

Optimal Investments for Price Processes with Long Memory

Zsolt NIKA

Supervisors:

Dr. Miklós RÁSONYI

Dr. Péter SZOLGAY

A thesis presented for the degree of
Doctor of Philosophy



Roska Tamás Doctoral School of Sciences and Technology
Pazmány Péter Catholic University

Budapest, Hungary
2020

Contents

1	Introduction	1
1.1	Investment decisions	1
1.1.1	Possible Objectives of an Investment	1
1.1.2	Early results	2
1.1.3	Portfolio	2
1.1.4	Self-financing, admissibility, short-long position	4
1.1.5	Arbitrage	5
1.1.6	Liquidity	5
1.2	Utility functions and performance measures	6
1.2.1	Log-optimal investments	7
1.2.2	Growth, Annual yield	8
1.3	Memory	8
1.4	Leverage effect	9
1.5	Fundamental properties of asset prices	11
1.5.1	Stylized facts we include in our models	14
1.5.2	Stylized facts we neglect	14
1.6	Research goals	15
2	Stock Price Models	16
2.1	Introduction	16
2.2	Long memory	17
2.3	Conditionally Gaussian Model Class	18
2.4	AR(1), MA(∞) and ARMA(p,q)	19
2.5	Discrete Gaussian Stochastic Volatility model	20
2.5.1	Conditionally Gaussian	22
2.5.2	Autocorrelation	22
2.5.3	Mean and Variance	24
2.5.4	Skewness and Kurtosis	26

2.5.5	Parameter Choice	27
2.5.6	Histogram	28
2.6	Bilinear ARCH models	29
2.6.1	Autocorrelation	30
2.7	Some Remarks on the above Models	31
3	Investment Strategies	34
3.1	Introduction	34
3.2	Log-optimal Strategy	36
3.3	Markovian strategy	39
3.4	Approximative strategies	40
3.5	Approximative Markovian strategies	42
3.6	Threshold Strategy	43
3.7	Constantly rebalancing portfolios	46
4	Log-Optimal Portfolios for Conditionally Gaussian Processes	48
4.1	Introduction	48
4.2	Log-Optimal Solution	49
4.3	Analysing the Role of Memory	51
4.4	Approximation of the Log-Optimal Strategy	53
4.5	Conclusion	55
5	Adaptive Learning of Log-Optimal Strategies	57
5.1	Introduction	57
5.2	Stochastic Approximation	59
5.2.1	The Robbins–Monro algorithm	59
5.2.2	The Kiefer–Wolfowitz algorithm	60
5.3	Numerical Results	62
5.3.1	Analysing the Expected Growth Function	62
5.3.2	Performance of the Kiefer–Wolfowitz algorithm	64
5.3.3	Modifications	65
6	Optimal Trading of Fractional Brownian Motion	67
6.1	Model	67
6.2	Optimal strategy	69
6.3	Numerical results	71
6.4	Extension	73

A Proof of Stationary Solution and Ergodicity-like Assumption	78
A.1 DGSV model	78
A.2 BARCH model	80
B Moments of the Discrete Gaussian Stochastic Volatility	81
B.1 Conditional mean	82
B.2 Unconditional mean	82
C Conditional Probability	84
C.1 Conditional expected value in general	84
C.2 Numerical implementation of expectations	85

Preface

In this thesis, we will discuss algorithmic trading in which the financial assets used during investments have prices with *long memory*.

Trading refers to the sale and purchase of financial assets. These financial assets may be simple, such as the purchase of bonds or stocks, or they may be more complex; these are called derivatives. In algorithmic trading, the choice of which asset to buy and when to buy it, is not made directly by the investor. Instead, the choice is made by an algorithm, which may be written by the investor or another individual. These algorithms are almost invariably run on computers. Trading involves the transaction of immense sums. The wealth (asset under management, AUM) handled by the ten largest investment companies in 2018 was over 28 trillion US dollars, that is, more than 28 thousand billion. The importance of algorithmic trading is also demonstrated by Kissel's 2013 study [1], which describes an increase in the proportion of algorithmic trading from 15% to 85% between 2003 and 2012.

A subset of algorithmic trading is *high-frequency trading* (HFT), in which decisions on trading are made chronologically close together. There is no officially accepted definition for HFT; even daily transactions may be defined as such, while in many cases, only trading taking place at least every minute or second is considered high-frequency. According to another definition, the trading decision and its execution happen simultaneously without outside interference. Although we cannot speak of exact numbers, in all probability, at least a third of all trading is high-frequency.

Stock price processes are mathematical processes which determine the price of a financial asset at a given time. Clearly, this is an idealized definition, since a given product may be purchased by many buyers simultaneously and from different suppliers. Therefore, in reality, several values may be attributed to it, but we will disregard this fact. The uncertainty of the asset price can be handled by characterizing it with a random variable rather than by a given value.

The past decades have seen an enormous amount of statistical examination of the asset price as time series. As to be expected, the statistical examinations corroborate that the asset price's past values (or their changes or "time derivatives") affect the current asset price

and its increments. By 1971, Mandelbrot had already postulated that the increments of the price process not only have memory, but in fact, possess long memory. He used fractional Brown movement to model the price process, which was, however, mathematically difficult to fit with the concepts of the time. Subsequently, in 1977 [2], he showed using statistical methods on historical prices that the increments of the price process do indeed have long memory. To determine this, he examined the daily price of 200 stocks.

Thus, during trading, products are purchased and sold which can be best characterized by their price (or perhaps the purchased amount). Since asset price is mathematically modelled as a stochastic process and the investment is goal-oriented, i.e. maximizing a consumption or utility factor, the topic of investment theory belongs to the subject of *control theory* or *decision theory*. Within that, it belongs to stochastic, non-Markovian decision theory. Investment problems can be described corresponding to control theory: asset price processes can be described by the "control system" Markovian processes, where the controller is the investor's decision on the type and quantity of product to purchase, while the objective function is the investor's utility function. While we focus primarily on log-optimal utility functions in the scope of this thesis, we will also come to interesting conclusions on risk-neutral investments. In control theory, the Markovian problem has well-known algorithms that can find the optimal decision in most cases. This cannot be said for non-Markovian problems. For example, it has been established for almost 30 years now that given log-optimal investments– to word it simply– if the increase of the asset price is both stationary and ergodic, then there exists an optimal investment [3]. It is easy to test the stationarity, but proving ergodicity is more problematic.

Concerning the practical aspects of investment, it is clear that despite the obvious presence of strong memory and a paltry mathematical toolbox, trading algorithms are created and used by many. This is indicated by its trading revenue and also by the fact that there are more than 20,000 hits on Amazon books that contain the terms "investment" and "financial." Real trading algorithms attempt to handle past information despite severe lack of constructive theoretical results.

The main goal of this thesis is to bring the two worlds of theory and application closer: we will define long memory processes that are realistic and have statistical characteristics important for us, such as long memory in their asset price processes. Furthermore, we will find specific algorithms that optimize the objective of the investor.

In the following, we will assume the reader has an understanding of the basics of probability theory, including the meaning of expected value, correlation, autocorrelation, standard deviation, and variance. In the discussion, other moments will be considered as well, which

will be defined in the respective sections. Expected value has an especially important role in financial mathematics, therefore a separate chapter is dedicated to its precise interpretation in the Appendix. Nevertheless, the entirety of this thesis is comprehensible without these.

The research goals of this thesis can be found in Section 1.6 at the end of Chapter 1, after the introduction of basic investment concepts.

Chapter 1

Introduction

Before we introduce our methodology in the next chapter, in this one we provide the reader with the basic concepts of investment theory – focusing only on the relevant issues for us. The notations that we use here are changed in Chapter 6 but otherwise it is consistently used throughout the whole dissertation.

1.1 Investment decisions

1.1.1 Possible Objectives of an Investment

There are several different types of algorithmic investments depending on their objectives. For example, one objective may be to decrease the market impact; typically, this is necessary when there are larger purchases. In this case, the investor buys the desired amount in several smaller purchases. In this way, the total transaction fee is lower (for example, a larger purchase may increase the price). This is almost invariably necessary during high frequency trading.

A frequent objective is following benchmarks, in which we expect the performance of a portfolio to "follow" the performance of a benchmark. The indices (or stock indices) of the financial market are products that aim to measure specific states of the market; usually, these serve as the benchmarks. Such an index is the S&P 500 index, which measures the joint performance of its 500 stocks. Commonly, indices measure the performance of select market segments. In passive investments, the investor's objective is that the portfolio's value changes correspond to that of the index.

Likewise, another frequent goal is the investor's desire to maximize their expected wealth while taking into account the risk. Broadly speaking, the goal of the investor is maximizing some type of *utility function*. At the same time, it is up to each investor how they take risk into account. In this dissertation, *log-utility* function and *risk neutral* investment will be

discussed.

Hereafter, we will assume that the investor's goal is exclusively to maximize a given utility function.

1.1.2 Early results

Modern portfolio optimization started with Harry Markowitz's theory [4], barely 20 years after Kolmogorov had laid the foundations of probability theory. Markowitz handled risk simply as the variance of the return and targeted to maximize expected return subject to a constant variance. For this theory Markowitz won a Nobel prize in Economics sciences in 1990. While this theory is easy to use in practice, thanks to the fact that it operates with simple moments that can be easily estimated, it is still overly simplistic.

In 1956, J. L. Kelly formulated a strong connection between information rate and portfolio theory [5]. A better approach is to maximise the expected logarithm (called *utility function*) of the portfolio value. This idea originated from Daniel Bernoulli, see [6]. Famous investors who used log-optimal investments include Jim Simons and Warren Buffett from Renaissance Medallion as well as the economist John Maynard Keynes. While plain expectation is a generalization of the arithmetic mean, the logarithm transforms it into a geometric mean which is more sensitive to the fluctuations in the future value, i.e. to risk. Log-optimal portfolio beats *every* other portfolio in the long run, see e.g. Chapter 16 of [7].

1.1.3 Portfolio

A portfolio is a collection of financial assets (e.g. stocks, bonds) held by an investor whose intention is to optimise some kind of benefit from this ownership. The most important feature of a portfolio is its value, which we refer to as the *wealth* of the investor, denoted at time t by W_t . The wealth of an investor, or the value of their portfolio, is given by the sum of the financial assets possessed by the investor in the portfolio.

Financial assets may be fundamentally classed into two groups: they can be characterized by their price or by their deterministic or stochastic processes. Since stochastic processes typically have variance and variance is unpredictable in the future, i.e. risk, they are usually termed *riskless* and *risky* assets. An example of a deterministic process is a bond or cash; its value at t is B_t . For the stochastic process, our primary example is stock – S_t .

The evolution of stock prices in time can be regarded as a stochastic dynamical system where the state variables are described by e.g. a (stochastic) differential equation. Such models implicitly assume that it is enough to know the structure of the dynamics and the actual state to make calculations for the future (i.e. the Markovian property is assumed). The theory of portfolio management is well-studied in this case, where the current state

(or the past few states) contains all the information we can have about the distribution of future prices. On the other hand, why shouldn't we take into account all the information from past data, as empirical evidence suggests (see e.g. [8, 9, 10, 11, 12])?

In the following sections we will see that it is more convenient to use the *log-return* process H_t , i.e. the increments of the log-price instead of using the stock price. We will mostly consider non-Markovian dynamics. For an overview about the statistical properties of prices, see [13] and the following sections. The most important property for us is the structure of time correlation which was studied already in the 90s, see e.g. [10, 11, 12].

We use discrete time models and the time parameter will be $t \in \mathbb{Z}$ or $t \in \mathbb{N}$. For simplicity, we deal only with one risky and one riskless asset. We denote the stock (risky asset) price by S_t and the bond (riskless asset) price by B_t at time t , with initial prices s_0 and b_0 , respectively.

Investment starts at $t = 0$ and the bond price evolves in time deterministically as $B_t = b_0(1+r)^t$, $t \in \mathbb{N}$, where $r \geq 0$ is the fixed interest rate. It is convenient to describe the stock dynamics by its log-return

$$H_t := \log(S_t/S_{t-1}) \quad (1.1)$$

whose statistical properties will be discussed in detail in the following sections. Sometimes we refer to log-return as return to avoid the multiple use of a word if its meaning is clear. *Return* is defined as $R_t := (S_t - S_{t-1})/S_{t-1} = S_t/S_{t-1} - 1$. That is $1 + R_t = \exp(H_t)$. First order Taylor-expansion of the exponential function say $\exp(x) \approx 1 + x$ if $|x|$ is small. Therefore on a daily time-scale return and log-return is very close to each other.

The investor's strategy will be described by a $[0, 1]$ -valued adapted process π_t , $t \in \mathbb{N}$ which represents the proportion of wealth allocated to the stock. More precisely, for initial wealth $W_0 = w_0$ is a constant and $w_0 > 0$ and for any strategy π_t , the investor's wealth W_t^π at time t follows the evolution

$$W_t^\pi = W_{t-1}^\pi \left((1 - \pi_{t-1}) \frac{B_t}{B_{t-1}} + \pi_{t-1} \frac{S_t}{S_{t-1}} \right), \quad t \in \mathbb{N}. \quad (1.2)$$

In our case, (1.2) is the most convenient form of the evolution of wealth in time. This is called *compounded capital*, because it is 'compounded' by the returns $(S_t/S_{t-1}$ and $B_t/B_{t-1})$. In the following chapters we always define the processes S_t and B_t explicitly.

There is another formalism which derives from its continuous counterpart, which uses additive evolution instead of the multiplicative we saw in (1.2). The portfolio value in this case is

$$W_t^\Delta - W_{t-1}^\Delta = \Delta_{t-1}^S (S_t - S_{t-1}) + \Delta_{t-1}^B (B_t - B_{t-1}), \quad (1.3)$$

where $\Delta_t^{S/B}$ denotes the *position* of the asset S or B at time $t - 1$, in other words, the

number of asset in the position at time $t - 1$.

Naturally, both equations are the same and could be used interchangeably, but they are more suitable in different contexts. With the choice

$$\pi_t = \frac{\Delta_{t-1}^S S_t}{W_{t-1}} \quad (1.4)$$

we end up in the same portfolio. The multiplicative formalism is especially useful when the utility function is logarithmic. Moreover, it is a more natural description, because asset prices typically increase or decrease exponentially. Usually, economists are interested in what percentage the price changed relative to a former value (*return*) and not the difference, the absolute change. On the other hand, the latter one is more often used in the 1) continuous-time case and 2) when friction is involved. Also, it is easier to understand (1.3), the change of the portfolio value comes from the changes in the assets' prices.

1.1.4 Self-financing, admissibility, short-long position

Portfolios can be grouped on the basis of whether they use external sources during the investment or solely the assets in the portfolio. External sources can be other financial assets that are not included in the portfolio or sources from investors and lenders.

We refer the reader for further reading on the following topic to consult [14].

In so-called *self-financing* portfolios, investors redistribute their wealth in each instance when they make a decision.

Definition 1.1.1. *A portfolio W_t^Δ is called self-financing if it can be written in the form of (1.3).*

An equivalent definition.

Definition 1.1.2. *A portfolio W_t^π , according to (1.2), is called self-financing.*

When investors sells more stock than they have, or buy more than they could, it is called short, respectively, long. Short and long position is when $\pi_t < 0$ and $\pi_t > 1$. In course of this document we assume that investor never holds short or long position, that is, $\pi_t \in [0, 1]$ and uses self-financing portfolio.

As a result of excluding external effects, we can focus on our aim, that is, the relation of the memory of the stock prices and the optimal decision.

Yet, we have not defined the dynamics of the stock prices which would define the probability space, but let us suppose that the stock price S_t is defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathbb{Z}}, \mathbb{P})$.

Definition 1.1.3. Let the trading strategy $\{\pi_t\}_{t \in \mathbb{N}}$ generate the portfolio W_t^π . The strategy is called admissible if it satisfies three conditions

1. The stochastic process π_t is adapted to the filtration \mathcal{F} . That is, the investor uses only information up to time t to make a decision.
2. The generated portfolio $W_t^\pi \geq 0$. This requirement is also called limited liability.
3. The created portfolio is self-financing.

1.1.5 Arbitrage

Arbitrage is the possibility of realizing profit without taking risk. It is usually a consequence of mispriced assets, i.e. market inefficiency. Many investors try to use arbitrage opportunities to take advantage of this 'glitch' in financial industry. If arbitrage exists and we find it, there is nothing to do about it: riskless profit would solve the problem of investment decisions. However, in efficient market it is always necessary to take risk into consideration.

There are several different types of arbitrage.

Definition 1.1.4. An admissible strategy π_t is called an arbitrage strategy, if

$$W_0^\pi = 0 \quad \text{and} \quad \mathbb{P}(W_t^\pi \geq 0) = 1, \mathbb{P}(W_t^\pi > 1) > 0$$

for some $t > 0$.

This is clearly not the case when applying time evolution (1.2), since $W_0 = 0$ would result in a portfolio $W_t = 0$, $t > 0$ almost sure.

We only need arbitrage-free market to secure that there is no profit without taking risk, hence making the investment problem well-posed. The continuous case is more difficult, in Chapter 6 we handle it.

1.1.6 Liquidity

Liquidity and other related effects like transaction costs are omitted here, except in Chapter 6 where these effects are in the spotlight.

In Chapters 4 and 5 we assume that investors can buy any stocks at any time without the obligation of paying for the transaction and we also assume that any trading does not have effect on stock prices.

The question we would like to answer is complex enough even under these assumptions. For the same reason, when we include illiquidity in Chapter 6, we have to make several assumptions to be able to find optimal strategy.

1.2 Utility functions and performance measures

Since stock prices are random variables in time, hence the future wealth of the investor, it is obvious that we have to transform the outcome of the investment into some ordered set to be able distinguish which is better. It follows that a well-chosen expected value can describe the problem. Actually, the idea of the expected value was probably more important in the past than probability. Probability as a number between 0 and 1 was only mentioned for the first time in 1713 by Jacob Bernoulli [15], while expected outcome of gambling games was used long before.

The first remarkable result from the 20th century about performance measures belongs to John von Neumann and Oskar Morgenstern [16] who established modern utility theory by giving axiomatic description to preferences. Their theory helps to construct well-defined utility functions, which are the expected value of some function of the outcome of the investment. Beside this, there is another extensively studied branch of performance measure, the so-called *mean-variance portfolio theory* by Harry Markowitz [4]. In this theory the risk is described by the variance of the investment profit and the goal is to maximize the expected profit minus the risk. A few decades later in 1979 a new approach, *prospect theory*, was born by Daniel Kahneman and Amos Tversky. Prospect theory based on the idea that risk aversion is asymmetric. This concept uses loss aversion which means that investors prefer avoiding losses to acquiring equivalent gains.

One of the main differences between investors is how sensitive they are to risk. Risk sensitivity usually is described by *risk aversion*. In modern portfolio theory risk aversion is regulated by how much the investor takes into account the variance as a penalty, while in utility theory, risk aversion is described by which function they choose as utility function.

At first let us assume, the investor is interested in optimizing the objective function at a given maturity/horizon time T . Let $u(x)$ be a real-valued and order-preserving function, called utility function (see more on its existence [17]). The investor's objective is to find a strategy $\{\pi_t\}_{t=1}^T$ that maximizes

$$\mathbb{E}[u(W_T^\pi)]. \quad (1.5)$$

Some popular choices are the power utility $u(x) = x^\theta/\theta$, for some $\theta < 1$ and $\theta \neq 0$ or the logarithmic utility $u(x) = \log(x)$. Risk-aversion is controlled by θ .

Some investors have no preferences about risk, in this case we have a *risk-neutral* investment problem. The objective function is simply

$$\mathbb{E}[W_T^\pi], \quad (1.6)$$

and the utility function is $u(x) = x$. This problem can be easily ill-posed and leads to

infinite utility. In Chapter 6, this investment problem will be investigated in the presence of illiquidity effect which makes the problem well-posed.

If there are no risk preferences and one wishes to use the risk-neutral setup, it is still desirable to qualify somehow the portfolio's performance. William F. Sharpe¹ introduced a performance measure in [18] and detailed in [19]. The ratio is defined in the case of zero risk-free return as

$$SR := \frac{\mathbb{E}[W_T^\pi]}{\sqrt{\text{Var}(W_T^\pi)}} \quad (1.7)$$

at a time horizon $T > 0$. Later other similar measures have been introduced similar to the Sharpe ratio, like the Sortino rate or the Treynor rate. These are more helpful in more economics-related problems.

1.2.1 Log-optimal investments

In the last 50 years the theory of log-optimal investments has been developed in a wide range of models, but with very little practical implication for the non-Markovian case. The handbook Maclean et al. [20], collected many important papers concerning the topic. We also mention the milestone paper [3], who characterized the log-optimal portfolio for stationary and ergodic log-return processes. However, this is a theoretical result that cannot be easily implemented in general.

Practical results for log-optimal investments in a distribution-free setting were investigated by Györfi et al., where the optimal strategy is achieved by machine learning methods, for a survey, see [21].

As opposed to the finite-horizon problem, infinite horizon problems describe the "average" behavior of the investment over time. We also refer to this kind of investments as *long-term* investments. In the case of logarithmic utility function, the infinite horizon problem is to maximize the time-average of the utility

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} [\log(W_T^\pi)]. \quad (1.8)$$

In the next chapters our approach, in contrast to the distribution-free solution, uses parametric models, where general properties of stock prices can be taken into consideration. Also, with parameters we are able to characterize the memory effect in the models in a simple way. We are also able to characterize how the optimal strategy or the value function depends on the memory.

¹Harry Markowitz, William F. Sharpe and Daniel Kahneman were awarded the Nobel Prize in Economics.

1.2.2 Growth, Annual yield

This is strongly connected to the log-optimal portfolios which are also called growth-optimal portfolios. The growth rate is only interpretable if wealth grows exponentially in time. If the limit exists, the growth rate G is defined as

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log(W_T^\pi) = G. \quad (1.9)$$

Usually it is very difficult to prove that this limit exists, in Chapter 3 and 4 we look at this problem.

Nevertheless, the growth rate G is a very practical concept, too. Mathematically loosely speaking

$$W_T^\pi \approx \exp(GT). \quad (1.10)$$

Annual yield (AY) (or yearly interest rate) equals to $\exp(250G) - 1$; that is, how much the portfolio's value grows in one year with growth-rate G .

When there are more than one stock or there are more realizations in a simulation, it is reasonable to speak of *Mean Annual Yield* (MAY):

$$MAY := \sum_i \exp(250G(i)) - 1, \quad (1.11)$$

where the sum goes takes every realization i , $G(i)$ is the growth of the i -th realization of the asset price process. (In financial practice it can also refer to the i -th asset of a portfolio with multiple assets, where MAY measure the performance of the whole portfolio.)

1.3 Memory

The main evidence for assuming non-Markovian dynamics for the log-return process of a stock price is coming from its time-correlation structure. It was already studied in the 90s, see e.g. [10, 11, 12].

Past information in the return is not presented in a linear relationship, as the autocorrelation function of H_t makes it clear. On the other hand, considering the autocorrelation function of the absolute log-return $|H_t|$ reveals how strong the memory effect really is. Most stock prices' returns have significant time-correlation if we disregard from returns' sign. The same true if one inspects the autocorrelation of the squared return H_t^2 or any positive power $|H_t|^p$, $p > 0$.

The covariance of two random variables X and Y with finite second moments is defined

as

$$\text{Cov}(X, Y) := \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]. \quad (1.12)$$

Autocorrelation of a stationary stochastic process X_t with finite second moment is defined as

$$\gamma_X(\tau) \equiv ACF_{X_t}(\tau) := \text{Cov}(X_{t+\tau}, X_t) / \text{Cov}(X_t). \quad (1.13)$$

An important property of correlation that it is only able to detect linear relationship between the variables but not a more complex structure.

Since memory effect is the cornerstone of the present thesis, therefore we investigated the daily returns of those 1250 stocks from New York Stock Exchange that have the most significant time-correlation structure. Of course, there are some stock returns without strong memory effect, but most of them do have. The webpage <http://www.long-memory.com/> collected important papers about long memory, especially in the case of financial problems until 2005.

Figures in 1.1 shows the evidence of long memory in log-returns. Usually the autocorrelation of the return H_t is insignificant, while for the absolute return $|H_t|$ it is significant.


A plot is called *log-log plot* if both axes are on logarithmic scales. This display mode helps to reveal power-law relationships, because a linear function on a log-log plot is equal to a power function on a linear-linear plot (when there is no scaling of the axes). A *log-lin plot* refers to the case, when only the vertical axis is rescaled logarithmically. A linear relationship on a log-lin plot suggests exponential behavior.

Most stock returns are similar to the Figure 1.1 (a) when the autocorrelation of $|H_t|$ are significant and decays with a power-law or similar to the Figure 1.1 (c) when power law decreasing transforms into an exponential decay. It is also possible to have a better exponential decay fit like in the case of Figure 1.1 (b).


1.4 Leverage effect

The name *leverage* comes from the name of the simple machine 'lever' used in physics to amplify one's strength to move a physical object. In finance, it refers to a technique where investors amplify the expected returns of their investments in some way. Speaking of firms, it usually means that a company purchases borrowed funds (debt) in the interest of buying financial assets. While this activity increases the expected return, it decreases the risk.


In financial mathematics it is commonly called as *leverage effect* and described slightly differently. Leverage effect is defined as the correlation between returns and future volatility. It was first introduced in [22]. Later it was studied exhaustively on real data [23]. They



(a) Typical behavior of the autocorrelation function (NYSE:ISF stock). There is no significant autocorrelation of H_t , while its absolute value $|H_t|$ decays by a power-law function, suggesting long memory effect.



