], the authors generalize robust positive expectation property of the gain-loss function to the more general case of a Proportional-Integrative (PI) controller, as discussed in the firt part of the previous section.

However, both in the case of a static controller and in the case of a dynamic time-invariant controller, it is necessary to face the problem of tuning the controller's parameters. Typically, the tuning of the parameters of a trading strategy is done through backtesting using in-sample calibration data and out-of-sample data for their validation.

However, it is well documented in literature [64, 98, 84] that the financial series of market prices are strongly non-stationary and subject to sudden changes in their dynamics, and sometimes, extraordinary phenomena such as financial crisis, generate price dynamics never seen in historical data. Therefore a calibration validated over a given out-of-sample period is not necessarily optimal for a different future scenario.

A possible solution to this problem, proposed by the author in this thesis, is to use an adaptive time-variant controller capable of varying its value online and therefore adapting, more readily than a static controller, to changes in the dynamics of prices. A suitable adaptive control law must be defined for this purpose. In Chapter 4 the first innovative contribution of this thesis is described for this purpose, the adaptive control law used is called Extremum Seeking.

Problem 2 - robustness for applications with real data

The original version of the RPE Theorem holds for ideal price dynamics governed by simple geometric Brownian motion. However, the RPE Theorem does not provide guarantees of robustness for real market prices that exhibit complex dynamics such as non-stationarity, non-gaussianity [54], stochastic jumps [6]. As reported in the previous section, there have been attempts to extend RPE Theorem to brownian processes with time-varying parameters and to diffusion processes with stochastic jumps. However, although complex, even the latter are simplified models of reality. Therefore they represent a strong assumption on the nature of stock prices.

A possible solution to this problem, proposed as the second innovative contribution of this thesis and described in Chapter 5, is to make only a very mild assumption on prices, which is that of considering the returns of prices as being lowerly and upperly bounded in a given interval of values. Defining this interval as an uncertainty interval set, it is possible to treat the return of prices as parametric uncertainty of the model and exploits the framework of robust control to synthesize a controller with, for example, the H_{inf} approach, robust to all the possible variations of returns in the considered range of values.

Chapter 3

Trading a portfolio of assets

3.1 Introduction

In Chapter 2 we addressed the problem of trading individual financial assets. However, in the financial literature it is known that diversifying among various financial instruments and industries can reduce the risk associated with the investment since different sectors react differently to the same event [75]. In this chapter we will therefore deal with the trading of multiple financial assets that constitute what in finance is called a portfolio.

In this chapter, we distinguish portfolio optimization approaches in two macro-categories. The first category is characterized by a single-period formulation and will be discussed in detail in Section 3.2. The most famous single-period approach dates back to Markowitz [73] and is based on the tradeoff between the return of an investment and its associated risk. This approach, although its conceptual elegance, suffers from two serious drawbacks that prevented its practical application. First, the Markowitz model is based on the assumption of perfect knowledge of a series of parameters that in reality must be estimated from past data and it has been shown that the mean-variance framework is very sensitive to parameter estimation errors [34]. Second, variance is not an appropriate risk measure because it is a symmetrical risk measure that penalizes equally large losses and large returns, while it is clear that an investor would like to distinguish between upside and downside risk and ideally have positive skewed returns [66]. Over the years, various approaches to solving the aforementioned problems have been proposed in the literature. To mitigate the problems caused by errors in estimating the parameters of the model, in sections 3.2.7 and 3.2.8, we describe the approaches based on parameter regularization techniques and robust optimization. Risk measures of an investiment strategy, alternatives to the variance are instead described in section 3.2.10.

The second category of portfolio optimization approaches is characterized by a multi-period framework and will be described in Section 3.3. The basic idea is to formulate from the beginning the allocation problem on a time horizon composed of multiple decision steps. The objective is to minimize the total risk over the entire investment horizon and at the same time to respect real constraints on the portfolio composition of each stage. In particular, approaches based on predictive control will be described which allow a convex formulation of the optimization problem and therefore can be solved efficiently.

3.2 Single period optimization

The foundation of the Modern Portfolio Theory lies in the model introduced by Harry Markowitz in 1952 [73]; it was in fact the first quantitative approach to building a portfolio of financial assets. He had the intuition to treat the problem of allocating an investor's monetary resources as a problem of mathematical optimization with the aim of finding the best trade-off between expected return and risk associated with the portfolio and measured by its variance. Thanks to this fundamental contribution, Markowitz won the Nobel Prize in Economics in 1990 together with Merton Miller and William Sharpe.

The classic formulation of Markowitz is single-period in the sense that the investor is assumed to make allocation decisions once and for all at the beginning of a given period (e.g. one quarter or one year). Once made, the allocation decisions are not allowed to change until the end of the period and the impact of decisions arising in subsequent periods is not considered [33].

3.2.1 The model

In this section we will introduce the model for single-period portfolio optimization mainly adopting the notation in [27].

Consider a collection of of n risky assets $a_1, \dots, a_n$, and call with $r \in \mathbb{R}^n$ the random vector with elements the holding period returns of the n assets over an arbitrary fixed time interval (for example a day, a week or a month). Let's denote with $x \in \mathbb{R}^n$ the portfolio vector(or vector of positions or weights) which elements x_i representing the dollar (or equally any other monetary currency) value of the portion of investor's wealth allocated in asset i, where $x_i \geq 0$ stands for a long position in asset i and $x_i < 0$ stand for a short position in asset i, with $i = 1, \dots, n$. Assuming that the investor possesses an initial portfolio denoted by x_0 and that he wants to carry out trading operations to vary the composition of the initial portfolio, we define with $u \in \mathbb{R}^n$ the vector of the dollar value of the trades at the current market prices, where $u_i > 0$ means that the investiment in the asset i is increased, $u_i < 0$ means that is decreased and $u_i = 0$ means that the investment in asset i is held unchanged. The post-trade portfolio is denoted as:

$$x = x_0 + u \tag{3.1}$$

Denoting with $w_0 = \mathbf{1}^T x_0$ the initial wealth (or the initial total value of the portfolio in dollars) and with $w = \mathbf{1}^T x$ the post-trade wealth, the change in the total portfolio value after the trade is given by

$$w - w_0 = \mathbf{1}^T x - \mathbf{1}^T x_0 = \mathbf{1}^T u \tag{3.2}$$

Typically one assume that the portfolio is *self-financing* meaning that no external cash is put into or withdrawn from the portfolio so that the wealth is preserved ($\mathbf{1}^T x = \mathbf{1}^T x_0$). The self-financing condition can be ensured (assuming also no holding and transaction costs) imposing:

$$\mathbf{1}^T u = 0 \tag{3.3}$$

Finally we can define the portfolio return as

$$r_p = \sum_{i=1}^{n} x_i r_i = r^T x \tag{3.4}$$

Of course, the larger portfolio return the better, on the investor side.

3.2.2 Mean-variance trade-off optimization

The basic idea behind the Markowitz model is to find, through an optimization procedure, a trade-off between expected return and variance as a measure of performance and risk, respectively, associated with the portfolio. Denoting the expected value of the return vector as $\mu = \mathbb{E}[r]$ and the covariance matrix of the asset returns as $\Sigma = \mathbb{E}[(r-\mu)(r-\mu)^T] = (\sigma_{i,j}) \in \mathbb{R}^{n \times n}$, where $\sigma_{i,j} = \mathrm{Cov}(r_i, r_j)$ is the covariance between returns of assets i and j, the expected return of the portfolio is given by

$$\mu_p = \mathbb{E}[r_p] = \sum_{i=1}^{n} x_i \mu_i = \mu^T x \tag{3.5}$$

and his associate variance is instead given by

$$\sigma_p^2 = \mathrm{Var}[r_p] = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i \sigma_{ij} x_j = x^T \Sigma x \tag{3.6}$$

The problem of selecting a portfolio can be formally stated as a tradeoff between these two quantities. The mean-variance framework is based on the strong assumption of having perfect knowledge of the random quantities involved in the model (μ, Σ and derived quantities). In reality, these parameters are not known and need to be estimated using past realizations of asset prices and therefore subject to estimation errors.

There are three possible different formulations of this problem, the return maximization problem, the risk minimization problem, and the risk-adjusted return maximization problem.

Return maximization problem

The return maximization formulation aims at maximizing portfolio expected return while keeping the risk under control imposing a maximum admissible level of risk $\bar{\sigma}^2$:

$$\max_{u} \quad \mu^T x \tag{3.7a}$$

$$\text{s.t.} \quad x^T \Sigma x \leq \bar{\sigma}^2, \tag{3.7b}$$

$$x \in \mathcal{X} \tag{3.7c}$$

where $\mathcal{X}$ is the set of admissible portfolios and reflects constraint on portfolio composition, for example

$$\mathcal{X} = \{x \in \mathbb{R}^n : x = x_0 + u,\ \mathbf{1}^T x = \mathbf{1}^T x_0,\ x \geq 0\} \tag{3.8}$$

where the non-negativity condition $x \geq 0$ rule out short-selling. Given that the covariance matrix is positive definite ($\Sigma \geq 0$), the above optimization problem has a linear objective function, linear and quadratic constraints and thus is efficiently solvable.

Risk minimization problem

The risk minimization formulation aims at minimizing the portfolio risk imposing the desired expected portfolio return being above a given target γ:

$$\min_{u} \quad x^T \Sigma x \tag{3.9a}$$

$$\text{s.t.} \quad \mu^T x \geq \gamma, \tag{3.9b}$$

$$x \in \mathcal{X} \tag{3.9c}$$

This formulation has quadratic objective function, linear constraints and thus leads to a Quadratic Programming (QP) problem.

Risk-adjusted return maximization problem

This third formulation is obtained building the objective function through linear combination of the two objective functions above and aims to maximize a risk-adjusted return:

$$\max_{u} \quad \mu^T x - \zeta \cdot x^T \Sigma x \tag{3.10a}$$

$$\text{s.t.} \quad x \in \mathcal{X} \tag{3.10b}$$

where ζ is called risk-aversion parameter as the higher the ζ value, the greater the risk investor's aversion. Also this last formulation leads to a QP problem.

3.2.3 Numerical example 1 - Efficient frontier computation

Each of the above three formulations has the difficulty in choosing reasonable value for the involved parameters $\bar{\sigma}^2, \gamma$ or ζ. What is done in practice is to solve the optimization problems mentioned above for increasing values of $\bar{\sigma}^2, \gamma$ or ζ.

It can be shown that the different formulations result in the same trade-off curve (Pareto curve) of optimal values in the risk/return[I] plane called the *efficient frontier* in the finance literature. A portfolio x that lies on this curve is called an *efficient portfolio*. A portfolio is efficient if, among all the possible portfolio with a given level of return is the one with minimun risk, or equivalently if is the one with maximum expected return for a givel level of risk.

As an example we consider the allocation over the following $n = 11$ Exchange-Traded Funds (ETFs):

1. SPDR S&P 500 ETF, ticker SPY;
2. iShares 20+ Year Treasury Bond ETF , ticker TLT;
3. SDPR Select Sector Fund - Financial Sector, ticker XLF;
4. SDPR Select Sector Fund - Energy Sector, ticker XLE;
5. SDPR Select Sector Fund - Health Care Sector, ticker XLV;
6. SDPR Select Sector Fund - Consumer Staples Sector, ticker XLP;

[I]Typically, for illustrative purpose, the standard deviation is used in place of the variance as the risk measure

7. SDPR Select Sector Fund - Consumer Discretionary Sector, ticker XLY;
8. SDPR Select Sector Fund - Industrial Sector, ticker XLI;
9. SDPR Select Sector Fund - Technology Sector, ticker XLK;
10. SDPR Select Sector Fund - Materials Sector, ticker XLB;
11. SDPR Select Sector Fund - Utilities Sector, ticker XLU.

We used 2 years of past monthly data corresponding to the $2016 - 2017$ period to estimate the empirical average of the return vector and the empirical covariance matrix. The estimate of the expected value of the returns has a minimum value of 2.5% corresponding to the asset with ticker TLT, and a maximum value of 1.74% corresponding to the asset with the XLK ticker. We then considered 40 equi spaced values for the target return γ in the interval $[2.5 \times 10^{-3}, 1.74 \times 10^{-2}]$ and for each value of γ we solved problem 3.9 with the no short-selling constraint $x \geq 0$.

In figure 3.1 is shown the discrete point approximation of the efficient frontier.

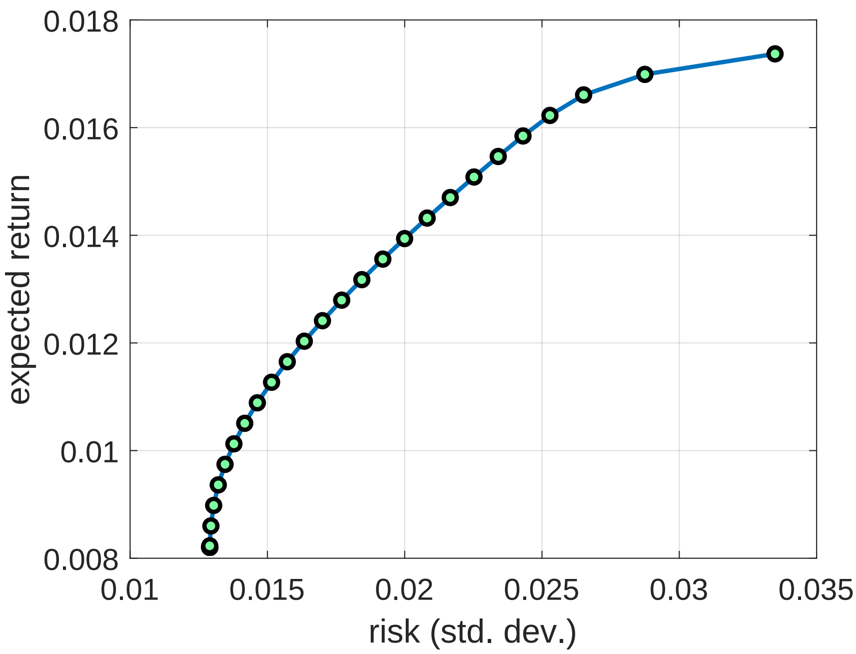

FIGURE 3.1: Efficient Frontier.

The figure 3.2 also shows with orange dots the performances that would be obtained with portfolios in which all the wealth is invested in the single asset highlighted by the corresponding label. What this figure tells us is that portfolios obtained as a solution to the optimization and located on the efficient frontier are characterized by a diversification of investments in different assets. In terms of performance, they dominate portfolios where all wealth is invested in a single asset. The only case that is an exception is the asset with ticker XLK, which is the one with the maximum expected return. To have a portfolio with maximum expected return, one must invest all the wealth in that asset.

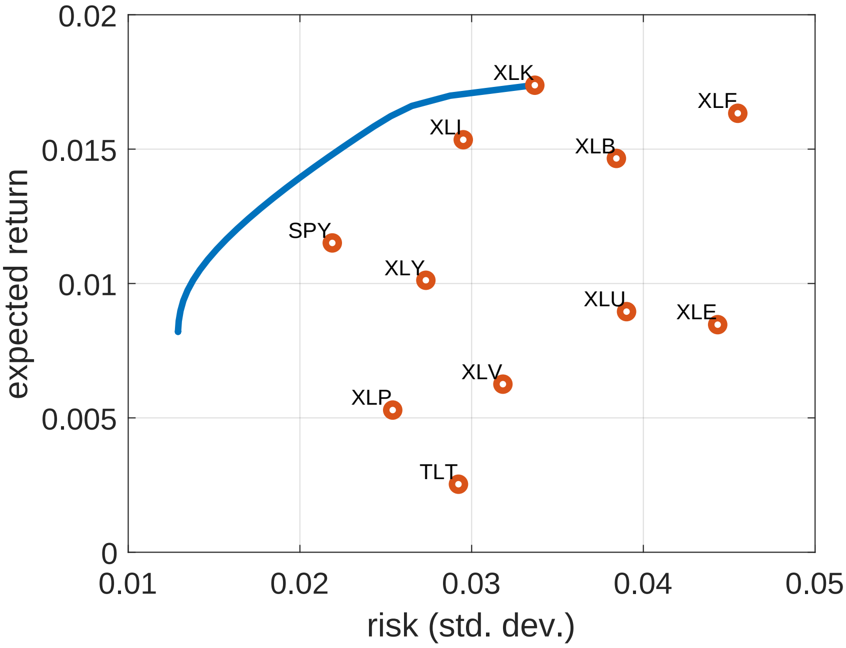

FIGURE 3.2: Efficient Frontier and individual assets performance.

Finally, figure 3.3 shows the size of the allocations in the assets of two different portfolios lying the efficient frontier, that we call x_1 and x_2, are taken into consideration, x_1 with expected return of 1.0% and standard deviation of 1.38% and x_2 with expected return 1.5% and standard deviation of 2.25%.

3.2.4 Sharpe ratio maximization

Markowitz mean-variance framework provides as solutions a collection of portfolios all lying on the Pareto efficient-frontier. The final choice on the single portfolio to be implemented is responsibility of the investor himself and depends on a combination of risk aversion and desired return. Typically an allocation strategy is evaluated based on a metric like the Sharpe ratio, and only a portfolio on the Pareto-optimal frontier achieves the maximum Sharpe ratio. As mentioned in section 2.5, Sharpe ratio quantifies the excess return (the expected portfolio return in excess of the risk-free rate, i.e., $\mu^T x - r_f$) normalized by the risk (the original definition consider the standard deviation of the portfolio $\sqrt{x^T \Sigma x}$ as the measure of risk), the

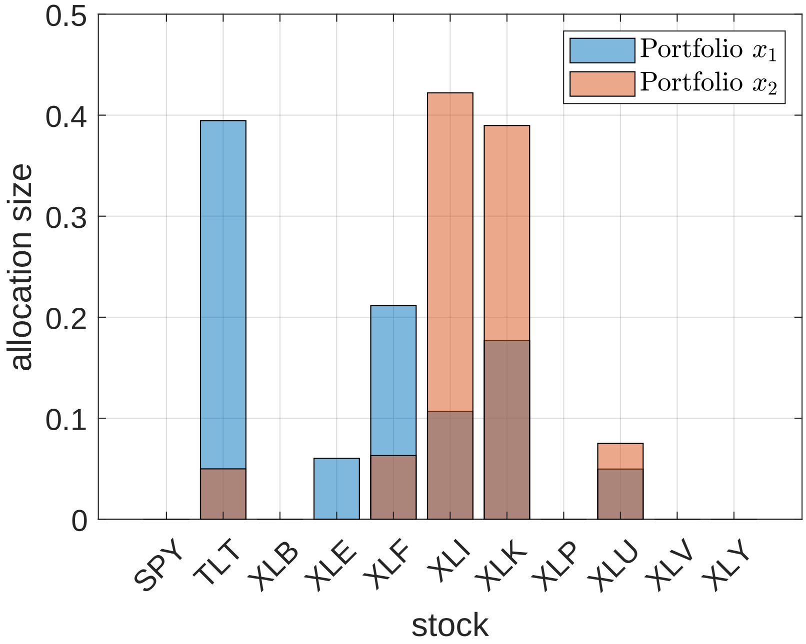

FIGURE 3.3: Size of allocations for two different portfolios x_1 and x_2.

higher the Sharpe Ratio of a portfolio, the greater the reward of the investment for the risk taken. The Sharpe Ratio maximization consists in the following optimization problem:

$$\max_{u} \quad \frac{\mu^T x - r_f}{\sqrt{x^T \Sigma x}} \tag{3.11a}$$

$$\text{s.t.} \quad \mu^T x > r_f, \tag{3.11b}$$

$$x \in \mathcal{X} \tag{3.11c}$$

Problem 3.11 is non-convex, however, through a series of adjustments it is possible to reformulate it as a Second Order Cone Program (SOCP) that can be solved efficiently, see [27] for details.

Geometrically the portfolio that maximizes problem 3.11 is the tangent point to the efficient frontier of the line passing through $(0, r_f)$ and called in the financial literature Capital Allocation Line (CAL) [47]. As a numerical example we solve problem 3.11 using the data of the numerical example of section 3.2.3 and assuming a risk-free rate $r_f = 2\times10^{-3}$. We obtain a Sharpe-optimal portfolio x^*, with expected return 1.17%, standard deviation 1.57% and Sharpe ratio 0.6138, lying on the efficient frontier at the point of tangency with the CAL line, as depicted in figure 3.4 with the red dot.

Finally, figure 3.5 shows the composition of the Sharpe-optimal portfolio x^*.

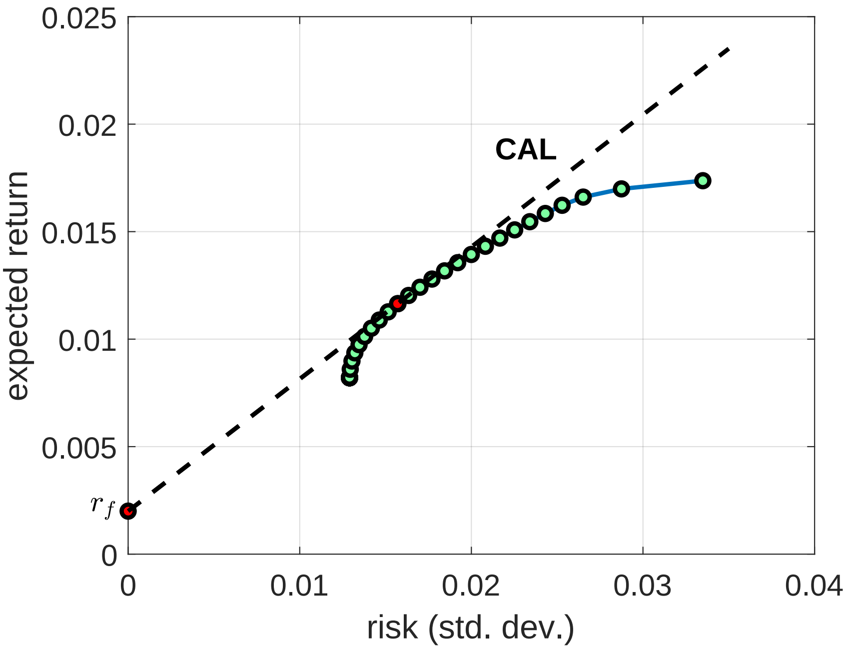

FIGURE 3.4: Sharpe-optimal portfolio for $r_f = 2 \times 10^{-3}$.

