The effect of a Tobin tax on the exchange rate volatility

Sandra Askegaard
Copenhagen Business School
Stud. M.Sc. Applied Economics & Finance

Flemming Søgaard
Copenhagen Business School
Stud. M.Sc. Applied Economics & Finance

Counselor:
Professor of Finance
Jesper Rangvid
Copenhagen Business School
Department of Finance

4th of March 2011

The thesis work contais 20 figures and tables of 800 characters
and text equal to 258,616 characters totalling 274,616 characters.
The word count can be found on the next page.

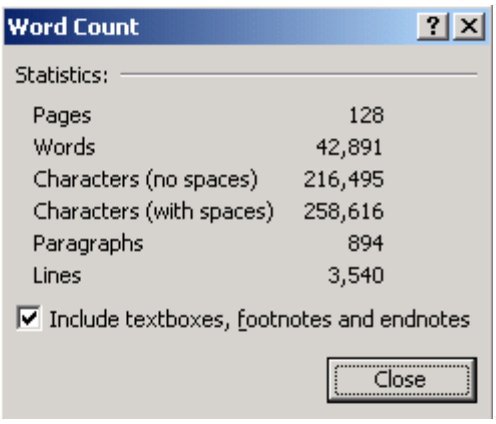
Word Count
Statistics:
Pages 128
Words 42,891
Characters (no spaces) 216,495
Characters (with spaces) 258,616
Paragraphs 894
Lines 3,540
Include textboxes, footnotes and endnotes
Close

Abstract

This thesis work analyses the impact of a Tobin tax on exchange rate volatility and evaluates the proposed sizes for such a tax. We furthermore suggest a possible solution to the inherent endogeneity problem between transaction costs and exchange rate volatility. Finally, we link our empirical results to the conclusions of a game theoretical model to comment on how the market composition corresponds to the contrasting theories underlying a Tobin tax. Using a multiple variable time-series model we estimate the effect of an increase in transaction costs, resembling the introduction of a Tobin tax, on the exchange rate volatility. We use a measure of cost based on the percentage bid-ask spread of the exchange rate analogous to a transaction tax, which furthermore does not require us to make any assumptions of the market efficiency. The measure of exchange rate volatility that we use captures the variation within a period by utilizing the minimum and maximum values of the exchange rate opposed to opening and closing values. As a possible solution to the endogeneity problem, we propose that the exchange rate can be used as an instrument for our transaction costs measure. We believe so because the exchange rate is unaffected by the volatility and because the exchange rate affects our measure of transaction costs. To test this, we set up a two-stage least squares model. From our multiple variable time-series model we conclude that the introduction of a Tobin tax into the currency market will have either a positive significant or negative insignificant effect on the exchange rate volatility. Furthermore, these results are emphasized by our analysis of our instrumental variable. Based on these results, we conclude that a Tobin tax is more likely to amplify than dampen volatility thereby contradicting the claim made by Tobin (1978) but agreeing with the view of Friedman (1953). Although comparing our empirical results to the game theoretical model suggest that the market is composed of a larger share of 'noise traders' than 'informed traders'. This indicates that the view of Friedman can exist even if the majority of traders are not informed. Additionally, we use our empirical results to measure the impact of two proposed taxes of 0.5% and 0.005%. The first is motivated by the desire to change market behaviour and the second is motivated only by the desire to generate revenue to cover the costs of financial crises without affecting the market. We find that the first tax is up to seventeen times larger than the average transaction costs and even if it worked as intentioned we find a tax of this size to be unrealistically large. In contrast, the revenue generated by the second tax only covers a fraction of the price of the latest crises and increasing it will only influence the market behaviour and thereby most likely increase volatility.

Table of Contents

1 Introduction

For decades exchange rate volatility has been a concern for governments. The highly unpredictable nature of floating exchange rates can cause governments to adjust their interest rates thereby making it difficult for them to keep long-run policies and improve employment, inflation and output. This makes the economies susceptible to speculative attacks, creating the foundation for possible financial crises. An often suggested instrument for curbing financial crises is the introduction of a *currency transaction tax*. Proponents of this instrument argue that the additional required foreign rate of return that the tax creates for traders makes the traders focus on longer term investment horizons; thereby decreasing destabilizing behaviour and excess volatility.

The idea of a security transaction tax was first proposed by John Maynard Keynes in 1936. Keynes believed that a tax would increase the amount of trading based on long term fundamentals and discourage trading based on the expectations of the market. Contrary to this, Milton Friedman argued in 1953 that traders are rational and stabilizing, wherefore a transaction tax would scare off traders thus making the share of rational traders smaller and the discovery of the fundamental price less efficient. In 1972, the period after the collapse of the so-called Bretton Woods system, James Tobin picked up Keynes' idea, and suggested a tax on currency transactions. During the movement away from a fixed exchange rate regime to a floating rate system, the proposal of a currency transaction tax could be seen as a new system for international currency stability. Since then, the proposal has generated a lot of interest both in academic circles as well as in international government institutions and is often referred to as the "*Tobin tax*".

Through the last three decades the Tobin tax has appeared on the agenda of several international institutions such as the *Bank of International Settlements* (BIS), the *Group of 20* (G20), the *American Economic Association* (AEA), the *International Monetary Fund* (IMF), the *European Union* (EU) and the *Asia-Pacific Economic Cooperation* (APEC). The discussion of a Tobin tax reappeared after the 1987 stock market crash, again in 1992 in relation to the European exchange rate crisis, in 1997 when the Asian financial crisis began and most recently in connection with the global financial crisis in 2007.

Since the Tobin tax has been suggested as an instrument for dampening financial crises so often, we find it interesting and important to test if the transaction tax has the intentional effect on volatility.

Others have previously studied the empirical effect of a transaction tax on volatility, although mainly on the stock market. Mulherin (1990), Ronan & Weaver (2001) and Hau (2006) look at the connection between transaction costs and volatility in the stock market and mainly find a positive significant relationship, although in a few cases the relationship is negative and insignificant. Whereas Aliber, Chowdhry & Yan (2003) and Lanne & Vesala (2006) look at the same relation but in the currency market and reach similar conclusions as the former, although they find a more definite positive relationship.

Our thesis work aims at examining the relationship between transaction costs and volatility in the four major currencies: Great British Pound (GBP), Swiss Franc (CHF), Japanese Yen (JPY) and Euro (EUR) against the US Dollar (USD). The majority of our significant results suggest that a Tobin tax on currency transactions will in fact result in an increase in the exchange rate volatility, contradicting the claim made by Tobin.

We link our results to the conclusions drawn from the game theoretical model proposed by De Long *et al.* (1990) and expanded further by Song & Zhang (2005). They conclude that a Tobin tax warrants both the views of Keynes and Friedman, where an increase in volatility is associated with major currencies, corresponding to our findings, and a decrease in volatility is related to minor currencies. Our results are again consistent with the views of Friedman, although based on Song & Zhang's (2005) model we find that the composition of the market is not in line with Friedman's assumption that the majority of traders are rational.

To empirically examine the relationship between transaction costs and exchange rate volatility, we use a multiple variable time-series model over the period 01-03-1990 to 10-22-2010. We therefore construct time series of weekly averages based on daily measurements for transaction costs, volatility and market volatility. We use the percentage bid-ask spread relative to the exchange rate as a measure of transaction costs, which, contrary to Frenkel & Levich (1975) and Roll (1984), does not require any assumptions as regards to the market efficiency. We furthermore use the percentage range between the maximum and minimum price compared to the mean price as a measure of volatility. The range measure is more robust to discrete changes in the exchange rate, contrary to the return standard deviation measure used by Mulherin (1990), Umlauf (1993), Bessembinder & Rath (2002), and Bessembinder (2003). For a benchmark of the general market volatility we use VIX, which is a measure of the implied volatility of S&P 500 index options.

Using this approach creates a theoretically based endogeneity problem between the bid-ask spread and volatility, which makes it difficult to determine the true causational relationship between transaction costs and volatility. The problem could have been circumvented by using a one-time change in transaction costs hereby making the change exogenous, although it should be noted that this method is statistically less powerful compared to the one we have chosen. The problem of endogeneity has been mentioned by Aliber, Chowdhry & Yan (2003) and Lanne & Vesala (2006) as well, but we conclude that their solutions are questionable and that they have failed to solve the endogeneity problem satisfactorily. We try to solve this with a two-stage least squares model with the exchange rate as an instrument variable for our cost measure. This method provides the same results as our multiple variable time-series model, hence adding emphasis to our original results.

1.1 Thesis statement

The overall aim of this thesis work is to determine the impact that a Tobin tax introduced into the currency market will have on the exchange rate volatility, and if it helps to stabilize or destabilize the currency. We address this issue from both a theoretical and an empirical point of view. We link our empirical results to the conclusions of a game theoretical model in order to examine if the *trader composition* corresponds to the theories underlying the Tobin tax. The questions we seek to answer are summarized below.

1) *How does a Tobin tax introduced to the currency market affect the exchange rate volatility?*
    - What market hypothesis do our empirical results indicate?
    - What are the motivations behind the proposed tax rates and do they seem reasonable?

2) *How can the endogeneity problem between transaction costs and volatility be solved?*
    - Does correcting for the endogeneity problem change the effect of the Tobin tax on exchange rate volatility?

3) *What does the effect of a Tobin tax on volatility say about the composition of the market?*
    - How does this fit with the underlying, contrasting theories of a Tobin tax?
    - How does this fit with our game theoretical model?

To our knowledge, research on this subject has not previously linked the contrasting theories underlying the Tobin tax, their empirical results and a detailed game theoretical model, which incorporates all the views of the contrasting theories. Furthermore, we believe this thesis work is the first to suggest a possible instrument as a proposed solution to the endogeneity problem.

1.2 Limitations

Throughout this thesis work we expect the reader to be familiar with basic financial and econometric theory and terminology. We will only focus on the OLS assumptions regarding heteroscedasticity, normality, multicollinearity, autocorrelation and exogeneity, as we expect the remaining to be met. We will address any problems or violations of the aforementioned assumptions we may encounter in the empirical analysis.

In the equations where we have not supplied the derivations, we have simply reported the reference in where it can be found. We have done so as we do not believe the derivations to be important for this thesis work as a whole, but that only the final result is relevant.

1.3 Thesis structure

This section aims to describe the structure of the thesis. Overall, the thesis consists of ten chapters, which are divided into sections. We will briefly explain the structure of the thesis and what the reader will find in each chapter.

Chapter 1: Introduction

The purpose of chapter 1 is to give the reader an introduction to the thesis topic and the problems addressed throughout the thesis work.

Chapter 2: Theory

The purpose of chapter 2 is twofold. The first part describes the contrasting theories underlying a Tobin tax. The reader should have a thorough understanding of the original idea behind the Tobin tax and of how the introduction of such a tax theoretically affects the market depending on the assumed market hypothesis. The second part aims to develop a game theoretical framework. After reading this, the reader should know how changes in the market, such as an added tax, affect the share of noise traders and informed traders in the market, and how this affects the volatility.

Chapter 3: Methodology

This chapter tries to explain the measures we chose to use in our models as well as the statistical model. We do this through a review of previous research in which we discuss the pros and cons of the different methods used. The reader will also obtain a basic knowledge of how the bid-ask spread is determined in the market and of the endogeneity problem originating from the interaction between transaction costs and exchange rate volatility.

Chapter 4: The data

In this chapter we describe our data and the variables that we have included in our forthcoming model as well as variables we would have liked to include but could not obtain. We evaluate our variables' stationarity and normality, and divide them into structural breaks.

Chapter 5: Diagnostics

This chapter aims at describing four of the main assumptions behind the ordinary least squares model and the bias it will cause if they are violated. Additionally, we will perform a granger causality test to obtain a preliminary idea of the endogeneity problem before running our OLS regressions.

Chapter 6: The model

In this chapter we setup our OLS regression model and compute our estimated coefficients for the variables that we have chosen to include. Before describing our results we point out critical arrears of our model and evaluate if our results are biased from violations of the model's OLS assumptions. We further compare our results to the results from different approaches used on the same dataset as well as the findings obtained throughout previous research.

Chapter 7: Analysis

In chapter 7 we look at the interpretation of our results and especially what they tell us about the relationship between transaction costs and exchange rate volatility. Furthermore, we evaluate the effect of two proposed Tobin taxes on the volatility, as well as the assumptions behind the Tobin tax. We also propose a possible solution to the endogeneity problem by using an instrument variable via a two-stage least squares model.

Chapter 8: Discussion

This chapter aims to discuss the effect a Tobin tax has on public welfare as well as its ability to postpone financial crises.

Chapter 9: Suggestions for further research

This chapter will propose four topics that we believe need to be researched more extensively, as this could give a more clear picture of a Tobin tax' effect on exchange rate volatility.

Chapter 10: Conclusion

Chapter 10 presents the overall conclusions of the thesis work and aims to answer the questions set forth in the thesis statement.

2 Theory

Throughout chapter 2 we will compare and contrast the different theories underlying a so-called Tobin Tax proposed in the seventies. Section 2.1 will give an introduction to the original proposal by Tobin and explain the implications a functioning transaction tax would have for the traders. Sections 2.2.1 and 2.2.2 will explain John Maynard Keynes' and Milton Friedman's arguments for their different views on how a Tobin Tax will affect the volatility of the market, respectively, and section 2.3 will give a discussion of the empirical evidence for both Keynes' and Friedman's theory, and provide a more critical view of their arguments. After having explained the assumptions behind the traders' behaviour we will give a thorough overview of the possible implications on the price development and volatility of the market through De Long *et al.*'s (1990) noise trader model in section 2.5, as well as the effect a transaction tax will have on the market (Song & Zhang (2005)).

2.1 James Tobin

When James Tobin first suggested a tax on exchange rate investments it was, as he put it, to throw some sand in the wheels of international finance (Tobin (1978)). Tobin's concern was that the international financial markets had become too efficient. Not in the strictly economic sense that all information was available at any time to anyone, but rather that the flow of available information moved swiftly around the world, the transaction costs had fallen and investments could be made quickly (Tobin (1978)).

The problem with the increasingly efficient financial market is that national economies cannot adjust quickly enough to rapid movements of liquid funds from one currency to another. Labour and goods move slower than liquid funds and it will therefore be difficult to adjust one's monetary policies quickly enough. This will in turn force governments to change their monetary policies, which could hurt employment, inflation and output (Tobin (1978)).

Ideally, an international currency across all boarders will help eliminate destabilizing speculation in currencies completely. As Tobin (1978) explains, one can imagine the trade between two states in America where both states trade in the US Dollar. Since the two states have the same currency no speculation can affect interest rates in either state. Another example made by Eichengreen, Tobin and Wyplosz (1995) is the Euro, or at that time the EMU, which eliminates many national currencies and replaces them with the same currency across countries. This will enable governments to keep long-term monetary policies with a focus on the country's economy instead of foreign currency speculation.

Tobin realized that a global currency laid far into the future and chose to suggest his second best solution: a tax on exchange rate trades. He assumed that the currency market consisted of a combination of *noise* and *informed* traders. Noise traders are defined as traders who do not buy and sell with knowledge of the currency's fundamental price, whereas informed traders are aware of this fundamental equilibrium and base their decision to buy or sell there on. Tobin believed that investments were primarily short-termed and irrational resulting in increased volatility in the different currencies (Tobin (1978)). He assumed that noise traders were speculating on the market's expectation of the future price of the currency and not the true equilibrium. In the case that speculation is based on the price of the currency the noise traders will have a tendency to bid up the price and move it away from its equilibrium, whereas, if the speculation is based on the currency's equilibrium the traders will always drive the price towards this equilibrium (Tobin (1978)), as we shall see in section 2.3.

Tobin believed that a percentage tax on each currency transaction would help drive away the short-term traders as the tax would consume a relatively larger amount of the traders gain on the investment, as opposed to long-term investments. The tax will furthermore be directly proportional to the amount of trades conducted.

Frankel (1996), amongst others, demonstrates how traders will require an increasingly large foreign rate of return on their investment as the tax rises and the time horizon shortens. Assuming that the trader has to pay the tax twice on a round trip investment Frankel (1996) finds that the minimum required foreign rate of return is equal to:

$$\mathrm{i}^* = \frac{i + 2\frac{t}{y}}{1 - t} \tag{1}$$

Where i is the domestic rate of return, i* is the foreign rate of return, t is the Tobin tax and y is the investment horizon. The derivations of equation (1) can be found in Frankel (1996). In Figure 1 we assume that the daily domestic rate of return is 0.05%, which results in the trader requiring a daily foreign rate of return of 2.07% at a Tobin tax of 1% and an investment horizon of 1 day. This requirement then falls as the investment horizon rises.

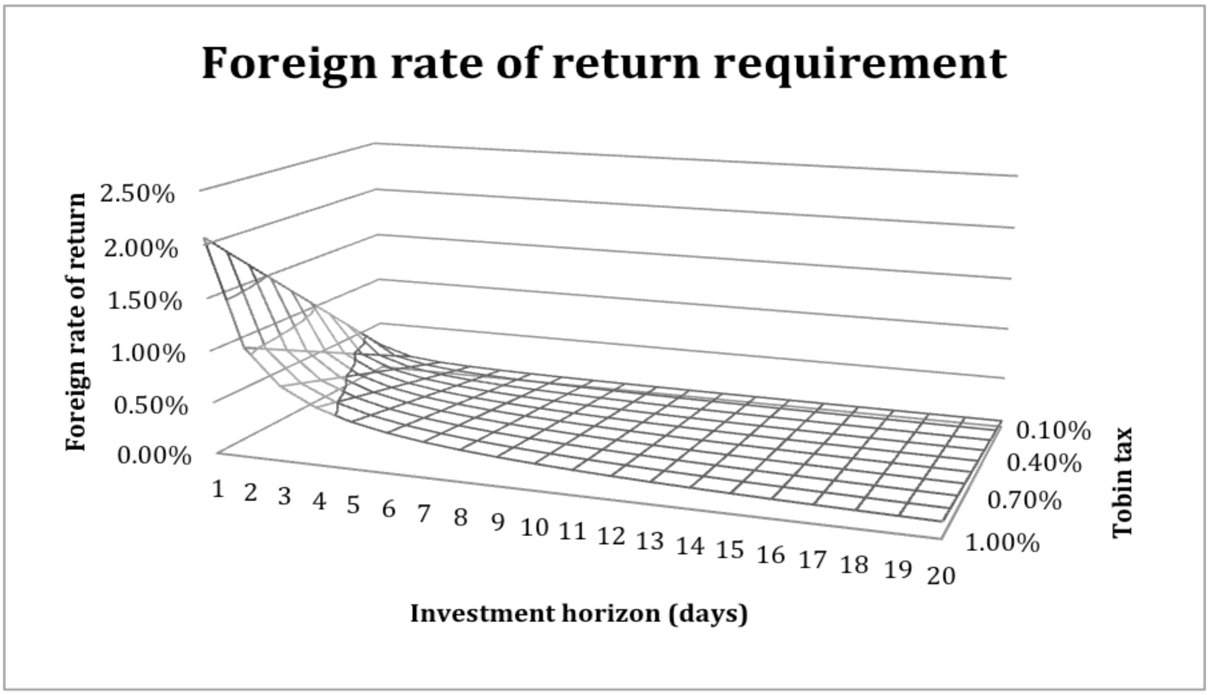

Figure 1 - A trader's daily foreign rate of return requirements with a daily domestic rate of return of 0.05%. Based on our own calculations.

For the Tobin tax to work, it should be applied to and be equal across all markets. If the tax is larger in one country compared to others or only present in one country, foreign exchange will move to other markets (Eichengreen, Tobin and Wyplosz (1995)). Umlauf (1993) finds that an introduction of a transaction tax of 1% that rises to 2% at a later point on all stock transactions in Sweden results in traders moving the majority of their trading with Swedish stocks to London, thereby circumventing the transaction tax.

Tobin first suggests an international tax of 1% in his article from 1978 but later Eichengreen, Tobin and Wyplosz (1995) change this proposal to a tax of 0.5%. The results from our models will give us an opportunity to evaluate the suggested sizes of the tax-rate later on.

The argument that a transaction tax should drive away short-term traders from the currency market and make noise traders focus more on long-term fundamentals is based on three assumptions (Haberer (2003)). First of all, one has to assume that short-term speculation is actually destabilizing and increases volatility. Second, that a Tobin tax will discourage this form of speculation and not to a larger extent harm stabilizing trading. Third, that a Tobin tax will actually turn traders to base their investment decisions on the macroeconomic fundamentals affecting the equilibrium price of currencies. In section 2.3 we will discuss the validity of these assumptions.

2.2 Opposing views: Keynes contra Friedman

When we look at the question of security transaction taxes, we find two contrasting theories underlying the effect that such a cost would have on the volatility of the traded security. First, John Maynard Keynes proposed his positive view on the introduction of a security transaction tax, which is the basic idea that Tobin picked up decades later. Second, we have Milton Friedman whose opposing view suggests the diametrically opposite results. Both Keynes and Friedman realize that the introduction of a transaction tax, or Tobin tax, would drive away traders, as described in section 2.1, but their similarities stop when comparing their views on what effects such a tax will have on the security's volatility.

Their different conclusions arise from their views on how actors in the market behave. On this subject, Keynes and Friedman differ in their opinions concerning if noise or informed traders drive the price development of the security and whether their speculation is either stabilizing or destabilizing.