(b) There are cases when the autocorrelation of $|H_t|$ decays exponentially.



(c) Log-returns of stock NYSE:LEG has weaker memory effect at first sight but the log-log plot justifies the power-law decay for the first 50 days lag.

Figure 1.1: Autocorrelation function of the log-returns of some daily stock prices (H_t) from the NYSE. Left panel shows the autocorrelation of H_t , while the middle and right panel show the autocorrelation of the absolute value of the return $|H_t|$.

The left and right panels are shown on lin-lin plot, while the right panel is shown on a log-log plot – suggesting power-law decay. Red dashed lines are the power-law fit on the data which is a straight line on a log-log plot.

found that the covariance function

$$\mathcal{L}(\tau) := \text{Cov}(H_{t+\tau}^2, H_t) \quad (1.14)$$

is negative and tends to 0 when $\tau \rightarrow \infty$. They investigated European stocks, but this relationship can be found in most stock prices. Therefore it is regarded now as a stylized fact of stock prices [13]. Leverage function could be defined as well as with the absolute value of the log-return

$$\tilde{\mathcal{L}}(\tau) := \text{Cov}(|H_{t+\tau}|, H_t). \quad (1.15)$$

What is important in the definitions is to eliminate the sign of H_t on one side. In the numerical results we refer to the function $\mathcal{L}(\tau)$ in (1.14) as *squared leverage function* and to the function $\tilde{\mathcal{L}}(\tau)$ in (1.15) as *absolute leverage function*.

Limit behavior is

$$\lim_{\tau \rightarrow \infty} \mathcal{L}(\tau) = \lim_{\tau \rightarrow \infty} \tilde{\mathcal{L}}(\tau) = 0. \quad (1.16)$$

The negative value of the correlation is related to the negative skewness of the return. The reason for this is that in the calculation of the skewness the correlation term in (1.14) appears. However, negative leverage function can result in positive skewness value as well. Typical values for daily trading ranges between (more or less) $[-2.4, 0.54]$.

Negative skewness indicate that the 'tail' of the distribution is on the left (negative) side. In practice it means that stock prices tends to increase more frequently but the size of the falls are bigger.

Descriptive statistics of daily log-return can be seen in the next section on Table 1.1. In 320 stocks out of the 1250 stocks that we have been investigated, skewness were positive. That is, while skewness is mostly negative, we cannot neglect when it is positive.


Some examples are shown on Figure 1.2.

As we will see in later chapters, leverage effect can be easily incorporated into the models using simply a negative coefficient for the volatility in the dynamics of the return.


1.5 Fundamental properties of asset prices

The two most important properties of asset prices have been presented in the previous sections: long memory and leverage effect. In [13] there is an exhaustive study of other properties of the stock prices. These are named as *stylized facts* and can be regarded as generally true characteristics of the return proces H_t .


Our analysis of historical log-return data can be found, besides of Figure 1.1 and 1.2, in Table 1.1 and Figure 1.3. We investigated daily (open) returns of 1250 stock prices from the



(a) Typical behavior of the leverage function. The negative exponential fit explains well the data. The skewness of the log-return of the analysed stock price NYSE:AED is around -0.99 .



(b) Leverage function of NYSE:ECA stock prices. In rare cases the function starts from a positive value ($\mathcal{L}(1)$ and $\tilde{\mathcal{L}}(1) > 0$). The skewness of the log-return is around 0.2 .



(c) Skewness of the log-return of NYSE:LEG is -0.44 .

Figure 1.2: Leverage function of three stock prices from the New York Stock Exchange. Left column shows the squared liquidity function (1.14), right column shows the absolute liquidity function (1.15). Red lines show the nonlinear fit $a \exp(b\tau)$, where a and b are the fitted parameters.

Statistics	Average	(0.05, 0.95) quantile
Mean	0.00014	(-0.00074, 0.00081)
St. dev.	0.02063	(0.00798, 0.04100)
Skewness	-0.65034	(-2.44097, 0.54823)
Kurtosis	41.3726	(5.19164, 157.17040)

Table 1.1: Statistical properties of log-returns investigated in 1250 stocks from New York Stock Exchange.




Figure 1.3: Histogram of historical log-return data from the New York Stock Exchange. The upper two plots show typical histograms: negative skewness and high kurtosis. While the bottom row displays deviant histograms: low kurtosis and close to 0 skewness. Stock prices: top row: NYSE:AED, NYSE:DDT, bottom row: NYSE:ECA, NYSE:LEG (these two were the deviant examples in Section 1.4).

New York Stock Exchange.

Uncertainty about stationary property of returns (see the following list) creates a trade-off: if we choose a large historical data set, then returns would likely be non-stationary; but if we choose a small data set, then it is not enough to make reliable conclusion about the data. For this reason we considered only those stock prices of NYSE that are long enough. This is the 1250 stocks we used. Furthermore, we chose that period when big financial crises did not happen, from 01/01/2010 for 1900 long daily data. (The end date can be different,

but a thumb of rule is that there is about 250 trading days in a year, more or less. It means the data reflects about 7.5 years.)

1.5.1 Stylized facts we include in our models

1. **Volatility clustering:** Mandelbrot described [24] for the first time that high volatility events are followed by more high volatility events, that is, volatility is clustered by its size in time. In [10] and [25] they describe that positive autocorrelation of $|H_t|$ is due to the volatility clustering and that volatility has long memory. Processes like ARCH [26] and GARCH [27] are used to model volatility clustering.
2. **Autocorrelation:** is insignificant for the return process H_t and significant for $|H_t|$. For more, see Section 1.3.
3. **Histogram of the returns:** higher moments of historical return data reveal that their distribution has fat/heavy tails, power-law or Pareto-like tail. The kurtosis of the distribution is extremely high and they usually have negative skewness, that is, more frequent small increasing in the price are followed by rare, but strong decreasing. For our results on the moments of log-returns from NYSE can be seen on Table 1.1.
4. **Leverage effect:** negative correlation between volatility and return. See Section 1.4 and Figure 1.2.
5. **Stationarity:** the question is whether the log-return process is stationary is difficult to answer. Common believe is to say yes, but it is only true for a few years period in real life. For example, a shock like financial crisis surely changes stationarity. Moreover, it is difficult to quantify stationarity. In statistics a unit-root test is used to answer this question, since if the unit-root equals to 1, then a linear time series model is non-stationary. Unfortunately this is not true in general.
6. **Ergodicity:** is a lot more difficult question than stationarity. The definition of ergodicity, to put it simply, is that a moment can be estimated by the process's time-average. On real data it is meaningless, since there is no way to prove that the data comes from an ergodic dynamics. Moreover, in most of the complex enough models it is very difficult to prove its ergodicity.

1.5.2 Stylized facts we neglect

Some of the stylized facts use variables that we do not aim to include in our models, such as trading volume. Some of them focus on temporal behavior instead of looking at the long term average behavior.

1. **Volume-volatility correlation:** there is a correlation between the size of a trade and the return's volatility. We assume that our investors have no impact on the market and we do not include volume in our models.
2. **Aggregational Gaussianity:** returns over longer period distributed normally. Returns of quarterly or yearly data often modelled as a Gaussian process.
3. **Intermittency:** probably the models we show in the following chapters describe this phenomena when return has big and irregular bursts.

1.6 Research goals

Our main goal is to create algorithms that find optimal investment strategies when the memory effect of stock prices are taken into consideration in parametric stock price dynamics.

In the previous sections we have seen that Algoet and Cover in their cornerstone study [3] assume stationarity and ergodicity, but this result is theoretical and does not propose a practical algorithm. Markovian methods, naturally, are not applicable in the long memory case. Another option could be to use distribution-free setting [21]. Using distribution free learning methods create the *black box problem*, a growing problem of today's machine learning algorithms. These procedures do not help in understanding the data and the decision.

We aim to create algorithms that help to understand the role of memory in investment and what kind of methods can be used. Algorithms described in Chapter 5 and 6 are very simple, but still, optimal in a suitable sense.

We heavily depend on numerical results. This proves the algorithm's feasibility and provides more information, helps to better understand the model.

In Chapter 2 we introduce the mathematical models of log-return of the stock prices, we work with in what follows. Investment strategies are defined and described in Chapter 3. Then in Chapter 4 we present and analyse the log-optimal portfolio with long memory. At this time, we suppose that the investor knows the 'whole past', while further on in Chapter 5 we suppose the minimum: the investor only knows past return values and we present a learning method that finds the optimal strategy. Chapters 2, 4 and 5 use discrete-time models without market impact and transactional costs; the investor uses logarithmic utility in the long-run case. On the other hand, Chapter 6 uses continuous-time model with transactional cost; the related strategy is described here instead of Chapter 3 because of the different nature of continuous processes. Because of the difficulty of the problem, in this case, the investor has no risk preferences and the optimal portfolio is evaluated and compared to real investments with analysing its Sharpe-ratio.

Chapter 2

Stock Price Models

At first, we give a definition of long memory. We could do so in several ways, but the one we use here is based on autocorrelation function, which makes it easy to detect on real data.

Then we define a general model class, named Conditionally Gaussian, that includes most well-known price dynamics model. Next we define some known stochastic processes, in particular with regard to long memory, and we also define our proposed model, called Discrete Gaussian Stochastic Volatility, that suits real data in many aspects. We study extensively its statistical properties: conditional/unconditional moments and autocorrelation functions.

2.1 Introduction

In 1900 L. Bachelier modelled financial processes as a Brownian Motion for the first time in [28]. Later, B. Mandelbrot extended it to Lévy processes with stable distribution in [29]. These models are lack of many, today regarded as fundamental properties, of asset prices that we described in Section 1.5, for example, volatility clustering and long memory.

Since then, price models have been extensively studied, summaries can be found. for example, in the handbook of Shepard [30] and Andersen [31]. There are two main types of models according to the handbook of [31], Stochastic Volatility models and GARCH-like models and we will present one of each type. Both types have a tremendous amount of examples already investigated but here we focus on those which capture the (long) memory effect in a simple and effective way. The basic difference between the two types of models is that in GARCH-like models the volatility is a deterministic function of past variables, while stochastic volatility models are not.

The theory of GARCH-like models started in 1982 with R. Engle's model called ARCH

(autoregressive conditional heteroskedasticity) process. In GARCH-like models the choice is clear: long memory is typically described by the Linear ARCH model (LARCH), defined by P. M. Robinson in [32], or its extension, the Bilinear ARCH model. For this reason we use the expression "GARCH-like", since LARCH and BARCH are not part of the GARCH models, but are typically mentioned as a variant of it. Therefore one of the two examples we give is the BARCH model. In Stochastic Volatility models long memory is described both in continuous time and discrete time. Two widely used examples are the Fractional Stochastic Volatility (FSV) model (in continuous time) by [8] and Long Memory Stochastic Volatility (LMSV) model by [33] (in discrete time, see also [34]). Our model proposition in the Stochastic Volatility framework can be regarded as a discretized version of the FSV as well as a more general formalization of LMSV; we call it Discrete Gaussian Stochastic Volatility (DGSV).

We not only define these two models in this section, but we also use heuristically chosen parameters based on real data and characterize the memory properties of the models. Clearly, our choice of models is somewhat arbitrary: yet it serves to demonstrate that even rather complex models easily fit into the model class we presented in Section 3.2.

A drawback of GARCH type models is that they typically require many parameters which are difficult to estimate; it is not easy to assign a financial meaning to them.

Continuous-time models, on the other hand, usually use fractional Brownian motion which, in contrast, captures the memory effect with only one parameter. The related portfolio optimization problems are, however, not solvable explicitly (except some very particular examples like in [35]).

Our model settings enables us to include memory in a simple manner in such a way that the memory can be characterized rather well by one parameter. At the same time, financial meaning can be assigned to the other parameters as well.

2.2 Long memory

First let us recall the definition in (1.13) of the auto-correlation function $\gamma_X(k)$ of a stationary stochastic process $\{X_t\}_{t \in \mathbb{Z}}$ with finite second moments:

$$\gamma_X(k) := \frac{\mathbb{E}[(X_t - \mathbb{E}[X_t])(X_{t+k} - \mathbb{E}[X_{t+k}])]}{\mathbb{E}|X_t - \mathbb{E}[X_t]|^2}. \quad (2.1)$$

We define the types of memory based on the tail-behaviour of the auto-correlation function. There is no universally agreed definition for long memory.

Definition 2.2.1. *A random process X_t with auto-correlation function $\gamma_X(k)$ has*

1. *long memory*, if $\sum \gamma_X(k) = \infty$;
2. *moderate memory*, if $\sum |\gamma_X(k)| < \infty$ and $\lim_k \gamma_X(k)k^l = c(l) > 0$ for some $l > 1$;
3. *short memory*, if $\sum |\gamma_X(k)| < \infty$ and $\lim_k \gamma_X(k)k^l = 0$ for every $l > 1$.

Processes with moderate memory are important because their autocorrelation decays as a power law function and not exponentially, therefore its effect is stronger. In some definition of short/long memory, moderate memory is regarded as short memory. In several time-series models ergodicity can only be proven in the case of moderate memory but not for the long one (see for example [36]).

2.3 Conditionally Gaussian Model Class

As we mentioned earlier, our aim is to build algorithms for parametric models. To do this, we define first a very wide class of stock price models. As we will see, this model class contains processes with statistically realistic properties suitable for creating log-optimal portfolios. The log-return process H_t is defined for $t \in \mathbb{Z}$.

Definition 2.3.1. *Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathbb{Z}}, \mathbb{P})$ be a probability space, where the σ -algebra \mathcal{F}_t is generated by $\{\varepsilon_j\}_{j \leq t}$ and $\{\eta_j\}_{j \leq t}$ where $\{\varepsilon_j\}_{j \in \mathbb{Z}}$ and $\{\eta_j\}_{j \in \mathbb{N}}$ are two i.i.d. sequences that are independent. We assume that $\varepsilon_j, j \in \mathbb{Z}$ are standard Gaussian.*

*The log-return process is called **Conditionally Gaussian** is it evolves in time*

$$H_t = F(Z_{t-1}, Y_t, \varepsilon_t, \eta_t), \quad (2.2)$$

where F is a measurable function, Y_t is a stationary Gaussian process such that Y_t is $\sigma(\varepsilon_j, j \leq t)$ -measurable.

Here Z_{t-1} contains the information about the past values of the log-return, so it is assumed that $Z_{t-1} = j(H_{t-1}, H_{t-2}, \dots)$ for some measurable $j : \mathbb{R}^{\mathbb{N}} \rightarrow \mathbb{B}$ for some Banach space \mathbb{B} . For instance, in the simplest case, we have $Z_{t-1} = H_{t-1}$ and $\mathbb{B} = \mathbb{R}$.

The process Y_t represents some economic factor, e.g. the *log-volatility*. In the following section we assume that Y_t is a causal linear process, i.e.

$$Y_t = \sum_{j=0}^{\infty} \beta_j \varepsilon_{t-j},$$

with some coefficients β_j satisfying $\sum_{j=0}^{\infty} \beta_j^2 < \infty$. We also assume $\beta_0 \neq 0$.

Notice that the random process Y_t , $t \in \mathbb{N}$ is also *conditionally Gaussian*, since the conditional distribution of Y_t on the σ -algebra \mathcal{F}_{t-1} is Gaussian. We remind the reader that the conditional law of Y_t is determined by the conditional expectation and the conditional variance processes:

$$\nu_{t-1} := \mathbb{E}[Y_t | \mathcal{F}_{t-1}] \quad \text{and} \quad \kappa_{t-1}^2 := \text{Var}(Y_t | \mathcal{F}_{t-1}), \quad t \in \mathbb{Z}. \quad (2.3)$$

Remark 2.3.1. It will become clear that for our arguments below it suffices to assume that Y_t is a conditionally Gaussian process. However, we stay with the simple linear case for the sake of a transparent presentation. Clearly, for the case of a linear process $\kappa_t = \kappa = \beta_0^2$ constant. We nevertheless keep the notation κ_t for a while to indicate that more general processes could be similarly treated.

As we can see, both variables (Z_{t-1} and Y_t) can bring about a memory effect (i.e. they can have memory, unlike a Markovian process).

For us, the main property of the models considered here is that they allow a non-Markovian structure, indicated many times by statistical investigations.

Such a model is general enough to describe important phenomena, such as leverage effect, correlation between volatility and driving noise, anti-persistence, the fit of the kurtosis and skewness of the returns etc. In the next two sections we are going to present two examples where these characteristics are discussed. Sampling from certain famous continuous-time models (see e.g. [9]) also leads to discrete-time models of the type treated here. Our main proposition of model is the Discrete Gaussian Stochastic Volatility (DGSV) dynamics which easily incorporates these features but we present another very useful dynamics, a type of heteroscedasticity model, Bilinear ARCH.

All the models in the following are part of the Conditionally Gaussian Dynamics model class. We start the description at the most basic models and proceed to the more elaborate ones: DGSV and BARCH. These two models are studied in the following sections in detail.

2.4 AR(1), MA(∞) and ARMA(p,q)

We only present these model briefly, firstly because they are well-known, secondly, they only serve as benchmark models, since none of them has long memory. AR(1) and MA(∞) processes play an important role in the case of learning algorithms. These algorithms are easier to understand if we investigate them on a simpler processes. While the ARMA(p,q) process will serve as a basic version in Chapter 6 of a more elaborated process, called ARFIMA(p,d,q).

Definition 2.4.1. *The process H_t is an Autoregressive (1) process, shortly AR(1) if*

$$H_t = \mu + \alpha H_{t-1} + \sigma \eta_t, \quad t \in \mathbb{N}. \quad (2.4)$$

The constants $\mu \in \mathbb{R}$, $\alpha \in (-1, 1)$ is the autoregressive or mean-reverting coefficient, while $\sigma > 0$ is the volatility. The noise η_t is an i.i.d. sequence.

This model is Markovian, it has no long memory, but the role of these constants will be similar in the following models.

It has stationary solution. The expected value and the variance of this process are

$$\mathbb{E}[H_t] = \frac{\mu}{1 - \alpha}, \quad \text{Var}(H_t) = \frac{\sigma^2}{1 - \alpha^2}. \quad (2.5)$$

If the driven noise η_t is Gaussian then H_t is Gaussian as well.

Definition 2.4.2. *The process H_t is a Moving Average process, shortly MA(∞) if*

$$H_t = \mu + \sum_{j=0}^{\infty} \beta_j \eta_{t-j}, \quad t \in \mathbb{N}. \quad (2.6)$$

The i. i. d. sequence $(\eta_j, j \leq t)$ is the driving noise of the process.

This process have long memory but in that case it also has long memory for both H_t and for $|H_t|$. The expected value and the variance of this process are

$$\mathbb{E}[H_t] = \mu, \quad \text{Var}(H_t) = \sum_{j=0}^{\infty} \beta_j^2 / 2. \quad (2.7)$$

Definition 2.4.3. *The process H_t is a Autoregressive Moving Average process, shortly ARMA(p, q) if*

$$H_t = \sum_{j=1}^p \phi_j H_{t-j} + \sum_{j=0}^q \beta_j \eta_{t-j}, \quad t \in \mathbb{N}, \quad (2.8)$$

for some coefficients $\phi_j, \beta_j, j \in \mathbb{N}$.

2.5 Discrete Gaussian Stochastic Volatility model

Discrete Gaussian Stochastic Volatility (DGSV) was conceived to have similar behaviour to the continuous-time Fractional Stochastic Volatility model (see [8] and [9]) ; at the same time, it is also element of the Conditionally Gaussian model class. For example, the log-volatility process is described as a stationary Gaussian process and has long memory. Furthermore,

the noise is also correlated with the volatility. The dynamics of the model is made up of two parts: the drift, which is a linear function of the past log-return H_{t-1} and the volatility.

Definition 2.5.1. *The process Y_t is called log-volatility. The dynamics of the Discrete Gaussian Stochastic Volatility model is:*

$$H_t = \mu + \alpha H_{t-1} + \sigma e^{Y_t} \left(\rho \varepsilon_t + \sqrt{1 - \rho^2} \eta_t \right), \quad \mu \in \mathbb{R}; \quad (2.9a)$$

$$Y_t = \sum_{j=0}^{\infty} \beta_j \varepsilon_{t-j}, \quad \beta_j \in \mathbb{R}. \quad (2.9b)$$

The constant $\rho \in [-1, 1]$ is the coefficient of the leverage effect, which causes correlation between the driving noises of the log-price and the volatility. We assume the mean-reverting coefficient $|\alpha| < 1$ and the volatility coefficient $\sigma > 0$. The two noise sequence ε_t and η_t are the same as in Definition 2.3.1. The β_j coefficients are such that $\sum_{j=0}^{\infty} \beta_j^2 < \infty$ and $\beta_0 \neq 0$.

The noise sequence $(\varepsilon_j, j \leq t)$ must be Gaussian if we wish the DGSV model be an element of the Conditional Gaussian class. The sequence $(\eta_j, j \leq t)$ is not necessarily Gaussian, but its distribution must be known in order to calculate the optimal strategy function $\hat{\pi}$ in Section 3.2

The driving noise in (2.9a) is constructed in this way in order to have a correlation between the log-volatility Y_t and the driving noise, hence resulting in leverage effect. The general form of the log-volatility describes a rather wide range of possible processes. In the numerical investigations we restrict its parameters to

$$\beta_j = b_0(1 + j)^{-b}, \quad j \in \mathbb{N}. \quad (2.10)$$

This simple form allow us to characterize the decay of the memory by one parameter, b . Long memory case is when b is between $(0.5, 1)$. If $b > 1$, it corresponds to the moderate memory case. The other constant b_0 also plays an important role: we will see that it does not only effect the strength of the memory but also the conditional skewness and kurtosis of the distribution of H_t . The condition that $b > 0.5$ is necessary to satisfy the condition $\sum_{j=0}^{\infty} \beta_j^2 < \infty$.

Remark 2.5.1. *An important property about this model, is that it can have long memory and can be ergodic (in some sense) at the same time. In the next chapter we will see how we are able to create a log-optimal portfolio with this dynamics. The proof of the ergodic-like property can be found in the Appendix A.*

Besides the other advantageous properties that we formerly mentioned, this model is also able to capture basic moment-related properties. In most stock price models with long

memory effect it is usually impossible to calculate the (conditional/unconditional) moments, while in our model it can be done, in addition, moments can be calculated explicitly.

2.5.1 Conditionally Gaussian

Within the Conditionally Gaussian class, the most important example we suggest is the DGSV model, therefore it is crucial to demonstrate their relationship.

Inside the Discrete Gaussian Stochastic Volatility model class we mostly consider the DGSV process.

We defined the log-return as $H_t = F(H_{t-1}, Y_t, \varepsilon_t, \eta_t)$ with $F(h, y, \varepsilon, \eta) = \mu + \alpha h + \sigma \exp(y)(\rho\varepsilon + \sqrt{1 - \rho^2}\eta)$. The conditional mean and variance of the log-volatility Y_t are

$$\nu_{t-1} = \mathbb{E}[Y_t | \mathcal{F}_{t-1}] = b_0 \sum_{j=1}^{\infty} (1+j)^{-b} \varepsilon_{t-j}, \quad (2.11a)$$

$$\kappa_{t-1}^2 = \text{Var}(Y_t | \mathcal{F}_{t-1}) = b_0^2. \quad (2.11b)$$

In order to numerically calculate the log-optimal trading strategy, we need to know these variables. Of course, in a real investment situation, without simulated prices, one needs to estimate these variables.

2.5.2 Autocorrelation

The autocorrelations of the simulated DGSV processes are realistic in the sense that both the simulated autocorrelations (Figure 2.1) and the empirical ones (Figure 1.1) decay similarly for the absolute values of the log-return $|H_t|$; also, the autocorrelation of the log-returns H_t are insignificant for lag > 0 . We will work with two parametrizations which are different in their autoregressive coefficient and their strength of memory. In the first parametrization where the autoregressive coefficient is bigger (0.3 versus 0.05) the autocorrelation of H_t is slightly bigger, while the autocorrelation of the absolute return is weaker. The value of b_0 also differs in the two parametrizations, which is $b_0 = 1/3$ in the first case and $b_0 = 0.55$ in the second case. Nevertheless, both diagrams on the log-log plot reveal a power-law decay of the correlations. Just like with real data, the power-law relationship exceptionally good for the first 50 lags.

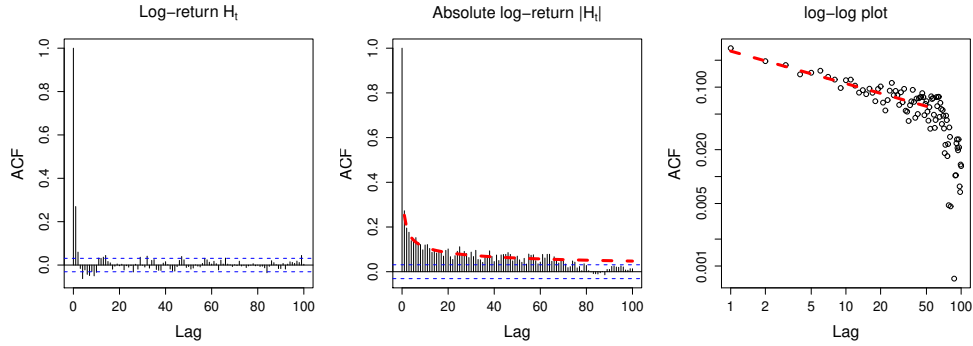
The parameters of the fitted lines are

$$\widehat{\gamma}_{|H|}(\tau) \approx 0.25 \tau^{-0.36} \quad (\text{DGSV-1 parametrization}) \quad (2.12a)$$


$$\widehat{\gamma}_{|H|}(\tau) \approx 0.34 \tau^{-0.39} \quad (\text{DGSV-2 parametrization}) \quad (2.12b)$$

Since the exponent in both cases are between -0.5 and 0 , they show the presence of long memory. There is a cut off in the autocorrelation on the log-log plot. It is clearly visible that around lag 50 the autocorrelation deviates from the fitted straight line. This phenomenon is due to the fact that there is only a finite amount of data.