3.2.5 Practical constraints and transaction cost

In practice, in the design of optimal portfolios, one must take into account mandatory constraints imposed by market regularizations, capital budgets and the cost associated with single transactions which can largely influence the solution to the optimization problem. Furthermore, discretionary constraints can be included in order to avoid certain undesirable portfolio compositions.

Below is a non-exhaustive list of real portfolio constraints, for a larger list, see for example[20].

Long-only

Imposing the simple constraint:

$$x \geq 0$$

ensures that only long asset positions are held. In addition to having an immediate economic interpretation in the fact that many investment funds and institutional investors are not allowed to do short-selling, this constraint also has the effect of creating more stable portfolios [60].

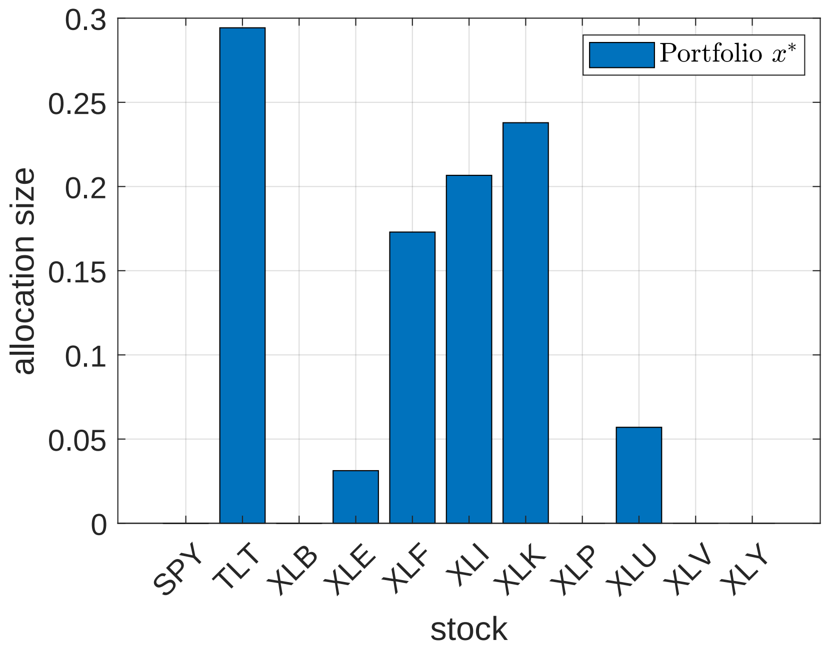

FIGURE 3.5: Size of allocations for the Sharpe-optimal portfolio x^*.

Turnover constraint

The turnover constraint is an optional linear constraint that enforces an upper bound on the average of purchases and sales. The turnover is defined by the 1-norm of the vector of the trades $\|u\|_1$, limiting it serves to limit the impact of transaction costs. We can impose the turnover constraint either on some particular asset i

$$|u| < \bar{u}_i$$

or on the whole portfolio

$$\|u\|_1 < \bar{u}$$

Sector bounds

It is possibile to limit the investment exposure in any individual asset or in a group of assets. The economic motivation of this constraint lies in the fact that the investor may not want to expose himself too much by investing in a particular sector made up, for example, by the first k assets and instead wants to maintain a certain degree of diversification in different sectors. This constraint can be imposed mathematically requesting that the total sum invested in the first k assets does not exceed a fraction α of the total wealth:

$$\sum_{i=i}^{k} x_i \leq \alpha \mathbf{1}^T x$$

Transaction costs

It is possible to take into account the presence of transaction cost in the model. We can consider a simple form of proportional transaction costs such that the investor incur in fee every time he buys or sells an asset:

$$\phi(u) = c\|u\|_1$$

where $c \geq 0$ is the unit cost of transaction. This model for transaction costs is easily generalizable, for instance, introducing proportional transaction cost depending on the sign of the trade or on the liquidity of each single asset. For example one can include a term bu to create asymmetry in the transaction cost function:

$$\phi(u) = c\|u\|_1 + bu$$

If $b = 0$ buying or selling generates the same transaction cost, if $b > 0$ it is cheaper to sell than to buy an asset. If transaction costs are taken into account the budget conservation condition must be modified to consider the money spent in the transaction:

$$\mathbf{1}^T u + \phi(u) = 0$$

3.2.6 Drawbacks of mean-variance optimization

The mean-variance framework represents one of the main achievements of quantitative methods in financial theory to the point of giving rise to a body of knowledge known as Modern Portfolio Theory [46]. However the Markowitz model presents 2 major drawbacks that have denied its practical application for many years:

1. it is high sensitive to parameter estimation errors. Estimating the covariance matrix Σ and especially the expected return μ is not an easy task [79]. This also has a strong negative effect on the out-of-sample performance of the method that could deviate greatly from the in-sample performance.

2. Variance and standard deviation are simmetric risk measures and therefore they penalize both under-average returns corresponding to unwanted high losses and above-average returns corresponding to welcome extra profits.

As a side note it also neglects multi period dynamics and long term investment objective. Regarding the first problem, the fact that the solution obtained with Markowitz formulation is highly unreliable, due to estimation errors, is well documented in literature [52], [53], [14], [34]. Typical solutions to the the first problem are the use of regularization techniques constraining portfolio norms to improve parameter estimation [65, 37] or the use of robust optimization techniques that naturally take into account the presence of uncertainty in the data [48]. Interestingly, in [22, 113] it is shown how the two approaches are actually linked. Typical solution to the second problem is, instead, the use of alternative asymmetric risk measures. In the next two sections, we will discuss the first problem in more detail and present the approaches based on regularization and subsequently on robust optimization, respectively. In the following section we will present a numerical example where the in-sample and out-of-sample performances of the optimal allocation strategies presented will be compared. Finally, in the section, we will introduce asymmetrical risk-based risk measures as possible solutions to the second problem discussed in this section.

3.2.7 Regularized portfolio optimization

Markowitz portfolio optimizatiom assume perfect knowledge of the expected value and covariance matrix of the assets return vector, that are fundamental inputs of the model. In reality we do not have access to these quantities, let alone the complete probability distribution of the random return vector. A possible solution is to use inferential statistics to estimate these parameters from a dataset of past realized returns. Given samples of realized returns r_t over time periods $t = 1, \dots, T$ the sample mean and the sample covariance are given by:

$$\hat{\mu} = \frac{1}{T} \sum_{k=1}^{T} r_t, \tag{3.12}$$

$$\hat{\Sigma} = \frac{1}{T-1} \sum_{k=1}^{T} (r_t - \hat{\mu})(r_t - \hat{\mu})^T \tag{3.13}$$

respectively. However, to have sufficiently accurate estimates using the sample estimators it would be necessary to have large T. If this is not the case, the sample mean produces noisy estimates [81]. In the case of estimation of the covariance matrix, having a large number of data is fundamental if the number of assets considered is high. Indeed,the covariance matrix consists of $n \times n$ entries (of which only half is necessary since the matrix is simmetric). It is clear that, if n is high, to correctly estimate a large quantity of parameters very long time series would be required, something that is not always possible, for example in the case of recent companies. What is even more limiting is that the lack of stationary in historical price series (due to changes in market conditions or the companies themselves) precludes the use of the complete time series [82].

In the literature, several methods have been proposed to deal with estimating the large number of elements of the covariance matrix in a reliable manner. A first trivial approach is to use higher frequency data such as daily market data [60]. A second approach is to assume some kind of structure on the estimator of the covariance matrix based on the fact that the market is driven by a limited number of common risk factors. For instances, the oil price should be a relevant risk factor for stock shares in the automotive industry, whereas this risk factor might play a less relevant role for stocks in the telecommunication industry [22]. On the basis of such considerations, in [32] a series of factor models are proposed which are able to reduce the number of parameters to be estimated and therefore mitigate the impact of possible estimation errors. Another approach to improve the quality of the estimation is based on *shrinkage estimators*. These types of estimators are based on the classical finding of Stein [105] that biased estimators may be superior to unbiased ones. For example, in [65] is proposed as an estimator a weighted average of the sample covariance matrix and the identity matrix. This has the effect of shrinking the sample covariance matrix towards the identity matrix. In [53, 60] it is shown that imposing a no short-selling constraint on the optimization problem is equivalent to shrinking the extreme elements of the covariance matrix.

More recently, it has been observed that the strategies developed in [65, 53, 60] were special cases of a more general framework in which portfolio weights are treated, rather than the moments of asset returns, as the objects of interest to be estimated [37]. This approach consists, instead of shrinking the moments of asset returns, in solving the mean-variance optimization problem by imposing constraints on the norm of the portfolio weights. It is shown that by imposing the constraint that the l_1 norm of the portfolio weights is less than 1 then the short-sale constrained portfolio of [53, 60] is obtained. On the other hand, by imposing a constraint on the l_2 norm of the portfolio vector, the portfolio developed in [65] is obtained.

This approach is closely related to the statistical literature of regression problems with regularization of parameters. In fact it is observed that the weights of the portfolio can be treated as coefficients of an Ordinary Least Square regression problem. By imposing a constraint on the l_1 norm of the portfolio vector to be less than a certain threshold, a problem is obtained which is equivalent to *lasso regression* [108], while imposing a constraint on the l_2 norm of the portfolio vector corresponds to *ridge regression* [58]. In [37] the so-called norm-constrained mean-variance portfolio is then defined as follows:

$$\begin{aligned}\min_u \quad & x^T \Sigma x && (3.14a)\\ \text{s.t.} \quad & \mu^T x \geq \gamma, && (3.14b)\\ & \|x\|_p \leq \delta, && (3.14c)\\ & x \in \mathcal{X} && (3.14d)\end{aligned}$$

where $\delta \geq 0$ is a certain threshold and $p = 1, 2$ tipically. Problem 3.14 represents a regularized version of problem 3.9 and is expressed in the so-called *Ivanov form*. Problem 3.14 can easily be rearranged in the so-called *Tikhonov form*:

$$\begin{aligned}\min_u \quad & x^T \Sigma x + \lambda \|x\|_p && (3.15a)\\ \text{s.t.} \quad & \mu^T x \geq \gamma, && (3.15b)\\ & x \in \mathcal{X} && (3.15c)\end{aligned}$$

where $\lambda \geq 0$ is a parameter that controls the strength of the regularization term $\|x\|_p$.

The first advantage of using the regularized version of the mean-variance portfolio optimization problem with respect to the original non-regularized version of Markowitz is that it returns a solution less sensitive to estimaion errors of the model's inputs. The second important advantsge is that it is possible to determine the values of the parameters δ and λ that optimize the out-o-sample performance by means of historical data and cross-validation [44]. In recent years this approach has been further developed, for example [50] considers even more complex penalty functions while [7] combine regularization and cross-validation proposing *performance-based-regularization* to optimize out-of-sample performance.

3.2.8 Robust portfolio optimization

In this section we will discuss the second possible solution to the first problem mentioned in section 3.2.6, which is based on the use of the robust optimization framework [13].

The basic idea is to take into account the uncertainty present in the data and to find portfolios compositions that are robust against such uncertainties. Over the years two techniques have been developed called stochastic programming and robust optimization. In stochastic programming one assumes that the uncertain parameters of the model follow a known probability distribution and the constraints of the model are converted from inequalities to probabilistic constraints [101]. However, the main problems related to this framework are that in reality the exact distribution of the uncertain parameters is not known, furthermore the problems of stochastic programming are typically difficult to solve. Robust optimization , on the other hand, does not characterize the uncertainty of the parameters through a stochastic characterization of uncertainty but rather through the belonging of these parameters to a bounded uncertainty set $\mathcal{U}$ that contains all scenarios one wants to safeguard against [13]. The solution to a robust optimization problem is feasible if all the constraints are satisfied for all possible variations of the uncertain parameters of the uncertainty set.

Suppose then that the expected value of the returns μ and the covariance matrix Σ belong respectively to the sets $\mathcal{U}_r$ and $\mathcal{U}_\Sigma$, the robust counterpart of the risk minimization problem has the following formulation:

$$\min_{u} \quad \sup_{\Sigma \in \mathcal{U}_\Sigma} x^T \Sigma x \tag{3.16a}$$

$$\text{s.t.} \quad \inf_{\mu \in \mathcal{U}_r} \mu^T x > \gamma \tag{3.16b}$$

$$x \in \mathcal{X} \tag{3.16c}$$

The difficulty in solving a robust optimization problem depends on the kind of uncertainty set we consider. Problem 3.16 can be solved efficiently only in the case of simple uncertainty set, for instance if $\mathcal{U}_r$ and $\mathcal{U}_\Sigma$ are interval (box) sets of the form:

$$\mathcal{U}_r = \{\mu : r_{min} \leq \mu \leq r_{max}\}, \tag{3.17}$$

$$\mathcal{U}_\Sigma = \{\Sigma : \Sigma_{min} \leq \Sigma \leq \Sigma_{max}, \Sigma \geq 0\} \tag{3.18}$$

then, under the further assumptions that $x \geq 0$ and $\Sigma_{max} \geq 0$, the double-layered minimax problem 3.16 reduces into the single-layer convex minimization problem:

$$\min_{u} \quad x^T \Sigma_{max} x \tag{3.19a}$$

$$\text{s.t.} \quad r_{min}^T x > \gamma \tag{3.19b}$$

$$\mathbf{1}^T u = 0, \tag{3.19c}$$

$$x = x_0 + u, \tag{3.19d}$$

$$x \geq 0. \tag{3.19e}$$

The price one pay to solve a robust mean-variance portfolio optimization problem is the fact that generally its solution is excessively conservative as the goal is to safeguard from the worst case scenario, and consequently it may not be possible to find a solution capable of guaranteeing the investment objective desired by the investor. Moreover, often such conservatism can lead to the robust counterpart being unfeasible, ie there may not be a solution that is able to safeguard against uncertainty at all.

Scenario approach to robust portfolio optimization

An approach that consists in a trade-off between an excessively optimistic approach such as the classic mean-variance optimization and an excessively conservative approach such as that of robust optimization is represented by the scenario approach [26] that can be interpreted as an approximation of the robust optimization approach. More precisely, the scenario optimization approach is a technique for obtaining solutions to robust optimization problems using random samples of the constraints called scenarios.

Suppose that the uncertainty sets $\mathcal{U}_r$ and $\mathcal{U}_\Sigma$ are governed by some probability distributions. Suppose we can be able to obtain N iid samples $(\mu^{(i)}, \Sigma^{(i)})$, $i = 1, \ldots, N$ distributed according to the probability distributions of the uncertainty sets. It should be pointed out that the actual distributions does not necessarily need to be known exactly. Indeed, one can resort in the exploitation of a parametric model or, more simply, the scenarios can be obtained by observing past outcomes. In this last case we talk about *data-driven optimization* [25]. The scenarios can be collected in the sets $\hat{\mathcal{U}}_r$ and $\hat{\mathcal{U}}_\Sigma$, that are:

$$\hat{\mathcal{U}}_r = \{\mu^{(1)}, \mu^{(2)}, \ldots, \mu^{(N)}\}, \tag{3.20}$$

$$\hat{\mathcal{U}}_\Sigma = \{\Sigma^{(1)}, \Sigma^{(2)}, \ldots, \Sigma^{(N)}\} \tag{3.21}$$

and can be seen as discrete approximations of the original uncertainty sets. In this sense scenario approach is an approximation of the robust approach. Then, we solve the scenario optimization problem:

$$\min_{x,t} \quad t \tag{3.22a}$$

$$\text{s.t.} \quad x^T \Sigma^{(i)} x \leq t, \qquad i = 1, \ldots, N, \tag{3.22b}$$

$$\mu^{(i)T} \geq \gamma, \qquad i = 1, \ldots, N, \tag{3.22c}$$

$$x \in \mathcal{X}. \tag{3.22d}$$

It must be pointed out that the solution obtainable from the scenario formulation 3.53 will be robust only with respect to the considered scenarios. In other words, there is no absolute guarantee that the solution obtained is robust for all the values of the original uncertainty set. However, it is shown in [27] that by choosing N sufficiently large the approximate solution obtainable by solving problem 3.53 is *probabilistically robust* up to a predetermined confidence level $\alpha \in (0,1)$, ie it is possible to fix a degree of probabilistic robustness α a priori and then choose N according to the rule:

$$N \geq \frac{2}{1-\alpha}(d+10) \tag{3.23}$$

where d is the dimensionality of the decision variable.

3.2.9 Numerical example 2 - Different allocation strategies comparison

The objective of this section is to present the results of a numerical experiment in which some of the optimal portfolio allocation strategies presented in this chapter are compared. To do this we use the same universe of $n = 11$ assets of section 3.2.3, but we consider a broader time horizon covering the period from 8/2002 to 7/2019, for a total of $T = 204$ monthly data returns for each asset.

Figure 3.6 shows the normalized monthly prices of the 11 financial assets considered and their evolution over the considered time horizon.

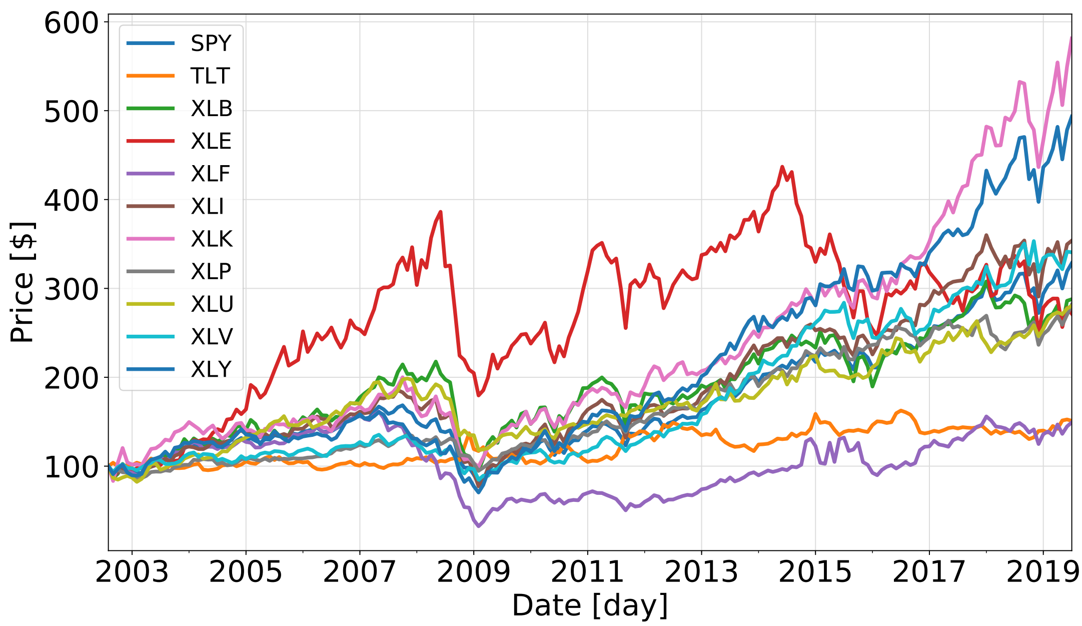

FIGURE 3.6: Normalized prices of financial instrument over time.

We will compare the following 4 different portfolio allocation strategies:

1. classical mean-variance (MV) portfolio, corresponding to optimization problem 3.9;

2. regularized mean-variance (ReMV) portfolio, corresponding to optimization problem 3.15 with regularization parameter $\lambda = 0.01$;

3. robust mean-variance (RoMV) portfolio, corresponding to optimization problem 3.19;

4. scenario based mean-variance (SMV) portfolio, correponding to optimization problem 3.22 with $N = 200$ scenarios.

We now describe a reliable methodology to compute the out-of-sample performances of the 4 different allocation strategy, procedure that has been introduced in [37].

Out-of-sample performance evaluation

The evaluation of the out-of-sample performance of a portfolio construction strategy can be performed through back-testing in a rolling-horizon fashion.