2.2.1 John Maynard Keynes

In his book from 1936, Keynes describes the noise traders' speculation as "*the activity of forecasting the psychology of the market*" and goes on to say that they are only "*interested in discovering what average opinion believes average opinion to be*" (Keynes (1936, p. 158 f.)). Keynes believes that the trading behaviour of the noise traders is affected by the pre-existing expectations of the market, or the "market mood", as well as emotions and psychology. It is, in essence, the objective of the trader to beat the market.

It is therefore believed that the noise traders have positive feedback to price moves (Haberer (2003) and Summers & Summer (1989)) and will try to determine in what way news will affect the market. Based on their belief of the "market mood", they will drive the prices up or down. This accumulative behaviour of the market trying to follow the rest causes the price of the security to follow either up- or downward trends away from the equilibrium price, as Schulmeister (2010) notes, and in that way increase the volatility. These quick reactions to news will, in the short-run, increase price fluctuation and over time the price development will follow more medium-term trends created by the market, which themselves deviate from the fundamental price over time.

Keynes' theory on how market participants act is displayed in Figure 2 by the line AKLO. At time t=1 new information is available to the market. Due to rigidity in the market, traders drive the price up slowly towards the new equilibrium price of P=104, but because of the trend created in the market, as explained above, the price continues upward. The same is the case when new information makes the price fall again at t=3 and rise at t=5 - a theory that Schulmeister (2010) dubs the "Bull-Bear-hypothesis".

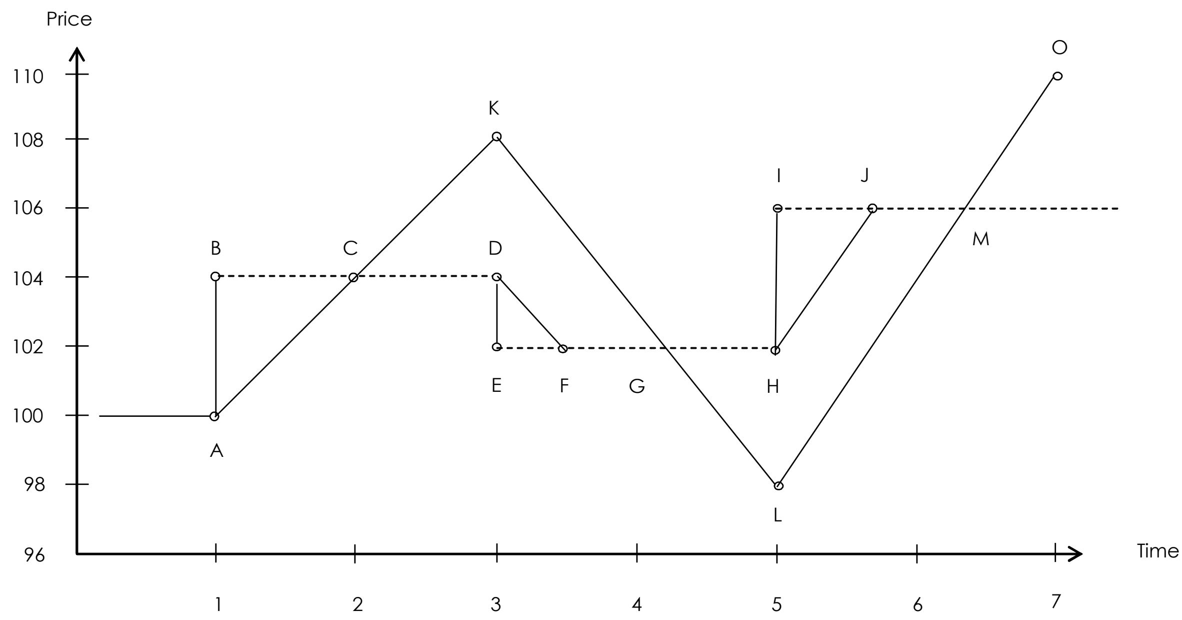

Figure 2 – Depiction of the different price developments under either Keynes' or Friedman's market hypothesis. Keynes' hypothesis of irrational market participants who drive prices away from their equlibrium is expressed by the line AKLO. Friedman's hypothesis of rational market participants who drive the price toward its fundamental is shown by the line ACDFHJ (Schulmeister (2010)).

The tax that first Keynes and later Tobin suggested would make traders demand an increasingly large foreign rate of returns on their currency investments, the shorter the time period they choose to invest, as explained in section 2.1. As these returns naturally become increasingly difficult to match, the tax will drive all but the most daring short-term traders away and shift focus from destabilizing speculation to more long-term investments (Schwert & Seguin (1993) and Song & Zhang (2005)).

2.2.2 Milton Friedman

Unlike Keynes, Friedman (1953) does not believe that speculation is destabilizing for a currency. His hypothesis is based on market participants having rational expectations (Schulmeister (2010)), though they are not aware of the other market participants' expectations. The price development that this entails is best explained by Figure 2, where the price would follow the path ACDFHJ. At time t=1 new information is available to the market. Since the market participants do not know each other's expectations the price moves smoothly toward its new equilibrium price at P=104 through a price discovery process. Friedman argues that it is exactly this type of speculation in price that helps the market discover the fundamental value of the currency. Since traders are rational they have negative feedback to changes in prices (Haberer (2003) and Summers & Summers (1989)) and will not let the price move away from its equilibrium because, as Friedman argues, if the trader knows if the currency is either over- or undervalued, he knows if it is time to sell or buy (Frankel (1996), Haberer (2003) and Summers & Summer (1989)). Since no rational trader would buy high and sell low, as this strategy would bankrupt him quickly, this mechanism will always drive the price towards its equilibrium and only new information would move the fundamental price as we see again in Figure 2 from point D to F and again from H to J (Carlson & Osler (1998)). The price would always be based on the traders' long-term expectations and move quickly towards the same point – the fundamental price – to eliminate profit opportunities, wherefore this situation could never create excessive volatility (Summers & Summers (1989) and Frankel (1996)), opposed to the bull-bear hypothesis where prices always move away from the equilibrium price creating excessive volatility. Introducing a Tobin tax to a market with primarily stabilizing traders would, according to Freidman, drive away traders and make the price discovery process less efficient thereby creating room for more volatility.

2.3 The empirical evidence on market hypotheses

In the following sections we will discuss the empirical research done on the trader composition of the market. In section 2.3.1 we will look at evidence supporting Friedman's point of view whereas in section 2.3.2 and 2.3.3 we look at arguments against Friedman and indications of Keynes' market view. Section 2.3.4 proposes a composition of the market, which can consist of both Keynes' and Friedman's views.

2.3.1 Rational expectations hypothesis

Friedman has not received much credit in the part of the academic world concerning itself with a Tobin tax. This may be because it is difficult to observe the fundamental price of a currency, and thereby test which of the market hypotheses are correct, although it is easy to measure the market's volatility (Dimson & Mussavian (2000)). It is, however, difficult to define if the volatility is "excessive", and when excessive volatility is present in the market or not (Schwert & Seguin (1993)). Schulmeister (2010) tries to define it as volatility from transactions which over- or undervalues the currency in respect to its fundamental value. This definition seems adequate to us since the goal of the Tobin tax is to drive speculation to focus more on fundamentals.

Proponents of Friedman argue that a Tobin tax to curb this excessive volatility might do more harm than good. Haberer (2003) argues that a currency dealer has a certain risk aversion that she wants to maintain. When a sudden order from a customer makes her be too long or short in a currency, she will try to restore her original risk aversion by passing on the unwanted position in the market. The possibility of passing on these unwanted positions in the market creates the possibility of getting rid of positions that are too risky and would thereby be reflected in the bid-ask spread faced by the customer. Since these "hot-potato" trades account for most of the trading done in the market, a Tobin tax would punish risk-lowering hedging (Haberer (2003) and Sarno & Taylor (2001)). This argument does not directly dismiss the first and second assumption of a Tobin tax, mentioned in section 2.1, that noise traders are destabilizing and that a Tobin tax would discourage this behaviour, but does suggest that the market consists of a much larger part of informed traders than noise traders, wherefore a tax would be more destabilizing than stabilizing.

2.3.2 Adaptive expectations hypothesis

As a proponent of a Tobin tax, Schulmeister (2010) argues that security prices, be it stocks or currencies, fluctuate around underlying trends, which again fluctuate around even longer trends. These small fluctuations that turn into trends, he says, are evidence of the Keynesian market hypothesis and the first assumption concerning the viability of a Tobin tax that short-term speculation creates excess volatility. Summers & Summers (1989) sum up some of the research done on what market hypothesis is dominating the stock market and conclude that there is little evidence to support that a tax discouraging short-term investing should lead to increased volatility, thus giving support to the second assumption of the Tobin tax. Frankel (1996) finds evidence on how a depreciation or appreciation of a currency affects the development of the currency after different time intervals. Over short periods of a week, he comes to the same results as Schulmeister (2010), namely that a change in the market leads to trends, but he finds that after a month the effect is dampened and after 3-12 months the effect becomes negative. This could indicate that traders are affected by news in the short run but follow the development in the currency's fundamental price in the long run.

2.3.3 Evidence against rational expectations

Based on evidence from the 1987 stock market crash, Summers & Summers (1989) note that traders were reluctant to trade according to the fundamentals when they noticed deviations in the form of an undervaluation of the market. Summers & Summers (1989) believe that the traders' reluctance to trade was due to the possibility that the market would depreciate more before it would appreciate, which would suggest that the majority of traders are destabilizing and increase volatility. In the short run, this is comparable to Frankel's (1996) and Schulmeister's (2010) results that we described in section 2.3.2. Schulmeister (2010) argues that if prices move smoothly from one equilibrium to another, as Friedman proposes, then profit-seeking traders will develop strategies based on these up- or downward trends. Frankel (1996) further describes how these profit seeking trends, that Keynes criticizes and Friedman deems a wrong interpretation of how the market works, can lead to rational speculative bubbles where traders are destabilizing the market without losing money. The theory of speculative bubbles assumes that the bubbles are created on the same basis as the market trends, described by Keynes (1936), Tobin (1978), Schulmeister (2010) and others. The rational aspect of the bubbles lies in what Summers & Summers (1989) noticed during the crash of the stock market, where it becomes rational for a trader not to try changing the trend by buying stocks for the fear of the market falling further, making that individual lose money.

There seems to be more evidence based on logical arguments made from events in history to support the assumptions that traders' speculation is destabilizing and that a transaction tax would discourage this speculation.

2.3.4 The market's trader composition

The fact might be, though, that the market consists of both noise and informed traders. It then becomes a question of the composition of the market, and whether the largest proportion of traders are noise traders or informed traders.

Frankel (1996) argues that the speculative bubbles are started by noise traders' responses to informed traders' reactions. It has been suggested that the two forms of traders work against each other thereby keeping the security's price in check. While noise traders increase volatility, the informed traders will stabilize the market by trading against the noise traders (Song & Zhang (2005)). The fact that informed traders have long-run investment horizons and noise traders have short-run investment horizons may result in excess volatility in the short run but a reversion towards the fundamentals in the long run (Frankel (1996), Sarno & Taylor (2001) and Haberer (2003)).

It seems that it is often forgotten that we cannot easily distinguish between an informed trader and a noise trader. If they both influence the market, the informed traders might trade as often as the noise traders because informed traders bet against the noise traders as they move the price away from its equilibrium (Eichengreen, Tobin & Wyplosz (1995)).

2.4 Summary

James Tobin wants to dampen destabilizing short-term speculation in the currency market by introducing a transaction tax. His point is that more stable currencies will give national economies room to focus on long term monetary planning instead of adjusting interest rates to cope with depreciations and appreciations of their currency in the short-run affecting the stability of employment, inflation and output. His idea is based on John Maynard Keynes' thoughts on how speculation affects prices and volatility. Keynes believes that traders base their investment decisions on their assumption of how the market will react. This behaviour is destabilizing and will accumulatively create market trends in the short- and medium-term, which will deviate from the true equilibrium price of the currency.

Contrary to this, Milton Friedman argues that traders are rational and informed. They are aware of when a currency is over- or undervalued and therefore know when to sell and buy, respectively. This behaviour will always drive the price towards its equilibrium and stabilize volatility.

It is difficult to say anything conclusive on how traders behave. Summers & Summers (1989) and Schulmeister (2010) argue that they can observe that traders behave like the noise traders described by Keynes. Frankel (1996), on the other hand, finds that in the short-term of 1 to 4 weeks there is a positive relation between an appreciation in a currency in one week on the next four weeks, but that this effect becomes negative over a period of 3-12 months. He goes on to say that the overshooting trends that create bubbles in the currency market can be started by a change in the informed traders' price signals, which are then picked up by the noise traders. Eichengreen, Tobin & Wyplosz (1995), Song & Zhang (2005) and Haberer (2003) express a similar point, namely that the market might very well consist of a combination of informed and noise traders. It then becomes a question of the size of each of these types of traders' fraction of the collective pool of traders to validate if the third assumption of the Tobin tax holds - that traders base their investment decisions on the market fundamentals or not. From the results of our statistical models later on, we hope to be able to explore this question further.

2.5 Game theoretical explanations

The purpose of this section is to try and explain how noise traders and informed traders react to changes in the market, such as an added tax and how they themselves affect the market and thereby the volatility. From the results of our forthcoming model in chapter 6, we will try to describe the composition of the market based on the conclusions from the game theoretical model, which again can be compared to the theories of Keynes, Tobin and Friedman.

From section 2.5.1 to 2.5.6 we will explain how the composition of noise traders and informed traders can affect the assets' price and volatility through a mathematical model proposed by De Long *et al.* (1990). Section 2.5.1 formulates the traders' utility function as well as their perception of prices. Section 2.5.2 and 2.5.3 will explain the traders' demand for the risky asset and their relative expected returns, respectively, when both noise traders and informed traders are trying to maximize their returns in a market with no fundamental risk. Section 2.5.4 introduces fundamental risk and section 2.5.5 looks at how the behaviour of the traders changes when they want to maximize utility instead of return. In the last section, section 2.5.6, we evaluate how the results of the previous sections changes when we introduce a Tobin tax and liquidity risk into the market.

2.5.1 Behaviour

One way to show how noise traders (n) and informed traders (i) affect the currency market is to formalize models to describe the traders' behaviour on a strictly mathematical basis. De Long *et al.* (1990) do this and reach many of the points made in the previous sections. They assume that the market consists of a share of noise traders (μ) and informed traders ($1-\mu$) who, respectively, will invest a share λ_t^n and λ_t^i in an unsafe asset (u). They further assume that the traders only live through two periods: t where the traders are young, and t+1 where traders are old, and that the market is solely made up of the unsafe asset and a safe asset (s) which both pay the dividend (r) in the absence of noise traders and that the market has no fundamental risk.

In the case that both the safe and the unsafe asset's prices reflected the present value of all future dividends they would be perfect substitutes and always sell at the fundamental value of 1, although this is not the case as the price of the unsafe asset at time t (p_t) is affected by the presence of noise traders. The noise traders misperceive the price of the risky asset by an i.i.d. normal random variable:

$$\rho_t \sim N\left(\rho^*, \sigma_\rho^2\right) \tag{2}$$

where ρ* is a measure of the average misperception by the noise traders and σ_ρ^2 is the variance of their misperception. It is this average misperception that we described as the "market mood" in section 2.2.1, and the present misperception (ρ_t) of equation (2) that describes the noise traders' reactions to news.

Additionally, it is assumed that both types of traders have a constant absolute risk aversion function defined as:

$$\mathrm{U} = -e^{-(2\gamma)w} \tag{3}$$

where w is wealth in period t+1 and γ is the coefficient of absolute risk aversion. Maximizing their expected utility of equation (3) is comparable to maximizing:

$$\mathrm{E(U)} = \bar{w} - \gamma\sigma_w^2 \tag{4}$$

where $\bar{w}$ is the expected value of wealth and σ_w^2 is the variance of wealth in the next period. From equation (4) $\bar{w}$ can be expressed as a function of the expected value of the traders' return on their unsafe asset, which is defined as $r + {}_tp_{t+1} - p_t$ and the safe asset defined as r:

$$\bar{w} = c_0' + \lambda_t\left(r + {}_tp_{t+1} - p_t\right) + (1 - \lambda_t p_t)r \tag{5}$$

where an anterior subscript denotes the time at which an expectation is taken. The second term represents the expected return from their investment in the unsafe asset, and the third term is the return from the share invested in the safe asset. When rearranging this we get:

$$\overline{w} = c_0 + \lambda_t^i \left(r + \ {}_t p_{t+1} - p_t(1+r)\right) \tag{6}$$

where $c_0 = c_0' + r$ and represents the first period's labour income. From equation (6) and the risk associated with this investment - the one-period ahead expected variance of p_{t+1} defined as: ${}_t\sigma_{p_{t+1}}^2 = E_t\left\{\left(p_{t+1} - E_t(p_{t+1})\right)^2\right\}$, we can now express equation (4) for both informed traders and noise traders as equation (7) and (8), respectively.

$$\begin{aligned} \mathrm{E}(\mathrm{U}_i) &= \overline{w} - \gamma\sigma_w^2 \\ &= c_0 + \lambda_t^i \left(r + \ {}_t p_{t+1} - p_t(1+r)\right) - \gamma\left(\lambda_t^i\right)^2\left\{ {}_t\sigma_{p+1}^2\right\} \end{aligned} \tag{7}$$

$$\begin{aligned} \mathrm{E}(\mathrm{U}_n) &= \overline{w} - \gamma\sigma_w^2 \\ &= c_0 + \lambda_t^n \left(r + \ {}_t p_{t+1} - p_t(1+r)\right) - \gamma(\lambda_t^n)^2\left\{ {}_t\sigma_{p_{t+1}}^2\right\} + \lambda_t^n\{\rho_t\} \end{aligned} \tag{8}$$

The last term of equation (8) reflects the noise traders' misperception of the price of asset u. It clearly shows that the traders' utility increases as the return enlarges, and decreases as the third term, the risk, increases, exhibiting the nature of the risk averse trader.

2.5.2 Demand for the risky asset

If we differentiate equation (7) and (8), setting them equal to 0 and solving for λ_t^i and λ_t^n, respectively, yield the following maximizing shares to invest:

$$\lambda_t^i = \frac{r + \ {}_t p_{t+1} - (1+r)p_t}{2\gamma\left\{ {}_t\sigma_{p_{t+1}}^2\right\}} \tag{9}$$

and

$$\lambda_t^n = \frac{r + \ {}_t p_{t+1} - (1+r)p_t}{2\gamma\left\{ {}_t\sigma_{p_{t+1}}^2\right\}} + \frac{\rho_t}{2\gamma\left\{ {}_t\sigma_{p_{t+1}}^2\right\}} \tag{10}$$

From equation (9) and (10) we see that the traders' demand for the unsafe asset is increasing in its numerator, the return, and decreasing in its denominator, the risk. The additional term in equation (10) shows that the noise traders' demand is also proportional to their misperception of the assets' price.

Since the holdings of the risky asset have to equal 1, a pricing formula can be derived from the following relationship, using equations (9) and (10):

$$1 = (1-\mu)\lambda_t^i + \mu\lambda_t^n \Leftrightarrow$$

$$p_t = \frac{1}{1+r}\{r + {}_tp_{t+1} - 2\gamma\{{}_t\sigma_{p_{t+1}}^2\} + \mu\rho_t\} \tag{11}$$

We see that the price of asset u at period t is a function of the riskless rate of return (r) the expected one-period ahead price of u, the variance of the expected price and the noise traders' misperception. The fact that ${}_tp_{t+1}$ is our expected price, at time t, of the asset one period ahead, means that our expectation of future prices is always defined by our available information. It will also be the case that the expected price two periods ahead is based on our information at period t+1 (${}_{t+1}p_{t+2}$) and so on, wherefore we can express the price of the unsafe asset by applying the law of iterated expectations:

$${}_tp_{t+1} \equiv E(P_{t+1}|\mathrm{I_t}) \tag{12}$$

and therefore:

$$E(E(\mathrm{P_{t+i}}|\mathrm{I_{t+i}})|\,\mathrm{I_t}) = \mathrm{E}(\mathrm{P_{t+i}}|\mathrm{I_t}) \tag{13}$$

where I is the information set available at time t or t+i. By applying this rule and recursively solving equation (11) (De Long *et al*. (1990)) we find:

$$p_t = 1 + \frac{\mu(\rho_t - \rho^*)}{1+r} + \frac{\mu\rho^*}{r} - \frac{2\gamma}{r}({}_t\sigma_{p_{t+1}}^2) \tag{14}$$

As only the misperception of the noise traders create variance in the price and this is constant over time cf. equation (2), we can express the last term of equation (14) as follows:

$${}_t\sigma_{p_{t+1}}^2 = \sigma_{p_{t+1}}^2 = \frac{\mu^2\sigma_\rho^2}{(1+r)^2} \tag{15}$$

This result originates from equation (14) and is rather trivial when realizing that ρ_t is the only stochastic variable and we already know its variance from equation (2). De Long *et al.* (1990) do not supply the derivation but we have come up with the following:

$$\mathrm{Var}\left(1+\frac{\mu(\rho_t-\rho^*)}{1+r}+\frac{\mu\rho^*}{r}\right)$$

$$=\mathrm{Var}(1)+\mathrm{Var}\left(\frac{\mu\rho_t}{1+r}\right)-\mathrm{Var}\left(\frac{\mu\rho^*}{1+r}\right)+\mathrm{Var}\left(\frac{\mu\rho^*}{r}\right)$$

$$=\mathrm{Var}\left(\frac{\mu\rho_t}{1+r}\right) \tag{16}$$

$$=\mathrm{E}(\rho_t)^2 Var\left(\frac{\mu}{1+r}\right)+E\left(\frac{\mu}{1+r}\right)^2\mathrm{Var}(\rho_t)+Var\left(\frac{\mu}{1+r}\right)\mathrm{Var}(\rho_t)$$

$$=\frac{\mu^2\sigma_\rho^2}{(1+r)^2}$$

We can now express equation (14) without any expectations of the variance in the price one period ahead:

$$p_t=1+\frac{\mu(\rho_t-\rho^*)}{1+r}+\frac{\mu\rho^*}{r}-\frac{2\gamma\mu^2\sigma_\rho^2}{r(1+r)^2} \tag{17}$$

From the first term of equation (17) we see that the fundamental price of the unsafe asset is 1. The following three terms describe how it is affected by the misperceptions of noise traders. The second term of equation (17) is the present generation of noise traders' effect on the price. If their misperception of the price is either over- or underestimated compared to the average mispricing they will drive the price either up or down, respectively, just as Summers & Summers (1989) and Haberer (2003) noted in section 2.3. It is clear that the larger a fraction of noise traders (μ) populate the market the more mispriced the asset price will be and therefore more volatile. The third term is the mean misperception, which can be interpreted as the price trend or "price pressure" as De Long *et al.* (1990) label it. The fourth term represents the uncertainty about next period's noise traders' perception of risk. As this risk (σ_ρ^2) increases it drives down the price in the present period hence increasing the return of the asset from period t until t+1. It is exactly this compensation that informed traders' request to invest in asset (u), since the absence of risk created by noise traders would have made the safe and unsafe asset perfect substitutes (De Long *et al.* (1990)).