Of course, because of the finiteness of the data, the autocorrelation function has a cut off (probably the same happening with real data: the length of the available historical return limit the power-law behavior).



(a) DGSV-1 parametrization.



(b) DGSV-2 parametrization

Figure 2.1: Autocorrelation of the Discrete Gaussian Stochastic Volatility model with parameters shown in Table 2.2 on page 28. DGSV-2 model has smaller α parameter which is responsible for the less significant autocorrelation of H_t and it has more significant autocorrelation for $|H_t|$ because the memory effect in the log-volatility is stronger (higher b_0 value, 0.55 instead of $1/3$) as Table 2.1 shows. Both DGSV models have long memory in the volatility process.

Other interesting question about the autocorrelation is how the parameter b influence the memory of the log-return. The long memory property of the log-volatility Y_t does not

b	0.6	0.7	0.8	0.9	1.0	1.1
exponent	-0.39	-0.46	-0.56	-0.69	-0.87	-0.98

Table 2.1: Relationship between parameter b of the DGSV model and the power-law decay of the autocorrelation function of $|H_t|$.

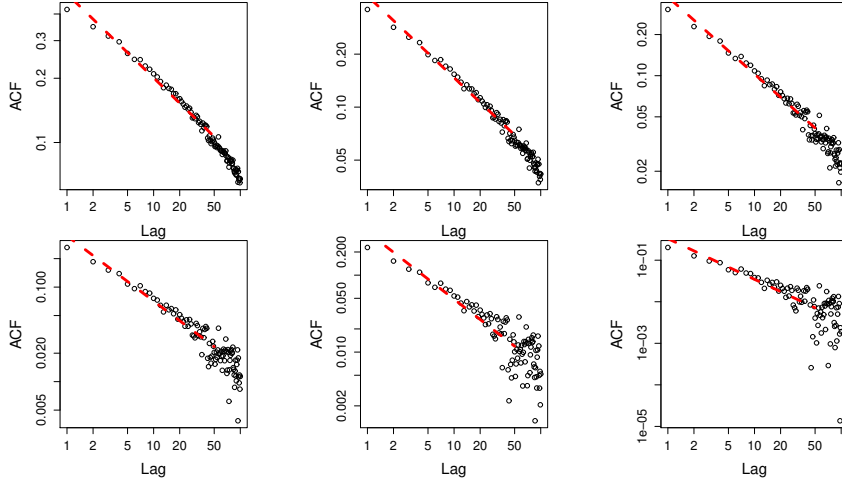


Figure 2.2: The effect of the memory parameter b of the DGSV model on the autocorrelation of $|H_t|$. From left to right, top to left the values of $b = 0.6, 0.7, \dots, 1.1$. The log-log plots suggest power-law decay in all cases, but the memory from $b = 0.8$ is only moderate.

necessarily imply that the log-return process H_t also has long memory. However, as Figure 2.1 shows, when $b = 0.6$ then the log-return process also has long memory.

Figure 2.2 shows how changes to b affect power-law decay. The result of the linear fit of the log-log plot, that is, the exponent of the power functions are displayed in Table 2.1. As the table shows, log-return does not have long memory on the whole $b \in (0.5, 1)$ interval. DGSV model has long memory only for small values of b , for example $b = 0.6$ or 0.7 .

In the simulation we used 500 realizations (generated paths) with $T = 5000$ timesteps and the numerical value of infinity in the sums was 1000. (These figures should not be directly compared with the autocorrelation of single stocks. This is an aggregated result, using multiple realizations. Moreover, the time horizon $T = 5000$ means a 20 year long daily data.)

2.5.3 Mean and Variance

Interestingly, both the conditional and unconditional moments of the marginal distributions of a DGSV model can be calculated. Moreover, we can do this explicitly. In most of the long memory models this is impossible. Unfortunately, even for the second moment, their

formulas become very complex.

These moments are important because in stock price models they are generally impossible to calculate. The conditional moments in the DGSV model give a very good approximation of the log-optimal strategy in the numerical results, as we will see in Section 4.4.

To calculate the k -th moment it is enough to know $\mathbb{E}[e^{K\varepsilon\varepsilon^j}]$ for every positive integer $j \leq k$, $K \in \mathbb{R}$, where $\varepsilon \sim \mathcal{N}(0, 1)$. By a simple substitution of the variable in the integral of the expression of expected value, we get that $\mathbb{E}[e^{K\varepsilon\varepsilon^j}] = e^{K^2/2}\mathbb{E}[(\varepsilon + K)^j]$, which can be calculated explicitly. E.g. for $j = 1$ it is equal to $e^{K^2/2}K$, for $j = 2$ it is equal to $e^{K^2/2}(1 + K^2)$. The calculations are elementary but long; derivation of the moments' formulas can be found in Appendix B.

The *conditional mean* and the *mean* of the DGSV model are:

$$E[H_t|\mathcal{F}_{t-1}] = \mu + \alpha H_{t-1} + \rho\sigma\kappa e^{\nu_{t-1} + \kappa^2/2}, \quad (2.13a)$$

$$E[H_t] = \frac{\mu + b_0\rho\sigma e^{\sum_{j=0}^{\infty} \beta_j^2/2}}{1 - \alpha}. \quad (2.13b)$$

Here $\kappa = \kappa_{t-1} = \beta_0 = b_0$ but since it is a constant, we do not need the time parameter in the notation.

In the conditional moments we prefer to use the ν, κ variables to emphasize the conditions, while in the unconditional moments we use the β_j description to incorporate the memory. We remind the reader that $\nu_{t-1} = \sum_{j=1}^{\infty} \beta_j \varepsilon_{t-j} = b_0 \sum_{j=1}^{\infty} (1+j)^{-b} \varepsilon_{t-j}$.

The last term of both equations, which contains the exponential term and ρ , are the consequences of the leverage effect. That is, if ρ is not 0 then there is a relationship between the (conditional) expected value of the log-return H_t and the log-volatility Y_t via the β_j parameters. Moreover, as is often observed, the log-return decreases if the volatility increases, which is $\rho < 0$ in our case.

To calculate any conditional moment, the crucial point is to calculate the value of $\mathbb{E}[e^{Y_t}|\mathcal{F}_{t-1}]$, which is $\kappa e^{\nu_{t-1} + \kappa^2/2}$. The equation of the unconditional moment is achieved by taking the expected value of both sides in the conditional moment in (2.13a). Of course, this only makes sense if the log-return process H_t is stationary. The existence of the stationary solution is not self-evident and mathematically challenging to prove. The proof can be found in Appendix A.

The values of the *conditional variance* and the *variance* are:

$$D^2(H_t|\mathcal{F}_{t-1}) = \sigma^2 e^{2\nu_{t-1} + \kappa^2} \left(e^{\kappa^2} - \kappa^2 \rho^2 + 4\kappa^2 \rho^2 e^{\kappa^2} \right), \quad (2.14a)$$

$$D^2(H_t) = \frac{\mu^2 + \sigma^2 e^{2\sum_0 \beta_j^2} (1 + 4b_0^2 \rho^2) + 2(\mu\alpha M + R(\alpha M + \mu))}{1 - \alpha^2} - M^2, \quad (2.14b)$$

where $M := E[H_t]$ and $R := b_0 \rho \sigma e^{\sum_0 \beta_j^2/2}$.

The value of the *conditional second moment* is:

$$E[H_t^2 | \mathcal{F}_{t-1}] = \mu^2 + \alpha^2 H_{t-1}^2 + \sigma^2 (1 + 4\rho^2 \kappa^2) e^{2\nu_{t-1} + 2\kappa^2} + 2 \left(\mu \alpha H_{t-1} + \kappa \mu \rho \sigma e^{\nu_{t-1} + \kappa^2/2} + \alpha H_{t-1} \kappa \rho \sigma e^{\nu_{t-1} + \kappa^2/2} \right). \quad (2.15)$$

Remark 2.5.2. *The (un)conditional first two moments are used to determine approximate strategies in the next chapter, therefore they play a crucial role.*

2.5.4 Skewness and Kurtosis

While it is possible to calculate the (un)conditional skewness and kurtosis, due to their complexity we do not present them here. Instead of investigating the higher moments of H_t , we use an approximating log-return process \bar{H}_t , defined as

$$\bar{H}_t = \sigma e^{b_0 \varepsilon_t} (\rho \varepsilon_t + \sqrt{1 - \rho^2} \eta_t). \quad (2.16)$$

We see that with this, log-return could be defined as $H_t = \mu + \alpha H_{t-1} + e^{\nu_{t-1}} \bar{H}_t$. The skewness is not linear, therefore we cannot say that the skewness of H_t equals to the sum skewness of the autoregressive part $\mu + \alpha H_{t-1}$ and $e^{\nu_{t-1}} \bar{H}_t$, but our assumption is that the autoregressive part is of a smaller order than the volatility part \bar{H}_t . Otherwise the autocorrelation of the log-return would look different. In this way we are able to approximate the skewness and kurtosis of H_t by the simpler formula we derived from \bar{H}_t .

The reason why we want to give a more or less simple, but approximative formula to calculate the skewness and kurtosis is that it helps us to find a good parameter choice that results in realistic simulated path. In the next section we show, how it can help.

Practically speaking, we approximate the skewness and kurtosis of H_t with \bar{H}_{t-1} . that is, when autoregressive effect is negligible and we replace ν_{t-1} with its expected value, i. e. 0, which results in \bar{H}_{t-1} .

The conditional skewness and kurtosis of \bar{H}_t

$$\text{skew}_{t-1}(H_t) = \mathbb{E} \left[\left(\frac{e^{\kappa \varepsilon - \kappa^2/2} (\rho \varepsilon + \sqrt{1 - \rho^2} \eta) - \kappa \rho}{\sqrt{e^{\kappa^2} - \kappa^2 \rho^2 + 4\kappa^2 \rho^2 e^{\kappa^2}}} \right)^3 \right], \quad (2.17a)$$

$$\text{kurt}_{t-1}(H_t) = \mathbb{E} \left[\left(\frac{e^{\kappa \varepsilon - \kappa^2/2} (\rho \varepsilon + \sqrt{1 - \rho^2} \eta) - \kappa \rho}{\sqrt{e^{\kappa^2} - \kappa^2 \rho^2 + 4\kappa^2 \rho^2 e^{\kappa^2}}} \right)^4 \right]. \quad (2.17b)$$

Remark 2.5.3. *(2.17a) and (2.17b) show that conditional skewness and kurtosis only de-*

pend on b_0 (that is, κ) and ρ .

Figure 2.3 shows the numerically calculated values of the conditional skewness and kurtosis as function of ρ and b_0 . Skewness is a monotone function of ρ and it is 0 for $\rho = 0$ (this would be an AR(1) process); real data has negative skewness, therefore ρ should be less than zero. Kurtosis is an even function of ρ and increases in b_0 .

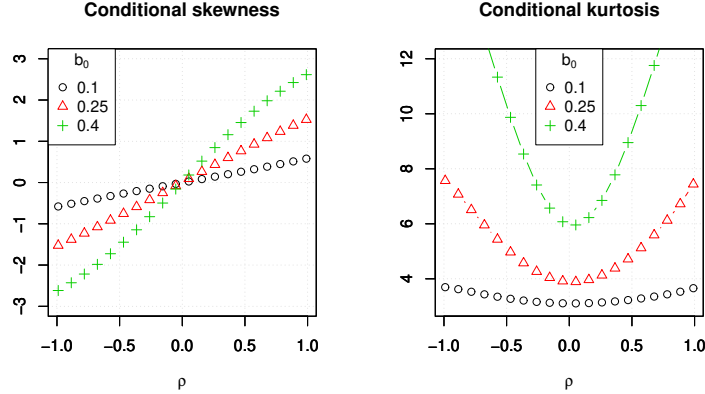


Figure 2.3: Conditional skewness and kurtosis of log-return in the DGSV model with parameters set to $\sigma = 1, \nu = 0$. Values around $\rho \approx -0.4$ and $b_0 \approx 0.25$ give similar skewness and kurtosis to those that can be observed in typical stock returns.

2.5.5 Parameter Choice

The log-optimal strategy we give in the next chapter is analysed by using of simulated data. In order to make the numerical simulations reasonable, we have to find the model parameters that result in realistic log-return data. Of course, a statistical estimation approach is impossible, the DGSV model is too complex. Fitting the parameters would require a whole study in itself and it is out of scope of this thesis.

Therefore we use a heuristic argument for our choice of parameters in order to make our choice clear. We are using the above mentioned moments as a help.

The only parameter that does not need to reflect realistic property is μ . In our case we assume that the interest rate $r = 0$. Its effect on the expected value is clear, it add a constant $\mu/(1 - \alpha)$ to the expected value.

The value of α is very small, but usually believed to be positive. Therefore we investigate the two cases when $\alpha = 0.3$ or 0.05 . Both cases result in a negligible autocorrelation of H_t . Its value is not crucial for the present thesis.

The value of ρ and b_0 is more interesting. In the continuous Fractional Stochastic Volatility models the average parameter is thought to be around -0.4 or -0.2 . As we mentioned,

we use the conditional skewness and kurtosis values to set the values of ρ and b_0 , since these conditional moments of \bar{H}_t a good approximation of the unconditional moments of H_t .

By looking at Figure 2.3, the values $(\rho, b_0) = (-0.4, 0.25)$ on both diagrams result in a realistic skewness (cca -0.6) and kurtosis (cca 5). If we choose a higher values, for example $(\rho, b_0) = (-0.15, 0.4)$ then the conditional skewness and kurtosis are around -0.6 and 7.

The most important parameter is obviously b that causes the long memory effect. The process Y_t as well as H_t has long memory if $b < 1$. The smaller its value, the stronger the memory effect, therefore we use $b = 0.6$.

Table 2.2 contains the chosen parameters for two different parametrization of the DGSV model: one with slightly stronger autoregressive part and the other one with slightly stronger memory. The table also contains the statistics of one-one simulated pahts from each model.

DGSV-1									
Params	μ	α	b	b_0	σ	ρ	Stat.	DGSV-1	DGSV-2
Value	1/550	0.3	0.6	1/3	0.01	-0.4	Mean	-0.00023	0.00006
DGSV-2							St. dev.	0.017	0.025
Params	μ	α	b	b_0	σ	ρ	Skew.	-1.39	-0.68
Value	1/550	0.05	0.6	0.55	0.008	-0.15	Kurt.	10.0	28.2
BARCH-1 – moderate memory							BARCH-1 BARCH-2		
Params	μ	α	b	b_0	a		Mean	0.00022	0.00025
Value	10^{-5}	0.3	1.7	-0.55	1/120		St. dev.	0.011	0.016
BARCH-2 – long memory							Skew.	-0.66	0.05
Params	μ	α	b	b_0	a		Kurt.	19.3	9.3
Value	10^{-5}	0.3	0.6	-0.25	1/60				

Table 2.2: Simulation parameters (left) and the statistical results of one realization (right) in the two models DGSV and BARCH.

2.5.6 Histogram

The histogram of the DGSV based on simulated data is very similar to the ones we can see with real data, see results on the NYSE data 1.3. This similarity is a consequence of the fact that the DGSV model can have very similar moments like the real data have.

A more difficult question is how the histogram changes for different memory. Since we are not able to calculate the distribution function of the DGSV model, we have to use Monte Carlo simulation. Figure 2.4 shows three histograms from the DGSV-2 parametrization with the different b values. The deviation from each other is so small that on printed pages it is




Figure 2.4: Histogram of the DGSV model with 3 different parameter settings. DGSV-2 parametrization, with $b = 0.6, 0.8, 1$.

barely visible. The only considerable difference is between the kurtosis of the distribution. Higher b value results in smaller kurtosis.

The numerical results regarding statistical data are very mutable, therefore we used 50 different realizations and each realization was iterated for 5000 times and infinity in the indices was 1000. We used longer time series instead of more realizations to be certain that the model reached stationarity. (Since the stationary distribution is unknown, we are not able to start the process from a stationary distribution.)

2.6 Bilinear ARCH models

Log-returns are more easily accessible than volatilities—this is one of the main reasons why conditional heteroscedasticity models became popular. In these models, to forecast the unknown return of the price, investors only have to know past data of the returns and estimate model parameters. The most representative feature of stock prices, mentioned in Section 1.3, 1.4 and 1.5 is that although log-return has small auto-correlation, its absolute value or its powers do show long memory. While ARCH or GARCH cannot capture long memory, Linear ARCH and Bilinear ARCH processes can. Not only long memory is captured but anti-persistence as well.

Recalling the observation that memory is carried within the volatility and not in the drift

([9]), we apply some restrictions on the original definition of the Bilinear ARCH process from [36]. For example, the original model has an $\text{AR}(\infty)$ part, which is replaced in our case by an $\text{AR}(1)$ dynamics – further autoregressive part would only complicate the autocorrelation function needlessly.

Definition 2.6.1. *The stochastic process H_t is a Bilinear ARCH process if*

$$H_t = \mu + \alpha H_{t-1} + \eta_t \sigma_t, \quad t \geq 0, \quad (2.18a)$$

$$\sigma_t = a + \sum_{j=1}^{\infty} \beta_j H_{t-j}, \quad (2.18b)$$

where μ, α, a, β_j are real parameters.

The constants μ and α are responsible for the autoregressive part. The term σ_t is the *conditional variance*, while $\mu + \alpha H_{t-j}$ is the *conditional mean*. The white noise $\eta_t \sim \mathcal{N}(0, 1)$ are i.i.d. variables. The assumption on the distribution is not necessary to be Gaussian, but we have to know it to be able to find the optimal strategy.

The term $\mu + \alpha H_{t-j}$ is also called the *drift*, while σ_t is the *volatility*.

With the proper choice of these parameters both H_t and σ_t will be ergodic processes with long memory. These choices keep the conditional mean part simple. We imagine α is small, because larger α values result in a stronger autocorrelation of the log-return, which should be negligible. Nevertheless, it is more difficult to find a proper parametrization of the models if α has a low value. For these reasons we concluded we should use $\alpha = 0.3$ in the numerical simulations (we indicate where we change this setting).

In the notation of Section 2.3 in (2.2) we can ignore the noise ε as well as Y_t , so $H_t = F(Z_{t-1}, \eta_t)$, where $Z_{t-1} = (H_{t-1}, H_{t-2}, \dots)$ is an infinite length vector of past log-returns. The function is $F(z, \eta) = \mu + \alpha z_1 + (a + \beta^T z)\eta$, where $\beta = b_0(2^{-b}, 3^{-b}, 4^{-b}, \dots)$ and z are the infinite vectors and z_1 is the first component of z .

The model has stationary solution and ergodicity can be proven for a set of parameters, see Appendix A. Figure 2.5 shows for which parameters the BARCH model is ergodic. When $b < 1$, the model is not ergodic.

2.6.1 Autocorrelation

There are two parametrizations in Table 2.2 for the BARCH model. The first one has moderate memory, while the other one has long memory. Their autocorrelations can be seen on Figure 2.6.

Ergodicity criteria in BARCH model,




Figure 2.5: Analysing the condition in (A.5). The left figure shows the function $b \rightarrow \frac{b_0(\zeta(b)-1)}{1-\alpha}$ at different b_0 values for $\alpha = 0.3$, the model is ergodic when this function is < 1 . We see that if there is stronger memory, i.e. larger b_0 or smaller b values, ergodicity may fail. The right figure shows a region on the plane of the memory parameters (b, b_0) where the model is ergodic for a given α . Outside of the ergodic region or when $b < 1$ (long memory) ergodicity has not been proven, although the opposite has not been proven either. Based on numerical simulations, it is likely that ergodicity fails.

The parameters of the fitted lines are

$$\widehat{\gamma}_{|H|}(\tau) \approx 0.16 \tau^{-1.1} \quad (\text{BARCH-1 parametrization}) \quad (2.19a)$$


$$\widehat{\gamma}_{|H|}(\tau) \approx 0.32 \tau^{-0.39} \quad (\text{BARCH-2 parametrization}) \quad (2.19b)$$

Indeed, the first one has moderate memory, while the second one has long memory.

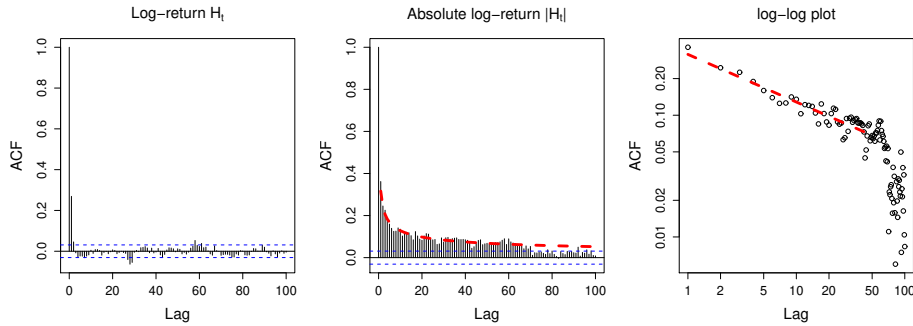
2.7 Some Remarks on the above Models

Based on simulations the role of the drift and volatility is clear: the larger the mean-reverting coefficient α , the more significant the autocorrelation of the log-return H_t , while the strength of the memory (higher b_0 or lower b) is related to the significance of the autocorrelation of the absolute log-return $|H_t|$. To put it in a simple way: drift causes autocorrelation of H_t , while volatility causes the autocorrelation of $|H_t|$. Since we are looking for stock price models which describe memory effect well, we have to choose models with lower α and higher b_0 or lower b .

In the introduction of the models we went into detail only about the DGSV model, which, we think, is an ideal model of real log-returns with long memory effect. Though, both long



(a) BARCH-1 parametrization.



(b) BARCH-2 parametrization

Figure 2.6: Comparison of the moderate and long memory BARCH process with exponent $b = 1.1$ and $b = 0.65$ with their autocorrelation function of log-returns and their absolute value in the Bilinear ARCH model. Parameters are shown in Table 2.2. BARCH-1 has moderate memory, BARCH-2 has long memory. The red dashed lines are the power function fit on the first 50 Lag of the data.

memory models performed well regarding their autocorrelation structure, we could achieve a better fit on the moments for DGSV. These models are very complex to fit their parameters statistically. Regarding the statistics of the generated data in Table 2.2, they are very similar to the statistics of the 1250 stocks from New York Stock Exchange in Table 1.1.

Based on the moments and the autocorrelation functions, we can say that the heuristic fit is sufficient to produce a realistic stock price; on this, we can apply numerical algorithms.

Remark 2.7.1. *The most important difference between the two models is that **ergodicity-like attribution can be proved for the DGSV model, even in the long memory case.** Unlike DGSV, the **BARCH model is not ergodic in the long memory case** and only ergodic for a certain range of the parameters. Without ergodicity or a similar property Assumption 3.2.3 is not satisfied. Therefore we cannot guarantee theoretically the convergence of the log-optimal portfolio only with numerical simulations.*

In both models, we gave two different parametrization: one with a weaker and one with a stronger memory effect. In the DGSV model the autoregressive coefficient (α) is stronger or weaker, while in the BARCH model the difference is in the β_j coefficients. The two parametrization help us to understand the role of the memory in the following chapters. Moreover, those model that have long memory, are the only models which could create realistic autocorrelations.

It is important to note that ergodicity and long/short memory properties do not coincide. In the BARCH model, long memory depends on whether b is smaller or bigger than 1. While its ergodicity is determined by (A.5).

Both models are able to capture the leverage effect, which we want to include in our model but we do not intend to investigate its role. The BARCH-2 parametrization failed to produce negative skewness. Possibly, it is due to the lack of a proper fit.

A small remark at the end: since we eliminate the impact of interest rate, that is, $r = 0$, the presented model settings have smaller expected value than real data. This could be easily amended by choosing bigger μ .

Chapter 3

Investment Strategies

This chapter contains the definitions of all the investment strategies we use in Chapter 4 and 5. Convergence is proved only for the log-optimal portfolio in Section 3.2. The rest of the strategies can be viewed as some kind of simplification of the log-optimal strategy based on some restriction, for example a reduced strategy class in threshold strategies in Section 3.6.

The definition of the investment strategy we use in the continuous time case for the Fractional Brownian Motion is described in Chapter 6.

3.1 Introduction

In what follows we are looking for a sequence of investment decisions that are optimal in some sense. We emphasize the difference between the sequence of optimal decisions $\{\pi_t^*\}_{t=0}^T$, where $T > 0$ can be infinity as well and between the function $\hat{\pi}(\cdot)$ that uses past information (i. e. \mathcal{F}_{t-1} -measurable variables) to construct the optimal decision.