First of all, we choose as the metric to assess the validity of the allocation strategies the out-of-sample portfolio Sharpe Ratio. Assuming to have a dataset available made of T observations of past returns, the procedure consists in initially determining a window of τ observations that will be used as in-sample data to make estimates of the expected value and covariance matrix of returns using the sample estimators, thanks to these etimates we build the optimal portfolio and then, if of interest, we compute the in-sample evaluation metrics. The optimal portfolio allocations are kept unchanged for the following period (can be a day, a week, a month, ...) and out-of-sample performances are recorded using the real returns occurred in this period. Subsequently, the estimation window is moved forward so that latter out-of-sample returns are included in the in-sample data and eventually the least recent data are discarded. The updated in-sample data window is then used to re-compute the sample estimates of expected returns and covariance and form a new optimal portfolio. This procedure is repeated as long as the end of the dataset is reached. At the end of the process we have generated $T - \tau$ portfolio weight vectors, that is x_k for $k = \tau, \dots, T-1$. Holding the portfolio x_k for one period gives the out-of-sample return at time $k + 1 : r_{p,k+1} = r_{k+1}^T x_k$ where r_{k+1} denotes the asset returns. With the time series of returns and weights it is possible to compute out-of-sample Sharpe ratio :

$$\hat{\sigma}^2 = \frac{1}{T - \tau - 1} \sum_{k=\tau}^{T-1} (r_{k+1}^T x_k - \hat{\mu}_p)^2 \tag{3.24}$$

where

$$\hat{\mu}_p = \frac{1}{T - \tau} \sum_{k=\tau}^{T-1} r_{k+1}^T x_k \tag{3.25}$$

and

$$\hat{\mathrm{SR}} = \frac{\hat{\mu}_p}{\hat{\sigma}} \tag{3.26}$$

Out-of-sample performance results

In our experiment, we set as the initial investment date 1/2010. With this choice, we have a moving window of lenght $\tau = 89$ of past in-sample data. Every month we use the in-sample data to estimate the sample mean and sample covariance of returns for the MV and ReMV strategies. For the SMV strategy we bootstrap the in-sample data to generate $N = 100$ scenarios $(\mu^{(i)}, \Sigma^{(i)})$, $i = 1, \dots, 100$ of the mean and covariance of returns. For the RoMV, we use the N generated scenarios and the same technique of [109] to estimate the lower and upper bound r_{min},Σ_{max} of the asset returns and covariance matrix respectively. We solve all the 4 allocation strategies considered for 20 increasing values of target return γ to compute the in-sample efficient frontiers and the optimal weights. We then move the window 1 month forward and we repeat the same procedure for a total of $T - \tau = 115$ times. This allows us to compute 115 in-sample-efficient frontiers for every tested allocation strategies, that we aggregate together to form Figure 3.7. From the graph of Figure 3.7 it can be seen that the strategy that allows to obtain the best in-sample performances is the MV strategy followed by the ReMV strategy. The portfolios obtained with the RoMV and SMV strategies show, instead, much more pessimistic results. The maximum in-sample SRs obtained with the 4 strategies are 0.2479 for the MV portfolio, 0.2430 for the ReMV portfolio, 0.1096 for the SMV portfolio and 0.0538 for the RoMV portfolio.

As explained in the previous section using the time series of optimal weights of the 4 strategies and the real returns we are able to calculate the 115 out-of-sample efficient frontiers, which are aggregated and shown in Figure 3.8. From the graph of Figure 3.8 graph we see an almost opposite situation compared to the in-sample results, in fact, the MV portfolio is the one that provides worse out-of-sample performances. The maximum out-of-sample SRs obtained with the 4 strategies are 0.2254 for the MV portfolio, 0.2459 for the ReMV portfolio, 0.2594 for the SMV portfolio and 0.2435 for the RoMV portfolio.

3.2.10 Alternative risk measure for portfolio design

In this section we will discuss the typical solution to the second problem of the classical Markowitz framework highlighted in section 3.2.6.

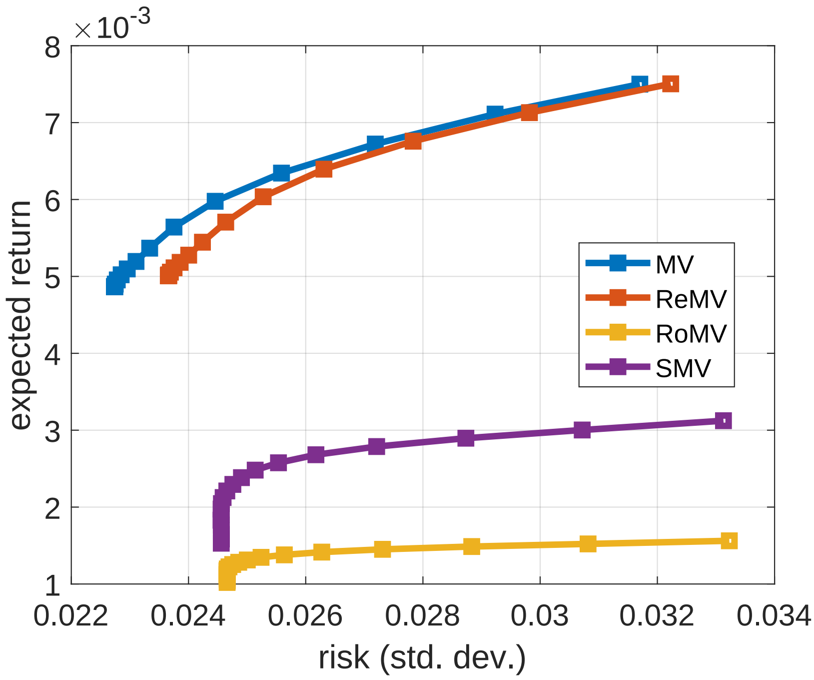

FIGURE 3.7: In-sample efficient frontiers for the 4 tested strategies.

In finance, the expected return of a portfolio is a fundamental quantity to be taken into consideration as it quantifies the average benefit of the investment. However, in reality, returns are not the only entity of interest as they are typically not sufficiently constant to allow the investor not to worry about the likelihood of going bankrupt. For this reason we use measures that quantify how risky an investment strategy is. The most classic risk measure of an investment, cosidered by Markowitz himself, is variance, it quantifies the variability in the performance of an investment strategy, a large variance can mean high earnings but at the same time large losses.

However, Markowitz himself, a few years after his most famous work, pointed out the inadequacy of variance as a risk measure of a portfolio strategy [74]. In fact, as mentioned in the section, the variance being a symmetrical risk measure penalizes both undesirable losses and desirable gains. Clearly, this fact is a strong limitation of the mean-variance optimization framework as each investor would just like to penalize losses.

To overcome this negative practical aspect of variance over the years, a series of alternative and asymmetrical risk measures have been proposed in the financial literature. Among the most famous

- Value-at-risk (VaR)
- Conditional Value at Risk (CVaR)
- Downside Risk (DR)

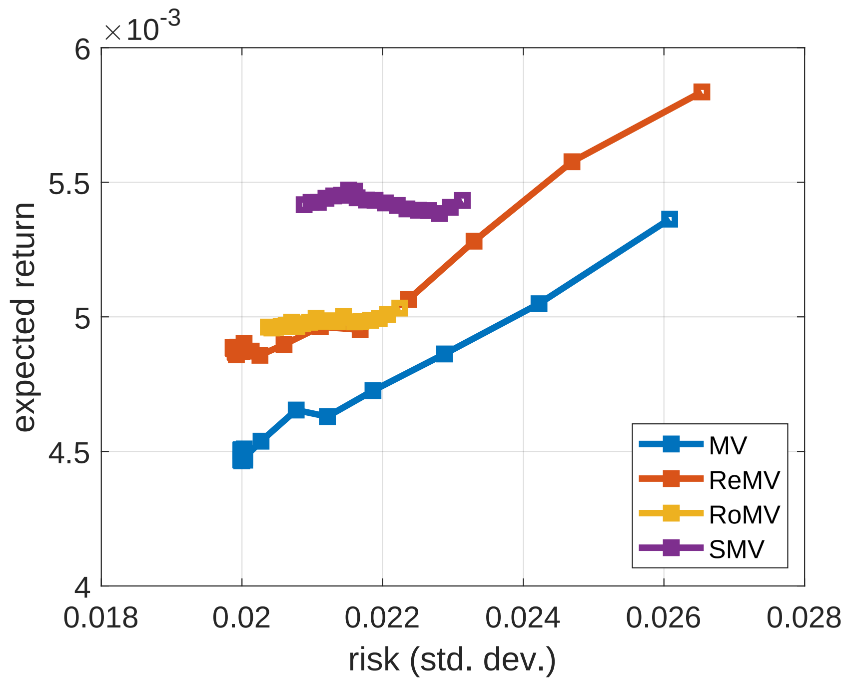

FIGURE 3.8: Out-of-sample efficient frontiers for the 4 tested strategies.

Value-at-Risk (VaR)

Value-at-Risk (VaR) is a quantile-based risk measure that was developed at J.P. Morgan made public in 1997. The idea behind it is to focus only on the negative tail of the distribution that is where the losses occur.

Let ξ be a random variable that represents the loss of a portfolio x over some period of time, where by loss we mean the negative return $\xi = -r^T x$. The VaR of portfolio x with confidence level $\alpha \in (0, 1)$, is defined as

$$\mathrm{VaR}_\alpha = \inf_{\xi_0} : \mathrm{Prob}(\xi \leq \xi_0) \geq \alpha \tag{3.27}$$

That is, the probability that the investment results in a greater loss than VaR_α is not greater than α, say, $\alpha = 0.95$. Although this risk measure has been widely used in practice it does not take into account the losses that exceed the VaR value and is also not convex.

Conditional Value-at-Risk (CVaR)

To overcome the limitations of VaR, but at the same time preserving some of its desirable characteristics, it is to consider conditional expectation on the tail. This led Rockafellar and Uryasev [95] to the definition of an alternative risk measure called Conditional Value-at-Risk (CVaR), or Expected Shortfall (ES).

CVaR is defined as the conditional expectation of losses exceeding the VaR_α

$$\text{CVaR}_\alpha = \mathbb{E}[\xi|\xi \geq \text{VaR}_\alpha] \tag{3.28}$$

The added benefit of CVaR is that it is a convex risk measure.

Downside Risk (DR)

Recalling that we indicate with r_p the random variable that models the return of a portfolio, it is possible to easily define an asymmetric risk measure directly based on variance, and called semi-variance or in financial terms downside risk, and given by

$$\text{SV} = \mathbb{E}[(\max\{0, \mathbb{E}[r_p] - r_p\})^2] \tag{3.29}$$

It should be noted that semi-variance is actually a special case of Lower Partial Moments (LPM)

$$\text{LPM}_\nu = \mathbb{E}[(\max\{0, \mathbb{E}[r_p] - r_p\})^\nu] \tag{3.30}$$

In practice, we consider only negative deviations with respect to the expected value. The idea can be generalized and made more flexible introducing negative deviations with respect to a minimun target that we wish to achieve, such negative deviations are called shortfall amounts. For instance, denoting with γ the desired target wealth, the expected shortfall is given by

$$\text{ES} = \mathbb{E}[(\max\{0, \gamma - w\})^\nu] \tag{3.31}$$

Shortfall is zero if we achieve or exceed the target, is insted a positive value if the target is not achieved. In this sense, this risk measure is asymmetric.

This risk measure is particularly convenient if we have a finite set of N scenarios available. In fact it is possible to approximate the expectation discretizing it in the following way

$$\hat{\text{ES}} = \frac{1}{N}\sum_{i=1}^{N}(\max\{0, \gamma - w^{(i)}\})^\nu \tag{3.32}$$

Common choices for *nu* are $\nu = 1$ and $\nu = 2$, with the intuition that the higher the degree the higher the risk aversion of the investor.

3.3 Multi-period portfolio optimization

In this section we will discuss an important extension to the optimization models discussed so far.

Indeed, the approaches described in section 3.2 are conducted in a single-period framework, ie they are based on the optimization of some portfolio performance to be obtained at the end of one single period. In realistic situations the investor is forced to re-address the allocation problem at the end of the period, furthermore it is reasonable to suppose that the investor has long-term objectives. The simplest solution would be to execute a sequence of single-period optimization problems, but this, as explained in [83] would be beneficial only in the absence of transaction costs and if the price returns were temporally independent. This strategy would also be short-sighted towards long-term objectives, for which it is useful to reformulate the problem of allocating resources over a time horizon composed of multiple periods with the aim, for example, of minimizing the total risk during the investment horizon at the same time time satisfying a series of constraints on the composition of the portfolio and on the expected return during the intermediate stages.

At first, multi-period formulations of the portfolio optimization problem were conceived as a stochastic control problem where the trader is allowed to vary over time the number of shares of the financial assets that make up his own portfolio. Particularly noteworthy are the contributions given by Merton which proposed a continuos-time stochastic dynamic programming approach to the multi-period allocation problem [77]. With no constraints or transaction cost, and under some additional assumptions, Merton derived a closed-form expression for the optimal policy. In the same period, Samuelson derived the discrete time version of Merton's approach [96]. However, the dynamic programming approach suffers from the curse of dimensionality and therefore their practical application is reported only in the case of few securities and very short time horizons. Furthermore, such approaches do not take into consideration realistic market features such as constraints on portfolio composition and transaction costs. The addition of these ingredients excludes the possibility of having solutions in closed form and makes the problem even more difficult from a computational point of view.

It is precisely in this context that the methods based on predictive control, although they provide suboptimal policy for the multi-period investment problem, are successful. What makes these methods attractive in the context of multi-period portolio optimization is the fact that they can naturally handle constraints on the composition of the portfolio and that they are suitable for practical implementations since they can be solved in useful times. Some formulations can even be solved as convex quadratic programs.

Examples of application of predictive control techniques for multi-stage financial problems such as multi-period portfolio allocation, wealth tracking, pairs trading, can be found in [41, 42, 57, 56], [104] and [114] respectively.

Most predictive control formulations applied to finance problems use an open-loop on-line optimization, however some formulations use some form of feedback control, such as the approaches of Calafiore in [29, 30, 31] that will be discussed in detail in the following sections because of particular interest for this dissertation.

3.3.1 Model and portfolio dynamics

In this paragraph, we will describe the model that will be used for multi-period portfolio optimization. We will follow the formulation in [29, 30, 31].

We consider a universe of n assets $a_1, \dots, a_n$ and we define with $p_i(k)$ the market value of the asset a_i at time k. We then define with $x_i(k)$ the monetary amount allocated in the asset a_i at time k. A value of $x_i(k) \geq 0$ represents a long position while a value $x_i(k) < 0$ represents a short position. The vector $x(k) \in \mathbb{R}^n$ with elements $x_i(k), i = 1, ..., n$ defines the portfolio of the investor.

From a control point of view, the amount of money $x_i(k)$ allocated in the individual assets represents the state variables of the model while the portfolio represents the state vector. The investor total wealth at time k is given by:

$$w(k) = \sum_{i=1}^{n} x_i(k) = \mathbf{1}^T x(k) \tag{3.33}$$

The control variables of the model $u_i(k)$ collected in the vector $u(k)$ are the monetary amount of the asset i that is sold (if $u_i(k) > 0$) or bought (if $u_i(k) < 0$) at the beginning of the period k. The disturbance that affects the system are the price returns $r_i(k)$, random variable that captures the performances of asset a_i between time $k-1$ and k and are defined as follows:

$$r_i(k) = \frac{p_i(k) - p_i(k-1)}{p_i(k-1)}, \quad i = 1, \dots, n; \quad k = 1, \dots, T \tag{3.34}$$

For the writing of the model it is also necessary to define the gross returns or gains as:

$$g_i(k) = \frac{p_i(k)}{p_i(k-1)} = 1 + r_i(k), \quad i = 1, \dots, n; \quad k = 1, \dots, T \tag{3.35}$$

The gains at time period k are collected in the diagonal matrix $G(k) \in \mathbb{R}^{n \times n}$:

$$G(k) = \begin{bmatrix} g_1(k) & 0 & 0 & \dots & 0 \\ 0 & g_2(k) & 0 & \dots & 0 \\ 0 & 0 & g_3(k) & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & g_n(k) \end{bmatrix} \tag{3.36}$$

and which role will be clear soon.

The model can be enriched by a set of constraints on states and control variables. For example, it is typically assumed that the portfolio is self-financing, this behavior can be modeled in a simple way by adding the constraint:

$$\sum_{i=1}^{n} u_i(k) = \mathbf{1}^T u(k) = 0 \quad k = 0, \dots, T-1 \tag{3.37}$$

In certain circumstances, short-selling could be prohibited, to model this it is sufficient to require the components of the portfolio to be positive or null:

$$x(k) \geq 0 \quad k = 0, \dots, T-1 \tag{3.38}$$

For other types of constraints on the composition of the portfolio, please refer to section 3.2.5.

Defined all the above quantities, it is possible to derive the dynamic model of the portfolio. The dynamic is due to the fact that the monetary value invested in one asset changes over time due to changes in price between one period and another. The transition equation of the state is given by:

$$x(k+1) = G(k+1)[x(k) + u(k)], \quad k = 0, \dots, T-1 \tag{3.39}$$

3.3.2 Multi-period mean-variance optimization

In [29] the author derives a multi-period version of the classical Markowit's mean-variance optimization model.

The total investment risk is quantified as a weighted sum of all stage wealth volatilities, measured as wealth variances:

$$J(T) = \sum_{k=1}^{T} \delta(k) \mathrm{Var}[w(k)] \tag{3.40}$$

where $\delta(k) \geq 0$ are given weights. Choosing $\delta(k) = 1/T, k = 1, \ldots, T$ we equally penalize the volatility of wealth along all the stages of the investment horizon, while choosing $\delta(T) = 1$ and $\delta(k) = 0, k = 1, \ldots, T-1$ we penalize only large values of final wealth variance.

Similarly to the optimization problems discussed in section 3.2.2 we choose optimal allocation strategies that minimizes risk $J(T)$ while satisfying a lower-bound constraint for the total return at the final stage:

$$\mathbb{E}[w(T)] \geq \gamma \tag{3.41}$$

and that at the same time satisfy any constraints on the composition of the portfolio such those discussed in section 3.2.5.

Open-loop formulation

A first, 'naive' version of the multi-period mean-variance optimization problem described in [29] is an open loop strategy in which the whole sequence of control actions (decision variables) $u = (u(0), \ldots, u(T-1))$ is computed at time $k = 0$ and held unchanged over the investment horizon.

Finally, in [29] it is shown that being the wealth variance at each stage a quadratic function of the decision variables $u(k)$ and the portfolio expectations affine functions of the decision variables $u(k)$, the open-loop multi-period optimization problem can be casted as the following convex quadratic programming problem:

$$\min_{u(0),\ldots,u(T-1)} \quad \sum_{k=1}^{T} \delta(k)\mathrm{Var}[w(k)] \tag{3.42a}$$

$$\text{s.t.} \quad \mathbb{E}[w(T)] \geq \gamma, \tag{3.42b}$$

$$\mathbf{1}^T u(k) = 0, \qquad k = 0, \ldots, T-1 \tag{3.42c}$$

$$x \in \mathcal{X}(k) \qquad k = 0, \ldots, T-1 \tag{3.42d}$$

where $\mathcal{X}(k)$ is the set of admissible portfolios (considering only satisfying portfolio dynamics and no constraints):

$$\mathcal{X}(k) = \{x(k) : x(k) \in \mathbb{R}^n,\ x(k+1) = G(k+1)[x(k) + u(k)]\} \tag{3.43}$$

Closed-loop formulation

The open-loop strategy described in the previous paragraph is not able to exploit the sequential nature of the problem, indeed only the first control action $u(0)$ must be implemented immediately, while the subsequent ones, $u(1), \ldots, u(T-1)$, can be progressively recomputed taking into consideration new market observations. At $k \geq 1$ we have observed $k-1$ return realizations for each single asset, a natural strategy to exploit this new information is to look for control actions within a parametric class $\mathcal{U}$ of causal functions of the observed returns, called *policies*, to conceive the closed-loop version of the multi-period portfolio optimization problem. The class of functions $\mathcal{U}$ must be wide enough to contain effective policies but at the same time sufficiently simple to allow its exploration in useful times for practical implementations.

The choice adopted in [29] is that of restricting the class of functions from which to look for policies to that of the affine functions of the most recent observed returns so that they can be easily computed through convex optimization techniques. That is, control actions of the following form are considered:

$$u(0) = \bar{u}(0) \tag{3.44}$$

$$u(k) = \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k)) \tag{3.45}$$

where $\bar{u}(0)$ is the initial control action made before observing the firt realization of market gains, $g(k)$ is the vector of gains at time k and $\bar{g}(k)$ is a given estimate of the expected value of $g(k)$, $\bar{u}(k) \in \mathbb{R}^n$ and $\Theta(k) \in \mathbb{R}^{n \times n}$ are the policy parameters. The policy has the following interpretation, if the market behaves as expected the control action is reduced to the nominal action $\bar{u}(k)$, if the market behavior deviates from that expected the nominal control action is corrected by a term proportional to the innovation of the returns, $g(k)-\bar{g}(k)$, and that depends on the policy parameters contained in the matrix $\Theta(k)$.

Given the policy structure 3.44, the state transition equations may be rewritten as:

$$x(1) = G(1)[x(0) + \bar{u}(0)], \tag{3.46}$$

$$x(k+1) = G(k+1)[x(k) + \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k))], \quad k = 1, \ldots, T-1 \tag{3.47}$$

It can be shown that the total wealth variance $\mathrm{Var}[w(k)]$ is a convex quandratic function of the decision variables $\bar{u}(k)$ and $\Theta(k)$ [29].