The demand functions, equation (9) and (10), and the pricing formula, equation (17), reveal many of the theoretical points discussed in section 2.3. From the demand functions, we see how noise traders are destabilizing by demanding more of the unsafe asset when their misperception of the price is positive (ρ > 0) and less when it is negative (ρ < 0), where informed traders are stabilizing. This is also shown by the fact that all variance in the model is created by the noise traders' misperceptions of the assets' price. The assumption of no fundamental risk does stress the destabilizing effect of noise traders but the assumption itself seems unrealistic and introducing a fundamental risk into the model will change some of the results, as we shall see later. The other assumption, that noise traders' and informed traders' investment horizons are both equal and short, from time t to t+1, is not necessarily the case either, and, as De Long *et al.* (1990) note, could affect informed traders choices. If informed traders could invest over longer periods of time, they could buy low more securely, knowing that over time the price of the asset would revert to its fundamental price. It does, however, correspond to the point made by Eichengreen, Tobin & Wyplosz (1995), as we discussed in section 2.3.4, that informed traders might trade exactly as often as noise traders although against them.

2.5.3 Relative expected return

Contrary to what was proposed by Friedman, as we mentioned in section 2.2.2, noise traders have the possibility of receiving higher expected returns than informed traders with this model because the increased risk weeds out informed traders as the risk of betting against noise traders becomes too large (De Long *et al.* (1990)). We shall look at what factors affect this possibility of higher expected returns. Equation (18) shows the difference in returns between noise and informed traders (ΔR_{n-i}), which is given by subtracting the second term of equation (8) from the second term of equation (7) since these represent the returns of the noise and informed traders:

$$\Delta R_{n-i} = \left(\lambda_t^n - \lambda_t^i\right)\left(r + p_{t+1} - p_t(1+r)\right) \tag{18}$$

From the demand functions, equation (9) and (10), and equation (15) we know that the difference in demands can be expressed as follows:

$$\lambda_t^n - \lambda_t^i = \frac{\rho_t}{(2\gamma)\ {}_t\sigma_{p_{t+1}}^2} = \frac{\rho_t}{(2\gamma)\dfrac{\mu^2\sigma_\rho^2}{(1+r)^2}} = \frac{(1+r)^2\rho_t}{(2\gamma)\mu^2\sigma_\rho^2} \tag{19}$$

From equation (19) we see that when the share of noise traders fall, i.e. μ decreases, the difference in demand increases. This can only be explained by noise traders and informed traders betting against each other, i.e. taking positions of the opposite sign. The reason for this is that as noise trader risk diminishes it becomes safer for each type of trader to bet against the other to try to take advantage of the arbitrage created by the different beliefs of the price movement (De Long *et al.* (1990)), just as we argued in section 2.3.4.

From the second term of equation (18) and our price function, equation (11), as well as equation (15), we can express the expected return on the risky asset as follows:

$$
\begin{aligned}
&{}_t\big(r + p_{t+1} - p_t(1+r)\big) \\
&= \left(r + {}_tp_{t+1} - (1+r)\frac{1}{1+r}\{r + {}_tp_{t+1} - 2\gamma\{{}_t\sigma^2_{p_{t+1}}\} + \mu\rho_t\}\right) \\
&= \big(r + {}_tp_{t+1} - \{r + {}_tp_{t+1} - 2\gamma\{{}_t\sigma^2_{p_{t+1}}\} + \mu\rho_t\}\big) \qquad (20) \\
&= 2\gamma\{{}_t\sigma^2_{p_{t+1}}\} - \mu\rho_t \\
&= \frac{(2\gamma)\mu^2\sigma^2_\rho}{(1+r)^2} - \mu\rho_t
\end{aligned}
$$

The expected return is thus increasing in the risk factor and decreasing in the noise traders' misperception of prices. Since traders require a higher risk premium when the risk (σ^2_ρ) increases, as expressed by the first term of equation (20), the return will increase. The second term of the equation explains that, as noise traders' misperception of prices rise, the price of the asset rises, resulting in a lower expected return. The increasing misperception will attract more noise traders but deter informed traders from demanding the unsafe asset, resulting in the noise traders and the remaining informed traders being willing to pay a higher price for the asset. Combining equation (19) and (20) gives us equation (18) expressed as:

$$
{}_t(\Delta \mathrm{R}_{n-i}) = \frac{(1+r)^2\rho_t}{(2\gamma)\mu^2\sigma^2_\rho}\left(\frac{(2\gamma)\mu^2\sigma^2_\rho}{(1+r)^2} - \mu\rho_t\right) = \rho_t - \frac{(1+r)^2\rho^2_t}{(2\gamma)\mu\sigma^2_\rho} \qquad (21)
$$

Taking the expected value of equation (21) enables us to evaluate the situations in which noise traders earn a higher expected return than informed traders. De Long *et al.* (1990) do not show these calculations but it is easily done when realizing the following:

$$
\begin{aligned}
\sigma_\rho^2 &= E[(\rho_t - \rho^*)^2] \\
&= E(\rho_t^2) - E(2\rho^*\rho_t) + E((\rho^*)^2) \\
&= E(\rho_t^2) - 2(\rho^*)^2 + (\rho^*)^2 \\
&= E(\rho_t^2) - (\rho^*)^2 \\
&\leftrightarrow E(\rho_t^2) = (\rho^*)^2 + \sigma_\rho^2
\end{aligned} \tag{22}
$$

By substitution we get that:

$$
\begin{aligned}
E(\Delta \mathrm{R}_{n-i}) &= E\left(\rho_t - \frac{(1+r)^2\rho_t^2}{(2\gamma)\mu\sigma_\rho^2}\right) \\
&= E(\rho_t) - \frac{(1+r)^2}{(2\gamma)\mu\sigma_\rho^2}E(\rho_t^2) \\
&= E(\rho_t) - \frac{(1+r)^2}{(2\gamma)\mu\sigma_\rho^2}\left((\rho^*)^2 + \sigma_\rho^2\right) \\
&= \rho^* - \frac{(1+r)^2(\rho^*)^2 + (1+r)^2\sigma_\rho^2}{(2\gamma)\mu\sigma_\rho^2}
\end{aligned} \tag{23}
$$

Equation (23) incorporates many of the aforementioned effects in describing the conditions necessary for noise traders to earn a higher expected return than informed traders. There are basically two effects working against them and two working for them. The first term of equation (23) is the mean misperception of the noise traders. As this increases, noise traders' holdings of the risky asset increases and their return becomes larger compared to the informed traders who will not hold more of the asset. The reason for this is that it becomes too risky for informed traders to bet against the noise traders.

When noise traders' misperception of prices increases it also increases the price of the asset resulting in an overall decrease in the expected return. This is what the first term of the numerator expresses and exactly what we saw in our pricing formula, equation (17).

The second term of the numerator describes how noise traders' relative expected return decreases as the variance in their misperceptions of the price increases, although this is offset somewhat by the denominator of equation (23). The expression in the denominator works in the same way as it did in equation (17). As the noise traders' risk increases the informed traders are driven out of the market to a higher extent than the noise traders themselves. This furthermore diminishes the two effects of the numerator as the price is driven down because of the increased risk. Overall, this adds to an increase in the noise traders' return relative to the informed traders'.

The possibility of receiving higher expected returns as a noise trader suggests that it might be attractive for new traders to shift strategy to become either a noise trader or an informed trader depending on what strategy has performed best in the past.

De Long *et al.* (1990) express this imitation of beliefs, as they call it, in the following way:

$$\mu_{t+1} = max\left(0, min\left(1, \mu_t + \zeta(R_n - R_i)\right)\right) \tag{24}$$

Equation (24) says that the number of noise traders in the next period is a function of the difference in return between the noise traders and the informed traders as well as some rate (ζ) at which new traders become noise traders, and that $0 \leq \mu_{t+1} \leq 1$. The possibility of imitation makes the static share of traders (μ) in all previous equations stochastic (μ_t). De Long *et al.* (1990) assume ζ to be close to 0 to keep their equations simple. When ζ increases, young traders in period t have to estimate the effects of their realized returns on young traders' decisions in period t+1. Keeping ζ close to 0 enables us to approximate μ_{t+1} with μ_t. The former relative expected return of equation (21) then becomes:

$$E_t(\Delta \mathrm{R}_{n-i}) = \rho_t - \frac{(1+r)^2 \rho_t^2}{(2\gamma)\mu_t \sigma_\rho^2} \tag{25}$$

From equation (25) we see that as μ_t decreases informed traders' expected return becomes relatively larger than noise traders' expected return. The reason being that informed traders will bet larger amounts against the positions of the noise traders and in the long run crowd out noise traders. The opposite is the case when μ_t becomes larger, and in the long run noise traders will crowd out informed traders.

2.5.4 Introducing fundamental risk

If we relax the assumption that there is no fundamental risk, and assume instead that asset (u) pays an uncertain dividend:

$$r + \varepsilon_t \tag{26}$$

where ε_t is serially independent and normally distributed with zero mean and constant variance and uncorrelated with ρ_t, we will see that the additional risk term σ_ε^2, which represents the fundamental risk, is added to the risk of the unsafe asset. This increases the total risk of asset (u) and puts a lower bound on the informed traders' demand for the risky asset and an upper bound on the noise traders' demand, as opposed to the case with no fundamental risk where their demands were unbounded. We know from equation (19) that the traders' demands for the risky asset are unbounded in the absence of a fundamental risk; however the introduction of a fundamental risk bounds the traders' demands as μ approaches 0, as we see in equation (27):

$$\lambda_t^n - \lambda_t^i = \frac{\rho_t}{(2\gamma)\left({}_t\sigma_{p_{t+1}}^2 + \sigma_\varepsilon^2 \right)} = \frac{\rho_t}{(2\gamma)\left[\sigma_\varepsilon^2 + \frac{\mu\sigma_\rho^2}{(1+r)^2}\right]} \tag{27}$$

As the share of noise traders approaches 0, the difference in demand for the risky asset becomes at least: $\frac{\rho_t}{(2\gamma)\sigma_\varepsilon^2}$. This is an intuitive result as the number of noise traders approaches 0 they cannot hold the entire share of the risky asset due to the increasingly large holding of risk. Therefore, the informed traders are bound to have a positive holding of the risky asset. This leads to a decline in the fundamental price of the risky asset, which becomes:

$$p_t = 1 + \frac{\mu(\rho_t - \rho^*)}{1+r} + \frac{\mu\rho^*}{r} - \frac{(2\gamma)}{r}\left[\sigma_\varepsilon^2 + \frac{\mu\sigma_\rho^2}{(1+r)^2}\right] = 1 - \frac{(2\gamma)}{r}\sigma_\varepsilon^2 \tag{28}$$

Thereby it gives a larger expected return than the safe asset because p_t will be less than 1, cf. equation (28). The boundary that the fundamental risk creates ensures that $E(\Delta \mathrm{R}_{n-i}) \leq \rho^*$ where it previously could be arbitrarily large. This can be seen by introducing the fundamental risk into equation (23):

$$E(\Delta \mathrm{R}_{n-i}) = \rho^* - \frac{(\rho^*)^2 + \sigma_\rho^2}{(2\gamma)\left[\frac{\sigma_\varepsilon^2}{\mu} + \frac{\mu\sigma_\rho^2}{(1+r)^2}\right]} \tag{29}$$

De Long *et al.* (1990) show that certain situations warrant noise traders to earn larger expected returns than informed traders, as Frankel (1996) also mentions in section 2.3.3. De Long *et al.* (1990) further go on to prove that as ζ approaches 0 the share of noise traders in the market approaches 1. This proof suggests that under certain conditions noise traders have an influence on the price development in the long run.

2.5.5 Implications of expected utility

Until now the points made in the past sections have all been based on the assumption that traders want to maximize their return when in fact they might want to maximize their utility as we mentioned in the beginning of section 2.5. Since informed traders want to maximize their true expected utility, any rise in the noise traders' expected returns must be followed by an increase in risk large enough to result in a lower level of utility, if not, informed traders would follow the strategy of the noise traders. De Long *et al.* (1990) explain that noise traders' utility falls or rises the same amount when ρ^* either rises or falls by some value. Furthermore, when $\rho^* > 0$ noise traders have the possibility of earning higher expected returns than informed traders but when $\rho^* < 0$ noise traders will earn both lower expected returns and utility compared to informed traders.

If traders focused on maximizing their utility they would also base their imitations of belief on the previous utility levels of the two different trading strategies. As we saw in equation (3), the traders have concave utilities, which mean that a decline in returns is punished more severely in the level of utility than an increase. Therefore, new traders would switch away from a strategy in much larger numbers in the case of past low returns, than they would switch to a strategy in case of past high returns. As we described before, informed traders will on average maximize their true expected utility over time, which in the long run will give them a higher utility than noise traders since their misperceptions are normally distributed and their negative misperceptions are punished more than their positive misperceptions are rewarded in realized utility. De Long *et al.* (1990) show that when ζ approaches 0 the share of noise traders in the economy will disappear.

2.5.6 Introducing a security transaction tax

From De Long *et al.*'s (1990) paper, which was discussed in the last sections, we know how both noise traders and informed traders behave under certain conditions. Song & Zhang (2005) expand De Long *et al.*'s (1990) model with fundamental risk to include both a liquidity risk (σ_θ^2), to capture uncertainties about the liquidity of the risky asset from changes in the trader composition as well as a security transaction tax (τ) to see how this affects the volatility of the risky asset. The liquidity risk is meant to resemble the view of Friedman, namely that a decrease in trading makes the price discovery process more rigid thereby creating more volatility. Song & Zhang (2005) find that the price volatility of the risky asset (V) can be expressed as:

$$V \equiv {}_t\sigma^2_{p_{t+1}} = \frac{\mu^2\sigma^2_\rho}{(1+r)^2(1+\tau)^2} + \frac{4\gamma^2\sigma^2_\theta}{(1+r)^2(1+\tau)^2}\left({}_t\sigma^2_{p_{t+1}} + \sigma^2_\varepsilon\right)^2 \tag{30}$$

Rearranging equation (30), Song & Zhang (2005) obtain an expression that relates the ratio of noise traders (μ) to volatility (V):

$$\mu^2\sigma^2_\rho = (1+r)^2(1+\tau)^2 V - 4\gamma^2\sigma^2_\theta(V+\sigma^2_\varepsilon)^2 \tag{31}$$

Plotting this nonlinear function between μ and V into a graph at three different tax levels, 0, 0.5 and 1 per cent, we get Figure 3:

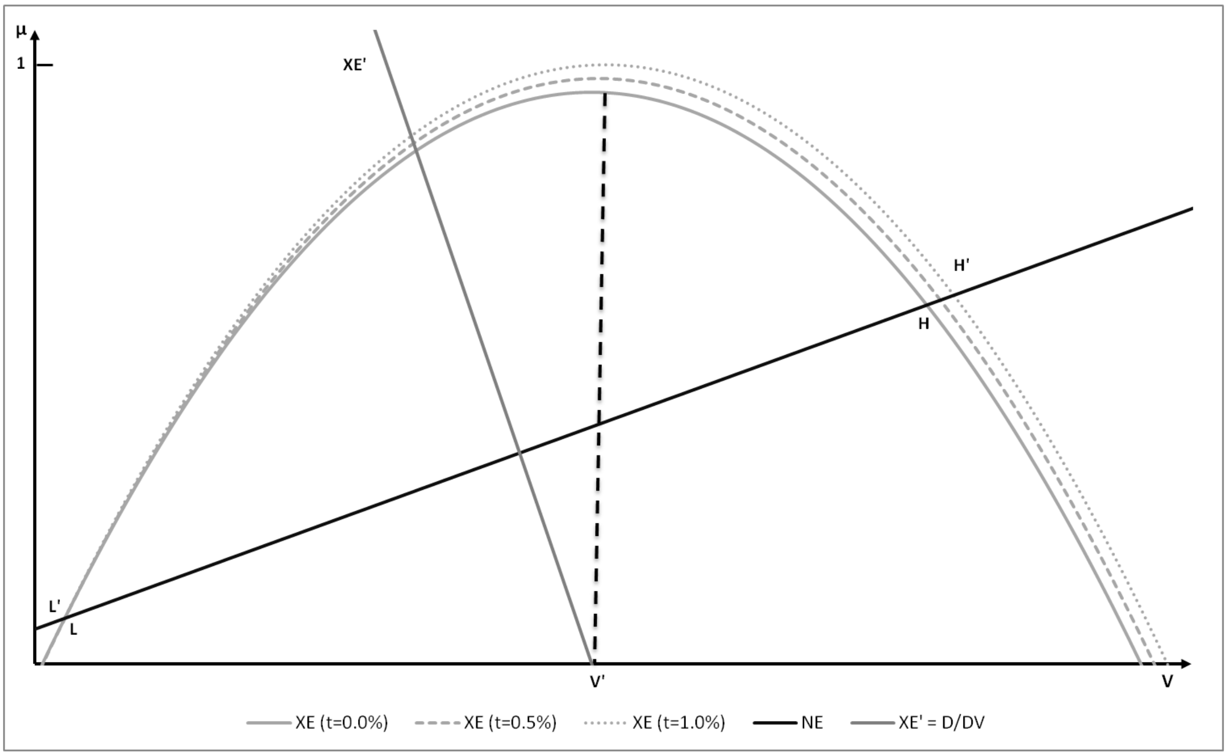

Figure 3 - The ratio of noise traders (μ) to volatility (V) with a 0, 0.5 and 1 per cent tax level.

Looking at Figure 3 and equation (31) it is noticeable that when V approaches 0 the tax (τ) loses its impact, but we assume that Song & Zhang (2005) do not expect that the volatility can reach 0, and have therefore not commented on it. Furthermore, from Figure 3 we see that when τ increases the XE curve shifts upward, which is also expressed by equation (30). We also find that on the left branch of the XE curves an entry of additional noise traders increases volatility while on the right branch we see the opposite. When XE is at its maximum the share of noise traders equal 1. Song & Zhang (2005) argue that there are two forces affecting the number of new entrants and the volatility of the asset. An increase in noise traders on the market also leads to an increase in the amount of noise, exactly as we saw in the second term of equation (14). The increasing noise makes it more difficult for the existing traders to predict the future value of the asset and so the marginal noise trader entering the market has a *negative externality effect* on the pre-existing traders. As the number of noise traders rise and their negative effect decreases, as seen by the line XE', the ability to share the added risk among the noise traders increases. Therefore, if looking at Figure 3 it appears that for a low level of volatility the negative externality effect dominates and for levels of volatility larger than V' the *risk-sharing effect* outweighs the negative externality effect. In this case speculative trading can become stabilizing, as volatility decreases in line with the increased share of noise traders.

Song & Zhang (2005) find the optimal entry decisions for the noise traders, by setting their modified version of equation (29) equal to 0:

$$E(\Delta \mathrm{R}_{n-i}) = \rho^* - \frac{\mu\left({\rho^*}^2 + \sigma_\rho^2\right)}{2\gamma(V + \sigma_\varepsilon^2)} \tag{32}$$

They further assume that a noise trader enters if $E(\Delta \mathrm{R}_{n-i}) > 0$, that he will not enter if $E(\Delta \mathrm{R}_{n-i}) < 0$ and that the noise trader will be indifferent between entering or not when $E(\Delta \mathrm{R}_{n-i}) = 0$. Solving for the share of noise traders (μ) gives us:

$$\mu = \frac{2\rho^*\gamma}{(\rho^*)^2 + \sigma_\rho^2}(V + \sigma_\varepsilon^2) \tag{33}$$

Plotting this into Figure 3, denoted by NE, we see that above the line the noise traders earn higher returns than the informed traders and below the line we see the opposite. At the line, where NE intersects the XE curve, the noise traders are indifferent between entering or not, as the return of both noise and informed traders are the same. This is illustrated by the two equilibriums L and H.

