Tilburg University

Department of Econometrics and Operations Research

Master's Degree in Quantitatve Finance and Actuarial Science

# Intraday Seasonality in the FX Swap Market

MASTER DEGREE THESIS

*Author*
Wasim van Houtum

*Supervisor*
dr. Nikolaus Schweizer

*Co-supervisor*
Tjerk Methorst

ACADEMIC YEAR 2023/2024
August 18, 2024

Wasim van Houtum

# Intraday Seasonality in the FX Swap Market

*A thesis submitted in partial fulfillment of the requirements for the degree of Quantitative Finance and Actuarial Science at Tilburg University.*

*A research conducted as part of the PGGM AXP program.*

TILBURG UNIVERSITY

August 2024

“*It’s Difficult to Make Predictions, Especially About the Future.*”

- Lots of smart people

# Abstract

*Intraday Seasonality in the FX Swap Market*

This thesis investigates intraday seasonality in the FX Swap market, focusing on EUR/USD three-month forward points. We model the conditional mean of forward points returns, emphasizing the identification and analysis of intraday seasonal patterns. A key contribution is the decomposition of returns into sign and magnitude components, treated as separate stochastic processes. Dependencies between these components are modeled using copula theory, offering a rigorous framework for understanding return dynamics.

Our analysis reveals intraday seasonality in return volatility, particularly during the mid-morning session. However, evidence for seasonality in return sign is inconclusive, with small coefficients and limited predictive accuracy. Despite these challenges, the decomposition model provides a promising framework for enhancing predictive performance, particularly through the use of copula methods to model dependencies between return components.

# Contents

# List of Figures

# List of Tables

# Acknowledgements

I am profoundly grateful to *dr. Nikolaus Schweizer* for his steadfast support and expert guidance throughout the completion of this thesis. Our bi-weekly meetings have significantly enriched my research journey and insights.

I wish to extend my heartfelt thanks to *Tjerk Methorst* for his invaluable assistance and encouragement during the preparation of this thesis. His insights, patience and involvements have played a crucial role in shaping my understanding of the financial market and in my academic pursuits, and I am truly appreciative of his mentorship.

I also want to thank *Wiebe Kuper* and *Angelo Barisano* for their support during my internship at PGGM.

# Chapter 1

# Introduction

With an average daily turnover of US$ 3.8 trillion, the Foreign Exchange (FX) Swap stands as the most traded FX instrument globally (BIS, 2022). An FX Swap involves exchanging one currency for another on one date and reversing the exchange on a future date, serving crucial roles in managing funding liquidity and hedging currency risks (ING, 2019). Despite its prominence and the high frequency of trading across the 24-hour, five-day-per-week FX market, information on FX Swaps remains relatively inaccessible due to their Over-the-Counter (OTC) nature, making data collection expensive and challenging (BIS, 2022).

Before defining an FX Swap (hereafter referred to as "swap"), let's consider a motivating example.

Imagine a Dutch investor who currently holds €1,000,000. The investor faces two options: to invest the funds in a European bank for three months at an interest rate of 4%, or to invest in a U.S. bank for the same period at a 5% interest rate. If the investor chooses the European option, the investment will grow to €1,040,000 after three months, calculated as €1,000,000 $\times$ 1.04.

Alternatively, if the investor opts to invest in the U.S., they would first need to convert the euros into dollars at the current spot exchange rate, $S_t$. Assuming the spot rate is $S_t = 1.08$, the investor would convert €1,000,000 into $1,080,000. This amount is then invested at the U.S. interest rate of 5%, resulting in a balance of $1,134,000 after three months.

However, this investment introduces a significant risk known as *currency risk* or *spot risk*. Since the investor ultimately requires euros, they will need to convert the dollars back into euros after the investment period. The exchange rate might fluctuate during these three months, potentially diminishing the investment's value when converted back to euros. For example, if the euro appreciates significantly against the dollar, the $1,134,000 could be worth much less in euros.

To mitigate this currency risk, the investor could enter into a forward rate agreement, which locks in an exchange rate at the time of the initial investment. This forward rate guarantees that the investor can convert dollars back to euros at a predetermined rate, thereby eliminating the spot risk. The forward rate is set to ensure that the investor is indifferent between investing in the U.S. or Europe, assuming a frictionless market with no transaction costs. This idea can be mathematically described by the following equality:

$$\frac{S_t \cdot (1 + r_{\$})}{F_{t,t+1}} = 1 + r_{€}, \tag{1.1}$$

where $F_{t,t+1}$ is the forward rate agreed upon at time $t$ that is applicable at time $t+1$, $S_t$ is the spot rate at time $t$, and $r_{\$}$ and $r_{€}$ are the dollar and euro one-period interest rates, respectively. This equality can be rewritten to the very well-known parity introduced first by Keynes (1923):

**Definition 1.1** (Covered Interest Rate Parity (CIP))**.**

$$F_{t,t+\delta} = S_t \cdot \frac{(1 + r_{\delta,\$})}{(1 + r_{\delta,€})}, \tag{1.2}$$

*The CIP describes the relationship between the future exchange rate of two currencies, considering the spot rate and their respective interest rates.*

Referring back to the example of the Dutch investor, it is now clear how the CIP ensures no-arbitrage in the market. Due to its simplicity, interpretability, and the profound macroeconomic theory underlying it, ensuring no-arbitrage in free markets, CIP is a cornerstone in valuing forwards.

Understanding how CIP relates to swaps is straightforward. A swap is an agreement between two parties to exchange a specified amount of one currency for a specified amount of another currency at a near future date (near leg). The transaction for the near leg is settled at the exchange rate that holds at the outset $t$, which is the spot rate $S_t$. At a future date (far leg), there will be a transaction to trade back the currencies, settled at the forward rate $F_{t,t+1}$ agreed upon at the outset. The concept of a swap is visualized in Figure 1.1 for better understanding.[1]


Figure 1.1: EUR/USD FX Swap Visualization

Given the substantial trading volumes and frequent transactions, one might expect the FX swap market to exhibit a high degree of efficiency compared to smaller, less liquid markets (Clarke et al., 2001). Efficient markets

[1]Instead of dividing by the forward rate, some may multiply by the *ask price* of the forward rate.

ensure that prices reflect all available information, thereby eliminating opportunities for risk-free profits—a principle validated in numerous studies (Oh et al., 2007). However, when violations of market efficiency, known as anomalies, occur, they become targets for exploitation by financial institutions and researchers alike (Kallianiotis, 2018). One such anomaly is the presence of predictable patterns throughout the trading day, which challenges the notion of market efficiency. While a substantial body of research has explored the anomaly of seasonality in the FX spot market, as discussed later, the literature on seasonality in the FX forward market remains sparse. This paper seeks to contribute to this underexplored area.

The primary distinction between analyzing FX spot and FX forward rates lies in their underlying theoretical frameworks, which are discussed in detail in Chapter 2. Spot rates operate under a floating rate system determined by supply and demand, whereas forward rate valuation is derived from a no-arbitrage framework grounded in the *Interest Rate Parity* (IRP) theory first introduced by Keynes (1923). The IRP, which establishes a relationship between interest rate differentials and forward premiums, serves as a benchmark for assessing perfect capital mobility across markets (Levich, 2011). Traditionally, any deviation from IRP indicated stress within the global financial system, as was notably the case during the 2008 financial crisis (Chatziantoniou et al., 2020). Despite the recovery from the crisis, persistent deviations from IRP continue to prompt research into their causes and implications for trading strategies (Chatziantoniou et al., 2020).

In examining FX swaps, this study focuses on understanding the dynamics of the difference between the spot rate and the forward rate—referred to as *forward points*—on an intraday basis. Building on the insights gained from this analysis, the paper attempts to model returns by decomposing them into two distinct components: return sign and return magnitude, as discussed in detail in Chapter 4. This thesis makes a threefold contribution to the literature:

1. Exploring seasonality in the FX swap market.
2. Modeling the conditional mean of EUR/USD three month forward points returns.
3. Decomposing returns into a sign component and a magnitude component, treating them as separate processes and accounting for their dependencies using copula theory.

The aim of this study is to provide insights into the intraday behavior of forward points, thereby contributing to a deeper understanding of FX market dynamics and potentially offering new avenues for trading strategy development.

# Chapter 2

# Literature Review

### Set-up

This chapter discusses past literature on intraday seasonality in the FX market, covering various aspects relevant to this research. Initially, the chapter provides foundational explanations of key topics, concepts, and terminology before delving into the corresponding literature. It begins with a discussion on seasonality in general and then narrows the focus to intraday seasonality, summarizing relevant findings, methods, and potential caveats. Subsequently, the chapter introduces the financial market, with a brief exploration of the FX market. The discussion extends to concepts like no-arbitrage and market efficiency, including stylized facts as defined by Challet et al. (2001). Finally, the forward points and their underlying framework, which remains crucial today, are introduced. The chapter concludes by linking these topics together, formulating a clear research question, and setting the stage for the research methods that will contribute to existing literature. A diagram is provided to visualize the research topic.

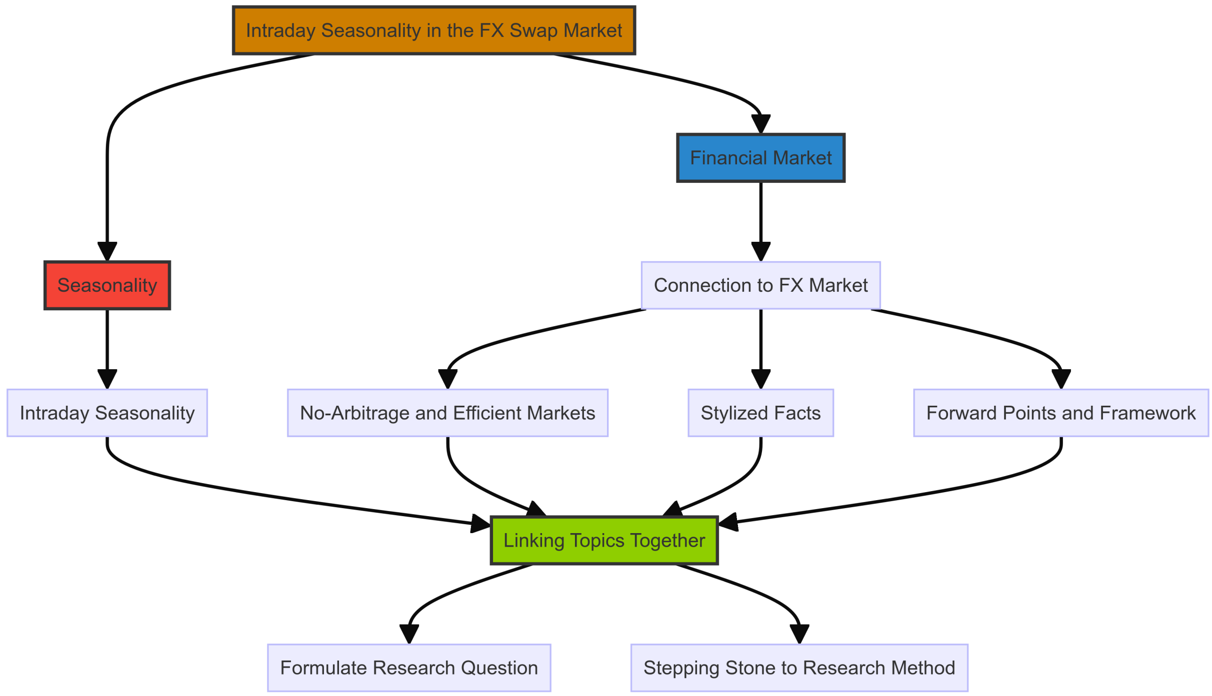


FIGURE 2.1: Literature Review Visualized

## 2.1 Seasonality

Davey and Flores (1993) define seasonality as a structured pattern of changes within a period. In other words, a time series exhibits seasonality if it shows recurring behavior after each period. Periods can range from as long as four years to as short as one hour. Granger (1978) defines yearly seasonality as a series having an observable component that repeats every 12 months. This has significant implications for interpreting and treating time series data, which should not be overlooked (Granger, 1978). The paper emphasizes that by properly investigating and accounting for seasonality in forecasting, econometric analyses and the understanding of data in general improve.

Historically, seasonality lacked a clear definition and methods for addressing it were inadequate. While early researchers like Winters (1960)[1] and Nerlove (1964) incorporated the concept of seasonality, it was not until 1978 that Clive Granger provided a mathematical definition and formalized time series models that exhibit seasonality. Although this paper will not delve into spectral analysis, its framework is solid and interpretable. Two widely-used models from Granger (1978) are presented below.

**Model 2.1.1** (Additive Model)**.** *Let $Y_t$ be a process without seasonality and let $S_t$ be a stochastic or deterministic seasonality process. Then,*

$$X_t = Y_t + S_t$$

**Model 2.1.2** (Multiplicative Model)**.** *Let $Y_t$ be a process without seasonality and let $S_t$ be a stochastic or deterministic seasonality process. Then,*

$$X_t = Y_t \cdot S_t,$$

[1]One of the founders of the Holt-Winters Exponential Smoothing framework

*where $S_t > 0$ to ensure non-negative processes.*

**Definition 2.1** (Deterministic Seasonality)**.** Let $n$ be the number of periods in a seasonal cycle. If the seasonal component $S_t$ is perfectly periodic, meaning $S_t = S_{t-n}$, then

$$S_t = \sum_{j=1}^{n} a_i \cdot d_j$$

where $d_{i,t}$ equals 1 in the $j^{\text{th}}$ period and 0 otherwise. Alternatively, if we let $\omega_s$ be the seasonal frequency, then

$$S_t = \sum_{j=1}^{n} a_i \cdot \sin(\omega_s + \theta_j)$$

From Definition 2.1, it is evident that deterministic seasonality can be applied in both models 2.1.1 and 2.1.2. The concept of deterministic seasonality, as defined in Definition 2.1, can be extended and adapted while retaining the core idea of a deterministic seasonal component. Each point in time belongs to a specific timeframe with varying characteristics, represented by the coefficient $a_i$ for time period $j$. These models suggest that, in theory, the seasonal component can be isolated. To achieve this, further decomposition of $Y_t$ is necessary. Shiskin (1967) describe one of the earliest methods, which decomposes a time series into a trend-cycle component, a seasonal component, a trading day component, and an irregular component. This method is captured in the X-11 model below.

**Model 2.1.3** (X-11 Consensus (Multiplicative))**.** *Let $C_t$ represent the trend component, $S_t$ the seasonality process, $TD_t$ the trading day component, and $I_t$ the irregular component. Then,*

$$X_t = C_t \cdot S_t \cdot TD_t \cdot I_t$$

The X-11 method also exists in an additive form. Its primary goal is to estimate the seasonal component using moving averages (Shiskin, 1967). The X-11 method revises seasonal factor estimates as new data becomes available, using asymmetric weights at the series' ends. However, this method can sometimes result in poor estimates and large revisions, which can undermine its credibility. The X-11 method serves as a foundational approach for researching economic time series. It should be noted that the method has evolved, with variants like STL by Cleveland et al. (1990) and MSTL by Bandara et al. (2021). This section aimed to explain seasonality within econometrics.

### 2.1.1 Intraday Seasonality

As discussed in the previous section, seasonality manifests at various frequencies. With a broad overview of seasonality provided, the focus now shifts to intraday seasonality. The literature on intraday seasonality within the FX spot market is extensive.[2] Before this literature emerged, research on intraday patterns in stock returns gained attention (Harris, 1986). Harris (1986) discuss two possible reasons for intraday patterns in the stock market:

[2] Note that this claim does not extend to FX swap markets.

- **Institutional Practices and Payment Schedules:** Institutional investors may receive inflows and outflows at specific times of the day. For example, Dutch pension funds typically trade during business hours, ruling out overnight trading and focusing on activity during business hours.
- **Systematic Timing of Information:** Many economic-related news announcements are scheduled at specific times and days, influencing trading behavior and potentially creating intraday patterns.

The FX market is highly sensitive to news announcements and is significantly influenced by these announcements and local trading sessions (Goodhart et al., 1993). Similar to the stock market, local trading behavior and news announcements create patterns throughout the day in the FX market. Research on intraday seasonality in FX returns typically divides the analysis into two components, both of which have become significant fields of study:

1. The level of the returns (Chang et al., 2008, Ranaldo, 2009, Zhang, 2018)
2. The (realized) volatility of the returns (Bollerslev, 1986; Baillie and Bollerslev, 1991; Payne et al., 1996; Martens et al., 2002; Ito and Hashimoto, 2006)

After extensive literature review, this paper concludes that the primary method for assessing seasonality in return levels is the Ordinary Least Squares (OLS) model. Researchers like Cornett et al. (1995), Chang et al. (2008), Ranaldo (2009), and Zhang (2018) investigate intraday seasonality by regressing returns (*regressand*) on dummies (*regressors*) and possibly additional explanatory variables. The model could take the form of:

$$y_t = \beta_0 + \sum_{j=1}^{n} \lambda_j \cdot d_{j,t} + \varepsilon_t \tag{2.1}$$

or with additional explanatory variables:

$$y_t = \beta_0 + \sum_{j=1}^{n} \lambda_j \cdot d_{j,t} + \sum_{i=1}^{k} \beta_i \cdot x_{i,t} + \varepsilon_t \tag{2.2}$$

In these equations, $d_j$ is a dummy that takes the value 1 if the observation $y_t$ occurs in a particular timeframe of the day, and $x_{i,t}$ are explanatory variables with assumed explanatory power over the regressand. The term $\lambda_j \cdot d_j$ serves as a proxy for the seasonal component $S_t$ defined in Definition 2.1.1. The coefficients $\{\lambda_j\}_{j=1}^{n}$ and $\{\beta_i\}_{i=1}^{k}$ are estimated, and their sign, magnitude, and significance are key to understanding the impact of each variable on the regressand (Heij, 2004). This implies that the seasonal coefficients ($\lambda_j$), which represent the seasonal component ($S_t$), are crucial in assessing seasonality within data. The main advantage of this model is its simplicity and interpretability. Existing research has empirically shown that these intraday patterns exist and persist. Another method of proxying the seasonal component is described by Andersen and Bollerslev (1997). Below is a summary of their framework:

$$R_{t,n} = \mathbb{E}[R_{t,n}] + \frac{\sigma_t s_{t,n} Z_{t,n}}{N^{\frac{1}{2}}} \tag{2.3}$$

- where $t$ represents the trading day,
- $n$ represents the $n^{\text{th}}$ interval on a particular day,
- $R_{t,n}$ represents the return on a particular trading day and interval,
- $\mathbb{E}[R_{t,n}]$ represents the unconditional mean of returns,
- $\sigma_t$ represents daily conditional volatility,
- $Z_{t,n}$ represents an i.i.d. random variable with mean 0 and variance 1,
- $N$ is the number of intervals per trading day,
- $s_{t,n}$ proxies the periodic component for the $n^{\text{th}}$ intraday interval.

Andersen and Bollerslev (1997) first described a two-step method involving:

1. Using a Fast Fourier Transform to approximate the seasonal component $s_{t,n}$, obtaining an estimator $\hat{s}_{t,n}$.
2. 'Filtering' the returns $R_{t,n}$ by dividing them by the estimated seasonal component $\hat{s}_{t,n}$.

Forecasting volatility remains an active area of empirical research (Mcmillan & Speight, 2012). Researchers have long recognized that the uncertainty of speculative prices, measured by (co)variances, changes over time. One of the first papers to explicitly model time-variation in second-order and higher moments was the Autoregressive Conditional Heteroskedastic (ARCH) model by Engle (1982). Since its publication, the ARCH model has received significant attention, leading to various extensions. One major extension is the Generalized Autoregressive Conditional Heteroskedastic (GARCH) model. The goal of the GARCH$(p, q)$ model is to model the conditional variance, which depends on the past.[3] GARCH models are increasingly applied today due to their simple and effective method of describing changing volatility in financial time series (Xu et al., 2011). GARCH models can also produce good forecasts of volatility changes in financial time series (Bauwens et al., 2006). Additionally, Hansen and Lunde (2005) compared 330 ARCH-type models when researching exchange rate data and found no evidence that more sophisticated models outperform the GARCH$(1, 1)$. However, GARCH models should not be applied to raw seasonal data (Yan, 2021). This is because GARCH models impose functional restrictions on the autocorrelation function (Yan, 2021). For a more detailed discussion, this paper refers to Bollerslev (1986), where it is shown that (G)ARCH models imply a geometric decay in the autocorrelation structure. While GARCH models can describe heteroskedasticity, they compromise on capturing periodicity in conditional volatility. Instead, the two-step approach shown in Equation 2.3 and outlined by Andersen and Bollerslev (1997) can be used to remove periodicity. Besides this approach, Baillie and Bollerslev (1991) propose a more practical and interpretable approach by using a seasonal GARCH model to account for seasonality patterns in the conditional variance of exchange rate data. They do this by adding dummies for specific hours of the day in the variance equation, similar to

[3] Later in this chapter, the GARCH model will be discussed in detail. This section focuses on outlining relevant literature rather than formulating models.

the idea described in Equation ?? for return levels. In conclusion, intraday seasonality in conditional variance can be addressed by either accounting for seasonality using dummies or explicitly modeling the seasonal component.

## 2.2 Financial Market

The financial market brings buyers and sellers together, ensuring liquidity and price discovery. For some time, scholars have questioned the need for markets. The discussion of market efficiency, which began generations ago, is ongoing. Malkiel (1973) famously claimed that a blindfolded chimpanzee throwing darts at the Wall Street Journal could select a portfolio that would perform as well as experts. While some nuance exists in this statement, it highlights the Efficient Market Hypothesis (EMH), which posits that prices reflect all available information and that attempting to beat the market is futile. Over time, as statistical and economic methods advanced, the idea of efficient markets was challenged, and predictability became more plausible. Fender (2020) provides a balanced view of markets today, explaining that the EMH was widely accepted due to its strong economic principles and empirical support. However, numerous anomalies have been reported, seemingly incompatible with the EMH. As Richard Feynman once said:

"*It doesn't matter how beautiful your theory is, it doesn't matter how smart you are. If it doesn't agree with experiment, it's wrong.*"

While this is generally true in econometrics and mathematics, some nuance is needed when applying it to markets. The general consensus today, as described by Fender (2020) and Fama (1991), is that markets are efficient to a certain extent but not completely, due to the existence of anomalies.

### 2.2.1 Foreign Exchange Market

The Foreign Exchange market (FX market) operates 24 hours a day, almost seven days a week. Trades are settled within seconds, thanks to electronic trading, regardless of traders' locations. This was not always the case. One of the earliest recorded traders was likely a member of the Medici family (Donnelly, 2019). At that time, currencies were valued based on their intrinsic worth, meaning their value was tied to the metal they contained. The modern system began to take shape in the early 19$^{th}$ century, with early metallic standards like the bimetallic standard giving way to the gold standard in countries like the United States, Canada, and the United Kingdom. While the gold standard promoted confidence in currencies, it was inflexible. Governments controlled gold stocks, and US citizens, for example, were not allowed to own raw gold. The shortcomings of the gold standard, highlighted by the World Wars, led to the Bretton Woods system, where countries fixed their currencies to the dollar to promote international trade and financial stability. This system also ended in 1971 due to instability.

Today, most of the world uses a system of floating rates based on supply and demand. With an average daily trading volume exceeding $6 trillion, the FX market is the deepest and largest financial market (Krohn & Sushko, 2022). This is primarily due to over-the-counter (OTC) transactions. Another characteristic of the FX market is that derivatives are traded more frequently than spot products. Among these derivatives, the most commonly traded is the FX Swap. The popularity of FX Swaps is due to their ability to hedge currency risk and provide short-term liquidity (ING (2019); Krohn and Sushko, 2022; BIS, 2022). Figure 2.2[4] shows (from left to right) the absolute net daily averages from 2001 to 2019, the segmentation of the FX Market in percentages in 2016, and the segmentation of the FX Market in percentages in 2019. This figure emphasizes the crucial role of FX Swaps in today's financial landscape.

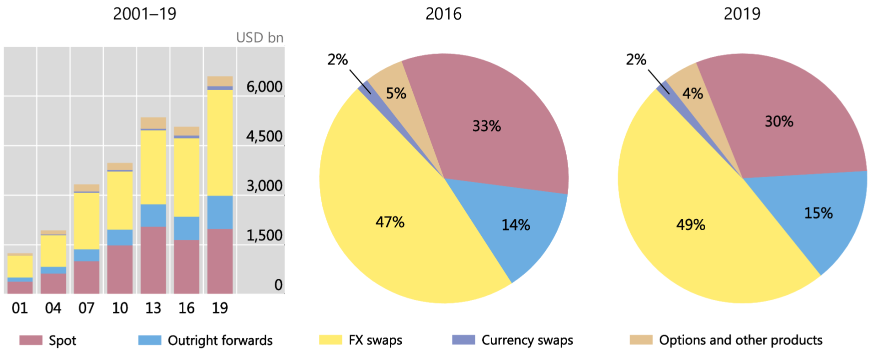


FIGURE 2.2: FX Market Segments in 2019. *Source: (BIS, 2019)*

### Forward Points

The CIP serves as a benchmark for assessing perfect capital mobility between markets (Levich, 2011). In practice, the CIP might not always hold, and deviations observed in the market are referred to as *basis*. Traditionally, any departure from the CIP indicated stress within the global financial system, notably during the 2008 financial crisis (Chatziantoniou et al., 2020). Despite the recovery from the crisis, persistent deviations from CIP continue, prompting research into their causes and implications for trading strategies (Chatziantoniou et al., 2020). Another aspect of the CIP that has puzzled researchers, both before and after the financial crisis, is the difference between the forward rate and the spot rate, known as the forward premium or forward points.