The closed-loop multi-period optimization problem can be casted as the convex quadratic programming problem:

$$\min_{u(0),\ldots,u(T-1);\Theta(1),\ldots,\Theta(T-1)} \quad \sum_{k=1}^{T} \delta(k)\mathrm{Var}[w(k)] \tag{3.48a}$$

$$
\begin{aligned}
\text{s.t.} \quad & \mathbb{E}[w(T)] \geq \gamma, & (3.48b)\\
& \mathbf{1}^T \bar{u}(k) = 0, \qquad k = 0, \ldots, T-1 & (3.48c)\\
& \mathbf{1}^T \Theta(k) = 0, \qquad k = 1, \ldots, T-1 & (3.48d)\\
& x \in \mathcal{X}(k) \qquad k = 0, \ldots, T-1 & (3.48e)
\end{aligned}
$$

where $\mathcal{X}(k)$ is the set of admissible portfolios:

$$
\mathcal{X}(k) = \{x(k) : x(k) \in \mathbb{R}^n,\ x(k+1) = G(k+1)[x(k) + \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k))]\} \tag{3.49}
$$

3.3.3 Scenario-based multi-period portfolio optimization

The multi-stage mean-variance optimization model described in the last section suffers from the same drawbacks of his single-stage counterparts, it relies on noisy estimates of the covariances $\Sigma(k)$ and expected gains $\bar{g}(k)$ by means of past time series analysis and/or elicted by expert advice, and uses the variance of the final stage wealth as a measure of risk, that is a simmetric risk measure.

In this section we describe the Calafiore approach in [31] that seeks to overcome the aforementioned limitations using the scenario approach of Section 3.2.8 and using the discretized version of the Expected Shortfall described in Section 3.2.10 as a risk measure.

It is now assumed that a scenario generation mechanism is available to generate N independent and identically distributed sample paths that produce N scenarios of gains:

$$
g(k)^{(i)}, G(k)^{(i)}, \qquad k = 1, \ldots, T, i = 1, \ldots, N \tag{3.50}
$$

At this point it is possible to formulate the counterparts of problems 3.42 and 3.48 based on the scenario approach and with the expected shortfall as risk measure.

Open-loop formulation

$$
\begin{aligned}
\min_{u(0),\ldots,u(T-1)} \quad & \frac{1}{N} \sum_{i=1}^{N} \max(0, \gamma - w^{(i)}(T))^2 & (3.51a)\\
\text{s.t.} \quad & x^{(i)} \in \mathcal{X}^{ol}(k) \qquad k = 0, \ldots, T-1,\ i = 1, \ldots, N & (3.51b)\\
& \mathbf{1}^T u(k) = 0, \qquad k = 0, \ldots, T-1 & (3.51c)
\end{aligned}
$$

with

$$\mathcal{X}^{ol}(k) = \{x(k) : x(k) \in \mathbb{R}^n,\ x(k+1) = G(k+1)[x(k) + u(k)]\} \tag{3.52}$$

Closed-loop formulation

$$\min_{\bar{u}(0),\dots,\bar{u}(T-1);\Theta(1),\dots,\Theta(T-1)} \quad \frac{1}{N}\sum_{i=1}^{N} \max(0, \gamma - w^{(i)}(T))^2 \tag{3.53a}$$

$$\text{s.t.} \quad x^{(i)} \in \mathcal{X}^{cl}(k) \qquad k = 0, \dots, T-1,\ i = 1, \dots, N \tag{3.53b}$$

$$\mathbf{1}^T u(k) = 0, \qquad k = 0, \dots, T-1 \tag{3.53c}$$

$$\mathbf{1}^T \Theta(k) = 0, \qquad k = 1, \dots, T-1 \tag{3.53d}$$

$$\tag{3.53e}$$

with

$$\mathcal{X}^{cl}(k) = \{x(k) : x(k) \in \mathbb{R}^n,\ x(k+1) = G(k+1)[x(k) + \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k))]\} \tag{3.54}$$

Chapter 4

Stock trading via feedback control: an extremum seeking approach

4.1 Introduction

In this chapter, we describe the first contribution of this dissertation to the literature on individual asset trading and in particular to stock trading via feedback control addressed in Chapter 2.

In 2, we reported an overview of the innovative approach to stock trading presented by B. Ross Barmish in [8] and follow-up works (e.g., [70, 71]). It has been shown that gaining money against unpredictable price variations can be reformulated as a *control design problem* with a disturbance rejection goal. Within such a framework, the disturbance, namely the price variations, does not need to be modeled, thus the approach can be referred to as *model-free* [10, 12].

Notwithstanding the basic idea is as clever as simple, the tuning of the controller defining the best investment level for a single stock is all but a straightforward task. In 2.4 we reported a description of a two degrees of freedom controller named *Simultaneous Long-Short (SLS)* controller designed to combine both a long and a short strategy. One control block is tuned to implement a long position, that is to perform well in all scenarios with a rising price (also known as *bull moments*), where the trader will profit from first buying and then selling the stock. The other control block is aimed to maximize the return when the price decreases (during the so-called *bear moments*), thus implementing a short investment position, where the trader profits by first selling and then buying the stock (uncovered sell). In [10], proportional controllers are used. The proportional gain value selection is really a critical task and it is performed by means of simulations using past stock prices. The resulting control law is simple, but does not adapt to market changes and may provide bad worst-case performance as evidenced in [9]. In [72], as reported in section 2.7, the same authors extend their work considering a Proportional-Integral

(PI) controller to regulate the investment function, nevertheless the time-invariant nature of the controller might not adapt well to inversions of price trend.

In this chapter, a different approach for control design in feedback stock trading is proposed, based on the *Extremum Seeking* rationale [4]. Such an approach appears to be very suitable for the problem at hand, for the following reasons: (i) its aim is to maximize the output of a system whose dynamics is unknown, like the excess return; (ii) it is intrinsically model-free; (iii) it provides a time-varying feedback gain, so it may adapt to market time-varying conditions; (iv) unlike many other adaptive methods, it is theoretically guaranteed to converge to a local optimum [62]. At the end of the chapter, a real case study and extensive experiments on a significant number of stocks show that the proposed approach may largely outperform the standard *Buy & Hold* strategy as well as the feedback scheme in [10] with a time-invariant SLS controller. Backtest overfitting [5] is avoided tuning the parameters of the methods on in-sample price data and testing them on out-of sample price data from the same stocks but coming from a different period of time, preventing ambiguity with respect to generalization properties.

The organization of the chapter is as follows. In Section 4.2, the Extremum Seeking functioning is described. Section 4.3 illustrates how the Extremum Seeking rationale can be applied to the problem at hand, while Section 4.4 shows experimental achievements.

4.2 Extremum Seeking control

Although there are many effective techniques for model-based control design, not all real physical systems can be adequately controlled by such techniques [23]. First, for many systems it may not be possible to derive a model, or if available, the model may not be suitable for control. Second, there may be progressive changes to the system that modify the underlying dynamics, and it may be difficult to measure or model these effects. This is the case of financial systems where stock price dynamics vary due to socio-economic effects, where the most extreme cases are represented by stochastic jumps.

The field of *adaptive control* seeks to overcome these challenges by employing time-varying control laws that can adapt to changes in the dynamics of such systems. Extremum Seeking control is a particularly attractive adaptive feedback control methodology for complex systems, that does not rely on a model of the process under control, to achieve the maximum (or the minimum) of an arbitrary cost function of the measured output [23]. Furthermore, Extremum Seeking can

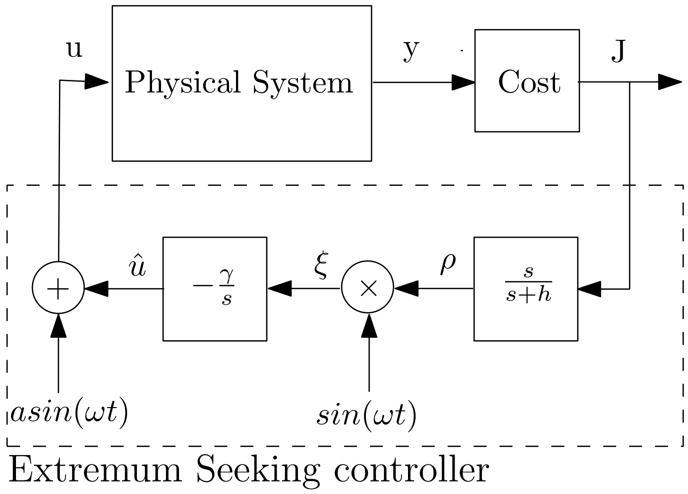

FIGURE 4.1: Schematic illustrian of an Extremum Seeking controller.

track maximum of an objective function despite external disturbances and varying system parameters.

The key idea behind the approach is that of estimating the gradient of the objective function of the output with respect to the input variables by perturbating *run-time* such variables, and thanks to the gradient direction information steer the actuation towards the value optimizing the objective.

The Extremum Seeking control architecture is depicted in figure4.1, the following is a description of his main components:

- the *dithering signal* $asin(\omega t)$, a small amplitude sinusoidal signal aimed at perturbing the control input. This is useful to find out the direction of improvement of the performance;
- the *unknown physical system*, the unknown system that generates the output signal y;
- the *cost function block*, which generates, as a function of the ouput, the objective signal J to optimize;
- the *washout filter*, described by the transfer function $s/(s+h)$, where h is the filter frequency, to reject the continuous component of the objective signal [16]. It outputs the zero-mean signal ρ;
- the *multiplicative node*, aimed to demodulate ρ multiplying it with the dithering signal and outputting the signal ξ;

- the *integrator* block γ/s, which integrate the demodulated signal producing the best guess $\hat{u}$ for the optimizing input u. For $\gamma > 0$ the objective is minimized, whereas it is maximized for $\gamma < 0$.

The algorithm works as follows. A sinusoidal perturbation is added to the estimate of the input $\hat{u}$ that maximizes the objective function:

$$u = \hat{u} + a\sin(\omega t) \tag{4.1}$$

The perturbed input passes through the dynamics of the system and through the block of the cost function that produces at the output the signal J that varies sinusoidally about some mean value. The objective function J is high-pass filtered to remove his mean (the continuos component) resulting in the oscillatory signal ρ. The high-pass filtered output is then multiplied by the dithering signal resulting in the demodulated signal ξ:

$$\xi = \sin(\omega t)\rho \tag{4.2}$$

The signal ξ is mostly positive if the input u is on the left of the optimal value, while is mostly negative if u is on the right of the optimal value. Finally the demodulated signal ξ is integrated into the best estimate of the optimizing value $\hat{u}$:

$$\frac{d}{dt}\hat{u} = -\gamma\xi \tag{4.3}$$

so that it is steered towards the optimal input. From the mathematical passages above it is not clear where the information on the gradient of the objective function appears. This is actually appreciable considering constant plant dynamics such that J is simply a function of the input:

$$J(u) = J(\hat{u} + a\sin(\omega t)) \tag{4.4}$$

Expanding $J(u)$ in the dithering signal, which is assumed to be small, yelds:

$$J(u) = J(\hat{u}) + \left.\frac{\partial J}{\partial u}\right|_{u=\hat{u}} \cdot a\sin(\omega t) + O(a^2) \tag{4.5}$$

The leading term in the high pass filtered signal is $\rho \approx + \left.\frac{\partial J}{\partial u}\right|_{u=\hat{u}} \cdot a\sin(\omega t)$. Averaging $\xi_{avg} = \sin(\omega t)\rho$ over one period yelds:

$$\xi_{avg} = \frac{\omega}{2\pi}\int_0^{\frac{\omega}{2\pi}} \sin(\omega t)\rho dt \tag{4.6}$$

$$= \frac{\omega}{2\pi}\int_0^{\frac{\omega}{2\pi}} \left.\frac{\partial J}{\partial u}\right|_{u=\hat{u}} a\sin^2(\omega t)dt \tag{4.7}$$

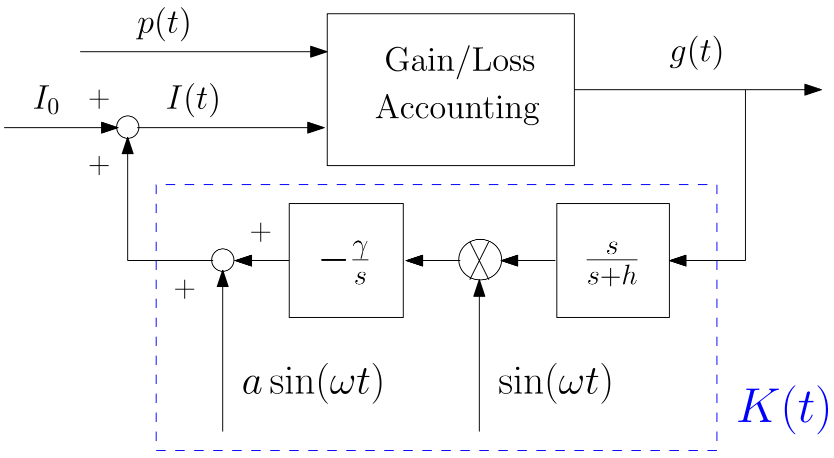

FIGURE 4.2: The Extremum Seeking control scheme for feedback stock trading.

$$= \frac{a}{2} \frac{\partial J}{\partial u}\bigg|_{u=\hat{u}} \int_0^{\frac{\omega}{2\pi}} a\sin^2(\omega t)dt \tag{4.8}$$

Thus, with trivial plant dynamics the signal ξ_{avg} is proportional to the gradient of the objective function J with respect to the input u.

The stability proof in [62] shows that the solution will be within a neighborhood of the maximum (or minimum) of the output function but it will never reach it, due to the persistent perturbation signal $asin(\omega t)$. To get an estimation as close as possible to the optimum value, it will be needed to leverage on the parameters involved in the algorithm.

4.3 Stock trading via Extremum Seeking control

The methodology is naturally suited for all systems, like the one considered here, in which the dynamics (of the price) is unknown or highly uncertain and vary over time. Moreover, it allows the design parameters, i.e., the controller gain K in this case, to be time-varying and self-adapted on-line.

Figure 4.2 depicts the new scheme for stock trading via feedback control that exploits the Extremum Seeking rationale. In the present setting, the unknown physical system is represented by the gain/loss accounting block, where the inputs are the investments $I(t)$ and the price $p(t)$, while the output to maximize is $g(t)$ itself. Notice that the resulting time-varying feedback gain $K(t)$ generated by composing the blocks in the blue box of Figure 4.2 must also here be composed by two elements: $K_S(t)$ for the short investment strategy and $K_L(t)$ for the long investment strategy. Therefore, the Extremum Seeking scheme is here applied to both the investment strategies and composed similarly to the SLS case of section 2.4 (the overall control gain will again be obtained as the sum of the two actions as $K(t) = K_L(t) + K_S(t)$). This new scheme proposed in this paper will be called Extremum Seeking Controller (ESC). Notice that a, ω, γ and h are tuning knobs and need to be determined. In particular, γ must have positive values when applied for the long controller and negative values for the short one.

To suitably select such parameters, [4] suggests some qualitative guidelines:

ω: perturbation frequency, it must be within the interval $[0, \pi]$ and it can be chosen depending on the closed loop system bandwidth;

a: perturbation amplitude, it must be small enough to obtain small changes in the output function but large enough to assure reliable measure of the gradient of g;

h: High pass filter, it must be designed such that $0 < h < 1$ and should be smaller than ω so that the filter removes the continuous component of the output without corrupting the estimation of the gradient.

γ: the gain of the Extremum Seeking scheme; large values will speed up the convergence rate as well as the possibility of saturation conditions. For this reason, we should select γ large enough to obtain satisfying results under both ideal conditions and on real data, but also small enough to avoid steady state situations, namely values of gain K constant over the time [4].

Simply put, the tuning of the above knobs could be summarized into a managing problem of the trade-off between convergence rate and accuracy of the optimizer. In other words, a rapid convergence will imply that the investment will spend a shorter amount of time very close to the optimal investment strategy; on the other hand, a better tracking of the optimal trading policy can be achieved at the cost of a slower rate of adaptation when the price dynamics change. In this work it was decided to set the aforementioned parameters only once, at the beginning of the experiment, however it would be interesting, and will be subject to further study, the study of the sensitivity of the performances by varying the tuning frequency of these parameters that can be set also run-time in an adaptive way.

It should be here also remarked that other dithers could be used instead of the sine wave as the perturbation signal. For instance, the square wave excitation is proved in [85] to provide the best convergence rate among all dithers of the same amplitude and frequency. Therefore, this signal will be used in the experimental examples of the next section.

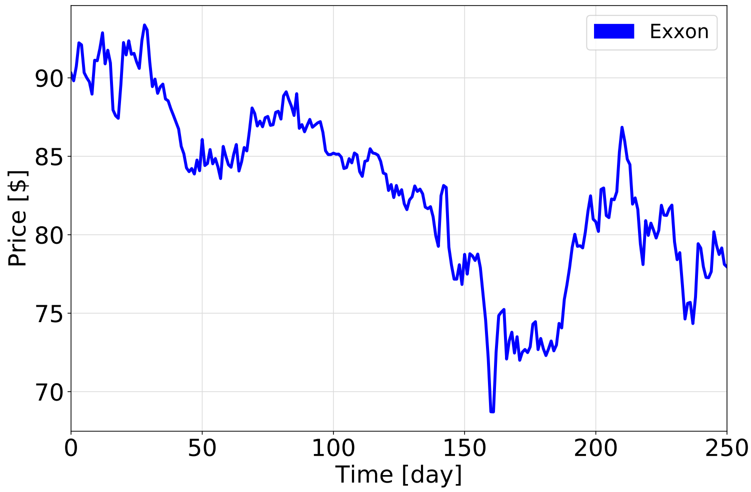

FIGURE 4.3: Daily closing prices for the Exxon (XOM, NYSE) stock from January 2015 until September 2015.

Finally, notice that the higher level of flexibility of such a time-varying control strategy is paid in terms of theoretical guarantees like Theorem 2.1, which no longer apply. Nonetheless, convergence analysis of extremum seeking schemes is indeed feasible (see again [85]) and will be object of future works. Moreover, as illustrated in the following example, there are already some cases in which the assumptions of the theory with time-invariant controllers are not satisfied, and in such cases the extremum seeking control appears as the most suited solution.

4.4 Experimental results

In this section, a numerical case study considering experimental data taken from a real stock is presented. Such an example has been selected within a round-trip period, where the traditional strategy is known to encounter some difficulties. After the case study, a more comprehensive simulation campaign, taking into account all DJIA's (Dow Jones Industrial Average) stocks between 2013 and 2015, is discussed, with the aim of providing a statistical assessment of the method.

4.4.1 The Exxon Mobil case study

Consider the Exxon Mobil Corporation (XOM, NYSE) stock from January 2015 until September 2015 in Figure 4.3. During this period, the equity presents different trends (both positive and negative), and, at the end of the period, the price is very close to the initial value.

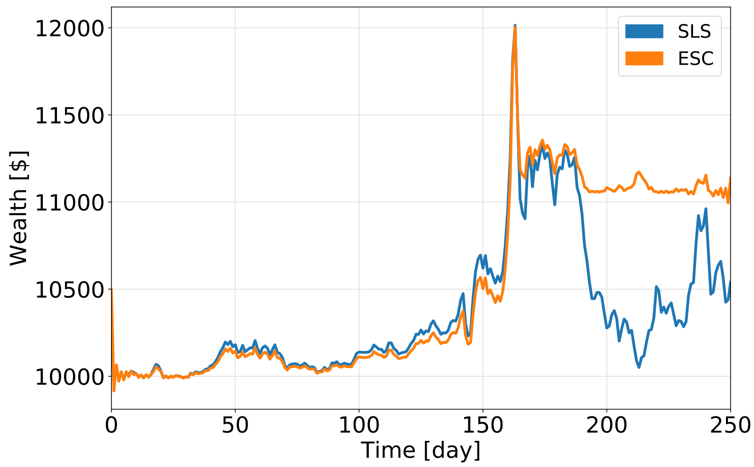

FIGURE 4.4: Wealth values for the Exxon (XOM, NYSE) example: SLS control and ESC control.

Two reactive control strategies are here applied: the SLS control of Section 2.4 and the approach proposed in this chapter. For a fair comparison, the same values for the initial wealth $w_0 = 10000\$$, the initial investment $I_0 = 5000\$$ and the maximum investment level $I_{max} = 20000\$$ are considered for both the techniques. Obviously, for the short-selling scheme, the initial and the limit investments must be considered with the negative sign. For the SLS control, $K_{SLS} = 4$ is selected via backtesting using as in-sample data the previous year of prices of the same stock. For the Extremum Seeking controller, the following parameters are instead chosen with grid-search, by means of the same backtesting strategy: $\alpha = 0.04$, $\omega = 0.4\pi\ rad/sample$, $\gamma = 4$ and $h = 0.8$. When the SLS controller is applied, the wealth value in Figure 4.4 is obtained. As expected, the controller increases the wealth in presence of the initial negative trend. However, at a time when the share price suddenly reverses its trend around the day $175th$, the SLS suffers a strong monetary loss in the days to follow. This is due to the fact that the SLS controller cannot rapidly change the short strategy into a long one. Ideally, the sign of the investment should change (from negative to positive) around the $175th$ day. The ESC, on the contrary, despite the sudden change in price trends, is able to limit monetary losses.

In Figure 4.5, the investment levels corresponding to the accounts of Figure 4.4 are illustrated (a split in long and short investment is instead shown in Figure **??**).

Such a plot confirms that the ESC approach addresses more promptly the change of the trend, basically stabilizing the wealth notwithstanding the unexpected behavior.

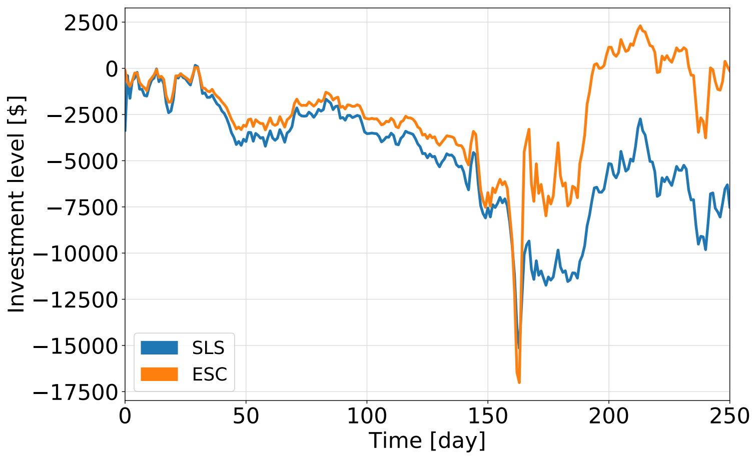

FIGURE 4.5: Investment levels for the Exxon (XOM, NYSE) example: SLS control and ESC control.