As described in section 2.2.2 an increase in the tax (τ) tends to reduce the demand from both the noise traders and the informed traders. This can also be seen from equations (9) and (10) if the tax term (τ) is added:

$$\lambda_t^i = \frac{d + {}_t p_{t+1} - (1+r)(1+\tau)p_t}{2\gamma\{ {}_t\sigma^2_{p_{t+1}} + \sigma^2_\varepsilon\}} \tag{34}$$

and

$$\lambda_t^n = \frac{d + {}_t p_{t+1} - (1+r)(1+\tau)p_t}{2\gamma\{ {}_t\sigma^2_{p_{t+1}} + \sigma^2_\varepsilon\}} + \frac{\rho_t}{2\gamma\{ {}_t\sigma^2_{p_{t+1}} + \sigma^2_\varepsilon\}} \tag{35}$$

Even though this affects both types of traders it will affect the informed traders more as they are more sensitive to changes in τ. As λ_t^i decreases more than λ_t^n we see from equation (11) that the share of noise traders must increase, since the holdings of the unsafe asset has to equal 1. The increase in the share of noise traders to informed traders brings about an increase in volatility due to the added uncertainty created by the increased misperception of prices. This affect is dubbed the liquidity effect by Song & Zhang (2005), which corresponds to the theory of Friedman (1953).

From Figure 3 we see that an increase in the tax shifts the EX curve upward, causing new equilibriums where noise traders are indifferent to entering the market or not. In the equilibrium L, the effect of a higher tax will lead to a decrease in the share of noise traders and lower the volatility creating a new equilibrium L'. When moving from equilibrium H to H' the effect will increase volatility described by the liquidity effect earlier in this section.

Song & Zhang (2005) conclude that in the presence of noise trader risk, fundamental risk and liquidity risk the view of both Keynes and Friedman can exist when introducing a Tobin tax, depending on the trader composition, as we just saw from Figure 3. On one hand they argue that when the share of noise traders in the market is low and the negative effect of a marginal noise trader dominates the risk-sharing effect, a Tobin tax will drive some of the existing noise traders out of the market, moving the equilibrium from L to L'. This will reduce the volatility since there will be less risk to share between the remaining traders. When the risk decreases it becomes safer for informed traders to enter the market and bet against noise traders, wherefore the share of informed traders compared to noise traders increases and μ decreases, as seen in equation (19). On the other hand they argue the opposite if moving from equilibrium H to H', where the share of noise traders dominate the informed traders. Based on the market behaviour explained in this section it is clear that the tax effect market volatility in different ways depending on the market composition and the pre-existing level of volatility. Song & Zhang (2005) further point out that minor economies often are associated with lower volatility and a smaller share of noise traders, whereas major economies are often related to higher volatility.

2.5.7 Summary

The first and most basic model proposed by De Long *et al.* (1990) was based on the assumption of no fundamental risk in the market. This showed how noise traders alone could create risk by their misperception of prices. The destabilizing nature of the noise traders means that the larger a fraction of noise traders the more volatile the market will be, as proposed by Tobin and Keynes, cf. sections 2.1 and 2.2.1, respectively. Furthermore, De Long *et al.* (1990) assume that both noise traders' and informed traders' investment horizons are equally long, which contradicts Tobin's and Keynes' view of the market where noise traders trade more frequently than informed traders. Eichengreen, Tobin & Wyplosz (1995) although points out that this might be the case, since informed traders could use noise traders' frequent mispricing to bet against them.

When noise traders and informed traders bet against each other in a market with no fundamental risk, De Long *et al.* (1990) shows that noise traders will receive a larger expected return compared to informed traders when the noise traders mean misperception of the price is positive, and vice versa when their mean misperception is negative. De Long *et al.* (1990) assume that if traders are return maximizing, contrary to utility maximizing, new traders will imitate noise traders or informed traders with equal probability. This is due to noise traders' misperception of prices being normally distributed giving equal probability to positive or negative returns. They go on to show that when fundamental risk is introduced to the market, noise traders' will tend to earn higher expected returns because of the added risk, which informed traders will be reluctant to take on. This results in noise traders being able to receive higher expected returns in the long run, wherefore imitating new traders will adopt the noise trading strategy and crowd out informed traders over time. If traders on the contrary are utility maximizing then informed traders will crowd out noise traders in the long run, as noise traders are punished more severely for their negative returns. Although, De Long *et al.* (1990) believe that traders are most likely return maximizing.

Song & Zhang (2005) introduces liquidity risk and a transaction tax to the model. They find two equilibriums in the relationship between volatility and the share of noise traders, where noise traders and informed traders expected returns are equal. One for the case of low volatility (L), which Song & Zhang (2005) associates with minor economies and one for the case of high volatility (H), which they associate with major economies. When introducing a tax the share of noise traders will rise in H. This is due to a larger decrease in informed traders' demand than noise traders', which increases the share of noise traders creating more misperception and uncertainty of the unsafe asset's price, i.e. increasing the volatility. For the equilibrium L the opposite applies.

3 Methodology

The following subsections of chapter 3 will introduce the reader to the variables used in the forthcoming regression model based on a discussion of previous research methods and their statistical strength. Section 3.1 discusses the pros and cons between different methods of analysis, such as time series or inter-temporal. Sections 3.2.1 and 3.2.2 discuss the pros and cons between different measures of transaction cost and volatility, respectively. In section 3.3 we describe what factors determine the bid-ask spread and in section 3.4 we note the problem of endogeneity and describe why we find previous approaches to solving this problem insufficient as well as what methods that possibly could solve the problem.

3.1 Method of study

Even though Tobin's suggestion was intended for currency transactions, many studies on stocks have been made. The reason for this is that information regarding stocks is easily accessible. Since stocks are traded on organized exchanges one can easily obtain data on volume, turnover, prices and much more, which is not publicly available for currencies. This of course requires the researcher to assume that the currency and stock market are comparable and that volatility in the currency market reacts to changes in costs in the same manner as stocks. We do not see any reason why the volatility in currencies should react in the opposite manner as stocks and therefore we find it reasonable to compare the results on studies based on stocks with our expectations for the results of our own study on currencies.

Since Tobin proposed his security transaction tax in 1972, several attempts have been made to determine the statistical relationship between the transaction costs of trading and the liquid asset's volatility, thereby validating some of the more theoretical arguments from section 2.3. These studies can be divided into two categories.

First, there is the single data point or inter-temporal method used by, amongst others, Umlauf (1993), Jones and Seguin (1997), Ronen and Weaver (2001) and Bessembinder (2003). In these type of studies the volatility before and after some arbitrary event is compared to see if there is any change. Umlauf (1993) measures the volatility around the introduction of first a 1% tax on transactions and later an increase of the tax to 2% as mentioned in section 2.1. In Jones and Seguin (1997) the arbitrary event is a one-time decrease in the brokers' commission on the New York Stock Exchange (NYSE) or the American Stock Exchange

(AMEX). Ronen and Weaver (2001) and Bessembinder (2003) both look at a sudden change in the tick size on stock prices.

The historical findings from the first method all find a positive relationship between transaction costs and volatility - although some of the results seem to be rather weak. Umlauf (1993) only finds that an introduction of a tax on the Swedish stock market does not decrease volatility on a set of Swedish stocks compared to the same set of stocks traded on the London stock exchange. This cannot be considered a direct positive relationship between cost and volatility, but neither a negative one. Jones and Seguin (1997), on the other hand find a positive relationship between cost and volatility on stocks traded on NYSE and AMEX, although they realise that their use of brokers' commissions as a proxy for transaction costs might not be sufficient. Furthermore they find the same reduction in volatility on Nasdaq, which they use as a control market. Ronen and Weaver (2001), and Bessembinder (2003) who use one-time tick size changes also find that a reduction in the tick size results in a reduction in volatility. It should be mentioned that a reduction in tick sizes gives way for a reduction in the bid-ask spread, which can be used as a proxy for transaction costs (Harris (1999)). We will explain why this is the case in section 3.2.1.

Second, there are studies based on time-series where the arbitrary event is not determined by a single point in time, but rather a characteristic of the asset itself, such as price. Bessembinder (2000), Bessembinder and Rath (2002) and Hau (2006) also use tick size changes but in these studies the change in tick size is dependent on the price of the stock. Bessembinder (2000) looks at the volatility of stocks on Nasdaq when the prices of the stocks are either above or below $10 where the tick size changes from 1/8 dollar to 1/32 dollar, respectively. Hau (2006) exploits the same situation on the French stock market where the tick size increases from 0.1 French francs to 1 French franc when the stock price increases above 500 French francs. Mulherin (1990) evaluates stock price volatility over a period from 1934 to 1987 by using the margin requirement as a measure of costs and Aliber, Chowdhry & Yan (2003) define discrepancies in the interest rate parity as transaction costs of currency futures.

Bessembinder (2000), Bessembinder and Rath (2002), Aliber, Chowdhry & Yan (2003) and Hau (2006) who all use changes that are independent of time come to the same conclusion; that an increase in costs will result in an increase in volatility. The same is the case for Mulherin (1990) - although he describes his results as somewhat weaker.

These different approaches have pros and cons, which are usually a trade-off between ease of implementation and the statistical power of their results. The following discussion of the pros and cons of the methods relies a great deal on Hau's (2006) critique of the methods that preceded his own study. The first method of using a single data point has the advantage of being very easy to implement. There are quite a few examples in history of one time regulatory changes, which is also shown by the amount of studies that utilize this method. The problem is that the method assumes that nothing but the one regulatory change affects volatility. This assumption might be valid if the volatility is measured very closely before and after the event, though over time, especially after the event, other factors such as changes in the underlying market will most likely also influence the volatility. Since the change is not measured separately from a fixed point in time it is difficult to conclude anything about the effects on time-varying volatility.

The second method can be divided into two parts. Namely the use of tick size changes alone and the additional use of other explanatory variables e.g. market volatility and volume. The use of only tick size changes is a stronger statistical method than the one described above. The reason for this is that the analysis is based on characteristics of the stock itself and is in this way independent of time. Even though the method is stronger statistically than the single data point method, it is still difficult to conclude that the change in tick size should be the only explanation for the change in volatility. Hau (2006) tries to correct this problem by using a model based on panel data where he studies the effect of a tick size change as well as the effect of other variables such as lagged volatility and several dummies. Aliber, Chowdhry & Yan (2003) also set up a time-series model where they regress volatility on their measure of transaction costs, lagged volatility, volume and a time trend. The difficulty with this sort of method lies in obtaining the data for all the explanatory variables as well as in the increase in possible statistical problems such as multicollinearity.

The results obtained from previous research seem to all paint the same picture, of a positive relationship between transaction costs and volatility in currency and stock prices. These results point in the direction that informed traders dominate the market, which is somewhat different from the majority of researchers' conclusions on the subject, as we saw in section 2.3.

We choose to base our forthcoming model on a multiple variable time-series model, as this, according to the above discussion, is the strongest statistical method. We would, although, expect to see a similar result as the aforementioned papers and find either a positive significant or a negative insignificant relationship.

3.2 Method of measurement

In this section we will discuss the pros and cons of different measures of cost and volatility used in previous research, which will lead to the measure that we will use in our forthcoming models.

3.2.1 Cost

When using a time-series model where our measurements of cost and volatility are not dependent on a fixed point in time, but rather accepts that volatility as well as cost is time varying there are many different ways to measure these variables.

In the early years after Tobin's suggestion of a security transaction tax Frenkel & Levich (1975, 1977 & 1979) used deviations in the theoretical triangular arbitrage model as a measurement for transaction costs when trading currency. From the standard equilibrium model we know that

$$\frac{F-S}{S}=\frac{i-i^*}{i+i^*} \tag{36}$$

where F is the forward exchange rate, S is the spot exchange rate and i and i* are the domestic and foreign interest rate, respectively. Frenkel & Levich (1975) then introduce a cost, defined as:

$$\Omega=(1-t)(1-t_S)(1-t^*)(1-t_F) \tag{37}$$

into the equilibrium which creates the possibility of a neutral band around the original equilibrium. The total amount of costs (Ω) is a multiple of a transaction cost on the sale of domestic securities (t), a purchase of foreign currency at spot price (t_S), a purchase of foreign securities (t*) and a sale of foreign currency at forward price (t_F) (Frenkel & Levich (1975)). Without going into too much detail about Frenkel & Levich's derivation of their final model it can be written that they end up with the following lower and upper limit for when it is profitable to trade in currencies with a transaction cost (Frenkel & Levich (1977)):

$$\frac{\Omega(1-i)-(1+i^{*})}{1+i^{*}} \leq p \leq \frac{(1-i)-\Omega(1+i^{*})}{\Omega(1+i^{*})} \tag{38}$$

where p is the forward premium, or left hand side of equation (36). To be able to define the band of equation (38) as the cost of trading Frenkel & Levich have to assume that the currency market is efficient enough to eliminate all profit opportunities in the market (Frenkel & Levich (1975)). When doing so, discrepancies in the equilibrium of equation (36) can be assumed to be transaction costs alone.

An empirical look at the difference between the left and right hand side of equation (36) would reveal any transaction costs; cf. Frenkel & Levich's assumptions, which could then be held up against the volatility of the currencies in question. We find this method of measurement for our costs rather troublesome compared to some of the other methods we will discuss shortly. As McCormick (1979) and Aliber, Chowdhry & Yan (2003) mention, a grave error in Frenkel & Levich's research is that they have not managed to use synchronous recorded prices across markets. For Frenkel & Levich this makes their results much less reliable as McCormick (1979) shows by using the same method on synchronously recorded prices. It is therefore possible to obtain synchronously recorded prices but other measurements of costs, which we will discuss in the following, will eliminate that complication completely. Further, Frenkel & Levich's assumption about the efficiency of the financial market can be avoided.

In Roll (1984) a new measure of the implied bid-ask spread, to be used as a measurement for transaction costs, is proposed. Roll (1984) assumes, as Frenkel & Levich (1975, 1977 & 1979) did, that the financial market is efficient and that only new information can change prices. Based on this assumption Roll is right to conclude that a fall (rise) in price can only be followed by a subsequent rise (fall), as no buyer (seller) would sell (buy) at a lower (higher) price than they bought (sold) it for – at least without any new information available to the market. On the basis of this assumption and the probability of the next price being either higher or lower than the previous one, he finds the covariance of price changes to be (Roll (1984)):

$$Cov(\Delta p_t, \Delta p_{t-1}) = \frac{1}{8}(-S^2 - S^2) = \frac{-S^2}{4} \tag{39}$$

where p is the price of the currency and s is the bid-ask spread faced by dollar-weighted average traders. Rearranging equation (39) we can express the spread as a function of the covariance:

$$\frac{S^2}{4} = -Cov(\Delta p_t, \Delta p_{t-1}) \Leftrightarrow S = 2\sqrt{-Cov(\Delta p_t, \Delta p_{t-1})} \quad (40)$$

From equation (40) we see that the assumption that the covariance is always negative is an important one or else the spread would not be calculable, which is also Aliber, Chowdhry & Yan's (2003) problem for some of their observations. Again, we find the assumptions that the measure is based on too weak and moreover possible to avoid. Roll (1984) notes that a comparison between his daily and weekly-based results suggest informational inefficiency, which is one of the main assumptions he makes.

Roll (1984) is right in using a proxy for the bid-ask spread as a measure for transaction costs as we shall see in a moment, and as he mentions himself, the true bid-ask spread would have been his best choice. The problem with studies done in around 1984 is that the available data was limited or expensive to obtain.

If we look at research done through the first decade of the new millennium, we see that Ronen & Weaver (2001), Hau (2006) and Lanne & Vesala (2006) are using the quoted bid and ask prices publicly available to compute the bid-ask spread as a substitute for transaction costs. One criticism of the bid-ask spread is put forth by Aliber, Chowdhry & Yan (2003). It states that the quoted bid and ask prices are in reality not those charged by the banks, and that non-bank customers might also endure other costs such as the dealers fee, possible costs for acquiring expertise, exposure management services and trading systems. As an alternative, they propose a variant of Frenkel & Levich's method where they look at the arbitrage between different forward rates and the corresponding risk free rates of interest.

Although, one reason to consider bid-ask spreads instead of any arbitrage model is that it is much easier to implement, as it is sufficient with two variables compared to four and we do not have to worry about non-synchronous recording of prices as Frenkel & Levich (1975, 1977 & 1979), as the bid and ask prices are not measured across exchanges. As done in Ronen & Weaver (2001) and Lanne & Vesala (2006), partly in Hau (2006) we look at the bid-ask spread as a measure of transaction costs. We use the spread to calculate the percentage costs (C) of the cost free fundamental price (S). To do this we assume that the spread between the bid-price (S_b) and S is proportional with the spread between the ask-price (S_a) and S so that we have:

$$\frac{S - S_b}{S} = C \Leftrightarrow S_b = S(1 - C) \tag{41}$$

$$\frac{S_a - S}{S} = C \Leftrightarrow S_a = S(1 + C) \tag{42}$$

To calculate C we have to eliminate S as it is not a given figure. To do so we isolate S in equation (42) and replace it in equation (41):

$$\begin{aligned} S_b &= \frac{S_a(1 - C)}{(C + 1)} \\ \Leftrightarrow C &= \frac{S_a - S_b}{S_b + S_a} \end{aligned} \tag{43}$$

The first thing noticeable with this measure is that the costs are calculated as a percentage of the cost-free price. This means that when we regress volatility on costs we can easily see from the estimated beta value what will happen with volatility given a one-percentage point change in costs. This will further correspond to introducing a Tobin tax, which is a percentage of the traded value. It should also be noted that the bid-ask spread requires no assumption about the market efficiency as was required by Roll (1984) and Frenkel & Levich (1975, 1977 & 1979).

3.2.2 Volatility

It seems that mainly two methods have been used to calculate the volatility in either stocks or currencies. Namely the return standard deviation suggested by French, Schwert & Stambaugh (1987) and used by Mulherin (1990), Umlauf (1993), Bessembinder & Rath (2002) and Bessembinder (2003) and the log range estimator suggested by Alizadeh, Brandt and Diebold (2002) and used by Hau (2006).

The return standard deviation is based on classic additions of variance as seen in equation (44). It usually uses daily returns on the security in question, be it stocks or a currency. One of the advantages of this measure is that it easily allows for additions of variances over a week or a month of data. Furthermore, it is a volatility measure that is commonly used in a number of financial areas (French, Schwert & Stambaugh (1987)):

$$Var(Y) = \sum_{i=1}^{N_t} Var(r_{it}) + 2\sum_{i=1}^{N_t-1} Cov(r_{it}, r_{i+1,t})$$
$$= \sum_{i=1}^{N_t} E[(r_{it} - \bar{r})^2] + 2\sum_{i=1}^{N_t-1} E[(r_{i+1,t} - \bar{r})(r_{it} - \bar{r})] \tag{44}$$

French, Schwert & Stambaugh (1987) assume the mean daily return denoted by $\bar{r} = 0$, which we find reasonable, wherefore it is easy to see that we arrive at the following:

$$\sigma_{mt}^2 = \sum_{i=1}^{N_t} r_{it}^2 + 2\sum_{i=1}^{N_t-1} r_{it} r_{i+1,t} \tag{45}$$

In equation (44) variance is denoted by σ^2, r is the daily return and N_t is the number of days in the period, e.g. month.

The other volatility measure is the log range estimator, which is based on the high and low prices within a given time period, e.g. an hour, as in Hau (2006), and originally expressed in Alizadeh, Brandt and Diebold (2002) as:

$$\ln\left|f\left(s_{iH,(i+1)H}\right)\right| = \ln\left(\sup_{iH<t<(i+1)H} s_t - \inf_{iH<t<(i+1)H} s_t\right) \tag{46}$$

where iH and (i+1)H are the time interval and sup s and inf s are the maximum and minimum prices of the security within the interval, respectively. This measure has several advantages over the return standard deviation. First, as Alizadeh, Brandt and Diebold (2002) argue, it is more efficient because it measures the variance based on the highest and lowest prices within a time interval instead of the opening and closing prices. In this way, it captures high volatility within the time interval even if the opening and closing prices are close to each other. Second, the natural logarithm gives the range near Gaussian qualities. Third and last, the range is a more robust measure in dealing with the so-called bid-ask bounce. The bid-ask bounce is mostly present in the stock market where tick sizes are defined externally, opposed to the currency market's PIPS, which are defined by the market, as we will explain in section 3.3. Ticks or PIPS result in small biases from the true price in every transaction as the bid or ask price will resemble the closest tick or pip to the true price (Hau (2006)), which we will explain more detailed in section 3.3. We will not go into details, but Alizadeh, Brandt and Diebold (2002) show that as transactions tend to go back and forth between the ask and the bid price, this deviation from the true price drives up the volatility although far less when using the range estimate than with the standard deviation.

Much of the research in this section uses some form of transaction costs measure, which is either a one-time change in regulations or a tick size change prompted by a change in the value of the security. Either way, Umlauf (1993), Jones and Seguin (1997), Ronen and Weaver (2001) and Bessembinder's (2003) measures of volatility could be compared before and after any of these events to test for differences. Our time-series method requires us to use comparable measures of both transaction costs and volatility to enable us to interpret how changes in the former effects the latter. On this account the log range estimator of volatility and its many advantages have inspired us to use a variant of said method.