We wish to optimize the wealth of the investor such that risk is taken into consideration. We recall the time-evolution of the wealth from (1.2), but using the log-return and that the interest rate $r = 0$, hence $B_t/B_{t-1} = 1$.

$$W_t^\pi = W_{t-1}^\pi \left((1 - \pi_{t-1}) + \pi_{t-1} e^{H_t} \right), \quad t \in \mathbb{N}. \quad (3.1)$$

The objective of the investment is to maximize the utility function by choosing the strategy π_t adequately. Typically two cases can be optimized in the log-optimal portfolio problem: (i) finite horizon problem, where we want to optimize the utility function at a fixed time $T > 0$; or (ii) the long-run problem, where we optimize the time-average of the log-wealth. The optimal strategy $\pi^*(T) := \{\pi_t^*(T)\}_{t=0}^{T-1}$ for the finite horizon problem (if

exists) satisfies

$$\mathbb{E} \left[\log(W_T^{\pi^*(T)}) \right] = \sup_{\pi} \mathbb{E} [\log(W_T^{\pi})], \quad (3.2)$$

where the supremum is over $[0, 1]$ -valued adapted processes $\pi = (\pi_0, \dots, \pi_{T-1})$. The long-run problem is defined similarly, the optimal strategy $\pi^* := \{\pi_t^*\}_{t=0}^{\infty}$ is such that

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\log(W_T^{\pi^*}) \right] = \max_{\pi} \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} [\log(W_T^{\pi})], \quad (3.3)$$

where $\pi_t, t \in \mathbb{N}$ ranges over adapted $[0, 1]$ -valued processes. The characterization of optimal strategy will be obtained thanks to the conditionally Gaussian dynamics we assumed about Y_t .

We will furthermore use ergodic arguments for the long-run case. For the price evolution of the stock, we assume that the log-return process is a Conditionally Gaussian process (2.2), $H_t = F(Z_{t-1}, Y_t, \varepsilon_t, \eta_t)$. In the log-optimal case, we prove that the algorithm is optimal indeed, but we also present other strategies that rely on some simplification of either the strategy or the objective function.

W. l. o. g. let the initial wealth of the investor $w_0 = 1$, thus

$$W_T^{\pi} = \prod_{t=1}^T \frac{W_t^{\pi}}{W_{t-1}^{\pi}} =: \prod_{t=1}^T \exp(G_t), \quad (3.4)$$

then we get that

$$\mathbb{E}[\log W_T^{\pi}] = \sum_{t=1}^T \mathbb{E}[G_t^{\pi}] = \sum_{t=1}^T \mathbb{E}[\mathbb{E}[G_t^{\pi} | \mathcal{F}_{t-1}]], \quad (3.5)$$

where $G_t^{\pi} := \log(W_t^{\pi}/W_{t-1}^{\pi})$ is the growth in the period t .

In order to optimize (3.3) we need to assure two things:

1. The optimal strategy π_t^* optimize the conditional probability $\mathbb{E}[G_t^{\pi} | \mathcal{F}_{t-1}]$, in π more explicitly:

$$\mathbb{E} \left[\log \left((1 - \pi)(1 + r) + \pi e^{H_t} \right) \middle| \mathcal{F}_{t-1} \right] \quad (3.6)$$

2. The optimal strategy exists and the objective function converges.

In the next section we show that for the Conditionally Gaussian model class it is tractable, then in the following sections we present the simplifications that perform well on simulated data and are also a base for learning algorithms.

3.2 Log-optimal Strategy

The log-optimal portfolio is also called *growth-optimal portfolio* because it is attained by optimizing the expected logarithm of the growth function at each investment decision. We introduce the *expected growth* function g^π :

$$g^\pi(z, \nu, \kappa) := \mathbb{E} \left[\log \left((1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} \right) \right] \quad (3.7)$$

Here ν and κ are real numbers, $\pi \in [0, 1]$, $z \in \mathbb{B}$; the random variables (ε_0, η_0) have the same distribution as (ε_t, η_t) , $t \geq 1$.¹

Notice that the logarithmic term in the expected value in (3.7) has the same distribution as the growth G_t conditionally on \mathcal{F}_{t-1}

$$G_t := \log \left(\frac{W_t}{W_{t-1}} \right) = \log \left((1 - \pi) + \pi e^{H_t} \right). \quad (3.8)$$

Though, in the following we prove the log-optimality for the case when $r \neq 0$, it could be omitted if we still assume that the interest rate is a constant. Let us denote by $\tilde{r} = \log(B_t/B_{t-1})$ the continuously compounded interest rate. In this case, the growth equals to

$$G_t = \tilde{r} + \log \left((1 - \pi) + \pi e^{H_t - \tilde{r}} \right). \quad (3.9)$$

This problem is the same as in the case of zero interest rate beside apart from a constant shift on the log return which would be $H_t - \tilde{r}$.

The following lemmas and assumptions are needed for the existence of the optimal strategy.

Lemma 3.2.1. *There exists a constant $c_1 > 0$ such that, for all $\pi \in [0, 1]$,*

$$\left| \log \left((1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} \right) \right| \leq c_1 (1 + |F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)|).$$

Proof. Clearly,

$$\begin{aligned} (1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} &\leq \\ 1 + r + e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} &\leq \\ 2e^{\max\{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0), \log(1+r)\}} &\leq \\ 2e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0) + \log(1+r)}. & \end{aligned}$$

¹Strictly speaking, κ sho range over $[0, \infty)$. However, the law of ε_0 being symmetric, one may equally well let κ run over \mathbb{R} .

If $\pi > 1/2$ then

$$(1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} \geq e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} / 2.$$

If $\pi \leq 1/2$ then

$$(1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} \geq (1 + r)/2.$$

The statement follows easily from these estimates. \square

From now on we concentrate on the case when $Z_{t-1} = H_{t-1} \in \mathbb{B} := \mathbb{R}$, but the following lemmas in the finite dimensional case $Z \in \mathbb{R}^k$ would be similar. We will briefly comment on the general case later.

Assumption 3.2.1. *Assume that F is continuous, $Z \in \mathbb{R}$ and for each $N > 0$,*

$$\sup_{z, \nu, \kappa \in [-N, N]^3} |F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)| \leq X(N)$$

with some integrable random variable $X(N)$.

Lemma 3.2.2. *Let Assumption 3.2.1 be in force. Then the function*

$$(z, \nu, \kappa, \pi) \rightarrow g^\pi(z, \nu, \kappa)$$

is continuous.

Proof. It is enough to prove continuity on every cube $[-N, N]^3$, $N > 0$. Now dominated convergence and Lemma 3.2.1 show the statement. \square

Lemma 3.2.3. *Under Assumption 3.2.1, there exists a measurable function $(z, \nu, \kappa) \rightarrow \hat{\pi}(z, \nu, \kappa) \in [0, 1]$ such that*

$$g^{\hat{\pi}(z, \nu, \kappa)}(z, \nu, \kappa) = \sup_{\pi \in [0, 1]} g^\pi(z, \nu, \kappa).$$

Proof. We will use Proposition D.5 of [37] with the choice $\psi(z, \nu, \kappa) := [0, 1]$ and let us define $v(z, \nu, \kappa, \pi) := g^\pi(z, \nu, \kappa)$. Note that ψ is trivially a compact-valued Borel-measurable multifunction and that v is continuous by Lemma 3.2.1 so the cited proposition indeed applies and we can conclude. \square

Lemma 3.2.3 guarantee the existence of the log-optimal function. We also refer to this function as *log-optimal selector function*, since for each argument it "selects" the optimal investment strategy.

The importance of the Conditionally Gaussian models is that with the help of the conditionally known parameters: Z_t, ν_t and κ_t , we are able to handle the problem similarly as if it was a Markovian decision problem.

The theorem below connects the log-optimal strategy (π_t^*) and the log-optimal selector function ($\hat{\pi}$). The optimal solution for the investment problem relies on a day-by-day (on-line) optimization procedure where the investor uses the information \mathcal{F}_{t-1} to achieve optimality at time t .

Theorem 3.2.2. *Let us regard a log-optimal portfolio problem defined in (3.2) with the dynamics given by (2.2). The following strategy is log-optimal, on finite as well as on infinite horizons:*

$$\pi_t^* = \hat{\pi}(Z_{t-1}, \nu_{t-1}, \kappa_{t-1}), \quad t = 1, 2, \dots, T. \quad (3.10)$$

Remark 3.2.1. *In order to use the above theorem, the investor needs the quantities Z_{t-1}, ν_{t-1} and κ_{t-1} but these are available information at $t - 1$. For calculating $\hat{\pi}$, the investor has to evaluate the function g^π at every value of $\pi \in [0, 1]$ but in practice it is enough to do this along a finite partition of the interval. Calculation of g^π is easy, since it is the expected value of a function of a two-dimensional random variable (ε_0, η_0) .*

Proof. We need to prove that $\mathbb{E}[\log(W_T^{\pi^*})] = \sup_{\pi} \mathbb{E}[\log(W_T^\pi)]$ for the strategy π^* defined in the statement of Theorem 3.2.2. W. l. o. g. let $w_0 := 1$. Using the identity

$$W_T^\pi = \prod_{t=1}^T \frac{W_t^\pi}{W_{t-1}^\pi} =: \prod_{t=1}^T \exp(G_t),$$

we get that

$$\mathbb{E}[\log W_T^\pi] = \sum_{t=1}^T \mathbb{E}[G_t^\pi] = \sum_{t=1}^T \mathbb{E}[\mathbb{E}[G_t^\pi | \mathcal{F}_{t-1}]].$$

Since the sum is finite it is enough to maximize each term in the sum individually, i.e.

$$\mathbb{E}[G_t^\pi | \mathcal{F}_{t-1}] \rightarrow \max.$$

Notice the identity

$$\text{Law}[G_t^\pi | \mathcal{F}_{t-1}] = \text{Law} \left[\log \left((1 - \pi)(1 + r) + \pi e^{F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)} \right) \right] \Big|_{z=Z_{t-1}, \nu=\nu_{t-1}, \kappa=\kappa_{t-1}}.$$

Hence it is enough to maximize the expected value of the right-hand side above which is exactly what $\hat{\pi}(Z_{t-1}, \nu_{t-1}, \kappa_{t-1})$ does. \square

As we can see, the theorem above only proves for finite T . The assumption below

ensures that the log-optimal strategy converges, that is, it is also the solution of the long-run problem. Moreover, in the finite case the optimality is true for the expected value, while the next **assumption ensures that π_t^* is log-optimal for almost every realization.**

Assumption 3.2.3. *Let Assumption 3.2.1 be valid. Denote*

$$Q_t := g^{\hat{\pi}(Z_t, \nu_t, \kappa_t)}(Z_t, \nu_t, \kappa_t), \quad t \in \mathbb{N}$$

where $\hat{\pi}$ is as in Lemma 3.2.3. The limit $\lim_{N \rightarrow \infty} (\sum_{i=0}^{N-1} \mathbb{E}Q_i)/N$ exists, the limit

$$\lim_{N \rightarrow \infty} \left(\sum_{i=0}^{N-1} Q_i \right) / N$$

also exists almost surely and the two limits are equal.

When we mention 'ergodicity-like' property, what we really mean is the fulfilment of Assumption 3.2.3.

Remark 3.2.2. *The (optimal) growth-rate G^* , that is,*

$$G^* := \lim_{N \rightarrow \infty} \left(\sum_{i=0}^{N-1} \mathbb{E}Q_i \right) / N. \quad (3.11)$$

Of course, in a case when H_t is stationary, we have $G^* = \mathbb{E}Q_i$.

3.3 Markovian strategy

The Markovian version of the algorithm would be where we optimize the expected growth conditionally on the previous value:

$$g^\pi(z) := \mathbb{E} \left[\log \left((1 - \pi)(1 + r) + \pi e^{H_t} \right) \mid H_{t-1} = z \right], \quad (3.12)$$

if the distribution of Y_t is known.

Of course, there is a wide range of possible algorithms to achieve optimality for the Markovian case. We do not wish to cover this topic further. The reason why we mention it here is that we are using a similar approach in Chapter 5 where the optimal solution is learned throughout the incoming stream of the data.

3.4 Approximative strategies

The log-optimal solution is achieved by optimizing the expected value in (3.7). To calculate the expected value with respect to ε_0 and η_0 , we need to numerically evaluate a two-dimensional integral for a given π . When there are N stocks – a case which we will not touch on in this paper – the numerical integral becomes $2N$ -dimensional. To overcome this difficulty it is desirable to approximate the log-optimal strategy with a sub-optimal one which is computationally more feasible. We will see that these approximations also help in understanding the log-optimal strategy. For simplicity, we choose the interest rate $r = 0$.

In the literature a second order approximation has been defined in [38] under the name *semi-log-optimal* strategy. In [39] and [40] the method was investigated numerically in a non-parametric model setting. We do not use this terminology since we will show another kind of approximation (first-order) as well, and it might be confusing as to what the semi-log-optimal portfolio refers to.

As we have seen, the key moment is to maximize the conditional expected value of the growth that is to find a constant $\pi \in [0, 1]$ such that

$$\mathbb{E}[\log(1 - \pi + \pi e^{H_t}) | \mathcal{F}_{t-1}] \quad (3.13)$$

reaches its maximum.

In the previous section we could guarantee the existence of the optimal strategy and also its convergence. To do that, in every instance we had to calculate (3.7), which contains a 2-dimensional Gaussian integral. Of course, this can be done without any problem, but if we use approximative strategies, we end up in simple function evaluations instead of numerical integrals which are calculated much faster.

In the next chapter we will see that the approximative strategies perform very well in the numerical simulations.

The basic idea is to approximate the the function $\log(\cdot)$ and $\exp(\cdot)$ by their Taylor polinomial in first or second order in the growth (3.8).

For example, in the linear approximation

$$\log(1 - \pi + \pi e^{H_t}) \approx \log(1 - \pi + \pi(1 + H_t)) = \log(1 - \pi H_t) \approx \pi H_t. \quad (3.14)$$

In the first step we approximated $\exp(x) \approx 1+x$ and in the last step we used that $\log(1-x) \approx x$. The second order approximation goes similarly:

$$\log(1 - \pi + \pi e^{H_t}) \approx \dots \approx \pi(H_t + \frac{1}{2}H_t^2) - \frac{1}{2}\pi^2 H_t^2. \quad (3.15)$$

Since the typical value of H_t is really small on a daily trades (in the 90% of the time this approximation has smaller than 1% error), these approximations are good.

The approximated growth functions in a general stock price model with log-return H_t have the following form:

$$g_{lin}^\pi := \mathbb{E}[\pi H_t | \mathcal{F}_{t-1}], \quad g_{quad}^\pi := \mathbb{E} \left[\pi \left(H_t + \frac{1}{2} H_t^2 \right) - \frac{1}{2} \pi^2 H_t^2 | \mathcal{F}_{t-1} \right]. \quad (3.16)$$

We call the first one as *linear approximation* and the second one as *quadratic approximation*. These expressions are polynomial functions of π , therefore their optima can be calculated easily. In the linear case the approximative optimal strategy is:

$$\bar{\pi}^{lin} = \begin{cases} 1, & \text{if } \mathbb{E}[H_t | \mathcal{F}_{t-1}] > 0 \\ 0, & \text{otherwise.} \end{cases} \quad (3.17)$$

In the quadratic case the approximative optimal strategy is:

$$\bar{\pi}^{quad} = \frac{\mathbb{E}[H_t | \mathcal{F}_{t-1}]}{\mathbb{E}[H_t^2 | \mathcal{F}_{t-1}]} - \frac{1}{2} \quad (3.18)$$

restricted on the $[0, 1]$ interval.

The restriction of the strategy on the $[0, 1]$ interval was arbitrary in the previous section. Any bounded interval could be adequate. Here, this is not generally true, since if the value of πH_t was big, that is, π was big, then the approximations would not work well.

Remark 3.4.1. *These conditional expectation formulas are not much easier to calculate than the original formula in (3.7). Moreover, these formulas are more general, since they also work on any model class where the conditional moments are easy to calculate.*

In the Conditionally Gaussian model class however, where $H_t = F(H_{t-1}, Y_t, \varepsilon_t, \eta_t)$, the conditional expectation can be reduced in the same way to an unconditional expectation. We have

$$\mathbb{E}[H_t | \mathcal{F}_{t-1}] = \mathbb{E}[F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)] \Big|_{z=H_{t-1}, \nu=\nu_{t-1}, \kappa=\kappa_{t-1}}, \quad (3.19)$$

which is, again, a two-dimensional integral in ε_0 and η_0 .

In specified parametric models, these formulas can become simpler. We will study the two models, DGSV and BARCH, in detail where these approximative strategies can be calculated as a simple function evaluation.

Remark 3.4.2. *In the linearly approximated log-optimal portfolio the optimal strategy is extreme, i.e. its value is 0 or 1, but there is no mixed strategy. The strategy in the Conditionally Gaussian model class is a function of h, ν and κ : $\bar{\pi}_t^{lin} = \bar{\pi}_t^{lin}(h, \nu, \kappa)$. For several*

models, as we will see later, we can define a 0–1 barrier function denoted as $\tilde{h}(\nu, \kappa)$ and approximate π_t^* by

$$\bar{\pi}_t^{lin} = \begin{cases} 1, & \text{if } H_{t-1} > \tilde{h}(\nu_{t-1}, \kappa_{t-1}), \\ 0, & \text{otherwise.} \end{cases} \quad (3.20)$$

In practice, the conditional expectation of a return can be easily calculated for most of the parametric models. This approximation works for other model class, not only for Conditionally Gaussian type process.

Example 3.4.1. (*AR(1)*) Since $\mathbb{E}[H_t | \mathcal{F}_{t-1}] = \mu + \alpha H_{t-1}$ if H_t is an *AR(1)* process (see Definition 2.4.1), the 0-1 barrier function is a constant:

$$\tilde{h} = -\frac{\mu}{\alpha}. \quad (3.21)$$

Example 3.4.2. (*DGSV*) The conditional mean has been shown in (2.13). Therefore, the 0-1 barrier function of a *DGSV* process is

$$\tilde{h}(\nu, \kappa) = -\frac{\mu + \rho\sigma\kappa_{t-1}e^{\nu+\kappa^2/2}}{\alpha}. \quad (3.22)$$

Remark 3.4.3. There is an interesting parallel between the strategy of the second order approximation and the log-optimal solution when stock price is a diffusion in continuous time. If the stock price in the diffusion case is described by $dS_t/S_t = \mu(S_t)dt + \sigma(S_t)dW_t$, $S_0 = s_0$, for some real-valued functions $\mu(\cdot)$ and $\sigma(\cdot)$ and a Brownian motion W_t , then the optimal solution is $\pi_t^{*,diffusion}$ equals $\mu(S_t)/\sigma(S_t)$ (here, it is not restricted between 0 and 1). See the optimal strategy in Subsection 1.2 at page 6 in [41]. It is very similar to the quadratic approximative solution in (3.19).

3.5 Approximative Markovian strategies

Similarly to the log-optimal strategies, the approximative strategies too can be used with the Markovian assumption. This model is not interesting per se, but we make some conclusion which will be helpful in establishing the learning method. To be discussed later.

Our aim is still to optimize the conditional expected growth with 1) the Markovian assumption about the information (Section 3.3) and 2) the approximation of the $\log(\cdot)$ and $\exp(\cdot)$ functions (Section 3.4). The conditional growth

$$\mathbb{E} [\log ((1 - \pi)(1 + r) + \pi \exp(H_t)) | \mathcal{F}_{t-1}]$$

with the Markovian condition (3.12) and the approximations (3.14) and (3.15) is:

$$g_{M,lin}^\pi(h) := \mathbb{E}[\pi H_t | H_{t-1} = h], \quad (3.23a)$$

$$g_{M,quad}^\pi(h) := \mathbb{E} \left[\pi (H_t + \frac{1}{2} H_t^2) - \frac{1}{2} \pi^2 H_t^2 | H_{t-1} = h \right]. \quad (3.23b)$$

The sub/superscript M refers to the Markovian condition. Of course, the approximative Markovian strategies are also very similar:

$$\pi_t^{M,lin} = 1, \quad \text{if } \mathbb{E}[H_t | H_{t-1}] > 0. \quad (3.24)$$

$$\pi_t^{M,quad} = \frac{\mathbb{E}[H_t | H_{t-1}]}{\mathbb{E}[H_t^2 | H_{t-1}]} - \frac{1}{2}. \quad (3.25)$$

3.6 Threshold Strategy

The approximative strategies are based on the simplification of the objective function. However, we can approach the problem from another point of view: what if we simplify the strategy class? Yet, our only assumption was about the strategy π_t is to be a $[0, 1]$ -valued adapted process, $t \in \mathbb{N}$.

We have seen that, the first order approximation simplifies somehow the strategy class: it takes values only 0 or 1, but not as a restriction.

$$\bar{\pi}^{lin} = \begin{cases} 1, & \text{if } \mathbb{E}[H_t | \mathcal{F}_{t-1}] > 0 \\ 0, & \text{otherwise} \end{cases}$$

or briefly $\pi_t = \mathbb{1}_{\{\mathbb{E}[H_t | \mathcal{F}_{t-1}] > 0\}}$. The indicator function $\mathbf{1}_{\{x > 0\}}$ is 1 if $x > 0$ and 0 otherwise.

We have also seen that in many cases there is a 0-1 barrier function, because the conditional expectation can be calculated easily for most of the models and we know when $\mathbb{E}[H_t | \mathcal{F}_{t-1}] > 0$ holds. In the Conditionally Gaussian model class, there is a two-variable function $\tilde{h}(\nu, \kappa)$ which separates the possible (h, ν, κ) space into two regions, $\tilde{h}(\nu, \kappa) > 0$, where the strategy is either 0 or 1.

The threshold strategy in Chapter 5 will be *learnt* as the data arrives. Therefore we formalize the strategy in a way, to be more applicable for data. Instead of using the σ -algebra notation, the conditional expectation is conditioned on a stochastic process X_{t-1} .

Threshold type strategy is where the strategy can only has two possible values based on a decision:

$$\pi_t = \mathbb{1}_{\{f(X_{t-1}) > 0\}}. \quad (3.26)$$

The variable X_t is an \mathcal{F}_t adapted stochastic process, symbolizes past data. In practice, it can be any variable that are easy to get from data. For example, it can be the average of the past three returns, or similar. It also can be an estimation of the volatility.

The conditional growth function in this case is

$$\tilde{g} = \mathbb{E}[\log(1 - \mathbb{1}_{\{f(X_{t-1}) > 0\}} + \mathbb{1}_{\{f(X_{t-1}) > 0\}} e^{H_t} | X_{t-1})]. \quad (3.27)$$

This function is still a random variable because it is a function of past data.

Of course, the choice of the event $\{f(X_{t-1}) > 0\}$ is somewhat arbitrary. It only makes sense if there is indeed a clear distinction between the region of $\pi = 0$ and $\pi = 1$.

With the indicator function, the growth function simplifies to

$$\tilde{g} = \mathbb{E} [H_t \mathbb{1}_{\{f(X_{t-1}) > 0\}} | X_{t-1}]. \quad (3.28)$$

For simplicity, the function $f()$ can be a linear function. However, the variable X_{t-1} still can be arbitrarily complex, hence information is in the choice of X_{t-1} . Then the investors can choose what information they would like to use.

With the linear $f()$ function the growth is

$$\tilde{g}(\theta) = \mathbb{E} [H_t \mathbb{1}_{\{X_{t-1} > \theta\}} | X_{t-1}], \quad \text{or} \quad \tilde{g}(\theta) = \mathbb{E} [H_t \mathbb{1}_{\{X_{t-1} < \theta\}} | X_{t-1}]. \quad (3.29)$$

Which relation sign should we use is depend on the correlation between X_{t-1} and H_t . For positively correlated case, one should use $>$ and for the negatively correlated case, the $<$.

The so-called Kiefer–Wolfowitz algorithm is to search for the maximum of an expected value depending on a parameter. Our goal in Chapter 5 is to find the maximum of the expected growth as a function of the threshold parameter θ

$$g(\theta) := \mathbb{E}[H_t \mathbb{1}_{\{X_{t-1} > \theta\}}], \quad (3.30)$$

similarly in the negatively correlated case

$$g(\theta) := \mathbb{E}[H_t \mathbb{1}_{\{X_{t-1} < \theta\}}]. \quad (3.31)$$

Let us define the *threshold strategy function* $\tilde{\pi}(x, \theta)$ such that

$$\tilde{\pi}(x, \theta) := \mathbb{1}_{\{x > \theta\}}. \quad (3.32)$$

Then the optimal threshold strategy is

$$\pi_t^* = \tilde{\pi}(H_{t-1}, \theta^*), \quad (3.33)$$

where θ^* is defined as

$$\theta^* := \arg \max_{\theta \in \mathbb{R}} g(\theta). \quad (3.34)$$

The following theorem states when the optimal θ exists.