**Definition 2.2** (Forward Points)**.**

$$F_{t,t+\delta} - S_t = S_t \cdot \frac{(1 + r_{\delta,\$})}{(1 + r_{\delta,€})} - S_t = S_t \cdot \left( \frac{(1 + r_{\delta,\$})}{(1 + r_{\delta,€})} - 1 \right)$$

*The $\delta$-forward points represent the difference between the spot rate and the $\delta$-forward rate, where $\delta$ stands for an arbitrary time period.*

The forward points have not been extensively researched as a standalone subject. A common field of research where forward points have appeared is the forward premium anomaly. This anomaly arises because economic models typically imply that domestic currency depreciates when domestic interest

[4]This graph is publicly available and published by BIS (2019)

rates are higher than foreign interest rates (Bansal & Dahlquist, 2000). When this is not the case, as empirical findings suggest, the literature refers to it as a puzzle—the 'forward premium puzzle' introduced by Fama (1984). Baillie and Bollerslev (2000) argue that the anomaly is 'not as bad as you think' by presenting a model suggesting that it may be merely a statistical artifact. Instead of solving the puzzle, these findings have only fueled its complexity. Another area of interest is the relationship between the spot rate and the forward rate, particularly the hypothesis that the forward rate is an unbiased predictor of the future spot rate. Engel (1996) thoroughly discusses this in his survey, considering past literature and essentially concluding that the hypothesis can be rejected.

Surprisingly, as far as the author is aware, no prior research has focused on the behavior of forward points over time, either on a daily or long-term basis. Despite their potential importance, the behavior of forward points has not been explored. A study by Joseph and Hewins (1992) comes close, examining interday seasonality in the FX market and evaluating whether day-of-the-week returns for spot rates and forward rates differ. While this has already been done for spot rates by Levi (1978), McFarland et al. (1982), and others, Joseph and Hewins (1992) were the first to explore seasonality in the FX forward market. Their study identified seasonality in the forward market, but this finding may be flawed due to the definition of returns used. For spot rates, returns are straightforward and can be taken as $\log\left(\frac{\{}{P}_{t+1} P_t\right)$ (Zhang, 2018). For forward rates, this is less straightforward due to how they are defined, as shown in Definition 2.2. Since forward rates comprise two interest rate components and a spot component, studies relying solely on raw forward rates may produce misleading results. Therefore, the observed seasonality in the forward market may be entirely driven by seasonality in the spot market. An improvement to the methodology used in Joseph and Hewins (1992) would involve isolating the influence of the interest rate components by removing the spot component from the forward rate.

### Stylized Facts in the FX Market

Stylized facts are empirical findings that are so consistent that they are believed to hold approximately. In econometrics, stylized facts summarize empirical observations about the distribution. This section briefly discusses stylized facts in the intraday FX market. While stylized facts are readily available in the financial market, they are less common in the FX market. The stylized facts listed here have been documented by Guillaume et al. (1997) and are mainly focused on spot products.

A well-known fact about financial returns, in general, is that they exhibit fat tails and usually reject normality. These traits point to a *leptokurtic* distribution, which holds in the FX market as well. Some researchers claim the distribution resembles a stable Paretian, while others believe it is closer to the Student-t distribution (Guillaume et al., 1997). The distributions are usually symmetric, with finite variance and an existing third moment that supports the claim of no skewness.

Guillaume et al. (1997), and many others he cites, claim that seasonal patterns in the FX market are apparent. They correspond to the hour of the day, the day of the week, or the presence of traders in the three major geographical trading zones. Although he does not explicitly mention these trading zones, it is assumed that he refers to the Asian, European, and North American trading sessions. Seasonality is found to exist in FX literature for, among other things, the variance and the directional change. In line with the conclusion in Section 2.1.1, Guillaume et al. (1997) also observes that the first and most straightforward way to treat seasonality is by using seasonal dummies. He also mentions the framework introduced in Equation 2.3, where seasonality is modeled explicitly, as the second profound method.

The goal of this literature review was to gain an understanding of the existing methods for treating seasonality, to understand the FX market and the framework of FX swaps, and to present past literature relevant to this topic. To conclude this chapter, this paper attempts to answer whether seasonality exists in FX forward points and whether this can be exploited. To maximize the efficiency of our methods, we draw on past research, leveraging the stylized facts that significant information about seasonality can be extracted from realized volatility and that directional changes hold valuable insights. Before presenting the proposed methodology in Chapter 4, an in-depth analysis and discussion of the data will be provided in Chapter 3.

# Chapter 3

# Data

### Set-up

This chapter presents relevant information and findings on the data. It begins with an introduction to the data, followed by visualizations and statistical reporting. Adjustments to the raw data are explained and justified. Stylized facts present in the data are also mentioned. Additionally, the chapter reveals insights through statistical analysis and supports statements with statistical tests.

## 3.1 Data Preparation and Adjustments

This paper examines the three-month EUR/USD forward points, representing the difference between the spot rate at time $t$ and the three-month forward rate at time $t$. PGGM provided the data via their dat provider New Change FX (NCFX). NCFX is an FCA-authorised benchmark administrator providing high-quality, real-time FX reference rates. The data is provided at a 1-minute frequency. Starting from January 1$^{st}$ 2021 to March 21$^{th}$ 2024. After adjustments, 112.560 data points remain, which translates to 670 trading days.

### 3.1.1 Exclusion of Overnight Data

The FX market operates 24 hours a day, 5 days a week, but due to liquidity issues, non-business hours are often uninformative or misleading, with increased trade risk and wider bid-ask spreads. According to Bjønnes and Rime (2005) and Bjønnes et al. (2005), market-making banks primarily provide intraday liquidity and typically offload inventories before the end of the trading day due to associated risks. This behavior is a reaction to the lack of demand overnight, which leads to reduced liquidity and activity during these hours. Market makers remain at their desks if there is sufficient demand, but the low demand overnight results in decreased liquidity. Therefore, overnight data is excluded from the dataset to avoid distortions caused by these liquidity constraints. Observations between 8PM GMT and 5:59AM GMT are excluded from the analysis.

### 3.1.2 Adjustment for Daylight Saving Time

Additionally, to account for the effects of Daylight Saving Time (DST), all timestamps were converted to UTC, and trading hours were adjusted to reflect local time changes. DST transitions can disrupt regular trading patterns due to changes in traders' schedules and behavior. By accounting for DST, we can standardize the timestamps and maintain consistent trading hours throughout the year, avoiding distortions.

### 3.1.3 Tick Frequency

The data provided is at 1-minute level. Following Andersen and Bollerslev (1997, 1998) and Ranaldo (2009), the forward points data is resampled to 5-minute level by selecting the close price of every 5-minute interval. This resampling reduces noise in the data while preserving essential information.

### 3.1.4 Holidays

Trading patterns on holidays differ from regular trading days, often exhibiting lower liquidity and higher volatility. According to Ranaldo (2009), market participation is skewed on holidays, leading to less efficient price discovery and wider bid-ask spreads. Additionally, holidays are characterized by significantly reduced market activity, similar to weekends, resulting in a lack of reliable return observations and subdued quoting activity (Andersen & Bollerslev, 1998). By excluding these days[1], we avoid distortions caused by atypical trading behavior (Andersen & Bollerslev, 1998; Ranaldo, 2009).

### 3.1.5 Spot Component

As outlined in Definition 2.2, the forward points include a spot component. To isolate the forward points, divide them by the spot rate, which fluctuates throughout the day. This adjustment ensures that spot rate variations do not obscure our analysis of the forward points.

### 3.1.6 Outliers

After removing overnight observations, no clear outliers have been detected.

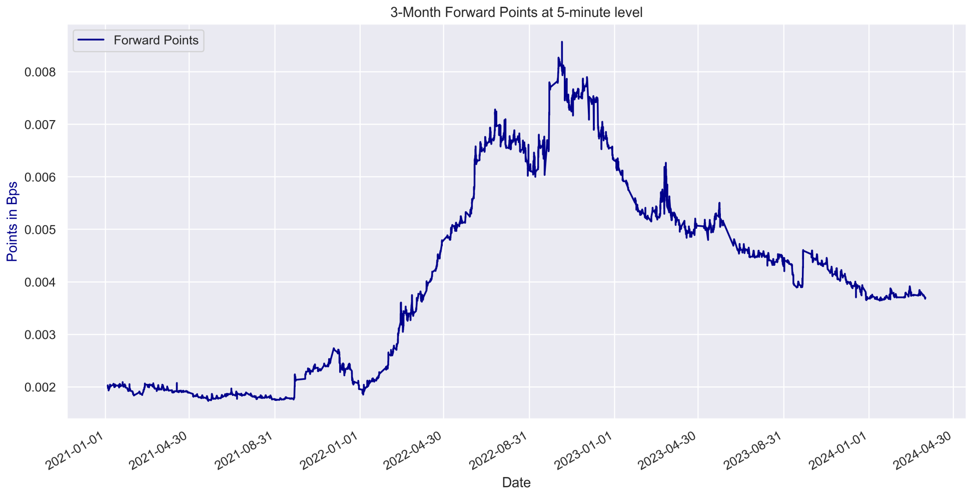


Figure 3.1: EUR/USD 3-Month Forward Points: Level

From January 2021 to mid-2022, forward points were stable, fluctuating between 10 and 20 basis points (Bps), indicating calm market conditions and predictable monetary policy.

In May 2022, forward points surged, peaking at around 80 Bps by September 2022. This spike likely reflects unexpected monetary policy changes by the Federal Reserve (FED) and the European Central Bank (ECB). The USD 3-month swap Overnight Index Swap (OIS) and EUR 3-month OIS[2], which proxy dollar and euro interest rates respectively, are shown in Figure A.5. The increasing trend in forward points in early 2022 coincides with the interest rate hikes.

From January 2023, forward points declined, stabilizing around 40 Bps by July 2023, suggesting a return to market stability, as interest rate hikes diminished. Volatility also declined, returning to pre-hike levels by January 2024.

[1] Appendix ?? shows the holidays that were excluded

[2] The EUR and USD 3-month OIS are overnight interest swap rates for the dollar and euro respectively

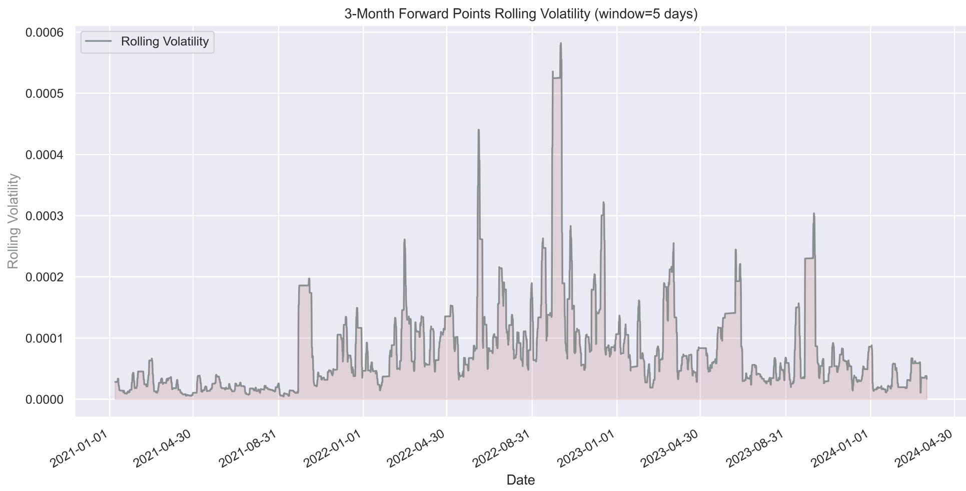


FIGURE 3.2: EUR/USD 3-Month Forward Points: Rolling Volatility (5-day window)

### 3.1.7 Stationarity

From Figure 3.1 and Figure 3.2, it is clear that the forward points do not show a constant mean and variance. These varying statistical properties over time indicate a non-stationary process. A stochastic process is considered (weakly) stationary if its first and second moments remain constant over time[3]. In order to obtain robust estimates and ensure constant statistical properties, stationarity of a time series is a desired property (Heij, 2004). To test for stationarity within a time series, we apply the Augmented Dickey–Fuller (ADF) test that check for a stochastic trend. The null-hypothesis claims that there is a unit root present in the time series. Upon rejection of the null-hypothesis, a frequently used transformation applied to the time series is that of first differencing (Heij, 2004). The test is documented in Appendix B. Let $y_t$ be the forward points time series, then the series of the first differences are defined in equation 3.1

$$\Delta y_t = y_t - y_{t-1}, \text{ for } t \in \{1, \dots n\} \tag{3.1}$$

For both series, stationarity is tested using the ADF test. The results are given in Table 3.1. We find that after taking first differences, the series becomes stationary. From now on, the difference series defined in equation 3.1, will be referred to as *returns* series. The returns series will be analyzed. Practitioners also take logarithmic returns to obtain stationary series, we however take first differences for two reasons. The series are expressed in Bps (hence: very small) and applying a logarithmic transformation inflates the series. Thereby, observing that the process is fairly linear, differencing is a reasonable measure to take here.

[3]Appendix ?? stationarity

TABLE 3.1: ADF Test Results

| Series | ADF Statistic | p-value | $\alpha$ |
|---|---|---|---|
| Forward Points | -1.186 | 0.679 | 5% |
| $\Delta$ Forward Points | -51.331 | 0.000 | 5% |

## 3.2 Returns

Figures 3.3 and 3.4 display the 5-minute returns and absolute returns, respectively. Table 3.2 provides the descriptive statistics for both series.

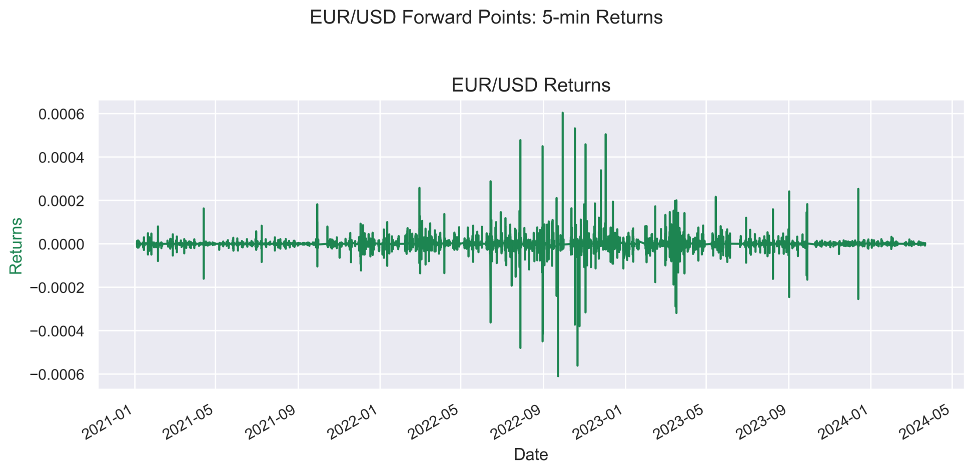


FIGURE 3.3: EUR/USD Forward Points: 5-min Returns

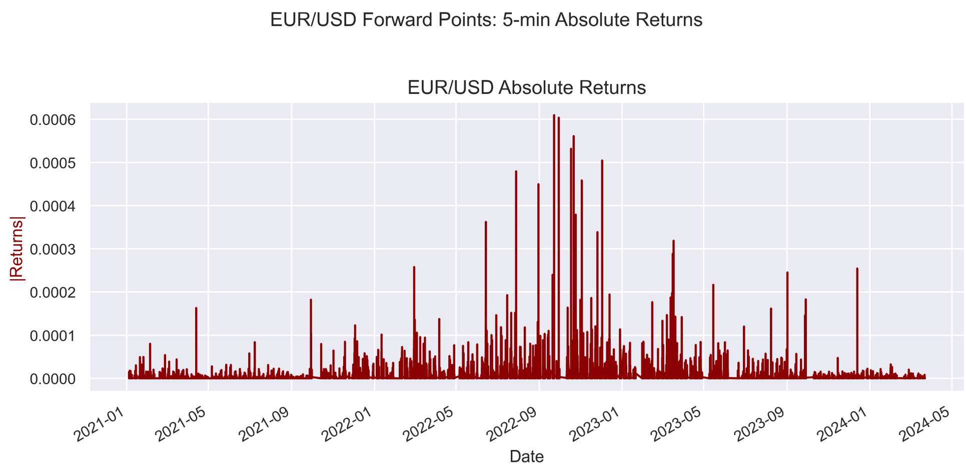


FIGURE 3.4: EUR/USD Forward Points: 5-min Absolute Returns

The returns plot indicates a stationary series with periods of increased volatility, which is a known stylized fact in financial markets. Volatility can be proxied by the absolute value of returns.[4] Figures 3.3 and 3.4 show the autocorrelation function (ACF) and partial autocorrelation function (PACF) for returns and absolute returns, respectively.

The descriptive statistics of returns reveal several stylized facts. A high kurtosis is observed, indicating leptokurtocity. The return distribution is also symmetric and centered around zero.

TABLE 3.2: Descriptive Statistics for Returns and Absolute Returns (in Bps)

| **Statistic** | **Returns (Bps)** | **Absolute Returns (Bps)** |
|---|---|---|
| Count | 112,560 | 112,560 |
| Mean | 0.0003 | 0.0214 |
| Std Dev | 0.0918 | 0.0893 |
| Min | -6.0965 | 0.0000 |
| 25% | -0.0060 | 0.0002 |
| 50% | 0.0000 | 0.0060 |
| 75% | 0.0060 | 0.0190 |
| Max | 6.0377 | 6.0965 |
| Kurtosis* | 1244 | 1363 |
| Skewness* | 0.8070 | 29.2380 |

* Skewness and Kurtosis are not converted to Bps.

The ACF for absolute returns shows significant autocorrelation at initial lags, especially at lag one, indicating volatility clustering. Periodic spikes around every 168 lags suggest possible daily seasonality. The PACF plot supports this with diminishing influence over time.

The ACF for returns exhibits minor spikes, notably at lag one, suggesting low overall autocorrelation and indicating the series is largely white noise. The PACF plot shows minimal initial spikes and weaker autocorrelation. As expected, absolute returns demonstrate more significant autocorrelation and partial autocorrelation compared to returns, emphasizing volatility clustering and potential daily seasonality. We leave additional information on the distributional properties for the appendix (A).

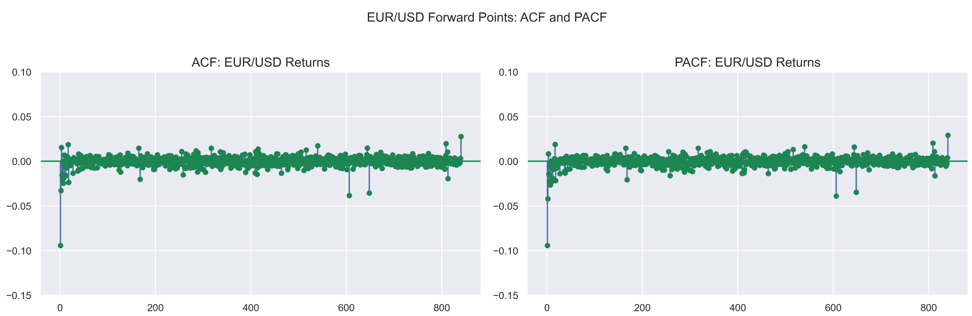


FIGURE 3.5: EUR/USD Forward Points: ACF and PACF Returns

[4]See the literature review for a broader discussion on this.

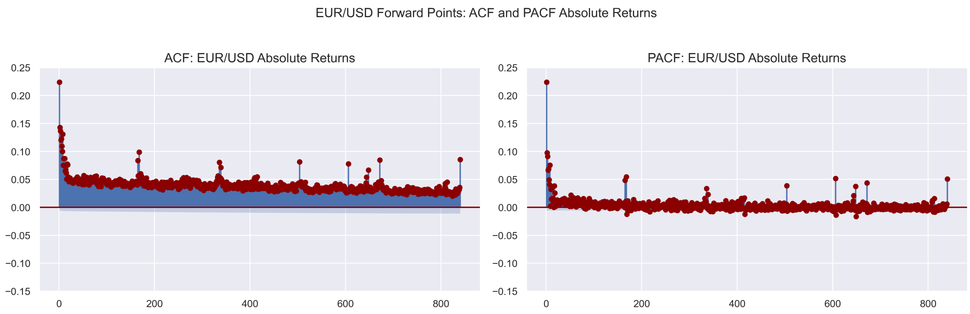


FIGURE 3.6: EUR/USD Forward Points: ACF and PACF Absolute Returns

### 3.2.1 Intraday Returns

Figure 3.7 shows the mean return per hour and volatility for that hour. The intraday plots reveal distinct patterns in returns. Positive mean returns are observed in the morning, peaking around 11 AM, while negative mean returns dominate the early afternoon, especially at 2 PM, indicating a possible daily seasonality. The highest levels of absolute returns, indicating peak volatility, occur at 11 AM and 2 PM, coinciding with major trading periods in the European and US markets.

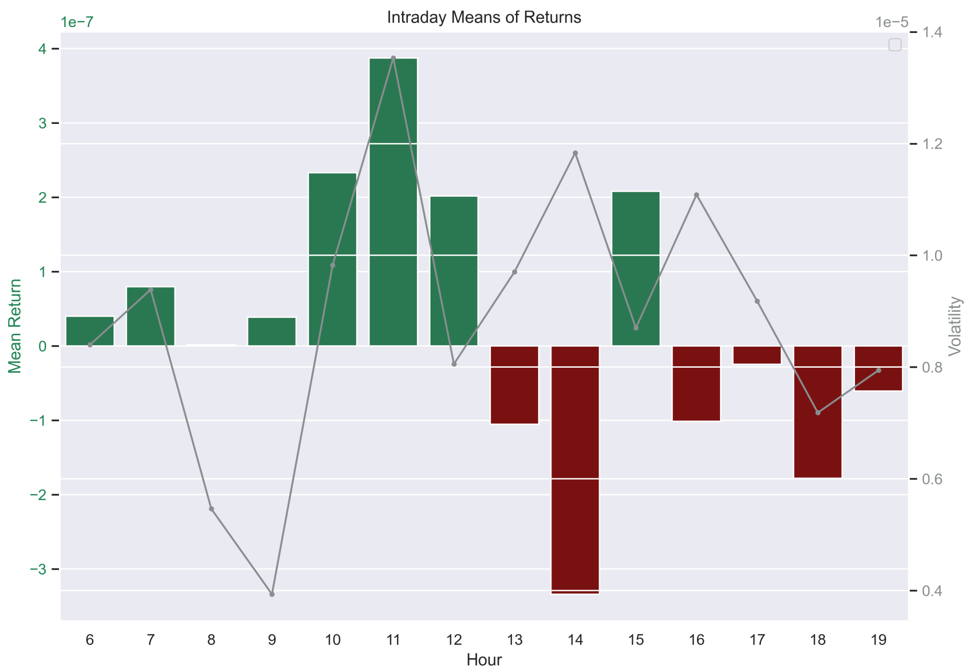


FIGURE 3.7: Intraday Mean Return per Hour

Following Iwatsubo et al. (2018) and Ranaldo (2009) and based on figure 3.7, we can divide the day into sessions. Table 3.3 shows how a trading day is divided.

TABLE 3.3: Time Bins

| **Bin** | **Time (London)** | **Session Description** |
|---|---|---|
| H1 | 6:00 AM - 7:59 AM | Pre London Market |
| H2 | 8:00 AM - 9:59 AM | London Open |
| H3 | 10:00 AM - 11:59 AM | Mid-Morning |
| H4 | 12:00 PM - 1:59 PM | Lunch |
| H5 | 2:00 PM - 3:59 PM | NY Open |
| H6 | 4:00 PM - 5:59 PM | London Close |
| H7 | 6:00 PM - 7:55 PM | Post Market |

Intraday patterns remain quite consistent across years with mean return around zero, suggesting a stable and balanced data-generating process. Volatility fluctuates across years. Market volatility is highest during lunch and the New York open, and the market is very tight during the mid-morning session as shown in A.3.

To statistically test for equality in mean returns, we perform Welch's t-test[5] as outlined in A.1.3. Welch's t-test is chosen because the trading sessions exhibit unequal variances, as shown by Levene's test, which is described in A.1.2. This choice ensures a more accurate assessment of mean differences between sessions. We find that variances are different across all trading sessions and the mean for the mid-morning (post market) session is significantly higher (lower).

### Trading Session Effect

To analyze the impact of different time bins on returns, we conducted an OLS regression with dummy variables for each time bin. The regression equation is as follows:

$$r_t = \beta_0 + \beta_1 \cdot H2_t + \beta_2 \cdot H3_t + \beta_3 \cdot H4_t + \beta_4 \cdot H5_t + \beta_5 \cdot H6_t + \beta_6 \cdot H7_t + \varepsilon_t, \quad (3.2)$$

where $Hi_t = 1$ if $r_t$ is observed in trading bin $Hi_t$ and $Hi_t = 0$ otherwise. The results of the regression are summarized in Table 3.4.

TABLE 3.4: OLS Results Return on Trading Session Dummies

| **Variable** | **Coefficient** | **Standard Error** | **p-value** |
|---|---|---|---|
| const | 5.985e-08 | 7.02e-08 | 0.394 |
| H2 | -3.954e-08 | 7.96e-08 | 0.620 |
| H3 | 2.506e-07 | 1.17e-07 | 0.032 |
| H4 | -1.159e-08 | 9.94e-08 | 0.907 |
| H5 | -1.228e-07 | 1.08e-07 | 0.255 |
| H6 | -1.228e-07 | 1.07e-07 | 0.249 |
| H7 | -1.79e-07 | 9.22e-08 | 0.052 |

[5] Also known as the unequal variances t-test

The regression results indicate that the coefficient for the H3 (mid-morning) time bin is positive and statistically significant at the 5% level. Compared to the baseline (H1) returns are higher in the mid-morning session, but this difference is very small.