To keep the notation and structure of our transaction costs measure we continue with Hau's (2006) range estimate:

$$RANGE_t = \frac{max_{(s \in I_t)}\, S - min_{(s \in I_t)}\, S}{max_{(s \in I_t)}\, S + min_{(s \in I_t)}\, S} \quad (47)$$

where $max_{(s \in I_t)}\, S$ is the maximum price over a fixed interval I_t and $min_{(s \in I_t)}\, S$ is the minimum price. Equation (47) is the non-logged variant of equation (46) expanded to give us the percentage volatility of the mid-price, S^{mid}, within the interval, I_t. We assume that the distance between S^{max} and S^{mid} equals the distance between S^{min} and S^{mid}, the same method used in our cost measure. This is why we do not find it necessary to multiply the range by two like in Hau (2006). Moreover, Hau (2006) also takes the natural logarithm of the range, as Alizadeh, Brandt, and Diebold (2002). For now, we will continue with the range found in equation (47) as taking the natural logarithm only complicates the interpretation of the forthcoming models' results, and further that it only features a near normal distribution. We will look into the matter of the natural logarithm in section 6.4.2.3.

3.3 Bid-ask spread determination

The determination of the bid-ask spread is based on several factors. The two main factors are the dealer fee and the cost following the risk of holding the asset. The dealer fee covers the service the dealer provides, possible costs for acquiring expertise, exposure management services and trading systems. The costs of holding the asset can be seen as a risk premium to the dealer to compensate for the risk he takes while the transaction occurs. Furthermore, size and number of trades for a given currency are also potential factors when measuring the risk the dealer imposes by holding the asset. The more liquid the market for a given currency is the easier it is to sell, and therefore the risk of holding it decreases and thereby the bid-ask spread narrows.

Sarno & Taylor (2001) add a third factor: the cost of adverse selection. They believe that an adverse-selection problem arises because the dealers cannot distinguish between liquidity-motivated and insider traders so the dealers have to widen the spread. This results in a gain from trading with the liquidity-motivated traders, which should cover the loss from trading with the insider traders, and thereby work as a defence against adverse-selection. Even though this argument seems very interesting, Sarno & Taylor (2001) adds that this is not something that can be measured and therefore has not been proven.

Furthermore, the way the spread widens and narrows is dependent on pips. Pips are the smallest price variation that a particular exchange rate can process; in other words, it is the minimum increment in which the exchange rate can change. For example, 1 pip = 0.0001 for EUR/USD, or 0.01 for USD/JPY. This means that if the true fundamental exchange rate in the case of EUR/USD is priced with more than four decimals the bid and ask prices would have to be rounded up or down to the nearest pip - either the maximum or the minimum pip. This will widen or narrow the spread, and therefore be biased compared to the true spread matching the true amount of risk where more than four decimal might be used. The scenario where the spread will widen because the nearest pip lies below the bid price and above the ask price is displayed in Figure 4.

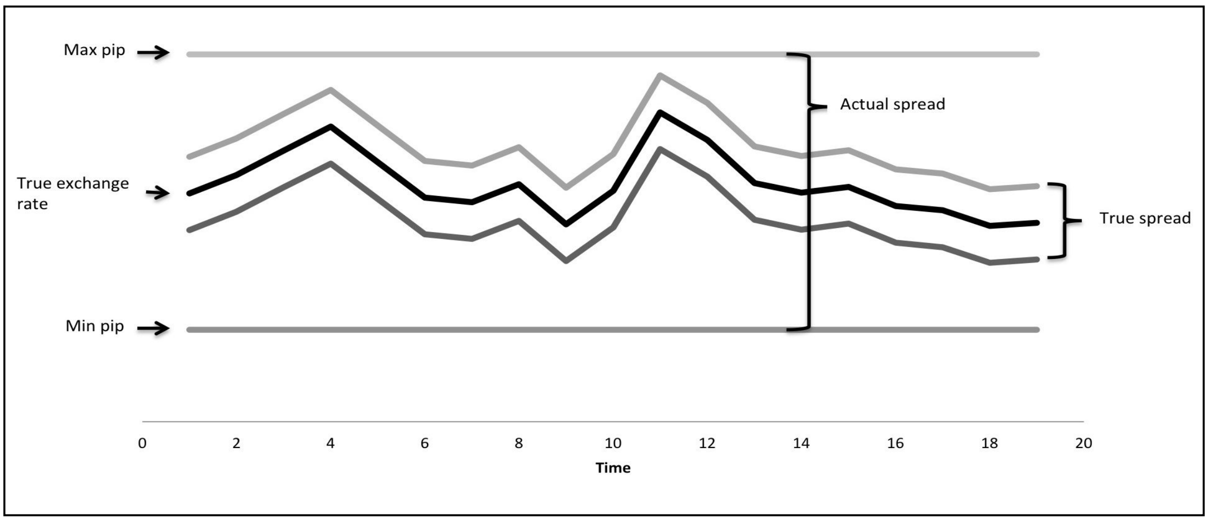

Figure 4 - Bid-ask spread. The figure shows how the true spread in this case will widen if the exchange rate is priced with more than four decimals.

3.4 The problem of endogeneity

A general problem when researching the relationship between transaction costs and volatility in the currency market is the existence of endogeneity. On the one hand a change in the transaction costs, and thereby a change in the bid-ask spread, will lead to a change in the behaviour of the market participants, and the following increase or decrease in trading will lead to a change in volatility as explained in section 2.5. On the other hand, a change in volatility will induce dealers to adjust the bid-ask spread to cover the new risk as we explained in section 3.3. The inherent relationship between costs and volatility is shown in Figure 5.

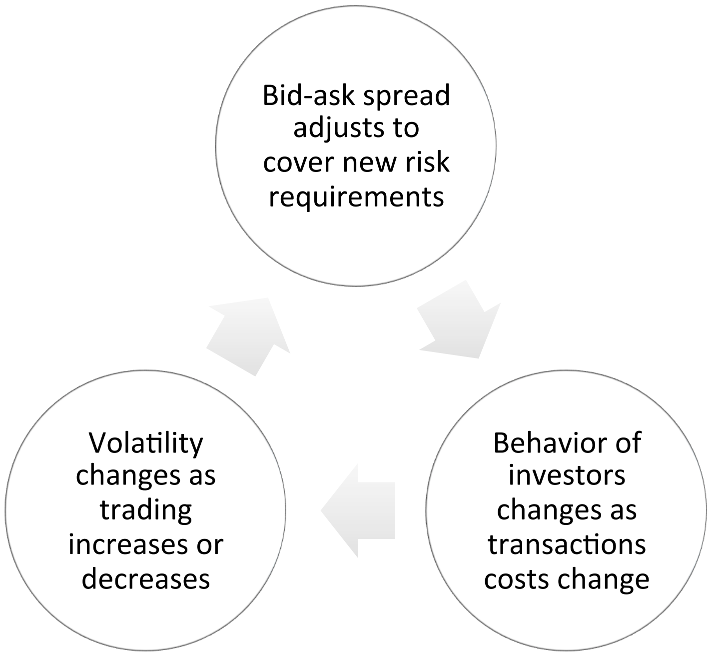

Figure 5 – The endogeneity problem inherent in the relation between transaction costs and volatility in the currency market.

When endogeneity between the dependent and independent variable is a natural result of their theoretical relationship, it becomes increasingly difficult to conclude anything about the direction of causality and the regression results become difficult to interpret.

The problem mainly exists when both volatility and transaction costs are considered as time varying and are measured over time. Inter-temporal research methods or tick-size change methods does not display this problem as the change in transaction costs is exogenous.

Few previous attempts at correcting for this problem have been done in time-series analyses, and the attempts made seem insufficient to us. Aliber, Chowdhry & Yan (2003) who use discrepancies in the interest rate parity as a measure of costs, argue that the transaction costs of exchange rate futures ($C_{f,t}$) could be affected by the fundamental volatility (σ_t), i.e. the non-excessive volatility, and some exogenous change in transaction costs, i.e. the error-term in equation (48).

$$C_{f,t} = a_f + b_f\sigma_t + e_{f,t} \tag{48}$$

They go on to say that they can model the exchange rate futures' volatility as a function of fundamental volatility as well as the exogenous change in transaction costs, based on the assumption that volatility is also explained by transaction costs.

$$\sigma_{f,t} = \alpha_f + \beta_f\sigma_t + \gamma_f e_{f,t} + \varepsilon_{f,t} \tag{49}$$

They assume that the same relations are true for the spot rate, where the symbols in equations (50) and (51) are the same as before although denoted by a subscript s to signify their representation of the spot rate.

$$C_{s,t} = a_s + b_s\sigma_t + e_{s,t} \tag{50}$$

$$\sigma_{s,t} = \alpha_s + \beta_s\sigma_t + \gamma_s e_{s,t} + \varepsilon_{s,t} \tag{51}$$

Since the fundamental volatility is not observable in the market they use the spot rate volatility as a proxy in the modelling of the exchange rate futures volatility.

$$\sigma_{f,t} = \theta_0 + \theta_C C_{f,t} + \theta_s\sigma_{s,t} + {\varepsilon'}_{f,t} \tag{52}$$

They further assume that exogenous changes in transaction costs have zero mean and some undefined standard deviation. We understand this to mean that exchange rate futures transaction costs are only affected by fundamental volatility in the long run and that speculating in futures is synonymous with the trading behaviour of informed traders'. This assumption seems reasonable for the futures market as it covers a longer time horizon than speculation in the spot rate market, and therefore speculation in the futures market might be more in line with fundamental values cf. section 2.3. Contrary, it seems questionable to assume the same about spot rate speculation, which, according to empirical results, is more likely to consist of noise trading.

A general problem with the proposed endogeneity solutions in the literature, as well as with Aliber, Chowdhry & Yan's (2003), is that they assume it to be sufficient to use some form of fundamental volatility as a correction for the problem. We do not see how this solves the problem. Introducing fundamental volatility into the model only seems to split up the effect transaction costs and fundamental volatility has on the excess volatility.

Lanne & Vesala (2006) follow the research method of Aliber, Chowdhry & Yan (2003) although they try to simplify the solution to endogeneity proposed by Aliber, Chowdhry & Yan (2003) by using the money-market headline news as a proxy for fundamental volatility. Even though this method is simpler than the previous one explained it adopts the same problems and does not help solve the problem.

Dealing with endogeneity in the model is best done by the use of instrumental variables. To do so in this case, one would have to find some variable only correlated with transaction costs and not the error term of the model, thereby affecting volatility through transaction costs only.

An instrument must satisfy two assumptions. It has to be relevant, meaning that it has to be correlated with the endogenous variable of the OLS model – in our case cost. Second, it has to be valid, meaning that it has to be exogenous and uncorrelated with the dependent variable of the OLS model (Johnston & DiNardo (1997) and Verbeck (2004)).

It is difficult to test for endogeneity because it is often a theoretically based intuition. In section 5.2 we will test for granger causality between the variables in our models. From the results of our granger tests we hope to be able to conclude at least something about the statistical causality between the variables. In chapter 6 we will run traditional OLS regression models to determine the relationship between cost and volatility, and be able to compare our results with previous research. In section 7.3 we will propose our own solution to the endogeneity problem using an instrumental variable.

4 The data

During chapter 4 we will look at from where we have obtained our data and what it consists of (section 4.1) as well as how it behaves. In section 4.2, 4.3 and 4.4 we will describe the structural breaks during the time period from which we have data, as well as normality and stationarity for cost, volatility and VIX, respectively. In section 4.5 we will briefly explain the attractiveness of volume and transactions as variables, but regret to inform that we have not been able to obtain this data. Section 4.6 gives an overview of the variables' development over time and in relation to each other.

4.1 Collected data

Our data has been collected through Thomson Reuters via Datastream. Thomson Reuters is a collaboration of over 1,100 banks throughout the world and their data consists of interbank transactions. Interbank transactions account for roughly 95% of the world's total currency transactions (Aliber, Chowdhry & Yan (2003)), this is the reason we believe that this gives us the best estimate of the true image.

We have collected data for four of the most traded currencies in the world: GBP, CHF, JPY and EUR where all of them are traded against the USD. The currency pairs we are looking at are therefore: GBP/USD, CHF/USD, JPY/USD and EUR/USD.

The collected data comprises of daily observations for the minimum and maximum traded prices pr. day as well as the daily closing bid and ask prices. This gives us 5,429 observations for each currency from 01-01-1990 to 10-22-2010. Since the Euro was first introduced to the market on the 01-01-1999 we have 3,081 observations for this currency.

Throughout chapter 4 we will focus on the treatment of GBP and only comment on the other currencies if they differ substantially from how GBP acts. Going through all the currencies we found that most of their characteristics were the same, and therefore the treatment of CHF, JPY and EUR will be found in Appendix 1 to Appendix 5.

4.2 Costs

In this section we will quickly recap the measure of cost and discuss the data's statistical properties such as normality and stationarity, which are important assumptions of a multiple variable time-series model. Furthermore, we will determine the structural breaks in cost over time, which we will use to divide our forthcoming model into time periods.

4.2.1 The cost measurement

Based on our review of previous research and because transaction costs are not explicitly quoted in the exchange rate market we have decided to use the method also used by Ronen & Weaver (2001) and Lanne & Vesala (2006) described in section 3.2.1 to calculate our transaction costs:

$$C = \frac{S_a - S_b}{S_b + S_a} \tag{53}$$

Considering our data, we do not have the problem regarding additional transaction costs described by Aliber, Chowdhry & Yan (2003), as it is comprised of wholesale quotes charged between banks and not between banks and non-bank customers (Reuters (2010)). This means that costs implied by our bid-ask spread is strictly transaction costs.

In Figure 6 the exchange rate, the average weekly cost measure and the daily bid-ask spread are displayed. Looking at the spread we find that from 01-03-1990 to around 04-30-2002 it seems that the spread keeps changing mainly between three different bands. From 04-30-2002 to the end of 2010, we see a little more variation but the spread still moves in bands. Because of this small variation, the daily cost measure will behave in approximately the same way, and thus we have decided to take a weekly average of the costs. Furthermore as a moving average works as a simplified Gaussian filter (Smith (1997)) we believe that the weekly average to some degree works in a similar way. The reason why we do not use the moving average is to avoid having information from observations repeated in each data point. The weekly average costs now consist of 1,099 observations for GBP, CHF and JPY and 623 observations for EUR. The weekly average cost for GBP is displayed in the second graph of Figure 6, where we see that it now moves with more variation compared to the spread because of our cost measure and the weekly average. From here on, we will refer to "average weekly costs" as "costs". Looking at the costs of CHF, JPY and EUR in Appendix 1, we only find that EUR behave the same way as GBP. CHF and JPY also lie in bands but they behave more smoothly compared to GBP and EUR.

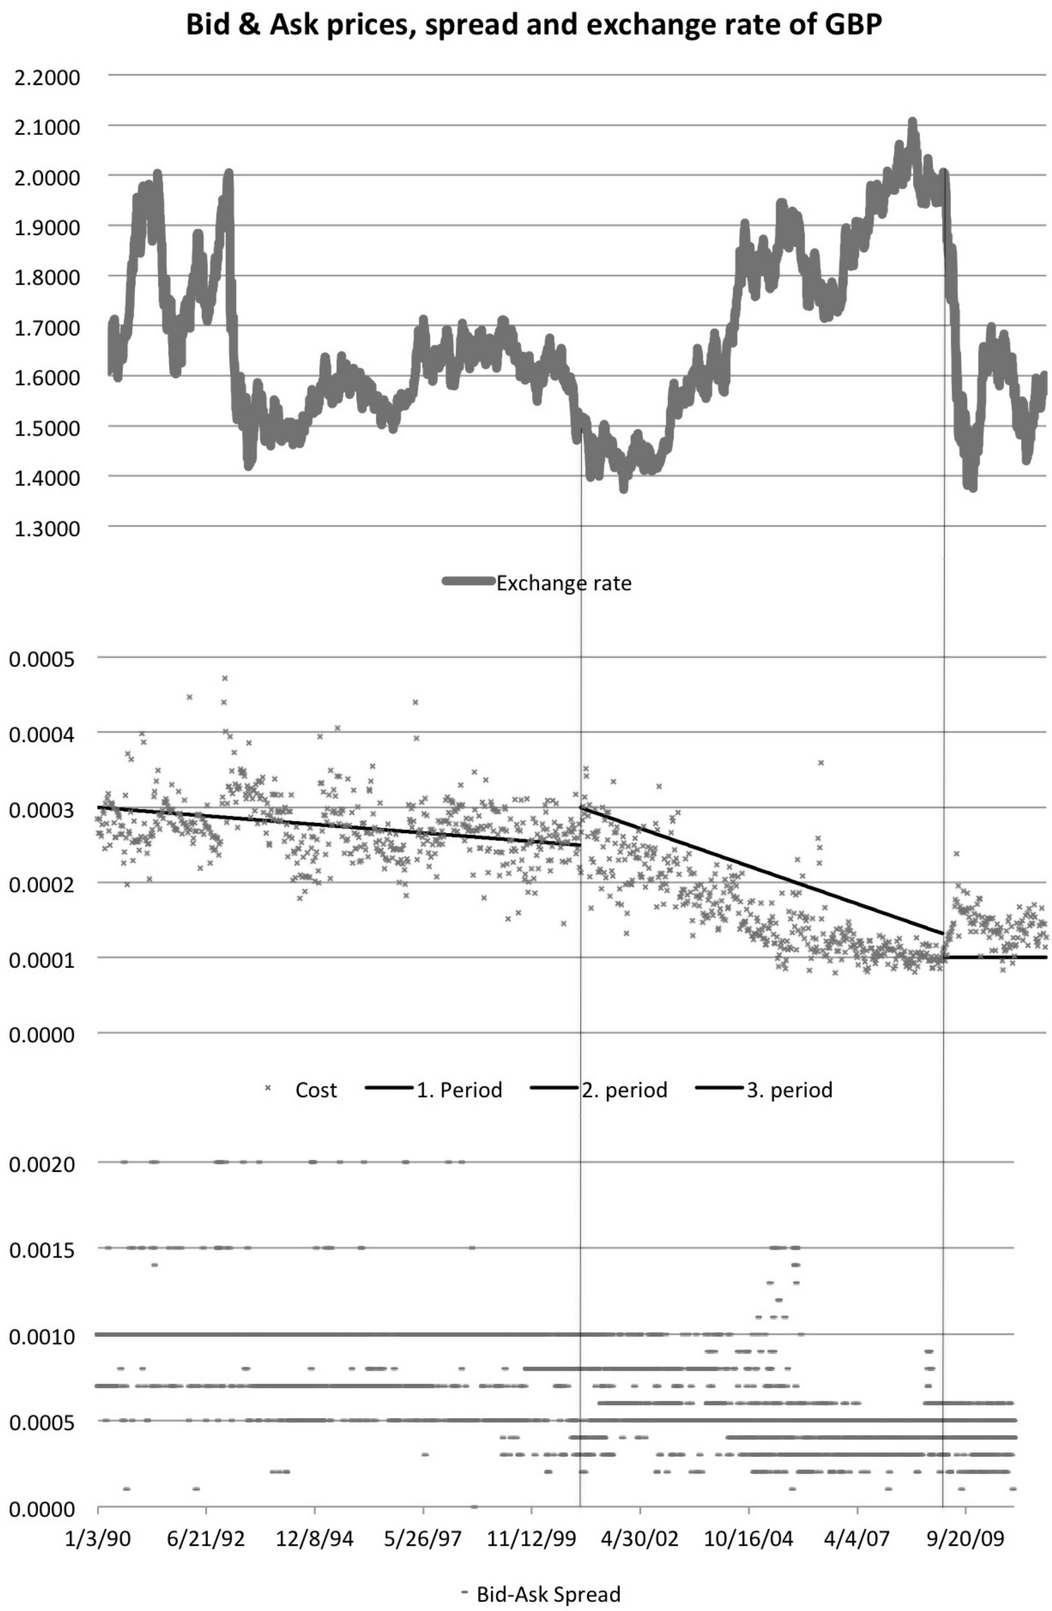

Figure 6 – Exchange rate, average weekly costs and bid-ask spread for GBP

4.2.2 Structural breaks

The period we are looking at stretches over approximately 21 years, except for EUR, which only stretches over approximately 12 years. Over such a long period, it is natural that changes in the structure occur. In Figure 6 we have displayed GBP costs over time as well as trend lines to capture the structural breaks, the exchange rate and the daily spread. It can be seen from the vertical lines how the breaks are constructed. Looking at Figure 6, we find that the break between the first and second period is caused by an appreciation in the exchange rate of GBP. This makes sense as our cost measure is calculated as a percentage of the cost-free price, and when this appreciates without a significant rise in the spread, the spread then becomes a smaller percentage of that price and vice versa. The break between the second and third period is due to the affect of the financial crises and the reason why we see a variation in the structure of the exchange rate and thereby also in the costs. It is obvious that the three periods differ in their structure, and therefore we believe that an overall model would fit worse than if divided into three periods.

As our thesis statement concentrate on how cost affects volatility we consider cost to be our primary explanatory variable, later on the thesis work we will therefore use these structural breaks based on the costs to divide the remaining variables into three different periods.

The structural breaks are not the same for all four currencies. They all move the same way by having a first period that is relatively flat, a steep declining second period and a third period that has a positive slope in the beginning and then later starts to decline again. In Table 1 the sub-periods for the different currencies are displayed.