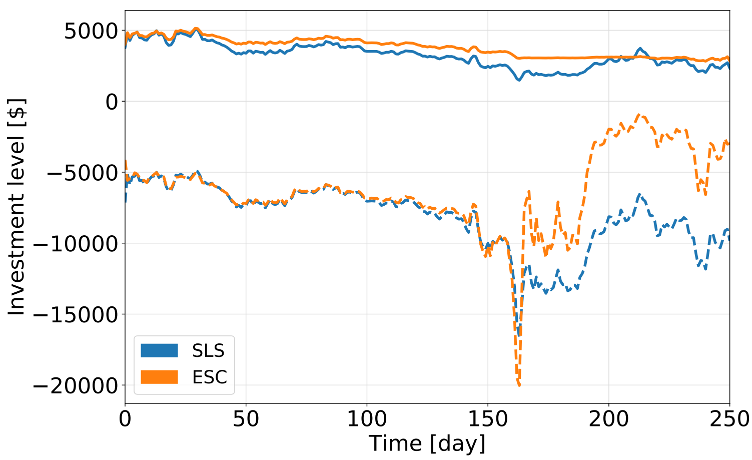

FIGURE 4.6: Split of long/short investment levels for the Exxon (XOM, NYSE) example.

4.4.2 The DJIA's stocks between 2013 and 2015

The case study in the previous section is definitely of interest to show that the proposed approach *might* work well in some situations, but indeed is not general. Since when K is time-varying, the theoretical properties proven in [10] do not hold even within idealized markets, in this subsection a statistical evaluation of the approach will be proposed over all the DJIA's stocks in the period 2013-2015.

To do that, ESC and SLS schemes have been designed for each stock following the same approach adopted for the Exxon (XOM, NYSE) example. The final results of the ESC scheme are reported in Table 4.1, in terms of average values of the final gain as compared to the SLS approach and a traditional Buy & Hold, but also considering the maximum and the minimum relative gain with respect to the other strategies.

Interestingly, the simulations show that not only the ESC strategy provides better results in terms of mean gain, but it is never (significantly) lower than the non-adaptive strategy [10]. At the same time, it may outperform it - in the lucky case - of about +50%. The performance improvement is even more significant as far as the traditional Buy & Hold approach is employed.

	B&H	SLS
Mean	+9.26%	+3.76%
Upper limit	+63.17%	+49.78%
Lower limit	-4.05%	-2.66%

TABLE 4.1: Account values obtained by the Extremum Seeking approach (ESC) in comparison to traditional SLS and Buy and Hold.

Chapter 5

A robust design strategy for stock trading via feedback control

5.1 Introduction

In this fifth chapter we describe the second innovative contribution of this dissertation to the line of research of stock trading via feedback control described in Chapter 2.

In the previous chapter we dealt with one of the main limitations we encountered in the original scheme of section 2.4, the SLS controller. In particular, we pointed out that a time-invariant controller could not be adequate to control a system whose dynamics are governed by market prices, which, by their nature are non-stationary stochastic processes and therefore subject to changes in their dynamics for socio-economic reasons. For this reason a time-invariant controller does not seem to be the ideal choice. Furthermore, in the theory of stock trading via feedback control there are no guidelines for tuning the controller values. We therefore decided to overcome this challenges by using an approach derived from adaptive control and based on the Extreme Seeking rationale. We have designed a new control scheme that we have called the Extremum Seeking Controller, characterized by a pair of time-varying controllers in which the control law varies reactively with respect to changes in price dynamics. Furthermore, we have provided guidelines for the calibration of the characteristic parameters of the controller. This adaptive approach proved to be effective in numerical experiments having exceeded in performance the two benchmarks, the original SLS scheme and the Buy & Hold strategy.

In this chapter we do not question the practical effectiveness of the ESC, however we note the fact that it does not provide guarantees of robustness and risk control features that typically are important qualities in the eyes of an investor. Furthermore we note a drawback that unites the two schemes mentioned above, neither the ESC nor the SLS in any way exploit information regarding prices and instead treat them as external disturbances completely unknown, while, on the other hand, exploiting it in a mild manner could be beneficial without undermining the model-free nature of the approaches.

In light of these fact, in this chapter, we will design a new investment strategy based on another control methodology, called *robust control*, able to provide guarantees of robustness and able to take advantage of the information available on price, and in particular on their returns. Indeed, we propose a reformulation of the reactive trading scheme described in Chapter 2, in which the return (i.e., the normalized price trend) is not treated as an unknown disturbance, but rather as an uncertain parameter within a limited range of values. The key observation motivating this analysis is that, especially in high frequency trading, the return can be well approximated as an *Uncertain But Bounded* (UBB) parameter, as we will also show via an extensive empirical study. It follows that, by assuming a very mild knowledge of the process dynamics (i.e., the return bounds), a robust controller can be designed, which not only provides good average performance, but also robustness guarantees. Specifically, to formulate the design problem as a H_∞ problem, a trend following scheme will be employed, where the desired gain is reasonably selected by considering the current situation.

The organization of this chapter is as follows. In Section 5.2, a reasonable assumption on price returns is described. . The proposed robust control design scheme is introduced and analyzed in Section 5.2.1. The effectiveness of the proposed strategy is illustrated in Section 5.3.

5.2 An additional (but reasonable) assumption

In this section, supported by the study of the financial literature, we discuss what reasonable assumptions can be made about price returns.

The behaviour of price returns of financial assest has been studied for years [76]. First studies date back to the early 1900s, in which the returns of the assets were assumed to be independent and identically distributed and modeled using normal distributions. From the early 1960s, the normality hypothesis began to be questioned, since empirical returns distributions showed fatter tails then the Gaussian model [54]. For these reasons, in the following years, some authors tried to find some new distributions that could better explain price returns, e.g., the stable Paretian distribution [86], the Student distribution or t distribution in [17], the exponential power distribution in [54], the Laplace distribution in [61, 68] and the logistic distribution in [54]. Nowadays, there is no formal agreement on the probabilistic distribution that best describes the return of financial assets [76]. This is also confirmed by the empirical study illustrated in Fig. 5.1 on 8 single stocks from the American market on a time frame that goes from 2010 to 2017. In this study, we used daily closing prices to compute price returns for a total of 14044 data points and it actually seems that the empirical probability distribution are not Gaussian showing fatter tails as stated in [54]. Therefore the assumption of gaussianity is to be excluded.

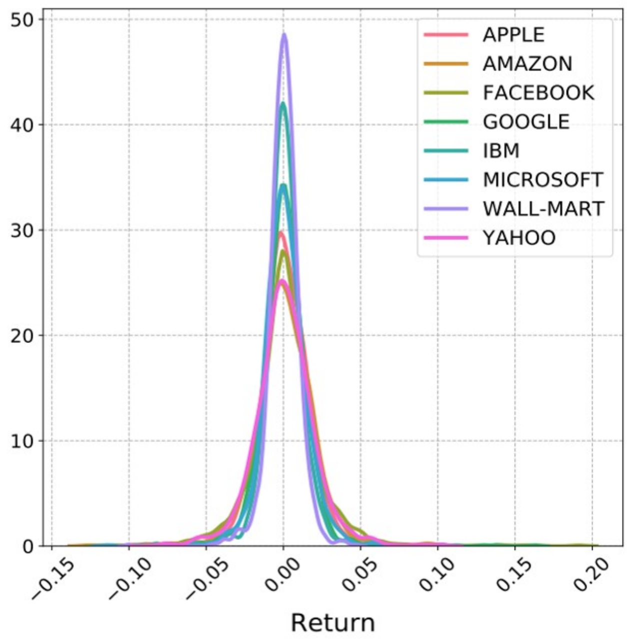

FIGURE 5.1: Empirical return distributions from 1-1-2010 to 31-12-2017 of 8 American famous assets.

The non-stationarity of prices, and in particular the heteroskedasticity of returns, is widely documented in the financial literature [64, 98, 84]. Heteroskedasticity (or heteroscedasticity) happens when the standard errors / variance of a variable, monitored over a period of time, are non-constant. A visual example of this property can be seen in figure 5.2, which shows the realzation of returns of the Facebook market price over the 2010-2017 period. It follows taht it is necessary exclude assumptions on stationary of price returns.

However, by looking at the boxplots in Fig. 5.3, we can show an important property: all the returns of the considered stocks are bounded to a narrow range of values within the interval of ±20%. Although conservative, the above number could be used as a limit for the return that, in turn, could represent a model uncertainty within a feedback scheme, as illustrated next.

In light of this analysis we consider as reasonable the assumption that

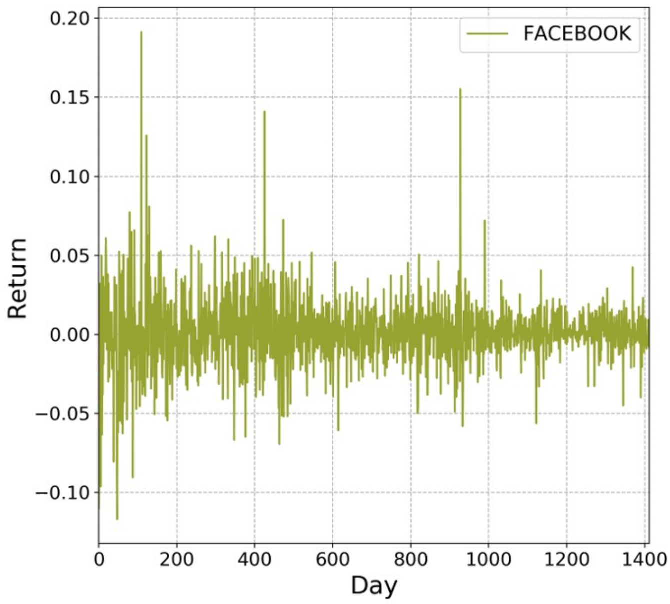

FIGURE 5.2: Facebook return realization from 1-1-2010 to 31-12-2017.

return $\rho(t)$ can be treated as an uncertain parameter belonging to the interval set $\mathcal{U}$ (*uncertain but bounded*):

$$\mathcal{U} = \{\rho(t) : \rho_{min} \leq \rho(t) \leq \rho_{max}\} \tag{5.1}$$

For deriving the control policy of this paper, we assume that a mild knowledge of the range of possible returns is available. This is not a strong assumption, especially in high frequency trading.

5.2.1 Robust controller design for optimal gain tracking

In this work, based on the uncertain but bounded assumption, the reactive trading problem is reformulated as a reference tracking problem, according to the scheme depicted in Fig. 5.4, where g_0 denotes the *desired gain*, whose selection will be discussed later.

The system P is the gain/loss function and can be described by the state space model:

$$\begin{cases} \dot{g}(t) = \rho(t)I(t) \\ y(t) = g(t) \\ e(t) = g_0(t) - g(t) \end{cases} \tag{5.2}$$

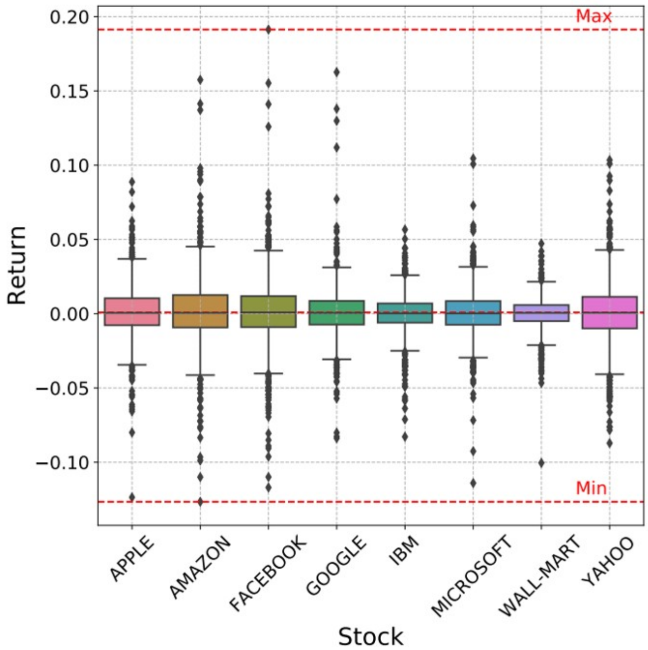

FIGURE 5.3: Empirical return distributions from 1-1-2010 to 31-12-2017 of 8 American famous assets.

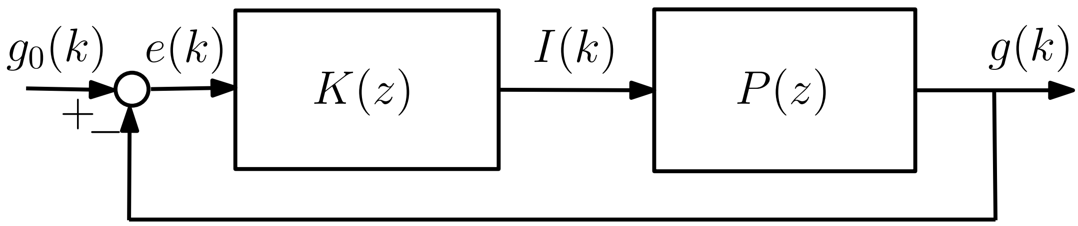

FIGURE 5.4: Reference tracking scheme.

More specifically, assuming to work with sampled data, we will work from now on with the more appropriate discrete time form:

$$\begin{cases} g(k+1) = g(k) + \rho(k)I(k) \\ y(k) = g(k) \\ e(k) = g_0(k) - g(k) \end{cases} \tag{5.3}$$

In (5.3), $\rho(k)$ is treated as a parametric uncertainty that affects the system. To synthesize the controller it was decided to use the H_∞ approach. Defining $T_{eg_0}(z, \theta)$ as the closed-loop transfer function between the reference gain $g_0(k)$ and the difference between the reference gain and the actual gain $e(k) = g_0(k) - g(k)$, the parametric robust structured H_∞ control problem consists in computing a structured controller $K(z, \theta^\star)$ with the following properties:

- *robust stability:* $K(z, \theta^\star)$ stabilizes the system for every possible variation of $\rho(k)$;
- *robust performance:* $K(z, \theta^\star)$ is the controller among all the others robust stabilizing controllers $K(z, \theta)$ with the same structure that satisfies:

$$\max \|T_{eg_0}(z, \theta^\star)\|_\infty \leq \max \|T_{eg_0}(z, \theta)\|_\infty$$

It should be noted that the H_∞ approach provides stability and robustness guarantees when the parameters are UBB and constant. If the parameters are time-varying, the H_∞ theoretical guarantees may not hold anymore (some additional assumption may be required, like for example sufficiently slow variation). This consideration also holds for robust performance. The efforts of future studies will be aimed at proving stability and robustness guarantees in a probabilistic way up to a predetermined confidence level assuming that price returns are governed by simple distributions such as the Gaussian distribution.

The choosen controller structure is that of a Proportional/Derivative (PD) controller. The integral action for zero steady-state error is not needed as the system already contains an integrator. More specifically, two branches for long and short trading with two PDs are implemented, according to the SLS rationale, the resulting sheme is is called Robust SLS (RSLS) and is depicted in Fig. 5.5. To realize this scheme, it is assumed that the uncertain $\rho(k)$ parameter can only take positive values for the long branch, while it can only vary among negative values for the short branch. The optimization problem was solved using the Matlab Robust Control Toolbox. The *systune* function, which implements the nonsmooth optimization algorithms described in [3] was used to find the otpimal parameters of the controller. The real uncertain parameter $\rho(k)$ was modeled using the function *ureal* from the same library.

The reference generation block is a tuning knob of the approach. In this work, as a criterion for the design of the reference signal g_0, a *trend following* strategy is implemented. The resulting scheme is depicted in Fig. 5.6 and it works as follow, if the price is increasing, the long gain will increase forcing the scheme to increase the long position, on the contrary the short gain will decrease, forcing the scheme to decrease the short position. If prices decrease, the situation is reversed. To this purpose, our choice is the inclusion of a second feedback of the same signal, the gain/loss function, obtained via multiplication with a new parameter α aimed to decoupling the inner and outer feedback and to pilot the scheme towards the right direction, without following too strictly the price dynamics. In this way $g_{0,l}(k) = \alpha g_l(k)$ for the long position and $g_{0,s}(k) = \alpha g_s(k)$ for the short one. The absolute value of such a parameter dictates the agressiveness of the investment.

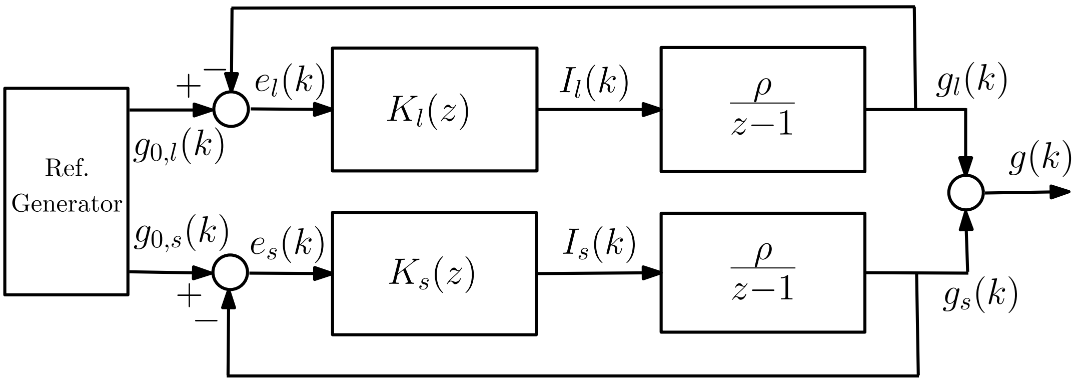

FIGURE 5.5: Robust SLS (RSLS) scheme.

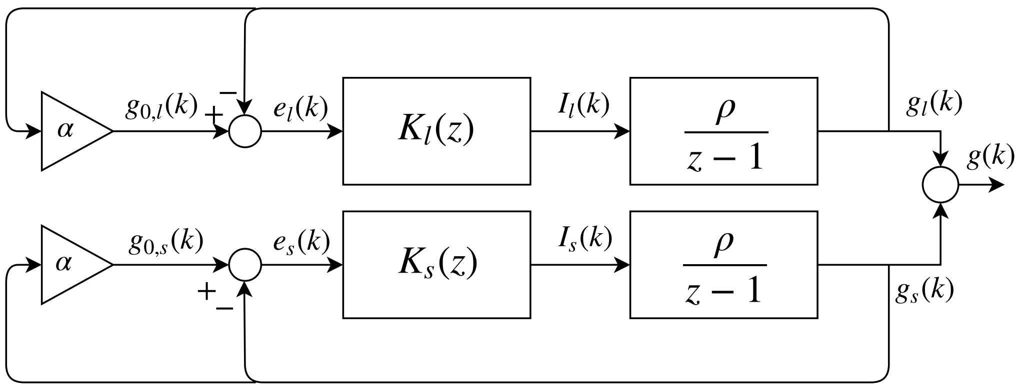

FIGURE 5.6: Robust SLS (RSLS) scheme with trend following.

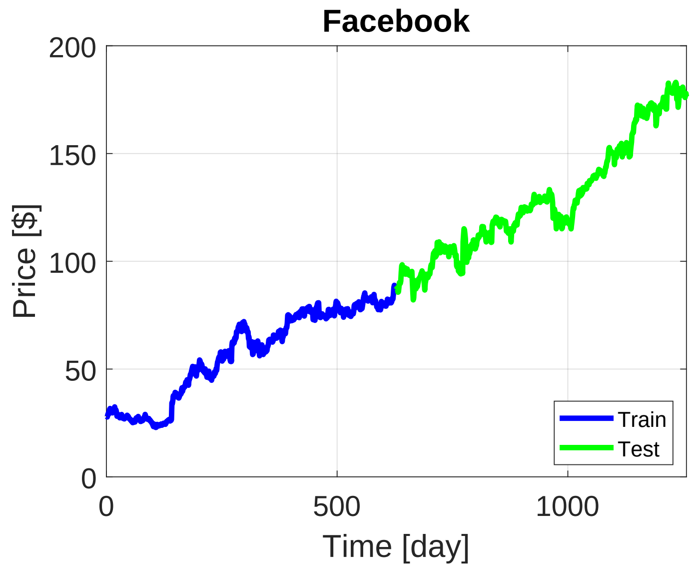

FIGURE 5.7: Facebook stock price from 1-1-2013 to 31-12-2017.

5.3 Experimental results

To show the effectiveness of RSLS control against standard SLS, real-world daily trading of Facebook and Google stocks are considered. Specifically, past data are used covering a period of 5 years, from the beginning of 2013 to the end of 2017. The data were then separated into two parts: from January 2013 to June 2015, data are used for training (blue lines in Fig. 5.7 for Facebook and Fig. 5.8 for Google), i.e. to select the optimal values of the hyper-parameters of the two algorithms, the feedback gain K of the SLS scheme and the α parameter of the Robust SLS scheme. The remaining two and a half years (green lines in Fig. 5.7 for Facebook and Fig. 5.8 for Google) are instead employed as a test set for evaluating the performances of the algorithms.