Other time bins do not show statistically significant coefficients at the 5% level. This is contrary to the finding for H7 in the Welch's t-test. Therefore, we conclude that only the mid-morning session produces significant positive returns. Heteroskedastic consistent error estimators proposed by White (1980) are used in this models and other OLS models in this paper.

#### Monday Morning Effect

Previous literature suggests that FX spot returns might be different during Monday mornings due to various market dynamics (Aloud et al., 2013; Singh, 2019). To investigate the Monday morning effect for FX swap returns, we conducted an OLS regression. The results of model 3.3 are summarized in table A.3. The regression equation is as follows:

$$r_t = \beta_0 + \beta_1 \cdot MM_t + \varepsilon_t, \tag{3.3}$$

where $MM_t = 1$ if $r_t$ is observed on a Monday morning and $MM_t = 0$ otherwise.
No significant effect is found for different returns during Monday mornings.

#### London and NY Opens

A known stylized fact, as shown by Aloud et al. (2013), suggests that returns in the FX spot market behaves different around the London open and NY open times. To examine these in the FX swap market, we performed an OLS regression where we regress returns on dummies for the London and NY opening sessions. The results of model 3.4 are summarized in table A.4. The regression equation is as follows:

$$r_t = \beta_0 + \beta_1 \cdot LO_t + \beta_2 \cdot \text{NYOpen}_t + \varepsilon_t, \tag{3.4}$$

where $LO_t = 1$ and $NY = 1$ if $r_t$ is observed during the London open and the New York open respectively.
No significant effect is found for different returns during the London and NY opening sessions.

### 3.2.2 Absolute Returns

There is only a limited amount of information to be extracted from returns. The autocorrelation structure contains no information about possible intraday seasonality. Welch's t-test and regression model 3.2 reveal some signficant activity during the mid-morning session. The absolute value of returns however show intraday seasonality in the autocorrelation structure, as shown in figure 3.8. We observe a weak U-shape that has been well documented for different spot exchange rates (Andersen & Bollerslev, 1997, 1998; Ito & Hashimoto, 2006).

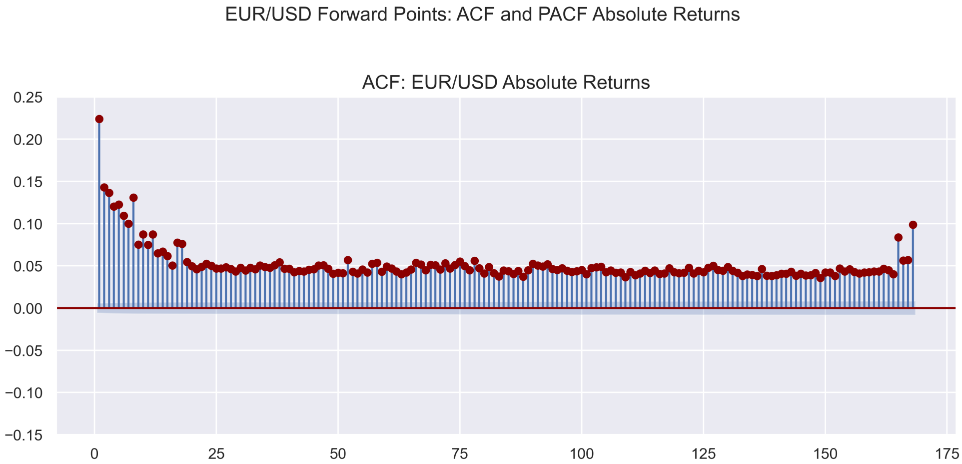


FIGURE 3.8: EUR/USD 3-Month Forward Points: Intraday ACF

We test the effects of different trading sessions on absolute returns throughout the day using the regression equation given in equation 3.5. The results, presented in Table 3.5, show that most trading sessions have a statistically significant impact on absolute returns, though the effects are small. The only exception is the London close session (H6).

$$|r_t| = \beta_0 + \beta_1 \cdot H2_t + \beta_2 \cdot H3_t + \beta_3 \cdot H4_t + \beta_4 \cdot H5_t + \beta_5 \cdot H6_t + \beta_6 \cdot H7_t + \varepsilon_t \quad (3.5)$$

TABLE 3.5: OLS Results Absolute Return on Trading Session Dummies

| **Variable** | **Coefficient** | **Standard Error** | **p-value** ($\alpha = 5\%$) |
|---|---|---|---|
| const | 1.979e-06 | 6.85e-08 | 0.000 |
| H2 | -4.328e-07 | 7.71e-08 | 0.000 |
| H3 | 5.099e-07 | 1.14e-07 | 0.000 |
| H4 | 3.267e-07 | 9.65e-08 | 0.001 |
| H5 | 4.132e-07 | 1.05e-07 | 0.000 |
| H6 | -1.048e-07 | 1.04e-07 | 0.316 |
| H7 | 4.381e-07 | 8.88e-08 | 0.000 |

# Chapter 4

# Methodology

### Set-up

This chapter introduces a framework designed to forecast returns by leveraging findings from Chapter 3. Our analysis of intraday raw returns revealed no significant seasonal patterns. However, we found notable differences in returns variances across trading sessions and identified seasonality in absolute returns through the autocorrelation structure depicted in Figure 3.6. The OLS model results in Equation 3.5 confirmed that absolute returns vary significantly throughout the day.

Given that financial returns are known to exhibit non-linearity, non-stationarity and conditional heteroskedasticity[1], our information set may not capture these non-linear dynamics adequately. By combining our empirical findings with methods from past research, we aim to incorporate hidden non-linearities and optimally utilize the statistical information and dependencies found in absolute returns. The remainder of this chapter focuses on the following approaches to construct forecasting models:

- Modeling the returns series using classical time series approaches to obtain a linear benchmark model.
- Decomposing returns into sign and magnitude components and modeling these separately, assuming independence between sign and magnitude.
- Incorporating possible dependencies in the decomposition model by using copulas.

[1] Bollerslev (1986), Geweke (1986), and Korkie et al. (2002) discuss this in detail

## 4.1 SARIMA-GARCH

As benchmark model, an SARIMA-GARCH model is estimated. The ARMA model is a model for conditional mean whereas GARCH is a conditional variance model. The AR(I)MA model was first introduced by (Box et al., 1970) and involves an iterative fitting procedure, collectively referred to as the *Box-Jenkins approach.*

The Box-Jenkins approach involves identifying, estimating, and diagnosing models to best represent a time series. This is done by iteratively applying differencing, autoregressive, and moving average components.

### SARIMA Model

An autoregressive (AR) model incorporates past returns behavior to model future returns, while the moving average (MA) model regresses current returns on previous errors. Upon combining the AR and MA models, the ARMA model is obtained. The ARMA$(p,q)$ model for returns is specified as follows:

$$r_t = \mu + \sum_{i=1}^{p} \phi_i \cdot r_{t-i} + \sum_{j=1}^{q} \theta_j \cdot \varepsilon_{t-i} + \varepsilon_t, \tag{4.1}$$

where $r_t$ is the observed return at time $t$, $\mu$ is the constant term, $\phi_i$ and $\theta_i$ are the $i$th AR and MA coefficients respectively. Last, $\varepsilon_t$ is the error term at time $t$, where the assumption is made that $\varepsilon_t \sim WN(0, \sigma^2)$. This can be rewritten as follows:

$$\Phi_p(L) r_t = \Theta_q(L) \varepsilon_t, \tag{4.2}$$

where $L$ is the lag-operator $\Phi_t$ and $\Theta_t$ are the AR and MA polynomials respectively given below.

$$\Phi_p(L) = 1 - \phi_1 L - \cdots - \phi_p L^p,$$
$$\Theta_q(L) = 1 + \theta_1 L + \cdots + \theta_q L^q,$$

An ARMA$(p,q)$ model becomes an ARIMA$(p,d,q)$[2] model upon applying the lag operator $d$ times to the time series (Box et al., 1970). Differencing is essential in case one works with non-stationary time series. In our case, saying the returns can be represented by an ARMA$(p,q)$ model is equivalent to saying the forward points are represented by an ARIMA$(p,1,q)$ model.

Seasonal differencing eliminates seasonality in our returns. Our seasonal period is m=168 intervals of five minutes. Similar to the regular AR and MA polynomials used in equation 4.2, the seasonal AR and MA polynomials are defined below.

$$\tilde{\Phi}_p(L^s) = 1 - \tilde{\phi}_1 L^s - \cdots - \tilde{\phi}_p L^s p,$$
$$\tilde{\Theta}_q(L^s) = 1 + \tilde{\theta}_1 L^s + \cdots + \tilde{\theta}_q L^s q,$$

, where $\tilde{\theta}_q(L)$ and $\tilde{\phi}_p(L)$ denote the seasonal AR and MA coefficients respectively.

Combining equations 4.2 with the seasonal polynomials results in the SARIMA$(p,0,q)(P,D_s,Q)$, which combines the ARIMA model with seasonal

[2]ARIMA stands for Autoregressive Integrated Moving Average

components. As we construct a model for returns that are stationary already, the lag operator is zero ($d = 0$) such that we can reduce the ARIMA part to an ARMA model. The SARMA$(p, q)(P, D_s, Q)$ is specified as follows:

$$\Phi_p(L)\tilde{\Phi}_p(L^m)r_t = \Theta_q(L)\tilde{\Theta}_q(L^m)\varepsilon_t, \tag{4.3}$$

Now that we have a rough structure of the model, we continue below with the Box-Jenkins approach to find a suitable benchmark model.

### 4.1.1 Model Identification

To identify the order of a SARMA model, the ACF and PACF are used to identify AR and MA terms respectively. The ACF and PACF for returns are given in As described in Box et al. (1970) we look for a cutoff in the decay of the ACF plot. We find two lags to be both significant at the 5% level and exhibit a decaying pattern. Hence, the model includes two AR terms.

For the AR terms, we look for significant lags in the PACF. Although the fifth lag is just significant at the 5% level, it is very close to being insignificant. Therefore we define the cutoff to be the fourth lag. Hence, the model includes four MA terms.

For the seasonal AR and MA polynomials, only the primary seasonal AR and MA terms will be included. The identified SARMA model is given below.

$$\Phi_2(L)\tilde{\Phi}_1(L^{168})r_t = \Theta_5(L)\tilde{\Theta}_1(L^{168})\varepsilon_t, \tag{4.4}$$

which we refer to as the mean equation. The mean equation written out fully is given as:

$$\begin{aligned} r_t = {} & \phi_1 r_{t-1} + \phi_2 r_{t-2} + \tilde{\phi}_1 r_{t-168} - \phi_1\tilde{\phi}_1 r_{t-169} - \phi_2\tilde{\phi}_1 r_{t-170} \\ & + \varepsilon_t + \theta_1\varepsilon_{t-1} + \theta_2\varepsilon_{t-2} + \theta_3\varepsilon_{t-3} + \theta_4\varepsilon_{t-4} \\ & + \tilde{\theta}_1\varepsilon_{t-168} + \theta_1\tilde{\theta}_1\varepsilon_{t-169} + \theta_2\tilde{\theta}_1\varepsilon_{t-170} + \theta_3\tilde{\theta}_1\varepsilon_{t-171} + \theta_4\tilde{\theta}_1\varepsilon_{t-172}. \end{aligned} \tag{4.5}$$

### 4.1.2 GARCH

As shown in chapter 3, returns are heteroskedastic. Therefore the general SARMA model, assuming constant variance, most likely fails at producing residuals that resemble white noise. We construct a model for conditional variance as well. Conditional variance is defined as $\sigma_t^2 = \mathbb{V}_{t-1}(\varepsilon_t) = \mathbb{E}_{t-1}(\varepsilon_t^2)$ and the GARCH$(p, q)$ model is defined below.

$$\sigma_t^2 = \omega + \sum_{i=1}^{p} \alpha_j \cdot \sigma_{t-i}^2 + \sum_{j=1}^{q} \beta_j \cdot \varepsilon_{t-j}^2 \tag{4.6}$$

Note that $\varepsilon_t$ is inferred from the the SARMA model in equation 4.4 and specified as follows:

$$\varepsilon_t = \sigma_t z_t,\ z_t \sim t_\nu(0, 1) \tag{4.7}$$

This paper will implement the GARCH$(1, 1)$-model of Bollerslev (1986) for conditional variance. Errors are assumed to follow Student's t distribution as it is more appropriate for leptokurtic processes. This is the most widely used specification of GARCH models due to its great performance. Hansen and Lunde (2005) compared volatility models to find out if anything could beat a GARCH(1, 1). They found that in the analysis of exchange rate data, no model could outperform the GARCH(1, 1)-model.

Below, we specify the complete SARIMA(2,0,4)(1,0,1)-GARCH(1,1) model.

$$\begin{aligned} \Phi_2(L)\tilde{\Phi}_1(L^{168})r_t = \Theta_5(L)\tilde{\Theta}_1(L^{168})\varepsilon_t, \\ \sigma_t^2 = \omega + \sigma_{t-1}^2 + \varepsilon_{t-1}^2 \\ \varepsilon_t = \sigma_t z_t, \ z_t \sim t_\nu(0,1) \end{aligned} \tag{4.8}$$

### 4.1.3 Parameter Estimation

The parameter set of the mean equation (4.4) $\varphi = \{\mu, \phi_1, \phi_2, \theta_1, \ldots, \theta_4, \tilde{\phi}_1, \tilde{\theta}_1\}$. are estimated using MLE. We impose Student's t-distribution on the innovations; $\varepsilon_t \sim t_\nu(0, \sigma_t^2)$. The seasonal components are simply additional AR and MA terms such that the estimation procedure remains the same. To verify whether the GARCH(1, 1) is validated, we test for ARCH effects (Engle, 1982). This test is included in Appendix B. It essentially tests the squared residuals of the model to exhibit autocorrelation which indicates volatility clustering (Francq & Zakoian, 2019). If the presence of ARCH effects is not rejected, we employ a GARCH(1,1) model for the conditional variance. We can estimate its parameters using MLE.

### 4.1.4 Model Diagnostics

#### Residuals Inspection

After estimating, the first step is to examine the residuals of the model. Ideally, the residuals should resemble white noise, which means they should have a mean of zero, constant variance, and no autocorrelation. The following diagnostic checks are performed:

- **ACF of Residuals:** We plot the ACF of the residuals to check for any remaining autocorrelation. If the model is correctly specified, the residuals should show no significant autocorrelation at any lag. We look for a pattern in the ACF plot; a well-specified model will have residuals whose autocorrelations lie within the confidence bounds, indicating that no significant autocorrelation remains.
- **PACF of Residuals:** The PACF of the residuals is also examined to ensure that any significant autocorrelation in the original series has been adequately captured by the model. As with the ACF, the PACF should not show any significant spikes if the model is well-fitted.
- **Ljung-Box Test:** To statistically assess the absence of autocorrelation, we perform the Ljung-Box test (Ljung & Box, 1978) on the residuals. The null hypothesis of this test is that the residuals are independently distributed. A high p-value from this test would indicate that we cannot reject the null hypothesis, suggesting that the model's residuals are indeed uncorrelated.

#### Heteroskedasticity Test

Since the SARIMA-GARCH model explicitly models the conditional variance, the ARCH-LM (Lagrange Multiplier) test is used to check for remaining

ARCH effects in the residuals. The null hypothesis is that there are no ARCH effects, meaning that the conditional variance model has successfully removed volatility clustering from the residuals. A non-significant test result indicates that the GARCH(1,1) model has sufficiently captured the conditional variance structure.

#### Parameter Significance

Each parameter in the SARIMA-GARCH model is tested for statistical significance by calculating its corresponding p-value. The p-values are derived from the t-tests, which are based on the estimated parameters and their standard errors. A parameter is considered statistically significant if its p-value is less $\alpha = 0.05$.

## 4.2 Decomposition Model

To start off, following Christoffersen and Diebold (2006), the returns are rewritten as the product of two separate components: a continuous component and a discrete component. This decomposition is represented as:

$$y_t - y_{t-1} \equiv r_t = \text{sign}(r_t) \times |r_t| \tag{4.9}$$

where $y_t$ is the observation for the forward point at time $t$, $\text{sign}(r_t)$ represents the direction of the return (positive or negative) and $|r_t|$ represents the magnitude of the return (absolute value) at time $t$. Note that these two components are not necessarily independent. The main advantage of decomposing returns as shown in Equation 4.9 is that it allows us to inspect the components separately and later jointly to incorporate possible dependencies.

### 4.2.1 Empirical Motivation and Findings

#### Intraday Patterns in Returns

In our empirical analysis of intraday returns, we observed significant intraday seasonality in the absolute value of returns, even though the raw returns themselves exhibited minimal predictable patterns. This is consistent with the literature (Guillaume et al., 1997). The ACF plots of absolute returns, as well as the OLS regression analysis incorporating time-of-day dummies, revealed these patterns.

#### Predictability of Returns Components

Research has shown that both the sign and magnitude of returns can be predictable. For instance, Breen et al. (1989), Hong and Chung (2003), and Nyberg (2011) demonstrated that the directional change (sign) of returns is, to some extent, predictable. Additionally, volatility, which is closely related to the magnitude of returns, is also highly predictable as shown by Cao and Tsay (1992), Figlewski (1997), and Alford and Boatsman (1995).

Moreover, Ghysels et al. (2006) and Forsberg and Ghysels (2007) found that absolute returns are strong predictors of future volatility, supporting the idea that modeling the absolute value component enhance forecasts. Ding et al. (1993) also highlighted that absolute returns might be more effective than squared returns at capturing specific volatility components. Summarized, modelling absolute returns is promising as previous literature suggests it is predictable, and the data analysis and tests in this research reveals that absolute returns contain information.

### 4.2.2 Theoretical Justification for Decomposition

Christoffersen and Diebold (2006) introduced the decomposition of returns into sign and absolute value components, explaining that both components exhibit persistent dynamics, making them forecastable, whereas the returns themselves are not easily forecastable. This decomposition allows for a more granular analysis of market behavior and improves forecasting models by leveraging the distinct information provided by each component.

Korkie et al. (2002) already supported this idea in the past by showing that sign prediction directly affects heteroskedasticity in asset returns, increases prediction precision and can reduce this heteroskedasticity. Their findings suggest that a nonlinear return generating model, which separately models return signs and magnitudes, can effectively capture the dual contributions of these components.

### 4.2.3 Mathematical Framework

The returns $r_t$ are decomposed as:

$$r_t = sign(r_t) \times |r_t| \tag{4.10}$$

Modelling strategy is split up into a model for sign($r_t$)and a model for $|r_t|$. The first model assumes independence and hence

#### Sign Component

The sign component of the returns is mathematically defined as:

**Definition 4.1** (Sign component)**.**

$$\text{sign}(r_t) = \begin{cases} -1 & \text{, if } r_t < 0 \\ \\ 1 & \text{, if } r_t \geq 0 \end{cases}$$

, the sign of returns is hence binary valued. Let $\mathscr{F}_t = \sigma\{(r_s, x_s), s \leq t\}$ represent the sigma field containing information up to time $t$.[3] It then follows that,

$$\mathbb{I}_{\{r_t>0\}}|\mathscr{F}_{t-1} \sim B(p_t), \tag{4.11}$$

with density of sign$(r_t)|\mathscr{F}_{t-1}$ given as:

$$f_{\{I_{r_t>0}\}|\mathscr{F}_{t-1}}(x_t|\mathscr{F}_{t-1}) = p_t^{x_t}(1-p_t)^{1-x_t} \tag{4.12}$$

A binary response model is suitable here. The most common approaches are the linear probability model (LPM), logit, and probit models (Horowitz & Savin, 2001; Vasisht, 2007). LPM often predicts conditional probabilities

[3] We use $r_t$ instead of sign$(r_t)$ since that includes information on both the sign and magnitude

outside the $[0, 1]$ interval due to its linear nature. In contrast, the logit and probit models constrain predicted probabilities between 0 and 1 by using a nonlinear relationship between the dependent and explanatory variables. This results in an S-shaped cumulative distribution function (CDF) that better fits the data. We briefly discuss the implications below.

### Nonlinear Probability Model

Equation 4.11 implies that $P_{t-1}[sign(r_t) = 1] = p_t$ and $P_{t-1}[sign(r_t) = 0] = 1 - p_t$. We focus on the probability $p_t$ given the information at $t - 1$. As discussed by Horowitz and Savin (2001) and Vasisht (2007), $p_t$ is related to a linear function $\pi_t$ of explanatory variables $x_t$ through a transformation to keep probabilities between 0 and 1. Common transformations are the CDFs of the normal distribution ($\Phi(\cdot)$) for the probit model and the logistic distribution ($\Lambda(\cdot)$) for the logit model. This nonlinear approach ensures probabilities remain within bounds. The conditional expectation and probabilities are linked as:

$$\begin{aligned} \mathbb{E}_{t-1}\Big[sign(r_t)\Big] &= 0 \cdot P_{t-1}\Big[sign(r_t) = 0\Big] + 1 \cdot P_{t-1}\Big[sign(r_t) = 1\Big] \\ &= P_{t-1}\Big[sign(r_t) = 1\Big] = \Lambda(\pi_t) = p_t \end{aligned} \tag{4.13}$$

A probability close to 0 predicts a negative return, while a probability close to 1 predicts a positive return. We refer to $p_t$ as the probability of success. The probability of success and the linear transformation function are defined as follows:

$$p_t = \Lambda(\pi_t) = \frac{e^{\pi_t}}{1 + e^{\pi_t}}, \tag{4.14}$$

where

$$\pi_t = \omega + \delta_2 \pi_{t-1} + \delta_1 \mathbb{I}_{\{r_{t-1}>0\}} + \sum_{j=1}^{p} \beta_j x_{t-1,j} \tag{4.15}$$

We refer to the parameter set as $\Theta_L = \{\omega, \delta_1, \delta_2, \beta_1 \ldots \beta_p\}$ Contrary to the standard logit model, we include lagged sign as well in to otbain a dynamic logit model (Korkie et al., 2002). This should account for correlation between $\mathbb{I}_{t-1}$ and $\mathbb{I}_t$ (Nyberg, 2010, 2011). Table **??** shows the results for the LPM, which serves as an indication of what explanatory variables are relevant to include.

TABLE 4.1: Linear Probability Model Results

| **Variable** | **Coefficient** | **Standard Error** | **p-value** ($\alpha = 5\%$) |
|---|---|---|---|
| Intercept | 0.42269 | 0.0036627 | 0.000 |
| H2 | -0.0029237 | 0.0049506 | 0.555 |
| H3 | 0.030365 | 0.0049495 | 8.549e-10 |
| H4 | 0.037383 | 0.0049591 | 4.793e-14 |
| H5 | 0.051929 | 0.0049794 | 1.873e-25 |
| H6 | 0.040222 | 0.0049595 | 5.103e-16 |
| H7 | 0.059864 | 0.0049896 | 3.806e-33 |
| $\mathbb{I}_{\{r_{t-1}<0\}}$ | -0.19502 | 0.0033946 | 0.000 |
| $ir_3^*$ | 0.012988 | 0.00040661 | 5.332e-223 |

$^*ir_3$ is the ratio of the EUR 3M OIS and USD 3M OIS rates respectively.

The dummy for the second trading session of the day is excluded due to its coefficient being insignificant. The dynamic logit model for the sign component is given as:

$$p_t = \Lambda(\pi_t) = \frac{e^{\pi_t}}{1 + e^{\pi_t}} \tag{4.16}$$

$$\pi_t = \omega + \delta_1 \mathbb{I}_{\{r_{t-1}>0\}} + \delta_2 \pi_{t-1} + \beta_1 \cdot ir_3 t - 1 + \sum_{j=3}^{7} \beta_{j-1} H j_t$$

and the

From equation 4.12, we can specify the likelihood function to then estimate the parameters. The (conditional) log-likelihood function is given as:

$$l_t(\theta) \sum_{t=1}^{T} x_t \log(\Lambda(\pi_t)) + \sum_{t=1}^{T} (1 - x_t) \log(1 - \Lambda(\pi_t)), \tag{4.17}$$

where $\pi_t$ is specified in equation 4.15. Kauppi and Saikkonen (2008) have shown that although the function is complicated and nonlinear, solving it by numerical methods is very straightforward.

### Absolute Value Component

When the conditional distribution of returns is assumed to follow a Student's t-distribution, omitting the sign of the returns and observing only the absolute values results in what is known as a *folded Student's t-distribution*[4] (Psarakis & Panaretoes, 1990). Because we earlier assumed the conditional distribution of returns to have a third moment equal to zero, the density of the folded t-distribution is given below. Below, Figure 4.1 is added that should clarify the folding process. We visually compare the quantiles from the data against those from a very heavy-tailed Folded Student's t-distribution and a Standard Normal distribution, as shown in Appendix A.7. The QQ-plot indicates that the half-Student's t-distribution is a reasonable fit for most of the data. However, there are slight deviations in the tails, suggesting that it may not fully capture extreme values.