Period	**GBP**	**CHF**	**JPY**	**EUR**
First period	199000	199000	199000	199901
Second period	200032	199515	200045	200241
Third period	200830	200825	200823	200823
Ending data	201042	201042	201042	201042

Table 1 - Table of sub-periods for the currency

4.2.3 Normality in cost

The time-series we are dealing with are financial and there is substantial evidence that financial time series, and therefore our measure of costs, are not normally distributed but rather leptokurtic i.e. they have high peaks and fat tails (Westerfield (1977), McFarland *et al*. (1982) and Boothe & Glassman (1987)).

As mentioned in section 4.2.1, taking a weekly average of the daily costs should work as a normal distribution filter. To test the impact of this statement we compare the Anderson-Darling values of daily costs and average weekly costs for the full period, and find that it decreases from 81.19 to 14.64. Despite of this filter we still find high peaks and fat tails, which result in a kurtosis value of 6.30 - a value that should equal 0 if kurtosis where not present. Furthermore we find that the full period GBP data is positively skewed with a value of 1.09, which ideally should also equal 0 in a normal distribution.

The first period of GBP's cost provides a better basis for a normal distribution even though the Anderson-Darling value is still too high with a value of 3.20. We still find traces of skewness and kurtosis with values of 0.68 and 2.44 respectively. The second period behaves like the full period with high peaks, positive skewness and an Anderson-Darling test-statistic of 9.44, which exceeds the critical value. The values of kurtosis and skewness are calculated to be -0.21 and 0.75, respectively. The third period is on the other hand normally distributed with an Anderson-Darling value of only 0.29. The distribution analysis for CHF, JPY and EUR can be seen in Appendix 2.

Besides taking the weekly average we have also in very few cases removed a small number of extreme outliers from the data counting for a maximum of 1.8%. This has only been done if the Anderson-Darling values have improved significantly and otherwise we have used the original data. The outliers removed can be seen in Appendix 3.

The reason why we focus on getting the data normally distributed is due to the testing of the forthcoming models. If data is normally distributed the residuals of a linear regression model will be normally distributed, which will make the null hypothesis of the beta values correct and unbiased (Overø & Gabrielsen (2005)).

4.2.4 Stationarity in cost

When making empirical work on time series data, it is assumed that the underlying time series are stationary (Gujarati (2003)). The reason for this is that otherwise the residuals of the model become non-stationary, the OLS estimator will not be normally distributed and as a consequence the standard results for OLS do not hold. Furthermore, regression models for non-stationary variables have the possibility of giving spurious results.

To test if the GBP's full cost period and the sub-periods are stationary, we will use an Autocorrelation Function (ACF) correlogram for graphic interpretation and the Augmented Dickey-Fuller (ADF) test to support the graphic results.

In the full period of GBP, the graphic inspection of the ACF-correlogram indicates that the period is non-stationary i.e. it has a unit root, since the ACF decays very slowly. In this particular case we find that the ADF-test has contradictory results compared to the ACF-correlogram. The ADF-test finds that for the first 5 lags of difference terms in a random walk with drift the null hypothesis of non-stationarity is rejected. If we include more than 5 lags of difference terms we find that the error term becomes more and more correlated, which is why we believe that the first difference, denoted by I(1), will make the full period stationary. After taking the first difference of the full period we find that the results strongly indicate stationarity.

Looking at Figure 6's cost, it seems that the first period does not have a unit root, which we from here on will denote I(0), as the mean and variance seem constant over time. Testing this, we find that the results support this belief. In the correlogram we find that the ACF decrease rapidly and even if 20 lags of difference terms in the ADF-test are included the error terms are still uncorrelated and the null hypothesis is rejected. For the second and third period of GBP costs, the ACF and ADF test results show that they are stationary after taking the first difference. For simplicity, Table 5 on page 73 displays which series do not have a unit root, denoted by I(0) and which ones that do. For these we have taken the first difference, denoted by I(1).

4.3 Volatility

In this section, we will quickly recap the volatility measure and discuss the data's statistical properties such as normality and stationarity for the same reasons described in the introduction to cost in section 4.2.

4.3.1 The volatility measurement

For measuring volatility we have, based on previous research, been inspired by Hau's (2006) range measure as described in section 3.2.2:

$$RANGE_t = \frac{max_{(s \in I_t)} S - min_{(s \in I_t)} S}{max_{(s \in I_t)} S + min_{(s \in I_t)} S} \quad (54)$$

We focus on this measure and not others because the range is more robust, it lets us use all observation within the interval, and the method is the same as used in our cost measure described in section 4.2.1.

Furthermore, it is important for the forthcoming model that linearity is present in the parameters, meaning that there should be a linear relationship between the dependent and independent variables. We find this to be the case and the graphical test for all currencies can be found in Appendix 4.

4.3.2 Normality in volatility

As the volatility measure is based on exchange rate data i.e. financial data, we generally find traces of high peaks and fat tails through the distribution analysis cf. section 4.2.3. Further applying a simple normality distribution filter by taking the weekly average of the daily volatility measures improves the normality of the volatility as we saw with costs as well. This can be seen by comparing the full periods' Anderson-Darling values from before and after the weekly average, which are 195.36 and 43.98, respectively. The decrease in the Anderson-Darling values also applies when looking at the sub-periods. Moreover, we find that the full period have a positive skewness of 2.47, and a kurtosis of 9.96.

For the first, second and third sub-period we find Anderson-Darling values of 14.61, 2.14 and 5.58, respectively. Meaning none of the sub-periods are normally distributed but still better than the full period. The tendency of positive skewness is visible through all the sub-periods. In the first period the value of skewness is 1.97 and in the second and third the values are 0.59 and 1.48, respectively. As pointed out in the beginning of this section, all periods suffer from high peaks and fat tails i.e. kurtosis. The figures for the sub-periods for this measure are: 7.42, 0.58 and 2.41, respectively.

4.3.3 Stationarity in volatility

We use the same procedure to test for stationarity in volatility as in section 4.2.4. For the full, first and second period we find it hard to determine whether the series are stationary or have a unit root. E.g. looking at the full period we find that the ADF-test in all cases (up to 20 lags) rejects the null hypothesis, but the ACF decreases very slowly and reaches 0 in around 100 lags i.e. 9% of the observations. Because of our doubt we will continue this thesis work by looking at both I(0) and I(1) for these particular periods. In the third period we find that the time series of volatility have a unit root and the period will thereby be denoted by I(1), cf. Table 5.

4.4 VIX

In this section we will describe the VIX measure, which we will use as a benchmark for the currency market volatility. Furthermore, we will discuss its statistical properties such as normality and stationarity for the same reasons described in the introduction to cost in section 4.2.

4.4.1 The VIX Measurement

VIX is the ticker symbol for the Chicago Board Options Exchange Market Volatility Index and was introduced in 1993 by Professor Robert Whaley. It is a measure of the implied volatility of S&P 500 index options and measures the market's expectation of stock market volatility over the next 30-day period. The VIX is calculated and disseminated in real-time by the Chicago Board Options Exchange and it is a weighted combination of prices for a range of options on the S&P 500 index. The VIX is quoted in percentage points and translates, roughly, to the expected movement in the S&P 500 index over the next 30-day period, which is then annualized. For example, if the VIX is 17 this represents an expected annualized change of 17% over the next 30 days (CBOE (2011)).

We have decided to include this measure in the forthcoming model and use it as a benchmark to our calculated volatility. We use VIX as an indication of how the market moves and by including this we find if our volatility measure is only affected by costs or also by the market volatility in general, even though VIX resembles the stock market volatility.

4.4.2 Normality in VIX

Looking at the distribution analysis for the full period and the three sub-periods of VIX we find Anderson-Darling values of 27.09, 7.81, 6.34 and 6.31, respectively. This means that none of the periods are normally distributed, even though we find much better results after applying the normal distribution filter. Furthermore, we can report positive skewness and kurtosis for all periods. For the full period we find the positive skewness to be 1.94 and a kurtosis value of 6.55. For the three sub-periods we find positive skewness values of 0.96, 0.73 and 1.32 and kurtosis values of 1.05, -0.06 and 1.08, respectively.

4.4.3 Stationarity in VIX

We used the same procedure to test for stationarity in VIX as seen in section 4.2.4. Looking at the ACF correlograms and the ADF tests for VIX we find a unit root in all periods. This should also be expected when looking at Figure 7, where it seems clear that VIX does not move around a constant mean and does not have a constant variance. The results of stationarity can be seen for all currencies in Table 5.

4.5 Volume and transactions

As the size and the number of trades affect the bid and ask prices and thereby our cost measure, cf. section 3.3, we believe that these two variables could be significant in our forthcoming model. This data is unfortunately very difficult to obtain, as it is not publicly available. To get an estimate of these two variables we have contacted some of Denmark's largest financial institutions i.e. Danske Bank, Nordea, JyskeBank, Sydbank, Nykredit, Alm. Brand Bank and National Banken.

Our idea was to sum the size of the trades per day divided by the number of trades for each currency cross. This would give us an estimate of the average size of a transaction on a given day. In addition to this, we would try to make an estimate of how the amount of trades moves over the chosen period. This can be done by summarizing the trades per week across the financial institutions for each currency and then standardize these as an index. By doing this, we would get an estimate of how the amount of trades move from week to week and then use this as an indication of how the total market moves.

Other authors such as Ronan & Weaver (2001) and Aliber, Chowdhry & Yan (2003) have also incorporated trades in their models and they find that the estimates of the beta values are significant. It is noticeable though that they find the beta value to be positive, which is the opposite of what the theory described in section 3.3 would suggest. Further, they do not mention anything about the correlation between their costs measure and the trade variables, which could be an issue.

Unfortunately our request has been rejected by all participants except for Alm. Brand Bank who has been so kind as to provide us with the necessary data. However, we do not believe that data from only one participant is sufficient to form a basis for qualified estimates. This is due to the fact that Alm. Brand Bank is a small actor on the currency exchange market and therefore we believe that their data, to a higher extent is determined by a few customers trading than the market in general. Based on this, we have decided that the data cannot be used as a realistic proxy for how the global market moves.

4.6 Overview of the movement in the variables

In Figure 7 the standardized time series of volatility, cost and VIX for GBP are displayed. These will be the variables used in the forthcoming GBP models. From the figure we can see that in the first period, which ends on 08-11-2000, the cost and volatility move alike with cost above volatility, compared to the second period, where cost and volatility crosses. In the third period, which starts on 08-01-2008, volatility increases and now lies above cost. Looking at VIX it seems that it only follows volatility in the beginning and end of the full period. Based on these findings we believe that the explanatory power of our forthcoming models will vary from period to period. CHF, JPY and EUR are displayed in Appendix 5.

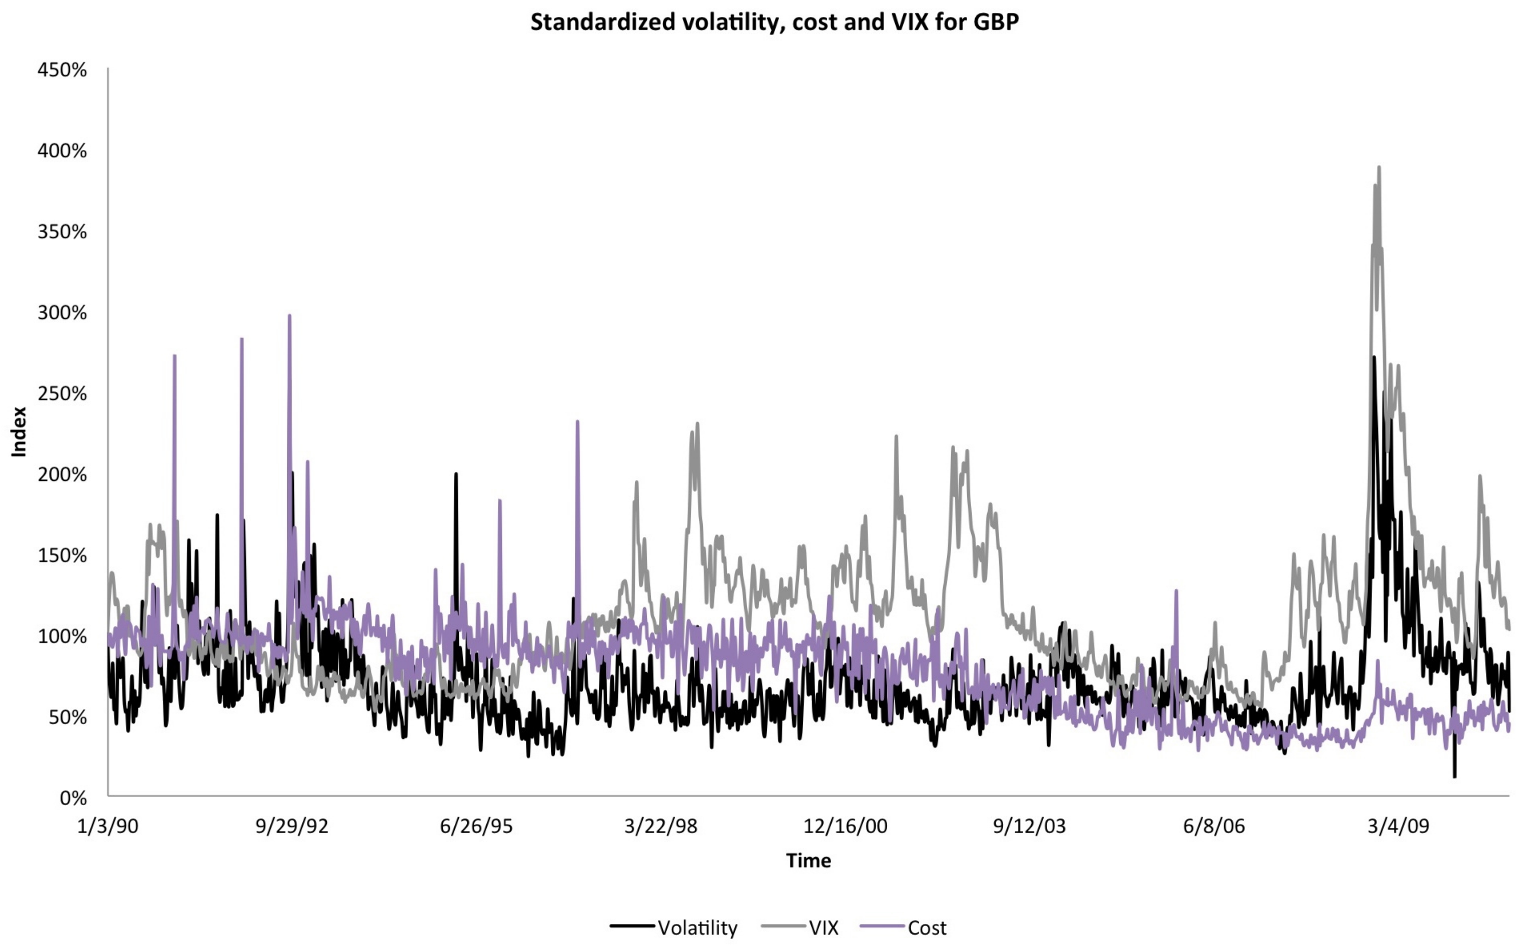

Figure 7 - Standardized volatility, cost and VIX measures for the full period of GBP.

5 Diagnostics

Before testing any statistical relationships between our data we give an overview of the main OLS assumptions behind the statistical tests used and how it affects our results if the assumptions are not met. We will do this in section 5.1. In section 5.2 we will investigate the granger causality between cost and volatility in order to determine a possible direction of causality before setting up our models in chapter 6. We do this to examine if there is any pattern in the statistical causality between cost and volatility and if we can thereby say anything about the endogeneity discussed in section 3.4.

5.1 Data assumptions

In section 5.1 we will describe the consequences of violating four of the main OLS assumptions, which we will refer to as the test statistics, behind the regressions that we will be running. A violation can either affect the distribution of our data or the variance in the estimated coefficients, which can lead to wrongly interpreted results from t-tests or F-tests. We do not believe the OLS assumption concerning exogeneity to be met in our models, cf. section 3.4. We will ignore this for now and return to it in section 7.3. The remaining OLS assumptions we do not describe in this section, we assume to be met.

5.1.1 Normality

If the assumption of normality in the error terms is violated the t-test or F-test will be less powerful (Gujarati (2003)). Either skewness or high or low kurtosis will lead to wrong test-values, or rather the data will not fit the distribution associated with the test and will therefore not correspond to the critical values proposed by said test. If the distribution has high kurtosis a higher percentage of the observations will lie outside the critical area for accepting the null hypothesis and the test will more often reject the hypothesis and vice versa with low kurtosis. The presence of positive skewness will lead to more frequently accepted null hypotheses whereas with negative skewness will lead to rejected null hypotheses more frequently. The reason for this stems from the high concentration of observation within the tests' distribution in the former case, and the majority of observations lying outside the distribution in the latter case. The effect of violating one of the reasons for non-normality, of course, depends on the size of the violation. In sections 5.2.3 and 6.3.1 we will use the Anderson-Darling test for normality to test if our data is normally distributed. The Anderson-Darling test tests the null hypothesis that data is normally distributed so by rejecting the null hypothesis we are accepting that data is not normally distributed.

5.1.2 Autocorrelation

Autocorrelation in the residuals is usually due to two errors. If a variable has been omitted from the model the residuals will include the missing variable and possibly be autocorrelated. The other error could be a misspecification of the functional form of the model (Gujarati (2003)). As stated in chapter 4, we have assumed that the relationship between the variables is linear. Another case of misspecification could be taking a moving average to reduce seasonal effects where there were none, which might be the case with our weekly averages of the variables.

Autocorrelation in the residuals does not affect the estimated beta values, although it does affect the variance of the estimates and therefore the F and t-value. In case of positive autocorrelation the standard error of the estimates will be underestimated and we will reject the null hypothesis more often due to the overestimated F and t-values. Another error originating from positive autocorrelation is an increased R^2 value for the model. In case of negative autocorrelation in the error terms, the bias in the model will point in the opposite direction as it did for positive autocorrelation (Gujarati (2003)).

5.1.3 Multicollinearity

Multicollinearity is the case of correlated variables within the model. The presence of multicollinearity does not affect the models' predictive power as a whole, but will make the correlated variables' parameters untrustworthy. If two variables are correlated they will, for all practicalities, tell us the same thing about the dependent variable. This results in the explanatory power of one of the variables being divided amongst them (Gujarati (2003)).

When variables are correlated the standard errors of their estimated coefficients tend to be larger than they would have been with no correlation. This results in lower F and t-values, which leads us to accept the null hypothesis more often than usual (Gujarati (2003)).

To determine if multicollinearity is a problem for our models we will use the value of the condition index (CI). A rule-of-thumb is that when CI is less than 10 multicollinearity is not present. When it is between 10 and 30, multicollinearity is moderate, and a value above 30 indicates severe multicollinearity (Gujarati (2003)).

5.1.4 Heteroscedasticity

The fourth assumption involves the variance in the residuals, which by the assumption of OLS has to be homoscedastic (Gujarati (2003)). One simple way to test for homoscedasticity is to look for any system between the residuals and both the dependent and the explanatory variables. A more formal test is the Breusch-Pagan test, which tests for homoscedasticity in the model. In the case where the test finds that the independent variables are jointly significant in explaining the estimated residuals of the model, it rejects the null hypothesis and the residuals are heteroscedastic.

The presence of unequal variance in the variables does not affect the estimated beta values of the model. It does, however, affect the variance in the parameters and gives us a wrong estimate of the population variance or true variance. In this case our F and t-values will be biased due to a wrongly estimated variance and we cannot for sure say whether our hypothesis is rejected or accepted (Gujarati (2003)).

The effect heteroscedasticity has on the t and F-values depend on the size of the violation. Although it is not possible for us to comment on the direction of the bias – we can only evaluate the size of the Breusch-Pagan value to consider if the t-test or F-test is still usable.

5.2 Granger causality

In section 5.2.1 we will evaluate the goodness-of-fit of our granger models with different lags to determine the appropriate number of lags to include in our granger causality models. Through section 5.2.3 to 5.2.6 we will describe if the results of our granger model are biased from violations of the four main assumptions underlying the regression analysis and our interpretation of the granger models' results. In section 5.2.7 we discuss the results from our granger models. We will use the results of our granger models to give us an indication of the presence of the endogeneity problem.

Before setting up our model we want to see if we can say anything about the direction of the causality between cost and volatility over time. To do this, we examine the granger causality between the lagged and non-lagged variables. We test both whether cost granger causes volatility and vice versa in order to obtain a preliminary idea of what to expect from the results of our model. It is important to note that granger causality does not imply true causality but merely tells us if one variable helps predict future values of the dependent variable, even when lagged values of the dependent variable are included in the model (Gujarati (2003)). We hope to be able to say something about the presence or size of the endogeneity problem as explained in section 3.4. The rejection of the hypothesis that one variable does not granger cause another only implies a statistical causality. If for example, the dependent variable and the independent variables in the model are in some way explained by an omitted variable, granger causality tests could still show a causal relationship where there in reality is none.

To test for granger causality we set up four regression models for each stationary scenario as seen in Table 5. Two where we regress volatility at time t on lagged values of cost and volatility, and two models where we regress cost at time t on lagged values of cost and volatility:

$$V_t = \beta_0 + \sum_{i=1}^{m} \beta_i V_{t-i} + \sum_{i=1}^{m} \alpha_i C_{t-i} + \varepsilon_t \tag{55}$$

$$V_t = \beta_0 + \sum_{i=1}^{m} \alpha_i C_{t-i} + \varepsilon_t \tag{56}$$

And

$$C_t = \beta_0 + \sum_{i=1}^{m} \beta_i V_{t-i} + \sum_{i=1}^{m} \alpha_i C_{t-i} + \varepsilon_t \tag{57}$$

$$C_t = \beta_0 + \sum_{i=1}^{m} \beta_i V_{t-i} + \varepsilon_t \tag{58}$$

where volatility and cost are denoted by V and C, respectively. The granger test uses an F-test to determine if one model is better than the other. In our case, we want to compare equation (55) with equation (56), and equation (57) with equation (58). Before running the regressions we have to decide the number of lags for each variable to include in the models. We base this decision on the Akaike Information Criterion as explained in section 5.2.1.