Proposition 3.6.1. *Let the log-return process H_t defined on the probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathbb{N}}, \mathbb{P})$ and let the X_t adapted on \mathcal{F}_t . Assume, there is only one root of the differentiable function $\phi(x) := \mathbf{E}[H_t | X_{t-1} = x]$ and that $\phi(x) > 0$ if $x > 0$. If H_t is stationary, then the root of $\phi(x)$ is the unique optimal threshold:*

$$\theta^* = \{x | \phi(x) = 0\}. \quad (3.35)$$

Proof. For the sake of simplicity assume that X_{t-1} has a pdf. The conditional expectation of the growth is

$$g(\theta) = \mathbb{E}[H_t \mathbb{1}_{\{X_{t-1} > \theta\}}] = \mathbb{E}[\mathbb{E}[H_t | X_{t-1}] \mathbb{1}_{\{X_{t-1} > \theta\}}].$$

Since $\mathbb{E}[H_t | X_{t-1}]$ is a function of X_{t-1} , call it $v(X_{t-1})$ and denote the pdf of X_{t-1} as $f_X(x)$, the expected value is

$$\int_{\theta}^{\infty} v(y) f_X(y) dy.$$

The integral has optimum where

$$-v(y) f_X(y) = 0.$$

Since $f_X(y)$ is non-negative therefore the optimal threshold is where $v(y) = 0$ which conclude our statement. □

Remark 3.6.1. *The main message of the theorem is that only those information can be used in the optimization algorithm which are not mean-independent from the price process. The concept of mean independence is well-known in econometrics which is a stronger property than uncorrelation but weaker than the stochastic independence. For the definition of mean independence, see Section 18.2 of [42].*

Remark 3.6.2. *A conclusion of Proposition 3.6.1 is that the linear approximative strategy is log-optimal if the strategy can only be 0 or 1. This is only true in the univariate case.*

Example 3.6.1 (AR(1)). Let H_t defined as in Definition 2.4.1. The conditional expectation is

$$\mathbf{E}[H_t | H_{t-1} = x] = \mu + \alpha x.$$

Its root, that is, the optimal threshold is

$$\theta^* = -\frac{\mu}{\alpha}.$$

When $\alpha < 0$, then the assumption in Theorem 3.6.1 about $\phi(x) > 0$ if $x > 0$ is false, but the optimality is true if we change the inequality sign in (3.26) to $\pi_t := \mathbb{1}_{\{X_{t-1} < \theta\}}$. (We remind the reader that the expected value is different, $\mu/(1 - \alpha)$.)

As we can see from the example, to determine the threshold we either need to estimate μ and α from a long enough sample or either we learn the value of θ^* by using Stochastic Gradient. In a more realistic dynamics there are more than two parameters that needed to be estimated. Furthermore the threshold is very sensitive to the estimation error of α .

Example 3.6.2 (DGSV). Let the log-return H_t be a DGSV process according to (2.9). Its conditional expectation is

$$\mathbf{E}[H_t | H_{t-1} = x] = \mu + \alpha x + \sigma \rho \mathbf{E}[e_t^Y | H_{t-1}].$$

The conditional expectation is unknown but we will see later in the numerical results that there is a unique solution.

3.7 Constantly rebalancing portfolios

In the literature, constantly rebalancing portfolio can have different meanings. What we mean here is that the strategy is constant in time. That is, the investor uses the same strategy over and over again. Despite being a very simple strategy, if the return process was an i. i. d. sequence, this strategy could perform log-optimally.

However, for us, the aim is to use this strategy as a benchmark model. In the next chapter we contrast the log-optimal and the approximative strategies with the constantly rebalancing portfolio.

Let us denote this strategy as $\pi_t^{reb} \equiv \pi^{reb}$.

It is defined similarly as before, but instead of the conditional expectation of the growth at t it uses unconditional expectation:

$$\pi^{reb*} := \arg \max_{\pi^{reb} \in [0,1]} \left\{ \lim_{T \rightarrow \infty} \mathbf{E}[\log(W_T^{\pi^{reb}})/T] \right\}. \quad (3.36)$$

Since H_t is stationary by our assumption, it is well-defined. For this strategy, the first order approximation is quite simple: the investor decides to invest the whole wealth at the initial moment, then does not do anything.

We do not use, only mention the second order approximation:

$$\pi_t^{reb,quad} = \frac{\mathbb{E}[H_t]}{\mathbb{E}[H_t^2]} - \frac{1}{2}. \quad (3.37)$$

Chapter 4

Log-Optimal Portfolios for Conditionally Gaussian Processes

This chapter is about the numerical results of the log-optimal portfolios and their approximations. In Section 4.2 we present that the log-optimal solution reaches the optimal portfolio for almost every realization, while in Section 4.3 we analyse how the parameters that are related to the long memory are in connection with the investment.

Section 4.4 focus on how good if the performance of the approximative solutions. We do not have theoretical results about their goodness, but the numerical simulations show that they approximate well.

4.1 Introduction

As we have mentioned in Section 3.2, the long horizon optimization problem in (3.3) can be solved going along only one realization (see (A.1)), concluded from Assumption 3.2.3. Therefore, we can denote the long horizon value function, the growth as $G^* := \lim_{T \rightarrow \infty} \log(W_T^*)/T$. Numerical simulations were performed until the finite time horizon $T = 4000$, its value function is denoted as $\log(W_T^*)/T$. (It should be noted that $\log(W_T^*)/T$ is not the solution of the short horizon problem in (3.2), but the optimizer strategy is the same until T). We recall here that the optimal strategy function (3.10) is based on the function $\hat{\pi}(z, \nu, \kappa)$:

$$\pi_t^* = \hat{\pi}(Z_{t-1}, \nu_{t-1}, \kappa). \quad (4.1)$$

It is worth comparing the outcomes of the log-optimal strategy with that of another, well-known strategy, and also with the approximative strategies described in Section 3.4.

For this purpose, we chose the *constantly rebalanced portfolio* as a frequently used strategy, where the investors use a constant strategy function $\pi_t = \pi^{reb} \in [0, 1]$, see Section 3.7. The *best constantly rebalanced* investor optimizes the expected value function: $\pi^{reb*} := \arg \max_{\pi^{reb} \in [0, 1]} \{\lim_{T \rightarrow \infty} \mathbb{E}[\log(W_T^{\pi^{reb}})/T]\}$. In practice we determine this strategy by Monte-Carlo simulation of the models. The difference $\Delta_T := (\log(W_T^*) - \log(W_T^{\pi^{reb*}}))/T$ between the log-optimal and the best constantly rebalanced portfolio gives the measure of how much better the log-optimal strategy performs.

The convergence in Assumption 3.2.3 works for the approximative strategies and the constantly rebalanced strategy as well, therefore, the value function converges to a constant in every case. Let us denote the value function by G^* , G^{lin} , G^{quad} and G^{reb} respectively for the log-optimal, the first and second order approximative and the constantly rebalanced strategies.

To better understand the value function of each strategy, which can be considered as a daily compound rate, we convert the value function to *annual yield* (AY) (or yearly interest rate), which equals $\exp(250G) - 1$; that is, how much the portfolio's value grows in one year with value function (growth-rate) G .

Of course, in the numerical results we use a finite T to approximate the growth-rate. Because it is a random variable, the mean of the annual yield in the numerical simulations is relevant. We already introduced the so-called *mean annual yield* (MAY) in (1.11). For example, for the log-optimal investor its mean annual yield is

$$MAY^* := \sum_i e^{(250 \frac{1}{T} \log(W_T^*(i)))} - 1,$$

where i denotes the i -th realization (of which we took 100) and $T = 4000$ is the number of time-steps in the simulation.

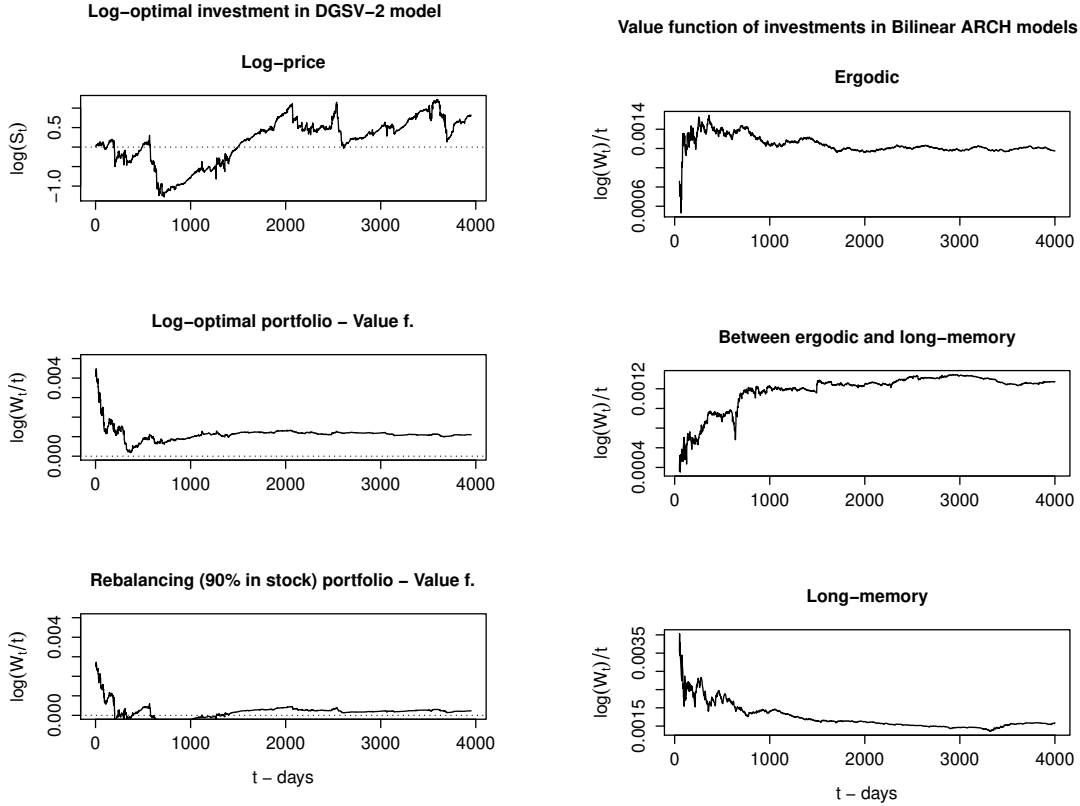
In every figure we used the same parameter settings as in Table 2.2. Whenever we changed the parameter, it is explicitly indicated. In every figure we used the same realization(s) to allow us to compare the effect of the parameters. We refer to different realisations as *seeds*.

4.2 Log-Optimal Solution

The two most important questions to ask about a new optimization algorithm are: 1) Is the optimum achieved on a 'reasonable time horizon'? 2) Is this optimum essentially better than competing strategies like constant rebalanced portfolios? The answer for the first question is shown in Figure 4.1. The value function becomes stable around $t = 1000$ in the DGSV model and around $t = 1500$ in the BARCH model, which is about 4-6 years assuming daily

portfolio rebalanced. It is reasonable, since log-optimal investments are typically used for long term portfolios. (In our simulation $T = 4000$, i.e. roughly 16 years.)

Value function and yearly interest rates are very high in the results of the simulation. Clearly, transaction fees of all types are excluded, which is why such excellent returns can be achieved.



(a) Value function of the DGSV model. Value functions converge around day 1000. Log-optimal portfolio exceeds by far the best constantly rebalanced strategy ($\pi^{reb} = 0.9$). Value functions of the log-optimal investor and the constantly rebalanced investor on the plotted realization are: 0.0009 and 0.00009, in annual yield: 26.6% and 2.2%. Parameters are shown in Table 2.2.

(b) Value function of the log-optimal investor in different parametrization of the BARCH model. Parameters for the first and third plot as the same as in Table 2.2 (left) for moderate and long memory. The middle plot has the same parameter as the ergodic one, but with $b = 1$. The annual yield in the three cases are: 26.5%, 37.4% and 62.3%.

Figure 4.1

The other important question is answered in Figure 4.2. With the same parameter settings as in Table 2.2, this figure shows how the log-optimal solution exceeds the best constantly rebalanced strategy for 100 seeds. The results are consistent and different driving noise did not affect them. We remind the reader that the two main differences between

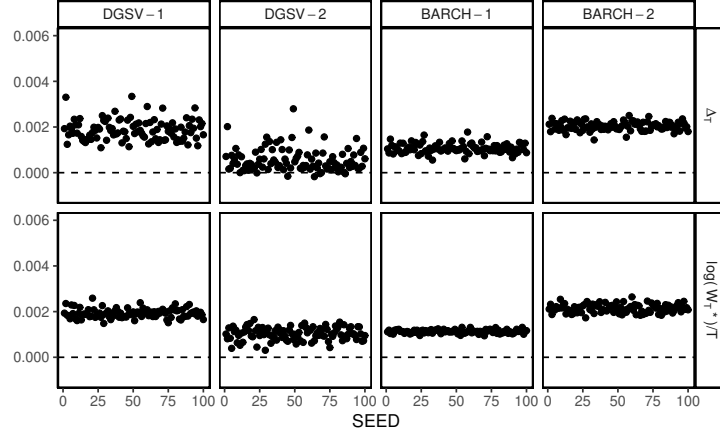


Figure 4.2: Value function at $T = 4000$ in every model for different realizations - because of the ergodicity, almost every realization should converge to the same number within a model. DGSV-2 has a more volatile stock price, it needs more time to converge, its deviation is higher. Surprisingly BARCH-long also shows convergence, although it is not proved to be ergodic.

DGSV-1 and DGSV-2 models is that the latter one has smaller autoregressive coefficient (α) and stronger memory (b_0 is higher). The result of this difference is that the convergence of the DGSV-2 is slower, hence the higher deviation on the figure.

Surprisingly, the BARCH-2 model with long memory also shows convergence, despite being unable to prove any kind of ergodicity (and Assumption 3.2.3 is also not proved). The BARCH-1 model has lower standard deviation than the others, which is obvious, since its memory is only of moderate type.

The variable Δ_T shows how the log-optimal investor overcomes the constantly rebalanced one. In the case of DGSV-2 we see that $T = 4000$ was not enough long to beat the benchmark strategy in 5 out of 100 realizations (we note that, by the ergodic assumption, Δ_T converges to a positive constant a.s.).

4.3 Analysing the Role of Memory

Figures 4.3 and 4.4 investigate the relationship between the strength of the memory and the performance of the log-optimal investor for 5 different scenarios (different random sequences). Regarding the strength of the memory, we are focusing on the exponent b , which is the decay of the past information. We remind the reader that when $b \leq 1$, the model has long memory, and when $b > 1$, it has moderate memory (both have infinite memory, but the former one is stronger).

Figure 4.3 shows the value function for the four models described in Section 2.5 and 2.6

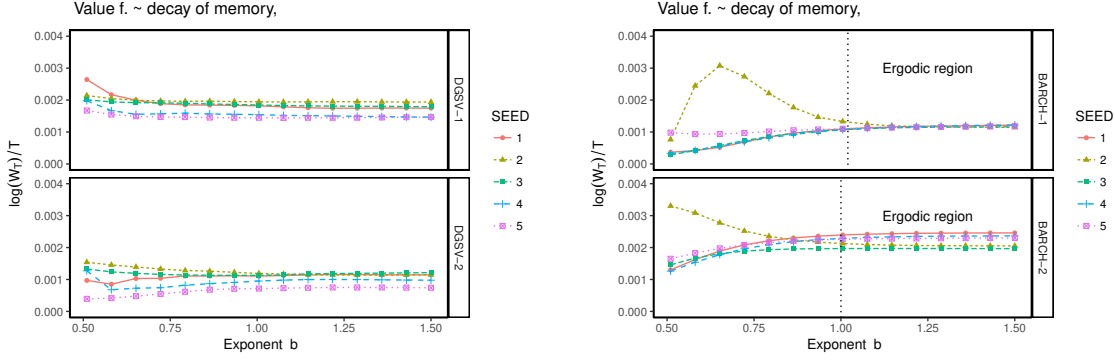


Figure 4.3: Relation between the value function of the two DGSV and the two BARCH models and the memory decay parameter b . The parameters are identical to the ones in Table 2.2 except for the exponent b . The model BARCH is not ergodic on every value of $b > 0.5$ unlike the model DGSV, therefore a dashed vertical line shows where the BARCH model is ergodic.

but for different b values. While in the previous sections the model BARCH-1 referred to the moderate memory case and BARCH-2 to the long memory, here it is different, since b determines the type of the memory. Interestingly, in those cases where the ergodicity assumption can be proven for the models, the value function is independent of b and the value function of different realizations are close to each other. We can see this on Figure 4.3. In the cases of DGSV-1/2 when $b \approx 0.5$, this effect deviates a little. In the long memory BARCH case ($b < 1$) the relationship between the value function and the memory decay parameter b cannot be described.

We have seen in Section 2.5 that every unconditional moments of the DGSV model only depends on the memory through b_0 and $\sum \beta_j^2$ (where $\beta_j = b_0(1+j)^{-b}$).¹ Based on this, we wanted to investigate how the value function depend on b_0 while the sum $\sum \beta_j^2$ is kept constant. We found that the value function has an approximately linear dependency structure on b_0 . Figure 4.4 displays their relation on a log-log plot and shows a linear relationship between them. Since the linear fit on the data (black line) performs well, it suggests a power law dependency between the value function and the decay parameter.

As an interesting fact we mention that the expected value of the log-return in model DGSV is a linear function of b , and in the continuous time-model the value function of the log-optimal investor is also a linear function of the mean. Unfortunately there is no evidence that it is indeed true in the discrete case (or at least in our case).

¹In the DGSV model the term $\sum \beta_j^2$ appears as the log-volatility. The formulation of this term in continuous time could be formulated as the variance of the log-volatility.

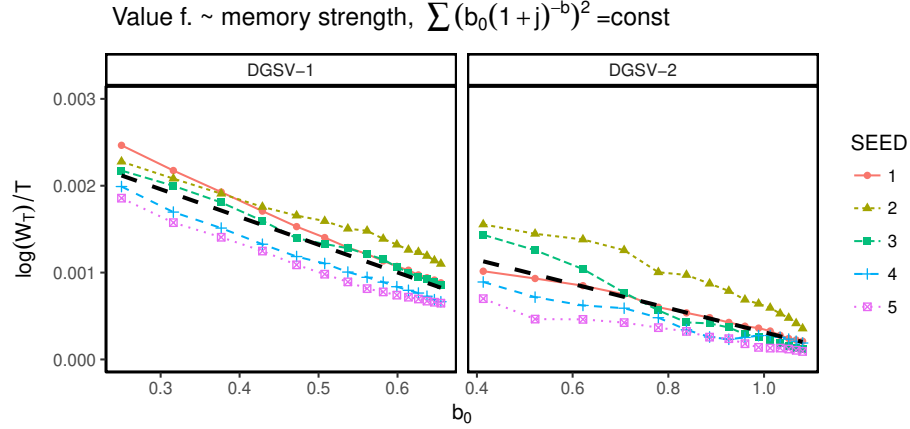


Figure 4.4: We show how the parameter b_0 effects the value function when the term $\sum \beta_j^2$ is kept constant. It appears that their relationship is linear, as we show with a fitted dashed black line on the five different realizations at once.

4.4 Approximation of the Log-Optimal Strategy

In the numerical simulations we found that log-optimal strategy is extreme in the sense that the investor only buys a riskless or risky asset every time but hardly uses mixed strategies (Figure 4.5). In a negligible part of the time there can be mixed decisions but this is so rare that it cannot even be seen on the histogram. In the DGSV-1 parameter settings the log-optimal investor used mixed strategy only 61 times out of the 4000 investment days, while in the DGSV-2 case the log-optimal investor used mixed strategy 338 times out of the 4000 investment days. The reason for this is that in the second case the autoregressive coefficient α is smaller (0.05 instead of 0.3) and that the memory is stronger ($b_0 = 0.55$ instead of $1/3$), which results in a much higher volatility.

We have seen that the optimal strategy π_t^* is determined by the optimal selector function $\hat{\pi}_t(h, \nu, \kappa)$ as (3.10) shows and the approximated portfolios approximate this log-optimal strategy. The first order approximation from Section 3.4 answers the question of why the log-optimal strategy is mostly 0 or 1. Since the standard deviation of the log-return is around 0.02, even the first order approximation performs well where the optimal strategy is exclusively 0 or 1 (see (3.20)). Figure 4.6 (a) shows the log-optimal strategy selector function $\hat{\pi}$ for the DGSV-2 model on a blue coloured scale, while the red curve on the figure shows the 0–1 barrier from (3.22). The first order approximated strategy is 1 above the red curve and 0 under it. We can see that the red curve approximates well where the log-optimal strategy is closer to 0 or 1.

Figure 4.6 (a) shows a small region on the plane (h, ν) where the log-optimal investor uses mixed (between 0 and 1) decisions. The second order approximation explains this deficiency

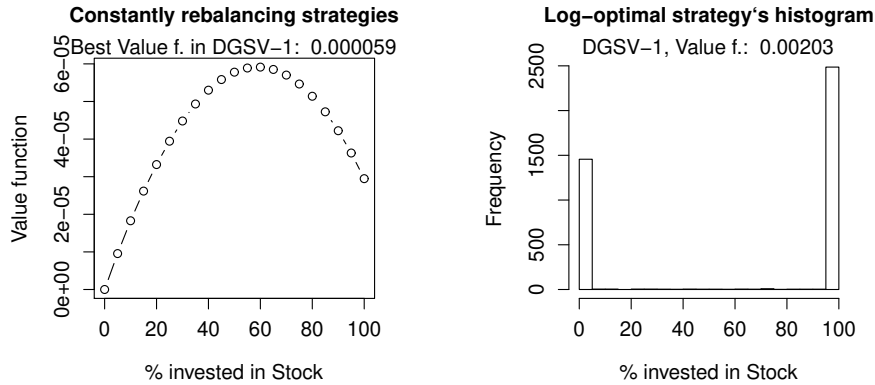


Figure 4.5: Comparing the Value function of the constantly rebalanced investors (left) versus the log-optimal investor (right, histogram of their strategy). The parameters are the same as in Table 2.2, DGSV-1 model. The best constantly rebalanced portfolio was determined by MC simulation of 200 paths with $T = 4000$. The value function reaches its maximum when we invest 60% in the risky asset. The yearly average yield of the constantly rebalanced and the log-optimal investor is: 1.5% and 66%. The log-optimal investor mostly uses 0 or 1 strategy.

and Figure 4.6 (b) shows that this approximation is almost impossible to distinguish from the log-optimal one on the figure. The size of the mixed strategy region is proportional to the memory. The higher the b_0 or smaller the b or smaller the α , the bigger this region is.

Of course, the approximated strategies cannot be the same as the log-optimal one. In the model DGSV-1, where the autoregressive part is also significant ($\alpha = 0.3, b_0 = 1/3$), the approximations perform well, while in the DGSV-2 model, where $\alpha = 0.05, b_0 = 0.55$, the mixed strategy region is bigger, hence the approximated strategies perform more poorly.

Numerical results of the competing strategies are displayed in Table 4.1. We observed that the quadratic approximation performs well in each case. In fact, in the case of BARCH-2, the approximation exceeds expectations (where log-optimal solution is not proved). The reason for this good performance is that the quadratic approximation simply uses function evaluation, whereas in the log optimal case, we use numerical integration and we need to discretize the strategy's interval. Linearly approximated strategy works well, but slightly worse in the case of DGSV-2, where the memory is stronger and the AR coefficient α is small. The best constantly rebalanced investor performs poorly in all cases compared to the log-optimal or the approximated log-optimal cases.

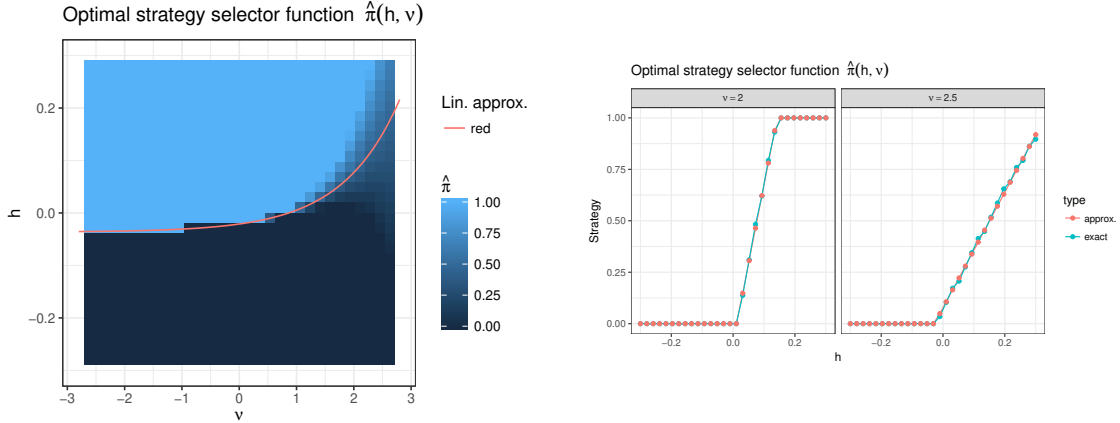


Figure 4.6: Log-optimal strategy and its approximations. Figure (a) compares the log-optimal strategy with the first order approximation (Lin. approx). Above the red curve the approximated strategy is 1, and below it, it is 0. Figure (b) compares the log-optimal strategy with the quadratic approximation. Instead of the whole $\hat{\pi}(h, \nu)$ two-variable function, we only show results for high values of ν , where the difference between the two strategies is more significant. Nevertheless, it is difficult to see the difference between the strategies on the plots.

Portfolio	Mean Annual Yield			
	DGSV-1	DGSV-2	BARCH-1	BARCH-2
Log-optimal	62.080%	29.641%	32.568%	70.966%
2nd order approx.	62.075%	29.623%	32.567%	70.971%
1st order approx.	62.051%	28.384%	32.569%	70.966%
Best const. rebal.	2.0946%	14.053%	1.618%	2.673%

Table 4.1: Comparing the profitability of different portfolios (different investment strategies) based on simulation of 100 paths for $T = 4000$ investment days. Since the bank's interest rate $r = 0$, every yearly interest rate above 0 is profitable.

4.5 Conclusion

In summary, numerical simulations show that our algorithm performs well in the models we considered. Not only can we prove that we have indeed found the optimal strategy, we can also evaluate its performance (which would not be possible in the general setting of e.g. [3]). The optimal strategy significantly beats constantly rebalanced strategies. Furthermore, the optimal strategy can be calculated in a simple way via numerical integration or further simplified with well performing approximations.