The idea behind these experimental tests is to calibrate the parameters of the two schemes, the SLS and the RSLS, using the train data and then using this parameterization to compare the performances of the two schemes both on the train data and on the test data. The objective is to verify whether the new scheme actually proves to be more robust to any changes in price dynamics than the original scheme.

The hyper-parameters of the two models, in particular the feedback gain K of the SLS scheme and the α parameter of the Robust SLS scheme were selected through a sensitivity analysis from a grid of possible values, minimizing the cost function suggested in [87], namely the trading MSE using the train data to perform the optimization. For Facebook $K = 11$ and $\alpha = 1.25$ were selected. For Google $K = 8$ and

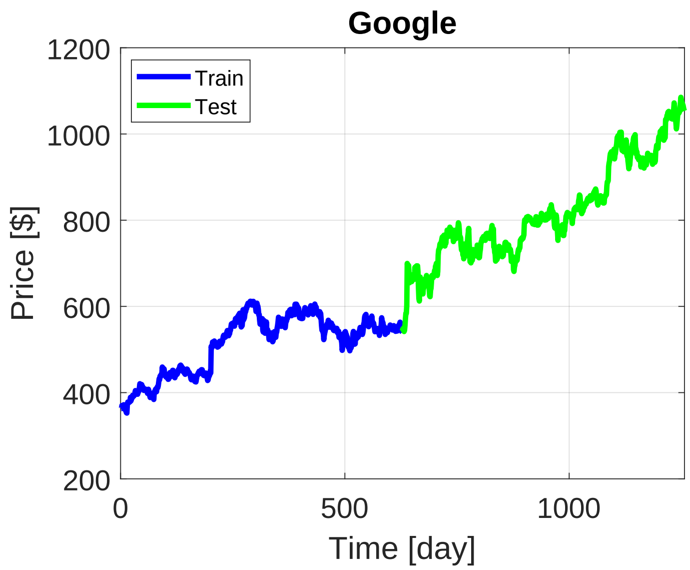

FIGURE 5.8: Google stock price from 1-1-2013 to 31-12-2017.

$\alpha = 1.1$ were selected.

The limits of the range of variation of the uncertain ρ parameter were chosen conservatively in such a way as to emphasize the characteristics of robustness and the new control scheme. For this reason a range of $\rho(k) \in [-0.2, 0.2]$ has been selected that is most likely able to encompass the great majority of possible variations of the return of the to stocks.

The resulting robust controller obtained thanks to the H_∞ synthesis procedure is given by the transfer function:

$$K(z) = K_P + K_D \frac{1}{Tf + T_s/(z-1)} \tag{5.4}$$

with $K_P = 284$, $K_D = -4.13e+04$, $T_f = 150$, $T_s = 1$ for the long controller and $K_P = -284$, $K_D = 4.13e+04$, $T_f = 150$ and $T_s = 1$ for the short controller.

In Fig. 5.9, it can be seen that the SLS scheme outperforms RSLS on the training data. However, if the same controller is used on the test data, see Fig. 5.10, the situation is reversed, as the stochastic characteristics of the price have changed and the gains of the SLS are no longer optimal.

The same procedure is carried out with Google data. By testing the algorithms on train data, it is observed from Fig. 5.11 that the performance of the two schemes is comparable in the first 500 days while in the final trading days the performance of the RSLS is less affected by the lack of price trends than those of the SLS. And again, observing the results of the test in the Fig. 5.12, it is clear that from 350-th day onwards the robust scheme is able to exploit in a more effective way the rising trend

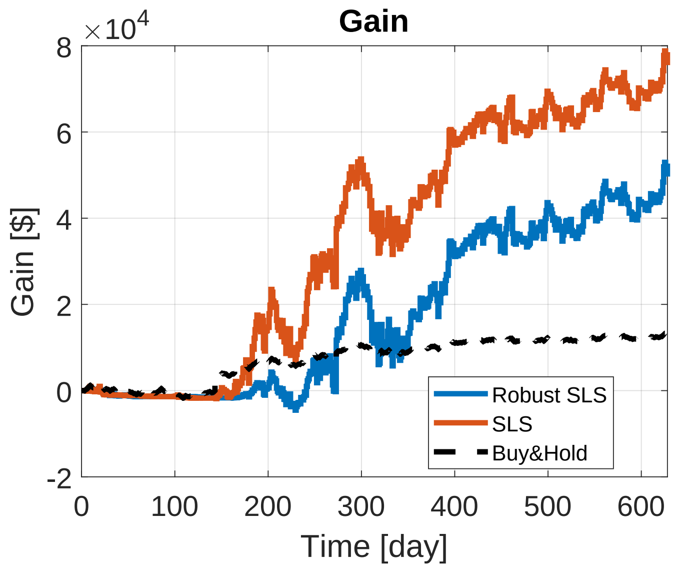

FIGURE 5.9: Comparison between gain-loss functions for the SLS, Robust SLS and Buy & Hold on the training part of Facebook data.

and consequently at the end of the training period the RSLS largely outperform the SLS.

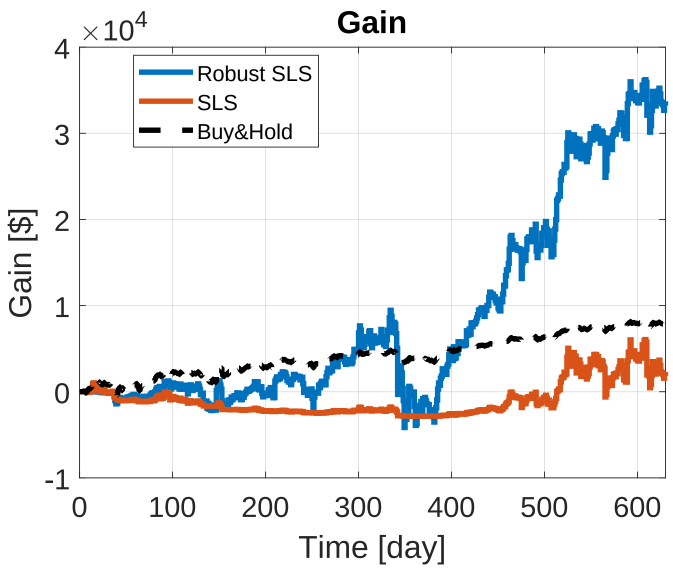

FIGURE 5.10: Comparison between gain-loss functions for the SLS Robust SLS and Buy & Hold on the test part of Facebook data.

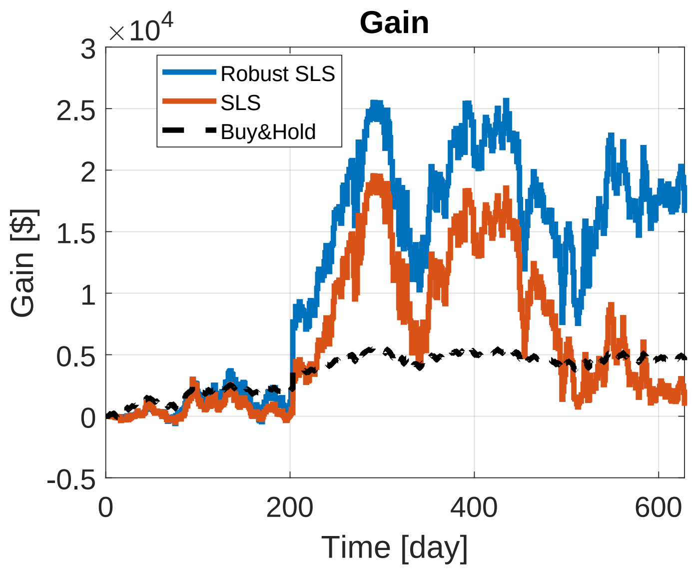

FIGURE 5.11: Comparison between gain-loss functions for the SLS Robust SLS and Buy & Hold on the training part of Google data.

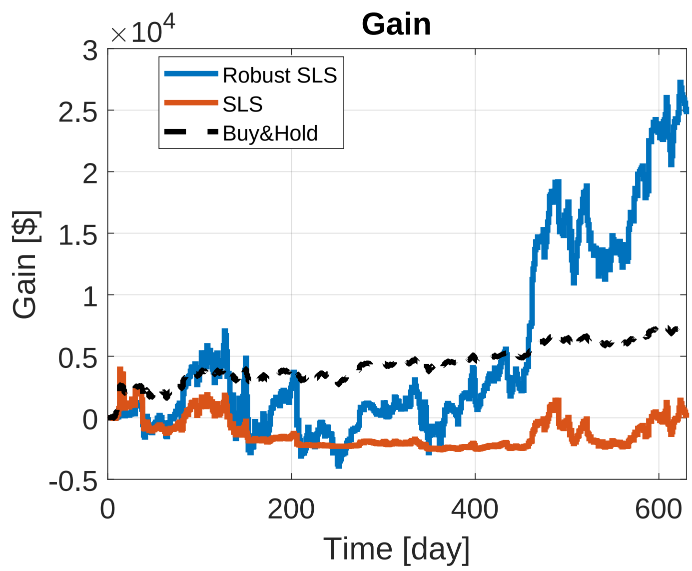

FIGURE 5.12: Comparison between gain-loss functions for the SLS Robust SLS and Buy & Hold on the test part of Google data.

Chapter 6

Multi-period asset allocation with kernel-based control policy

6.1 Introduction

This chapter presents an innovative contribution of this dissertation to the research field of multi-period portfolio optimization introduced in section 3.3.

In chapter 3 we addressed the problem of optimal allocation of an investor's monetary resources in a portfolio of financial assets. In section 3.2, we discussed the classic single-period portfolio optimization framework dating back to Markowitz and subsequent approaches based on robust optimization and asymmetric risk measures aimed at improving the original model.

In section 3.3 we introduced an important extension of the single-period framework based on a multi-period decision-making approach aimed at satisfying long-term investment objectives of an investor. In multi-period asset allocation, optimization is used to plan a sequence of trades to carry out over a set of future periods.

We first mentioned the classical literature on multi-period portfolio optimization based on dynamic programming and dating back to the work of Merton [77] and Samuelson [96]. However, using dynamic programming for asset allocation has proved impractical, except for short investment horizons and limited numbers of assets, due to the "curse of dimensionality" [21]. We have therefore introduced recent and promising convex formulations based on Model Predictive Control techniques which, although they provide sub-optimal control policies, have the advantage of being able to naturally incorporate constraints on portfolio composition and can also be resolved in useful times for practical implementations. Of particular interest for this dissertation are the approaches presented in sections 3.3.2 and 3.3.3, able to exploit the dynamic nature of the multi-stage problem by introducing closed-loop policies which are parametric functions of past market return realizations. To preserve the convexity of the optimization problem, affine parametrized policies are adopted.

In this chapter we generalize the latter approach by employing techniques of statistical learning theory that allow extending the policies's class to nonlinear functions while preserving the convexity of the optimization problem and therefore ensuring reasonable computation times for computing the optimal allocations.

6.2 Problem statement

In this section, first of all, we recall the main variables involved in the multi-period allocation problem discussed in section 3.3. Assuming n tradable assets, we denote by $x(k) \in \mathbb{R}^n$ the portfolio vector whose elements represent the portion of total wealth, measured in Euros/Dollars, allocated in the individual assets that make up the investor's portfolio. With $w(k) = \sum_{i=1}^{n} x_i(k)$ we indicate the investor's total wealth at time k.

We denote by $g(k) \in \mathbb{R}^n$ the vector of assets gains over the k-th period and with $G(k) \in \mathbb{R}^{n \times n}$ the diagonal matrix having the elements of $g(k)$ in the diagonal. We assume that the future investment horizon is made up of $T \geq 1$ decision-making stages during which the investor has the opportunity to vary the allocation of his resources between the traded assets. Finally, we indicate with $u(k) \in \mathbb{R}^n$ the vector of portfolio adjustements whose elements quantify the monetary amount invested or disinvested in the respective assets.

The recursive equation modeling the portfolio dynamics is given by:

$$x(k+1) = G(k+1)[x(k) + u(k)], \quad k = 0, \dots, T-1 \tag{6.1}$$

Following the scenario approach described in section 3.3.3, we assume that we have a mechanism for generating N iid scenarios of the future gains $\{G^{(j)}(k),\ k = 1, \dots, T\},\ j = 1, \dots, N$. We will use these sampled scenarios to minimize some portfolio's empirical performance measure $J(u(k))$ associated with the investment strategy. For example, in section 3.3.3 we measured the portfolio's performance by the empirical shortfall at the final stage:

$$J(u(k)) = \frac{1}{N} \sum_{j=1}^{N} \max(0, \gamma - w^{(j)}(T))^{\nu} \tag{6.2}$$

where γ is the desired target wealth at the end of the investment horizon and ν is typically set to $\nu = 1$ or $\nu = 2$ based on the risk aversion attitude of the investor. With this setting the multi-period portfolio optimization problem can be cast as the following dynamic optimization problem:

$$\min_{u(k)\in\mathcal{U}} \frac{1}{N}\sum_{j=1}^{N} \max(0, \gamma - w^{(j)}(T))^{\nu} \tag{6.3a}$$

$$\text{s.t.} \quad \mathbf{1}^T u(k) = 0, \qquad k = 0, \ldots, T-1 \tag{6.3b}$$

$$x^{(j)}(k) \in \mathcal{X}(k) \qquad k = 0, \ldots, T-1,\ j = 1, \ldots, N \tag{6.3c}$$

Constraint 6.3b models the self-financing condition, while constraint 6.3c ensures that, for each scenario j and for each decision stage k, the portfolio vector $x^{(j)}(k)$ respects all the constraints on its composition contained in the set $\mathcal{X}(k)$, including in the first place the dynamic equation 6.1. Finally, the portfolio adjustments $u(k)$ are elements of the set $\mathcal{U}$. In the simplest case $u(k)$ is a vector of unconditional adjustments and $\mathcal{U}$ corresponds to the $\mathbb{R}^n$ space, in a more complex case $u(k)$ are control functions, called policies, and the set $\mathcal{U}$ is a generic set of functions.

Note that the cost function 6.3a is not differentiable due to the presence of the $\max(\cdot,\cdot)$ function. To address this issue it is possible to add slack variables and reformulate problem 6.3 in the following way:

$$\min_{u(k)\in\mathcal{U}} \frac{1}{N}\sum_{j=1}^{N} z_j^{\nu} \tag{6.4a}$$

$$\text{s.t.} \quad \mathbf{1}^T u(k) = 0, \qquad k = 0, \ldots, T-1 \tag{6.4b}$$

$$x^{(j)}(k) \in \mathcal{X}(k), \qquad k = 0, \ldots, T-1,\ j = 1, \ldots, N \tag{6.4c}$$

$$z_j \geq 0, \qquad j = 1, \ldots, N \tag{6.4d}$$

$$z_j \geq \gamma - w^{(j)}(T) \qquad j = 1, \ldots, N \tag{6.4e}$$

As discussed in section 3.3.2, we want to generate investment policies that exploit the dynamic nature of the problem by progressively using the new information that becomes available and represented by the observation of returns, or equivalently the market gains. Assuming that, at time $k \geq 1$, we have at our disposal the sequence of market realizations $g(1), \ldots, g(k)$ we look for policies that are functions of the observed gains, and therefore having the following structure:

$$u(k) = f_k(g(1), \ldots, g(k)) \tag{6.5}$$

With this setting, optimization problem 6.4 consists in looking for the optimal function in the potentially infinite dimensional functional space $\mathcal{U}$ and therefore finding a solution is generally unfeasible. For this reason in [29, 30] the author adopts a sub-optimal strategy based on a finite-dimensional affine parameterization of the policies, allowing for a tradeoff between policy complexity and numerical tractability of the optimization procedure. In particular, in [29], the autors shows that appropriately choosing cost function and portfolio composition constraints and adopting an affine policy with a single-stage memory depth of the form:

$$u(k) = \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k)) \tag{6.6}$$

where $\bar{g}(k)$ are the expected gains, $\bar{u}(k) \in \mathbb{R}^n$ and $\Theta(k) \in \mathbb{R}^{n \times n}$ are the decision variables, optimization problem 6.4 is convex and therefore solvable in an efficient way.

In [30] the same author generalize this result allowing for control policies with full memory of the past return history:

$$u(k) = \bar{u}(k) + \sum_{\tau=1}^{k} \Theta_\tau(k)(g(\tau) - \bar{g}(\tau)) \tag{6.7}$$

showing that the full memory extension does not compromise the convexity of the optimization problem.

6.3 Generalization to non-linear control policy

It is reasonable to ask whether a linear parameterization of the control law is sufficient to capture any complex market dynamics, therefore the objective of this section is to extend the space of functions among which to look for the optimal policy to the space of more generic non-linear functions.

We will study two different approaches, the first aims at obtaining non-linear policies through the pre-processing of policy inputs [1], the second is based on kernel methods [97].

6.3.1 Nonlinearity by preprocessing

We start assuming, for each asset i, the following linear control policy with memory depth of τ:

$$u_i(k) = \bar{u}_i(k) + \beta_i^T(k)\mathrm{g}(k) \tag{6.8}$$

where $\mathrm{g}(k) \in \mathbb{R}^{(n \cdot \tau)}$ is called *regressors vector* and has the following structure:

$$\mathrm{g}(k) = [g_1(k) \ \cdots \ g_1(k - \tau + 1) \ \cdots \ g_n(k) \ \cdots \ g_n(k - \tau + 1)]^T \tag{6.9}$$

while $\beta_i(k) = [\beta_{i,1}(k)\, \beta_{i,2}(k) \, \cdots \, \beta_{i,(n\cdot\tau)}(k)]^T$ is a vector of weights associated with the regressors.

A first possible approach to move to nonlinear control actions is to preprocess policy's inputs (in our case the regressors vector $\mathbf{g}(k)$) into an high-dimensional inner product space $\mathcal{F}$ called *feature space* as described in [1]. Let's consider for now, for simplicity, $\mathcal{F} = \mathbb{R}^d$. This mapping from the original space to the feature space is performed by means of a *feature map* ϕ:

$$\phi : \mathbb{R}^{(n\cdot\tau)} \longrightarrow \mathbb{R}^d, \tag{6.10}$$

$$\mathbf{g}(k) \mapsto \phi(\mathbf{g}(k)) \tag{6.11}$$

The feature map ϕ can be seen as a collection of arbitrary functions that are applied to the inputs forming the *feature vector*:

$$\phi(\mathbf{g}(k)) = [\varphi_1(\mathbf{g}(k)), \ldots, \varphi_d(\mathbf{g}(k))]^T \tag{6.12}$$

where:

$$\varphi_j : \mathbb{R}^{(n\cdot\tau)} \longrightarrow \mathbb{R}, \qquad j = 1, \ldots, d \tag{6.13}$$

$$\mathbf{g}(k) \mapsto \varphi_j(\mathbf{g}(k)), \qquad j = 1, \ldots, d \tag{6.14}$$

The structure of the control policy, with the mapping applied to the inputs, becomes:

$$u_i(k) = \bar{u}_i(k) + \beta_i^T(k)\phi(\mathbf{g}(k)) = \bar{u}_i(k) + \sum_{j=1}^{d} \beta_{i,j}^T(k)\varphi_j(\mathbf{g}(k)) \tag{6.15}$$

where the weight vector $\beta_i(k)$ is now an element of $\mathbb{R}^d$.

The resolution of the optimization problem 6.4 with policies given by 6.15 will provide a linear control function in the feature space $\mathbb{R}^d$, but since the chosen feature map ϕ is generally non-linear, a linear control function in the space $\mathbb{R}^d$ corresponds to a nonlinear policy in the original space [100].

The key to the success of this approach depends on a good choice of the feature map, i.e. the function ϕ in making the image of its inputs to be more informative in the feature space. Choosing such a mapping generally requires expert knowledge of the given task. Possible insights for the choice of the feature map can be taken from the financial literature, for example there are studies ([99, 2]) that show empirically that there is a non-negligible autocorrelation between squared value of the market returns, therefore choosing a feature map that performs a quadratic transformation could potentially increase the exploitable information contained in the original data. Often, however, since it is not possible to know in advance which are the features that improves the investment performance, generic mappings are used, such as polynomial type mappings.

In general it is possible to select a feature map ϕ that maps the original space to some Hilbert space, so that:

$$u_i(k) = \bar{u}_i(k) + \langle \beta_i(k), \phi(\mathbf{g}(k)) \rangle \tag{6.16}$$

where $\langle \cdot, \cdot \rangle$ is the inner product of the Hilbert space.

However, generic feature maps such that the range of ϕ is a high dimensional space can generate too complex feature vectors and the learning could be prone to overfitting and thus bad generalization properties. This issue can be tackled adding a complexity control term, a penalty on the norm of the policy weight vector $\|\beta_i(k)\|^2$, which lead to the regularized cost functional:

$$J_{reg}(u(k)) = J(u(k)) + \lambda \sum_{i=1}^{n} \|\beta_i(k)\|^2 \tag{6.17}$$

where $\lambda \geq 0$ is a parameter which controls the importance of the regularization term.

By studying the literature about Model Predictive Control it is possible to find other suggestions of regularization penalties to be imposed directly on the policy. In particular, in [90], the authors suggest adding the following regularization term:

$$\|u_i(0) - u_i(T-1)\|^2 + \sum_{k=1}^{T-1} \|u_i(k-1) - u_i(k)\|^2 \tag{6.18}$$

where the first term of 6.18 penalizes large deviations between the initial control action and the control action at the last step of the horizon while the second term penalizes large differences between control actions of adjacent steps.

