LUND UNIVERSITY
School of Economics and Management

Pivot Point Trading in the Foreign Exchange Market

A Test for Market Efficiency

by

David Frykmer & Ola Johnsson

June 2019

NEKN02 & NEKP01

Master Essay

Supervisor: Anders Vilhelmsson

Abstract

The foreign exchange (FX) market is today the world's broadest, most active, financial market. One of many technical analysis techniques amongst professional FX traders, trying to narrow down the best entry and exit points in this market, is to utilize pivot points. The ability of this technique has, as far as we know, never been rigorously evaluated. In collaboration with the Malmö-based hedge fund Century Analytics, we undertake such a test. The purpose is to investigate if this technique can help predict intraday interruptions for exchange rates. To this end, we apply both a trading strategy approach with the aim of establishing whether or not a certain trading strategy gives rise to exploitable profit opportunities, and a bootstrap resampling method. Studying three pairs of currencies, we find that the predictive power of applied pivot points varies across the pairs. Our main result is that we find no support for the hypothesis that pivot points reveal useful information in their traditional way about intraday price patterns in the FX market. Instead, our findings show that the technique can be profitable if it is implemented inversely for two of the pairs.

Keywords: Pivot points, technical analysis, foreign exchange market

Acknowledgements

First and foremost, we wish to extend our gratitude to Century Analytics, who provided us with insight into the hedge fund industry, assistance in the data collection process as well as valuable advice and guidance during the course of this thesis. We would also gratefully acknowledge the inputs and feedback from Anders Vilhelmsson and Adrian Mehic throughout the process.

Table of Contents

Glossary

Behavioral finance … proposes psychology-based theories to explain financial market anomalies. The purpose is to identify and understand why people make certain financial choices. Within behavioral finance, it is assumed the information structure and the characteristics of market participants systematically influence individuals' investment decisions as well as market outcomes.

Data snooping … refers to statistical inference that the researcher decides to perform after looking at the data (as contrasted with pre-planned inference, which the researcher plans before looking at the data).

Exchange rate … is the value (or price) of one nation's currency versus the currency of another nation or economic zone. For example, the EUR/USD rate says how many euros it takes to buy one U.S. dollar.

Exotic currency pair … is a major currency, paired alongside the currency of an emerging economy, such as the Mexican Peso, Russian rubles, and many currencies from countries outside the Euro-area. These pairs are usually not traded as often due to high transaction cost (which surpasses those seen in major or minor currencies) as a result of the absence of liquidity in these markets.

Floor trader … is an individual who trades on any contract market for his or her own account.

Intraday … literally means "within the day", and refers here to price movements of a given asset over the course of one day of trading.

Level 1 quote … refers to the first level in the orderbook which only includes the best bid price and ask price, and sometimes also the accumulated volumes displayed as bid size and ask size. This can be compared to a level 2 quote that provides additional information, such as multiple levels of bids and offers. The additional information may indicate which side is more eager or more powerful, and may predict the short-term direction of the market price.

Major currency pairs … is a major currency, paired alongside another major currency. Major currencies are simply the most traded currencies and they include the euro (EUR), the U.S. dollar (USD), the British Pound sterling (GBP), the Australian dollar (AUD), the Japanese yen (JPY) and the Swiss franc (CHF).

Market maker … is a company or an individual that quotes both a buy and a sell price in a financial instrument or commodity held in inventory, hoping to make a profit on the bid-offer spread, or turn.

Mid-price … is the exact mid-point of the current bid-ask spread quoted for an asset.

Noise traders … is a term used to describe investors who make buy and sell decisions without the support of professional advice or advanced fundamental analysis. These investors typically follow trends and overreact to good and bad news. Trading by noise traders tends to be impulsive and based on irrational exuberance, fear or greed.

Pivot levels … refer to the support and resistance levels that are calculated from the pivot points (P_t). In this study there are 6 pivot levels: S_1, S_2, S_3, R_1, R_2 and R_3, where $S_1 - S_3$ are support levels and $R_1 - R_3$ are resistance levels. The pivot levels indicate potential trading ranges for the next trading session. For further explanation, see page 15-16.

Pivot points … are used to determine trading levels at which the price trend is inclined to change direction and head to possible support and resistance levels. They are designed to be trend-predicting indicators instead of lagging indicators. In this study, the pivot points (P_t) are calculated as the average of the previous trading day's high, low and close. For further explanation, see page 15-16.

Resistance levels … are levels or areas on the chart under the market where selling interest is sufficiently strong to overcome buying pressure. As a result, an increase is halted and prices turn back again. For example, a FX trader who publishes a resistance level of €1.30/$ would claim that the EUR/USD is likely to stop increasing if it reaches €1.30/$, and from that level turn back down again. Support and resistance levels are together referred to as pivot levels.

Stop-limit order … a stop-limit order can be executed at a specified (or potentially better) price, after a given stop price has been reached. Once the stop price is reached, the stop-limit order becomes a limit order to buy or sell at the limit price or better.

Stop-loss order … is designed to limit an investor's loss on a long or short position. For example, setting a stop-loss order for 10% below the price at which you bought the asset will limit your loss to 10%.

Support levels … are levels or areas on the chart under the market where buying interest is sufficiently strong to overcome selling pressure. As a result, a decline is halted and prices turn back up again. For example, a FX trader who publishes a support level of €1.20/$ would claim that the EUR/USD is likely to stop falling if it reaches €1.20/$, and from that level turn back up again. Support and resistance levels are together referred to as pivot levels.

Tick data … or tick-by-tick data, is information that is used to watch the current price movements at its most detailed level. Usually, a tick refers to the change in the price of an asset from trade to trade. Tick data is much more detailed compared to, for example, daily data. The number of tick observations in one single day in a liquid market is equivalent to the number of daily data within 30 years (Dacorogna et al., 2001:6).

Trading day … time frame that a financial market is open for making transactions. Most stock and futures exchanges have a limited trading day and only allow transactions during the week, but the FX market instead has trading sessions that correspond to normal Tokyo, London and New York business hours since the FX market trades around the clock from Sunday afternoon to Friday evening EST.

Win rate … the number of trades won out of total trades, to determine the probability of a trader's success. A win rate above 50% is usually favorable.

Win/loss ratio … or the "success ratio", is the ratio of the total number of winning trades to the number of losing trades. A win/loss ratio above 1.0 is usually favorable.

1 Introduction

The predictability of financial markets remains a widely debated topic. One of the major gulfs between academic finance and industry practice is the separation that exists between technical analysts and their academic critics (Lo et al., 2000). Academics, in general, are inclined to support a more efficient view of financial markets. The reason for this is partly empirical evidence and partly convenience. For one thing, it would be very difficult to develop financial theories if the asset price was not tied to the intrinsic value of an asset (Byström, 2004:196-197). Secondly, the behavior of market participants induces risk-adjusted returns that obey the efficient market hypothesis, otherwise there would exist a "money-machine" producing unlimited wealth, which cannot occur in a stable economy (Timmermann & Granger, 2004).

Despite the overpowering logic of efficient financial markets, many economists and statisticians came to believe that financial markets are at least somewhat predictable based on technical analysis as well as certain fundamental valuation metrics in the early 2000's (Malkiel, 2003). In contrast to fundamental analysis, which was quick to be adopted by the scholars of modern quantitative finance, technical analysis – which involves the prediction of asset price movements from inductive analysis of past movements – has been an orphan from the start (Lo et al., 2000). Although the discipline of technical analysis still lacks support by any economic theory, practitioners in especially the FX market have for a long time believed it to be very useful and powerful. Already two decades ago, 25 to 30 percent of FX traders based most of their trades on technical trading signals (Cheung & Chinn, 1999; Cheung & Wong, 2000). More broadly, technical analysis has been used as either a primary or secondary source of trading information by more than 90 percent of FX market participants in London (Allen & Taylor, 1992) and Hong Kong (Lui & Mole, 1998).

The FX market is today the world's broadest, most active, financial market, with an average daily volume of about $5 trillion. It functions 24 hours a day for 5.5 days a week, opening on Sunday afternoon in Sydney (Australia) and closing on Friday, along with the NY market (Fernández, 2019). The FX market consists almost entirely of professional traders (Sager & Taylor, 2006), and many of them utilize an arsenal of technical analysis techniques in order to decide whether to buy or sell at certain points in time. According to Osler (2000), one of the

techniques common to technical FX traders is to employ support and resistance levels into price charts of currency pairs. A support level is the straight line that delimits the price trend from below, and the resistance level is the one that delimits it from above. The idea of these levels is to identify a range at which a price trend is likely to stop and may be reversed. Once a range of support and resistance has been identified, the general belief among traders is that it provides valuable points for entry or exit. Usually, the support level indicates an entry point, and the resistance level indicates a point to sell, or, at least, not be long the asset (Osler, 2000).

Another technical analysis technique, which is closely related to the technique of support and resistance levels, is referred to as "pivot points". In its original form, the support and resistant levels technique has only one support level and one resistance level, creating a single trading range. The pivot points-technique, on the other hand, is an extended tier-based approach with multiple support and resistance levels. The pivot point at time t, which in its essence is a specific quote of the exchange rate, is calculated from observed price information at $t-1$. Then, the tiered pivot support and resistance levels, or pivot levels, are calculated from the value of the pivot point. In practice, the pivot levels are assumed to reveal where the price has a higher probability of trading as the market changes. They can also help in identifying stop and profit levels (Williams, 2015). Often, the best trading tools are the simplest, and according to Bhandari (2012) pivot points qualify to this category.

Despite the wide use of pivot points in short-term FX forecasting, the ability of these trading signals to predict intraday trend interruptions has, as far as we know, never been rigorously evaluated. In collaboration with the Malmö-based hedge fund Century Analytics, we undertake such a test, applying the pivot points-technique on three currency pairs – the EUR/USD, the USD/CHF and the GBP/JPY. The central idea underlying this study is that if the FX market is efficient in the short-term, one should not be able to use publicly available information to predict future short-term changes in exchange rates and hence to make abnormal profits. In other words, if the FX market is efficient in the short-term the pivot points-technique should fail.

The purpose of this thesis is to investigate if pivot points help predict intraday interruptions for exchange rates.

The data samples used in this study are provided by Century Analytics and consist of level 1 quoted tick data of the EUR/USD and the USD/CHF, ranging from 1st June 2017 to 29th

December 2017, and level 1 quoted one-minute data of the GBP/JPY, ranging from 2nd January 2003 to 13th October 2017. A level 1 quote refers to the first level in the orderbook which only includes the best bid price and ask price for each timestamp. For all three currency pairs, we calculate the mid-price by dividing the sum of the bid price and ask price by two for each timestamp. The mid-prices are used as the market prices (or spot prices). The reason for selecting the above listed currency pairs is primarily because they are all seen as major currency pairs and secondarily because we believe they make up an interesting mix based on their data types, different activity levels, daily traded volumes and perhaps most intensive trading hours during a trading day. By currency pairs, the EUR/USD is the world's most actively traded (Dunis & Williams, 2002) and is likely to be much more efficient than other assets (Byström, 2004:197). The USD/CHF and the GBP/JPY are not as actively traded as the EUR/USD, but nevertheless not exotic.

In order to evaluate the pivot points-technique, we apply two different methods. First, we apply a trading strategy approach that simulates the actions of a trader according to a high-frequency trading strategy with the aim of establishing whether or not this strategy gives rise to exploitable profit opportunities, which could be seen as evidence against market efficiency. Alongside this approach, the binomial test is used to rule out the possibility of luck and the likelihood ratio (LR) test for independence is used to investigate the independence of profitable and unprofitable trades. Second, we involve the application of a bootstrap resampling method, which boils down to evaluating the distributional behavior of the exchange rates in comparison with the distribution of thousands of pseudo exchange rates.

The results show no support for the use of pivot points in their traditional way. For two of the currency pairs, the support levels defined from the pivot points should be used as selling signals and the resistance levels should be used as buying signals. In other words, pivot points should be implemented inversely. Moreover, we find that the success of the inversely defined trading strategy based on pivot points are time dependent for the GBP/JPY, but not for the other currency pairs. Lastly, a comparison between the actual price evolutions that hit the pivot levels and resampled data supports the results of the trading strategy, but also implies that the variance is higher around the pivot levels. The higher variance might imply that there indeed is higher activity in the FX market as the price fluctuates around the pivot levels.

The remainder of this thesis is structured as follows. Chapter 2 provides a review of the applicable research to this thesis. In chapter 3, the relevant theories regarding the objective of

this thesis are presented. Chapter 4 presents the data used in this study. Chapter 5 introduces the chosen methodology. In Chapter 6, the results are presented and analyzed. Chapter 7 provides a conclusion as well as suggestions for further research in the area.

2 Previous Research

Over the past twenty years or so, international financial economists have increasingly turned their attention to the study of technical analysis in an attempt to understand both the behavior of asset prices and the behavior of market participants; so much so, in fact, that quite an extensive literature has developed on this topic (Menkhoff & Taylor, 2007). Overall, the empirical evidence is mixed and inconclusive as to whether technical approaches can generate superior performance. In an extensive survey by Park & Irwin (2007), the authors review the evidence on the profitability of technical analysis. They find that modern studies indicate that technical trading strategies consistently generate economic profits in a variety of speculative asset markets at least until the early 1990s, but the profits decline substantially or disappear altogether in subsequent years. Among a total of 95 modern studies, 56 studies find positive results regarding technical trading strategies, 20 studies obtain negative results, and 19 studies indicate mixed results. Despite the positive evidence on the profitability of technical trading strategies, most empirical studies are subject to various problems in their testing procedures, such as data snooping and difficulties in estimation of risk and transaction costs.

In summary, there does not seem to be a clear consensus regarding the profitability and usefulness of technical analysis. However, other interesting findings and suggestions from previous research give a different view to the occasionally success of technical analysis strategies. A classical explanation of the profitability of technical analysis methods given by Merton (1948) states that technical analysis can work as a self-fulfilling prophecy when a large enough group of investors (but not necessarily all of them) uses technical analysis and follows its implied signals. Another explanation presumes that, with the tools of technical analysis, traders can anticipate the movements of noise traders and make profits on that basis (De Long et al., 1990). Lo (2004) takes it one step further and introduces the concept of the Adaptive Markets Hypothesis, arguing that arbitrage opportunities do exist from time to time, but as profitable trading opportunities are exploited, they disappear. At the same time, changing business conditions induce new profit opportunities that exist until they – too – eventually collapse. Using a similar analogy, Grossman & Stiglitz (1980) argues that the assumptions that all markets, including the market for information, are always in equilibrium and always

perfectly arbitraged are inconsistent when arbitrage is costly. Asset price reactions to information are not expected to be occur instantaneously when the costs of obtaining and exploiting information are significant. Instead, asset price adjustment due to information must allow traders who invest resources in processing information to earn a normal rate of return on their investment.