[4]Also known in literature as half Student's t-distribution.

$$f(x) = \begin{cases} \frac{2\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}, & x \geq 0 \\ 0, & x < 0 \end{cases} \tag{4.18}$$

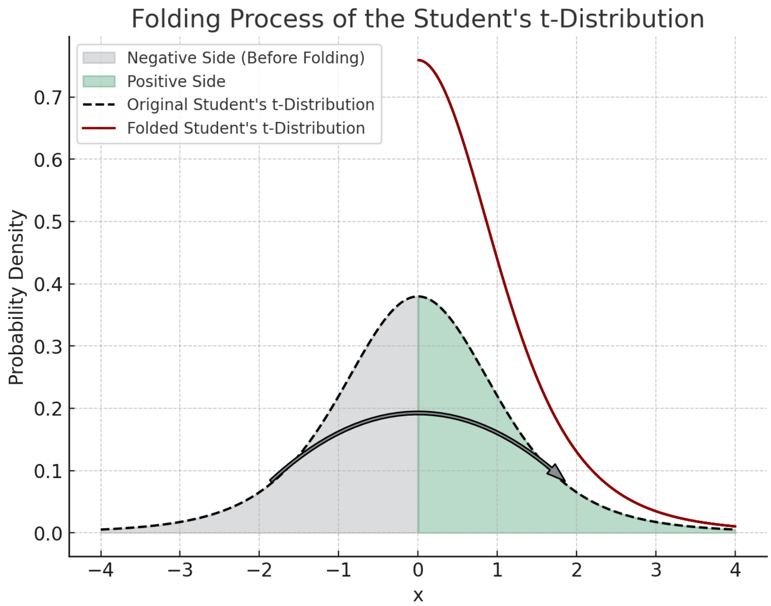


FIGURE 4.1: Folding Process of the Student's t-Distribution ($\nu = 1.9$)

In Chapter 3, we concluded that stationary absolute returns exhibit significant correlations at non-zero lags, indicating that the time series has an autocorrelation structure that could be captured by ARMA components (Box et al., 1970). Additionally, Table 3.5 demonstrated that the time of day significantly impacts absolute returns. Therefore, we consider an ARMAX model for the mean equation to address these two findings.

The ARMAX$(p, q)$ model for returns is specified as follows:

$$|r_t| = \mu + \sum_{i=1}^{p} \phi_i \cdot r_{t-i} + \sum_{j=1}^{q} \theta_j \cdot \varepsilon_{t-j} + \sum_{l=1}^{m} \sum_{k=1}^{p} \beta_{k,l} \cdot x_{t-k,l} + \varepsilon_t, \tag{4.19}$$

where $|r_t|$ is the observed absolute return and $\{\mu, \phi_i, \theta_j, \varepsilon_t,\}$ are the same parameters as in equation 4.1 with the extension to let an exogenous[5] variable $x_t$, affect $|r_t|$. The method of estimation remains unchanged, except that we include an extra term in the mean equation.

We expect, based on findings in chapter 3, that absolute returns shows some form of heteroskedasticity. After estimating the parameters for the ARMAX model using ML, we check the residuals for heteroskedasticty using the ARCH test. If ARCH effects cannot be rejected, we model conditional volatility by GARCH(1,1).

[5]Hence X in ARMAX.

Following Box et al. (1970) again to identify the correct ARMA order, we see a clear cutoff at the 1st in the (P)ACF plot in Figure 3.6. Hence, one AR term and one MA term are included. Another approach could be to estimate various ARMAX models and opt the one that gives the lowest AIC (Akaike, 1973). However, the parameter estimates and estimated standard errors from the marginal distribution will not be valid for the decomposition model. This means that the best model for absolute returns, is not per definition the best specified model in the decomposition model. The final ARMAX(1,1) model for absolute returns is given below:

$$r_t| = \mu + \phi_1 \cdot r_{t-1} + \theta_1 \cdot \varepsilon_{t-1} + \beta_1 \cdot H1_t + \beta_2 \cdot H2_t + \beta_3 \cdot H3_t + \beta_4 \cdot H4_t + \beta_5 \cdot H5_t + \beta_6 \cdot H7_t + \varepsilon_t \tag{4.20}$$

$$\varepsilon_t \sim \text{folded-}t(0, \sigma^2),$$

Then, since equation **??** is linear in $\varepsilon_t$, the conditional density of $|r_t|$ is given as:

$$f_{|r_t||\mathscr{F}_{t-1}}(x) = \begin{cases} \frac{2\,\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\sigma\,\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{(x-\mu_t)^2}{\nu\sigma^2}\right)^{-\frac{\nu+1}{2}}, & \text{if } x \geq 0, \\ 0, & \text{if } x < 0, \end{cases} \tag{4.21}$$

The conditional mean was modelled successfully as we were able to remove autocorrelation in the residuals. The test for ARCH effects shows that a volatility model is preferable. Upon estimating the ARMAX(1,1)-GARCH(1,1) model, we still cannot reject that ARCH effects are not present, which means the latter model is only able to partially describe the data. These findings have been documented in A.8. Below is an output of the estimation results for the model and the final model specification.

$$|r_t| = \mu + \phi_1 \cdot r_{t-1} + \theta_1 \cdot \varepsilon_{t-1} + \beta_1 \cdot H1_t + \beta_2 \cdot H2_t + \beta_3 \cdot H3_t + \beta_4 \cdot H4_t + \beta_5 \cdot H5_t + \beta_6 \cdot H7_t + \varepsilon_t$$

$$\sigma_t^2 = \omega + \alpha \cdot \sigma_{t-1}^2 + \beta \cdot \varepsilon_{t-1}^2 \tag{4.22}$$

$$\varepsilon_t \sim \text{folded-}t(0, \sigma^2)$$

Table 4.2: ARMA Model Results

| **Variable** | **Coefficient** $\times 10^{-4}$ | **t-Statistic** | **p-value** ($\alpha = 5\%$) |
|---|---|---|---|
| $\mu$ | 17.191 | 931.56 | 0.000 |
| $\phi_1$ | 0.88104 | 763.36 | 0.000 |
| $\theta_1$ | -0.77745 | -200.61 | 0.000 |
| $\beta_1$ | -1.1719 | -152.56 | 0.000 |
| $\beta_2$ | -3.1066 | -242.85 | 0.000 |
| $\beta_3$ | 2.2615 | 366.75 | 0.000 |
| $\beta_4$ | 5.7792 | 509.73 | 0.000 |
| $\beta_5$ | 3.4851 | 225.86 | 0.000 |
| $\beta_6$ | 9.8241 | 648.11 | 0.000 |

Table 4.3: GARCH Model Results

| **Variable** | **Coefficient** | **t-Statistic** | **p-value** ($\alpha = 5\%$) |
|---|---|---|---|
| $\omega$ | 0.21862 | 20.491 | 2.596e-93 |
| $p$ | 0.37134 | 315.11 | 0.000 |
| $q$ | 0.62866 | 21.66 | 4.862e-104 |
| $\nu$ | 2.096 | 428.55 | 0.000 |

### Copula

To capture the interaction term between $\text{sign}(r_t)$ and $|r_t|$, we combine them with a copula. A copula is a function that links multivariate distribution functions to their respective one-dimensional marginal distribution functions (Nelsen, 2006). Sklar's theorem, first discussed briefly in French Sklar (1959) and later on in more detail by Schweizer and Sklar (2011), is the foundation in modelling copulas. The theorem for a two-dimensional copula is given below.

**Theorem 4.2.1** (Sklar's Theorem)**.** *Let $X_1$ and $X_2$ be two random variables with marginal CDFs $F_1(\cdot)$ and $F_2(\cdot)$. Let $F(\cdot)$ be their joint distribution. Sklar's theorem then claims that the joint distribution of two marginal distributions $F_1$ and $F_2$ can be specified through a copula that takes as inputs the two marginal CDFs:*

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)), \tag{4.23}$$

*where $C(F_1(x_1), F_2(x_2)) : [0,1] \times [0,1] \mapsto [0,1]$ is the copula. We refer to $C(\cdot)$ as the copula of $F(\cdot)$. Given the bivariate copula in 4.23, the corresponding density function is given by:*

$$f(x_1, x_2) = \frac{\partial^2}{\partial x_1 \partial x_2} C(F_1(x_1), F_2(x_2)) f_1(x_1) f_2(x_2), \tag{4.24}$$

*which follows from the chain rule.*

From equation 4.24 it is clear that Sklar's theorem was designed for continuous random variables. Anatolyev and Gospodinov (2010) extended Sklar's theorem to be applicable to mixed bivariate distribution, as in this case. Key aspects of their proof include using finite differences instead of partial derivatives for the discrete random variable (in this case, the sign). Additionally, they were able to explicitly write out the possible finite differences because the sign variable only takes on two possible values.

**Theorem 4.2.2.** *Let $X_1$ and $X_2$ be a continuous and Bernoulli distributed random variable with marginal CDFs $F_1(\cdot)$ and $F_2(\cdot)$. Let $F(\cdot)$ be their joint distribution. Anatolyev and Gospodinov (2010) then claim that the joint density $f(\cdot)$ is given as:*

$$f(x_1, x_2) = f_1(x_1) \varrho_t(F_1(x_1))^{x_2} \Big(1 - \varrho_t(F_1(x_1))\Big)^{1-x_2}, \tag{4.25}$$

*where $\varrho_t(z) = 1 - \frac{\partial C(z, 1-p_t)}{\partial u_1}$.*

As the joint density $f(x_1, x_2)$ is available, the likelihood function can now be specified (Bain & Engelhardt, 1992):

$$\prod_{j=1}^{T} f_1(x_1)\varrho_t(F_1(x_1))^{x_2}\Big(1-\varrho_t(F_1(x_1))\Big)^{1-x_2},$$

ans the log-Likelihood function as well:

$$\begin{aligned}\log\left(\prod_{j=1}^{T} f(x_1,x_2)\right) &= \sum_{t=1}^{T}\log\left(f_1(x_{1,t}|\mathscr{F}_{t-1})\right)+\sum_{t=1}^{T}\left(1-\varrho\left(F_1(x_{1,t}|\mathscr{F}_{t-1})\right)\right)^{1-x_{2,t}}\\ &\quad+\sum_{t=1}^{T}\log\left(\varrho\left(F_1(x_{1,t}|\mathscr{F}_{t-1})\right)\right)^{x_{2,t}}\\ &= \sum_{t=1}^{T}\log\left(f_1(x_{1,t}|\mathscr{F}_{t-1})\right)\\ &\quad+\sum_{t=1}^{T}\Big(1-\mathbb{I}_{\{r_t>0\}}\Big)\log\left(1-\varrho(F_1(x_{1,t}|\mathscr{F}_{t-1}))\right)\\ &\quad+\sum_{t=1}^{T}\mathbb{I}_{\{r_t>0\}}\log\left(\varrho(F_1(x_{1,t}|\mathscr{F}_{t-1}))\right)\end{aligned} \tag{4.26}$$

Note that the returns decomposition can be written as $|r_t|\cdot sign(r_t)=|r_t|\cdot(2\mathbb{I}_{\{r_t>0\}}-1)$ and future returns can be predicted by:

$$\mathbb{E}_t(r_{t+1})=2\mathbb{E}_t(|r_{t+1}|\mathbb{I}_{\{r_{t+1}>0\}})-\mathbb{E}_t(|r_{t+1}|)=\hat{r}_{t+1}, \tag{4.27}$$

In the case of conditional independence of he two components, this simplifies to

$$\mathbb{E}_t(r_{t+1})=2\mathbb{E}_t(|r_{t+1}|)\mathbb{E}_t(\mathbb{I}_{\{r_{t+1}>0\}})-\mathbb{E}_t(|r_{t+1}|)=2(\hat{p_{t+1}}-1)|\widehat{r_{t+1}}| \tag{4.28}$$

When incorporating dependence, $\mathbb{E}_t(|r_{t+1}|)$ can just be computed directly from $f_1(|r_{t+1}||\mathscr{F}_t)$ and the conditional expectation of the cross-product $\xi_{t+1}=2\mathbb{E}_t(|r_{t+1}|\mathbb{I}_{\{r_{t+1}>0\}})$ can be computed as follows: First, by the law of iterated expectations:

$$\begin{aligned}2\mathbb{E}_t\Big(|r_{t+1}|\mathbb{I}_{\{r_{t+1}>0\}}\Big) &= 2\mathbb{E}_t\Big(\mathbb{E}_t\Big[|r_{t+1}|\mathbb{I}_{\{r_{t+1}>0\}}\mid\mid r_{t+1}\mid\Big]\Big)\\ &= 2\mathbb{E}_t\Big(|r_{t+1}|\mathbb{E}_t\Big[\mathbb{I}_{\{r_{t+1}>0\}}\mid\mid r_{t+1}\mid\Big]\Big)\end{aligned}$$

Hence, we need the density of $\mathbb{I}_{\{r_{t+1}>0\}}$ given $r_{t+1}$, given as:

$$f_{\mathbb{I}_{\{r_{t+1}>0\}}||r_{t+1}|}(x_{2,t+1}\mid |r_{t+1}|)=\frac{f_{r_{t+1}}(r_{t+1},x_{2,t+1})}{f_{|r_{t+1}|}(\mid r_{t+1}\mid)}= \tag{4.29}$$

$$\varrho_t(F_1(x_1))^{x_{2,t+1}}\Big(1-\varrho_t(F_1(x_1))\Big)^{1-x_{2,t+1}},$$

and the conditional expectation of $\mathbb{I}_{\{r_{t+1}>0\}}$ given $\mid r_{t+1}\mid$ is

$$\mathbb{E}_t\Big[\mathbb{I}_{\{r_{t+1}>0\}}\mid\mid r_{t+1}\mid\Big]=0\cdot f_{\mathbb{I}_{\{r_{t+1}>0\}}||r_{t+1}|}(0\mid|r_t|)+1\cdot f_{\mathbb{I}_{\{r_{t+1}>0\}}||r_{t+1}|}(1\mid|r_t|)$$

$$= \varrho_{t+1}(F_1(|r_{t+1}|\mathscr{F}_t). \tag{4.30}$$

Using equations 4.29 and 4.30, the expectation of the cross-product is given as

$$\xi_{t+1} = 2\mathbb{E}_t(|r_{t+1}|\mathbb{I}_{\{r_{t+1}>0\}})$$
$$= \int_0^\infty u f_{|r_{t+1}||\mathscr{F}_t)}(| r_{t+1} |) \cdot (\varrho_{t+1}(F_{|r_{t+1}|}(|r_{t+1}||\mathscr{F}_t))du \tag{4.31}$$

As 4.31 cannot be solved for analytically (Anatolyev & Gospodinov, 2010), we make an attempt at approximating it numerically. In contrast to the latter paper, we do not perform a change of variable and approximate the improper integral directly. As the sign model and the absolute returns model are both specified, completing the decomposition model means specifying the copula. Using a copula of the Archimedean family implies a simple form with closed-form solution for the dependence structure (Nelsen, 1997). This is highly desirable when one is working with complex marginals. Besides simplicity, it also takes into account upper and lower tail dependence which is necessary considering the data used in this research (Joe, 1993). The definition of an Archimedean copula is given below.

**Definition 4.2** (Archimedean Copula Form)**.** Archimedean copulas are specified through their respective generator function, $\phi(\cdot)$:

$$C(F_1(x_1), F_2(x_2)) = \phi^{-1}(\phi(F_1(x_1) + \phi(F_2(x_2)),$$

where the *generator function* $\phi(\cdot)$ must satisfy the following three conditions (Boateng et al., 2022):

1. $\phi$ is a continuous, strictly decreasing, and convex function, mapping from the domain $[0,1]$ to $[0,\infty)$.
2. $\phi(0)$ is equal to infinity, i.e., $\phi(0) = \infty$.
3. $\phi(1)$ is equal to zero, i.e., $\phi(1) = 0$.

The Archimedean family of copulas is widely regarded for its flexibility, with the Clayton, Frank, Gumbel and Farlie-Gumbel-Morgenstern copulas being among the most popular options. Given that sign and absolute return show modest correlation, the FGM copula (Émile & Gumbel, 1960) is particularly well-suited, as it effectively captures dependence while maintaining a relatively simple analytical form. The definition of the FGM copula can be found below.

**Definition 4.3** (FGM Copula)**.** The Farlie-Gumbel-Morgenstern (FGM) copula is given by:

$$C(F_1(x_1), F_2(x_2)) = F_1(x_1)F_2(x_2)\left[1 + \theta(1 - F_1(x_1))(1 - F_2(x_2))\right],$$

where $\theta \in [-1, 1]$ is the dependence parameter.

We can now derive $\varrho_t(\cdot)$ as defined in Equation 4.25. For reference, this derivation is detailed in Equation D.1. The full decomposition model, which consists of the conditional distribution of the sign, the conditional distribution of the absolute value of returns, and a copula term, is specified below through its log-likelihood function, as presented in Equation 4.26. During the estimation procedure, we minimize the negative log-likelihood.

$$\begin{aligned} l(\theta|\mathscr{F}_t) = &\sum_{t=1}^{T} \log\left(f_1(x_{1,t}|\mathscr{F}_t)\right) \\ &+ \sum_{t=1}^{T} \left(1 - \mathbb{I}_{\{r_t>0\}}\right) \log\left(1 - \varrho(F_1(x_{1,t}|\mathscr{F}_t))\right) \\ &+ \sum_{t=1}^{T} \mathbb{I}_{\{r_t>0\}} \log\left(\varrho(F_1(x_{1,t}|\mathscr{F}_t))\right) \end{aligned} \tag{4.32}$$

where:

- $f_1(x_{1,t}|\mathscr{F}_{t-1}) = f_{|r_t|}(|r_t| \mid \mathscr{F}_{t-1}) = \frac{2\cdot\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\cdot\Gamma\left(\frac{\nu}{2}\right)\cdot\sigma_t\cdot\left(1+\frac{(|r_t|-\mu_t)^2}{\nu\sigma_t^2}\right)^{\frac{\nu+1}{2}}}$
- $F_1(x_{1,t}|\mathscr{F}_{t-1}) = F_{|r_t|}(|r_t| \mid \mathscr{F}_{t-1}) = 2\cdot T_\nu\left(\frac{x-\mu_t}{\sigma_t}\right) - 1$
- $\mu_t = \mu + \phi_1 r_{t-1} + \theta_1\varepsilon_{t-1} + \beta_1 H1_t + \beta_2 H2_t + \beta_3 H3_t + \beta_4 H4_t + \beta_5 H5_t + \beta_6 H7_t$
- $\varrho_t(z) = (1-p_t)\left[1+\theta p_t(1-2z)\right]$, where $\theta = 0$ implies independence.
- $p_t = \Lambda(\pi_t)$

# Chapter 5

# Results

### Set-up

This chapter presents the results from the SARIMA(2,0,4)(1,0,1)-GARCH(1,1) model and the decomposition model, both estimated using MATLAB. MATLAB's advanced algorithms effectively minimize the negative log-likelihood and approximate improper integrals, ensuring robust model estimation. The chapter begins with an analysis of the benchmark SARIMA-GARCH model, detailing the estimated parameters and the fitted conditional mean. The model diagnostics are conducted according to the Box-Jenkins methodology.

Following this, the decomposition model is examined, focusing on the estimated coefficients and the conditional mean fit through in-sample predictions. We also evaluate the out-of-sample performance using one-step-ahead forecasts. To compare the models, we assess their accuracy by calculating the root mean squared error (RMSE). Additionally, a confusion matrix is presented to evaluate the sign accuracy of the decomposition model. Finally, the chapter reviews seasonality by analyzing the seasonal coefficients and their significance in the benchmark conditional mean model and the decomposition model.

## 5.1 SARIMA-GARCH

The initial SARIMA(2,0,4)(1,0,1)-GARCH(1,1) is not very informative. It confirms the heavy tails by the degrees of freedom close to 2, but all other estimated coefficients are insignificant. This model might be inadequate to capture the returns dynamics. To confirm this, we tested the residuals using the Ljung-Box test which rejects the null hypothesis of no autocorrelation in the model's residuals. Below is the ACF plot added from which it also becomes evident that residuals still exhibit autocorrelation. Therefore we decide to remove the seasonal parameters and reduce the model to ARIMAX(2,0,4)-GARCH(1,1). The reason that $\beta$ is not removed is because there is evidence that on average, returns during the mid-morning session (H3) are statistically different from other sessions during the day. Thereby, we aim to obtain the most parsimonious model and the seasonal parameters do not to contribute to that goal in this model. The constant term is also removed as it estimated to be very close to zero and insignificant. Below the estimation results for the SARIMAX(2,0,4)(1,0,1)-GARCH(1,1) are added.

| **Parameter** | **Value** | **Standard Error** | **T-Statistic** | **P-Value** |
|---|---|---|---|---|
| **ARIMAX(2,0,4) Conditional Mean Model** | | | | |
| Constant | -1.4736e-06 | 3.2223e-05 | -0.045731 | 0.96352 |
| AR1 | -0.80348 | 0.0010368 | -774.95 | 0 |
| AR2 | -0.50709 | 0.0012874 | -393.9 | 0 |
| SAR168 | -0.0017823 | 0.010357 | -0.17209 | 0.86337 |
| MA1 | 0.75965 | 0.0010802 | 703.23 | 0 |
| MA2 | 0.46839 | 0.001372 | 341.39 | 0 |
| MA3 | -0.026115 | 0.0024954 | -10.465 | 1.2482e-25 |
| MA4 | -0.0015464 | 0.0017931 | -0.86245 | 0.38844 |
| SMA168 | 0.0017462 | 0.010364 | 0.16848 | 0.8662 |
| Beta(1) | 1.5171e-05 | 8.6442e-05 | 0.17551 | 0.86068 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 2e-07 | 3.3189e-08 | 6.0261 | 1.6801e-09 |
| $\alpha$ | 0.70333 | 0.0011681 | 602.1 | 0 |
| $\beta$ | 0.29667 | 0.0038674 | 76.711 | 0 |
| $\nu$ | 2.5117 | 0.0083938 | 299.23 | 0 |
| AIC | -655375 | | | |

Table 5.1: Estimation Results for SARIMAX(2,0,4)(1,0,1)-GARCH(1,1) Model with t Distribution

We proceed in the same manner by fitting the reduced ARIMAX(2,0,4)-GARCH(1,1) model to the data. The estimation results are given below:

| Parameter | Value | Standard Error | T-Statistic | P-Value |
|---|---|---|---|---|
| **ARIMAX(2,0,4) Conditional Mean Model** | | | | |
| $\phi_1$ | 0.53505 | 0.031684 | 16.887 | 0.0000 |
| $\phi_2$ | 0.29168 | 0.03525 | 8.2745 | 0.0000 |
| $\theta_1$ | -0.57883 | 0.031697 | -18.261 | 0.0000 |
| $\theta_2$ | -0.2719 | 0.033888 | -8.0235 | 0.0000 |
| $\theta_3$ | 0.014627 | 0.0025944 | 5.638 | 0.0000 |
| $\theta_4$ | 0.0037454 | 0.0018167 | 2.0616 | 0.0000 |
| $\beta$ | 2.0786e-06 | 8.1706e-06 | 0.25439 | 0.79919 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 2e-07 | 3.3203e-08 | 6.0235 | 0.000 |
| $\alpha$ | 0.70364 | 0.0011668 | 603.07 | 0.000 |
| $\beta$ | 0.29636 | 0.0038668 | 76.642 | 0.000 |
| $\nu$ | 2.5108 | 0.0083872 | 299.37 | 0 |
| AIC | $-656397$ | | | |

Table 5.2: Estimation Results for ARIMAX(2,0,4)-GARCH(1,1) Model with t Distribution

Firstly, the AR and MA coefficients are all highly significant. The coefficient for the exogenous dummy variable for the mid-morning session, is not significant. We see that the ARIMAX(2,0,4)-GARCH(1,1) provides a better fit. Thereby, according to the AIC, the ARIMAX(2,0,4)-GARCH(1,1) is the preferred model. Hence, following the Box-Jenkins approach, we continue with this model.

### Model Diagnostics

We have checked for parameter significance. Now we review the model by checking (i) the ACF plot, (ii) the PACF plot, (iii) performing the Ljung-Box test and (iv) performiing the ARCH-LM test.

- **ACF Plot:** The ACF plot shows very little autocorrelation, but still significant.
- **PACF Plot:** The PACF confirms what the ACF found. Although correlations have declined significantly, they are still significant.
- **Ljung-Box Test:** The Ljung-Box test confirms autocorrelation in the model's residuals at the 5% significance level.
- **ARCH-LM Test:** The ARCH-LM test fails to reject that there is no heteroskedasticity present in the residuals at the 5% significance level.

The ARIMAX-GARCH(1,1) model successfully captures and reduces a significant portion of the autocorrelation in the data. However, residual analysis reveals persistent autocorrelation, which suggests that the model may not be fully specified. Statistical tests confirm that the null hypothesis of no autocorrelation in the residuals is rejected, indicating potential misspecification of the model.

Moreover, the conditional volatility model incorporated within the GARCH(1,1) framework fails to adequately capture the heteroskedasticity present in the data. This is evidenced by the results of the ARCH-LM test, which indicate that significant volatility clustering remains unexplained.

While the model does not identify significant intraday seasonality effects, this result should be interpreted cautiously. The lack of detected seasonality may point to potential misspecification of the model rather than a definitive absence of seasonality. It is possible that the model, at least in its current form, does not fully capture the complexities of the underlying time series, particularly with respect to intraday patterns and conditional volatility.

#### Out-of-sample performance

As the benchmark model is not well specified, meaning it fails capturing the returns' dynamics, we want to ensure that the changing market characteristics are not a reason for this. We split the data up in four equal parts. Then, we estimate the simplified models separately on 80% of the quartered data, respectively. This results in four distinct models, allowing us to assess whether the models perform better over shorter time horizons. The remaining 20% of the data from each year is reserved for testing out-of-sample one-step-ahead forecasts.