Using the regularized cost functional in equation 6.17, multi-period portfolio optimization problem 6.4 becomes:

$$\min_{u(k)\in\mathcal{U}} \frac{1}{N}\sum_{j=1}^{N} z_j^{\nu} + \lambda \sum_{i=1}^{n} \|\beta_i(k)\|^2 \tag{6.19a}$$

$$\text{s.t.} \quad \mathbf{1}^T u(k) = 0, \qquad k = 0,\dots,T-1 \tag{6.19b}$$

$$x^{(j)}(k) \in \mathcal{X}(k) \qquad k = 0,\dots,T-1,\ j = 1,\dots,N \tag{6.19c}$$

$$z_j \geq 0, \qquad j = 1,\dots,N \tag{6.19d}$$

$$z_j \geq \gamma - w^{(j)}(T), \qquad j = 1,\dots,N \tag{6.19e}$$

where:

$$\mathcal{U} = \{u(k) : u_i(k) = \bar{u_i}(k) + \langle \beta_i(k),\ \phi(\mathrm{g}(k))\rangle, \quad i = 1,\dots,n\} \tag{6.20}$$

However the approach discussed in this section involves another kind of problem. Using generic nonlinear feature maps like the polynomial one can easily make the optimization problem computationally unfeasible if the order of polynomial features is and the size of the inputs are high, as the number of different monomial features of degree d is $\binom{d+(n\cdot\tau)-1}{d}$ [103]. This can therefore compromise the applicability of the described approach.

6.3.2 Implicit mapping via kernels

A second approach that represents a possible remedy to the computational problem described in the previous paragraph involves the use kernel methods.

Assume we are at time k and to have at our disposal the following N matrix of regressors:

$$\begin{bmatrix} [g_1(\tau)\ \cdots\ g_1(1)\ \cdots\ g_n(\tau)\ \cdots\ g_n(1)]^T \\ \vdots \\ [g_1(k)\ \cdots\ g_1(k-\tau+1)\ \cdots\ g_n(k)\ \cdots\ g_n(k-\tau+1)]^T \\ \vdots \\ [g_1^{(j)}(T)\ \cdots\ g_1^{(j)}(T-\tau+1)\ \cdots\ g_n^{(j)}(T)\ \cdots\ g_n^{(j)}(T-\tau+1)]^T \end{bmatrix} = \begin{bmatrix} \mathrm{g}(\tau) \\ \vdots \\ \mathrm{g}(k) \\ \vdots \\ \mathrm{g}(T)^{(j)} \end{bmatrix} \quad j = 1,\dots,N \tag{6.21}$$

where the first $k-\tau+1$ rows of matrix 6.21 contain observed market realizations while the rest contain scenarios of future market realizations.

Given the feature map ϕ we define the kernel function as follows:

$$K : \mathbb{R}^{(n\cdot\tau)} \times \mathbb{R}^{(n\cdot\tau)} \longrightarrow \mathbb{R}, \tag{6.22}$$

$$(\mathrm{g}(s), \mathrm{g}(t)) \mapsto K(\mathrm{g}(s), \mathrm{g}(t)) \tag{6.23}$$

where

$$K(\mathrm{g}(s), \mathrm{g}(t)) = \langle \phi(\mathrm{g}(s)), \phi(\mathrm{g}(t)) \rangle \tag{6.24}$$

and $\tau \leq s, t \leq T$.

One can interpret K as a measure that specifies a sort of correlation between regressors vectors at different time instant s, t. The main advantage of this approach is that it is neither necessary to specify the transformed inputs nor to specify the feature map [97].

As proved in [106], we can apply the following theorem:

Theorem 6.1 (Representer theorem, [97]). *Assuming ϕ is a mapping to an Hilbert space, then, there exist a vector $\alpha_i(k)$ such that $\beta_i(k) = \sum_{t=\tau}^{k} \alpha_{i,t}(k)\phi_i(\mathrm{g}(t))$ is an optimal solution of problem 6.19.*

On the basis of the representer theorem we can optimize problem 6.19 with respect to the coefficients $\alpha_i(k)$ instead of the coefficients $\beta_i(k)$ as follows. Given $\beta_i(k) = \sum_{t=1}^{T} \alpha_{i,t}(k)\phi_i(g(t))$ we have that for all k:

$$\langle \beta_i(k), \phi_i(\mathrm{g}(k)) \rangle = \left\langle \sum_{t=\tau}^{k} \alpha_{i,t}(k)\phi_i(\mathrm{g}(t)), \phi_i(\mathrm{g}(k)) \right\rangle \tag{6.25}$$

$$= \sum_{t=\tau}^{k} \alpha_{i,t}(k) \langle \phi_i(\mathrm{g}(t)), \phi_i(\mathrm{g}(k)) \rangle \tag{6.26}$$

$$= \sum_{t=\tau}^{k} \alpha_{i,t}(k) K(\mathrm{g}(t), \mathrm{g}(k)) \tag{6.27}$$

Similarly, the regularization term can be rearranged as:

$$\|\beta_i(k)\|^2 = \left\langle \sum_{t=\tau}^{k} \alpha_{i,t}(k)\phi_i(\mathrm{g}(t)), \sum_{t=1}^{k} \alpha_{i,t}(k)\phi_i(\mathrm{g}(t)) \right\rangle \tag{6.28}$$

$$= \sum_{t,k=\tau}^{k} \alpha_{i,t}(k)\alpha_{i,k}(k) \langle \phi_i(\mathrm{g}(t)), \phi_i(\mathrm{g}(k)) \rangle \tag{6.29}$$

$$= \sum_{t,k=\tau}^{k} \alpha_{i,t}(k)\alpha_{i,k}(k) K(\mathrm{g}(t), \mathrm{g}(k)) \tag{6.30}$$

Instead of solving problem 6.18 we can solve the equivalent problem:

$$\min_{u(k)\in\mathcal{U}} \frac{1}{N} \sum_{j=1}^{N} z_j^{\nu} + \lambda \sum_{i=1}^{n} \sum_{t,k=\tau}^{k} \alpha_{i,t}(k)\alpha_{i,k}(k) K(\mathrm{g}(t), \mathrm{g}(k)) \tag{6.31a}$$

$$\text{s.t.} \quad \mathbf{1}^T u(k) = 0, \qquad k = 0, \dots, T-1 \tag{6.31b}$$

$$x^{(j)}(k) \in \mathcal{X}(k), \qquad k = 0, \dots, T-1,\ j = 1, \dots, N \tag{6.31c}$$

$$z_j \geq 0, \qquad j = 1, \ldots, N \tag{6.31d}$$

$$z_j \geq \gamma - w^{(j)}(T) \qquad j = 1, \ldots, N \tag{6.31e}$$

where:

$$\mathcal{U} = \{u(k) : u_i(k) = \bar{u}_i(k) + \sum_{t=\tau}^{k} \alpha_{i,t}(k) K(\mathbf{g}(t), \mathbf{g}(k)) \ i = 1, \ldots, n\} \tag{6.32}$$

To solve the optimization problem 6.31 we do not need direct access to elements in the feature space, we only need to know how to compute the kernel function. In fact, to solve optimization problem 6.31, we only need the elements of the Gram matrix $\mathbf{K}$, which is a positive semidefinite symmetric matrix, such that:

$$\mathbf{K}_{s,t} = K(\mathbf{g}(s), \mathbf{g}(t)) \tag{6.33}$$

Given the Gram matrix $\mathbf{K}$ we could rewrite the regularization term as:

$$\lambda \sum_{i=1}^{n} \sum_{t,s=\tau}^{k} \alpha_{i,t}(k) \alpha_{i,s}(k) K(\mathbf{g}(s), \mathbf{g}(t)) = \lambda \sum_{i=1}^{n} \alpha_i^T \mathbf{K} \alpha_i \tag{6.34}$$

The advantage of using kernels rather than looking for the optimal value of $\beta_i(k)$ in the feature space is that often the feature space is very large while implementing a kernel function is a simpler task [100]. Some examples of popular kernels are given in the following.

- **Polynomial kernel**. The d-degree Polynomial kernel is defined as:

$$K(\mathbf{g}(s), \mathbf{g}(t)) = (\mathbf{g}(s)^T \mathbf{g}(t) + c)^d \tag{6.35}$$

  where $c \geq 0$ is a free parameter trading off the influence of higher-order versus lower-order terms in the polynomial.

- **Gaussian Kernel**. The Gaussian kernel is defined as:

$$K(\mathbf{g}(s), \mathbf{g}(t)) = e^{-\frac{\|\mathbf{g}(s) - \mathbf{g}(t)\|^2}{2\sigma}} \tag{6.36}$$

  where $\sigma > 0$ is scale parameter. It can be shown [100] that the Gaussian kernal corresponds to an inner product in an infinite dimensional feature space. Therefore solving optimization problem 6.29 corresponds to solving the optimization problem 6.18 with infinite dimensional feature vectors.

6.4 Numerical experiments

In this section we present a numerical test based on real financial data. We considered a multi-period allocation problem involving $n = 7$ assets and a planning horizon of $T = 12$ periods, each period having the duration of one trading month (about 21 days per month). Historical data of monthly asset returns were downloaded from the Yahoo Finance website covering a period from 5/2007 to 5/2019 (May 2019 included).

The assets taken into consideration for the composition of the portfolio are the same used in [27], and are the following:

1. SPDR Dow Jones Industrial Average ETF, ticker DIA;
2. iShares Transportation Average ETF, ticker IYT;
3. iShares U.S. Utilities ETF, ticker IDU;
4. First Trust Nasdaq-100 Ex-Technology Sector Index Fund, ticker QQXT;
5. SPDR Euro Stoxx 50 ETF, ticker FEZ;
6. iShares 20+ Year Treasury Bond ETF, ticker TLT;
7. iShares iBoxx USD High Yield Corporate Bond ETF , ticker HYG.

Figure 6.1 shows the normalized prices of the 7 financial assets considered.

The dataset has been divided into two parts, a first part of in-sample data covering the period from 5/2007 to 12/2015 and a second part of out-of-sample data covering the period from 1/2016 to 5/2019.

The in-sample data were used to generate 2 different sets of scenarios, both composed of $N = 200$ realizations of $T = 12$ periods each, for a total of $N * T = 2400$ scenario samples per set. The first set of scenarios was used as a train set to compute the optimal allocation of the tested allocation strategies. The second set was used as a validation set to perform the tuning of the hyperparameters of the tested allocation strategies.

Finally, using the out-of-sample data, a third set of scenarios was generated, also consisting of 200 realizations of 12 periods each and was used to test the implemented strategies and compute the out-of-sample performances.

All the scenarios were generated using the Bootstrap method (sampling with replacement from the hystorical returns [45]) to preserve the probability distribution of the real returns.

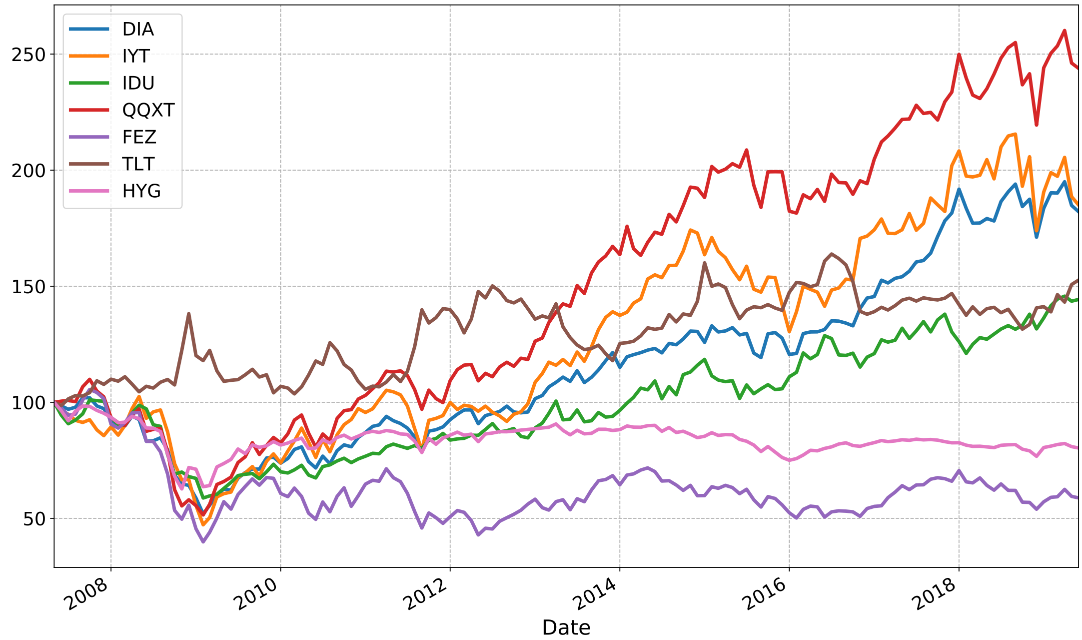

FIGURE 6.1: Normalized prices of financial instrument over time.

6.4.1 Implemented strategies

Four different allocation strategies were tested and compared with a benchmark. The allocation strategies implement the following policies:

1. Affine policy with single-stage memory depth:

$$\Pi_1 : u(k) = \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k)) \tag{6.37}$$

2. Affine policy with double-stage memory depth:

$$\Pi_2 : u(k) = \bar{u}(k) + \Theta_1(k)(g(k) - \bar{g}(k)) + \Theta_2(k)(g(k-1) - \bar{g}(k-1)) \tag{6.38}$$

3. Policy with single-stage memory depth and squared trasformation of the inputs:

$$\Pi_3 : u_i(k) = \bar{u}_i(k) + \sum_{j=1}^{n} w_{i,j}^T(k)\varphi_j(g(k)) \quad i = 1, \dots, n \tag{6.39}$$

where:

$$\varphi_j : \mathbb{R}^n \longrightarrow \mathbb{R}, \qquad j = 1, \dots, n \tag{6.40}$$

$$g_j(k) \mapsto g_j(k)^2 \qquad j = 1, \dots, n \tag{6.41}$$

4. Policy with Gaussian kernel:

$$\Pi_4 : u_i(k) = \bar{u}_i(k) + \sum_{t=\tau}^{k} \alpha_{i,t}(k) K(\mathrm{g}(k), \mathrm{g}(t)) \quad i = 1, \ldots, n \tag{6.42}$$

where $K(\mathrm{g}(k), \mathrm{g}(t))$ is given by 6.36.

The strategy used as a benchmark is the *naive* portfolio diversification rule, the strategy in which a fraction $1/n$ of investor's wealth is allocated to each of the n assets at each rebalancing stage. Although its simplicity, the naive strategy has proved to be extremely difficult to outperform out-of-sample [38].

The constraints on the portfolio composition have been set to $\mathcal{X}(k) = \{x(k) : x(k) \geq 0, k = 1, \ldots, T\}$, while the initial holdings were set as $x_i(0) = 1/n$ for $i = 1, \ldots, n$.

Two instance of optimization problem 6.4 with policy Π_1 and Π_2 were solved. To improve the out-of-sample performance a regularization term implementing a penalty on the squared norm of the decision variable was added to cost function 6.4a with $\lambda = 0.1$ and $\lambda = 0.5$ for policy Π_1 and Π_2 respectively. One instance of optimization problem 6.18 was solved with policy Π_3 and $\lambda = 0.125$. Finally one instance of optimization problem 6.26 was solved with policy Π_4 and memory depth $\tau = 12$, $\sigma = 0.9$ and $\lambda = 0.5$.

All the optimization problems were solved with portfolio performaces measured by the expected shortfall given by 6.2 and with $\nu = 2$. The values of the hyperparameter were choosen through cross validation using the set of validation scenarios.

6.4.2 Out-of-sample results

Figure 6.2 shows the results of a first numerical test where we compted the out-of-sample multi-period efficient frontiers for the implemented allocation strategies and the benchmark. It represents the optimal trade-off curve of the minimal expected shortfall obtainable for a given value of target final wealth γ. We considered 20 equi-spaced values of the target final wealth γ in the interval $[1, 1.2]$. From the figure it is observed that all the 4 strategies that implement the closed-loop policies outperform the trivial allocation strategy. Policy Π_2 with double-step memory does not give a substantial improvement in performance compared to policy Π_1 with a single-step memory. Policy Π_3 with single-step memory and square returns provides an improvement in performance compared to policy Π_1 and policy Π_2. Overall policy Π_4 which implements the kernel-based strategy dominates the others in terms of performance.

Figure 6.3 shows the results of a second numerical test where we computed the out-of-sample histograms of the final wealth achieved by each strategy setting the final target wealth to $\gamma = 1.08$ (which corresponds to a return of 8% at the end of a 12-month investment period). For each panel in figure 6.3 the red vertical dashed line represents the target wealth of $\gamma = 1.08$ while the black dashed line represents the average value of the wealths obtained by the implemented strategies for the different out-of-sample scenarios. The average values of wealth obtained with the strategies are respectively for policy Π_1 a value of 1.057, for policy Π_2 a value of 1.058, for policy Π_3 a value of 1.061, for policy Π_4 a value of 1.077, and for the naive policy a value of 1.052.

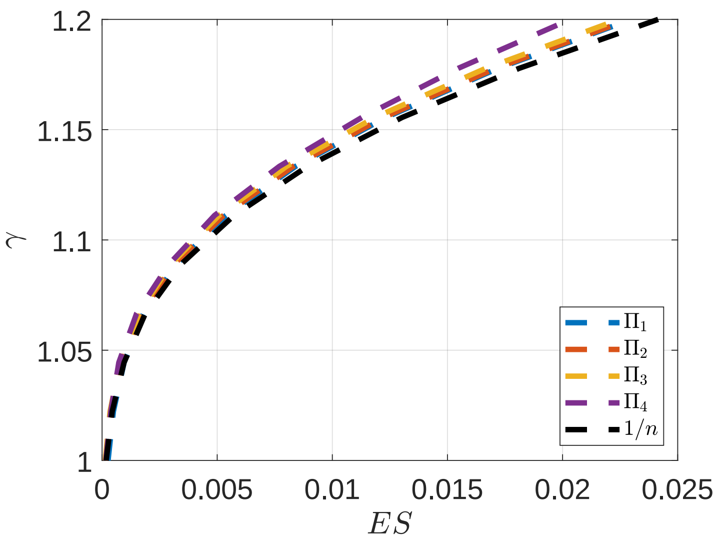

(A) Efficient frontier in the range $\gamma \in [1, 1.2]$.

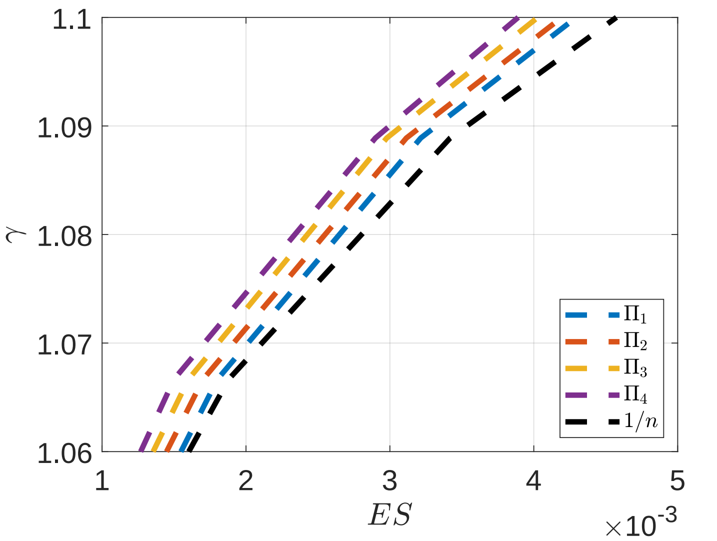

(B) Efficient frontier zoomed on the range $\gamma \in [1.06, 1.1]$.

FIGURE 6.2: Comparison of the out-of-sample multi-period efficient frontiers of the tested allocation strategies and the benchmark.