Previous research of technical analysis more specifically in the FX market vary quite a lot in terms of undertaken tests, technical trading techniques and trading rules. Despite the variation, many studies indicate that technical trading strategies can be profitable. Gradojevic & Lento (2015) finds that technical indicators constructed from order flows can be profitable for the CAD/USD. In Neely et al. (1997), the authors use a genetic program as a search procedure for identifying optimal technical trading rules and find evidence of economically significant out-of-sample excess returns to those rules for each of six currency pairs over the period 1981-1995. Levich & Thomas (1993) implements a testing procedure based on bootstrap methodology and finds evidence on the profitability and statistical significance of technical trading rules in the FX market using currency futures contracts. Chang & Osler (1999) finds that a trading strategy based on the head-and-shoulders chart pattern is profitable for daily USD exchange rates vis-à-vis the German mark (DEM) and the JPY, although not for four other USD exchange rates. To a broader extent, Qi & Wu (2006) investigates the universe of 2,127 parameterizations of technical trading rules, which are by far the largest in number and the most comprehensive in scope in FX studies. The authors find that performances vary a lot depending on trading rules, performance measures and time periods. Furthermore, they suggest that there is a potential change in the dynamics of exchange rate series and that the FX market becomes increasingly more efficient over time.

The specific conclusion that exchange rates tend to stop trending at pivot levels has no precedent in the academic literature. The closest point of comparison found are four existing studies of support and resistance levels – Brock et al. (1992); Zapranis & Tsinaslanidis (2012); Curcio et al. (1997); and Osler (2000). Brock et al. (1992) tests the hypothesis that prices tend to move rapidly once the support and resistance levels are breached and finds that the hypothesis is true for daily movements of the Dow Jones stock index from 1897 to 1986, but the profits from exploiting this feature may not be sufficient to offset transaction costs. Zapranis & Tsinaslanidis (2012) evaluates whether support and resistance levels are efficient trend-reversal predictors, and if they can help generate systematic abnormal returns on the US stock markets

over the last two decades. The authors' findings are aligned with the efficient market hypothesis and more concretely, support levels outperform resistance ones in predicting trend interruptions, but they fail to generate excess returns when they are compared with simple buy-and-hold strategies.

Curcio et al. (1997) analyzes the profitability of technical trading rules based on pre-defined price ranges using hourly rates for the DEM, the JPY, and the GBP against the USD, and concludes that on average the profits are not significant, even before transaction costs were taken into account. However, the trading rules may hold true during periods of strong trending. Osler (2000) investigates the predictive power of support and resistance levels for high frequency (one-minute) data of three currencies relative to the USD: the DEM, the JPY, and the GBP (from January 1996 through March 1998). The author fails to report profitability after taking transaction costs into account.

3 Theoretical Review

This chapter provides a theoretical background to this thesis. Section 3.1 describes the Efficient Market Hypothesis which is associated with the idea that asset price movements are random. Section 3.2 presents the more evolutionary alternative to market efficiency, namely the Adaptive Market Hypothesis. In section 3.3, a brief introduction to the field of technical analysis is given and the techniques to identify support and resistance levels as well as pivot points are explained.

3.1 Efficient Market Hypothesis

The Efficient Market Hypothesis (EMH) is essentially an investment theory that states that asset prices in financial markets fully reflect all available information. Formally, this implies that in an efficient financial market it is impossible to "beat the market" consistently on a risk-adjusted basis since market prices should only react to new information. The EMH is associated with the idea of a "random walk", which is a term used to characterize a price series where all subsequent price changes represent random deviations from previous prices. The logic of the random walk idea is that if the flow of information is unimpeded and immediately reflected in asset prices, tomorrow's price change will reflect only tomorrow's news and will be independent of the price changes today (Malkiel, 2003). By assuming a random walk model for asset prices completely rules out the probability of strategies based on technical analysis (Smith et al., 2016).

Timmermann & Granger (2004) states that the EMH is based on the overpowering logic that if risk-adjusted returns were predictable, many investors would use them to generate unlimited profits. The behavior of market participants induces risk-adjusted returns that obey the EMH, otherwise there would exist a "money-machine" producing unlimited wealth, which cannot occur in a stable economy. According to Byström (2004:196), markets are almost always without exception explicitly, or implicitly, assumed to be efficient. The reason for this is partly

convenience and partly empirical evidence. If the asset price was not tied to the intrinsic value of the asset, it would be very difficult to develop financial theories.

To answer the question of whether markets are efficient (i.e. random) or not the concept of efficiency must be defined. According to Fama (1970): "A market in which prices always fully reflect all available information is to be seen as efficient". To test whether this statement holds or not, it is also necessary to define what is meant by 'all available information'. The classic taxonomy of information sets, due to Roberts (1967), distinguishes between three forms of efficiency:

- **Weak Efficiency:** The price reflects only historical prices and returns themselves.
- **Semi-Strong Efficiency:** The price reflects all information known to all the market participants (publicly available information).
- **Strong Efficiency:** The price reflects all the information known to any market participant, including insider information.

If abnormal asset return with respect to the chosen information set is unforecastable, and in this sense random, then the hypothesis of market efficiency is not rejected. The abnormal return is the difference between the return on an asset and its expected return on a risk-adjusted basis, where the expected return is usually assumed to be constant over time. (Campbell et al., 1997:22).

The opinions regarding the EMH are divided. Byström (2004:197) argues that academics, in general, are more inclined to support a more effective view of capital markets than practitioners (stockbrokers, fund managers, analysts, and the like) who are more in favor of thoughts that market prices can be predicted. Surrounding the determination and prediction of exchange rates, there are certain core financial theories that traditionally have been used by economists to help explain the long-run trends in exchange rates. These theories link exchange rates, price levels, inflation rates and interest rates and they are usually referred to the (international) parity conditions (Eiteman et al., 2010:164).

According to Eiteman et al. (2010:164), the parity conditions do not always work out to be "true" when compared to what practitioners observe in the real world, but they are central to any understanding of how international business is conducted and funded in the world today. One of these theories is called the international Fisher effect. This theory states that the spot exchange rate should change in an equal amount but in the opposite direction to the difference

in interest rates between two currencies (Eiteman et al., 2010:165). An extension to the international Fisher effect is the theory of interest rate parity (IRP). This theory states that the difference in the national interest rates for securities of similar risk and maturity should be equal to, but opposite in sign to, the forward rate discount or premium for the foreign currency, except for transaction costs. When the spot- and forward exchange rates are not in equilibrium as described by the IRP, the potential for risk-free or arbitrage profit exists. This is called the covered interest arbitrage (CIA) (Eiteman et al., 2010:184).

According to Yun (2017), nominal exchange rates are slow to converge to the long-run parity conditions. In a comprehensive investigation of the forecasting performance of a large number of exchange rate models, Cheung et al. (2005) concludes that no model consistently outperforms a random walk at short horizons, thus confirming that random walk forecasts of the exchange rates generally outperform alternative models extracted from economic theories, particularly at short horizons.

3.2 Adaptive Market Hypothesis

Especially since the dawn of behavioral finance in the 1980s the EMH has come under relentless attack. In an attempt to offer a new perspective that reconciles the two opposing schools of thought, Lo (2004) suggests an evolutionary alternative to market efficiency. The proposed reconciliation, which is called the Adaptive Markets Hypothesis (AMH), provides a ground where the EMH can co-exist alongside behavioral finance in an intellectually consistent manner. In this hypothesis, market efficiency is not an all-or-none condition but is a characteristic that varies continuously over time and across markets. Prices reflect as much information as dictated by the combination of environmental conditions and the number and nature of distinct groups of market participants, each behaving in a common manner.

Lo (2004) states that within a single market where multiple market participants are competing for rather scarce resources, that market is likely to be highly efficient. An example is the market for 10-year U.S. Treasury notes, where most relevant information is incorporated into prices within minutes. On the other hand, if a small number of market participants are competing for rather abundant resources in a given market, that market will be less efficient. An example of such a market is the market for oil paintings from the Italian Renaissance. The main point is thus that market efficiency cannot be evaluated in a vacuum, but is highly context dependent

and dynamic, just as insect populations advance and decline as a function of the seasons, the number of predators and prey they face, and their abilities to adapt to an ever-changing environment.

3.3 Technical Analysis

Technical Analysis (TA), also known as “charting”, has been part of financial practice for many decades, but this discipline has not received the same level of academic scrutiny and acceptance as more traditional approaches, such as fundamental analysis. One of the main drawbacks is the highly subjective nature of TA. The presence of geometric shapes in historical price charts is often in the eyes of the beholder (Lo et al., 2000).

According to Pring (1991:2), the technical approach to investment is essentially a reflection of the idea that prices move in trends which are determined by the changing attitudes of investors toward a variety of economic, political, monetary and psychological forces. The art of TA is to identify trend changes at an early stage and to maintain an investment posture until the weight of the evidence indicates that the trend has reversed.

Human nature remains more or less constant and tends to react to similar situations in consistent ways. TA is therefore based on the assumption that people will continue to make the same mistakes they have made in the past. Human relationships are extremely complex and never repeat in identical combinations. The financial markets, which are reflections of people in action, never duplicate their performance exactly, but the recurrence of similar characteristics is sufficient to enable technicians to identify major juncture points. By studying the nature of previous market turning points, it is possible to develop some characteristics that can help to identify major market tops and bottoms. However, since no single indicator has signaled, or indeed could signal, every cyclical market juncture, technical analysts have developed an arsenal of tools to isolate these points (Pring, 1991:2).

Many traders base their trading strategies, either partly or entirely, on TA. According to Pring (1991:3), most of these strategies can be broken down into three essential categories:

- Sentiment
- Flow-of-funds

- Market structure indicators.

Sentiment or *expectational indicators* monitor the actions of different market participants, e.g. mutual funds, hedge funds, and floor specialists. For example, insiders and hedge funds tend to be correct at market turning points; in aggregate, their transactions are on the buy side toward market bottoms and on the sell side toward tops. Other market participants, such as advisory services, represent the majority opinion and they are often wrong about these turning points. Since the consensus normally mistimes market turning points, these kind of indicators of market psychology are a useful basis from which to form a contrary opinion (Pring, 1991:3).

Flow-of-funds indicators analyze the financial position of various investment groups in an attempt to measure their potential capacity for buying or selling assets. The price at which an asset transaction takes place has to be the same for the buyer and the seller, so naturally the amount of money flowing out of the market must equal that put in. Flow-of-funds strategies are based on the before-the-fact balance between supply and demand. For example, if at a given price there is a preponderance of buyers over sellers, it follows that the price will have to rise to bring buyers and sellers into balance (Pring, 1991:3-4).

The set of the *market structure indicators* includes many of the strategies which are generally most associated with TA. These indicators monitor the trend of various price indices, market breadth, volume, cycles, etc., in order to evaluate the status of bull and bear markets. Many of them rise and fall together, but toward the end of market movements the paths of them diverge from the price. Such divergences offer signs of technical deterioration during advances, and technical strength following declines. Through careful observation of these signs of latent strength and weakness, technical analysts are alerted to the possibility of a reversal in the trend of the market itself (Pring, 1991:4). One of the basic modes of analysis in this field of indicators is that of support and resistance levels, which is further discussed in section 3.2.2.

3.3.1 Technical Trading of Currencies

The principles that underlie analysis of currencies from a technical aspect are basically the same as those used in any other financial market, such as the stock market or the commodity market. However, there are certain differences which potentially make technical trading strategies profitable in the FX market, but not in the stock market. According to Kubińska et al. (2016), the FX market is not operating in the form of a centralized exchange compared to a regulated

stock market. It is organized as over-the-counter (OTC) market, where transactions are bilateral; there is no single market price like at exchange-based multilateral markets. This means that transactions may (and do) occur simultaneously or near simultaneously in the market at different prices. Furthermore, information transparency is a crucial aspect of the FX market. Participants of the FX market have only partial knowledge about the trades of other investors; they are aware of liquidity, supply, and demand only for some segments of the market, not the entire market. This makes the process of price-information interaction sometimes difficult to observe and understand.

Another major difference is that the price, or level, of a currency is always relative to the prices of all other currencies. Hence, a rising trend in the EUR/USD will be bullish for Europeans, or any holder of EUR, and bearish for Americans, or any holder of USD, but when stocks or gold are declining sharply, they are bearish for everyone (Pring, 1991:465).

3.3.2 Support and Resistance Levels

A TA technique used to estimate the potential extent or duration of a trend is to employ support and resistance levels into price charts. A support level is the straight line that delimits the price trend from below, and the resistance level is the one that delimits it from above. There is little disagreement among researchers and analysts on the definition of support and resistance levels. According to Pring (1991:199), support and resistance levels represent a concentration of demand and supply sufficient to halt a price move at least temporarily. Murphy (1986:59) states that support is a level or area on the chart under the market where buying interest is sufficiently strong to overcome selling pressure. As a result, a decline is halted, and prices turn back again. Resistance is the opposite of support.

Osler (2000) states that the idea of support and resistance levels are to identify ranges at which a price trend is likely to stop and may be reversed. Once a range of a support level and a resistance level has been identified, the general belief among traders is that it provides valuable potential trade entry or exit points. Usually, the support level indicates an entry point (hence its name), and the resistance level indicates a point to sell, or, at least, not be long the asset (Osler, 2000). As the price reaches a point of support or resistance, it will do one of two things – bounce back into the range of the support and resistance levels, or violate the price level and continue in its direction. Many technical strategies are based on the belief that support and resistance levels will not be broken. It is, however, important to emphasize that a support and/or a

resistance level is not a hard limit. The price can go lower (higher) than the support (resistance) level (Maeda & Jacka, 2018). Whether the price is halted by the support or resistance level, or it breaks through, traders can "bet" on the direction by placing a trade and can quickly determine if they are correct. If the price moves in the wrong direction, the position can be closed at a small loss. If the price moves in the right direction, however, the profit may be substantial.

Figure 1 gives an illustrative example of how the support and resistance technique can be implemented. Potentially, a trader would assume that if the EUR/USD rate reaches the support level of approx. €1.161/$, the price would stop falling. Perhaps, he or she uses this level as the entry point. If the price moves below this level the position is closed, but if the price moves in the right direction the trader either closes the position at the resistance level of approx. €1.165/$ or puts in a stop-limit order at this level.