5.2.1 AIC

The Akaike Information Criterion is used to determine the trade-off between a model's simplicity and goodness-of-fit. Adding more variables to a model will always increase its R^2, but the model itself will become more complex. To measure the AIC of a model we use the following equation (Gujarati (2003)):

$$ln(AIC) = \left(\frac{2k}{n}\right) + ln\left(\frac{RSS}{n}\right) \tag{59}$$

where k is the number of regressors including the intercept, n is the number of observations and RSS is the residual sum of squares of the model. The AIC statistic has no meaning in itself; only when it is compared to other more or less complex models with the same dependent variables can we say if it is relatively better or worse. We therefore found the percentage difference between the models AIC as a way to determine the appropriate number of lags to include.

Period		Scenario	Lag	$\%\Delta_{1\text{-}2}$	$\%\Delta_{2\text{-}3}$	$\%\Delta_{3\text{-}4}$	$\%\Delta_{4\text{-}5}$	$\%\Delta_{5\text{-}6}$
Full	Volatility	I(1), I(0)	3	15.56%	1.34%	0.33%	0.96%	0.19%
	Cost	I(1), I(0)	3	7.70%	4.94%	1.42%	1.20%	2.00%
	Volatility	I(1), I(1)	3	2.33%	1.87%	1.76%	0.82%	0.11%
	Cost	I(1), I(1)	3	7.73%	5.12%	1.22%	1.54%	1.74%
First	Volatility	I(0) , I(0)	2	5.10%	0.00%	-0.35%	-0.17%	-0.60%
	Cost	I(0) , I(0)	2	1.91%	0.70%	-0.12%	-0.64%	-0.25%
	Volatility	I(0) , I(1)	2	3.99%	1.91%	1.24%	0.05%	0.74%
	Cost	I(0) , I(1)	2	2.74%	1.12%	0.80%	-0.27%	0.01%
Second	Volatility	I(1), I(0)	3	9.43%	3.77%	0.35%	0.81%	2.72%
	Cost	I(1), I(0)	2	5.01%	0.07%	1.10%	2.85%	1.17%
	Volatility	I(1), I(1)	3	8.18%	2.09%	0.43%	1.78%	2.64%
	Cost	I(1), I(1)	2	4.97%	0.07%	1.81%	2.15%	1.24%
Third	Volatility	I(1), I(1)	2	2.31%	-3.41%	-0.91%	-1.89%	-1.93%
	Cost	I(1), I(1)	3	1.60%	7.20%	-3.08%	-1.13%	-3.05%

Table 2 - The percentage difference in AIC-value for all scenarios and periods for GBP. The $\%\Delta_{1\text{-}2}$ for example means the percentage change in AIC-value from a model with 1 lag of cost and volatility to a model containing 2 lags of both cost and volatility. The value in column 4 is the number of chosen lags of both variables in the granger model. Column 2 tells us the dependent variable in the regression.

In Table 2 we see the different scenarios for both cost and volatility as the dependent variable for both the full and all the sub periods of GBP. The scenario in the third column explains if first cost has a unit root or not and next if volatility has. For example for the full period's first scenario we see that cost is differenced and volatility is not. The AIC results for CHF, JPY and EUR can be found in Appendix 6. Column four tells us the number of lags of both cost and volatility that we have chosen in our granger model, based on the percentage change in AIC value from models with increasing lags. If we look at our model with volatility as the dependent variable for the second period when cost is differenced and volatility is not, we see that we have chosen to include three lags of both variables since the AIC-value improves by 9.43% when increasing the number of lags from 1 to 2 and again increases by 3.77% when the number of lags increases from 2 to 3. From there on, adding additional number of lags does not improve the model noticeably.

5.2.2 Granger hypothesis testing

Having decided on the number of lags of each variable to include in our models, we can go on to test for granger causality between cost and volatility. More formally, we want to test if the following null hypothesis is accepted for equation (55) and (57):

$$H_0: \alpha_1 = \alpha_2 = \cdots = \alpha_m = 0$$
$$H_1: H_0^C \tag{60}$$

The granger causality test requires the time-series to be both linear and stationary. Based on our ARIMA models examined, we assume stationarity in the scenarios mentioned in section 4.2.4, 4.3.3 and 4.4.3, as well as linearity, cf. Appendix 4. These assumptions are required to hold so as to be able to perform the F-test, which is the basis for the granger causality test (Gujarati (2003)). Furthermore, the F-test requires the model to uphold the normal OLS assumptions of normality, no autocorrelation, no multicollinearity and homoscedasticity (Gujarati (2003)). In Table 3, we have reported the results of the test statistics for the four assumptions. The results for CHF, JPY and EUR can be seen in Appendix 7.

Period		Scenario	Lag	Skewness	Kurtosis	Anderson-Darling (t)	Auto-correlation	Multi-collinearity	Breusch-Pagan (p)
Full	Volatility	I(1), I(0)	3	2.0993	12.2112	22.8484	(0.0098)	9.7140	<.0001
	Cost	I(1), I(0)	3	4.3747	49.3124	49.8483	(0.0277)	9.7140	0.3934
	Volatility	I(1), I(1)	3	1.5541	10.4130	17.6082	(0.0192)	2.1834	0.0008
	Cost	I(1), I(1)	3	4.3405	49.1585	49.9708	(0.0294)	2.1834	0.2828
First	Volatility	I(0) , I(0)	2	1.6924	6.4595	9.8503	(0.0216)	22.0715	<.0001
	Cost	I(0) , I(0)	2	0.1087	3.3991	3.4783	(0.0200)	22.0715	0.4798
	Volatility	I(0) , I(1)	2	0.9840	5.9263	6.7520	(0.0351)	16.6250	<.0001
	Cost	I(0) , I(1)	2	0.6390	3.2781	4.1674	(0.0246)	16.6250	0.4524
Second	Volatility	I(1), I(0)	3	0.7790	0.9955	3.1595	(0.0252)	12.0924	0.5957
	Cost	I(1), I(0)	2	0.1087	3.3991	5.1230	(0.0257)	10.4172	0.3681
	Volatility	I(1), I(1)	3	0.5910	1.1337	2.2547	(0.0223)	2.3747	0.8045
	Cost	I(1), I(1)	2	0.1078	3.4029	5.1381	(0.0258)	1.7415	0.3095
Third	Volatility	I(1), I(1)	2	0.7435	2.8663	0.8768	(0.0162)	1.9943	0.1849
	Cost	I(1), I(1)	3	(0.5206)	1.1596	0.6077	(0.0143)	2.3747	0.3587

Table 3 – Test results of all granger models for all scenarios in all periods for GBP. Column 5-7 reports the skewness, kurtosis and the Anderson-Darling t-statistic for the granger models residuals. Column 8 reports the autocorrelation between the granger models residuals. Column 9 states the degree of multicollinearity between the explanatory variables represented as the condition index and column 10 reports the p-value from the Breusch-Pagan test for homoscedasticity.

5.2.3 Normality in Granger

In Table 3 we report the kurtosis, skewness and Anderson-Darling test-statistic to determine if the residuals of the models are normally distributed, and if not in what way the bias is directed. None of the periods, except the third, are normally distributed according to our Anderson-Darling test. The period in which the data least fits a normal distribution is the full period with an Anderson-Darling statistic of about 50 and 20 when cost or volatility, respectively, acts as the dependent variable. For the first and second period, the Anderson-Darling values range from approximately 2.25 to 9.85. It is the case for all the non-normally distributed periods that they have a relatively small skewness close to 0 and a positive kurtosis becoming leptokurtic as much financial data is, cf. section 4.2.3. The two sorts of bias' work against each other but since the Anderson-Darling test does not deem them normally distributed we must assume that the positive kurtosis leads to a bigger bias than the skewness. This means that when evaluating the results from the granger causality test, we have to be cautious about rejecting our null hypothesis too quickly.

5.2.4 Autocorrelation in Granger

In Table 3 we see the autocorrelation for all the cases of GBP. The autocorrelation for all the models is very close to 0 and of no considerable magnitude when we evaluate the granger causality models.

5.2.5 Multicollinearity in Granger

Since multicollinearity only affects the individual beta values but not the overall model as described in section 5.1.3, this will not be a matter of concern when we evaluate the results of our F-test for granger causality (Gujarati (2003)).

5.2.6 Heteroscedasticity in Granger

From Table 3 it shows that the models for equation (55) are highly heteroscedastic during the full and the first period. This is the case for all scenarios, both when cost or volatility is either differenced or non-differenced. If we look at the results for equation (57) during the same periods we see that these exhibit much less heteroscedasticity, which leads us to believe that the cause lies with the relationship between the volatility and its lagged observations. The second and third period of GBP, has a lower degree of heteroscedasticity compared to the full and first period. It is difficult to say why the results here, in the case of the results for equation (55), are so different from the first two periods. In the cases of heteroscedasticity we have to be wary about trusting our test results from the granger test. In a few cases it might be worth looking at the level of significance where the null hypothesis is either accepted or rejected, to determine if it might be too biased to conclude anything from the results.

5.2.7 Test results

The F-test we performed to test whether equation (56) is insignificantly different from equation (55) in the same manner as equation (58) are from (57), and hence if there are any granger causality between cost and volatility is stated as:

$$F = \frac{\frac{RSS_r - RSS_{ur}}{k-q}}{\frac{RSS_{ur}}{n-k}} \sim F(k-q, n-k) \tag{61}$$

where the residual sum of squares for the restricted and unrestricted models are denoted by RSS_r and RSS_{ur}, respectively. The number of regressors including the intercept for the restricted model is defined as q and for the unrestricted model k. The number of observations is n.

The results of the test are shown in Table 4. The column marked F-value tells us if the null hypothesis is either accepted or rejected, where the number of stars indicates at what level the hypothesis is rejected. One star means a 10% significance level, two stars are equal to 5% and finally, three stars mean 1%.

Looking at the full period in Table 4 we see that for all of the currencies, cost granger causes volatility while volatility does not granger cause cost. This is the case when cost is differenced and volatility is not, and only for GBP and JPY is the result the same when volatility is also differenced. In all the cases where the hypothesis is rejected for the full period it happens at a 5 or 1 per cent significance level. Considering this while keeping in mind the results of the test for the assumptions in Table 3 we should only have to worry about the presence of heteroscedasticity in the cases where volatility acts as the dependent variable and normality in all the cases.

The kurtosis of the residuals for the full period's scenarios is of considerable size compared to the individual sub-periods, although much larger when testing if volatility granger causes cost than vice versa. This should lead us to wrongly reject a null hypothesis.

What we can conclude about the granger causality of the full period is that the hypothesis that cost granger causes volatility is probably more likely than volatility granger causing cost. Due to the low F-values of the models with cost as the dependent variable and the high values of kurtosis we can almost certainly say that volatility does not granger cause cost, although the reverse relationship is harder to dismiss even with high kurtosis.

In Figure 7 we saw that our first period looked relatively calm, and cost and volatility seemed to follow each other more closely as opposed to the second and third period. The results of our granger test in Table 4 confirm the belief that we would expect to find some granger causality here. For all currencies, except EUR, we reject the hypothesis for one or more scenarios, i.e. we find some granger causality. In scenario I(0)I(0) for CHF and GBP we find evidence of a granger causality between cost and volatility that goes both ways, where for JPY only cost granger causes volatility and for EUR there is no granger causality. Looking at how violated the assumptions are in Table 3 we see that the tests give us a much better picture compared to the full period. For almost all the models we see that the residuals to a high degree are both homoscedastic and normally distributed, at least when cost works as the dependent variable and less when volatility does. Contrary to the full period, the first period shows larger signs of multicollinearity, although, since multicollinearity does not affect our F-test we find the causality relationships for the I(0)I(0) scenarios reliable.

In the first period of CHF's I(1)I(0) scenario we only find volatility to granger cause cost and for JPY's I(0)I(1) scenarios we find the opposite relationship. For EUR and GBP we find no granger causality.

The results for the second period is that the hypothesis is accepted for all scenarios, except I(1)I(0) for CHF where we find weak evidence of cost granger causing volatility and for JPY where we find evidence of the opposite relationship. The F-values are relatively far from their closest critical value and looking at Table 3 the only source of any bias would be the kurtosis. This amplifies our acceptance of our hypothesis and indicates that the F-value rejecting the hypothesis from CHF might be smaller than it is. From Figure 7 this would also seem to be the result that one would expect.

It is difficult to say anything general about the third period. For CHF we only find that volatility granger causes cost, for GBP we find that the granger causality goes both ways, with JPY we have that cost granger causes volatility at least when both variables are I(1), and last but not least we find no evidence of any granger causality for EUR. For this sub-period all assumptions seem to hold, thus the above stated results can be considered conclusive.

Period	LHS	Cost, Vol	Lags	N	K	Q	RSS_unres	RSS_res	F-Value	DF1	DF2	LHS	Cost, Vol	Lags	N	K	Q	RSS_unres	RSS_res	F-Value	DF1	DF2
							CHF											**JPY**				
Full	Volatility	I(1), I(0)	2	1096	5	3	1.9767E-03	1.9892E-03	3.46007**	2	1091	Volatility	I(1), I(0)	3	1095	7	4	2.9392E-03	2.9619E-03	2.80178**	3	1088
	Cost	I(1), I(0)	5	1093	11	6	2.5695E-06	2.5798E-06	0.87281	5	1082	Cost	I(1), I(0)	3	1095	7	4	5.8603E-01	5.8716E-01	0.70227	3	1088
	Volatility	I(1), I(1)	3	1095	7	4	2.0134E-03	2.0249E-03	2.06281	3	1088	Volatility	I(1), I(1)	3	1095	7	4	3.0686E-03	3.0949E-03	3.10739**	3	1088
	Cost	I(1), I(1)	5	1093	11	6	2.5701E-06	2.5798E-06	0.81515	5	1082	Cost	I(1), I(1)	4	1094	9	5	5.7096E-01	5.7276E-01	0.85616	4	1085
1	Volatility	I(0), I(0)	2	271	5	3	6.2105E-04	6.3504E-04	2.99764*	2	266	Volatility	I(0), I(0)	4	564	9	5	1.8766E-03	1.9046E-03	2.06516*	4	555
	Cost	I(0), I(0)	4	268	9	5	2.7366E-07	2.9045E-07	3.97284***	4	259	Cost	I(0), I(0)	2	566	5	3	1.7165E-06	1.7189E-06	0.3881	2	561
	Volatility	I(1), I(0)	2	270	5	3	6.2784E-04	6.3504E-04	1.52073	2	265	Volatility	I(0) , I(1)	4	563	9	5	1.9161E-03	1.9592E-03	3.1122**	4	554
	Cost	I(1), I(0)	5	267	11	6	2.7152E-07	2.9532E-07	4.48772***	5	256	Cost	I(0) , I(1)	2	565	5	3	1.7088E-06	1.7189E-06	1.65271	2	560
2	Volatility	I(1), I(0)	2	586	5	3	7.6806E-04	7.7452E-04	2.44117*	2	581	Volatility	I(1), I(0)	2	325	5	3	2.9979E-04	3.0157E-04	0.95027	2	320
	Cost	I(1), I(0)	2	586	5	3	1.1155E-06	1.1189E-06	0.89242	2	581	Cost	I(1), I(0)	3	324	7	4	4.3694E-03	4.5185E-02	3.60607**	3	317
3	Volatility	I(0), I(1)	3	119	7	4	0.00037092	0.00038647	1.56436	3	112	Volatility	I(0), I(1)	2	123	5	3	5.0958E-04	5.1605E-04	0.74969	2	118
	Cost	I(0), I(1)	2	120	5	3	5.79859E-07	6.03856E-07	2.37958*	2	115	Cost	I(0), I(1)	1	124	3	2	1.9303E-07	1.9402E-07	0.62031	1	121
	Volatility	I(1), I(1)	3	227	7	4	4.4513E-04	4.5226E-04	1.17466	3	220	Volatility	I(1), I(1)	5	192	11	6	5.8654E-04	6.4710E-04	3.73758***	5	181
	Cost	I(1), I(1)	2	228	5	3	1.0073E-06	1.0327E-06	2.81664*	2	223	Cost	I(1), I(1)	4	193	9	5	2.2985E-02	2.3527E-02	1.08467	4	184
							GBP											**EUR**				
Full	Volatility	I(1), I(0)	3	1095	7	4	1.6262E-03	1.6482E-03	4.91678***	3	1088	Volatility	I(1), I(0)	2	620	5	3	9.2742E-04	9.4293E-04	5.14277***	2	615
	Cost	I(1), I(0)	3	1095	7	4	2.4755E-01	2.4788E-01	0.47687	3	1088	Cost	I(1), I(0)	4	618	9	5	1.0569E-06	1.0628E-06	0.84518	4	609
	Volatility	I(1), I(1)	3	1095	7	4	1.6814E-03	1.7052E-03	5.12858***	3	1088	Volatility	I(1), I(1)	3	619	7	4	9.0030E-04	9.0039E-04	0.01931	3	612
	Cost	I(1), I(1)	3	1095	7	4	2.4707E-01	2.4788E-01	1.19142	3	1088	Cost	I(1), I(1)	4	618	9	5	1.0538E-06	1.0628E-06	1.29125	4	609
1	Volatility	I(0) , I(0)	2	551	5	3	8.4470E-04	8.6654E-04	7.05922***	2	546	Volatility	I(0), I(0)	4	194	9	5	2.5419E-04	2.5562E-04	0.25947	4	185
	Cost	I(0) , I(0)	2	551	5	3	6.6234E-07	7.1238E-07	20.62594***	2	546	Cost	I(0), I(0)	1	197	3	2	7.0712E-07	7.1416E-07	1.9319	1	194
	Volatility	I(0) , I(1)	2	550	5	3	9.5152E-04	9.5244E-04	0.26481	2	545	Volatility	I(0) , I(1)	3	194	7	4	2.7432E-04	2.7512E-04	0.18275	3	187
	Cost	I(0) , I(1)	2	550	5	3	7.1029E-07	7.1238E-07	0.80504	2	545	Cost	I(0) , I(1)	3	194	7	4	6.6736E-07	6.8869E-07	1.99283	3	187
2	Volatility	I(1), I(0)	3	416	7	4	2.4752E-04	2.4986E-04	1.28916	3	409	Volatility	I(1), I(1)	3	294	7	4	1.9134E-04	1.9221E-04	0.43655	3	287
	Cost	I(1), I(0)	2	417	5	3	4.2477E-02	4.2526E-02	0.23906	2	412	Cost	I(1), I(1)	4	293	9	5	2.3150E-07	2.3185E-07	0.10457	4	284
	Volatility	I(1), I(1)	3	416	7	4	2.6029E-04	2.6072E-04	0.22292	3	409											
	Cost	I(1), I(1)	2	417	5	3	4.2478E-02	4.2526E-02	0.23406	2	412											
3	Volatility	I(1), I(1)	2	116	5	3	3.8230E-04	4.1002E-04	4.02455**	2	111	Volatility	I(1), I(1)	1	124	3	2	4.1579E-04	4.1681E-04	0.2983	1	121
	Cost	I(1), I(1)	3	115	7	4	3.9259E-03	4.5045E-03	5.30557***	3	108	Cost	I(1), I(1)	2	123	5	3	9.8290E-08	9.9562E-08	0.76387	2	118

Table 1 – Granger results for all scenarios of all periods within all of the four currencies. The F-values from the granger test is marked with *'s to show how rejected the null hypothesis is. F-values without any stars mean an accepted null hypothesis of no granger causality between then variables. One star means that the null hypothesis is rejected at the 10%-level, two stars mean a rejection at the 5%-level and three stars mean a rejection at the 1%-level.

5.2.8 Summary

Overall, it seems that for the full period we find that cost granger causes volatility for all currencies. The first period shows signs of cost granger causing volatility for all currencies except for EUR. For CHF and GBP we also find that volatility granger causes cost. The main result from the second period is that there is no granger causality between cost and volatility. Finally, the third period indicates granger causality between both variables for GBP, from cost to volatility for CHF, the opposite for JPY and no causality for EUR.

Considering all the results of the granger causality test we find it difficult to conclude either that cost only granger causes volatility, and vice versa, or even that there at all times exists a granger causality both ways. Since our thesis statement asks the question of whether cost affects volatility we shall continue with this problem and focus on the situations where volatility is used as the dependent variable. As mentioned, our results do not indicate that the causality should go in one specific direction, which, if anything, indicates that the problem of endogeneity is present.

6 The model

In this chapter we will evaluate the results of our regression models. In section 6.1 we will explain what models we have used based on their stationary conditions explained in chapter 4, and how we have treated the outliers in the models. In section 6.2 we will setup our regression models for the four currencies' different periods. Section 6.3 quickly recaps the effects a violation of one of the main OLS assumptions have on the models' results as well as to what degree our models obey the assumptions. Section 6.4 describes the results of our own models and compares them to other approaches and other results from previous research.