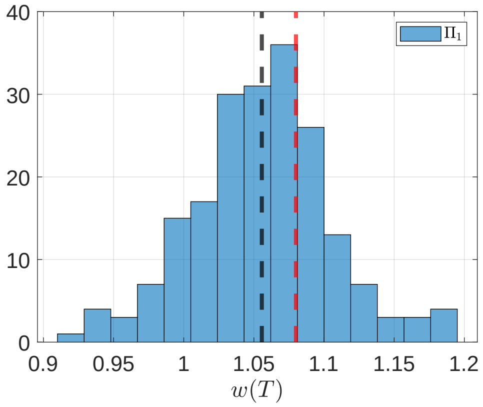

(A) Histogram of the final wealth obtained with policy Π_1

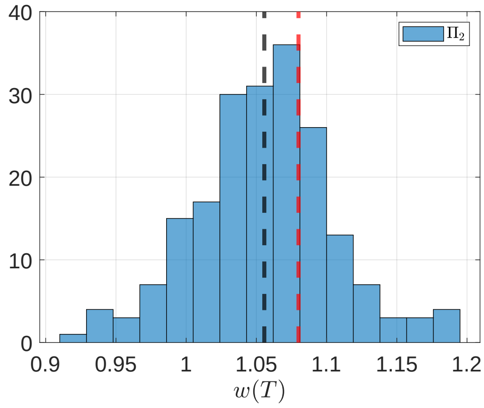

(B) Histogram of the final wealth obtained with policy Π_2

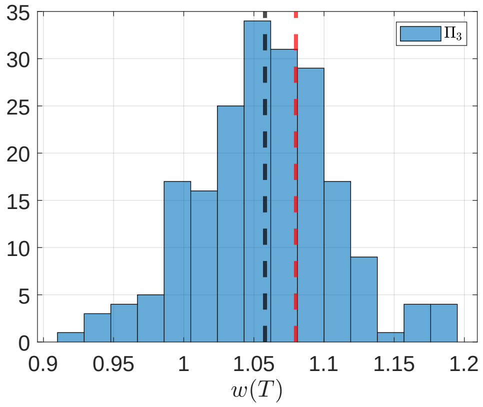

(C) Histogram of the final wealth obtained with policy Π_3

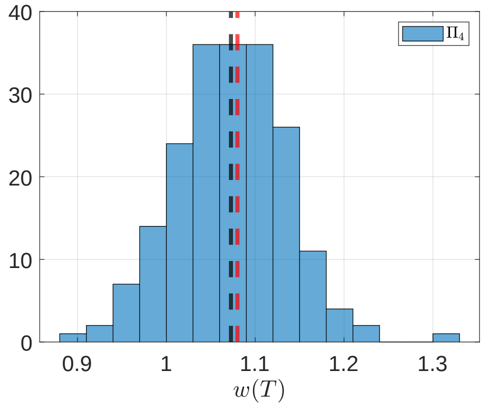

(D) Histogram of the final wealth obtained with policy Π_4

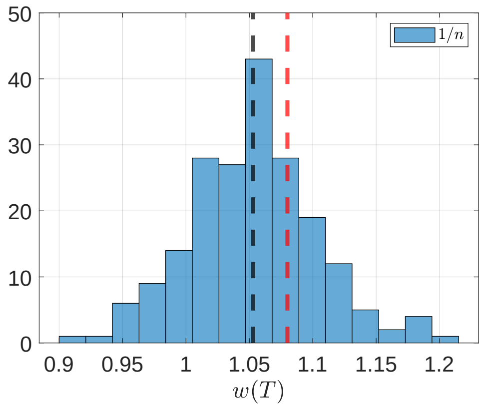

(E) Histogram of the final wealth obtained with naive policy

FIGURE 6.3: Histograms of the final wealths obtained with the implemented allocation strategies

Chapter 7

Conclusions

This thesis dealt with two of the main typical financial problems that are the trading of individual assets and the trading of a portfolio of multiple assets. In particular, two recent lines of research have been described that have the particularity of formulating such financial problems as control problems: reactive trading and multi-period portfolio optimization.

In Chapter 1 these two lines of research were introduced and their open problems were discussed.

In Chapter 2 we have described the line of research called reactive trading whose main innovation is to treat stock prices as external disturbances affecting the system. Thanks to a trading scheme capable of simultaneously implementing a long and a short investment strategy using two feedback loops running in parallel, called SLS, it is possible to guarantee certain performance levels under some market assumptions and independently of the direction of the market. Two main open problems of this approach were then highlighted, which are one of the reasons for the existence of this thesis.

In Chapter 3 we have described the literature concerning the trading of multiple financial assets that make up a portfolio. We started by describing the classic financial literature based on Markowitz's single-period optimization approach. We then highlighted the main problems related to this classical approach and new methodologies born over the years to deal with these problems. In the second part of Chapter 3 we have described an important extension of the classical approach which is multi-period optimization. In particular, we focused on recent approaches based on predictive control as they are able to be formulated as convex optimization problems that can be solved efficiently.

In Chapter 4 we described the first innovative contribution of this thesis to reactive trading. In the SLS the controllers that are on the feedback loops are time-invariant, we have instead proposed time-varying controllers based on the logic of the Extremum Seeking and able to adapt more readily to the variations of dynamics typical of the prices of market and caused by unpredictable economic and political events. Convergence analysis of the proposed scheme will be object of future works together with a sensitivity analysis of the algorithm's performance to the variation of the tuning frequency of the parameters of the model that can be set run-time in an adaptive way.

In Chapter 5 the second innovative contribution to reactive trading was described. Starting from the mere assumption that price returns are uncertain but bounded has allowed us to reformulate reactive trading as a problem of robust control of a system affected by parametric uncertainty and to synthesize robust controllers with ranges of variation in price returns. Future studies will aim to demonstrate, in a probabilistic, the stability of the proposed approach in the case of price returns governed by simple distributions such as the Gaussian distribution.

Finally, in Chapter 6, we moved to multi-period portfolio optimization. In particular we have described how it is possible to use typical techniques of statistical learning such as kernel-based methods to obtain non-linear investment policies able to better capture possible complex market dynamics rather than linear policies described in the literature. Future developments of this work will consist in using more specific kernel functions for temporal problems such as the methods described in [92], furthermore, massive experimental tests will be carried out to further validate the proposed methodologies.

Bibliography

[1] M. A. Aizerman, E. A. Braverman, and L. Rozonoer. Theoretical foundations of the potential function method in pattern recognition learning. In *Automation and Remote Control,*, number 25 in Automation and Remote Control,, pages 821–837, 1964.

[2] Vedat Akgiray. Conditional heteroscedasticity in time series of stock returns: Evidence and forecasts. *The Journal of Business*, 62(1):55–80, 1989.

[3] Pierre Apkarian, Minh Ngoc Dao, and Dominikus Noll. Parametric robust structured control design. *IEEE Transactions on automatic Control*, 60(7):1857–1869, 2015.

[4] Kartik B Ariyur and Miroslav Krstic. *Real-time optimization by extremum-seeking control*. John Wiley & Sons, 2003.

[5] David H Bailey, Jonathan Borwein, Marcos Lopez de Prado, and Qiji Jim Zhu. Pseudo-mathematics and financial charlatanism: The effects of backtest overfitting on out-of-sample performance. *Notices of the American Mathematical Society*, 61(5):458–471, 2014.

[6] Clifford A Ball and Walter N Torous. On jumps in common stock prices and their impact on call option pricing. *The Journal of Finance*, 40(1):155–173, 1985.

[7] Gah-Yi Ban, Noureddine El Karoui, and Andrew E. B. Lim. Machine learning and portfolio optimization. *Management Science*, 64(3):1136–1154, 2018.

[8] B. Ross Barmish. On trading of equities: a robust control paradigm. *IFAC Proceedings Volumes*, 41(2):1621–1626, 2008.

[9] B. Ross Barmish. On performance limits of feedback control-based stock trading strategies. *Proceedings of the 2011 American Control Conference*, pages 3874–3879, 2011.

[10] B Ross Barmish and James A Primbs. On a new paradigm for stock trading via a model-free feedback controller. *IEEE Transactions on Automatic Control*, 61(3):662–676, 2015.

[11] Michael Heinrich Baumann. On stock trading via feedback control when underlying stock returns are discontinuous. *IEEE Transactions on Automatic Control*, 62(6):2987–2992, 2016.

[12] Michael Heinrik Baumann. On stock trading via feedback control when underlying stock returns are discontinuous. *IEEE Transactions on Automatic Control*, 62(6):2987–2992, 2017.

[13] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton University Press, 2009.

[14] Michael J Best and Robert R Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *The review of financial studies*, 4(2):315–342, 1991.

[15] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.

[16] John H. Blakelock. *Automatic control of aircraft and missiles*. John Wiley & Sons, 1991.

[17] Robert C Blattberg and Nicholas J Gonedes. A comparison of the stable and student distributions as statistical models for stock prices. *The journal of business*, 47(2):244–280, 1974.

[18] Zvi Bodie, Alex Kane, and Alan J Marcus. Investments, 10th global edition. *Berkshire: Mc Graw-Hill*, 2014.

[19] Zvi Bodie and Robert C Merton. Finance.- prentice hall. *Upper Saddle River, NJ*, 2000.

[20] Stephen Boyd, Enzo Busseti, Steve Diamond, Ronald N Kahn, Kwangmoo Koh, Peter Nystrup, Jan Speth, et al. Multi-period trading via convex optimization. *Foundations and Trends® in Optimization*, 3(1):1–76, 2017.

[21] Stephen Boyd, Mark T Mueller, Brendan O'Donoghue, Yang Wang, et al. Performance bounds and suboptimal policies for multi–period investment. *Foundations and Trends® in Optimization*, 1(1):1–72, 2014.

[22] Paolo Brandimarte. *An Introduction to Financial Markets: A Quantitative Approach*. John Wiley & Sons, 2017.

[23] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.

[24] Giuseppe Calafiore and Marco C Campi. On two-stage portfolio allocation problems with affine recourse. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 8042–8047. IEEE, 2005.

[25] Giuseppe Calafiore and M.C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, Jan 2005.

[26] Giuseppe C Calafiore and Marco C Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

[27] Giuseppe C Calafiore and Laurent El Ghaoui. *Optimization models*. Cambridge university press, 2014.

[28] Giuseppe C Calafiore and Bruno Monastero. Triggering long-short trades on indexes. *International Journal of Trade, Economics and Finance*, 1(3):289, 2010.

[29] Giuseppe Carlo Calafiore. Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473, 2008.

[30] Giuseppe Carlo Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM Journal on Control and Optimization*, 48(4):2254–2274, 2009.

[31] Giuseppe Carlo Calafiore. Optimal dynamic asset allocation with lower partial moments criteria and affine policies. *International Journal of Financial Engineering and Risk Management*, 2(2):87–108, 2015.

[32] Louis KC Chan, Jason Karceski, and Josef Lakonishok. On portfolio optimization: Forecasting covariances and choosing the risk model. *The review of Financial studies*, 12(5):937–974, 1999.

[33] Nicolas Chapados. *Portfolio choice problems: An introductory survey of single and multiperiod models*. Springer Science & Business Media, 2011.

[34] Vijay K Chopra and William T Ziemba. The effect of errors in means, variances, and covariances on optimal portfolio choice. In *Handbook of the Fundamentals of Financial Decision Making: Part I*, pages 365–373. World Scientific, 2013.

[35] Andreas F Clenow. *Following the trend: diversified managed futures trading*. John Wiley & Sons, 2012.

[36] Min Dai, Qing Zhang, and Qiji Jim Zhu. Trend following trading under a regime switching model. *SIAM Journal on Financial Mathematics*, 1(1):780–810, 2010.

[37] Victor DeMiguel, Lorenzo Garlappi, Francisco J. Nogales, and Raman Uppal. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5):798–812, 2009.

[38] Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, 22(5):1915–1953, 2007.

[39] Atul Deshpande and B Ross Barmish. A generalization of the robust positive expectation theorem for stock trading via feedback control. In *2018 European Control Conference (ECC)*, pages 514–520. IEEE, 2018.

[40] Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. Deep learning for event-driven stock prediction. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.

[41] Vladimir Valentinovich Dombrovskii, Dmitrii Vladimirovich Dombrovskii, and Elena Aleksandrovna Lyashenko. Predictive control of random-parameter systems with multiplicative noise. application to investment portfolio optimization. *Automation and remote control*, 66(4):583–595, 2005.

[42] Vladimir Valentinovich Dombrovskii, Dmitrii Vladimirovich Dombrovskii, and Elena Aleksandrovna Lyashenko. Model predictive control of systems with random dependent parameters under constraints and its application to the investment portfolio optimization. *Avtomatika i Telemekhanika*, (12):71–85, 2006.

[43] VV Dombrovskiy, DV Dombrovskiy, and EA Lyashenko. Investment portfolio optimisation with transaction costs and constraints using model predictive control. In *Proceedings. The 8th Russian-Korean International Symposium on Science and Technology, 2004. KORUS 2004.*, volume 3, pages 202–205. IEEE, 2004.

[44] Bradley Efron and Gail Gong. A leisurely look at the bootstrap, the jackknife, and cross-validation. *The American Statistician*, 37(1):36–48, 1983.

[45] Bradley Efron and Robert J Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.

[46] Edwin J Elton and Martin J Gruber. Modern portfolio theory, 1950 to date. *Journal of Banking & Finance*, 21(11-12):1743–1759, 1997.

[47] Edwin J Elton, Martin J Gruber, Stephen J Brown, and William N Goetzmann. *Modern portfolio theory and investment analysis*. John Wiley & Sons, 2009.

[48] Frank J. Fabozzi, Petter N. Kolm, Dessislava A. Pachamanova, and Sergio M. Focardi. Robust portfolio optimization. *The Journal of Portfolio Management*, 33(3):40–48, 2007.

[49] Eugene F Fama. The behavior of stock-market prices. *The journal of Business*, 38(1):34–105, 1965.

[50] Björn Fastrich, Sandra Paterlini, and Peter Winker. Constructing optimal sparse portfolios using regularization methods. *Computational Management Science*, 12(3):417–434, 2015.

[51] Yiyong Feng, Daniel P Palomar, et al. A signal processing perspective on financial engineering. *Foundations and Trends® in Signal Processing*, 9(1–2):1–231, 2016.

[52] George M Frankfurter, Herbert E Phillips, and John P Seagle. Portfolio selection: the effects of uncertain means, variances, and covariances. *Journal of Financial and Quantitative Analysis*, 6(5):1251–1262, 1971.

[53] Peter A Frost and James E Savarino. For better performance: Constrain portfolio weights. *Journal of Portfolio Management*, 15(1):29, 1988.

[54] J Brian Gray and Dan W French. Empirical comparisons of distributional models for stock index returns. *Journal of Business Finance & Accounting*, 17(3):451–459, 1990.

[55] Selim S Hacısalihzade. *Control Engineering and Finance*. Springer, 2018.

[56] Florian Herzog, Gabriel Dondi, and Hans P Geering. Stochastic model predictive control and portfolio optimization. *International Journal of Theoretical and Applied Finance*, 10(02):203–233, 2007.

[57] Florian Herzog, Simon Keel, Gabriel Dondi, Lorenz M Schumann, and Hans P Geering. Model predictive control for portfolio selection. In *2006 American Control Conference*, pages 8–pp. IEEE, 2006.

[58] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

[59] Sesan Iwarere and B Ross Barmish. A confidence interval triggering method for stock trading via feedback control. In *Proceedings of the 2010 American Control Conference*, pages 6910–6916. IEEE, 2010.

[60] Ravi Jagannathan and Tongshu Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1683, 2003.

[61] Samuel Kotz, Tomasz Kozubowski, and Krzystof Podgorski. *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media, 2012.

[62] Miroslav Krstic and Hsin-Hsiung Wang. Stability of extremum seeking feedback for general nonlinear dynamic systems. *Automatica*, 36(4):595–601, 2000.

[63] Manish Kumar and M Thenmozhi. Forecasting stock index movement: A comparison of support vector machines and random forest. In *Indian institute of capital markets 9th capital markets conference paper*, 2006.

[64] Christopher G Lamoureux and William D Lastrapes. Heteroskedasticity in stock return data: Volume versus garch effects. *The journal of finance*, 45(1):221–229, 1990.

[65] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411, 2004.

[66] Hayne E Leland. Beyond mean–variance: Performance measurement in a nonsymmetrical world (corrected). *Financial analysts journal*, 55(1):27–36, 1999.

[67] Baruch Lev and S Ramu Thiagarajan. Fundamental information analysis. *Journal of Accounting research*, 31(2):190–215, 1993.

[68] Mikael Linden. A model for stock return distribution. *International Journal of Finance & Economics*, 6(2):159–169, 2001.

[69] Andrew W Lo, Harry Mamaysky, and Jiang Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The journal of finance*, 55(4):1705–1765, 2000.

[70] Shirzad Malekpour and B Ross Barmish. On stock trading using a controller with delay: The robust positive expectation property. In *IEEE 55th Conference on Decision and Control (CDC)*, pages 2881–2887. IEEE, 2016.

[71] Shirzad Malekpour, James A Primbs, and B Ross Barmish. On stock trading using a pi controller in an idealized market: The robust positive expectation property. In *IEEE 52nd Annual Conference on Decision and Control (CDC)*, pages 1210–1216. IEEE, 2013.

[72] Shirzad Malekpour, James A Primbs, and B Ross Barmish. A generalization of simultaneous long–short stock trading to pi controllers. *IEEE Transactions on Automatic Control*, 63(10):3531–3536, 2018.

[73] Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

[74] Harry M. Markowitz. *Portfolio Selection: Efficient Diversification of Investments.* Yale University Press, 1959.

[75] Harry M Markowitz. Foundations of portfolio theory. *The journal of finance*, 46(2):469–477, 1991.

[76] María del Mar López Martín, Catalina García García, and José García Pérez. Treatment of kurtosis in financial markets. *Physica A: Statistical Mechanics and its Applications*, 391(5):2032–2045, 2012.

[77] Robert C Merton. Optimum consumption and portfolio rules in a continuous-time model. In *Stochastic Optimization Models in Finance*, pages 621–661. Elsevier, 1975.

[78] Robert C Merton. Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3(1-2):125–144, 1976.

[79] Robert C Merton. On estimating the expected return on the market: An exploratory investigation. *Journal of financial economics*, 8(4):323–361, 1980.

[80] Robert C Merton and Paul Anthony Samuelson. Continuous-time finance. *The Journal of Finance*, 46(4):1567, sep 1990.

[81] Attilio Meucci. *Risk and asset allocation.* Springer Science & Business Media, 2009.

[82] Thomas Mikosch and Cătălin Stărică. Nonstationarities in financial time series, the long-range dependence, and the igarch effects. *Review of Economics and Statistics*, 86(1):378–390, 2004.

[83] John M Mulvey, William R Pauling, and Ronald E Madey. Advantages of multiperiod portfolio models. *The Journal of Portfolio Management*, 29(2):35–45, 2003.

[84] Daniel B Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, pages 347–370, 1991.

[85] D. Nešić. Extremum seeking control: convergence analysis. *European Journal of Control*, 15(3):331–347, 2009.

[86] John P Nolan. Modeling financial data with stable distributions. In *Handbook of heavy tailed distributions in finance*, pages 105–130. Elsevier, 2003.

[87] Joseph D O'Brien, Mark E Burke, and Kevin Burke. A generalized framework for simultaneous long-short feedback trading. *arXiv preprint arXiv:1806.05561*, 2018.

[88] Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

[89] Alba V Olivares-Nadal and Victor DeMiguel. A robust perspective on transaction costs in portfolio optimization. *Operations Research*, 66(3):733–739, 2018.

[90] Sina Khoshfetrat Pakazad, Henrik Ohlsson, and Lennart Ljung. Sparse control using sum-of-norms regularized model predictive control. In *52nd IEEE Conference on Decision and Control*, pages 5758–5763. IEEE, 2013.

[91] Cheol-Ho Park and Scott H Irwin. What do we know about the profitability of technical analysis? *Journal of Economic Surveys*, 21(4):786–826, 2007.

[92] Gianluigi Pillonetto, Francesco Dinuzzo, Tianshi Chen, Giuseppe De Nicolao, and Lennart Ljung. Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3):657–682, 2014.

[93] James A. Primbs. *Applications of MPC to Finance*, pages 665–685. Springer International Publishing, Cham, 2019.

[94] James A Primbs and B Ross Barmish. On robustness of simultaneous long-short stock trading control with time-varying price dynamics. *IFAC-PapersOnLine*, 50(1):12267–12272, 2017.

[95] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

[96] Paul A Samuelson. Lifetime portfolio selection by dynamic stochastic programming. In *Stochastic Optimization Models in Finance*, pages 517–524. Elsevier, 1975.

[97] Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond.* MIT press, 2001.

[98] G William Schwert and Paul J Seguin. Heteroskedasticity in stock returns. *the Journal of Finance*, 45(4):1129–1155, 1990.

[99] Martin Sewell. Characterization of financial time series, 2011.

[100] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms.* Cambridge university press, 2014.

[101] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory.* SIAM, 2009.

[102] William F Sharpe. Mutual fund performance. *The Journal of business*, 39(1):119–138, 1966.

[103] Alex J Smola and Bernhard Schölkopf. A tutorial on support vector regression. *Statistics and computing*, 14(3):199–222, 2004.

[104] Srikanth Sridharan, Divakar Chitturi, and Armando A Rodriguez. A receding horizon control approach to portfolio optimization using a risk-minimax objective for wealth tracking. In *2011 IEEE International Conference on Control Applications (CCA)*, pages 1282–1287. IEEE, 2011.

[105] Charles Stein. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. Technical report, Stanford University Stanford United States, 1956.

[106] Yuichi Takano and Jun-ya Gotoh. A nonlinear control policy using kernel method for dynamic asset allocation (< special issue> scope (seminar on computation and optimization for new extensions)). *Journal of the Operations Research Society of Japan*, 54(4):201–218, 2011.

[107] Stephen J Taylor. *Asset price dynamics, volatility, and prediction.* Princeton university press, 2011.

[108] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

[109] Reha H Tütüncü and M Koenig. Robust asset allocation. *Annals of Operations Research*, 132(1-4):157–187, 2004.

[110] Reha H Tütüncü and M Koenig. Robust asset allocation. *Annals of Operations Research*, 132(1-4):157–187, 2004.

[111] Tony Van Gestel, Johan AK Suykens, D-E Baestaens, Annemie Lambrechts, Gert Lanckriet, Bruno Vandaele, Bart De Moor, and Joos Vandewalle. Financial time series prediction using least squares support vector machines within the evidence framework. *IEEE Transactions on neural networks*, 12(4):809–821, 2001.

[112] Paul Wilmott. *Paul Wilmott on quantitative finance*. John Wiley & Sons, 2013.

[113] Huan Xu, Constantine Caramanis, and Shie Mannor. Robust regression and lasso. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1801–1808. Curran Associates, Inc., 2009.

[114] Yuji Yamada and James A Primbs. Model predictive control for optimal portfolios with cointegrated pairs of stocks. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 5705–5710. IEEE, 2012.