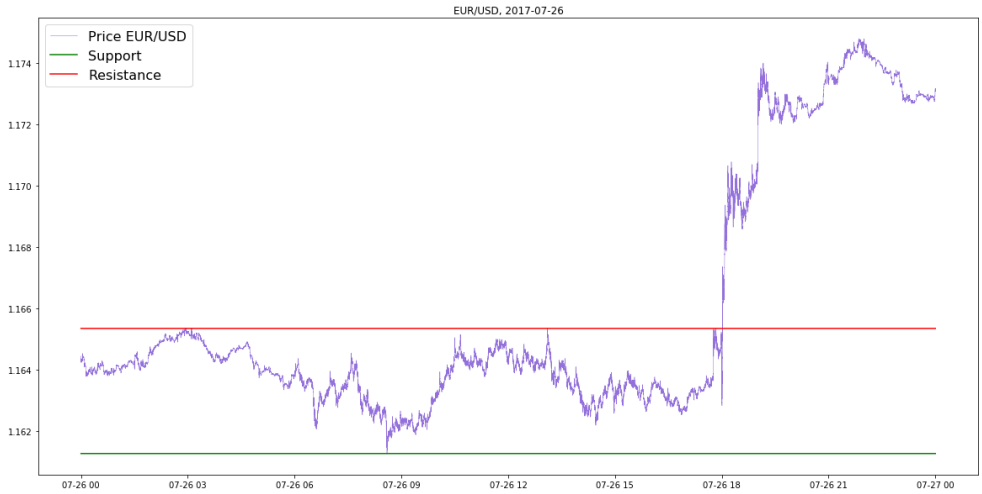

Figure 1: The figure gives an illustrative example of how the support and resistance technique can be implemented. The figure shows the intraday movements of the EUR/USD rate during the 26th July 2017. The horizontal, green line in the bottom shows the support level of approx. €1.161/$ and the horizontal, red line, shows the resistance level of approx. €1.165/$. *Source: Authors' calculations.*

According to Osler (2000), practicing technical analysts consult a variety of information inputs to identify the support and resistance levels relevant for the coming day. These include visual assessments of recent price performance, simple numerical rules based on recent price performance, inference based on knowledge about order flow, and market psychology.

3.3.3 Pivot Points

Before the age of personal computers, market makers and floor traders needed a quick method to determine whether asset prices were cheap or expensive. These prices were considered "pivotal" to the direction of prices. According to Williams (2015), traders developed a simple mathematical equation for this that is still in use today. This simple formula computes a tier-based approach to support and resistance in the market called "pivot points". Bhandari (2012) states that there are variations to the pivot point calculations, but the most popular among FX traders is:

$$P_t = \frac{(High_{t-1} + Low_{t-1} + Close_{t-1})}{3} \tag{1}$$

where P_t is the value of the pivot point at time t, and the nominator on the RHS is simply the sum of the intraday high, low and close prices at $t - 1$.

Then, the tiered pivot support and resistance levels, or pivot levels, are calculated from the P_t. The pivot levels are used to indicate potential trading ranges for the next trading session (Bhandari, 2012). On the support side, there are often two or three pivot levels which usually are defined as follows

$$S_1 = P_t - (High_{t-1} - P_t) \tag{2}$$

$$S_2 = P_t - (High_{t-1} - Low_{t-1}) \tag{3}$$

$$S_3 = Low_{t-1} - 2 \times (High_{t-1} - P_t) \tag{4}$$

Similarly, on the resistance side, the pivot levels are usually defined as

$$R_1 = P_t + (P_t - Low_{t-1}) \tag{5}$$

$$R_2 = P_t + (High_{t-1} - Low_{t-1}) \tag{6}$$

$$R_3 = High_{t-1} + 2 \times (P_t - Low_{t-1}) \tag{7}$$

In practice, all these levels are assumed to help reveal where the price has a higher probability of trading as the market changes. They can also help in identifying stop and profit levels. The most important level is the pivot point (P_t) itself. It is the level above or below which the price

is expected to advance toward R_1 or S_1. Furthermore, the inner tiers are assumed to help determine where the price is consolidating when it is hovering around a particular level. The outer pivot tiers can also reveal where the price is overextended and likely to reverse. In this way, they act as types of oscillators, showing where price is oversold and overvalued. This makes the outer pivot point tiers a place to take profits, adjust stop-limits and watch for potential countertrend setups (Williams, 2015).

Figure 2 gives an illustrative example of how the technique of pivot points can appear for an intraday of the EUR/USD. The pivot point and the six different pivot levels are calculated according to the equations $(1) - (7)$ above. In Figure 2, it can be seen that the EUR/USD trades above the pivot point (P_t) during this particular intraday.

Figure 2: The figure gives an illustrative example of how the pivot points-technique can be implemented. The figure shows the intraday movements of the EUR/USD rate during the 29th August 2017. The horizontal, green line in the middle shows the pivot level (P_t) itself of approx. €1.195/$. Above this level, the three different levels of resistance (all in grey tones) are shown and below the three levels of support (all in blue tones) are shown. *Source: Authors' calculations.*

Although pivot levels, according to Bhandari (2012), most often are calculated for a day's session, they are also used on an hourly basis. In another application, many traders apply the daily pivots on hourly, 30-minute, 15-minute and 5-minute charts. The theory behind this use is that the daily pivots are more reliable than pivots derived from shorter time frames, which are less accurate and significant.

4 Data

The data samples used in this study are provided by Century Analytics and consist of level 1 quoted tick data of the EUR/USD and the USD/CHF, ranging from 1st June 2017 to 29th December 2017, and level 1 quoted one-minute data of the GBP/JPY, ranging from 2nd January 2003 to 13th October 2017. A level 1 quote refers to the first level in the orderbook which only includes the best bid price and ask price for each timestamp. For all three currency pairs, we calculate the mid-price by dividing the sum of the bid price and ask price by two for each timestamp. The mid-prices are used as the market prices (or spot prices). In total, there are 175,120,345 tick observations of mid-prices for the EUR/USD and 116,104,191 for the USD/CHF, which are distributed over 31 weeks or 152 trading days. There are 5,358,976 one-minute observations of mid-prices for the GBP/JPY, which are distributed over 779 weeks or 3,839 trading days. Figure 3 shows the daily price development for the three currency pairs during their sample periods.

Figure 3: The figure shows the daily price development for the three currency pairs during their sample periods. The top left graph shows the EUR/USD over the period from 1st June 2017 to 29th December 2017 (31 weeks or 152 trading days). The top right graph shows the USD/CHF over the same period as the former. The bottom graph shows the GBP/JPY over the period from 2nd January 2003 to 13th October 2017 (779 weeks or 3,839 trading days). *Source: Bloomberg.*

The primarily reason for selecting the EUR/USD, the USD/CHF and the GBP/JPY is because they are all seen as major currency pairs, which means they are more liquid compared to any exotic currency pair. The second reason is that we believe these specific pairs make up an interesting mix due to their differences in data types, activity levels, daily traded volumes and perhaps most intensive trading hours during a trading day. We believe that the difference in activity levels potentially can have an impact on the performance of the pivot points-technique, which is something we will be able to evaluate.

The sample periods of the three currency pairs are determined based on data availability from Century Analytics. Of the available data for each of the currency pairs, the above listed sample periods are the closest to today's date. Worth mentioning is that tick data (and probably even one-minute data), to our knowledge, is not distributed free of charge compared to, for example, daily data or monthly data. More recent data, closer to today's date, of the specific data types would therefore imply a cost.

In Figure 3, it can be seen that the three data samples differ in lengths depending on the underlying data type. The reason for using tick data for two of the currency pairs is because it provides the possibility to study price movements at its most detailed level, near continuous observation. Tick data changes in the price from trade to trade, which is essential when evaluating technical trading strategies in a high-frequency domain. With this type of data, it is also possible to analyze during which minutes, hours or days there is most activity in the market. Figure 4 illustrates the hourly distribution of changes in the order book for the EUR/USD and the USD/CHF respectively.

A drawback of using tick data is that a substantial amount of data is needed. The number of observations in one single day of a liquid market is equivalent to the number of daily data within 30 years (Dacorogna et al., 2001:6). This implies that we limit ourselves to study a relatively short period using this type of data since it gets both cumbersome and time-consuming to manage any greater data sets. A sample of more than 175 million tick observations for the EUR/USD and 116 million for the USD/CHF is already challenging as it is. To be able to evaluate the pivot points-technique over a longer period, we choose therefore to include the GBP/JPY as an additional currency pair, but for this using one-minute data of mid-prices.

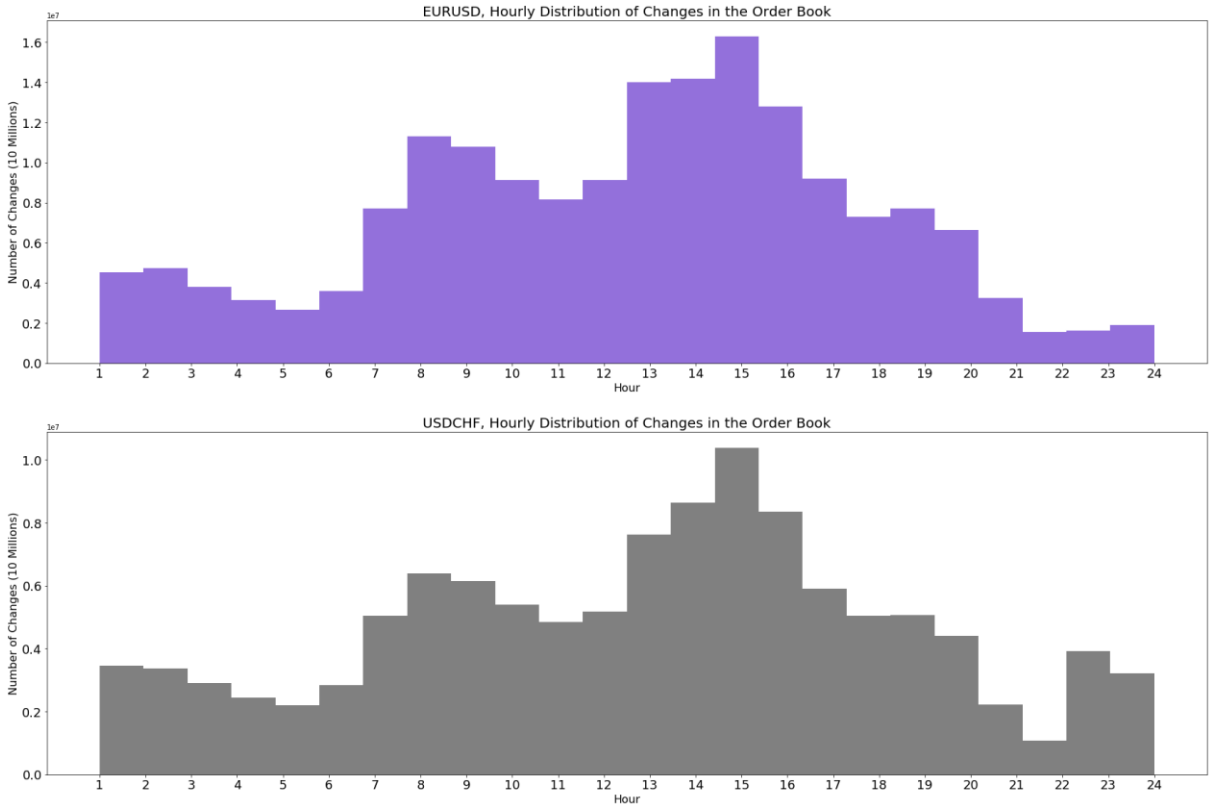

Figure 4: The top histogram shows the hourly distribution of changes in the order book for the EUR/USD over the period from 1st June 2017 to 29th December 2017 (31 weeks or 152 trading days). Similarly, the hourly distribution of changes in the order book for the USD/CHF is shown in the bottom histogram over the same period as the former. Both series correspond to the time-zone: UTC±0. In both histograms, it can be seen that most changes in the order book occur around the FX opening trading in Europe (07:00) and in New York (13:00). *Source: Authors' calculations.*

A summary of some descriptive statistics of all currency pairs is given in Table 1.

Table 1: This table provides a summary of some descriptive statistics of the data series, for each of the three currency pairs. The data samples for the EUR/USD and the USD/CHF cover tick observations of mid-prices over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The data sample for the GBP/JPY covers one-minute observations of mid-prices over the period from 2nd January 2003 to 13th October 2017 (779 weeks or 3,839 trading days).

Currency pair	Data type	*n*	Mean	Min.	Max.	Std. Dev.	Skewness	Kurtosis
EUR/USD	Tick	175 120 345	1.17	1.11	1.21	0.024	-0.85	-0.34
USD/CHF	Tick	116 104 191	0.97	0.94	1.00	0.014	0.36	-0.59
GBP/JPY	One-minute	5 358 976	170.39	116.83	251.17	34.07	0.24	-1.08

Due to the relatively low activity on Sunday afternoons compared to any full-day of trading, we realign the data observations for each Sunday by including them in the next-coming Monday's data. If we were to handle Sundays separately the spread between pivot levels for

the next-coming Mondays would be too low and most probably lead to erroneous results for the pivot points-technique. Figure 5 provides an illustration of how the pivot points-technique can appear for a given week after the Sunday adjustment.

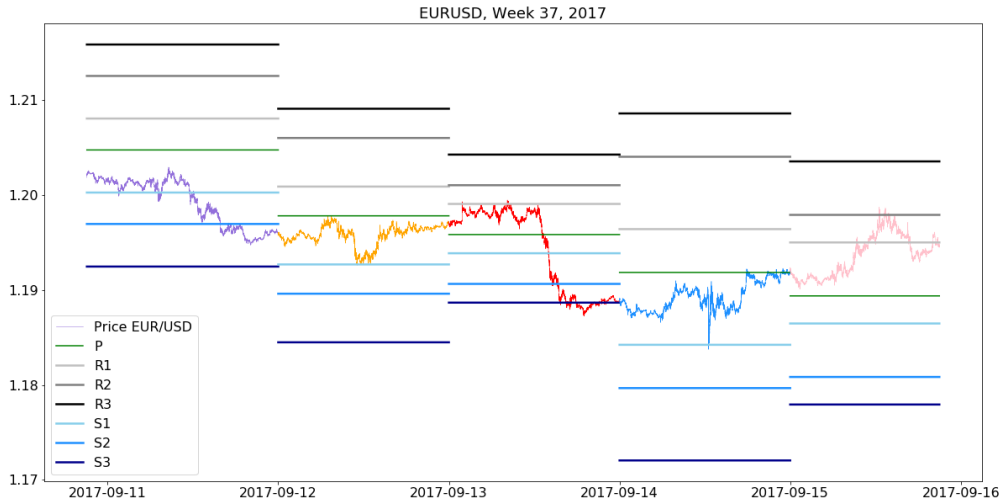

Figure 5: The figure shows the pivot points-technique implemented in the EUR/USD currency pair over week 37, 2017, after the Sunday adjustment (explained above). Due to the Sunday adjustment, it can be seen in the figure that the first trading day's session (i.e. the Monday's period) is longer compared to the other trading days' sessions during the week, which are all equally long. *Source: Authors' calculations.*

5 Methodology

This chapter presents the methodology used to satisfy the purpose of this thesis. Section 5.1 describes the approach taken in order to evaluate the FX market efficiency using a specific trading strategy. Section 5.2 explains the binomial test and the likelihood ratio (LR) test of independence. In Section 5.3, another approach frequently used to evaluate technical trading strategies is explained, namely the bootstrap resampling method.

5.1 Trading Strategy Approach

The first step in this study is to apply a trading strategy approach that mimics the actions of a trader according to a pre-determined trading strategy with the aim of establishing whether or not this strategy gives rise to exploitable profit opportunities, which could be seen as evidence against market efficiency. In other words, if a consistently profitable trading strategy can be developed based on historical prices, the market cannot by definition be efficient since all forms of EMH are rejected. Our trading strategy uses the pivot points-technique to define three tiered pivot support and resistance levels respectively. The tiered pivot support and resistance levels, also referred to as the pivot levels, are given by equations $(2) - (7)$ and depend on the calculated pivot points, which in turn are determined by the previous day's prices. In this study, a new day starts at UTC±0.