We were able to investigate the effect of the strength of memory on investment success

in a quantitative way. We found relationships between parameters related to memory and the value function of the log-optimal investment.

Comparison with real statistical data makes it plausible that it may indeed be profitable to apply this strategy in investment practice. Knowing the underlying model, $\hat{\pi}$ and $g^{\hat{\pi}}$ can be determined in advance and the data can be fed into it as time goes on. The model can easily be extended to include several assets and exogenous macroeconomical factors.

In the next chapter, we explain an algorithm that work online, without knowledge of the underlying model.

Chapter 5

Adaptive Learning of Log-Optimal Strategies

The numerical results in the previous chapter showed that approximative strategies perform well on a realistic investment problem. Building on this observation, in this chapter we introduce a learning algorithm that is able to find the (approximative) optimal strategies for the log-optimal investor in an online setting where data points arrive one-by-one and the investor adapts his portfolio to them. In contrast to most learning methods, the computational capacity of our algorithm is low therefore it seems applicable in situations where resources are constrained, such as in the case of high frequency trading. The optimization procedure goes along with the incoming data stream without the need of complex estimations, unlike most non-parametric methods. The method can be used on real data as well. In the case of simulated data, the optimal parameter of the algorithm can be calculated for parametric models (using numerical integration or Monte Carlo) and thus the performance of our procedure can be compared to them.

The systematic study of this method is fairly novel in the field of portfolio optimization. In this chapter we aim to study the theory and practice of Stochastic Gradient and related algorithms as applied to parametrized trading strategies.

5.1 Introduction

The goal in this chapter is to optimize the logarithmic utility based on continuously incoming information on financial assets (typically their prices). This study investigates the Kiefer–Wolfowitz variant of the Stochastic Gradient method, that is shown to converge in simulations to the log-optimal solution in the threshold-type, buy-and-sell strategy class (see

Section 3.6).

We demonstrate that there is an optimal threshold type strategy which can be learned. We check this for different choices of stock price dynamics: e.g. with stochastic volatility and long memory (DGSV), and on certain simpler models.

Subsequently, we numerically verify the convergence of the algorithm. Furthermore, we analyze in depth the typically problematic question of how to choose the hyperparameters (the parameters of the algorithm and not the parameters in the dynamics of the prices) without knowing anything about the price other than a small sample.

In high frequency trading the decisions must be made in seconds or even milliseconds [43]. Given the nature of algorithms, those that require huge computational capacities are not efficient since they are slow. For example, non-parametric methods or complex machine learning algorithms work well on big data sets with immense computer efforts (see a survey of non-parametric methods e.g. in [39] and a summary in machine learning methods [44]). On the other hand, parametric models that are based on dynamics of the prices or indices may give fairly good results if precise and accurate parameter estimations are available. To get satisfactory estimations, again one needs a big data set, and typically decisions are sensitive to the error of the estimations.

To resolve the above-mentioned problems we use the Kiefer–Wolfowitz method to choose the parameters of our *strategies* which lets us (i) make decisions immediately starting at the initial step; (ii) process new information/data as they arrive, without needing to wait until we have a big enough data set, as the strategy function improves in every step and (iii) there is no need to estimate the parameters of the *dynamics*. With this method, we aim to optimize *log-utility* investments (maximizing the expected value of the logarithm of the wealth). The method is also capable of tracking the changes of the market dynamics, which we ignore here in order to investigate the method itself in finance instead of market changes.

Stochastic Approximation [45] (or Robinson–Monro method) is an iterative method to find the root of $f(\theta) := \mathbb{E}[X(\theta)]$, such that $f(\theta^*) = m$, where $X(\theta)$ is a random variable, such that its distribution depends on θ and θ^* is the root of $f(\theta) - m$. The method assumes that we are able to generate i. i. d. samples from $X(\theta)$ with a given θ . Basically, it is a stochastic version of the Newton–Raphson method where there are consecutive observations of the functions loaded with randomness/noise. If the derivate of the function exists then the method can be used for optimization. When the derivate does not exist or is unknown, Kiefer and Wolfowitz proposed [46] a finite difference approximation based on consecutive observations. This method is a variant of the Stochastic Gradient method.

Nonetheless, as far as we know, stochastic approximation or Kiefer–Wolfowitz has not been used for directly optimizing the parameter of a parametrized family of strategies as we

do here. We hope that this introduction of the usage of the method in investment theory will lead to further developments. Other works dealt with different approaches, like learning the parametrized stopping time for American/Asian options [47] or the optimal stopping time of liquidation [48]. Other typical fields of applying this algorithm are the estimation of quantiles for CVaRs [49], [50] or [51]. There is also an intriguing study about the optimal splitting of orders in [52].

5.2 Stochastic Approximation

5.2.1 The Robbins–Monro algorithm

Stochastic approximation is an iterative method to find the root of a function where one can only take measurements of the function contaminated with noise. If we can make (noisy) measurement of the derivate of the function then it can be used for optimization.

The basic problem was described Robbins and Monro [45]. We assume that we are able to take i. i. d. samples X_1, X_2, \dots of $X(\theta)$ with a distribution function $D(x|\theta)$ defined as

$$\mathbb{P}(X < x|\theta) = \int_{-\infty}^x dD(x|\theta), \quad (5.1)$$

where the parameter θ runs over \mathbb{R} , for simplicity. We wish to solve the equation $f(\theta) = m$ where m is given and

$$f(\theta) = \mathbb{E}[X(\theta)] = \int_{-\infty}^{\infty} x dD(x|\theta). \quad (5.2)$$

The Robbins–Monro procedure is defined as

$$\theta_{t+1} = \theta_t + a_t(m - X_t), \quad (5.3)$$

for some arbitrary θ_0 initial value. Under mild conditions, the series θ_t converges in probability to a constant θ^* as $t \rightarrow \infty$ that satisfies the equation

$$f(\theta^*) = m. \quad (5.4)$$

In words, we can formulate the procedure as

$$\theta_{t+1} = \theta_t + \{\text{step size}\}\{\text{direction}\}. \quad (5.5)$$

Besides assuming the availability of i. i. d. samples, typical assumptions are:

1. The values of X_t must come from a bounded set, i.e. there exist a finite constant $c > 0$

such that $\int_{-c}^c dD(x|\theta) = 1$.

2. The real sequence $a_t > 0$ satisfies $\sum_t a_t^2 < \infty$ and $\sum_t a_t = \infty$.
3. The function $f(\theta)$ has only one root and $f(\theta) \leq m$ for $\theta < \theta^*$ and $f(\theta) \geq m$ for $\theta > \theta^*$.

The last assumption is similar to the assumption we had for the existence of optimal threshold in Proposition 3.6.1.

The Robbins–Monro algorithm can be formulated in a way to separate the random variable X and the parameter, let $f(\theta) := \mathbb{E}[F(X, \theta)]$ for a suitable function F . In this setting, at step t , we use $F(X_t, \theta_t)$ to determine the direction in the procedure.

The first version of the Robbins–Monro algorithm relies on the assumptions that we are able to make i. i. d. and unbiased samples of the function $f(\theta)$, but both assumption can be relaxed. Clearly, if f is the derivative of a function U then applying the Robbins–Monro algorithms with $m = 0$ then the minimum (maximum) of U can be found.

5.2.2 The Kiefer–Wolfowitz algorithm

The Robbins–Monro procedure can be used for optimization if noisy, unbiased estimation of the derivate is available. Otherwise the Kiefer–Wolfowitz algorithm or its variants can be used to approximate the derivate.

Let us assume that we want to maximize the function $g(\theta) = \mathbb{E}[G(X, \theta)]$, where X is a random variable and θ is a real parameter. In the Robbins–Monro algorithm (5.5) the direction should be the gradient, if one wishes to optimize the function $g(\theta)$. Since, it is not available, we have to use an approximation. In the Kiefer–Wolfowitz procedure, a finite difference is used as an approximation of the gradient.

The Kiefer–Wolfowitz procedure is

$$\theta_{t+1} = \theta_t + a_t \frac{G(X_t, \theta_t + c_t) - G(X_t, \theta_t - c_t)}{c_t}, \quad (5.6)$$

where the step-size a_t and the step-size of the finite difference c_t are real-valued sequences.

In our case, we want to optimize the growth function in (3.30) or (3.31), that is, $G(H_t, X_{t-1}; \theta) := H_t \mathbb{1}_{\{X_{t-1} > \theta\}}$. From the mathematical point of view, the particularity is that G does not only fail smoothness, it even fails continuity. We will see, however, that this does not prevent the algorithm from converging.

The Kiefer–Wolfowitz algorithm in our case is

$$\theta_{t+1} = \theta_t + a_t \frac{G(H_t, X_{t-1}; \theta_t - c_t) - G(H_t, X_{t-1}; \theta_t + c_t)}{c_t}, \quad (5.7)$$

Since the growth $G(\theta; \dots)$ is an indicator function of θ , therefore its finite difference can be simplified to the indicator function of a range. For greater clarity we denote the range $[x - c, x + c]$ as $[x \pm c]$. Then the algorithm can be written as

$$\theta_{t+1} = \theta_t + a_t \frac{H_t \mathbb{1}_{\{X_{t-1} \in [\theta_t \pm c_t]\}}}{c_t}. \quad (5.8)$$

This formalism will help us in the latter to better understand the usage of the method.

To prove convergence of the algorithm in such generality seems very challengin. But via examples in the Section 5.3 we will check *numerically* that this recursively defined process θ_t converges to the optimum what we showed in the previous section:

$$\theta_t \xrightarrow{MSE} \theta^*. \quad (5.9)$$

Convergence is shown in *numerical simulations* in the sense of the Mean Squared Error (MSE) which is a sample-based analogue of the L^2 -convergence of random variables.

Definition 5.2.1. (*Convergence in Mean Squared Error*)

Let $x_t(i)$, $t \in \mathbb{N}$ and $i = 1, 2, \dots, N$ be the i -th sample of a measurement at time t . Then x_t converges to the real number x in the sense of Mean Squared Error, if

$$\frac{1}{N} \sum_{i=1}^N (x_t(i) - x)^2 \xrightarrow{t \rightarrow \infty} 0, \quad (5.10)$$

that is, if the squared error at time t , which is $\frac{1}{N} \sum_{i=1}^N (x_t(i) - x)^2$ converges to 0.

We find that, typically, its convergence speed is a power function of the number of iterations.

In general, there is no straightforward way to choose the hyperparameters determining the choice of the sequences a_t, c_t . In Section 5.3 we present some ideas on the basis of which we can choose the hyperparameters in a reasonable way.

Since the algorithm iterates the θ parameter in every step, it gives the investor the possibility to modify the investment strategy. Using the threshold strategy function in (3.32) with the current parameter θ_t , the strategy is

$$\pi_t = \tilde{\pi}(H_{t-1}, \theta_t). \quad (5.11)$$

As θ_t converges to θ^* , the strategy improves during the investment.

In real life investment the financial environment is not static, the dynamics of prices can

change and new factors can appear/disappear, therefore optimal strategy changes as well. To this end, in practice investors use constant and very small step sizes a_t and c_t which can track down the changes of the optimal values. As before, we do not investigate the role of temporal financial effects here.

5.3 Numerical Results

The critical part of every algorithm is the choice of the hyperparameters. In their paper, J. Kiefer and J. Wolfowitz [46] also address the issue of parameter-choice though they were able to give exact and sufficient conditions in a simpler context. These conditions are typical requirements and our model satisfy them as well:

1. $c_t \rightarrow 0$.
2. $\sum_{t=1}^{\infty} a_t = \infty$, that is, the algorithm can reach any state.
3. $\sum_{t=1}^{\infty} a_t c_t < \infty$.
4. $\sum_{t=1}^{\infty} a_t^2 c_t^{-2} < \infty$.

For simplicity, we choose the two series $a_t = t^{-p}$ and $c_t = t^{-q}$ such that it staisfy the above mentioned criteria. With this choice, we are able to characterize the stepsize and the small change of the numerical difference by one-one parameter. The most used choice in theoretical studies is $p = 1$ and $q = 1/3$ as they give optimal *theoretical* convergence speed. We remark, however, that this is not always the optimal choice in practice. Unless stated otherwise, we use these values.

During the numerical results we often rely on Monte Carlo simulations. Of course, the aim of using iterative methods is to avoid using Monte Carlo algorithms. Monte Carlo procedure uses the whole past data to estimate an expected value, that is, it requires huge computational effort. In a high frequency trading environment there is no capacity for this and the iterative method requires few capacity and can be used right continuously as the data stream arrives. The reason why we use Monte Carlo here, is that on a big dataset it gives a very good estimate of θ^* and thus we can *assess* the performance of the Kiefer–Wolfowitz algorithm by comparing to the value thus calculated. Moreover, it helps us to estimate the expected growth function $g(\theta)$ to understand the K–W algorithm.

5.3.1 Analysing the Expected Growth Function

Let us analyse the expected growth function $g(\theta)$. Understanding its behavior helps us to construct the step-sizes in a suitable way, and also to understand how the algorithm works.

Figure 5.1 and (5.8) makes it clear that θ_t must stay in the same range as X_{t-1} , since $X_{t-1} \notin [\theta_t \pm c_t], \forall t \in \mathbb{N}$ would result in constant θ_t . In the numerical simulations we only show results about the $X_{t-1} := H_{t-1}$ case. On the examples shown on Figure 5.1 we can

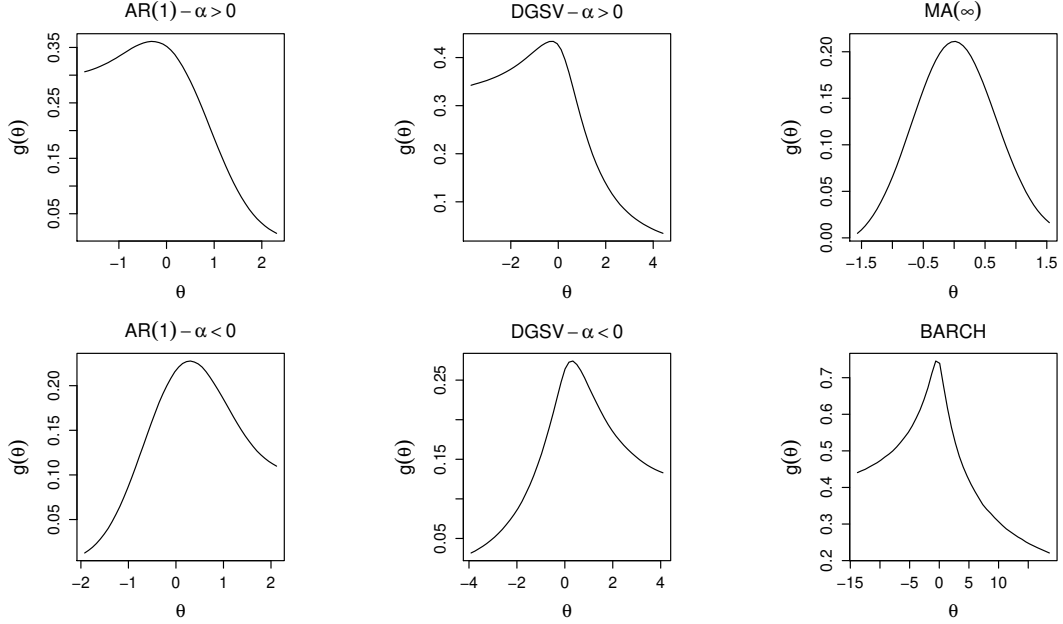


Figure 5.1: Expected growth as a function of threshold parameter θ . The Kiefer–Wolfowitz algorithm seeks its maximum. When $\alpha < 0$, H_t and H_{t-1} are negatively correlated, when $\alpha > 0$ and in the MA(∞) and BARCH cases they are positively correlated. The values of the function were generated by Monte Carlo simulation of the process.

make the following remarks:

- $g(\theta \rightarrow -\infty) = \mathbb{E}[H_t]$, low θ means that $\pi_t = 1$, that is, the wealth equals to the price of the stock.
- $g(\theta \rightarrow \infty) = 0$, high θ means that $\pi = 0$, the wealth equals to the price of the bond.
- In the negatively correlated case instead of (3.30) we should use (3.31). Then $g(\theta \rightarrow -\infty) = 0$ and $g(\theta \rightarrow \infty) = \mathbb{E}[H_t]$. This is the case when the autoregressive coefficient is negative, $\alpha < 0$.
- If θ_t takes value out of the typical value of H_t where the derivate of $g(\theta)$ is zero then it is hopeless for the algorithm to return and it stays there.
- Interestingly, the steepest $g(\theta)$ functions are produced by the DGSV and BARCH model, although $\alpha = \pm 0.5$ in every cases. The reason for this, is that extra information about H_t is hidden in the volatility because of the leverage effect.

- Steeper $g(\theta)$ function does not necessarily means faster convergence. The original K-W algorithm assumes independent samples, but in our case we have the log-return process $\{H_t\}_{t \in \mathbb{N}}$, which is time-correlated. This results in a slower convergence speed.
- In every presented dynamics there is a unique θ^* value. Moreover, the $g(\theta)$ function is smooth in every cases.
- As we can see that $g(\theta)$ flattens out towards the $\pm\infty$, it is desirable to choose the initial value θ_0 as close to the optimal θ^* as possible. One possible way to do this is to set θ_0 to a preliminary estimation of the mean.

To overcome on the problem of choosing a proper initial value θ_0 and to avoid the flat part of the curve $g(\theta)$ we apply two procedures. First, the initial value is set to the average of the first few values of the log-return

$$\theta_0 := \frac{1}{10} \sum_{t=1}^{10} H_t. \quad (5.12)$$

Second, a very frequently used method, is to project the θ_t values on a bounded set where $g(\theta)$ is not flat.

$$\theta_t = \begin{cases} \min_{1 \leq j \leq t} H_j, & \text{if } \tilde{\theta}_t < \min_{1 \leq j \leq t} H_j, \\ \tilde{\theta}_t, & \text{otherwise,} \\ \max_{1 \leq j \leq t} H_j, & \text{if } \tilde{\theta}_t > \max_{1 \leq j \leq t} H_j, \end{cases} \quad (5.13)$$

5.3.2 Performance of the Kiefer–Wolfowitz algorithm

The major question about the algorithm is weather the θ_t series converges to θ^* and in what sense? It seems challenging to prove mathematically for a general long memory process, but the results of the numerical simulations show that

$$\theta_t \xrightarrow{MSE} \theta^*, \quad (5.14)$$

that is, it converges in the Mean Squared Error (MSE) sense.

That is, we generated by Monte Carlo simulation N different realizations for the log-return process and for each process we applied the K-W procedure, resulting in $\{\theta_t(i)\}_{t=1, i=1}^{T, N}$, where T is the number of time-steps the procedure was running. The Mean Squared Error is

$$MSE_t := \frac{1}{N} \sum_{i=1}^N (\theta_t(i) - \theta^*)^2, \quad (5.15)$$




Figure 5.2: Convergence of the Kiefer–Wolfowitz algorithm in L^2 for the some process as in 5.1. The left figures shows 5-5 realizations, while the right figures shows the Mean Squared Error (MSE) of all the realizations.

where $\theta_t(i)$ is the i -th realization at the t -th timestep. It is also common to talk about the *Root Mean Square Error (RMSE)*, which is the square root of MSE , that is, $RMSE = \sqrt{MSE}$.

Another important question is if the Mean Squared Error converges, then how fast? In most cases we found MSE_t decays at a power speed, that is $MSE_t \approx t^{-r}$.

Theoretical studies also give such results: $\lim_{t \rightarrow \infty} \mathbb{E}[|\theta_t - \theta^*|^2] = 0$ is shown to hold under appropriate conditions (which are not met in our case). It has been shown that the best possible convergence rate, using Robbins–Monro algorithm is t^{-1} , while using the Kiefer–Wolfowitz algorithm is $t^{-2/3}$, (see [53], [54] and [55]).

Using the simple parametrization $a_t = t^{-1}$ and $c_t = t^{-1/3}$ Figure 5.2 shows the convergence of (5.15) in the case of AR(1) and DGSV process ($\alpha > 0$). The simulations were executed with $N = 25$ realizations and for $T = 50\,000$ time steps and the figure shows 5 realizations out of the 25.

The Mean Square Error is shown on a log-log plot on Figure 5.2. Power function on a log-log plot is a linear line. As we can see, in both cases, $t \rightarrow MSE_t$ is approximately linear on the log-log plot, that is, $MSE_t \sim t^{-r}$ for some $r > 0$.

5.3.3 Modifications

First, we investigate the effect, what happens if we leave projecting on a bounded set as in (5.13). The algorithm parameters were set to $a_t = t^{-1}$ and $c_t = t^{-1/3}$. The numerical simulations consisted of $T = 50.000$ timesteps and $N = 25$ realizations. We investigated two models, the AR(1) and DGSV, both with three different parametrizations. Table 5.2 shows the speed of convergence for each dataset with and without using the projection. The

parameter settings of each dataset is shown in Table 5.1. While 'dataset 1' does not show much difference for using the projection (the rate changes from 0.54 to 0.46 for the AR(1) process and from 0.31 to 0.30 for the DGSV process), in the case of the other two processes, the convergence speed more or less triple if we use the projection.

	μ	α	σ	ρ	b_0	b
dataset 1	0.15	0.5	0.75	-0.2	0.4	0.7
dataset 2	0.01	0.5	0.05	-0.2	0.4	0.7
dataset 3	0.005	0.3	0.05	-0.2	0.4	0.7

Table 5.1: Parameters of the three datasets used in Table 5.13. The AR(1) model only uses the μ, α and σ parameters, while DGSV uses all of them.

Model	dataset 1	dataset 2	dataset 3
AR(1)	-0.54 / - 0.46	-0.05 / - 1.35	-0.32 / - 1.40
DGSV	-0.31 / - 0.30	-0.86 / - 1.61	-0.51 / - 1.32

Table 5.2: Convergence speed of the Mean Square Error for AR(1) and the DGSV model. For each model and dataset the table shows the results 'without / with' projection.

Model	dataset 1	dataset 2	dataset 3
AR(1)	-0.16	-0.030	-0.013
DGSV	-0.008	-1.61	-1.18

Table 5.3: Convergence speed in the case, when θ_0 is set as the 0.995-th quantile of of the process.

The other settings which we can modify is how we choose the initialization θ_0 . Up to now, we estimated it from from a short sample, as the average of the first 10 elements (each realization started from different θ_0). Let us examine, what happens when θ_0 is started from the 99% percentile of the process. Results are presented in Table 5.3. In some cases the speed of convergence is not ruined by choosing a very sub-optimal θ_0 .

Chapter 6

Optimal Trading of Fractional Brownian Motion

This chapter is based on joint work with Paolo Guasoni and Miklós Rásonyi. The results of this chapter can be found in [56].

In contrast to the previous chapters, we alter our notations in some cases. From now on, we use continuous-time models and we include a friction term that has been so far neglected. That is, we assume that the investor faces liquidity constraints and has to pay an extra price for the trades, depending on trading speed.

The stock price here is considered to be a fractional Brownian motion, which may seem an arbitrary choice at first glance, but it captures the essence of continuous-time processes with long memory. Despite the fact that the stock price can have negative values, we explain that this investment model is financially realistic.

Because of the complexity of the continuous models with long memory, we solely focus on the risk-neutral preference case.

6.1 Model

Like in the previous chapters, we still assume that the investor uses only two assets: a riskless and a risky one. The latter will be assumed a *Fractional Brownian Motion* (FBM), while the riskless asset will have 0 interest rate, therefore one might think of it as cash. The assumption of friction is not only to make the model more realistic: without it, the problem is ill-posed.

The appearance of friction makes the stock illiquid, that is, the investor is not able to sell

it instantly for free. There is a finite "speed" of the transactions that restricts the investor, therefore it is meaningless to value the investment as the sum of the cash and stock. A reasonable question is how to maximize the cash position at the *investment/trading horizon* $T > 0$.

So that we avoid any misunderstanding, we use a different notation for the asset prices as before. Let (Ω, \mathcal{F}, P) be a probability space equipped with a filtration \mathcal{F}_t , $t \in [0, \infty)$. The stock price $S_t = \sigma B_t^H$ is a fractional Brownian motion B_t^H with some scaling factor $\sigma > 0$. The Hurst-exponent H characterizes the memory: when $H > 0.5$ the process has positive long memory, when $H < 0.5$, the memory is called *anti-persistent*. For the $H = 0.5$ case we have a Brownian motion without memory. The anti-persistent memory is also called negative-memory, since in this case, any non-overlapping increments are negatively correlated—which results in heavily wiggly paths. The fractional Brownian motion process is defined by three requirements: *i*) its initial value $B_0^H = 0$; *ii*) for every realization its path are continuous in time; *iii*) it is a Gaussian process and its covariance function is

$$\text{Cov}(B_t^H, B_s^H) = \frac{1}{2} (|t|^{2H} + |s|^{2H} - |t - s|^{2H}).$$

In the current case, the strategy ϕ_t is the *speed of trading*, which gives how fast the investor buys or sells its stocks. The strategy is a *feasible* if it is optional¹ and almost surely absolutely integrable on $[0, T]$. We are only looking for feasible strategies.