None of the models represent the data properly, indicating misspecification once more. The (P)ACF plots and estimation resutls have been added to the Appendix C. Also here, we reject the Ljung-Box test and the ARCH-LM test. The third model indicates that returns are significantly more positive during the mid-morning session. We do not see improvement from the previous model. Meaning, that the changing dynamics, if any, are not the issue for this model. The out-of-sample one-step ahead forecasts are shown below.

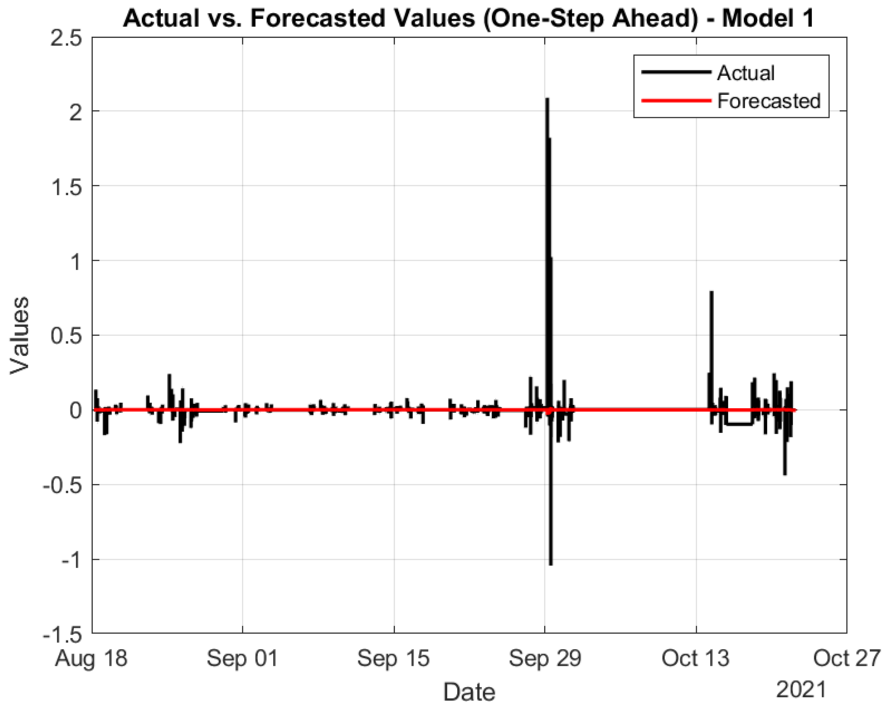


RMSE: 0.044548

FIGURE 5.1: Out-of-sample Benchmark Model 1: $r_t$ and $E_{t-1}(r_t)$

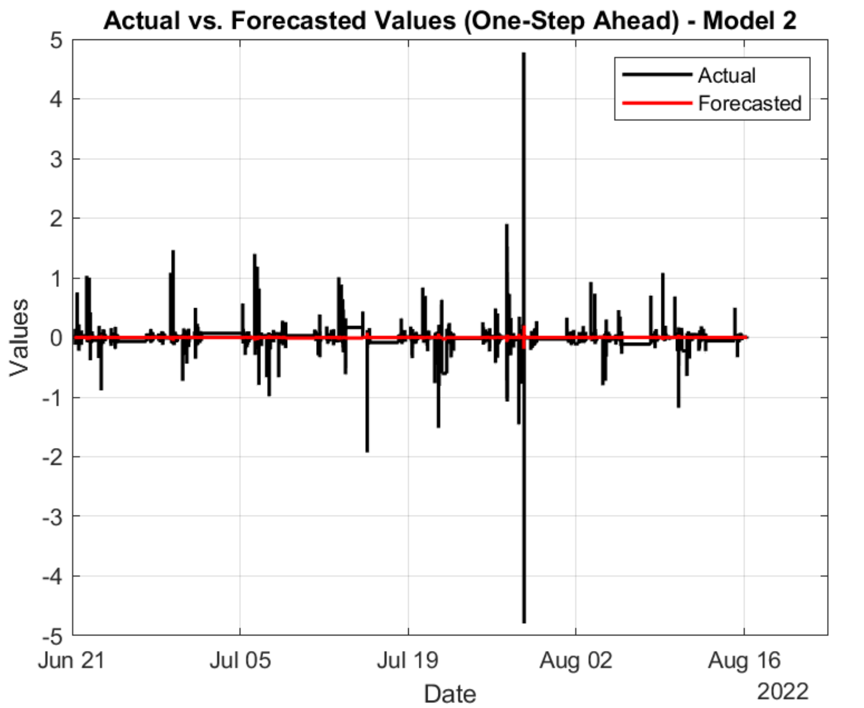


RMSE: 0.13355

FIGURE 5.2: Out-of-sample Benchmark Model 2: $r_t$ and $E_{t-1}(r_t)$

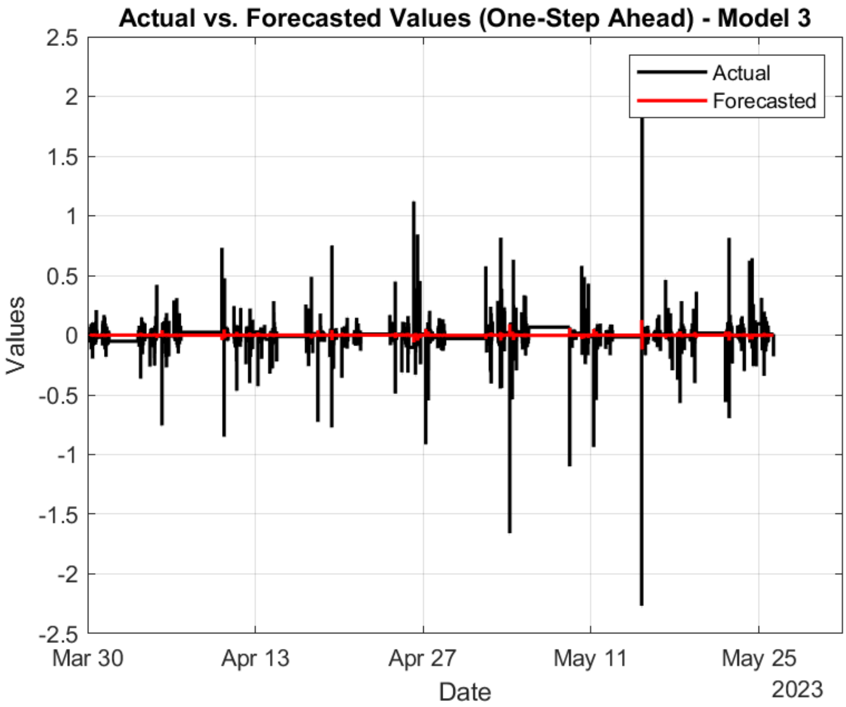


RMSE:0.082406

FIGURE 5.3: Out-of-sample Benchmark Model 3: $r_t$ and $E_{t-1}(r_t)$

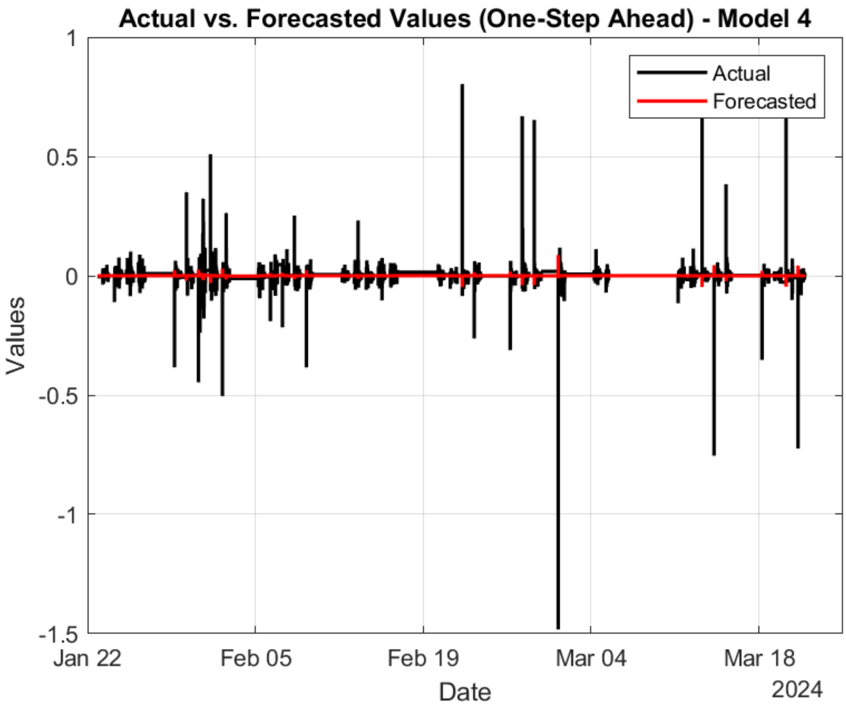


RMSE: 0.040332

FIGURE 5.4: Out-of-sample Benchmark Model 4: $r_t$ and $E_{t-1}(r_t)$

## 5.2 Decomposition Model

In the decomposition model, the sign component is represented by equation 4.16 and and the absolute value component is captured by equation 4.22. Predictions are derived using equations 4.27 and 4.28. The results, however, reveal no statistical significance for the estimated coefficients, which is counterintuitive when considering the findings from Chapter 3. This outcome may suggest that the model is overly complex. To investigate this further, we later simplify the parametric representation of the absolute value component by removing the dummy variables. Although seasonality is then modeled explicitly in the sign component, the joint density function described in equation 4.25 continues to account for seasonality through the copula term. The in-sample fit, in-sample probabilities, and a table with the estimation results are presented below.

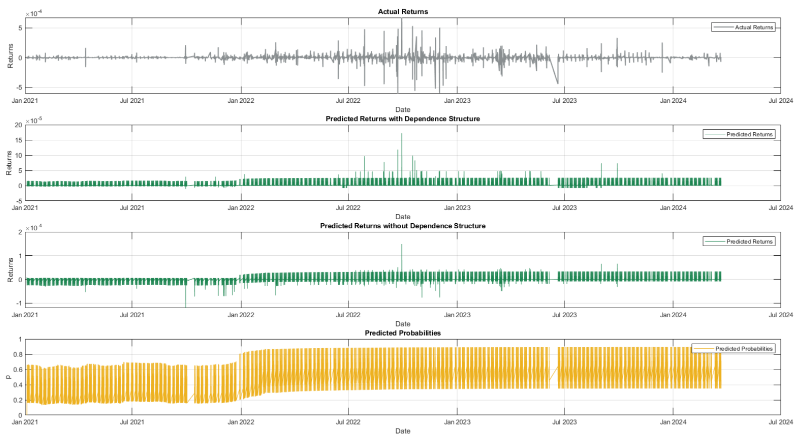


FIGURE 5.5: In-sample Fits: $r_t, E_{t-1}(r_t), p_t$

| | **Predicted Negative** | **Predicted Positive** | **Total** |
|---|---|---|---|
| **Actual Negative** | 5.582% | 44.541% | 50.123% |
| **Actual Positive** | 5.5268% | 44.351% | 49.878% |
| **Total** | 11.1088% | 88.892% | 100% |

TABLE 5.3: Confusion Matrix in Percentages (with Dependence Structure). Total accuracy: **0.49933**

| | **Predicted Negative** | **Predicted Positive** | **Total** |
|---|---|---|---|
| **Actual Negative** | 29.361% | 20.762% | 50.123% |
| **Actual Positive** | 29.275% | 20.603% | 49.878% |
| **Total** | 58.636% | 41.365% | 100% |

TABLE 5.4: Confusion Matrix in Percentages (without Dependence Structure). Total accuracy: **0.49963**

### Independence

The model that assumes independence provides a poor in-sample fit. The predicted returns are significantly smaller in magnitude compared to the actual returns and show a persistent underestimation, particularly during periods of high volatility. The model's predictions exhibit a cyclical pattern, likely due to overfitting caused by the inclusion of seasonal dummies. This overfitting results in the model capturing regular, seasonal noise rather than the true underlying market dynamics, leading to inaccurate and overly smoothed predictions. As a result, the model fails to capture the large positive and negative returns. The RMSE for this model is 0.13393.

### Incorporating Dependence

Incorporating dependence into the model yields a slightly better fit, with an RMSE of 0.12946. While this model captures some aspects of volatility clustering—such as during the interest rate hike period discussed in Chapter **??**—it still suffers from a significant underestimation of return magnitudes. The cyclical pattern in the predicted returns suggests that the model may be overfitting to seasonal patterns, driven by the seasonal dummies. This overfitting causes the model to focus on predictable, repetitive cycles rather than accurately responding to actual market fluctuations, particularly during periods of market stress. Although the model is somewhat better at identifying periods of increased volatility, it struggles to predict the true magnitude of returns.

To simplify the model representation and avoid overfitting, we remove the seasonal dummy variables in the absolute returns component except for $H3$, which was the dummy for the mid-morning session. The reason we leave $H_3$ in the conditional mean equation is because it was the only trading session that exhibited significantly different return, variance and absolute return. The model for absolute returns is then reduced to:

$$|r_t| = \mu + \phi_1 \cdot r_{t-1} + \theta_1 \cdot \varepsilon_{t-1} + \beta_1 + \beta_3 \cdot H3_t + \varepsilon_t \tag{5.1}$$

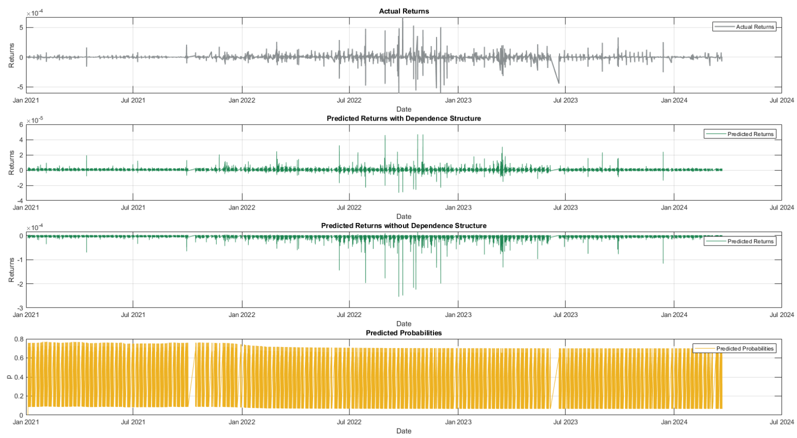


FIGURE 5.6: In-sample Fits: $r_t, E_{t-1}(r_t), p_t$

| | Predicted Negative | Predicted Positive | Total |
|---|---|---|---|
| **Actual Negative** | 25.178% | 24.945% | 50.123% |
| **Actual Positive** | 24.84% | 25.038% | 49.878% |
| **Total** | 50.018% | 49.983% | 100% |

Table 5.5: Confusion Matrix in Percentages (with Dependence Structure). Total accuracy: **50.216%**

| | Predicted Negative | Predicted Positive | Total |
|---|---|---|---|
| **Actual Negative** | 40.993% | 9.1299% | 50.123 |
| **Actual Positive** | 40.554% | 9.3231% | 49.877 |
| **Total** | 81.547% | 18.453% | 100% |

Table 5.6: Confusion Matrix in Percentages (without Dependence Structure). Total accuracy: **50.316%**

### Independence

The model that assumes independence provides a poor in-sample fit. It tends to overestimate negative sign predictions, particularly during volatile periods, indicating a bias toward predicting negative returns. This model has an RMSE of 0.11875.

### Incorporating Dependence

Incorporating dependence into the model results in a modest improvement, with a reduced RMSE of 0.10446. This version of the model captures some aspects of volatility clustering, such as during the interest rate hike period discussed in Chapter **??**. However, the model still primarily focuses on the central mean, leading to a significant underestimation of return magnitudes. While it generally follows the direction of returns, it struggles to predict extreme returns or significant directional changes accurately. The RMSE for the in-sample one-step-ahead predictions produced by this model is slightly higher at 0.12946, suggesting that, despite some improvements, the model still faces challenges in fully capturing market dynamics.

Although intraday returns and intraday absolute returns have remained relatively stable over the years, neither model adequately captures the dynamics necessary to model the conditional mean effectively. Exploiting dependence does slightly improve the model's performance, but not enough to overcome its limitations. To rule out the possibility that changing market dynamics are causing these shortcomings, we split the data up in four equal parts. Then, we estimate the simplified models separately on 80% of the quartered data, respectively. This results in four distinct models, allowing us to assess whether the models perform better over shorter time horizons. The remaining 20% of the data from each year is reserved for testing out-of-sample one-step-ahead forecasts. As we the model that incorporates dependencies has a slight edge over the model assuming conditional independence, the one-step ahead forecasts are based on the former.

### Estimated Models

We estimated four models:

- **Model 1:** 2021-01-04 - 2021-10-22
- **Model 2:** 2021-10-22 - 2022-08-15
- **Model 3:** 2022-08-16 - 2023-05-25
- **Model 4:** 2023-05-25 - 2024-03-21

We train the models on 80% of the data, and test it on the remaining 20% to see out-of-sample prediction. Below we provide the estimation results for all four models:

#### Model 1:

In the first time period, all coefficients for the absolute value component became significant. The logit model indicates no intraday seasonality during this period. However, the ARIMAX model reveals a significant decrease in return magnitude during the mid-morning session. It is important to note that, despite statistical significance, the estimated coefficient is very small, indicating that the effect, while present, is minimal.

#### Model 2:

The second time period suggests a slight shift in market dynamics. During this period, the mid-morning session is not significantly different in magnitude from the rest of the day. However, the logit model shows that all trading sessions have a significantly higher probability of positive returns compared to the pre-London open and London open sessions, indicating a potential shift in market sentiment during these periods.

#### Model 3:

In the third model, the results suggest that return magnitude is significantly lower during the mid-morning session. While all ARMAX coefficients are significant, the GARCH model's parameters are largely insignificant. This may be due to the overall high volatility throughout this period, where volatility remains elevated rather than increasing. Consequently, the model does not need to capture an upward trend in volatility. Additionally, the insignificance of $\beta_1$ suggests that interest rates have minimal explanatory power over return signs during this period, which aligns with the relatively stable interest rates observed. The logit model also indicates that the mid-afternoon and London-close sessions have significantly higher probabilities of positive returns.

#### Model 4:

The final model for the last segment of the data suggests that mid-morning sessions exhibit significantly lower return magnitude compared to other sessions throughout the day. The logit coefficients $\beta_3$ and $\beta_5$ indicate a higher probability of observing positive returns during lunchtime and the London-close session, suggesting a shift in market dynamics during these times.

Overall, the data likely exhibit changing dynamics that the decomposition model struggles to capture consistently across all periods. However, when applied to shorter timeframes, the decomposition model reveals several interesting findings. These results suggest that return signs and magnitudes may be separately predictable, even though the returns themselves exhibit limited predictability over longer horizons. The copula term $\beta_\varrho$, which is used to capture interdependencies, is highly significant across all models. This supports the decision to include the dependence structure in the model.

| **Model 1** | | | | |
|---|---|---|---|---|
| **Parameter** | **Value** | **Standard Error** | **T-Statistic** | **P-Value** |
| **ARIMAX(1,0,1)** | | | | |
| $\mu$ | 0.000000 | 0.0000102 | 11.1953 | 0.0000 |
| $\phi_1$ | 0.7847282 | 0.0203699 | 38.5239 | 0.0000 |
| $\theta_1$ | -0.7741531 | 0.0211684 | -36.5712 | 0.0000 |
| $\beta_3$ | -0.0000296 | 0.0000109 | -2.7218 | 0.0065 |
| **GARCH(1,1)** | | | | |
| $\omega_{garch}$ | 0.0000010 | 0.0000001 | 9.5326 | 0.0000 |
| $\alpha_1$ | 0.0003873 | 0.0000173 | 22.3242 | 0.0000 |
| $\beta_{garch}$ | 0.6364053 | 0.0210072 | 30.2946 | 0.0000 |
| **Logit** | | | | |
| $\omega_{logit}$ | -0.1769799 | 0.0280 | -6.3259 | 0.0000 |
| $\delta_1$ | 0.1649163 | 0.0213 | 7.7331 | 0.0000 |
| $\delta_2$ | 0.8652023 | 0.0200 | 43.1795 | 0.0000 |
| $\beta_1$ | -0.0036917 | 0.0032 | -1.1423 | 0.2533 |
| $\beta_2$ | 0.0098350 | 0.0241 | 0.4084 | 0.6830 |
| $\beta_3$ | 0.0007519 | 0.0153 | 0.0492 | 0.9608 |
| $\beta_4$ | -0.0026227 | 0.0171 | -0.1532 | 0.8782 |
| $\beta_5$ | 0.0202300 | 0.0134 | 1.5051 | 0.1323 |
| $\beta_6$ | 0.0139708 | 0.0211 | 0.6611 | 0.5085 |
| $\beta_\rho$ | 0.9999973 | 0.0250 | 39.9462 | 0.0000 |
| $\nu_t$ | 2.0000012 | 0.0326 | 61.4108 | 0.0000 |

TABLE 5.7: Model 1

| **Model 2** | | | | |
|---|---|---|---|---|
| **Parameter** | **Value** | **Standard Error** | **T-Statistic** | **P-Value** |
| **ARIMAX(1,0,1)** | | | | |
| $\mu$ | 0.0000000 | 0.00002708 | 0.7777 | 0.4368 |
| $\phi_1$ | 0.8091541 | 0.0062851 | 128.7425 | 0.0000 |
| $\theta_1$ | -0.6327137 | 0.0079712 | -79.3754 | 0.0000 |
| $\beta_3$ | -0.00002829 | 0.00003941 | -0.7178 | 0.4729 |
| **GARCH(1,1)** | | | | |
| $\omega_{garch}$ | 0.00000100 | 0.00000121 | 0.8299 | 0.4066 |
| $\alpha_1$ | 0.1347648 | 0.0071485 | 18.8521 | 0.0000 |
| $\beta_{garch}$ | 0.5008652 | 0.0187407 | 26.7260 | 0.0000 |
| **Logit** | | | | |
| $\omega_{logit}$ | -0.1226790 | 0.0069498 | -17.6523 | 0.0000 |
| $\delta_1$ | 0.1944689 | 0.0106114 | 18.3264 | 0.0000 |
| $\delta_2$ | 0.8964339 | 0.0076559 | 117.0909 | 0.0000 |
| $\beta_1$ | 0.0027111 | 0.0008204 | 3.3045 | 0.0010 |
| $\beta_2$ | 0.0435253 | 0.0065471 | 6.6480 | 0.0000 |
| $\beta_3$ | 0.0213757 | 0.0060848 | 3.5130 | 0.0004 |
| $\beta_4$ | 0.0300095 | 0.0060662 | 4.9470 | 0.0001 |
| $\beta_5$ | 0.0253480 | 0.0059046 | 4.2930 | 0.0000 |
| $\beta_6$ | 0.0302626 | 0.0071240 | 4.2480 | 0.0000 |
| $\beta_\rho$ | 0.6066349 | 0.0133965 | 45.2831 | 0.0000 |
| $\nu_t$ | 2.0000010 | 0.0619950 | 32.2608 | 0.0000 |

Table 5.8: Model 2

| **Model 3** | | | | |
|---|---|---|---|---|
| **Parameter** | **Value** | **Standard Error** | **T-Statistic** | **P-Value** |
| **ARIMAX(1,0,1)** | | | | |
| $\mu$ | 0.0000000 | 0.0009185 | 2.5465 | 0.0109 |
| $\phi_1$ | 0.7409081 | 0.3104105 | 2.3869 | 0.0170 |
| $\theta_1$ | -0.5902935 | 0.2720663 | -2.1697 | 0.0300 |
| $\beta_3$ | -0.0008271 | 0.0004003 | -2.0660 | 0.0388 |
| **GARCH(1,1)** | | | | |
| $\omega_{garch}$ | 0.00006412 | 0.0001070 | 0.5991 | 0.5491 |
| $\alpha_1$ | 0.0883221 | 0.2141695 | 0.4124 | 0.6801 |
| $\beta_{garch}$ | 0.3614523 | 0.3750 | 0.9639 | 0.3351 |
| **Logit** | | | | |
| $\omega_{logit}$ | -0.7330280 | 1.9945 | -0.3675 | 0.7132 |
| $\delta_1$ | -0.0643030 | 0.0910791 | -0.7060 | 0.4802 |
| $\delta_2$ | -0.4548490 | 0.9486944 | -0.4794 | 0.6316 |
| $\beta_1$ | 0.4470149 | 3.8287 | 0.1168 | 0.9071 |
| $\beta_2$ | 0.1656757 | 0.1674022 | 0.9897 | 0.3223 |
| $\beta_3$ | 0.3366087 | 1.3314 | 0.2528 | 0.8004 |
| $\beta_4$ | 0.5170429 | 0.3947712 | 1.3097 | 0.1903 |
| $\beta_5$ | 0.5605975 | 0.2410495 | 2.3257 | 0.0200 |
| $\beta_6$ | 0.6029269 | 0.2828521 | 2.1316 | 0.0330 |
| $\beta_\rho$ | 0.2758103 | 0.1285824 | 2.1450 | 0.0320 |
| $\nu_t$ | 2.0000012 | 5.9117 | 0.3383 | 0.7351 |

TABLE 5.9: Model 3

| **Model 4** | | | | |
|---|---|---|---|---|
| **Parameter** | **Value** | **Standard Error** | **T-Statistic** | **P-Value** |
| **ARIMAX(1,0,1)** | | | | |
| $\mu$ | 0.0000000 | 0.0007604 | 3.2304 | 0.0012 |
| $\phi_1$ | 0.6293925 | 0.1037539 | 6.0662 | 0.0000 |
| $\theta_1$ | -0.5659701 | 0.1055828 | -5.3604 | 0.0000 |
| $\beta_3$ | -0.0004260 | 0.0001749 | -2.4359 | 0.0149 |
| **GARCH(1,1)** | | | | |
| $\omega_{garch}$ | 0.00000310 | 0.00000303 | 1.0220 | 0.3068 |
| $\alpha_1$ | 0.0007105 | 0.00007366 | 9.6454 | 0.0000 |
| $\beta_{garch}$ | 0.9124666 | 0.0561699 | 16.2448 | 0.0000 |
| **Logit** | | | | |
| $\omega_{logit}$ | -4.6838070 | 0.0551703 | -84.8973 | 0.0000 |
| $\delta_1$ | -0.0909219 | 0.0560380 | -1.6225 | 0.1047 |
| $\delta_2$ | -0.4906196 | 0.0312001 | -15.7250 | 0.0000 |
| $\beta_1$ | 6.0108094 | 0.1421027 | 42.2990 | 0.0000 |
| $\beta_2$ | 0.2303069 | 0.1715186 | 1.3428 | 0.1794 |
| $\beta_3$ | 0.1965691 | 0.0615395 | 3.1942 | 0.0014 |
| $\beta_4$ | 0.3690088 | 0.2202590 | 1.6753 | 0.0939 |
| $\beta_5$ | 0.2490853 | 0.0638956 | 3.8983 | 0.0001 |
| $\beta_6$ | 0.4734219 | 0.2496743 | 1.8962 | 0.0580 |
| $\beta_\rho$ | 0.1584643 | 0.0374902 | 4.2268 | 0.0000 |
| $\nu_t$ | 2.0000023 | 0.0499164 | 40.0670 | 0.0000 |

Table 5.10: Model 4

### Out-of-sample performance

We tested the models on their out-of-sample performance using the remaining 20% of the data. While the decomposition models exhibit a reasonable ability to capture volatility, they consistently struggle with accurately predicting the direction of returns. Notably, the models tend to overestimate positive returns, which suggests a bias in the directional component.