The strategy itself is straight forward; if the price hits a support level, the price is expected to increase, and a long position is thus taken. If the price instead hits a resistance level, the price is expected to decrease, and a short position is taken. For simplicity, we are not taking any transaction costs or risk-free rates into account. The purpose of this thesis is to investigate if pivot points-technique can help predict intraday interruptions for exchange rates. If we are able to show that it can, further research will be needed to investigate how the technique best can be implemented to generate maximal risk-adjusted profits. The trading strategy is illustrated in Figure 6.

Figure 6: In this graphical example of the GBP/JPY during the 27th March 2003, the price first hits S_1 at point A, which is assumed to indicate buying pressure and the strategy is thus to go long the currency pair at this level. The dashed lines (in similar colors as S_1) above and below S_1 are the target level and stop-loss level respectively. At point B, the price hits the target level which means that the position is closed, and a profitable trade is registered. After point B, the price decreases. No new trade is taken when the price again hits S_1, because a certain pivot level can individually only generate one trade per day. As the price hits S_2 at point C, a new trade is taken. At point D, the price hits the target level and the position is closed, again with a profitable outcome. A final trade is taken as the price hits S_3 at point E, but this time it leads to a loss as the price hits the stop-loss level at point F. The ranges of the target and stop-loss levels are here 0.20% of the pivot point (P_t) which for the three trades correspond to holding periods between 1 to 3 hours. *Source: Authors' calculations.*

In order to decide when to close a position, the strategy needs certain trading rules. For this matter, we construct a target level and stop-loss level for each pivot level. These levels are determined as a fraction of the pivot point (P_t) for a given time t. For instance, if the price hits the first support level (S_1), the target level is ($S_1 + P_t \times X$) and the stop-loss level is ($S_1 - P_t \times X$), where $P_t \times X$ is a fraction of the pivot point. Obviously, if the price hits a support level, the target level is above the support level, and if the price hits a resistance level, the target level is below the resistance level. After discussions with Century Analytics, we have deemed it appropriate to define an interval such that it takes around two to three hours for the price to hit the target or stop-loss level. An interval of that length is considered suitable for intraday trading strategies for Century Analytics and potentially other hedge funds as well. Due to this, we test a number of different fractions in order to come close such interval. In addition, trying a number of different fractions allows us to examine the strategy more extensively. Worth noting is that there are quite substantial differences between the ratios of different currency pairs. For instance, the EUR/USD has a mean of 1.17 while the GBP/JPY has a mean of 170.39. Therefore, different fractions are expected to impact the currency pairs differently.

Each time the price hits one of the pivot levels, an equal-sized trade is taken. If the price subsequently hits the target level, the trade is registered as profitable. Instead, if the price hits a stop-loss level, the trade is registered as unprofitable. Since the distance between any pivot level and the target level is the same as the distance between any pivot level and the stop-loss level, the price should, in the abstract, hit the target level and stop-loss level equally often if the market is efficient. Consequently, the win rate, defined as the number of profitable trades divided by the number of total trades, should be very close to 50%. However, the tendency of the price to bounce toward the target level in actual data can potentially slightly exceed this benchmark for two reasons. First, changes in actual data have a fairly strong negative first-order autocorrelation, as noted by Goodhart & Figliuoli (1991). Second, the price can continue trending slightly after officially hitting the pivot level and still be considered as having "bounced" (Osler, 2000). Despite these two arguments we will consider a win rate of 50% as our benchmark. The trading strategy is considered profitable if the win rate is significantly higher than 50%, thus the market can be seen as inefficient. On the contrary, if the win rate is significantly less than 50%, the strategy can instead be implemented inversely, and the market can again be seen as inefficient.

Due to the substantial amount of data, the trading strategy approach is executed in Python to facilitate the calculations. For each day, the data is filtered to only contain observations for the day of interest. The program then determines if the price hits any pivot level, and if it does, it determines which support and/or resistance levels the price hits. Worth mentioning is that the price often hits more than one pivot level during a trading day (see for example Figure 2), but a certain pivot level, e.g. S_1 and/or R_2, can individually only generate one trade per day. Then, the data is again filtered to only include trades that were executed after the hit. Thereafter, it is determined if the price hits the target level or stop-loss level. If the price hits both levels, the timestamp of the both hits are extracted and compared, and the relevant level is the one that occurred first. Once the program has determined which pivot levels the price hits, and if the price followingly hits the target level or the stop-loss level, the trade is classified as either profitable or unprofitable.

For some days, the price hits a pivot level during the end of the day, say at $t = 0$, but neither the target nor the stop-loss level is hit later during that day. In such case, the program extracts the prices for the next day at $t = 1$ as well and determines if a target level or a stop-loss level is hit during that day, using the calculated pivot levels for $t = 0$. One exception is if the price

hits a pivot level during the end of a Friday. In such scenario, the program does not search the next Monday, but instead closes the trade at the end of the day and the trade is considered profitable if the closing price is closer to target level than to the stop loss level. Else, the trade is considered unprofitable.

5.2 Binomial Test and LR Test of Independence

The trading strategy approach, as explained in section 5.1, needs to be further analyzed in order to draw any sophisticated conclusions. Even if the trading strategy turns out to be consistently profitable, it is necessary to apply some kind of statistical tests to rule out the possibility of luck. In this thesis, we apply two statistical tests which we find suitable, namely the binomial test and the likelihood ratio (LR) test of independence.

The (two-sided) binomial test compares a sample proportion to a presumed value. Since each simulated trade is categorized as either profitable or unprofitable, the results can easily be translated into a binary series. The binary series can be illustrated by a series of zeroes and ones, where the zeroes represent unprofitable trades and the ones represent profitable trades. This kind of sequence follows a binomial distribution which means that the theoretical distribution is known. The binomial distribution is defined as (see for example; Hull (2015:271))

$$P(B = k) = \binom{n}{k} p^k (1-p)^{n-k} \tag{8}$$

where n is the sample size, k is the number of profitable trades and p is the probability of a profitable trade. Note that $\binom{n}{k}$ is a shorthand for $\frac{n!}{k!(n-k)!}$.

The EMH asserts that a speculative strategy based on historical prices should be profitable in 50% of the times. Consequently, the win rate of the trading strategy developed using the pivot point-technique can be compared to the expected win rate of 50%, which means that the null hypothesis of the binomial test is that the probability of observing k wins from a sample of n trades is 50% (or 0.5).

$$H_0: p = 0.5 \tag{9}$$

$$H_1: p \neq 0.5 \tag{10}$$

The binomial test is also executed in Python. If there are large, significant deviations from the expected win rate, statistical conclusions based on the binomial test can be drawn regarding consistent profitability of the strategy.

A drawback of the binomial test is that it only compares a sample proportion to a presumed value on an aggregate level. For instance, the win rate could be 75% for the first half of the sample period and 25% for the last half, which would still result in an aggregate win rate of 50%. If that is the case, there might be specific periods for which the strategy performs better (or worse). Therefore, it is of interest to determine if the win rates are independent of each other. Since the outcome of the trading strategy can be represented by a binary series, the likelihood-ratio (LR) test of independence proposed by Christoffersen (1998) is a suitable approach to test the win rate independency.

The LR test proposed by Christoffersen (1998) was originally intended for testing independence of value-at-risk breaches. However, the test is applicable to any binary series. The intuition behind the test is to examine if it is more likely to observe a profitable trade at time t if the trade at time $t-1$ was profitable. Or the other way around, is it more likely to observe an unprofitable trade at time t if the trade at time $t-1$ was unprofitable. The LR test is defined as:

$$LR = -2 * \log\left[\frac{L\left(\widehat{\prod}_2; I_1, I_2, \dots, I_T\right)}{L\left(\widehat{\prod}_1; I_1, I_2, \dots, I_T\right)}\right] \tag{11}$$

where

$$\widehat{\prod}_1 = \begin{bmatrix} \frac{n_{00}}{n_{00}+n_{01}} & \frac{n_{01}}{n_{00}+n_{01}} \\ \frac{n_{10}}{n_{10}+n_{11}} & \frac{n_{11}}{n_{10}+n_{11}} \end{bmatrix} \tag{12}$$

and

$$\widehat{\prod}_2 = \hat{\pi}_2 = (n_{01}+n_{11})/(n_{00}+n_{10}+n_{01}+n_{11}) \tag{13}$$

$\widehat{\prod}_2$ is the ML estimate, which under the null hypothesis can be represented by a first-order Markov chain. Furthermore,

$$\prod_2 = \begin{bmatrix} 1-\pi_2 & \pi_2 \\ 1-\pi_2 & \pi_2 \end{bmatrix} \tag{14}$$

$\prod_2$ corresponds to independence. The likelihood under the null hypothesis becomes

$$L(\prod_2; I_1, I_2, \dots, I_T) = (1-\pi_2)^{n_{00}+n_{10}} {\pi_2}^{n_{01}+n_{11}} \tag{15}$$

The $I_1, I_2, \dots, I_T$ represents the binary series and the n_{ij} corresponds to how many times i is followed by j in the binary series. The test statistic is asymptotically distributed as χ^2 with $(s-1)^2$ degrees of freedom. Since the test concerns a binary series, this implies that s equals 2.

5.3 Bootstrap Resampling

To further examine if pivot points can help predict intraday interruptions for exchange rates, we apply a bootstrap resampling method. In this study, we are interested in the behavior of the price once it hits any of the pivot levels. Therefore, we examine the price evolutions from the time the price hits one of the pivot levels and three hours forward. To illustrate, assume that the price hits the first support level (S_1) at 14:05 for a given day. Then, the rest of price observations between 14:05 and 17:05 for that particular day are extracted and saved separately. Note that this is done for all pivot levels. This leads to six different groups, one per pivot level, of price evolutions for each currency pair. For example, one group in which the price hits S_1 and the observed price evolutions for the following three hours, one group in which the price hits R_1 and the observed price evolutions for the following three hours, and so on. Figure 7 shows the EUR/USD price evolutions for R_1.

Followingly, each group is compared to randomly selected price series in order to examine whether the price seems to behave differently once it hits the different pivot levels. To perform the comparison, 1,000 three-hour intervals were randomly selected from the empirical distribution using the bootstrap resampling technique. Further, we calculate a statistic for each bootstrap resample and use the distribution of the simulated statistics to approximate characteristics of the population.

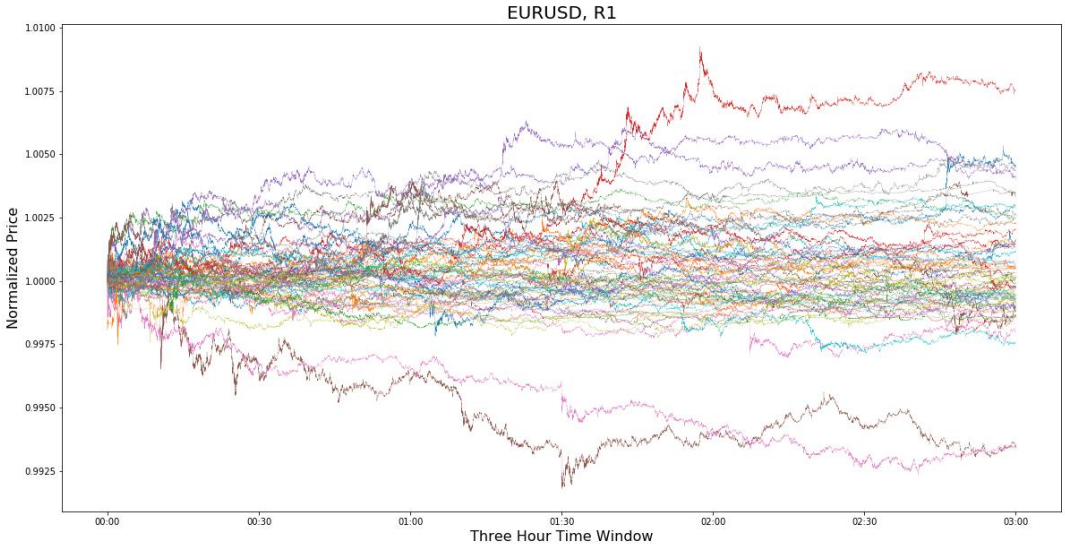

Figure 7: The graph shows all the normalized price evolutions for the EUR/USD from the point when the price hits R_1 and three hours forward. By assuming that the pivot points-technique works, one would expect the price to decrease more often than increase after hitting the first resistance level (R_1). If that would actually be the case, one would observe an intensive cluster of price evolutions below the normalized price of 1, which does not seem to be the case by looking at the figure. In fact, it seems to be the other way around. *Source: Authors' calculations.*

The bootstrap resampling method generates the desired number of new price series by randomly choosing 1,000 of single price observations from the original data set. The timestamp for each observation is extracted and the appropriate time interval is created by adding three hours to the timestamp and indexing the data based on the defined interval. This results in 1,000 randomly generated pseudo exchange rates which can be compared to the actual price evolutions of each group.

The comparison itself can be performed in a number of ways. In this study, the moments of the randomly generated price series are compared to the moments of the actual price series for each level and for each currency pair. For the three currency pairs, it takes on average 10.0 – 10.5 hours to hit either R_1 or S_1. Since there might be intraday volatility patterns, it is relevant to ensure that the randomly generated series are extracted from a similar time. This is done by using the timestamps from the days that each pivot level is hit and filtering the empirical data using the extracted timestamps. This means that the randomly generated data can be taken from any day, but the filtering makes sure that the resampled price evolutions are from a similar time interval.

6 Results and Analysis

In this chapter, results using the stated methodology are presented. Section 6.1 presents and evaluates the results from the trading strategy approach for each of the three currency pairs. Section 6.2 presents and evaluates the results generated by the bootstrap resampling method. In section 6.3, the results and analysis are summarized and further discussed.

6.1 Trading Strategy Approach

6.1.1 EUR/USD

Figure 8 shows the number of times the price hits each of the pivot levels during the 152 trading days in the sample period for the EUR/USD. As expected, S_1 and R_1 are hit the most while the price rarely hits S_3 or R_3. In fact, the price only hits S_3 a total of 7 times. The corresponding number for R_3 is 14. In Figure 8, it can also be seen that the price hits the resistance levels more often than the support levels. A potential explanation to that is because the EUR tended to appreciate against the USD during the sample period (see Figure 3), implying an upwards pressure.