The portfolio is characterized by two terms. The amount of cash at t is denoted by X_t^0 , while the number of the shares the investor holds in stock is denoted by X_t^1 . Their initial value at $t = 0$ is $(z^0, z^1) \in \mathbb{R}^2$. Since ϕ_t is the speed of trading, the number of shares equals

$$X_t^1(\phi) := z^1 + \int_0^t \phi_u du. \quad (6.1)$$

When $\phi_u < 0$, the investor sells and when $\phi_u > 0$, the investor buys stock. The amount of cash, on the other hand, changes in the opposite way with the speed of trading:

$$X_t^0(\phi) := z^0 - \int_0^t \phi_u S_u du - \int_0^t \lambda |\phi_u|^\alpha du. \quad (6.2)$$

The term $\int_0^t \lambda |\phi_u|^\alpha du$ is what the investor pays for the transaction. The typical value of the *illiquidity* λ is 0.01 – 0.001, while the *price-impact* $\alpha > 1$. The $\alpha = 1$ case is called the *constant bid-ask spread*, where the investor pays transactional cost proportional to the number of stocks he/she bought or sold. When $\alpha > 1$, the investor pays higher transactional

¹We do not give the definition of optional process, since it is rather technical and does not restrict a real investment in any way.

cost per stock for a bigger buy or sell. Notice that the first integral is price of stock that the investor has bought or sold on the interval $[0, t]$, because the term $\phi_u du$ equals the number of stocks that has been traded, while S_u is its price.

Because of friction, we investigate the monetary profits, as we noted, the terminal position of the risky asset must be sold away and the portfolio consist at the end only cash. It restricts the feasible-strategy class to

$$\mathcal{A}(T) := \{\phi \in \mathcal{S}(T) : X_T^1 = 0, E[(X_T^0(\phi))_-] < \infty\},$$

for $T \in (0, \infty)$.

The aim of the investor is to optimize the above mentioned trading problem with the objective function

$$u(T) := \sup_{\phi \in \mathcal{A}(T)} E[X_T^0(\phi)].$$

6.2 Optimal strategy

The optimization task would be ill-posed in the case of absence of frictions and arbitrage opportunity would rise. From Theorem 5.1 of [57] it follows that if

$$\int_0^T E|S_t|^{\frac{\alpha}{\alpha-1}} dt < \infty$$

for every $T > 0$, then $u(T) < \infty$, the problem is well-posed and there exists a trading strategy ϕ^* such that

$$u(T) = EX_T^0(\phi^*(T)).$$

Optimality can only be proven asymptotically. We classify the strategies based on the "speed": how fast the expected payoff function grows.

Definition 6.2.1. *A family of strategies $(\phi^T)_{T>0}$, $\phi^T \in \mathcal{A}(T)$ is asymptotically optimal if $E[X_T^0(\phi^T)] > 1$ for all T large enough and*

$$\lim_{T \rightarrow \infty} \frac{\ln E[X_T^0(\phi^T)]}{\ln u(T)} = 1. \quad (6.3)$$

In other words, a sequence of strategies is asymptotically optimal when it approaches the maximal rate of profit growth for large horizons.

We have to make two important remarks about the definition. First, the maturity T is a fixed parameter, it must be known at the beginning of the investment. Second, the logarithm function in the fraction means that if the expected payoffs only differs in a term

that are negligible with respect to $\ln u(T)$, then they are in the same family of strategies.

The following theorem shows the main theoretical statement of this chapter. The proof of this theorem can be found in [56]

Theorem 6.2.1. *Let $\lambda > 0, \alpha > 1, H \in (0, \frac{1}{2}) \cup (\frac{1}{2}, 1)$. Then for $S_t := B_t^H$,*

(i) *Maximal expected profits satisfy*

$$\limsup_{T \rightarrow \infty} \frac{u(T)}{T^{H(1+1/(\alpha-1))+1}} < \infty. \quad (6.4)$$

(ii) *For each $0 < \kappa < 1/(\alpha - 1)$, the strategies*

$$\phi_t(T, \kappa) := \begin{cases} \operatorname{sgn}(S_t(H - 1/2))|S_t|^\kappa, & t \in [0, T/2), \\ -\frac{1}{T/2} \int_0^{T/2} \phi_s ds, & t \in [T/2, T] \end{cases} \quad (6.5)$$

satisfy

$$\lim_{T \rightarrow \infty} \frac{EX_T^0(\phi(T, \kappa))}{T^{H(1+\kappa)+1}} > 0, \quad (6.6)$$

(iii) *Let $m > 1/(\alpha - 1)$ be an integer. There exists an increasing sequence $T_n, n \geq m$ tending to ∞ such that the trading strategies*

$$\phi^T = \phi(T, 1/(\alpha - 1) - 1/n), \quad T_n \leq T < T_{n+1}, \quad n \geq m; \quad \phi^T := 0, \quad 0 \leq T < T_m$$

are asymptotically optimal (and the T_n can be explicitly calculated).

(iv) *The strategies $\phi(T, \kappa)$ defined above satisfy*

$$\lim_{T \rightarrow \infty} \frac{\operatorname{Var}(X_T^0(\phi(T, \kappa)))}{T^{2H(\kappa+1)+2}} = D$$

for some constant $D = D(H, \alpha, \kappa, \lambda) \in (0, \infty)$, where $D(H, \alpha, \kappa, \lambda)$ is uniformly bounded in κ . Hence, their Sharpe ratios

$$\sup_{T > 0} \frac{EX_T^0(\phi(T, \kappa))}{\sqrt{\operatorname{Var}(X_T^0(\phi(T, \kappa)))}}$$

are uniformly bounded in T .

Instead of presenting the proof, here, we focus on the essential meaning of the theorem and in the next chapters we will show, based on numerical simulations, how the solutions work.

The first statement (i) of the theorem gives that the optimal strategy asymptotically grows as $T^{H(1+1/(\alpha-1))+1}$. Less rigorously, it means that the optimal solution is $u(T) \approx T^{H(1+1/(\alpha-1))+1}$ as T goes to ∞ .

The second statement (ii) shows a simple strategy that can approach arbitrary close the maximal expected profit. Since $0 < \kappa < 1/(\alpha - 1)$, the speed of $EX_T^0(\phi(T, \kappa))$ which is $T^{H(1+\kappa)+1}$ is always less than $T^{H(1+1/(\alpha-1))+1}$. We call the variable κ as the *trading intensity*.

Following the second statement a natural idea arises to construct an optimal strategy with a continuously increasing trading intensity, that is, κ depends on time and converges to $1/(\alpha - 1)$. This is stated in (iii), but as we will see in the next section, the previously presented and simpler strategy is sufficient.

The investment problem under consideration is particular because the price of the stock can be negative (with 1/2 probability). However, this is true e.g. in futures markets. Therefore we need to assure that the investment is financially reasonable. The fourth statement (iv) states that the Sharpe ratio of the investment is uniformly bounded in T . This is a regular way to qualify an investment. As the theorem showed, the Sharpe ratio is the fraction of the expected (excess) return and its standard deviation. For a brief and exhaustive introduction of the Sharpe ratio we refer the reader to [19].

6.3 Numerical results

This section takes a closer look at the performance of the strategies proven to be asymptotically optimal in Theorem 6.2.1. The questions that we answer here can be divided into 3 groups: 1) How good are the strategies for finite T , because the theorem works only for the limit $T \rightarrow \infty$? 2) How realistic is the Sharpe ratio of the investment? 3) What are the effects of the parameters? Reported results are obtained from Monte Carlo simulations of 10^4 paths of fractional Brownian motion, each of them with 10^3 discretization intervals. The number of paths and the discretization is sufficient, because with larger numbers we obtain the same result.

As we earlier noted, the parameters can be grouped by whether they describe the financial environment $(\alpha, \sigma, \lambda, H)$ or the strategy of the investor (κ, T) . For λ we tested values between 0.0001 and 0.01, settling for the upper bound 0.01. Likewise, results do not change significantly for horizons greater than the value $T = 1$ used hereafter. To match the variation typical of stock prices, we simulated the asset price $S_t = \sigma B_t^H$ with σ in the range 0.2–0.5 and found that results were robust, thereby settling for $\sigma = 0.5$. It is still a question what is the exact value of α and how it can be estimated. Most results suggest its value is between 1 and 2, but the case $\alpha = 1$ is ill-posed, therefore we investigated the case when α

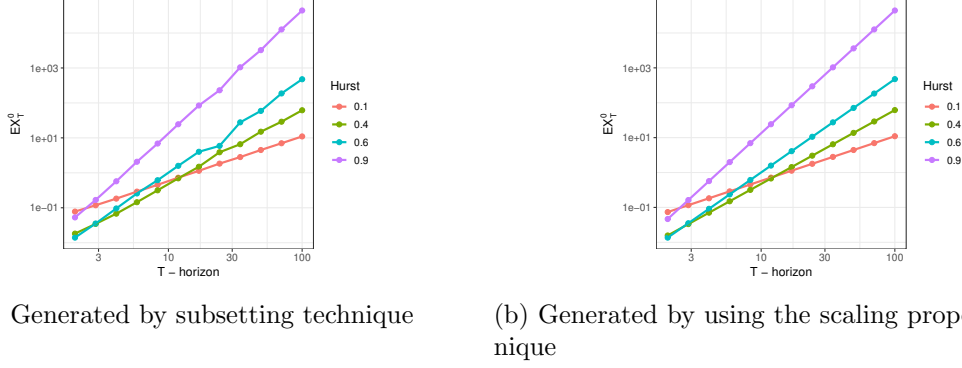


Figure 6.1: Expected payoff (vertical axis, in logarithmic scale) against horizon (horizontal axis, in logarithmic scale). $\alpha = 2$, $\sigma = 0.5$, $\lambda = 0.01$, $\kappa = 0.9 \frac{1}{\alpha-1}$.

is in the range 1.1–2. The possible values of κ depends on α , because of this we introduce the *critical trading intensity* variable $\kappa_c := 1/(\alpha - 1)$ and then we examine the role of the scaled variable κ/κ_c which is between 0 and 1.

Figure 6.1 studies the dependence of the Sharpe ratio on the time horizon considered. Theorem 3.2 asserts that the optimal growth rate for the expected value of our wealth accumulated till T is $T^{H(1+\kappa)+1}$ for large T , which means that $\mathbb{E}[X_T^0]$ as a function of T should be approximately linear on a log-log scale. Figure 6.1 confirms this implication, not only true for large T but also for relatively small horizons of a few years, as the slopes of the lines closely approximate the value $H(1 + \kappa) + 1$ in the theorem. The function $T \rightarrow \mathbb{E}[X_T^0]$ can be simulated in two different ways.

First, we generate the 10^4 paths split into 10^3 intervals/time steps, which refers to the entire time-domain $[0, T]$. Then every subset of the generated paths relates to the proportional time-interval. For example, the first $10^3/2$ time steps relate to the $[0, T/2]$ interval. This can be seen on 6.1a (*subsetting technique*). Second, we can use the scaling property of the fractional Brownian motion, $B_t = |a|^H B_{at}$ in distribution. This is shown on 6.1b (*scaling technique*). Both techniques can be viewed as an estimation of the desired expected value, but the latter one uses more data, therefore it is more precise.

The two questions that how realistic is the Sharpe ratio and how does the investment behave for different parameters are studied on Figure 6.2 and 6.3. We define the Sharpe ratio of the strategies $\phi(T, \kappa)$ again:

$$EX_T^0(\phi(T, \kappa)) / \sqrt{\text{Var}(X_T^0(\phi(T, \kappa)))},$$

which converges for T large in view of Theorem 6.2.1.

Figure 6.2 shows that such Sharpe ratio remains stable for price impact away from the




Figure 6.2: Sharpe Ratio (values in diagram) against Hurst exponent H (horizontal axis) and the price impact α (vertical axis), for the moderate range $H \in [0.1, 0.9]$ (left panel) and for extreme values (right panels, from left $H \approx 0$ to right $H \approx 1$). Parameters: $T = 1, \sigma = 0.5, \lambda = 0.01, \kappa = 0.9\kappa_c$.

boundary $\alpha = 1$ for moderate values of H (top panel). On the contrary, dependence on price impact becomes significant for extreme values of H (bottom panel), leading to very high Sharpe ratios for H near zero but to rather modest values for H near one. Figure 6.3 graphically summarizes these findings, showing how $\alpha = 1.1$ leads to mediocre Sharpe ratios, while $\alpha = 1.5$ and $\alpha = 2$ result in similarly large ratios.

The results in Figure 6.2 are obtained from strategies with $\kappa = 0.9\kappa_c$, and the right panel in Figure 6.3 studies the sensitivity of the Sharpe ratio to such parameter κ . The sensitivity to κ is higher for more nonlinear price impact (α closer to one) and very weak for linear impact (α close to two). Sensitivity is highest for a low Hurst exponent (negative return autocorrelation) combined with a highly nonlinear impact, in which low trading rates are penalized and hence lower value of κ/κ_c are more effective.

6.4 Extension

Without going into detail, we present one extension of the above mentioned model. Other extensions exist on the topic, for further details we suggest to read [56]. In this chapter until now we were trading with a stock price that is a fractional Brownian motion. Its autocorrelation structure was governed by the Hurst parameter which also determined the decay of the autocorrelation function. The main theorem can be extended to other types of processes where less strict requirements are required about the correlation function.

In the following assumption we introduce a definition of long memory Gaussian processes which include, for $H > 1/2$, B_t^H as well.

Assumption 6.4.1. *The process $S_t, t \geq 0$ is zero-mean Gaussian with stationary increments*

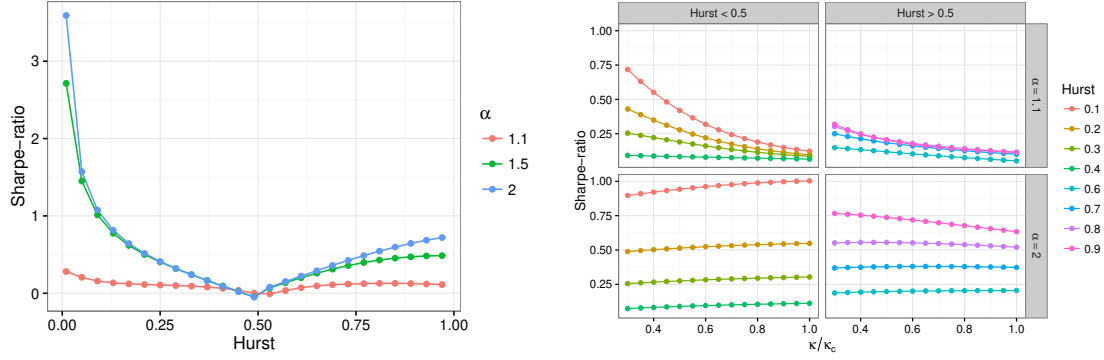


Figure 6.3: Left panel: Sharpe ratio (vertical axis) against Hurst exponent H for different price impact α . Right panel: Sharpe ratio (vertical axis) against strategy exponent $\kappa/\kappa_c \in (0, 1)$ (horizontal) for different price impact α and Hurst exponent H . Parameters: $T = 1, \sigma = 0.5, \lambda = 0.01$. Left panel: $\kappa = 0.9\kappa_c$.

such that

$$\text{cov}(S_u - S_t, S_t) \geq 0, \quad 0 \leq t \leq u. \quad (6.7)$$

Define the covariance function of its increments by

$$r(k) := \text{cov}(S_1 - S_0, S_{k+1} - S_k), \quad k \in \mathbb{N}.$$

Assume that it satisfies

$$r(0) = 1, \quad J_1 k^{2H-2} \leq r(k) \leq J_2 k^{2H-2}, \quad k \geq 1, \quad (6.8)$$

for some constants $J_1, J_2 > 0$ and $H \in (1/2, 1)$.

This assumption is extremely helpful if we wish to define a stock price S_t by its increments. Moreover, in numerical simulations it is enough to know the increments of the process without knowing the exact structure of the whole continuous process. In the following example we focus only on the definition of a discrete-time process which describe the stationary increments of the price process.

Example 6.4.1. Let the Gaussian process X_t be a ARFIMA(p, d, q) process, which is stationary when the memory parameter $d \in (-1/2, 1/2)$. The ARFIMA(p, d, q) process can be obtained by using an ARMA(p, q) process Z_t . Let $X_0 = 0$ and define recursively

$$(1 - B)^d X_t = Z_t, \quad t \geq 1. \quad (6.9)$$

The backshift operator defined by $BX_t = X_{t-1}$ and the term $(1 - B)^d$ is the fractional

difference defined by $(1 - B)^d := \sum_{k=0}^{\infty} \binom{d}{k} (-B)^k$.

This example does not satisfy the Assumption 6.4.1 because of the continuous path requirements. However, for practical reasons we will investigate in the following this process (along with the process DGSV). The process $ARFIMA(p, d, q)$ is frequently used in investment practice and its statistical properties are well-known.

There is a strong relation between the fractional Brownian motion and the process in the above example. The process $ARFIMA(0, d, 0)$ sometimes called fractional white noise equals to the increments of the fractional Brownian motion in distribution. More precisely,

$$B_t^H - B_{t-1}^H \sim ARFIMA(0, H - 1/2, 0). \quad (6.10)$$

As we can see, there is a one-to-one function between the Hurst parameter H and the memory parameter d . This relationship is for other values of p and q , because the autocorrelation function of an $X_t \sim ARFIMA(p, d, q)$ process is

$$r_X(k) \approx L(k)k^{2d-1}, \quad k \rightarrow \infty, \quad (6.11)$$

where $L(k)$ is a slow-varying function, from which we conclude that

$$H = d + 1/2 \quad (6.12)$$

regarding the decaying speed of the autocorrelation function. As the following theorem shows, the limit behavior of the autocorrelation function determines the strategy.

Theorem 6.4.2. *Let $\lambda > 0, \alpha > 1$. Let Assumption 6.4.1 be valid. Then*

(i) *Maximal expected profits satisfy*

$$\limsup_{T \rightarrow \infty} \frac{u(T)}{T^{H(1+1/(\alpha-1))+1}} < \infty. \quad (6.13)$$

(ii) *For each $0 < \kappa < 1/(\alpha - 1)$, the strategies*

$$\phi_t(T, \kappa) := \begin{cases} \operatorname{sgn}(S_t)|S_t|^\kappa, & t \in [0, T/2), \\ -\frac{1}{T/2} \int_0^{T/2} \phi_s(T, \kappa) ds, & t \in [T/2, T] \end{cases} \quad (6.14)$$

satisfy

$$\lim_{T \rightarrow \infty} \frac{EX_T^0(\phi(T, \kappa))}{T^{H(1+\kappa)+1}} > 0, \quad (6.15)$$

where T ranges over integers divisible by 4.

(iii) Let $m > 1/(\alpha - 1)$ be an integer. There exists an increasing sequence T_n , $n \geq m$ tending to ∞ such that the trading strategies

$$\phi^T = \phi(T, 1/(\alpha - 1) - 1/n), \quad T_n \leq T < T_{n+1}, \quad n \geq m; \quad \phi^T := 0, \quad 0 \leq T < T_m$$

are asymptotically optimal (and the T_n can be explicitly calculated).

Theorem 6.4.2 is almost identical to Theorem 6.2.1, but the assumptions are changed. The main messages of this theorem is that the same trading strategy needed and that the optimal expected profit EX_T^0 grows with a power-law in T . The speed of this power-law function solely depends on the characteristic of the autocorrelation of the process. It is worth investigating other processes besides fractional Brownian motion.

Numerical results

In the Example 6.4.1 the process $X_t \sim ARFIMA(p, d, q)$ was characterized by the memory term d instead of the Hurst-parameter. Because of the structure of the autocorrelation function in the limit $k \rightarrow \infty$, one would expect that if the Assumption 6.4.1 is true for an $ARFIMA(p, d, q)$ process, then the optimal expected profit is proportional to

$$T^{(d-1/2)(1+1/(\alpha-1))+1}, \quad (6.16)$$

while the expected profit in case of trading with intensity κ is

$$T^{(d-1/2)(1+\kappa)+1}. \quad (6.17)$$

Because of the lack of continuous paths, Assumption 6.4.1 is not valid, therefore we expect a deviation from these growth factors. Beside the process $ARFIMA(p, d, q)$, we also investigated the case when the increments come from a DGSV process.

Figure 6.4a and 6.4b show the power-law dependency of EX_T^0 on T . The straight lines show that in both model the expected payoff grow as a power of T . In the $ARFIMA(p, d, q)$ case the slopes are steeper the memory parameter d is higher, similarly to the Hurst parameter. The results are further studied and showed in Table 6.1. The exponent of T is still depends linearly on the intensity κ and almost linearly on the memory parameter d , but it is different from $(d - 1/2)(1 + \kappa) + 1$. Table 6.1 shows in the first two columns the Hurst parameter and the exponent in the fractional Brownian motion case, while the third and fourth column show d and the numerical results of the exponent. The exponent slightly differs from the fractional Brownian motion case: if $d < 0$, the exponent is higher, when $d > 0$, the exponent is lower than expected.

In the case of the DGSV process we do not have any connection between the *Hurst*




Figure 6.4: Expected payoff (vertical axis, in logarithmic scale) against horizon (horizontal axis, in logarithmic scale). $\alpha = 2$, $\sigma = 0.5$, $\lambda = 0.01$, $\kappa = 0.9 \frac{1}{\alpha-1}$.

parameter and the process's memory parameter b . The results with the DGSV model does not show any variety if b changes (despite the fact that if $b > 1$, the long memory behavior turns into short memory). This is because the DGSV model was constructed in a way, to have insignificant autocorrelation for H_t . That is, long memory effect in the volatility is not described by the method of this chapter.

Fractional Brownian motion		$ARFIMA(3, d, 5)$		κ/κ_c
H	$H(1 + \kappa) + 1$	d	slope of the function $\log(T) \sim \log(EX_T^0)$	
0.1	1.16	-0.4	1.35	0.3
0.4	1.64	-0.1	1.91	
0.6	1.96	0.1	1.85	
0.9	2.44	0.4	1.77	
0.1	1.22	-0.4	1.43	0.6
0.4	1.88	-0.1	2.14	
0.6	2.32	0.1	2.15	
0.9	2.98	0.4	2.25	
0.1	1.28	-0.4	1.50	0.9
0.4	2.12	-0.1	2.36	
0.6	2.68	0.1	2.51	
0.9	3.52	0.4	2.70	

Table 6.1: Results of the linear fit on the $\log(T) \sim \log(EX_T^0)$ plot for the model $ARFIMA(3, d, 5)$ based on numerical simulations (4th column). The result is compared for each value of d with the corresponding Hurst parameter and what the fit would be for the fractional Brownian motion. In the antipersistence case ($H < 0.5$ or $d < 0$), model ARFIMA outperforms the fractional Brownian motion case, while in the positive memory case ($H > 0.5$ or $d > 0$) is the opposite.

Appendix A

Proof of Stationary Solution and Ergodicity-like Assumption

In Chapter 2 we defined new and already existing long memory stock return models as well. However, their definitions do not guarantee that they define mathematically well-defined stochastic processes.

We also have not proven yet that these processes satisfy our ergodicity-like criteria in Assumption 3.2.3 which guarantees that the log-optimal portfolio exists. From this assumption it follows, that with optimal strategy π^* from Theorem 3.2.2,

$$\liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\log \left(W_T^{\pi^*} \right) \right] = \lim_{T \rightarrow \infty} \frac{1}{T} \log \left(W_T^{\pi^*} \right). \quad (\text{A.1})$$

That is, instead of calculating the expected value for every T and taking the limit, it is enough to determine the optimal strategy for only one realization.

In this chapter we show that these processes are well-defined and stationary processes in the strong sense. Moreover, we prove that these models satisfy Assumption 3.2.3.

A.1 DGSV model

The first theorem states that the stationary solution of the DGSV dynamics, defined in (2.9), exists almost surely.

Theorem A.1.1. *There exists a stationary process \tilde{X}_t , $t \in \mathbb{Z}$ satisfying*

$$\tilde{X}_{t+1} = \mu + \alpha \tilde{X}_t + \rho e^{Y_t} \eta_{t+1} + \sqrt{1 - \rho^2} e^{Y_t} \varepsilon_{t+1}, \quad (\text{A.2})$$

such that \tilde{X}_t is $\sigma(\eta_j, \varepsilon_j, j \geq t)$ -measurable. Defining

$$Q_t := g^{\hat{\pi}(\tilde{X}_t, \nu_t, \kappa_t)}(\tilde{X}_t, \nu_t, \kappa_t),$$

and

$$F(z, y, \varepsilon, \eta) := \mu + \alpha z + \rho e^y \eta + \sqrt{1 - \rho^2} e^y \varepsilon,$$

Assumptions 3.2.1 and 3.2.3 are valid.

Proof. We follow the arguments of [58] closely. Consider the mapping

$$(x, y, \eta, \varepsilon) \rightarrow L(x, y, \eta, \varepsilon) := \mu + \alpha x + \rho e^y \eta + \sqrt{1 - \rho^2} e^y \varepsilon,$$

where $x, y, \eta, \varepsilon \in \mathbb{R}$. This satisfies

$$|L(x_1, y, \eta, \varepsilon) - L(x_2, y, \eta, \varepsilon)| \leq |\alpha| |x_1 - x_2|$$

so it is a contraction. For each $n \geq 1$ and for each $t \in \mathbb{Z}$, define, by recursion,

$$X_{t-n}^{t,n} := 0, \quad X_{k+1}^{t,n} := L(X_k^{t,n}, Y_{k+1}, \eta_{k+1}, \varepsilon_{k+1}), \quad k \geq t - n.$$

Set $\tilde{X}_t^n := X_t^{t,n}$. Since

$$|\tilde{X}_t^n - \tilde{X}_t^{n+1}| \leq |\alpha|^n |\rho e^{Y_{t-n}} \eta_{t-n} + \sqrt{1 - \rho^2} e^{Y_{t-n}} \varepsilon_{t-n}|,$$

and the right-hand side is bounded in L^1 , it follows by the Markov inequality that

$$\sum_{n=1}^{\infty} P(|\tilde{X}_t^n - \tilde{X}_t^{n+1}| \geq \sqrt{|\alpha|^n}) \leq \sum_{n=1}^{\infty} C \frac{|\alpha|^n}{\sqrt{|\alpha|^n}} < \infty,$$

with some $C > 0$. The Borel–Cantelli lemma implies that the sequence \tilde{X}_t^n , $n \in \mathbb{N}$ is a.s. a Cauchy sequence, hence it is convergent almost surely to some \tilde{X}_t .