One possible explanation for this bias is that the models, while designed to estimate conditional means, might be overly influenced by the unconditional mean, which is positive. This could lead to an overemphasis on predicting positive returns. Mostly because of this, out-of-sample, the decomposition model performs worse than the benchmark model.

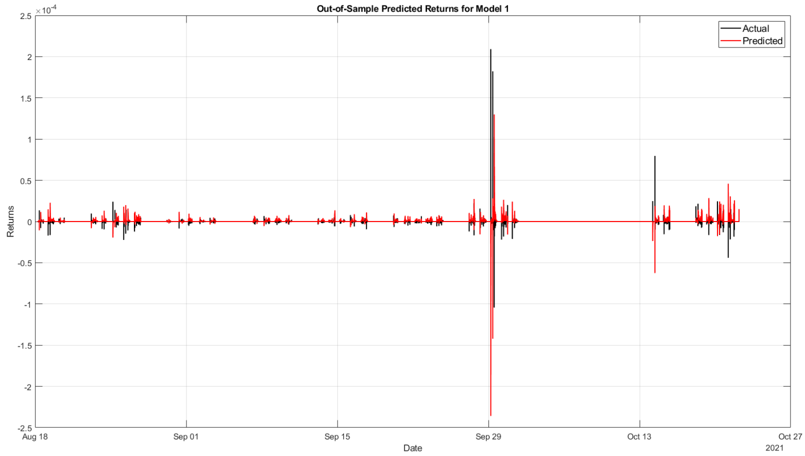


RMSE: 0.0912, Accuracy: 18.95%

FIGURE 5.7: Out-of-sample Model 1: $r_t$ and $E_{t-1}(r_t)$

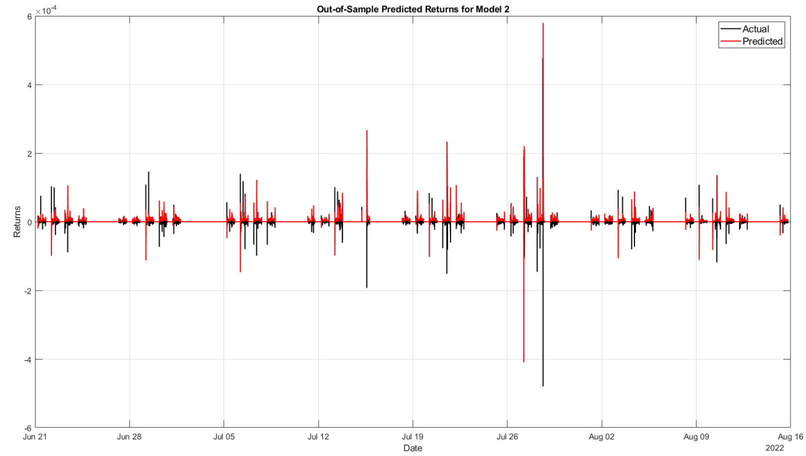


RMSE: 0.2746, Accuracy: 40.90%

FIGURE 5.8: Out-of-sample Model 2: $r_t$ and $E_{t-1}(r_t)$

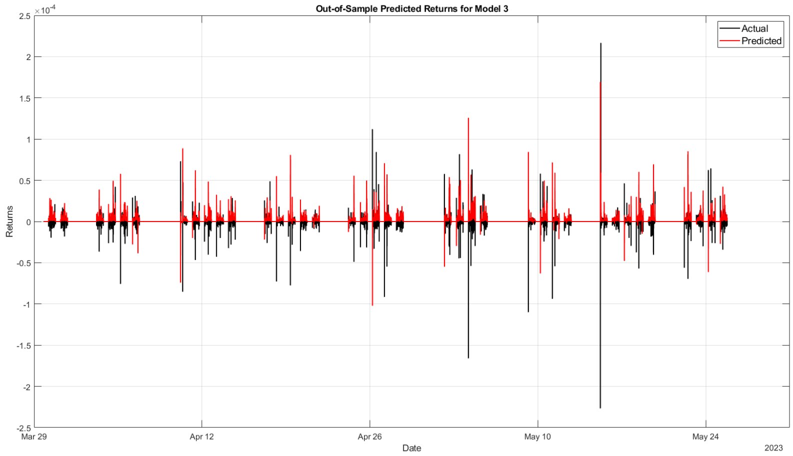


RMSE: 0.1415, Accuracy: 44.39%

FIGURE 5.9: Out-of-sample Model 3: $r_t$ and $E_{t-1}(r_t)$

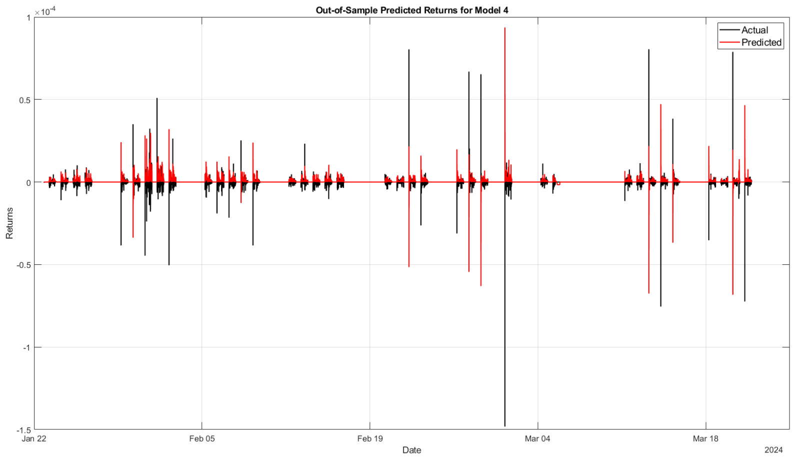


RMSE: 0.0627, Accuracy: 44.96%

FIGURE 5.10: Out-of-sample Model 4: $r_t$ and $E_{t-1}(r_t)$

# Chapter 6

# Conclusion

Given the findings in Chapter 3, we initially hypothesized that seasonality might be present in the 3-month EUR/USD forward points' returns. Our tests confirmed this hypothesis, revealing distinct variances across different trading periods, and suggested that returns during the mid-morning session were significantly higher than in other sessions. To model the conditional mean of returns, we employed a decomposition approach. As a benchmark, we also constructed a model following the Box-Jenkins methodology. However, the benchmark model was not entirely well-specified; it failed to capture the seasonality in returns, suggesting instead that returns did not exhibit any significant seasonal patterns. The linear ARIMA representation used for the conditional mean struggled to properly reflect the dynamics of the returns.

Chapter 3 also highlighted that both absolute returns and the sign of returns exhibit intraday patterns. These observations, coupled with previous literature[1], suggest that while returns themselves are challenging to predict due to their directionality, their components—magnitude and sign—might be more predictable. Initially, our decomposition of returns did not reveal any persistent seasonality in the EUR/USD FX 3m-swap market. However, upon considering the potential for changing return dynamics over time, the four models estimated on different parts of the dataset indicated the presence of varying seasonal patterns. Notably, all models except for the second suggested lower volatility during the mid-morning session, aligning with the findings discussed in Chapter 3 and Appendix **??**. These results imply that seasonal patterns in the data may shift over time, as further discussed in Chapter 5.

Despite these insights, we must acknowledge potential flaws in our findings. While the ARIMA-GARCH model, employing the folded Student's t-distribution, performed reasonably well in capturing the volatility dynamics of returns in out-of-sample predictions, the binary choice model used to predict the direction of returns performed poorly. Although the estimated coefficients

[1] See Chapter 4 for empirical motivation.

were statistically significant, the model's accuracy was low due to a bias towards predicting positive returns. This underlines the inherent difficulty in forecasting return direction—a longstanding challenge in financial modeling.

The decomposition model presents several attractive features, notably its ability to separate a particularly challenging component to predict (returns) into two potentially more predictable components (magnitude and sign). By using copula methodology to model the dependence structure between sign and magnitude, we found that this approach was preferable to assuming independence. Despite the challenges in accurately predicting return direction, the decomposition model offers a promising framework for improving predictive performance. It allows us to refine the models for the sign and magnitude components separately and then incorporate their interdependencies using copula theory. This approach, though not without its difficulties, offers a more nuanced and potentially more effective means of modeling financial returns compared to traditional linear time series models.

In conclusion, while our results suggest that some degree of seasonality exists in the volatility of returns, the evidence for intraday seasonality in the sign of returns remains inconclusive. We do not firmly reject the hypothesis of intraday seasonality in return sign, but we are cautious in claiming its presence, particularly given the poor out-of-sample performance and the very small magnitudes of the coefficients attached to seasonal dummies. These small coefficients imply that any potential profits from exploiting these inefficiencies would likely be minimal.

# Chapter 7

# Discussion

In this thesis, we explored intraday seasonality in the FX Swap market, specifically focusing on EUR/USD three-month forward points. While the use of dummy variables provided a straightforward method for modeling seasonal patterns, an alternative approach could involve proxying the seasonal component using the method described by Andersen and Bollerslev (1997). Their approach involves decomposing the returns into a predictable component and a stochastic component, where the predictable component captures intraday seasonality through Fourier series or similar periodic functions. This method could potentially offer a more refined capture of intraday seasonality compared to the simpler dummy variable approach.

Another consideration is the complexity of optimizing the highly nonlinear likelihood function used in our models. Bayesian inference offers a flexible approach to addressing this challenge, particularly through the use of Markov Chain Monte Carlo (MCMC) methods. As demonstrated by Liu and Luger (2015), MCMC can be effectively employed to estimate the decomposition model, even when dealing with non-differentiable likelihood functions arising from absolute value and indicator functions. The adaptive MCMC algorithm they utilized enhances the mixing of the Markov chain, allowing for more robust parameter estimation in complex, non-linear models.

Furthermore, the copula-based approach to modeling dependencies between return components proved useful in this analysis. However, exploring different types of copulas or even dynamic copula models might provide additional insights into the evolving dependencies between the sign and magnitude of returns. This could be particularly valuable in capturing the time-varying nature of relationships in financial markets.

Finally, the granularity of the data used in this analysis, while providing detailed insights, might also introduce noise that obscures broader seasonal patterns. Future research could consider using less granular data, such as hourly or daily observations, to determine whether intraday seasonality can still be detected. This approach might strike a balance between capturing significant patterns and reducing the noise associated with high-frequency data.

# Appendix A

# Distribution

## A.1 Returns

### A.1.1 Plots

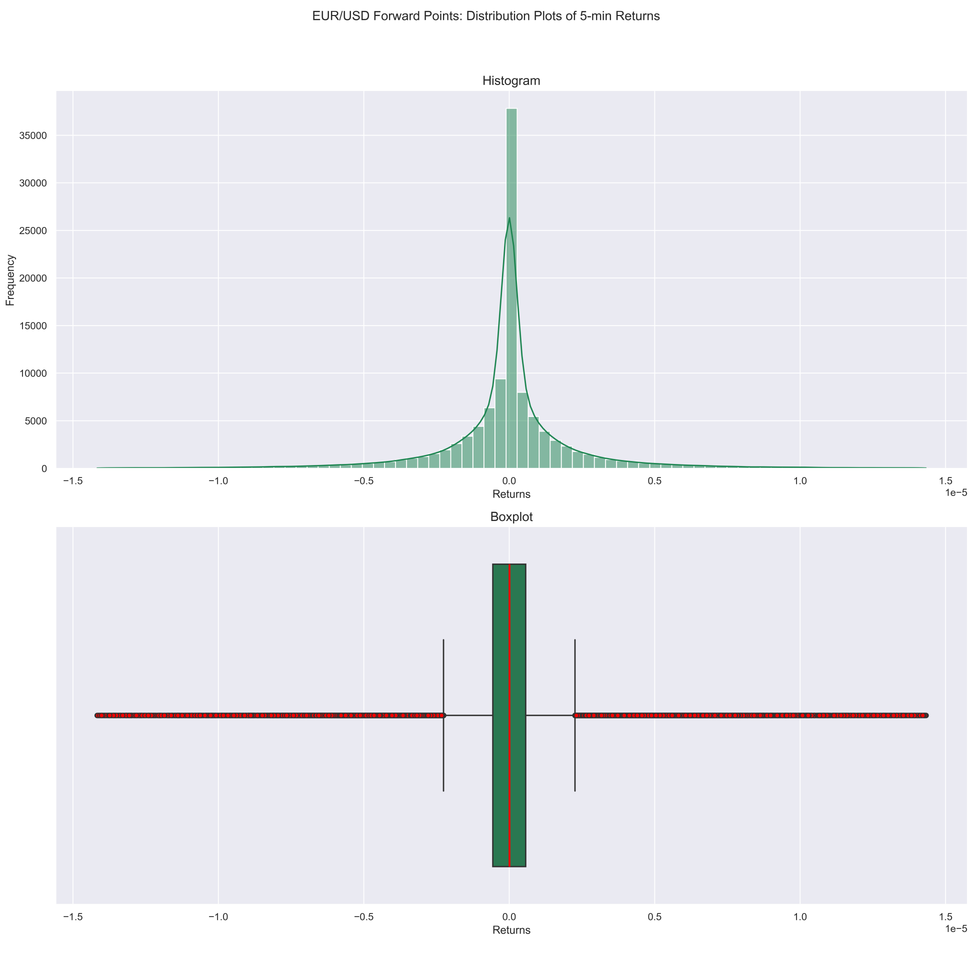


FIGURE A.1: EUR/USD Forward Points: Histogram (top) and Boxplot (bottom) Returns

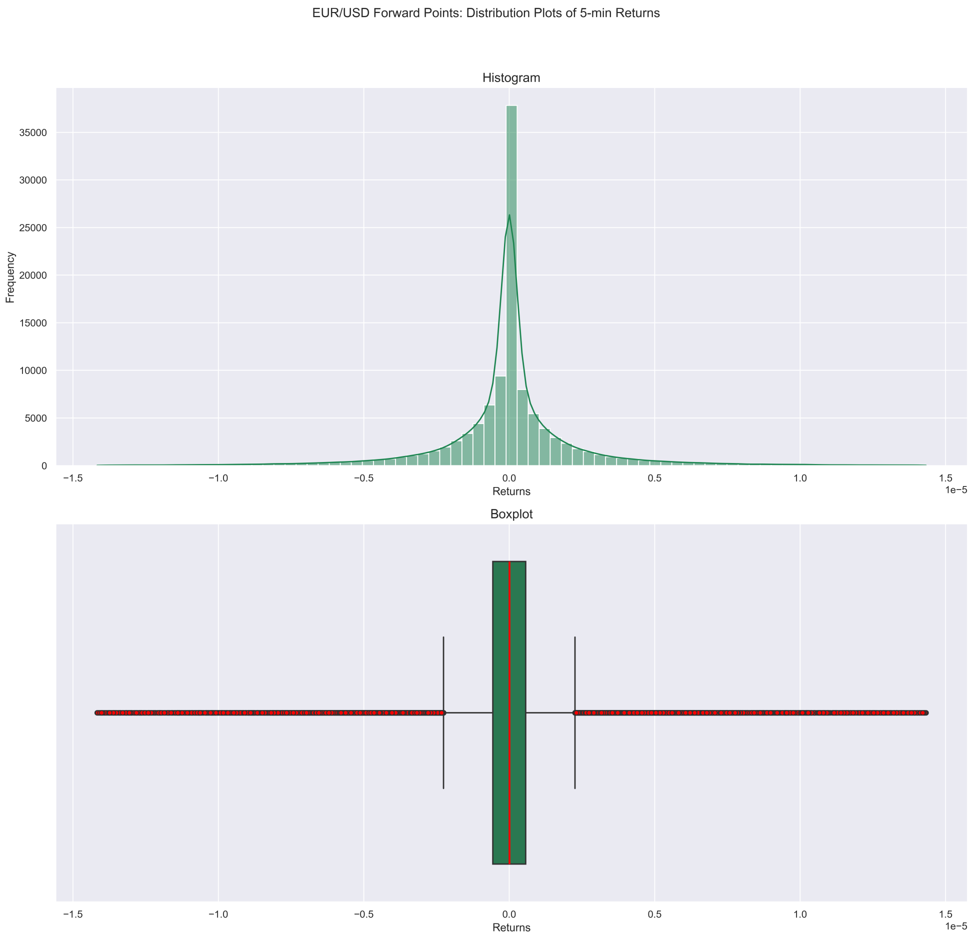


Figure A.2: EUR/USD Forward Points: Scatterplot up to one hour

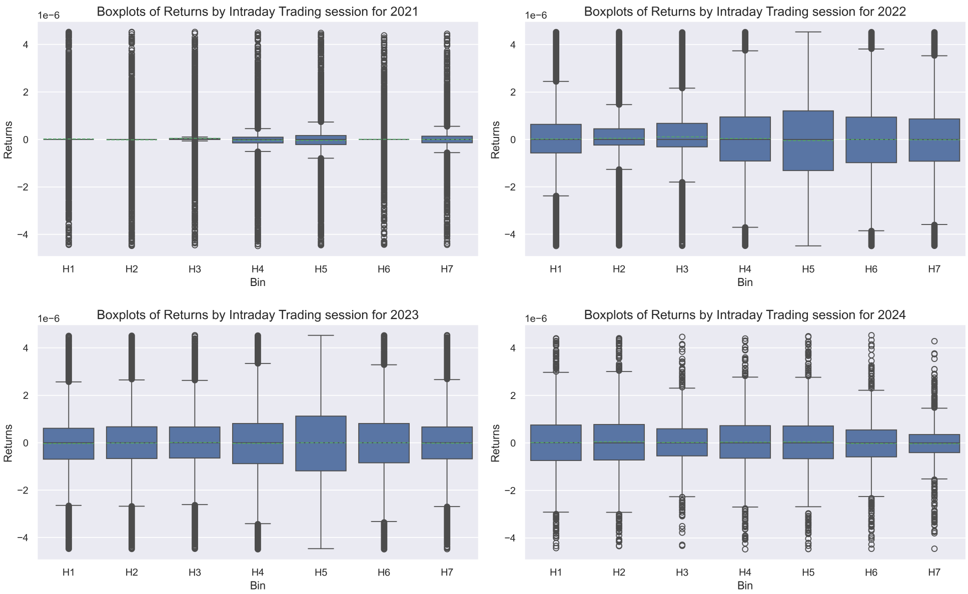


FIGURE A.3: EUR/USD Forward Points: Boxplot per Trading Session per Year

### A.1.2 Levene's Test for Equality of Variances

Levene's test is used to assess the equality of variances for a variable calculated for two or more groups. In this context, we compare the variances of returns during specific intraday trading sessions (H1 to H7) against the variances during all other sessions.

$$H_0 : \sigma_1^2 = \sigma_2^2 = \ldots = \sigma_7^2$$
$$H_1 : \sigma_p^2 \neq \sigma_q^2 \text{ for at least one pair } \{p, q\}$$

#### Test Statistic

Levene's test statistic is calculated as follows:

$$W = \frac{(N-k)}{(k-1)} \frac{\sum_{i=1}^{k} N_i (Z_{i\cdot} - Z_{\cdot\cdot})^2}{\sum_{i=1}^{k} \sum_{j=1}^{N_i} (Z_{ij} - Z_{i\cdot})^2}$$

where:

- $Z_{ij} = |X_{ij} - \bar{X}_i|$
- $Z_{i\cdot}$ is the mean of the $Z_{ij}$
- $Z_{\cdot\cdot}$ is the overall mean of the $Z_{ij}$

#### Results

The results of the Levene's test for the intraday trading sessions are summarized in the table below:

Table A.1: Levene's Test Results for Intraday Trading Sessions

| **Session** | **Levene Statistic** | **Levene p-value** ($\alpha = 5\%$) |
|---|---|---|
| H1 | 6.3536 | 0.0117 |
| H2 | 83.9325 | 0.0000 |
| H3 | 28.0967 | 0.0000 |
| H4 | 6.1986 | 0.0128 |
| H5 | 14.5771 | 0.0001 |
| H6 | 17.0438 | 0.0000 |
| H7 | 17.6373 | 0.0000 |

A significant p-value ($< 0.05$) indicates a statistically significant difference in variances of returns for the respective trading session.

### A.1.3 Welch's t-test

Based on Levene's test we conclude unequal variances. To ensure robustness we perform Welch's t-test to determine if the means of two independent samples are significantly different from each other.

$$H_0 : \mu_1 = \mu_2$$
$$H_1 : \mu_1 \neq \mu_2$$

**Test statistic**

The t-test statistic for unequal variances is calculated as follows:

$$t = \frac{\bar{X}_1 - \bar{X}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$$

, where $s_1^2, s_2^2$ are sample analogues for the variances.

**Results**

The results of Welch's t-test for the intraday trading sessions are summarized in the table below:

Table A.2: T-Test Results for Intraday Trading Sessions

| **Session** | **T-test Type** | **t-statistic** | **p-value** ($\alpha = 5\%$) |
|---|---|---|---|
| H1 | Welch's t-test | 0.4922 | 0.6226 |
| H2 | Welch's t-test | -0.1760 | 0.8603 |
| H3 | Welch's t-test | 3.3888 | 0.0007 |
| H4 | Welch's t-test | 0.3145 | 0.7531 |
| H5 | Welch's t-test | -1.2167 | 0.2237 |
| H6 | Welch's t-test | -1.2393 | 0.2152 |
| H7 | Welch's t-test | -2.5581 | 0.0105 |

A significant p-value (< 0.05) indicates a statistically significant difference in mean returns for the respective trading session.