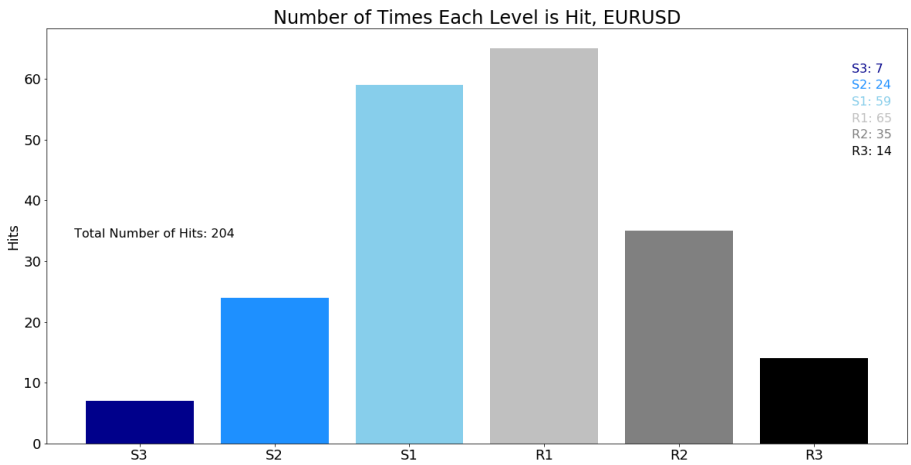

Figure 8: This figure shows the number of times each pivot level is hit for the EUR/USD over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The price hits the inner pivot levels, that is S_1 and R_1, more often compared to the outer pivot levels. *Source: Authors' calculations.*

The outcome of the trading strategy based on pivot points applied on the EUR/USD is presented in Table 2.

Table 2: This table summarizes the outcome of the trading strategy based on pivot points applied on the EUR/USD over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. The columns S_1 through R_3 correspond to the win rates for each pivot level, given the multiple of P_t in the first column. Moreover, the weighted win rate is reported for all multiples, calculated as the win rate for each pivot level multiplied by the number of times the pivot level is hit in total. The two rightmost columns contain the win rates if the strategy is implemented for only support levels or only for resistance levels respectively. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the binomial test.

Multiple of P_t (%)	S1	S2	S3	R1	R2	R3	Weighted Win Rate	Weighted Win Rate (S Only)	Weighted Win Rate (R Only)
0.01	45.8%	33.3%	100.0% **	30.8% ***	42.9%	50.0%	41.2% **	46.7%	36.8% ***
0.05	52.5%	50.0%	85.7%	41.5%	51.4%	21.4% *	47.5%	54.4%	42.1%
0.10	47.5%	54.2%	71.4%	41.5%	48.6%	21.4% *	45.6%	51.1%	41.2% *
0.15	45.8%	45.8%	28.6%	38.5% *	47.1%	21.4% *	41.4% **	44.4%	39.0% **
0.20	50.8%	45.8%	28.6%	36.7% **	44.1%	21.4% *	41.8% **	47.8%	37.1% ***
0.25	50.0%	54.2%	28.6%	28.1% ***	42.4%	28.6%	40.0% ***	49.4%	32.5% ***

The success of the strategy greatly depends on the multiple and pivot level. Moreover, the holding period is directly related to the multiple of P_t (see Appendix A, Table 11). For instance, given a multiple of 0.01%, the average holding period is merely 17 seconds whereas the average holding period for a multiple of 0.25% is about 7.5 hours. To achieve the desired holding period of 2-3 hours, a multiple of around 0.15% is necessary for the EUR/USD. However, given that multiple, no win rate above 50% is accomplished.

For each level and multiple, a binomial test is performed to determine whether the win rate is statistically different from 50%. In Table 2, it can be seen that there are large variations across the win rates, and rather few of them are individually significant. The highest win rate (100%), which is significant, is reported for S_3 with a multiple of 0.01%. However, we bear in mind that this specific level only includes 7 observations, which suggests that the result should be interpreted with great care. In addition, a multiple of 0.01% corresponds to an average holding period of a few seconds, making such strategy somewhat problematic to implement.

At the aggregate level, the strategy underperforms for a majority of the multiples. Even worse results are found if the strategy is implemented only for the resistance levels, implying that

profits can be generated by implementing the strategy inversely. Not only are the weighted win rates for resistance levels low, all of them except one are significant as well. At the individual level, i.e. for each pivot level, the most interesting results are the win rates for R_1. Given a multiple of 0.15% or more, the win rate for R_1 is significantly lower than 50%. While the binomial tests for R_3 are significant at a 10% level for a majority of the multiples as well, one should keep the small number of R_3 hits in mind.

For the LR test, the binary series of all profitable (ones) and unprofitable trades (zeros) is tested for independence. The test is performed with respect to the multiples of P_t and the results are presented in Table 3. Only the multiple of 0.05% indicates an independence issue in the binary series for the EUR/USD.

Table 3: This table summarizes the results of the LR test of independence for the EUR/USD over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the LR test.

Multiple of P_t (%)	***Test Stat.***	***P-value***
0.01	1.22	0.27
0.05	4.28	0.04 **
0.10	0.03	0.86
0.15	0.13	0.72
0.20	1.39	0.24
0.25	0.54	0.46

6.1.2 USD/CHF

Figure 9 illustrates how many times the price hits the pivot levels for the USD/CHF during the 152 trading days. Again, there are rather few hits for the outer pivot levels. During the sample period, the CHF tended to appreciate against the USD (see Figure 3), which could explain why the price hits the resistance levels more often than the support levels.

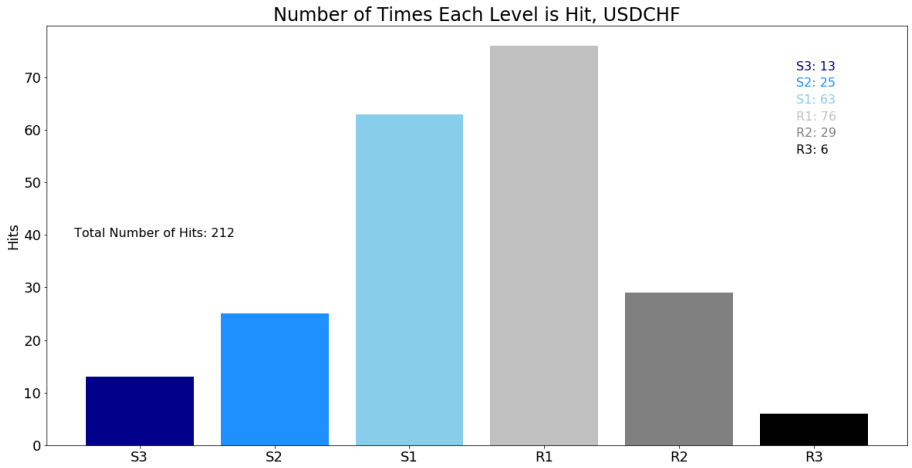

Figure 9: This figure shows the number of times each pivot level is hit for the USD/CHF over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). As also shown for the EUR/USD in Figure 8, the price hits the inner pivot levels, that is S_1 and R_1, more often compared to the outer pivot levels. *Source: Authors' calculations.*

The outcome of the trading strategy based on pivot points applied on the USD/CHF is presented in Table 4.

Table 4: This table summarizes the outcome of the trading strategy based on pivot points applied on the USD/CHF over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. The columns S_1 through R_3 correspond to the win rate for each support and resistance level, given the multiple of P_t in the first column. Moreover, the weighted win rate is reported for all multiples, calculated as the win rate for each pivot level multiplied by the number of times the pivot level is hit in total. The two rightmost columns contain the win rates if the strategy is implemented for only support levels or only for resistance levels respectively. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the binomial test.

Multiple of P_t (%)	S1	S2	S3	R1	R2	R3	Weighted Win Rate	Weighted Win Rate (S Only)	Weighted Win Rate (R Only)
0.01	38.1% *	28.0% **	46.2%	43.4%	37.9%	50.0%	39.6% ***	36.6% ***	42.3%
0.05	42.9%	40.0%	53.8%	47.4%	62.1%	16.7%	46.7%	43.6%	49.5%
0.10	49.2%	44.0%	69.2%	48.7%	65.5%	16.7%	50.9%	50.5%	51.4%
0.15	54.8%	44.0%	46.2%	51.3%	65.5%	50.0%	53.1%	51.0%	55.0%
0.20	55.7%	40.0%	50.0%	52.6%	65.5%	50.0%	53.6%	51.1%	55.9%
0.25	53.4%	36.0%	33.3%	48.6%	55.2%	16.7%	47.6%	46.5%	48.6%

Compared to the EUR/USD, the win rates tend to be generally higher for the USD/CHF, but there are only 2 cases for which the win rates are significantly different from 50%. In these 2 cases, the multiple is 0.01%, implying holding periods of seconds. Again, a multiple of around 0.15% would correspond to a holding period of 2-3 hours (see Appendix A, Table 12). With

this multiple, the win rate for S_1, R_1 and R_2 are above 50%, R_3 is exactly at 50% while S_2 and S_3 are lower than 50%. Weighting the win rates by the number of hits for the corresponding pivot level results in a win rate of approximately 53%. However, the binomial test is insignificant, implying that the outcome might be due to luck. If the strategy is implemented for only resistance levels and weighted by the number of hits using a multiple of 0.15%, the win rate is 55%, but the binomial test is once again insignificant.

The results of the LR test of independence are presented in Table 5. For the USD/CHF, there is no indications that unprofitable trades (zeros) and profitable trades (ones) come clustered together in a time-dependent fashion.

Table 5: This table summarizes the results of the LR test of independence for the USD/CHF over the period from 1th June 2017 to 29th December 2017 (31 weeks or 152 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the LR test.

Multiple of P_t (%)	***Test Stat.***	***P-value***
0.01	2.38	0.12
0.05	2.50	0.11
0.10	1.48	0.22
0.15	1.84	0.17
0.20	1.58	0.21
0.25	0.13	0.71

6.1.3 GBP/JPY

In Figure 10, the number of times each pivot level is hit for GBP/JPY during the 3,839 trading days is illustrated. Since the data we use for the GBP/JPY is of the one-minute type instead of tick type, there are significantly more hits for all levels compared to the other two currency pairs. The depreciation of the JPY against the GBP during the sample period (see Figure 3) might explain why there are a bit more hits for the support levels compared to the resistance levels.

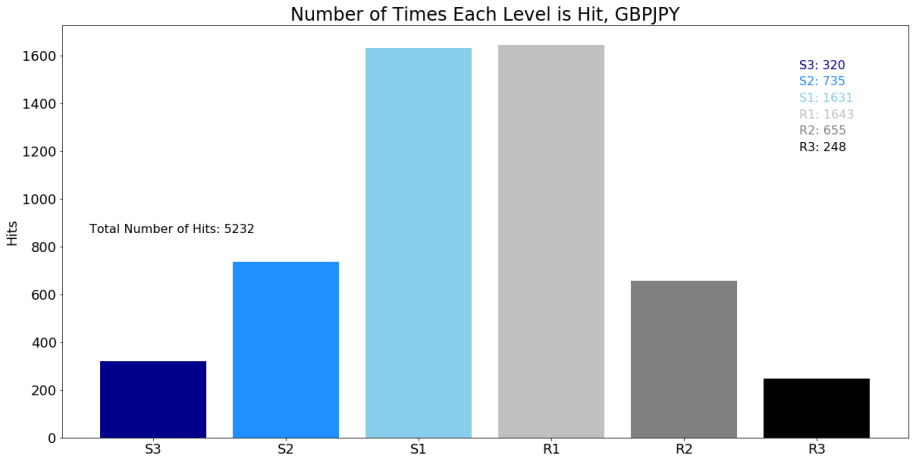

Figure 10: This figure shows the number of times each pivot level is hit for the GBP/JPY over the period from 2nd January 2003 to 13th October 2017 (779 weeks or 3,839 trading days). As also shown for the EUR/USD and the USD/CHF, the price hits the inner pivot levels, that is S_1 and R_1, more often compared to the outer pivot levels. *Source: Authors' calculations.*

In Table 6, the outcome of the trading strategy for the GBP/JPY is summarized. Due the large number of days in combination with deviations from 50%, almost all of the win rates pass the binomial test. In fact, only two of the win rates are insignificant. The weighted win rate for the strategy is significant at a 1% level for all multiples of P_t, indicating that the strategy can be implemented in reverse to generate profits. Worth noting is that the win rate is less than 50% for all pivot levels regardless of multiple.

Table 6: This table summarizes the outcome of the trading strategy based on pivot points applied on the GBP/JPY over the period from 2nd January 2003 to 13th October 2017 (779 weeks or 3,839 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. The columns S_1 through R_3 correspond to the win rate for each support and resistance level, given the multiple of P_t in the first column. Moreover, the weighted win rate is reported for all multiples, calculated as the win rate for each pivot level multiplied by the number of times the pivot level is hit in total. The two rightmost columns contain the win rates if the strategy is implemented for only support levels or only for resistance levels respectively. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the binomial test.

Multiple of P_t (%)	S1	S2	S3	R1	R2	R3	Weighted Win Rate	Weighted Win Rate (S Only)	Weighted Win Rate (R Only)
0.01	32.7% ***	34.0% ***	36.6% ***	35.8% ***	33.1% ***	32.3% ***	34.2% ***	33.5% ***	34.8% ***
0.05	40.4% ***	40.3% ***	39.1% ***	43.4% ***	40.5% ***	40.3% ***	41.3% ***	40.2% ***	42.4% ***
0.10	43.5% ***	40.7% ***	43.0% **	43.8% ***	41.8% ***	37.7% ***	42.7% ***	42.7% ***	42.7% ***
0.15	44.4% ***	42.3% ***	43.0% **	44.3% ***	43.0% ***	35.8% ***	43.4% ***	43.7% ***	43.1% ***
0.20	44.0% ***	44.8% ***	42.6% ***	44.1% ***	43.5% ***	40.7% ***	43.8% ***	44.1% ***	43.6% ***
0.25	44.5% ***	47.3%	45.7%	45.3% ***	42.4% ***	38.2% ***	44.7% ***	45.4% ***	43.9% ***

The best strategy for GBP/JPY would be to implement the trading strategy inversely with a multiple of 0.01%. Since this currency pair historically trades at a considerable higher exchange rate compared to the EUR/USD and the USD/CHF, the range between each target level and stop-loss level is higher as well. Consequently, the average holding period for the GBP/JPY with a multiple of 0.01% is 22 minutes, which may be viewed as a more manageable time frame compared to a couple of seconds. Nonetheless, to achieve the desired time horizon of 2-3 hours, a multiple of between 0.15% and 0.20% is necessary for the GBP/JPY (see Appendix A, Table 13).

The results of the LR test of independence is presented in Table 7. Compared to the LR test results of the EUR/USD and the USD/CHF, the LR test results of the GBP/JPY indicate severe clustering issues for the binary series of profitable/unprofitable trades.

Table 7: This table summarizes the results of the LR test of independence for the GBP/JPY over the period from 2nd January 2003 to 13th October 2017 (779 weeks or 3,839 trading days). The leftmost column is the multiple of the pivot point, i.e. the percentage of P_t that determines the range of the target and stop-loss level. ***/**/* correspond to the significance levels 1%/5%/10% respectively from the LR test.