Since \tilde{X}_t^n , $t \in \mathbb{Z}$ is easily seen to be a stationary process for each $n \geq 1$, the process \tilde{X}_t , $t \in \mathbb{Z}$ is stationary, too. Notice that $\tilde{X}_{t+1}^{n+1} = L(\tilde{X}_t^n, Y_{t+1}, \eta_{t+1}, \varepsilon_{t+1})$ and L is continuous in its first variable so we conclude that (A.2) indeed holds true. The claim about measurability is clear from the construction of \tilde{X} . Theorem 3.5.8 of [59] implies that \tilde{X} is ergodic. Actually, that theorem implies also that $G_t := (\tilde{X}_t, Y_t, \eta_t, \varepsilon_t)$, $t \in \mathbb{Z}$ is ergodic since G_t is a functional of $(\varepsilon_t, \eta_t, \varepsilon_{t-1}, \eta_{t-1}, \dots)$. Furthermore, as Q_t is a functional of $(\tilde{X}_t, Y_t, \eta_t, \varepsilon_t)$ it is also ergodic.

We now claim $\mathbb{E}|Q_t| < \infty$. Notice that, using the Cauchy inequality,

$$\mathbb{E}|\alpha z + \rho e^{\nu + \kappa \varepsilon_0} \varepsilon_0 + \sqrt{1 - \rho^2} \eta_0| \leq |z| + e^{2\nu + 2\kappa^2} \times [1 + \mathbb{E}\eta_0^2] \quad (\text{A.3})$$

and the sequence $e^{2\nu t}$ is clearly bounded in L^1 . Hence, by Lemma 3.2.1, it suffices to show that the sequence \tilde{X}_t^n is bounded in L^1 . Notice that

$$\mathbb{E}[X_{k+1}^{n,t}] \leq |\alpha| \mathbb{E}[X_k^{n,t}] + \mathbb{E}[e^{Y_{k+1}}[|\eta_{k+1}| + |\varepsilon_{k+1}|]] \leq |\alpha| \mathbb{E}[X_k^{n,t}] + C$$

for some $C > 0$ which easily implies that all the $X_k^{n,t}$ are bounded in L^1 , which shows our claim.

Now Assumption 3.2.1 simply holds as

$$\sup_{z, \nu, \kappa \in [-N, N]^3} |F(z, \nu + \kappa \varepsilon_0, \varepsilon_0, \eta_0)| \leq N + e^{N + \kappa |\varepsilon_0|} \times [|\varepsilon_0| + |\eta_0|]$$

which clearly has a finite expectation. □

A.2 BARCH model

We do not prove that there exists a stationary and ergodic solution for the model in (2.18) because the model has a more general form described in [36].

Theorem A.2.1. *There exists a stationary process \tilde{X}_t , $t \in \mathbb{Z}$ satisfying*

$$\tilde{X}_t = \mu + \alpha \tilde{X}_{t-1} + \left(a + b_0 \sum_{j=1}^{\infty} j^{-b} \tilde{X}_{t-j} \right) \eta_t, \quad (\text{A.4})$$

such that $\mu \neq 0$, $\eta \sim \mathcal{N}(0, 1)$ and $|\alpha| < 1$. We also assume that the constants satisfy the inequality

$$\left(\frac{b_0}{1 - \alpha} \zeta(2b) \right)^2 < 1, \quad (\text{A.5})$$

where the function $\zeta(\cdot)$ is the (Euler-)Riemann zeta function, which is defined as $\zeta(s) := \sum_{n=1}^{\infty} n^{-s}$. Moreover, this process is also ergodic and its expected value is $\mathbb{E}[\tilde{X}_t] = \mu / (1 - \alpha)$.

Proof. See Propositions 2.4 and 2.7 of [36]. □

Ergodicity of the log-return and its finite first moment is a necessary condition but not sufficient. The proof of the previous model does not work here, because the function $Z(\cdot)$ has infinitely many variables $(H_t, H_{t-1}, \dots, \eta_{t+1})$. Nevertheless, numerical results suggest that ergodicity of H_t in the BARCH model is also sufficient.

Appendix B

Moments of the Discrete Gaussian Stochastic Volatility

Typically, the more complex a stochastic process's dynamics, the more difficult it is to calculate its moments. In addition to the DGSV model having realistic properties, we are able to calculate its conditional and unconditional moments. Since the calculations are similar for every moment, only the conditional mean and mean are presented in details.

We recall the definition of the DGSV model:

$$H_t = \mu + \alpha H_{t-1} + \sigma e^{Y_t} \left(\rho \varepsilon_t + \sqrt{1 - \rho^2} \eta_t \right). \quad (\text{B.1})$$

In order to calculate the unconditional moments, we need to assume that the process is stationary. In Appendix A we proved that the stationary solution exists, therefore we can assume without loss of generality that the process H_t is stationary. (In simulations this is usually achieved by starting from an arbitrary point, then dropping the first several values when the process is far from its stationary distribution.)

As we mentioned in Subsection 2.5.3, the key aspect to calculating the moments is the appearance of the term $\mathbb{E}[e^{K\varepsilon} \varepsilon^k]$ during the derivation of the k -th moments, where K is a real number and ε is a Gaussian variable with zero mean and unit variance. We remind the reader about the following identity

$$\mathbb{E}[e^{K\varepsilon} \varepsilon^j] = e^{K^2/2} \mathbb{E}[(\varepsilon + K)^j].$$

For example, when $j = 1$, the expected value is equal to $e^{K^2/2} K$, and for $j = 2$, it is equal to $e^{K^2/2} (1 + K^2)$.

B.1 Conditional mean

We are interested in the conditional mean of the DGSV process H_t . Taking the conditional expectation of both sides in (B.1) we have

$$\mathbb{E}[H_t|\mathcal{F}_{t-1}] = \mu + \alpha\mathbb{E}[H_{t-1}|\mathcal{F}_{t-1}] + \mathbb{E}\left[\sigma e^{Y_t}\left(\rho\varepsilon_t + \sqrt{1-\rho^2}\eta_t\right)|\mathcal{F}_{t-1}\right],$$

where we made use of the fact that the conditional expectation operator is linear. Since H_{t-1} is \mathcal{F}_{t-1} -measurable and that $\mathbb{E}[\eta_t|\mathcal{F}_{t-1}] = 0$, we get that

$$\mathbb{E}[H_t|\mathcal{F}_{t-1}] = \mu + \alpha H_{t-1} + \sigma\rho\mathbb{E}[e^{Y_t}\varepsilon_t|\mathcal{F}_{t-1}].$$

The variable Y_t is defined as a linear function of the Gaussian i.i.d noises $\varepsilon_t, \varepsilon_{t-1}, \dots$, that is $Y_t := \sum_{j=0}^{\infty} \beta_j \varepsilon_{t-j}$. Moreover, Y_t equals $\nu_{t-1} + \kappa\varepsilon$ in distribution, where $\nu_{t-1} := \sum_{j=1}^{\infty} \beta_j \varepsilon_{t-j}$ and $\kappa = \beta_0$. It follows that

$$\mathbb{E}[e^{Y_t}\varepsilon_t|\mathcal{F}_{t-1}] = \mathbb{E}[e^{\nu_{t-1}+\kappa\varepsilon}\varepsilon|\mathcal{F}_{t-1}] = e^{\nu_{t-1}}\mathbb{E}[e^{\kappa\varepsilon}\varepsilon] = e^{\nu_{t-1}}\kappa e^{\kappa^2/2};$$

in the second equation we used the property that ε is independent of the σ -algebra \mathcal{F}_{t-1} .

From the above equations we can conclude that the conditional mean is

$$\mathbb{E}[H_t|\mathcal{F}_{t-1}] = \mu + \alpha H_{t-1} + \rho\sigma\kappa e^{\nu_{t-1}+\kappa^2/2}. \quad (\text{B.2})$$

B.2 Unconditional mean

So far, we have not made use of the process H_t being a stationary process. From the stationarity property, it follows that the mean is constant in time that is $\mathbb{E}[H_t] = \mathbb{E}[H_{t-1}]$ for all $t \in \mathbb{Z}$.

Using the *tower-rule* we can calculate the unconditional mean of the process H_t . For those readers who are unfamiliar with term tower-rule, we suggest Appendix C. From the tower-rule it follows that

$$\mathbb{E}[\mathbb{E}[H_t|\mathcal{F}_{t-1}]] = \mathbb{E}[H_t].$$

Let us take the expected value of both sides in (B.2) then we have

$$\mathbb{E}[H_t] = \mu + \alpha\mathbb{E}[H_{t-1}] + \mathbb{E}\left[\rho\sigma\kappa e^{\nu_{t-1}+\kappa^2/2}\right].$$

Since $\mathbb{E}[H_t]$ equals $\mathbb{E}[H_{t-1}]$

$$\mathbb{E}[H_t] = \mu + \alpha \mathbb{E}[H_t] + \mathbb{E} \left[\rho \sigma \kappa e^{\nu_{t-1} + \kappa^2/2} \right].$$

After grouping the terms which contain the expected value on the left hand side, we have:

$$\mathbb{E}[H_t] - \alpha \mathbb{E}[H_t] = \mu + \mathbb{E} \left[\rho \sigma \kappa e^{\nu_{t-1} + \kappa^2/2} \right].$$

We assumed that $|\alpha| < 1$, therefore the term $(1 - \alpha)$ is never zero and we can divide by it, hence

$$E[H_t] = \frac{\mu + \mathbb{E} \left[\rho \sigma \kappa e^{\nu_{t-1} + \kappa^2/2} \right]}{1 - \alpha}. \quad (\text{B.3})$$

In the numerator the constants ρ, σ, κ and $e^{\kappa^2/2}$ can be written outside of the expected value: $\rho \sigma \kappa e^{\kappa^2/2} \mathbb{E} [e^{\nu_{t-1}}]$. Since ν_{t-1} is a linear function of i.i.d. Gaussian variables with unit variance and zero mean, therefore

$$\mathbb{E} [e^{\nu_{t-1}}] = \mathbb{E} \left[e^{\sum_{j=1}^{\infty} \beta_j \varepsilon_{t-j}} \right] = \mathbb{E} \left[\prod_{j=1}^{\infty} e^{\beta_j \varepsilon_{t-j}} \right] = \prod_{j=1}^{\infty} \mathbb{E} \left[e^{\beta_j \varepsilon_{t-j}} \right] = \prod_{j=1}^{\infty} e^{\beta_j^2/2} = e^{\sum_{j=1}^{\infty} \beta_j^2/2}.$$

As $\kappa = \beta_0$ we conclude that the mean is

$$E[H_t] = \frac{\mu + \rho \sigma \beta_0 e^{\sum_{j=0}^{\infty} \beta_j^2/2}}{1 - \alpha}. \quad (\text{B.4})$$

The reason for using two different notations for the same constant, κ and β_0 , which are equal to each other, is to emphasize that the current value in the conditional moments depends on the conditional standard deviation of Y_t , namely κ , while in the unconditional moments we use the constants from the definition of that process, namely β_0 .

Appendix C

Conditional Probability

Because the concept of conditional probability is fundamental to understanding all chapters of this thesis, in this chapter we would like to briefly introduce the topic of conditional probability and expectation for the benefit of those who might not have a strong background in mathematics.

A substantial part of the thesis focuses on numerical results; therefore, it is essential to connect the probability theory concepts with their numerical implementations.

Section C.1 introduces the concept of conditional expected value at such a level that it is needed to understand for the statements of the thesis (especially in Chapter 3). Then, in section C.2, we describe how we calculated numerically certain concepts of probability theory.

C.1 Conditional expected value in general

We present the general definition in the case when the random variable X has finite second moment. In this case, conditional expectation is equivalent to an orthogonal projection.

Definition C.1.1. *Let X be a random variable defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with finite second moments. The conditional expectation $Z : \mathbf{E}[X|\mathcal{G}]$ is an orthogonal projection on the sub-space $(\Omega, \mathcal{G}, \mathbb{P})$.*

A property of this definition is, if the random variable Y is defined on $(\Omega, \mathcal{G}, \mathbb{P})$, then

$$\mathbb{E}[XY] = \mathbb{E}[ZY].$$

This formula is used to further generalize the definition which we exclude here.

Important properties of the conditional expectation:

- Linearity: $\mathbb{E}[aX_1 + bX_2|\mathcal{G}] = a\mathbb{E}[X_1|\mathcal{G}] + b\mathbb{E}[X_2|\mathcal{G}]$ for any $a, b \in \mathbb{R}$.

- Best prediction: the random variable $\mathbb{E}[X|\mathcal{G}]$ minimizes the expression $\mathbb{E}[(X - Y)^2]$ among all \mathcal{G} -measurable random variable Y .
- Tower rule: Let $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{F}$ increasing σ -algebras, then $\mathbb{E}[\mathbb{E}[X|\mathcal{G}|\mathcal{H}] = \mathbb{E}[X|\mathcal{H}]$.

The mostly used property in this dissertation is the last one. We note that it also includes the case when $\mathcal{H} = \mathcal{F}$, that is, instead of calculating the expected value, sometimes it is easier for first to calculate a conditional expected value then taking its expected value.

C.2 Numerical implementation of expectations

In numerical simulations it is often required to estimate the values of the moments, either conditional or unconditional. In the case of a random variable, the expected value is usually estimated by the average of an i.i.d. sample $X(1), X(2), \dots, X(N)$:

$$\frac{1}{N} \sum_{i=1}^N X(i).$$

The i -th realization of the random variable X is denoted as $X(i)$. In a general settings where an i.i.d. sample is not available it may fail to give an unbiased estimation of the expected value. In the formal chapters the numerical approximations gave unbiased estimations but we exclude its proof here.

For stochastic processes, the estimation consists of two things: average in time and average in realizations. Then the expected value of a stationary process X_t is estimated by

$$\frac{1}{TN} \sum_{t=1}^T \sum_{i=1}^N X_t(i).$$

Since we chose the initial value to be a constant, therefore it takes a while for the process to be stationary. For this reason we typically chose a much larger T than N .

For example, in Chapter 5 where we calculated the function $g(\theta) := \mathbb{E}[H_t \mathbb{1}_{\{H_{t-1} > \theta\}}]$, it was estimated as

$$\frac{1}{TN} \sum_{t=1}^T \sum_{i=1}^N H_t(i) \mathbb{1}_{\{H_{t-1}(i) > \theta\}},$$

for some values of θ .

The autocorrelation function was estimated in two way:

$$\frac{1}{T - \tau} \sum_{t=1}^{T-\tau} X_{t+\tau}(i) X_t(i) - \left(\frac{1}{T} \sum_{t=1}^T X_t(i) \right)^2,$$

if we were interested in the autocorrelation of only one realization (like in Chapter 4) and the other way was

$$\frac{1}{N} \sum_{i=1}^N \left[\frac{1}{T-\tau} \sum_{t=1}^{T-\tau} X_{t+\tau}(i)X_t(i) - \left(\frac{1}{T} \sum_{t=1}^T X_t(i) \right)^2 \right].$$

Bibliography

- [1] R. L. Kissell, *The science of algorithmic trading and portfolio management*. Academic Press, 2013.
- [2] M. T. Greene and B. D. Fielitz, “Long-term dependence in common stock returns,” *Journal of Financial Economics*, vol. 4, no. 3, pp. 339–349, 1977.
- [3] P. H. Algoet, T. M. Cover *et al.*, “Asymptotic optimality and asymptotic equipartition properties of log-optimum investment,” *The Annals of Probability*, vol. 16, no. 2, pp. 876–898, 1988.
- [4] H. Markowitz, “Portfolio selection,” *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
- [5] J. L. Kelly, “A new interpretation of information rate,” *Bell System Technical Journal*, vol. 35, no. 4, pp. 917–926, 1956.
- [6] D. Bernoulli, “Exposition of a new theory on the measurement of risk,” *Econometrica: Journal of the Econometric Society*, pp. 23–36, 1954.
- [7] T. M. Cover and J. A. Thomas, *Elements of information theory*, 2nd ed. Wiley-Interscience, Hoboken, NJ, 2006.
- [8] F. Comte and E. Renault, “Long memory in continuous-time stochastic volatility models,” *Mathematical Finance*, vol. 8, no. 4, pp. 291–323, 1998.
- [9] J. Gatheral, T. Jaisson, and M. Rosenbaum, “Volatility is rough,” *Quantitative Finance*, vol. 18, no. 6, pp. 933–949, 2018.
- [10] Z. Ding, C. W. Granger, and R. F. Engle, “A long memory property of stock market returns and a new model,” *Journal of Empirical Finance*, vol. 1, no. 1, pp. 83–106, 1993.
- [11] T. G. Andersen and T. Bollerslev, “Intraday periodicity and volatility persistence in financial markets,” *Journal of Empirical Finance*, vol. 4, no. 2, pp. 115–158, 1997.

- [12] T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys, “The distribution of realized exchange rate volatility,” *Journal of the American Statistical Association*, vol. 96, no. 453, pp. 42–55, 2001.
- [13] R. Cont, “Empirical properties of asset returns: stylized facts and statistical issues,” *Quantitative Finance*, vol. 1, no. 2, pp. 223–236, 2001.
- [14] M. M. Christensen, “On the history of the growth optimal portfolio,” *Machine Learning For Financial Engineering*, vol. 389, 2005.
- [15] J. Bernoulli, *Ars conjectandi*. Impensis Thurnisiorum, fratrum, 1713.
- [16] J. Von Neumann, O. Morgenstern, and H. W. Kuhn, *Theory of games and economic behavior (commemorative edition)*. Princeton university press, 2007.
- [17] G. Debreu, “Representation of a preference ordering by a numerical function,” *Decision processes*, vol. 3, pp. 159–165, 1954.
- [18] W. F. Sharpe, “Mutual fund performance,” *The Journal of business*, vol. 39, no. 1, pp. 119–138, 1966.
- [19] —, “The sharpe ratio,” *Journal of portfolio management*, vol. 21, no. 1, pp. 49–58, 1994.
- [20] L. C. MacLean, E. O. Thorp, and W. T. Ziemba, *The Kelly capital growth investment criterion: Theory and practice*. World Scientific, 2011, vol. 3.
- [21] L. Györfi, G. Ottucsák, and A. Urbán, “Empirical log-optimal portfolio selections: a survey,” in *Machine learning for financial engineering*. World Scientific, 2012, pp. 81–118.
- [22] F. Black, “Studies of stock market volatility changes,” *1976 Proceedings of the American Statistical Association Business and Economic Statistics Section*, 1976.
- [23] J.-P. Bouchaud, A. Matacz, and M. Potters, “Leverage effect in financial markets: The retarded volatility model,” *Physical review letters*, vol. 87, no. 22, p. 228701, 2001.
- [24] B. B. Mandelbrot, “The variation of certain speculative prices,” in *Fractals and scaling in finance*. Springer, 1997, pp. 371–418.
- [25] Z. Ding and C. W. Granger, “Modeling volatility persistence of speculative returns: a new approach,” *Journal of econometrics*, vol. 73, no. 1, pp. 185–215, 1996.

- [26] R. F. Engle, “Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation,” *Econometrica: Journal of the Econometric Society*, pp. 987–1007, 1982.
- [27] T. Bollerslev, “Generalized autoregressive conditional heteroskedasticity,” *Journal of econometrics*, vol. 31, no. 3, pp. 307–327, 1986.
- [28] L. Bachelier, “Théorie de la spéculation,” in *Annales scientifiques de l’École normale supérieure*, vol. 17, 1900, pp. 21–86.
- [29] B. Mandelbrot, “New methods in statistical economics,” *Journal of political economy*, vol. 71, no. 5, pp. 421–440, 1963.
- [30] N. Shephard, *Stochastic volatility: selected readings*. Oxford University Press, 2005.
- [31] T. G. Andersen, R. A. Davis, J.-P. Kreiss, and T. V. Mikosch, *Handbook of financial time series*. Springer Science & Business Media, 2009.
- [32] P. M. Robinson, “Testing for strong serial correlation and dynamic conditional heteroskedasticity in multiple regression,” *Journal of Econometrics*, vol. 47, no. 1, pp. 67–84, 1991.
- [33] F. J. Breidt, N. Crato, and P. De Lima, “The detection and estimation of long memory in stochastic volatility,” *Journal of Econometrics*, vol. 83, no. 1, pp. 325–348, 1998.
- [34] F. Comte, C. Lacour, and Y. Rozenholc, “Adaptive estimation of the dynamics of a discrete time stochastic volatility model,” *Journal of Econometrics*, vol. 154, no. 1, pp. 59–73, 2010.
- [35] P. Guasoni, Z. Nika, and M. Rásonyi, “Trading Fractional Brownian Motion,” *Preprint, SSRN:2991275*, 2017.
- [36] L. Giraitis and D. Surgailis, “ARCH-type bilinear models with double long memory,” *Stochastic Processes and their Applications*, vol. 100, no. 1, pp. 275–300, 2002.
- [37] O. Hernández-Lerma and J. B. Lasserre, *Discrete-Time Markov Control Processes*. Springer, 1996.
- [38] I. Vajda, “Analysis of semi-log-optimal investment strategies,” in *Prague Stochastics*, 2006, pp. 719–727.
- [39] B. Li and S. C. Hoi, “Online portfolio selection: A survey,” *ACM Computing Surveys (CSUR)*, vol. 46, no. 3, p. 35, 2014.

- [40] L. Györfi, A. Urbán, and I. Vajda, “Kernel-based semi-log-optimal empirical portfolio selection strategies,” *International Journal of Theoretical and Applied Finance*, vol. 10, no. 03, pp. 505–516, 2007.
- [41] M. H. Davis and S. Lleo, *Risk-sensitive Investment Management*. World Scientific, 2014.
- [42] A. C. Cameron and P. K. Trivedi, *Microeconometrics: methods and applications*. Cambridge University Press, 2005.
- [43] I. Aldridge, *High-frequency trading: a practical guide to algorithmic strategies and trading systems*. John Wiley & Sons, 2013, vol. 604.
- [44] M. L. De Prado, *Advances in financial machine learning*. John Wiley & Sons, 2018.
- [45] H. Robbins and S. Monro, “A stochastic approximation method,” *The annals of mathematical statistics*, pp. 400–407, 1951.
- [46] J. Kiefer, J. Wolfowitz *et al.*, “Stochastic estimation of the maximum of a regression function,” *The Annals of Mathematical Statistics*, vol. 23, no. 3, pp. 462–466, 1952.
- [47] Z. Zhang, G. Yin, and Z. Liang, “A stochastic approximation algorithm for american lookback put options,” *Stochastic Analysis and Applications*, vol. 29, no. 2, pp. 332–351, 2011.
- [48] G. Yin, Q. Zhang, F. Liu, R. Liu, and Y. Cheng, “Stock liquidation via stochastic approximation using nasdaq daily and intra-day data,” *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, vol. 16, no. 1, pp. 217–236, 2006.
- [49] Y. Chow and M. Ghavamzadeh, “Algorithms for cvar optimization in mdps,” in *Advances in neural information processing systems*, 2014, pp. 3509–3517.
- [50] S. Laruelle and G. Pagès, “Stochastic approximation with averaging innovation applied to finance,” *Monte Carlo Methods and Applications*, vol. 18, no. 1, pp. 1–51, 2012.
- [51] A. Kibzun and R. Lepp, “Discrete approximation in quantile problem of portfolio selection,” in *Stochastic Optimization: Algorithms and Applications*. Springer, 2001, pp. 121–135.
- [52] S. Laruelle, C.-A. Lehalle, and G. Pages, “Optimal split of orders across liquidity pools: a stochastic algorithm approach,” *SIAM Journal on Financial Mathematics*, vol. 2, no. 1, pp. 1042–1076, 2011.

- [53] D. Burkholder, “On a class of stochastic approximation processes,” *The Annals of Mathematical Statistics*, pp. 1044–1059, 1956.
- [54] V. Fabian, “Stochastic approximation,” in *Optimizing methods in statistics*. Elsevier, 1971, pp. 439–470.
- [55] J. Sacks, “Asymptotic distribution of stochastic approximation procedures,” *The Annals of Mathematical Statistics*, vol. 29, no. 2, pp. 373–405, 1958.
- [56] P. Guasoni, Z. Nika, and M. Rasonyi, “Trading fractional brownian motion,” *SIAM Journal on Financial Mathematics*, vol. 10, no. 3, pp. 769–789, 2019.
- [57] P. Guasoni, M. Rásonyi *et al.*, “Hedging, arbitrage and optimality with superlinear frictions,” *The Annals of Applied Probability*, vol. 25, no. 4, pp. 2066–2095, 2015.
- [58] P. Diaconis and D. Freedman, “Iterated random functions,” *SIAM Rev.*, pp. 45–76, 1999.
- [59] W. F. Stout, *Almost sure convergence*. Academic Press, New York, 1973.