### A.1.4 Monday Morning Effect: OLS Results

Table A.3: OLS Regression Results for Monday Morning Effect

| Variable | Coefficient | Standard Error | p-value |
|---|---|---|---|
| Constant | 1.743e-08 | 2.84e-08 | 0.539 |
| MM | 1.813e-07 | 1.04e-07 | 0.080 |

### A.1.5 London & NY Open Effect: OLS Results

Table A.4: OLS Regression Results for London and NY Opens

| Variable | Coefficient | Standard Error | p-value |
|---|---|---|---|
| Constant | 4.261e-08 | 3.69e-08 | 0.248 |
| LondonOpen | 4.867e-08 | 5.76e-08 | 0.398 |
| NYOpen | -1.183e-07 | 7.77e-08 | 0.128 |

## A.2 Absolute Returns

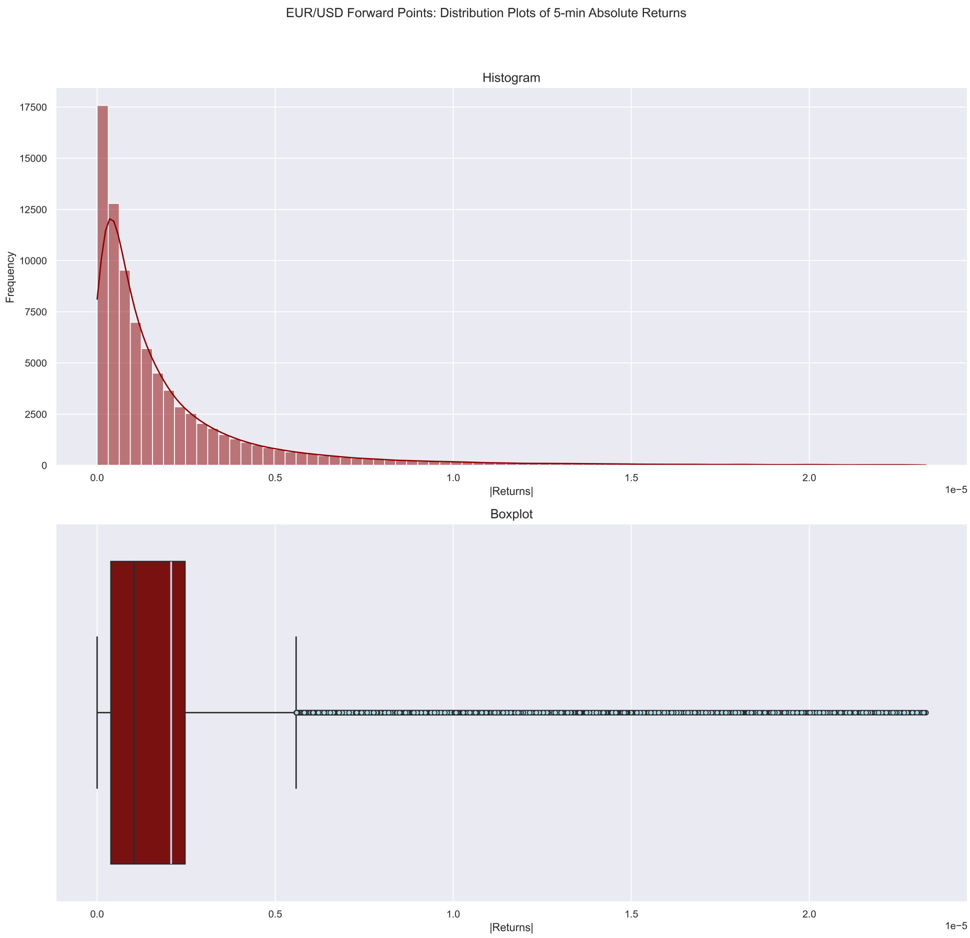


FIGURE A.4: EUR/USD Forward Points: Histogram (top) and Boxplot (bottom) Absolute Returns

## A.3 Explantory Variables

### A.3.1 3-Month Swap Rates

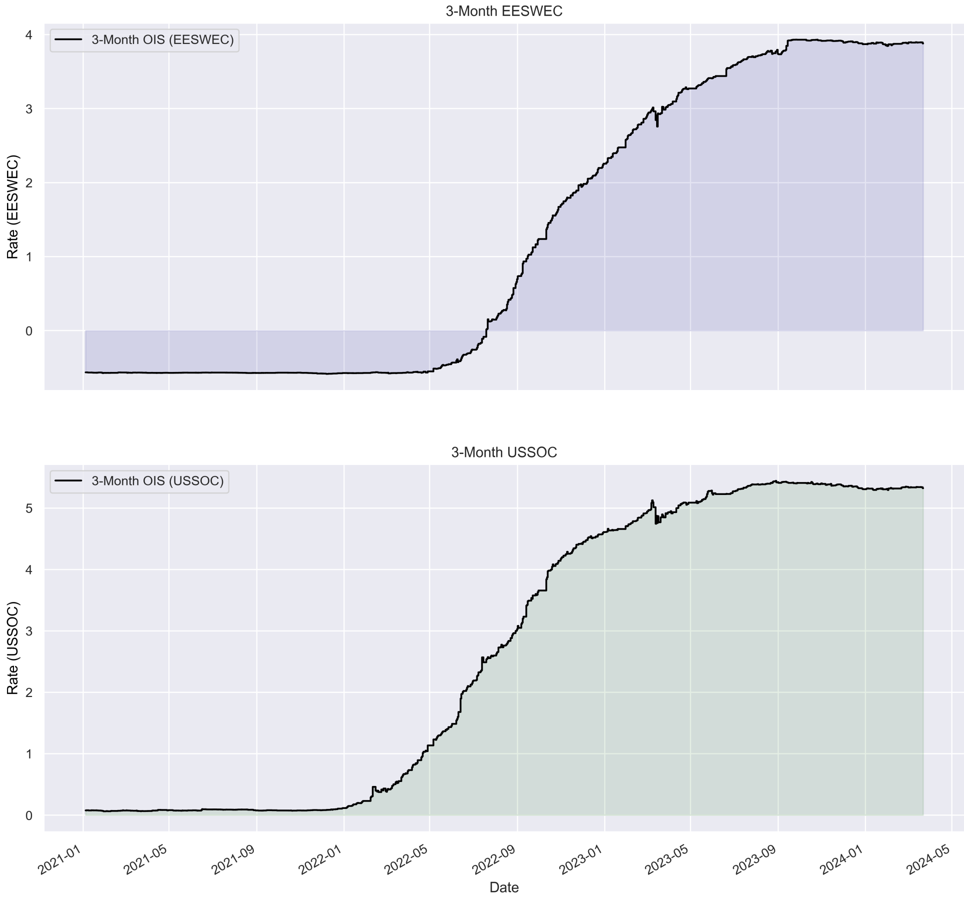


Figure A.5: 3-Month EUR interest rate (top) and 3-Month dollar interest rate (bottom)

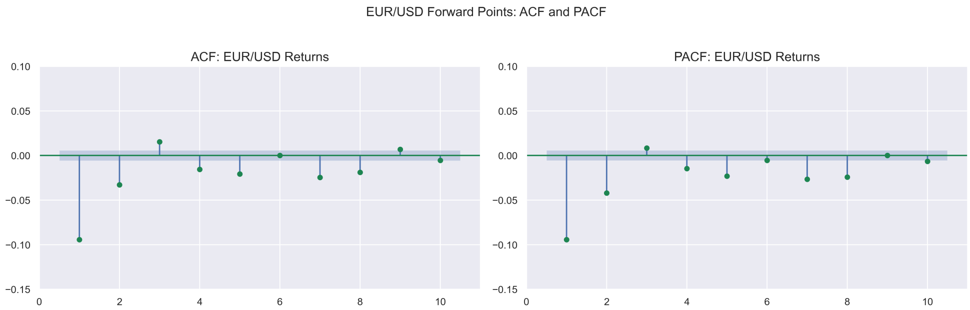


Figure A.6: EUR/USD Forward Points: ACF and PACF Returns (10 lags)

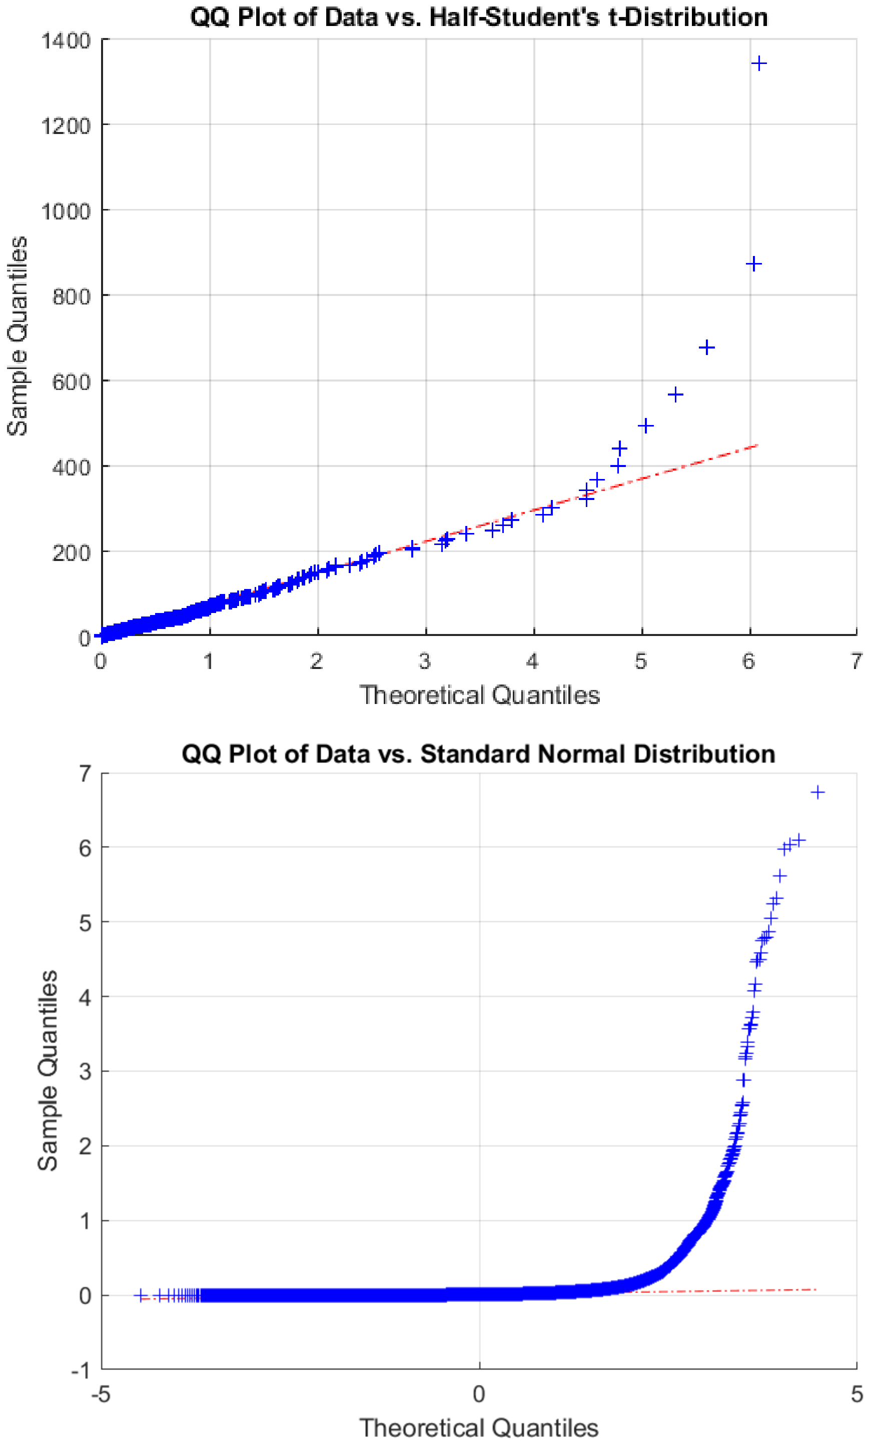


FIGURE A.7: Comparison of Sample Quantiles vs. Folded Student-t Quantiles (top) and Standard Normal Quantiles (bottom)

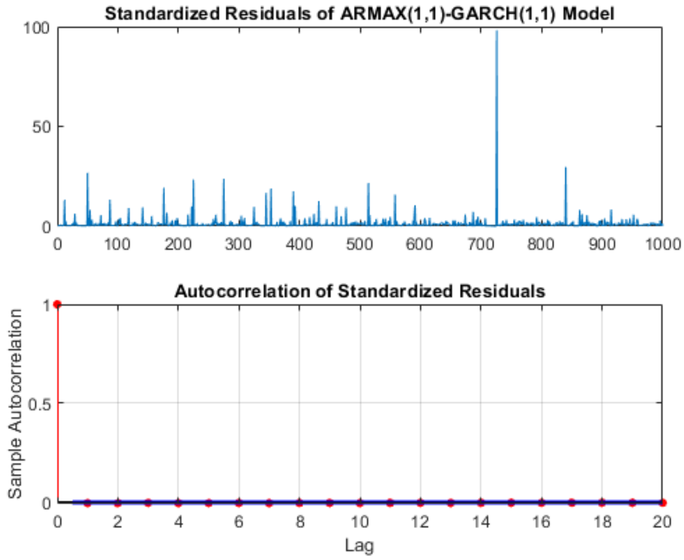


FIGURE A.8: Standardized Residuals ARMAX(1,1)-GARCH

# Appendix B

# Diagnostic Tests

### Augmented Dickey Fueller Test

The ADF test is used to determine whether a time series is stationary, specifically testing for the presence of a unit root. The null hypothesis states that the time series has a unit root (not stationary). The alternative hypothesis states that the time series is stationary. The test is performed as follows:

1. **Specify the Test Equation:** The ADF test is based on estimating the following regression model:

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} + \sum_{i=1}^{p} \delta_i \Delta y_{t-i} + \epsilon_t$$

   where $\Delta y_t$ is the first difference of the series, $\alpha$ is the constant (drift), $\beta t$ is the time trend, $y_{t-1}$ is the lagged level of the series, and $\sum_{i=1}^{p} \delta_i \Delta y_{t-i}$ represents the lagged differences to account for serial correlation.

2. **Estimate the Parameters:** Fit the above regression model to the data and estimate the parameter $\gamma$.

3. **Test Statistic:** The test statistic is the t-statistic for the coefficient $\gamma$:

$$\tau = \frac{\hat{\gamma}}{\mathrm{SE}(\hat{\gamma})}$$

   where $\hat{\gamma}$ is the estimated coefficient and $\mathrm{SE}(\hat{\gamma})$ is its standard error.

4. **Evaluate the Test Statistic:** Calculate the p-value associated with the test statistic $\tau$. The null hypothesis of a unit root is rejected if the p-value is less $\alpha = 0.05$

**Conclusion:** If the null hypothesis is rejected, it indicates that the time series is stationary. If the null hypothesis is not rejected, it suggests that the series has a unit root and is non-stationary, potentially requiring differencing to achieve stationarity.

### Test for ARCH Effects

The ARCH test is used to detect autocorrelation in the squared residuals of a time series model. The test is performed as follows:

1. **Fit the Model:** Estimate the appropriate time series model and obtain the residuals $\hat{\epsilon}_t$.

2. **Compute the Squared Residuals:** Calculate the squared residuals $\hat{\epsilon}_t^2$.

3. **Perform the ARCH-LM Test:** Regress the squared residuals on a constant and $q$ lagged values of $\hat{\epsilon}_t^2$:

$$\hat{\epsilon}_t^2 = \alpha_0 + \alpha_1 \hat{\epsilon}_{t-1}^2 + \alpha_2 \hat{\epsilon}_{t-2}^2 + \cdots + \alpha_q \hat{\epsilon}_{t-q}^2 + u_t$$

4. **Test Statistic:** The Lagrange Multiplier (LM) test statistic is calculated as:

$$\text{LM} = n \times R^2$$

   where $n$ is the number of observations and $R^2$ is the coefficient of determination from the regression in 3..

5. **Evaluate the Test Statistic:** The LM statistic follows a chi-square distribution with $q$ degrees of freedom. The null hypothesis ($H_0$) of no ARCH effects is rejected if the p-value is below $\alpha = 0.05$.

**Conclusion:** If the null hypothesis is rejected, it indicates the presence of ARCH effects, suggesting that the variance of the residuals is not constant and may be modeled more effectively with a GARCH model.

# Appendix C

# Benchmark Model

### Model 1

| Parameter | Value | Standard Error | T-Statistic | P-Value |
|---|---|---|---|---|
| **ARIMAX Conditional Mean Model** | | | | |
| $\phi_1$ | -0.32855 | 0.001138 | -288.7 | 0 |
| $\phi_2$ | 0.0026469 | 0.0011458 | 2.3101 | 0.020884 |
| $\phi_1$ | 0.28901 | 0.0011536 | 250.53 | 0 |
| $\phi_2$ | -0.020058 | 0.0011733 | -17.095 | 1.624e-65 |
| $\phi_3$ | -0.0035867 | 0.0020605 | -1.7407 | 0.081736 |
| $\phi_4$ | 0.00029402 | 0.0018981 | 0.1549 | 0.8769 |
| $\beta_1$ | 9.1637e-06 | 5.4697e-05 | 0.16754 | 0.86695 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 2e-07 | 3.4044e-08 | 5.8748 | 4.2343e-09 |
| $\alpha$ | 0.6821 | 0.0012969 | 525.96 | 0 |
| $\beta$ | 0.3179 | 0.0051937 | 61.21 | 0 |
| $\nu$ | 2.4097 | 0.008374 | 287.76 | 0 |
| AIC | -524152 | | | |

Table C.1: Estimation Results for ARIMAX-GARCH(1,1) Model with t Distribution

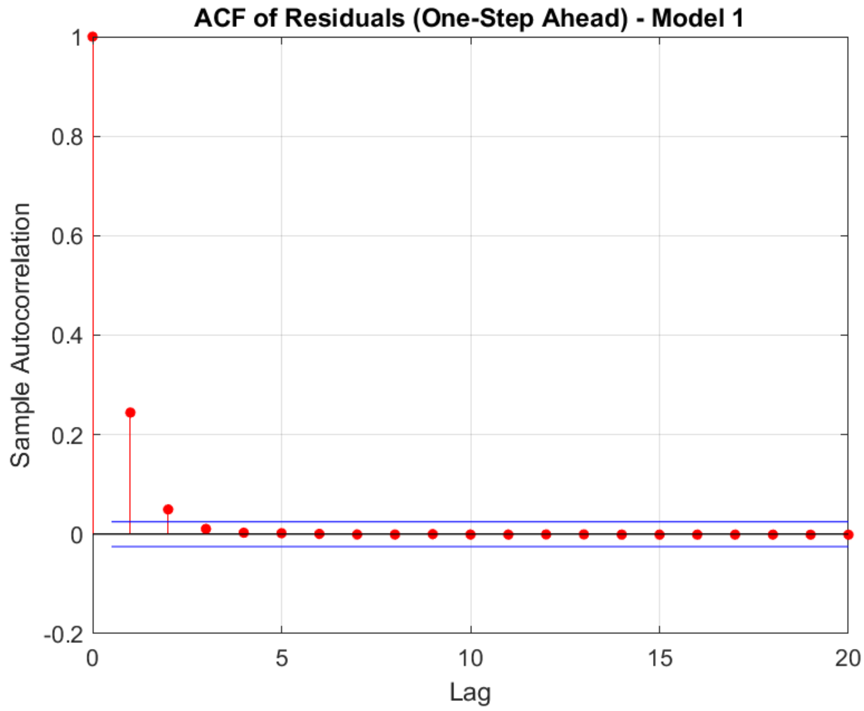


Figure C.1: Out-of-sample Benchmark Model 1: $r_t$ and $E_{t-1}(r_t)$

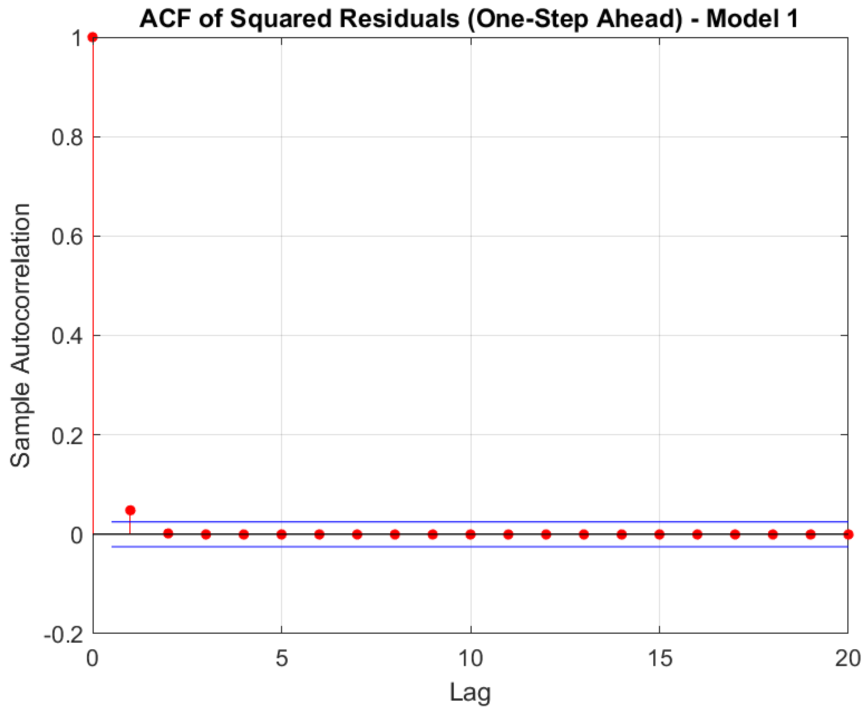


Figure C.2: Out-of-sample Benchmark Model 1: $r_t$ and $E_{t-1}(r_t)$

### Model 2

| Parameter | Value | Standard Error | T-Statistic | P-Value |
|---|---|---|---|---|
| **ARIMAX Conditional Mean Model** | | | | |
| $\phi_1$ | 0.76889 | 0.23518 | 3.2694 | 0.0010779 |
| $\phi_2$ | -0.71666 | 0.17409 | -4.1165 | 3.8463e-05 |
| $\theta_1$ | -0.8093 | 0.23525 | -3.4401 | 0.00058145 |
| $\theta_2$ | 0.74641 | 0.1781 | 4.1909 | 2.778e-05 |
| $\theta_3$ | -0.024664 | 0.010229 | -2.4112 | 0.0159 |
| $\theta_4$ | -0.0017592 | 0.0045488 | -0.38673 | 0.69895 |
| $\beta_1$ | 5.2487e-06 | 9.6327e-05 | 0.054488 | 0.95655 |
| $\nu$ | 2.4247 | 0.016353 | 148.28 | 0 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 2e-07 | 7.0282e-08 | 2.8457 | 0.0044316 |
| $\alpha$ | 0.64873 | 0.0025935 | 250.13 | 0 |
| $\beta$ | 0.35127 | 0.010884 | 32.275 | 1.5654e-228 |
| $\nu$ | 2.4247 | 0.016353 | 148.28 | 0 |
| AIC | -123169 | | | |

TABLE C.2: Estimation Results for ARIMAX-GARCH(1,1) Model with t Distribution

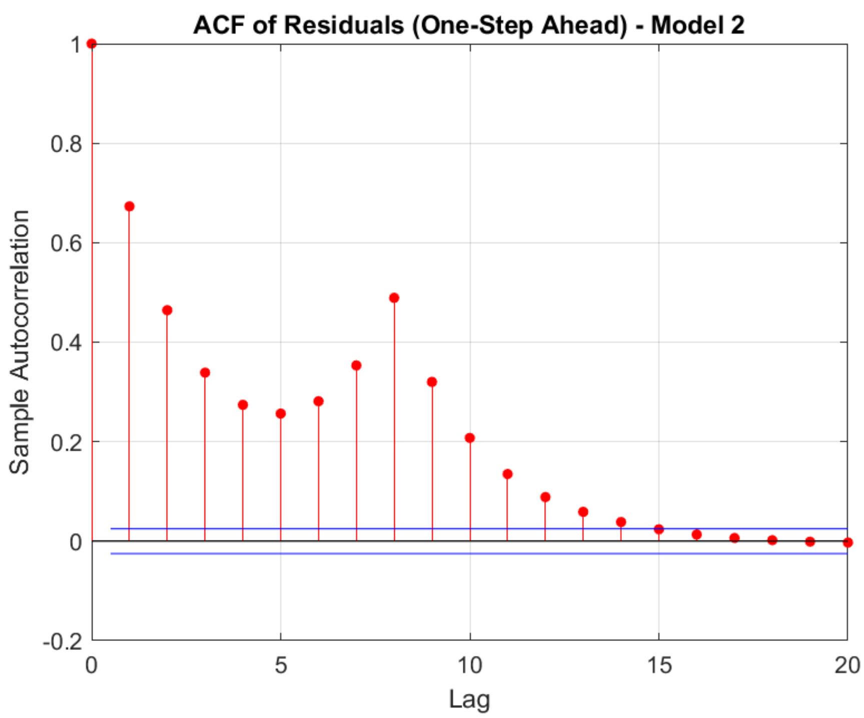


FIGURE C.3: Out-of-sample Benchmark Model 2: ACF

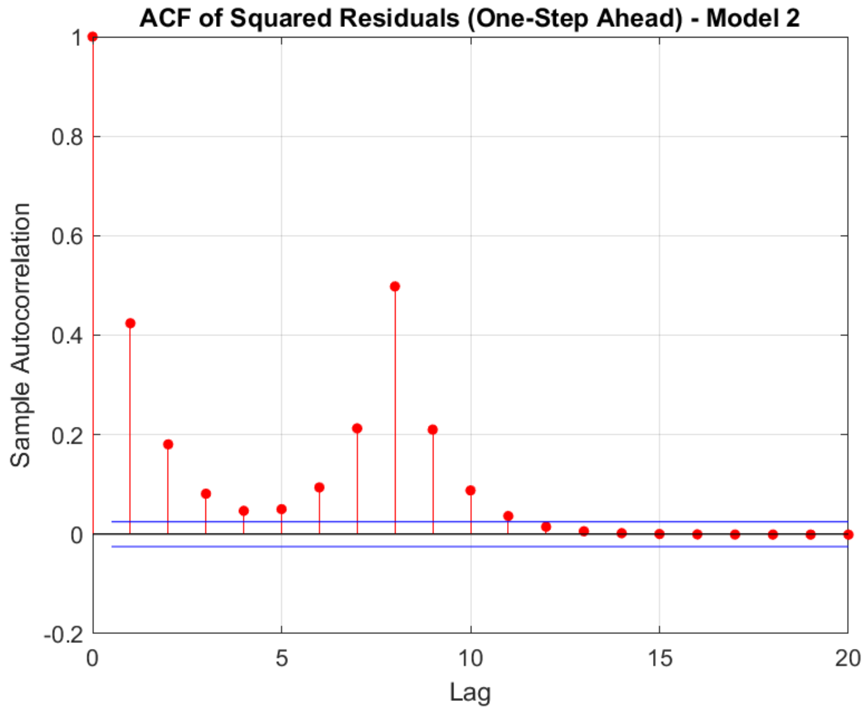


FIGURE C.4: Out-of-sample Benchmark Model 2: ACF

## Model 3

| Parameter | Value | Standard Error | T-Statistic | P-Value |
|---|---|---|---|---|
| **ARIMAX Conditional Mean Model** | | | | |
| $\phi_1$ | 0.60161 | 0.12656 | 4.7534 | 2.0006e-06 |
| $\phi_2$ | -0.79703 | 0.11674 | -6.8274 | 8.6491e-12 |
| $\theta_1$ | -0.65748 | 0.12669 | -5.1899 | 2.1044e-07 |
| $\theta_1$ | 0.82267 | 0.11982 | 6.8658 | 6.6138e-12 |
| $\theta_3$ | -0.039525 | 0.0086016 | -4.5951 | 4.326e-06 |
| $\theta_4$ | -0.0042925 | 0.0043488 | -0.98706 | 0.32362 |
| $\beta$ | 0.00091348 | 0.00045065 | 2.027 | 0.042661 |
| $\nu$ | 2.2939 | 0.023624 | 97.097 | 0 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 0.0012719 | 8.393e-05 | 15.155 | 7.038e-52 |
| $\alpha$ | 0.3034 | 0.0085468 | 35.499 | 5.1052e-276 |
| $\beta$ | 0.6966 | 0.050218 | 13.871 | 9.4431e-44 |
| $\nu$ | 2.2939 | 0.023624 | 97.097 | 0 |
| AIC | -89649.3 | | | |