Multiple of P_t (%)	***Test Stat.***	***P-value***
0.01	4.43	0.04 **
0.05	7.05	0.01 ***
0.10	4.24	0.04 **
0.15	8.74	0.00 ***
0.20	10.82	0.00 ***
0.25	13.65	0.00 ***

6.2 Bootstrap Resampling

6.2.1 EUR/USD

In Table 8, some descriptive statistics for the moments of the actual EUR/USD price series and the moments of the randomly generated price series are presented. In the actual data, the mean is below 1 for the support levels while the mean is above 1 for the resistance levels. This is

contrary to what is expected from the hypothesis. If prices would tend to increase (decrease) as the price hits a support (resistance) level, the means of the actual data would be above (below) 1; not the other way around. For all pivot levels, the variance is higher for the actual data compared to the resampled data. Potentially, this implies that there is some kind of market reaction when prices hit pivot levels compared to when they do not. A comparison of the distributions of changes in the order book for the actual data and the resampled data gives more reasons to believe that is the case (see Appendix B, Figure 11).

Even though the variances for the actual data and the resampled data in Table 8 differ on the sixth or seventh decimal, the price evolutions have been normalized to 1, resulting in lower variances overall. Moreover, the absolute value of the skewness is greater for the actual data. There are no major differences between the kurtosis of the actual and resampled data. In fact, the kurtosis is overall negative in both groups. The price evolutions used for calculating the actual and resampled moments are shown in Appendix C, Figure 13.

Table 8: This table shows some descriptive statistics for the moments of the actual EUR/USD price series and the moments of the randomly generated price series. The actual moments are calculated using the extracted data from when the price hits one of the pivot levels and three hours forward. The resampled moments are calculated using 1,000 randomly selected three-hour intervals to approximate characteristics of the population. The notations $(\times 10^{-6})$ and $(\times 10^{-2})$ mean that the reported values should be multiplied with 10^{-6} and 10^{-2} respectively.

	Actual moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	0.9999	0.8141	-5.1635	-0.3467
S2	0.9999	1.1667	3.5526	-0.4388
S3	0.9996	0.6335	9.6999	-0.5833
R1	1.0004	0.8019	-8.2792	-0.4509
R2	1.0001	0.6845	-2.9262	-0.6737
R3	1.0005	0.8227	-8.1843	-0.4701

	Resampled moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	1.0001	0.5666	0.0632	-0.5613
S2	1.0001	0.5646	0.4116	-0.5596
S3	1.0001	0.4948	-0.0992	-0.5526
R1	1.0000	0.4862	1.4121	-0.5915
R2	1.0001	0.5985	-1.8275	-0.5745
R3	1.0000	0.5526	3.8482	-0.5383

6.2.2 USD/CHF

Table 9 shows some descriptive statistics for the moments of the actual USD/CHF price series and the moments of the randomly generated price series. The means of the resampled data are overall much closer to 1 while the price series that hit some of the pivot levels tend to drift away

from the starting value. Again, the variance of the actual data tends to be somewhat higher. Indeed, Figure 12 in Appendix B shows that there seems to be some kind of market reaction when the prices hit pivot levels compared to when they are not.

In Table 9, it can be seen that the absolute value of the skewness is higher for the actual data. Moreover, the magnitude of the skewness seems to increase the further out the pivot level is from the pivot point. While the sign of the skewness is different for different pivot levels, the sign of the kurtosis is negative for all levels, both in the actual and resampled data. The price evolutions used for calculating the actual and resampled moments are shown in in Appendix C, Figure 14.

Table 9: This table shows some descriptive statistics for the moments of the actual USD/CHF price series and the moments of the randomly generated price series. The actual moments are calculated using the extracted data from when the price hits one of the pivot levels and three hours forward. The resampled moments are calculated using 1,000 randomly selected three-hour intervals to approximate characteristics of the population. The notations $(\times 10^{-6})$ and $(\times 10^{-2})$ mean that the reported values should be multiplied with 10^{-6} and 10^{-2} respectively.

	Actual moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	1.0000	0.7813	-21.0790	-0.1174
S2	0.9999	0.7336	-10.9490	-0.3984
S3	1.0003	0.5529	-27.3139	-0.4099
R1	1.0000	0.6445	0.7243	-0.0026
R2	0.9995	0.5917	16.8634	-0.2224
R3	0.9996	1.1573	-19.5048	-0.5749

	Resampled moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	1.0000	0.5565	0.0945	-0.5036
S2	1.0000	0.5895	-0.1829	-0.4774
S3	1.0000	0.6687	-1.2455	-0.4875
R1	1.0000	0.5756	1.4755	-0.5252
R2	1.0000	0.6571	-2.2295	-0.5123
R3	0.9999	0.5086	-2.0513	-0.2647

6.2.3 GBP/JPY

In Table 10, some descriptive statistics for the moments of the actual GBP/JPY price series and the moments of the randomly generated price series are shown. For the actual means, all support levels are below 1 and all resistance levels are above 1 which is contrary to what is expected from the hypothesis. In line with the other currency pairs, the variance is in general higher for the actual data and tends to increase the further out the pivot level is from the pivot point. Overall, the skewness tends to be more extreme in the actual data. Again, the kurtosis is negative for all levels for both the actual and resampled data. Also, the magnitude of the kurtosis

is quite similar, and no general major differences can be seen. The price evolutions used for calculating the actual and resampled moments are shown in Appendix C, Figure 15.

Table 10: This table shows some descriptive statistics for the moments of the actual GBP/JPY price series and the moments of the randomly generated price series. The actual moments are calculated using the extracted data from when the price hits one of the pivot levels and three hours forward. The resampled moments are calculated using 1,000 randomly selected three-hour intervals to approximate characteristics of the population. The notations $(\times 10^{-6})$ and $(\times 10^{-2})$ mean that the reported values should be multiplied with 10^{-6} and 10^{-2} respectively.

	Actual moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	0.9999	2.1371	-6.5906	-0.2523
S2	0.9999	3.0950	-7.1142	-0.3479
S3	0.9999	9.6329	-11.1353	-0.2140
R1	1.0001	1.5747	3.1556	-0.2629
R2	1.0001	1.7546	-0.5899	-0.2667
R3	1.0001	2.8791	-0.3847	-0.3858

	Resampled moments			
Pivot level	***Mean***	***Variance*** $(\times 10^{-6})$	***Skewness*** $(\times 10^{-2})$	***Kurtosis***
S1	0.9999	1.5492	-5.2325	-0.3815
S2	1.0000	1.6519	-0.9223	-0.3383
S3	1.0001	1.4296	-1.2588	-0.3807
R1	1.0000	1.4141	-3.6982	-0.3692
R2	0.9999	1.7916	-2.2747	-0.4167
R3	1.0001	1.6405	-3.5162	-0.3365

6.3 Summary of Results and Analysis

The results of the trading strategy approach vary between the currency pairs. For two of the currency pairs (EUR/USD & GBP/JPY), almost all win rates are significantly lower than 50%. The win rates for the third currency pair (USD/CHF) seem to be closer to 50%, indicating that the market for this currency pair is more efficient compared to the other pairs. For the EUR/USD, the best outcome would be to implement the trading strategy inversely, but only for the resistance levels. The fact that a trading strategy based on historical prices can be successful for one of the most liquid assets in the world is quite surprising. However, it is worth stressing that there are too few observations for the second and third pivot levels for both the EUR/USD and the USD/CHF to draw any statistical conclusions. For the GBP/JPY, the trading strategy should be implemented inversely, but is about equally successful for support and resistance levels, which implies that a weighted strategy could be used.

Following the results of the LR test of independence, the win rates for the EUR/USD and the USD/CHF are independent for almost all of the multiples. The independence implies that the

zeroes (losses) and ones (gains) are not clustered in a time-dependent manner. On the contrary, the test statistics for the GBP/JPY are significant for all multiples. The significance suggests that there are periods for which the trading strategy is more or less successful and the win rates for this currency pair might therefore be time dependent.

The comparison of the actual moments and resampled moments supports the results of the trading strategy. For the currency pairs with win rates that are statistically lower than 50%, there seems to be downward pressure on the price for support levels, and an upward pressure on the price for resistance levels. This follows from that the mean is lower than one for support levels, and higher than one for resistance levels. The variances of the actual data are consistently higher than the variances of the resampled data for all currency pairs, which might suggest that there indeed is higher activity around the pivot levels. The distributions of changes in the order book shown in Figure 11 and Figure 12 in Appendix B support this observation for the EUR/USD and the USD/CHF. In general, the skewness is larger in absolute value for the actual data but regarding the kurtosis, no major difference can be seen between the two.

7 Conclusion

The purpose of this thesis is to investigate if pivot points help predict intraday interruptions for exchange rates. To this end, we apply both a trading strategy approach and a bootstrap resampling method to three major currency pairs (EUR/USD, USD/CHF & GBP/JPY). The aim of the trading strategy approach is to establish whether or not a certain trading strategy gives rise to exploitable profit opportunities. To statistically rule out the possibility of luck for the trading strategy, we use the binomial test. We also investigate if the win rates generated by the trading strategy are time-dependent by performing the likelihood ratio (LR) test of independence. To complement the analysis, we apply a bootstrap resampling method in order to evaluate the distributional behavior of the exchange rates in comparison with the distribution of thousands of pseudo exchange rates.

The results show no support for the use of pivot points in their traditional way. Instead, for the EUR/USD and the GBP/JPY, the support levels derived from the pivot points should be used as selling signals and the resistance levels should be used as buying signals. In other words, pivot points should be implemented inversely. For the USD/CHF, the results indicate that the pivot points-technique is unable to help predict interruptions in the price, which can be interpreted as this currency pair being more efficient than the other two. Moreover, we find that the success of the inversely defined trading strategy based on pivot points are time dependent for the GBP/JPY, but not for the other currency pairs. Lastly, a comparison between the actual price evolutions that hit the pivot levels and resampled data supports the results of the trading strategy, but also implies that the variance is higher around the pivot levels. The higher variance might imply that there indeed is higher activity in the market as the price fluctuates around the pivot levels, which also Figure 11 and Figure 12 in Appendix B seem to show. Therefore, pivot points might serve as useful indicators for when the market will react and cause increased price volatility.

In an effort to explain the aforementioned results, we find it quite surprising that pivot points appear more useful in an opposite way to what the traditional technique suggests for two of the currency pairs. There can, of course, be many reasons for that. According to Williams (2015), one should keep in mind that it is important to use pivot points in ways that everyone else is

not. Most will look at the pivot points themselves but will fail to use them on larger time frames; that is when one finds new opportunities. Perhaps, we have found a new opportunity to profit from the pivot points-technique in the FX market. This can actually be the case if one would believe the argument by Grossman & Stiglitz (1980), which says that asset price adjustment to information must allow traders who invest resources in processing information to earn a normal rate of return on their investment. However, according to the AMH, the FX market is likely to be highly efficient because participants are competing for rather scarce resources, which implies that we should question our findings.

For the EUR/USD, we investigate a relatively short sample period. If we would extend this period, perhaps the results would differ to what we see now. For the GBP/JPY, on the other hand, we investigate a relatively long sample period. Covering longer sample period increases the risks of other external factors impacting the exchange rates, if not for the whole sample period, at least for certain intervals. Such factors include announcements by central banks, the impact of Brexit, changes in macro indicators, changes in inflation expectations, etc. This can be the explanation to why we see the trading strategy based on pivot points being time dependent for the GBP/JPY, but not for the other two currency pairs. Additionally, we want to emphasize that the trading strategy approach involves a vast amount of approach choices that could have an impact on the results, for better or for worse. The choices we make include trading rules determination, time zone specification for defining a new day, Sunday adjustments, optimal holding periods, number of allowed trades per day, and more. There is a possibility that an optimal trading strategy could improve the predictive power of applied pivot points and give different results to ours. Therefore, it is hard to generalize our findings beyond the scope of this study.

Moreover, we find it surprising that the results indicate that the USD/CHF is more efficient compared to the other two currency pairs. If there would be one currency pair being more efficient than the other two, it would most probably be the EUR/USD since it is the world's most actively traded pair. We argue that it is hard to draw any extensive conclusions from this finding. Our results are not sufficient to determine if a currency pair is more efficient than another. It is reasonable to believe that we would get different results if we, again, would compare the currency pairs investigated in this study over different or extended sample periods.

The ability of the pivot points-technique has, as far as we know, never before been rigorously evaluated. This thesis has undertaken such a test and the overall results are fairly in line with

findings in Curcio et al. (1997) and Osler (2000), which are the closest point of comparison to this study. We feel that we have only scratched upon the surface of this specific TA technique, but yet laid the foundation for future research in this area. In order to further investigate the usefulness of pivot points, and the field of TA in general, we believe it is of prime importance to conduct experiments that consider combinations of pivot points and other technical rules. As with most technical indicators, pivot points should perhaps not be used in isolation, but with combinations of, if not the whole universe of technical rules, at least the most celebrated of those.

References

Allen, H., & Taylor, M.P. (1992). The Use of Technical Analysis in the Foreign Exchange Market, *Journal of International Money & Finance*, vol. 11, no. 3, pp. 304-314.

Bhandari, B. (2012). Trading with Pivot Points, *Futures: News, Analysis & Strategies for Futures, Options & Derivatives Traders*, vol. 41, no. 3, pp. 38-41.

Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple Technical Trading Rules and the Stochastic Properties of Stock Returns, *The Journal of Finance*, vol. 47, no. 5, pp. 1731-1764.

Byström, H. (2004). *Finance – Markets, Instruments & Investments*. 3rd edition. Lund: Studentlitteratur AB.

Campbell, J.Y., Lo, A.W., & MacKinlay, A.C. (1997). *The Econometrics of Financial Markets*. 1st edition. Princeton, NJ: Princeton University Press.

Chang, P.H.K., & Osler, C.L. (1999). Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts, *The Economic Journal*, vol. 109, no. 458, pp. 636-661.

Cheung, Y.W., & Chinn, M.D. (1999). *Traders, Market Microstructure and Exchange Rate Dynamics*, NBER Working Paper No. 7416. Available at: https://www-nber-org.ludwig.lub.lu.se/papers/w7416 (Accessed: 14th April 2019).

Cheung, Y.W., Chinn, M., & Garcia Pascual, A. (2005). Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?, *Journal of International Money & Finance*, vol. 24, no. 7, pp. 1150-1175.

Cheung, Y.W., & Wong, C.Y.P. (2000). A Survey of Market Practitioners' Views on Exchange Rate Dynamics, *Journal of International Economics*, vol. 51, no. 2, pp. 401-419.

Christoffersen, P.F. (1998). Evaluating Interval Forecasts, *International Economic Review*, vol. 39, no. 4, pp. 841-862.

Curcio, R., Goodhart, C., Guillaume, D., & Payne, R. (1997). Do Technical Trading Rules Generate Profits? Conclusions from the Intra-day Foreign Exchange Market, *International Journal of Finance & Economics*, vol. 2, no. 4, pp. 267-280.