TABLE C.3: Estimation Results for ARIMAX-GARCH(1,1) Model with t Distribution

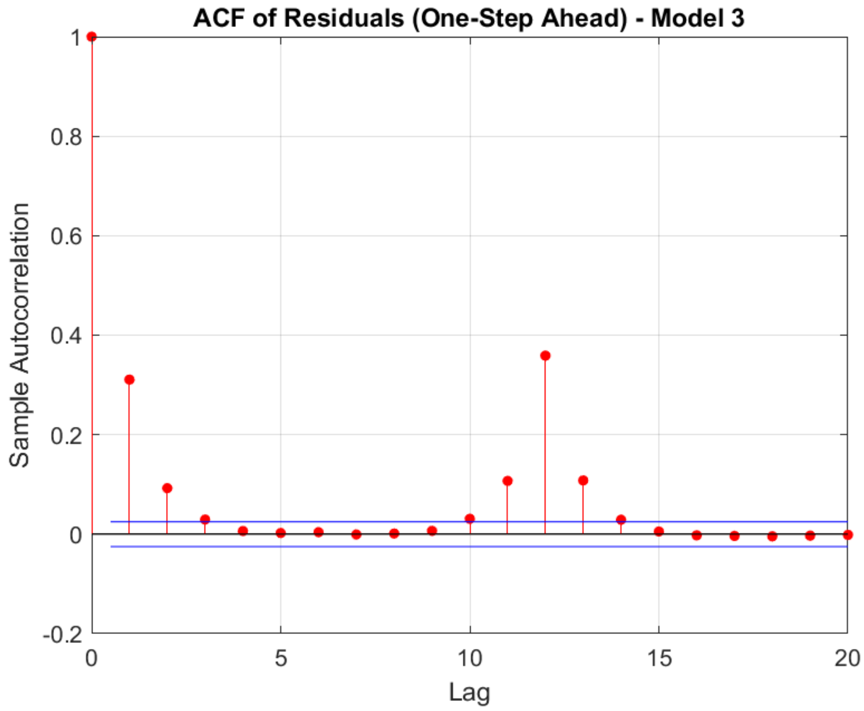


RMSE: 0.044548

FIGURE C.5: Out-of-sample Benchmark Model 3: $r_t$ and $E_{t-1}(r_t)$

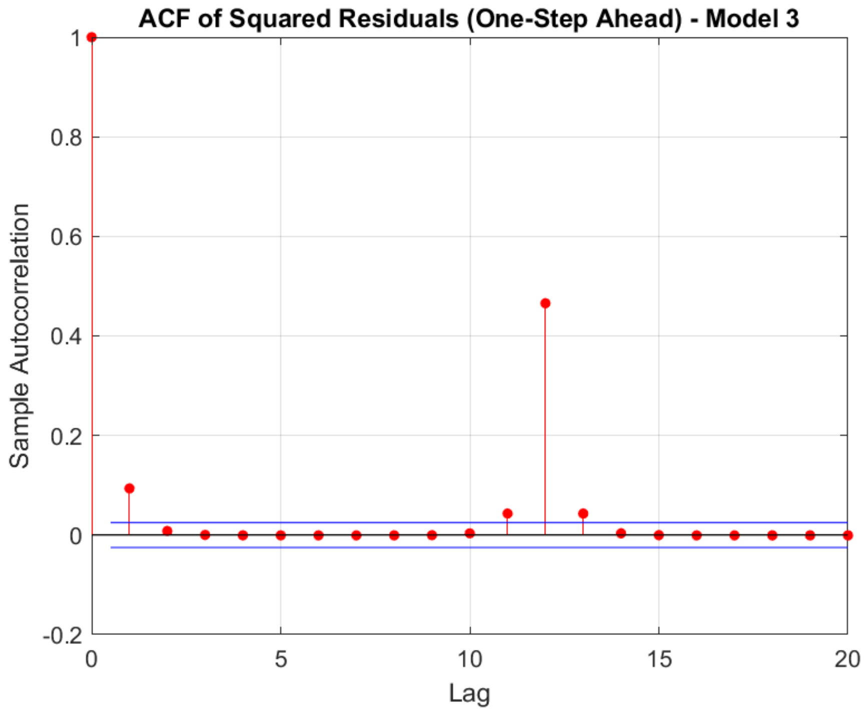


RMSE: 0.13355

FIGURE C.6: Out-of-sample Benchmark Model 3: $r_t$ and $E_{t-1}(r_t)$

### Model 4

| Parameter | Value | Standard Error | T-Statistic | P-Value |
|---|---|---|---|---|
| **ARIMAX Conditional Mean Model** | | | | |
| $\phi_1$ | 0.1518 | 1.0329 | 0.14696 | 0.88316 |
| $\phi_2$ | 0.0017897 | 0.47204 | 0.0037914 | 0.99697 |
| $\theta_1$ | -0.21061 | 1.0329 | -0.20391 | 0.83843 |
| $\theta_2$ | 0.0065901 | 0.52729 | 0.012498 | 0.99003 |
| $\theta_3$ | 0.0072141 | 0.027663 | 0.26079 | 0.79426 |
| $\theta_3$ | 0.0038703 | 0.0087997 | 0.43982 | 0.66007 |
| $\beta$ | -5.3596e-05 | 0.00017444 | -0.30725 | 0.75865 |
| **GARCH(1,1) Conditional Variance Model (t Distribution)** | | | | |
| $\omega$ | 0.00046441 | 4.0728e-05 | 11.403 | 4.0442e-30 |
| $\alpha$ | 0.21033 | 0.010554 | 19.928 | 2.3195e-88 |
| $\beta$ | 0.78967 | 0.076288 | 10.351 | 4.1309e-25 |
| $\nu$ | 2.2594 | 0.027222 | 83.001 | 0 |
| AIC | -132717 | | | |

TABLE C.4: Estimation Results for ARIMAX-GARCH(1,1) Model with t Distribution

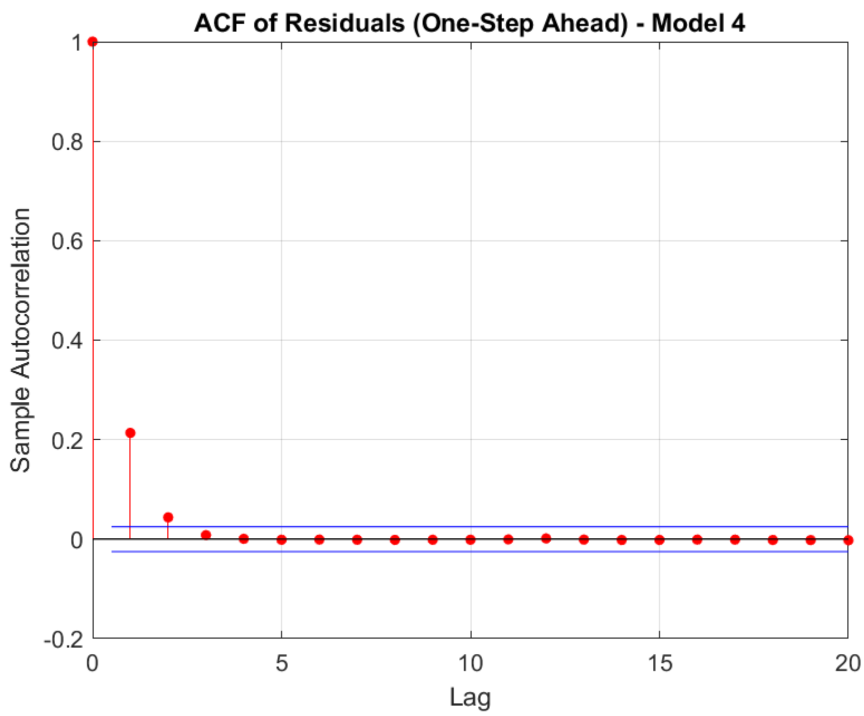


RMSE: 0.044548

FIGURE C.7: Out-of-sample Benchmark Model 4: $r_t$ and $E_{t-1}(r_t)$

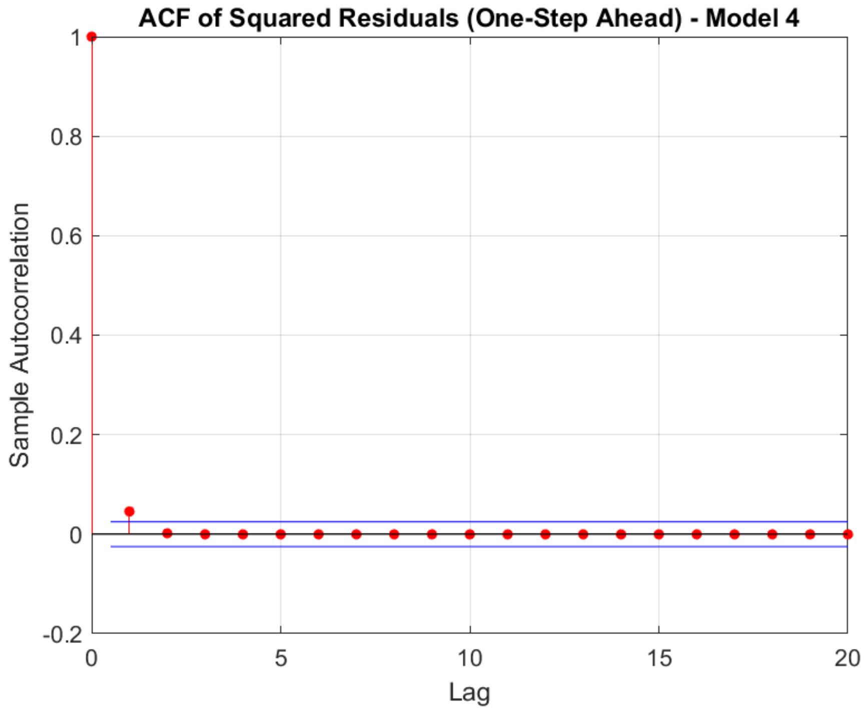


RMSE: 0.13355

FIGURE C.8: Out-of-sample Benchmark Model 4: $r_t$ and $E_{t-1}(r_t)$

# Appendix D

# Copula

$$\varrho_t(z) = 1 - \frac{\partial C(z, 1-p_t)}{\partial F_1(x_1)}$$
$$\frac{\partial C(z, 1-p_t)}{\partial F_1(x_1)} = \frac{\partial C(z, 1-p_t)}{\partial F_1(x_1)} \tag{D.1}$$

The Farlie-Gumbel-Morgenstern (FGM) copula is defined as:

$$C(u, v) = uv\,[1 + \theta(1-u)(1-v)]$$

We want to find the partial derivative $\frac{\partial C(u,v)}{\partial u}$.

$$\frac{\partial C(u, v)}{\partial u} = \frac{\partial}{\partial u}\,[uv\,(1 + \theta(1-u)(1-v))]$$

Expanding this:

$$\frac{\partial C(u, v)}{\partial u} = v\,[1 + \theta(1-u)(1-v)] + uv \cdot \frac{\partial}{\partial u}\,[1 + \theta(1-u)(1-v)]$$

$$\frac{\partial C(u, v)}{\partial u} = v\,[1 + \theta(1-u)(1-v)] + uv \cdot [-\theta(1-v)]$$

Simplifying the expression:

$$\frac{\partial C(u, v)}{\partial u} = v\,[1 + \theta(1-u)(1-v)] - uv\theta(1-v)$$

$$\frac{\partial C(u, v)}{\partial u} = v\,[1 + \theta(1-v) - \theta u(1-v)]$$

$$\frac{\partial C(u, v)}{\partial u} = v\,[1 + \theta(1-v)(1-u)]$$

Now, we evaluate the partial derivative at $(z, 1-p_t)$:

$$\frac{\partial C(z, 1-p_t)}{\partial z} = (1-p_t)\left[1+\theta(1-p_t)(1-z)\right]$$

Expanding and simplifying:

$$\frac{\partial C(z, 1-p_t)}{\partial z} = (1-p_t) + \theta(1-p_t)^2 - \theta(1-p_t)^2 z$$

Finally:

$$\frac{\partial C(z, 1-p_t)}{\partial z} = (1-p_t)\left[1+\theta p_t(1-2z)\right]$$

Then, $\varrho_t$ is given as:

$$1-(1-p_t)\left[1+\theta p_t(1-2z)\right]$$

# Bibliography

Akaike, H. (1973). Maximum likelihood identification of gaussian autoregressive moving average models. *Biometrika*, *60*(2), 255–265.

Alford, A. W., & Boatsman, J. R. (1995). Predicting long-term stock return volatility: Implications for accounting and valuation of equity derivatives. *Accounting Review*, 599–618.

Aloud, M., Fasli, M., Tsang, E., Dupuis, A., & Olsen, R. (2013). Stylized facts of trading activity in the high frequency fx market: An empirical study. *Journal of Finance and Investment Analysis*, *2*(4), 145–183.

Anatolyev, S., & Gospodinov, N. (2010). Modeling financial return dynamics via decomposition. *Journal of Business & Economic Statistics*, *28*(2), 232–245.

Andersen, T. G., & Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of empirical finance*, *4*(2-3), 115–158.

Andersen, T. G., & Bollerslev, T. (1998). Deutsche mark–dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *the Journal of Finance*, *53*(1), 219–265.

Baillie, R. T., & Bollerslev, T. (1991). Intra-day and inter-market volatility in foreign exchange rates. *The Review of Economic Studies*, *58*(3), 565–585.

Baillie, R. T., & Bollerslev, T. (2000). The forward premium anomaly is not as bad as you think. *Journal of International Money and Finance*, *19*(4), 471–488.

Bain, L. J., & Engelhardt, M. (1992). *Introduction to probability and mathematical statistics* (Vol. 4). Duxbury Press Belmont, CA.

Bandara, K., Hyndman, R. J., & Bergmeir, C. (2021). Mstl: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. *arXiv preprint arXiv:2107.13462*.

Bansal, R., & Dahlquist, M. (2000). The forward premium puzzle: Different tales from developed and emerging economies. *Journal of international Economics*, *51*(1), 115–144.

Bauwens, L., Laurent, S., & Rombouts, J. V. (2006). Multivariate garch models: A survey. *Journal of applied econometrics*, *21*(1), 79–109.

BIS. (2019). Triennial central bank survey: Foreign exchange turnover in april 2019. *Triennial Central Bank Survey*, 1–23.

BIS. (2022). Triennial central bank survey otc foreign exchange turnover in april 2022. https://www.bis.org/statistics/rpfx22_fx.pdf

Bjønnes, G. H., & Rime, D. (2005). Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics*, *75*(3), 571–605.

Bjønnes, G. H., Rime, D., & Solheim, H. O. A. (2005). Liquidity provision in the overnight foreign exchange market. *Journal of International Money and Finance*, *24*(2), 175–196.

Boateng, M., Omari-Sasu, A., Avuglah, R., & Frempong, N. (2022). A mixture of clayton, gumbel, and frank copulas: A complete dependence model. *Journal of Probability and Statistics*, *2022*(1), 1422394.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, *31*(3), 307–327.

Box, G. E., Jenkins, G. M., & Reinsel, G. (1970). Times series analysis forecasting and control. *BoxTime Series Analysis: Forecasting and Control Holden Day1970*.

Breen, W., Glosten, L. R., & Jagannathan, R. (1989). Economic significance of predictable variations in stock index returns. *The Journal of finance*, *44*(5), 1177–1189.

Cao, C. Q., & Tsay, R. S. (1992). Nonlinear time-series analysis of stock volatilities. *Journal of applied econometrics*, *7*(S1), S165–S185.

Challet, D., Marsili, M., & Zhang, Y.-C. (2001). Stylized facts of financial markets and market crashes in minority games. *Physica A: Statistical Mechanics and its Applications*, *294*(3-4), 514–524.

Chang, S.-C., Chen, S.-S., Chou, R. K., & Lin, Y.-H. (2008). Weather and intraday patterns in stock returns and trading activity. *Journal of Banking & Finance*, *32*(9), 1754–1766.

Chatziantoniou, I., Gabauer, D., & Stenfors, A. (2020). From cip-deviations to a market for risk premia: A dynamic investigation of cross-currency basis swaps. *Journal of International Financial Markets, Institutions and Money*, *69*, 101245.

Christoffersen, P., & Diebold, F. X. (2006). Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science*, *52*(8), 1273–1287.

Clarke, J., Jandik, T., & Mandelker, G. (2001). The efficient markets hypothesis. *Expert financial planning: Advice from industry leaders*, *7*(3/4), 126–141.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., Terpenning, I., et al. (1990). Stl: A seasonal-trend decomposition. *J. off. Stat*, *6*(1), 3–73.

Cornett, M. M., Schwarz, T. V., & Szakmary, A. C. (1995). Seasonalities and intraday return patterns in the foreign currency futures market. *Journal of Banking & Finance*, *19*(5), 843–869.

Davey, A., & Flores, B. (1993). Identification of seasonality in time series: A note. *Mathematical and computer modelling*, *18*(6), 73–81.

Ding, Z., Granger, C. W., & Engle, R. F. (1993). A long memory property of stock market returns and a new model. *Journal of empirical finance*, *1*(1), 83–106.

Donnelly, B. (2019). *The art of currency trading: A professional's guide to the foreign exchange market*. John Wiley & Sons.

Émile, M., & Gumbel, J. (1960). Distributions des valeurs extrêmes en plusieurs dimensions. *Annales de l'ISUP*, *9*(2), 171–173.

Engel, C. (1996). The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of empirical finance*, *3*(2), 123–192.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, 987–1007.

Fama, E. F. (1984). Forward and spot exchange rates. *Journal of monetary economics*, *14*(3), 319–338.

Fama, E. F. (1991). Efficient capital markets: Ii. *The journal of finance*, *46*(5), 1575–1617.

Fender, J. (2020). Beyond the efficient markets hypothesis: Towards a new paradigm. *Bulletin of Economic Research*, *72*(3), 333–351.

Figlewski, S. (1997). Forecasting volatility. *Financial markets, institutions & instruments*, *6*(1), 1–88.

Forsberg, L., & Ghysels, E. (2007). Why do absolute returns predict volatility so well? *Journal of Financial Econometrics*, *5*(1), 31–67.

Francq, C., & Zakoian, J.-M. (2019). *Garch models: Structure, statistical inference and financial applications*. John Wiley & Sons.

Geweke, J. (1986). Comment on: Modelling the persistence of conditional variances. *Econometric Reviews*, *5*(1), 57–61.

Ghysels, E., Santa-Clara, P., & Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, *131*(1-2), 59–95.

Goodhart, C. A., Hall, S. G., Henry, S. B., & Pesaran, B. (1993). News effects in a high-frequency model of the sterling-dollar exchange rate. *Journal of Applied Econometrics*, *8*(1), 1–13.

Granger, C. W. (1978). Seasonality: Causation, interpretation, and implications. In *Seasonal analysis of economic time series* (pp. 33–56). NBER.

Guillaume, D. M., Dacorogna, M. M., Davé, R. R., Müller, U. A., Olsen, R. B., & Pictet, O. V. (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and stochastics*, *1*, 95–129.

Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a garch (1, 1)? *Journal of applied econometrics*, *20*(7), 873–889.

Harris, L. (1986). A transaction data study of weekly and intradaily patterns in stock returns. *Journal of financial economics*, *16*(1), 99–117.

Heij, C. (2004). *Econometric methods with applications in business and economics*. Oxford University Press, USA.

Hong, Y., & Chung, J. (2003). Are the directions of stock price changes predictable. *Statistical Theory and Evidence. Cornell University.*

Horowitz, J. L., & Savin, N. (2001). Binary response models: Logits, probits and semiparametrics. *Journal of economic perspectives*, *15*(4), 43–56.

ING. (2019). Ing financial markets fx derivatives fx swap. https://www.ingwb.com/binaries/content/assets/solutions-and-products/hedging-solutions/foreign-exchange/ing_financial_markets-product_description-fx_swap_be-en.pdf

Ito, T., & Hashimoto, Y. (2006). Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system. *Journal of the Japanese and International Economies*, *20*(4), 637–664.

Iwatsubo, K., Watkins, C., & Xu, T. (2018). Intraday seasonality in efficiency, liquidity, volatility and volume: Platinum and gold futures in tokyo and new york. *Journal of Commodity Markets*, *11*, 59–71.

Joe, H. (1993). Parametric families of multivariate distributions with given margins. *Journal of multivariate analysis*, *46*(2), 262–282.

Joseph, N. L., & Hewins, R. D. (1992). Seasonality estimation in the uk foreign exchange market. *Journal of Business Finance & Accounting*, *19*(1), 39–71. https://doi.org/https://doi.org/10.1111/j.1468-5957.1992.tb00609.x

Kallianiotis, I. N. (2018). How efficient is the foreign exchange market? *Athens journal of business & economics*, *4*(3), 293–326.

Kauppi, H., & Saikkonen, P. (2008). Predicting us recessions with dynamic binary response models. *The Review of Economics and Statistics*, *90*(4), 777–791.

Keynes, J. M. (1923). *A tract on monetary reform*. Cosimo Classics.

Korkie, B., Sivakumar, R., & Turtle, H. (2002). The dual contributions of information instruments in return models: Magnitude and direction predictability. *Journal of Empirical Finance*, *9*(5), 511–523.

Krohn, I., & Sushko, V. (2022). Fx spot and swap market liquidity spillovers. *Journal of International Money and Finance*, *120*, 102476.

Levi, M. D. (1978). The weekend game: Clearing house vs federal funds. *The Canadian Journal of Economics/Revue canadienne d'Economique*, *11*(4), 750–757.

Levich, R. M. (2011). Evidence on financial globalization and crises: Interest rate parity.

Liu, X., & Luger, R. (2015). Unfolded garch models. *Journal of Economic Dynamics and Control*, *58*, 186–217.

Ljung, G. M., & Box, G. E. (1978). On a measure of lack of fit in time series models. *Biometrika*, *65*(2), 297–303.

Malkiel, B. G. (1973). *A random walk down wall street*. Norton, New York.

Martens, M., Chang, Y.-C., & Taylor, S. J. (2002). A comparison of seasonal adjustment methods when forecasting intraday volatility. *Journal of Financial Research*, *25*(2), 283–299.

McFarland, J. W., Pettit, R. R., & Sung, S. K. (1982). The distribution of foreign exchange price changes: Trading day effects and risk measurement. *the Journal of Finance*, *37*(3), 693–715.

Mcmillan, D. G., & Speight, A. E. (2012). Daily fx volatility forecasts: Can the garch (1, 1) model be beaten using high-frequency data? *Journal of Forecasting*, *31*(4), 330–343.

Nelsen, R. B. (1997). Dependence and order in families of archimedean copulas. *Journal of Multivariate Analysis*, *60*(1), 111–122.

Nelsen, R. B. (2006). *An introduction to copulas*. Springer.

Nerlove, M. (1964). Spectral analysis of seasonal adjustment procedures. *Econometrica: Journal of the Econometric Society*, 241–286.

Nyberg, H. (2010). Dynamic probit models and financial variables in recession forecasting. *Journal of Forecasting*, *29*(1-2), 215–230.

Nyberg, H. (2011). Forecasting the direction of the us stock market with dynamic binary probit models. *International Journal of Forecasting*, *27*(2), 561–578.

Oh, G., Kim, S., & Eom, C. (2007). Market efficiency in foreign exchange markets. *Physica A: Statistical Mechanics and its Applications*, *382*(1), 209–212.

Payne, R., et al. (1996). *Announcement effects and seasonality in the intraday foreign exchange market*. London School of Economics; Political Science, Financial Markets Group.

Psarakis, S., & Panaretoes, J. (1990). The folded t distribution. *Communications in Statistics-Theory and Methods*, *19*(7), 2717–2734.

Ranaldo, A. (2009). Segmentation and time-of-day patterns in foreign exchange markets. *Journal of Banking & Finance*, *33*(12), 2199–2206.

Schweizer, B., & Sklar, A. (2011). *Probabilistic metric spaces*. Courier Corporation.

Shiskin, J. (1967). *The x-11 variant of the census method ii seasonal adjustment program*. US Department of Commerce, Bureau of the Census.

Singh, V. K. (2019). Day-of-the-week effect of major currency pairs: New evidences from investors' fear gauge. *Journal of Asset Management*, *20*(7), 493–507.

Sklar, M. (1959). Fonctions de répartition à n dimensions et leurs marges. *Annales de l'ISUP*, *8*(3), 229–231.

Vasisht, A. (2007). Logit and probit analysis. *IASRI, Library Avenue, New Delhi–110*, *12*.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: journal of the Econometric Society*, 817–838.

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management science*, *6*(3), 324–342.

Xu, J., Zhang, Z., Zhao, L., & Ai, D. (2011). The application review of garch model. *2011 International Conference on Multimedia Technology*, 2658–2662.

Yan, X. (2021). Multiplicative component garch model of intraday volatility. *arXiv preprint arXiv:2111.02376*.

Zhang, H. (2018). Intraday patterns in foreign exchange returns and realized volatility. *Finance Research Letters*, *27*, 99–104.