Dacorogna, M.M., Gençay, R., Müller, U.A., Olsen, R.B., & Pictet, O.V. (2001). *An Introduction to High-Frequency Finance*. 1st edition. San Diego: Academic Press.

De Long, J.B., Shleifer, A., Summers, L.H., & Waldmann, R.J. (1990). Noise Trader Risk in Financial Markets, *Journal of Political Economy*, vol. 98, no. 4, pp. 703-738.

Dunis, C.L., & Williams, M. (2002). Modelling and Trading the Euro/US Dollar Exchange Rate: Do Neural Network Models Perform Better?, *Derivatives Use, Trading & Regulation*, vol. 8, no. 3, pp. 211-229.

Eiteman, D.K., Stonehill, A.I., & Moffett, M.H. (2010). *Multinational Business Finance*. 12th edition. Boston: Pearson Education, Inc.

Fama, E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work, *The Journal of Finance*, vol. 25, no. 2, pp. 383-417.

Fernández, F.J. (2019). *The Foreign Currency Market: What It is and How It Works*, BBVA. Available at: https://www.bbva.com/en/foreign-currency-market-work/ (Accessed: 1st April 2019).

Goodhart, C., & Figliuoli, L. (1991). Every Minute Counts in Financial Markets, *Journal of International Money & Finance*, vol. 10, no. 1, pp. 23-52.

Gradojevic, N., & Lento, C. (2015). Multiscale Analysis of Foreign Exchange Order Flows and Technical Trading Profitability, *Economic Modelling*, vol. 47, no. 1, pp. 156-165.

Grossman, S.J., & Stiglitz, J.E. (1980). On the Impossibility of Informationally Efficient Markets, *The American Economic Review*, vol. 70, no. 3, pp. 393-408.

Hull, J.C. (2015). *Risk Management and Financial Institutions*. 4th edition. Hoboken, NJ: John Wiley & Sons, Inc.

Kubińska, E., Czupryna, M., Markiewicz, L., & Czekaj, J. (2016). Technical Analysis as a Rational Tool of Decision Making for Professional Traders, *Emerging Markets Finance & Trade*, vol. 52, no. 12, pp. 2756-2771.

Levich, R.M., & Thomas, L.R. (1993). The Significance of Technical Trading-rule Profits in the Foreign Exchange Market: A Bootstrap Approach, *Journal of International Money & Finance*, vol. 12, no. 5, pp. 451-474.

Lo, A.W., Mamaysky, H., & Wang, J. (2000). Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation, *The Journal of Finance*, vol. 55, no. 4, pp. 1705-1765.

Lo, A.W. (2004). The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective, *Journal of Portfolio Management*, vol. 30, no. 1, pp. 15-29.

Lui, Y.H., & Mole, D. (1998). The Use of Fundamental and Technical Analyses by Foreign Exchange Dealers: Hong Kong Evidence, *Journal of International Money & Finance*, vol. 17, no. 3, pp. 535-545.

Maeda, J., & Jacka, S.D. (2018). *Modeling Technical Analysis*, Working Paper No. 1707.05253, arXiv.org. Available at: https://arxiv.org/abs/1707.05253 (Accessed: 4th April 2019).

Malkiel, B.G. (2003). The Efficient Market Hypothesis and Its Critics, *The Journal of Economic Perspectives*, vol. 17, no. 1, pp. 59-82.

Menkhoff, L., & Taylor, M.P. (2007). The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis, *The Journal of Economic Literature*, vol. 45, no. 4, pp. 936-972.

Merton, R.K. (1948). The Self-Fulfilling Prophecy, *The Antioch Review*, vol. 8, no. 8, pp. 193-210.

Murphy, J.J. (1986). *Technical Analysis of the Futures Market: A Comprehensive Guide to Trading Methods and Applications*. 1st edition. New York: Prentice Hall.

Neely, C., Weller, P., & Dittmar, R. (1997). Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach, *The Journal of Financial & Quantitative Analysis*, vol. 32, no. 4, pp. 405-426.

Osler, C. (2000). Support and Resistance: Technical Analysis and Intraday Exchange Rates, *Federal Reserve Bank of New York Economic Policy Review*, vol. 6, no. 2, pp. 53-68.

Park, C.H., & Irwin, S.H. (2007). What Do We Know About the Profitability of Technical Analysis?, *Journal of Economic Surveys*, vol. 21, no. 4, pp. 786-826.

Pring, M.J. (1991). *Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points*. 3rd edition. New York: McGraw-Hill, Inc.

Qi, M, & Wu, Y. (2006). Technical Trading-Rule Profitability, Data Snooping, and Reality Check: Evidence from the Foreign Exchange Market, *Journal of Money, Credit & Banking*, vol. 38, no. 8, pp. 2135-2158.

Roberts, H. (1967). "Statistical versus Clinical Prediction of the Stock Market", Unpublished manuscript, Center for Research in Security Prices, University of Chicago, May.

Sager, M.J., & Taylor, M.P. (2006). Under the Microscope: The Structure of the Foreign Exchange Market, *International Journal of Finance & Economics*, vol. 11 no. 1, pp. 81-95.

Smith, D.M., Wang, N., Wang, Y., & Zychowicz, E.J. (2016). Sentiment and the Effectiveness of Technical Analysis: Evidence from the Hedge Fund Industry, *Journal of Financial & Quantitative Analysis*, vol. 51, no. 6, pp. 1991-2013.

Timmermann, A., & Granger, C.W.J. (2004). Efficient Market Hypothesis and Forecasting, *International Journal of Forecasting*, vol. 20, no. 1, pp. 15-27.

Williams, B. (2015). Pivot Point Trading, *Modern Trader*, pp. 46-49. Available at: http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=109035617&site=eds-live&scope=site (Accessed: 5th April 2019).

Yun, T. (2017). Natural Exchange Rate and Its Implications for the Purchasing Power Parity Puzzle, *Emerging Markets Finance & Trade*, vol. 53, no. 11, pp. 2397-2418.

Zapranis, A., & Tsinaslanidis, P.E. (2012). Identifying and Evaluating Horizontal Support and Resistance Levels: An Empirical Study on US Stock Markets, *Applied Financial Economics*, vol. 22, no. 19, pp. 1571-1585.

Appendix A

Table 11: This table shows the average time in minutes from the point that the price of the EUR/USD hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2\ or\ R_3)$ to the point when it hits either the target level or the stop-loss level for that specific pivot level. In other words, it shows for how many minutes the price of the EUR/USD remains in the interval between the target level and the stop-loss level of a certain pivot level.

	Pivot level					
Multiple of P (%)	**S1**	**S2**	**S3**	**R1**	**R2**	**R3**
0.01	0.5	0.1	0.2	0.2	0.5	0.1
0.05	16.5	12.0	5.5	11.3	13.2	5.8
0.10	68.0	59.0	12.1	73.1	50.1	26.4
0.15	179.6	220.1	158.0	212.0	181.7	66.8
0.20	296.3	385.1	265.1	347.3	301.3	150.8
0.25	506.8	545.6	406.9	522.3	458.0	201.4

Table 12: This table shows the average time in minutes from the point that the price of the USD/CHF hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2\ or\ R_3)$ to the point when it hits either the target level or the stop-loss level for that specific pivot level. In other words, it shows for how many minutes the price of the USD/CHF remains in the interval between the target level and the stop-loss level of a certain pivot level.

	Pivot level					
Multiple of P (%)	**S1**	**S2**	**S3**	**R1**	**R2**	**R3**
0.01	0.7	0.5	0.3	0.6	0.4	0.1
0.05	24.8	10.4	9.3	13.8	12.0	18.0
0.10	60.2	48.2	47.2	77.3	72.3	62.0
0.15	163.3	97.7	217.2	163.2	122.4	137.9
0.20	257.2	144.6	404.1	311.6	211.8	249.8
0.25	390.8	368.3	583.2	497.3	394.8	449.6

Table 13: This table shows the average time in minutes from the point that the price of the GBP/JPY hits the specific pivot level ($S_1, S_2, S_3, R_1, R_2\ or\ R_3$) to the point when it hits either the target level or the stop-loss level for that specific pivot level. In other words, it shows for how many minutes the price of the GBP/JPY remains in the interval between the target level and the stop-loss level of a certain pivot level.

	Pivot level					
Multiple of P (%)	**S1**	**S2**	**S3**	**R1**	**R2**	**R3**
0.01	7.6	16.1	36.6	7.6	18.2	47.3
0.05	15.0	21.9	41.8	18.1	26.7	54.7
0.10	45.1	45.4	67.2	53.5	69.1	98.2
0.15	94.4	92.6	112.2	119.3	127.8	152.6
0.20	160.4	159.8	175.1	205.2	230.3	241.2
0.25	246.2	244.4	255.1	311.2	348.0	363.0

Appendix B

Figure 11: This figure presents the distributions of changes in the order book for the EUR/USD for the actual and resampled data shown in Appendix C, Figure 13. The histograms in the left column show the distributions of the changes in the order book, on an accumulated basis, for all the times the price hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2 \; or \; R_3)$ and three hours forward. The histograms in the right column show the distributions of the changes in the order book, also on an accumulated basis, for the 1,000 price evolutions randomly generated for each of the pivot levels. Note that the Y-axis in each of the histograms in the right column display units of 10 million changes. *Source: Authors' calculations.*

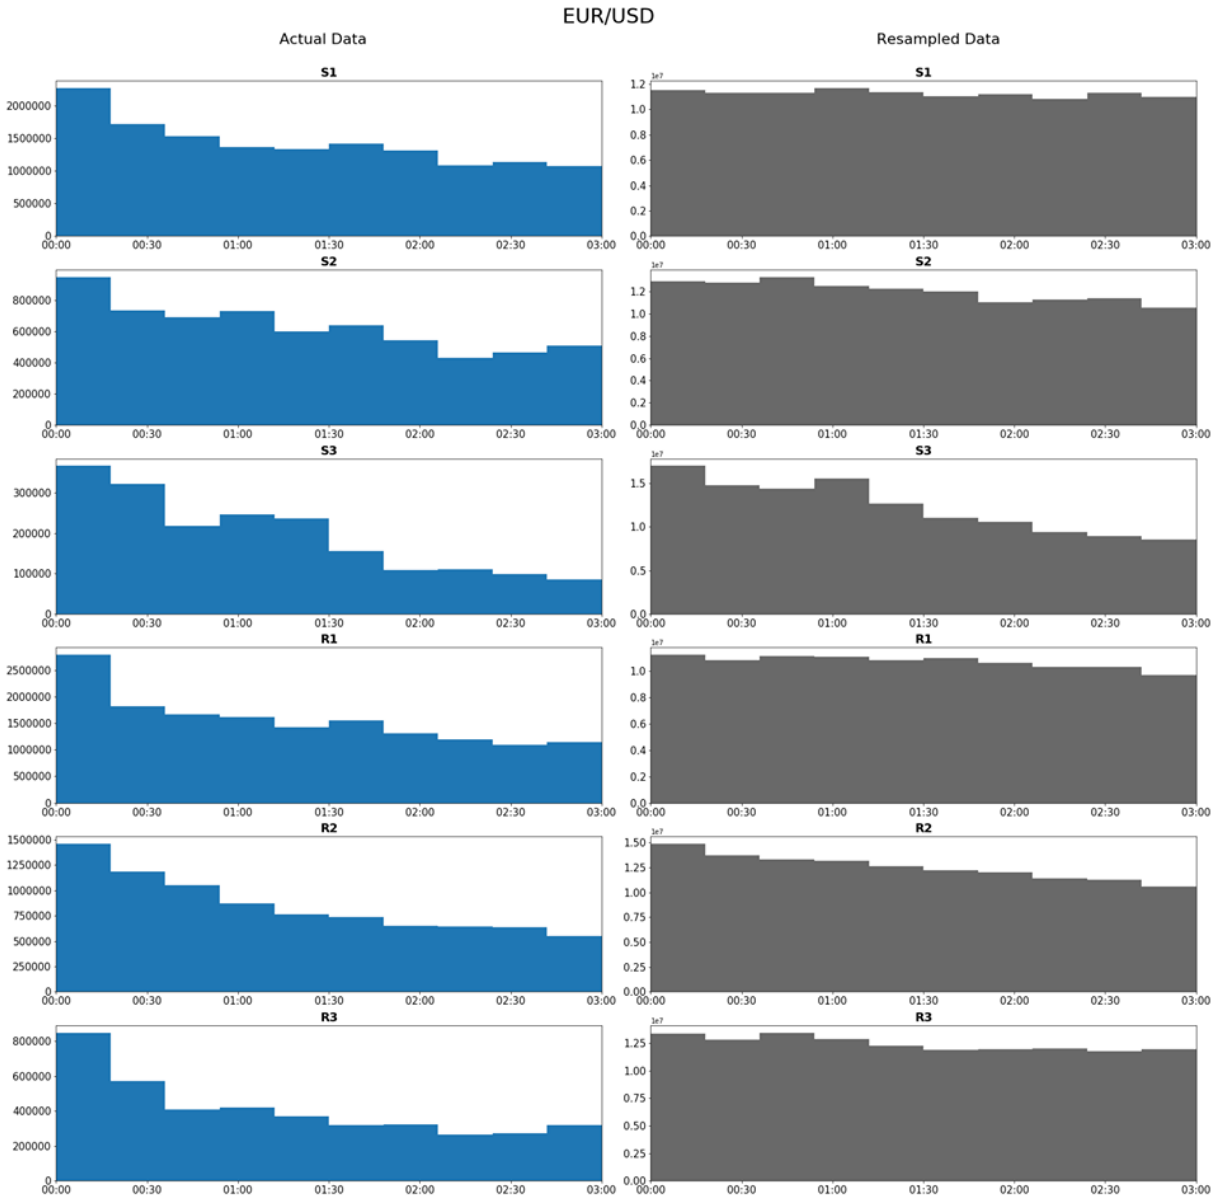

Figure 12: This figure presents the distributions of changes in the order book for the USD/CHF for the actual and resampled data shown in Appendix C, Figure 14. The histograms in the left column show the distributions of the changes in the order book, on an accumulated basis, for all the times the price hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2 \; or \; R_3)$ and three hours forward. The histograms in the right column show the distributions of the changes in the order book, also on accumulated basis, for the 1,000 price evolutions randomly generated for each of the pivot levels. Note that the Y-axis in the first and fourth histograms in the right column display units of 1 million changes. The Y-axis in the rest of the histograms in the right column display units of 10 million changes. *Source: Authors' calculations.*

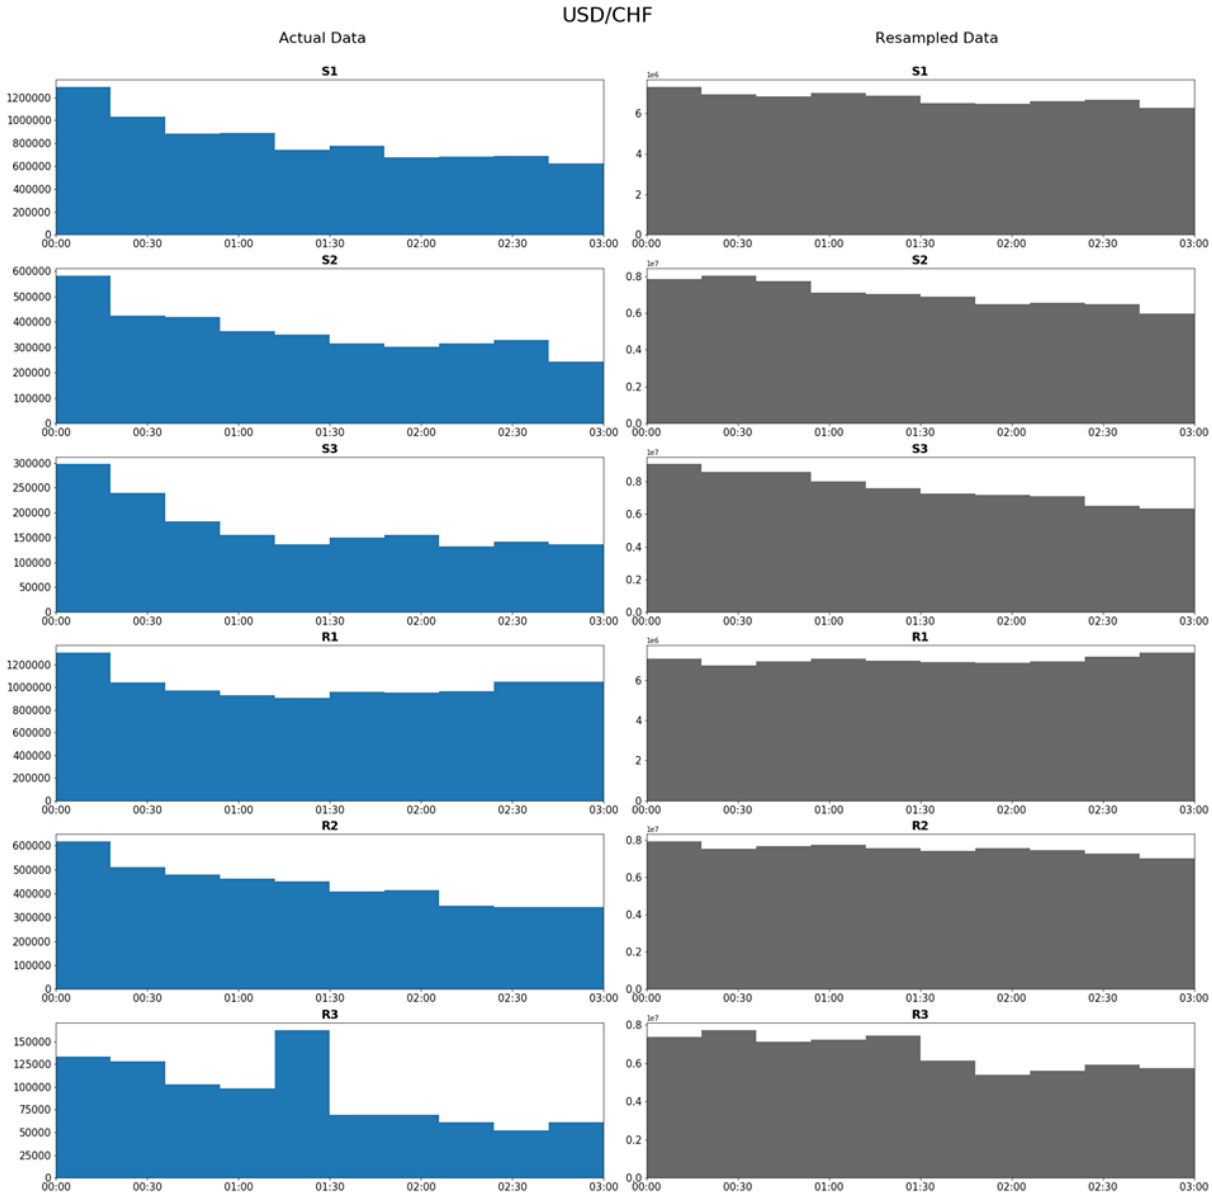

Appendix C

Figure 13: This figure provides an illustration of the bootstrap resampling method for the EUR/USD. The left column shows all the actual price evolutions from the time the price hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2 \; or \; R_3)$ and three hours forward. This leads to six different groups, one per each pivot level, of actual price evolutions. The right column shows 1,000 price evolutions, randomly generated for each of the pivot levels. The generated series are extracted from similar times as the actual price evolutions. The Y-axes in each of the twelve graphs show the normalized price and the X-axes show a three-hour time window. *Source: Authors' calculations.*

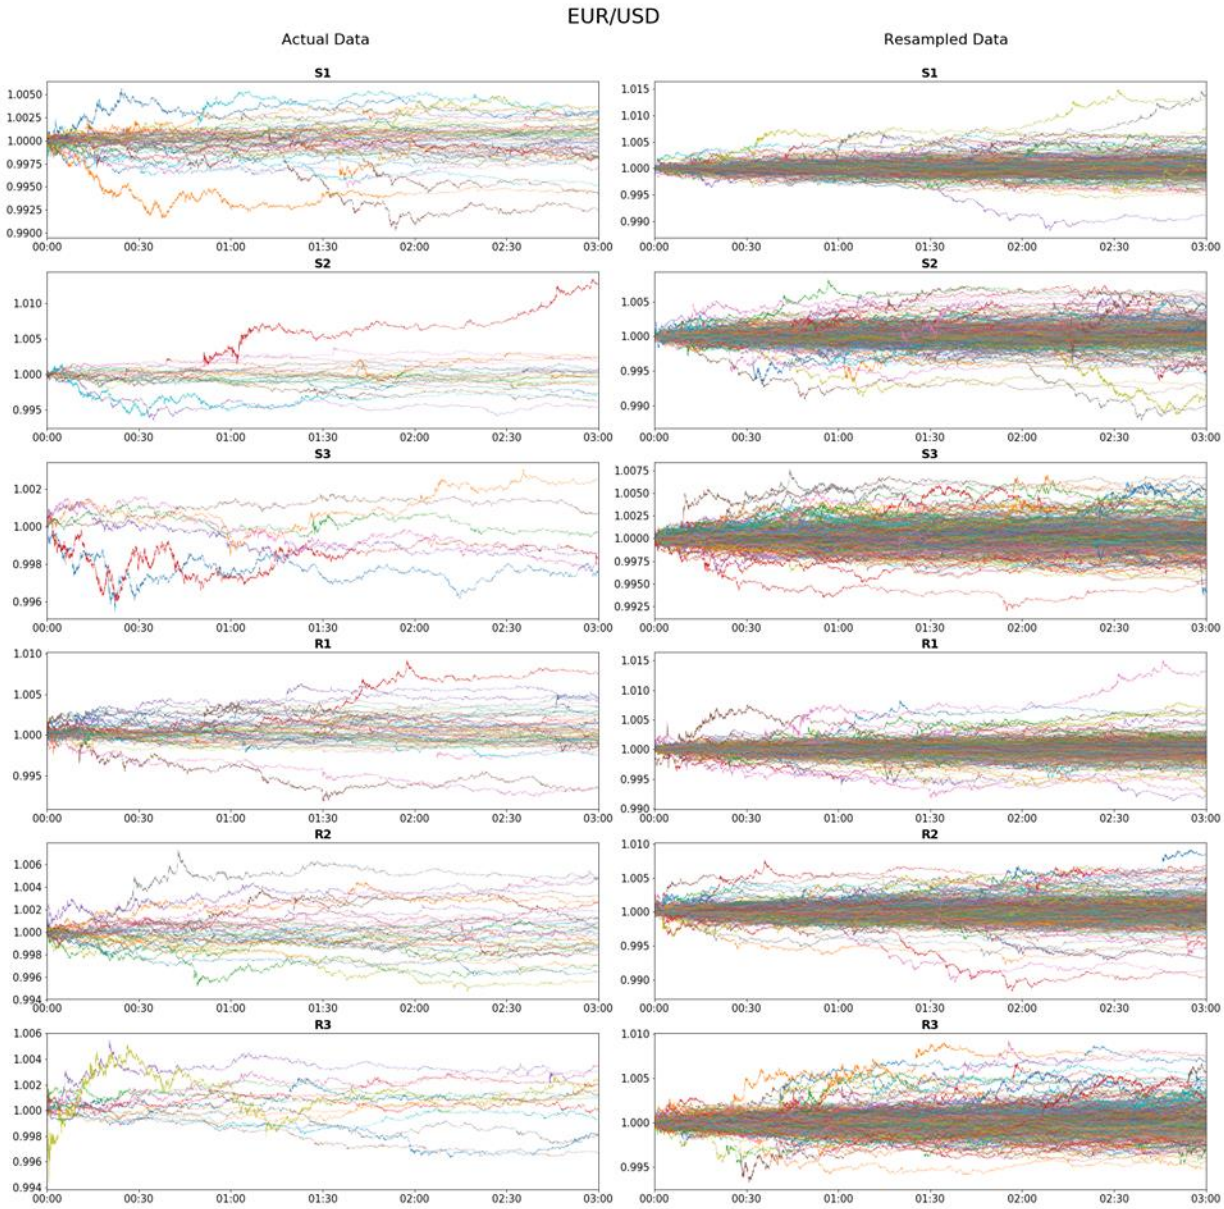

Figure 14: This figure provides an illustration of the bootstrap resampling method for the USD/CHF. The left column shows all the actual price evolutions from the time the price hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2 \; or \; R_3)$ and three hours forward. This leads to six different groups, one per each pivot level, of actual price evolutions. The right column shows 1,000 price evolutions, randomly generated for each of the pivot levels. The generated series are extracted from similar times as the actual price evolutions. The Y-axes in each of the twelve graphs show the normalized price and the X-axes show a three-hour time window. *Source: Authors' calculations.*

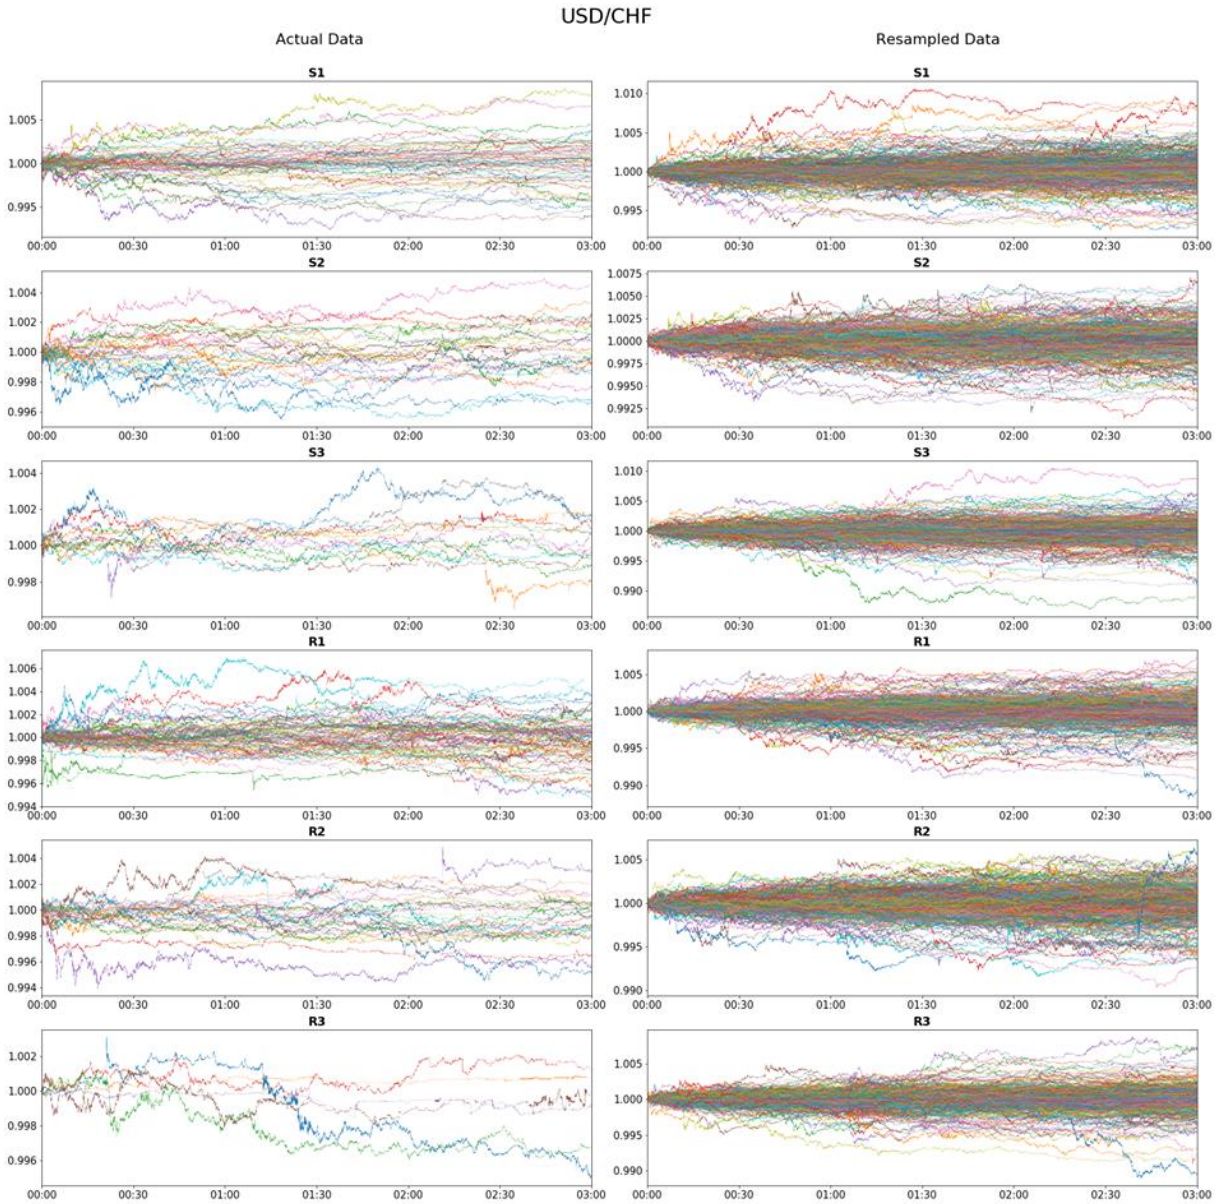

Figure 15: This figure provides an illustration of the bootstrap resampling method for the GBP/JPY. The left column shows all the actual price evolutions from the time the price hits the specific pivot level $(S_1, S_2, S_3, R_1, R_2\ or\ R_3)$ and three hours forward. This leads to six different groups, one per each pivot level, of actual price evolutions. The right column shows 1,000 price evolutions, randomly generated for each of the pivot levels. The generated series are extracted from similar times as the actual price evolutions. The Y-axes in each of the twelve graphs show the normalized price and the X-axes show a three-hour time window. *Source: Authors' calculations.*