6.1 Scenarios

In chapter 4 we looked at the GBP time series to determine if the data was normally distributed and stationary. We found that the data was only near normally distributed even though we transformed it via the normality distribution filter: the weekly average. This statement applied to all the four currencies. We have decided to continue with this data, as we believe that with a near normal distribution this will not have a major effect on our forthcoming results. We also looked at the currencies' time series to see if some of them had a unit root. In some cases the tests were unclear wherefore we decided to continue with both cases for the given time series. In Table 5 the stationarity results of the different periods for GBP, CHF, JPY and EUR are displayed.

		Period			
Currency	**Data**	*Full*	*First*	*Second*	*Third*
GBP	*Cost*	I(1)	I(0)	I(1)	I(1)
	Volatility	I(0) , I(1)*	I(0) , I(1)**	I(0) , I(1)*	I(1)
	Vix	I(1)	I(1)	I(1)	I(1)
CHF	*Cost*	I(1)	I(0) , I(1)**	I(1)	I(0) , I(1)*
	Volatility	I(0) , I(1)*	I(0)	I(0)	I(1)
	Vix	I(1)	I(1)	I(1)	I(1)
JPY	*Cost*	I(1)	I(0)	I(1)	I(0) , I(1)*
	Volatility	I(0) , I(1)*	I(0) , I(1)**	I(0)	I(1)
	Vix	I(1)	I(1)	I(1)	I(1)
EUR	*Cost*	I(1)	I(0)	I(1)	I(1)
	Volatility	I(0) , I(1)*	I(0) , I(1)**	I(1)	I(1)
	Vix	I(1)	I(1)	I(1)	I(1)

Table 5 - Stationarity table: The table describes if the different series do not have a unit root, denoted by I(0), or have a unit root, denoted by I(1). A model made for e.g. CHF second period will be denoted by I(1), I(0), I(1). *: No difference and **: Clear difference.

To find the best fitted model for the four currencies at the different periods, we have run regressions for all the stationary scenarios displayed in Table 5. The model used is given in equation (62). E.g. for GBP's full period we have run two models; The first model is where only cost and VIX are differenced, therefore the model is denoted by I(1)I(0)I(1). In the second model all the variables are differenced, wherefore the model is now denoted by I(1)I(1)I(1). After running the regressions for all the scenarios we have compared the results to discover if more models for the same currency and period have been run. We have described this using the stars in Table 5 that will be explained in section 6.1.1. For an easier interpretation of the forthcoming models the best scenario would be if all the variables included would have the I(0) form. It would then be possible to see exactly how cost affects volatility and not how Δcost affects Δvolatility, which would complicate the interpretation of the model.

6.1.1 Choosing the models based on stationarity

Looking at the scenarios under the full period in Table 5 we find that the volatility time series for CHF, GBP, JPY and EUR all have one star. In all of these cases we have decided to continue with I(1), because the outcome is the same as in I(0) and we find the interpretation of the model to be easier if all the variables have the same form when running the regression. For the same reasons just described we also chose to use differenced cost measures for the third period of CHF and JPY. We have also given the case of volatility during the second period for GBP one star. Contrary to the one star situations of the full and third period, we have chosen to continue with I(0) in this case. Although we would like the variables within the model to be the same for ease of interpretation, we would also like the models across the currencies to be consistent. This is also due to ease of interpretation when the models across the currencies are compared. We have further looked into both scenarios and the result is that it makes no difference if we choose I(0) or I(1), because either way the beta value of cost will not be accepted in the model. Choosing to continue with I(0) for GBP's volatility means that in the cases of GBP, CHF and JPY's second period all variables have the same form, contrary to EUR's second period where volatility conclusively had a unit root, which makes the model different from the other second period models.

Looking at the scenarios under the first period, we see that cost for CHF as well as volatility for GBP, JPY and EUR all have two stars. In these cases, we find a clear difference in the models' outcome. If we use I(1), the models simply do not work, meaning that the cost variables are not significant, and for I(0) cost is significant at a 1%-level in all currency cases. Furthermore, we have compared the test-statistics': normality, autocorrelation, multicollinearity and heteroscedasticity, of the models for the two scenarios and find no significant difference in these tests. With this in mind, we have decided to continue with the I(0) cases, because the study revolves around finding out what effect cost has on volatility.

6.1.2 Outliers

In some of the chosen models described above we find some extreme outliers looking at the residuals. In some of these cases we have been able to improve the models by adding an outlier-dummy or removing outliers.

Our first thought was to include an outlier-dummy, which is denoted by 1 in the weeks with the extreme outliers and 0 the rest of the weeks. The dummy was not significant in all cases, and is therefore only included in two scenarios: GBP's full period and JPY's full period, cf. Table 6.

In GBP's full period the significance of the beta value of cost increases from a 5 to a 1 per cent level compared to before including the dummy variable. Furthermore, the determination coefficient increases by 4.39 percentage points and the residuals get closer to being normally distributed. In JPY's full period the cost's coefficient does not become significant even after introducing the dummy, but the Anderson-Darling value decreases and the determination coefficient increases by 14.45 percentage points.

In three cases we have removed outliers: CHF's third period, JPY's second period and EUR's first period. This was not our first priority, as removing observations could give a misleading picture of what is really happening, but we have chosen to do so in these three cases because the outliers removed only count for around 1 - 2.5% of all the individual models' observations, and the exclusion improves the models significantly. Furthermore, the observations we have removed are not jointly connected.

In CHF's third period, the beta value of cost becomes significant at a 10%-level after removing only three observations. Moreover, the models' residuals are now normally distributed and the only negative thing is that the p-value of the Breusch-Pagan test drops from 0.66 to 0.01. In JPY's second period, the t-value of cost the coefficient does not improve, but the Breusch-Pagan and Anderson-Darling measures do. This means that the t-tests are more accurate even though cost still does not explain volatility. In EUR's first period, the improvement can be seen in the test of autocorrelation, which drops from -0.15 to -0.10, and in the Anderson-Darling value, which is now significant at a 5%-level. All the results described above as well as the final models are displayed in Table 6.

6.2 Ordinary least squares regressions

Based on section 3.1, we chose to use OLS regressions to estimate the relationship between cost and volatility. Our decision to use OLS is based on that this is what previous approaches have used, and our results will therefore be more easily comparable with previous research. Furthermore, we saw in section 4.3.1 that the relationship between the dependent and the independent variables are linear. To control if there are any other possible variables that might affect the determination of volatility we have included lagged values of volatility and VIX, as discussed in section 4.3 and 4.4, respectively, as well as the dummy variable explained in section 6.1.2. In section 6.3.5, we will touch upon some possible criticisms of our model.

Our four models for each of the four currencies can be expressed by matrix notation as in equation (62):

$$V = X\beta + u \tag{62}$$

where V and u are two 16x1 vectors, as seen in equation (63) and (64), consisting of the volatility measures at time t, which we are trying to estimate and the residuals, respectively.

$$V = \begin{bmatrix} \Delta V_{CHF,full}^{t} \\ V_{CHF,first}^{t} \\ V_{CHF,second}^{t} \\ \Delta V_{CHF,third}^{t} \\ \Delta V_{GBP,full}^{t} \\ V_{GBP,first}^{t} \\ V_{GBP,second}^{t} \\ \Delta V_{GBP,third}^{t} \\ \Delta V_{JPY,full}^{t} \\ V_{JPY,first}^{t} \\ V_{JPY,second}^{t} \\ \Delta V_{JPY,third}^{t} \\ \Delta V_{EUR,full}^{t} \\ V_{EUR,first}^{t} \\ \Delta V_{EUR,second}^{t} \\ \Delta V_{EUR,third}^{t} \end{bmatrix} \quad (63) \qquad u = \begin{bmatrix} u_{CHF,full} \\ u_{CHF,first} \\ u_{CHF,second} \\ u_{CHF,third} \\ u_{GBP,full} \\ u_{GBP,first} \\ u_{GBP,second} \\ u_{GBP,third} \\ u_{JPY,full} \\ u_{JPY,first} \\ u_{JPY,second} \\ u_{JPY,third} \\ u_{EUR,full} \\ u_{EUR,first} \\ u_{EUR,second} \\ u_{EUR,third} \end{bmatrix} \quad (64)$$

In equation (62) X is a 16x9 matrix as displayed in equation (65). X contains the explanatory variables in our models, where column 1 defines if the model has an intercept or not.

$$X = \begin{bmatrix}
1 & \Delta C_{CHF,full} & 0 & 0 & \Delta V_{CHF,full}^{t-1} & \Delta V_{CHF,full}^{t-2} & \Delta V_{CHF,full}^{t-3} & \Delta VIX_{CHF,full} & 0 \\
0 & C_{CHF,first} & 0 & 0 & V_{CHF,first}^{t-1} & V_{CHF,first}^{t-2} & 0 & \Delta VIX_{CHF,first} & 0 \\
1 & \Delta C_{CHF,second} & 0 & 0 & V_{CHF,second}^{t-1} & V_{CHF,second}^{t-2} & 0 & \Delta VIX_{CHF,second} & 0 \\
1 & 0 & \Delta C_{CHF,third}^{t-1} & 0 & \Delta V_{CHF,third}^{t-1} & \Delta V_{CHF,third}^{t-2} & \Delta V_{CHF,third}^{t-3} & \Delta VIX_{CHF,third} & 0 \\
1 & \Delta C_{GBP,full} & 0 & 0 & \Delta V_{GBP,full}^{t-1} & \Delta V_{GBP,full}^{t-2} & \Delta V_{GBP,full}^{t-3} & \Delta VIX_{GBP,full} & D_{GBP,full} \\
0 & C_{GBP,first} & 0 & 0 & V_{GBP,first}^{t-1} & V_{GBP,first}^{t-2} & 0 & \Delta VIX_{GBP,first} & 0 \\
1 & \Delta C_{GBP,second} & 0 & 0 & V_{GBP,second}^{t-1} & V_{GBP,second}^{t-2} & 0 & \Delta VIX_{GBP,second} & 0 \\
1 & 0 & 0 & \Delta C_{GBP,third}^{t-2} & \Delta V_{GBP,third}^{t-1} & 0 & 0 & 0 & 0 \\
1 & 0 & \Delta C_{JPY,full}^{t-1} & 0 & \Delta V_{JPY,full}^{t-1} & \Delta V_{JPY,full}^{t-2} & \Delta V_{JPY,full}^{t-3} & \Delta VIX_{JPY,full} & D_{JPY,full} \\
0 & C_{JPY,first} & 0 & 0 & V_{JPY,first}^{t-1} & V_{JPY,first}^{t-2} & 0 & \Delta VIX_{JPY,first} & 0 \\
1 & \Delta C_{JPY,second} & 0 & 0 & V_{JPY,second}^{t-1} & V_{JPY,second}^{t-2} & 0 & \Delta VIX_{JPY,second} & 0 \\
1 & 0 & \Delta C_{JPY,third}^{t-1} & 0 & \Delta V_{JPY,third}^{t-1} & \Delta V_{JPY,third}^{t-2} & 0 & \Delta VIX_{JPY,third} & 0 \\
1 & \Delta C_{EUR,full} & 0 & 0 & \Delta V_{EUR,full}^{t-1} & \Delta V_{EUR,full}^{t-2} & \Delta V_{EUR,full}^{t-3} & \Delta VIX_{EUR,full} & 0 \\
0 & C_{EUR,first} & 0 & 0 & V_{EUR,first}^{t-1} & 0 & V_{EUR,first}^{t-3} & \Delta VIX_{EUR,first} & 0 \\
1 & \Delta C_{EUR,second} & 0 & 0 & V_{EUR,second}^{t-1} & V_{EUR,second}^{t-2} & V_{EUR,second}^{t-3} & \Delta VIX_{EUR,second} & 0 \\
1 & \Delta C_{EUR,third} & 0 & 0 & \Delta V_{EUR,third}^{t-1} & 0 & 0 & \Delta VIX_{EUR,third} & 0
\end{bmatrix} \quad (65)$$

The remaining term in equation (62), β, is a 1x9 vector of β's from β_0 to β_8. The variables that have an anterior Δ correspond to the variables that have a unit root, as explained in section

6.1, wherefore they are differenced for the reasons explained in chapter 4. The results of equation (62) are displayed in Table 6.

	GBP				CHF				JPY				EUR			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1)	I(0) , I(0)	I(1) , I(0)	I(1) , I(1)	I(1) , I(1)	I(0) , I(0)	I(1) , I(0)	I(1) , I(1)	I(1) , I(1)	I(0) , I(0)	I(1) , I(0)	I(1) , I(1)	I(1) , I(1)	I(0) , I(0)	I(1) , I(1)	I(1) , I(1)
Intercept	0.0000	-	0.0014	0.0000	0.0000	-	0.0018	0.0000	-0.0001	-	0.0016	0.0000	0.0000	-	0.0000	0.0000
	(-1.06)	-	*(7.16***)*	*(0.03)*	*(-0.02)*	-	*(9.73***)*	*(-0.22)*	*(-1.36)*	-	*(7.40***)*	*(-0.02)*	*(0.12)*	-	*(0.14)*	*(0.18)*
Cost	1.8942	5.4734	-0.6999	-	-1.1238	7.6476	-0.8499	-	-	5.2806	-0.9075	-	-0.9873	2.8011	-2.4735	-3.2611
	*(2.89***)*	*(8.13***)*	*(-0.67)*	-	*(-1.62)*	*(6.80***)*	*(-0.99)*	-	-	*(7.88***)*	*(-0.84)*	-	*(-1.00)*	*(2.81***)*	*(-1.72*)*	*(-0.67)*
Cost lag1	-	-	-	-	-	-	-	2.7673	0.7938	-	-	4.3639	-	-	-	-
	-	-	-	-	-	-	-	*(1.84*)*	*(1.55)*	-	-	*(1.21)*	-	-	-	-
Cost lag2	-	-	-	-21.1668	-	-	-		-	-	-	-	-	-	-	-
	-	-	-	*(-2.85***)*	-	-	-		-	-	-	-	-	-	-	-
Volatility lag1	-0.6018	0.4217	0.3353	-0.5048	-0.6107	0.3523	0.4233	-0.6816	-0.5575	0.4711	0.4043	-0.5960	-0.6296	0.5966	-0.6190	-0.5841
	*(-20.64***)*	*(10.14***)*	*(7.24***)*	*(-6.52***)*	*(-20.85***)*	*(5.87***)*	*(11.83***)*	*(-7.43***)*	*(-21.79***)*	*(11.82***)*	*(8.27***)*	*(-7.42***)*	*(-16.22***)*	*(10.23***)*	*(-10.79***)*	*(-7.90***)*
Volatility lag2	-0.2509	0.2402	0.3260	-	-0.3488	0.2345	0.2330	-0.3690	-0.3248	0.2335	0.2439	-0.2680	-0.3488	-	-0.4054	-
	*(-7.71***)*	*(5.82***)*	*(7.01***)*	-	*(-10.63***)*	*(4.09***)*	*(6.45***)*	*(-3.40***)*	*(-11.69***)*	*(5.90***)*	*(4.96***)*	*(-3.34***)*	*(-8.03***)*	-	*(-6.42***)*	-
Volatility lag3	-0.1378	-	-	-	-0.1928	-	-	-0.2116	-0.2040	-	-	-	-0.2507	0.2448	-0.1829	-
	*(-4.78***)*	-	-	-	*(-6.61***)*	-	-	*(-2.26**)*	*(-8.01***)*	-	-	-	*(-6.53***)*	*(4.28***)*	*(-3.19***)*	-
Delta Vix	0.0001	0.0001	0.0001	-	0.0001	0.0002	0.0002	0.0000	0.0002	0.0002	0.0001	0.0002	0.0001	0.0001	0.0001	0.0001
	*(4.95***)*	*(3.46***)*	*(2.93***)*	-	*(5.66***)*	*(2.65***)*	*(7.01***)*	*(-0.10)*	*(9.27***)*	*(5.32***)*	*(4.18***)*	*(4.62***)*	*(4.14***)*	*(2.78***)*	*(2.26**)*	*(1.21)*
Outlier-dummy	0.0024	-	-	-	-	-	-	-	0.0086	-	-	-	-	-	-	-
	*(8.43***)*	-	-	-	-	-	-	-	*(16.64***)*	-	-	-	-	-	-	-
R^2	0.3132	-	0.3359	0.3371	0.3019	-	0.3719	0.3510	0.4322	-	0.3496	0.3900	0.3217	-	0.3066	0.3538
N	1095	551	418	116	1095	271	694	116	1095	566	393	123	619	193	294	124
Autocorrelation	0.0170	-0.0378	-0.0608	-0.0188	-0.0216	-0.0347	-0.0302	-0.0159	-0.0026	-0.0501	-0.0210	-0.0305	-0.0454	-0.0997	-0.0411	-0.0844
Breusch-Pagan	<.0001	<.0001	0.2411	0.2357	0.2652	0.1481	<.0001	0.0124	<.0001	<.0001	0.3248	0.0938	0.0816	0.0756	0.5069	0.1098
Condition index	2.1386	7.9320	10.4681	1.1613	2.1243	8.9813	9.8367	2.4172	1.9227	6.5205	10.4322	1.6478	2.1714	8.7989	2.1133	1.1415
Anderson-Darling	13.6685	9.2474	2.3839	0.8188	10.8477	3.9487	2.1553	0.9472	8.6226	11.6344	3.0448	0.8484	6.2013	0.9691	0.2698	2.5666
Skewness	1.5192	1.6777	0.7136	0.7161	1.2957	1.5590	0.3354	0.6881	1.7366	2.6790	0.5829	0.3864	0.4597	0.4163	0.0834	0.5641
Kurtosis	9.7306	6.2597	0.9731	2.8347	9.2430	5.7368	0.2132	1.7586	13.1393	20.2827	0.6613	1.3812	6.9471	-0.0119	0.3135	5.4393

Table 6 – The results of our model for GBP, CHF, JPY and EUR over all four periods. T-values are shown in parentheses and the test statistics are displayed beneath each model. One star means that the null hypothesis is rejected at the 10%-level, two stars mean a rejection at the 5%-level and three stars mean a rejection at the 1%-level.

6.3 Model bias

During chapter 5 we explained in detail the repercussions of violating the OLS assumptions of normality, multicollinearity, autocorrelation and homoscedasticity. The following section will describe if our models of volatility for each currency in each time period are affected by a violation of one or more of the four assumptions, and if so can we say anything about the direction of the bias? Answering this will give us a better understanding of the quality of our estimated coefficients, which we will comment on in section 6.4. Section 6.3.5 will highlight some possible points of criticism our model might be subjected to.

6.3.1 Normality in the models

When we discussed the implications of non-normality in section 5.1.1, we concluded that our hypotheses more often than not would be accepted in the case of low kurtosis and positive skewness, and would be rejected more often than not when high kurtosis and negative skewness were present.

In section 4.2.1 we talked about how our data was more leptokurtic than normally distributed which is also the case for many of our models of volatility. Normality in the residuals are however present in the third period for CHF, GBP and JPY, as well as in the first and second period of EUR.

For all models, we have skewness that is positive, which, ceteris paribus, means that our t-values are underestimated compared to what they would have been under a normal distribution. This direction of the bias makes it easier to conclude that our explanatory variables are significantly different from zero.

The kurtosis is relatively high for a lot of the models, such as the full and first period for CHF, GBP and JPY, as well as for the full and third period for EUR. The direction of the high kurtosis overestimates our t-values, meaning that our null hypotheses are more accepted than they look.

The models with the highest kurtosis are also the ones with the most positive skewness – two biases working in opposite directions. Although when considering the estimated beta values, we should probably consider the size of the t-values for the full period of CHF and GBP, as well as the full and first period of JPY and the full and third period of EUR where the kurtosis is more than seven times higher than the skewness.

6.3.2 Multicollinearity in the models

In section 5.1.3 we explained that multicollinearity between two or more variables in the model would lead us to accept our null hypotheses more often than usual. Since we are most interested in rejecting the null hypothesis, which is that any of our beta estimates are equal to 0, the bias caused by multicollinearity works to our advantage. Ceteris paribus, if no multicollinearity is present we can trust our t-values, and if we do find multicollinearity our hypotheses will be more rejected than it seems. Of course it still affects the estimated beta values of the correlated variables, and here we will be going by the common rule-of-thumb (Gujarati (2003)) and only consider multicollinearity to be a problem if the condition index rises above ten.

If we look at all the models in Table 6 for all the full and third periods for the four currencies ,as well as the second period for EUR, which are based on the scenario I(1)I(1)I(1), we see that the condition index floats around 2. In these cases we should not see an effect on either our t-values or the estimated beta values. The first period for all currencies based on scenario I(0)I(0)I(1) has a condition index between 6 and 9. According to the rule-of-thumb, this should not be a problem considering the estimated values of the betas. The most severe cases of multicollinearity occur in the second period for currencies: CHF, GBP and JPY where cost is differenced and volatility is not. The condition index for CHF is still below ten opposed to GBP and JPY where it is calculated to approximately 10.5. This could mean that some of the estimated beta values are not distributed correctly between the correlated variables. However, considering that the rule-of-thumb deems the range 10-30 a possible problem of multicollinearity, our most extreme values will probably not cause any real trouble.

6.3.3 Autocorrelation in the models

The correlation of the residuals with themselves does, as explained in section 5.1.2, underestimate the standard errors of the coefficients in case of positive autocorrelation and overestimate them in case of negative autocorrelation.

Looking at Table 6, we see that for all sixteen models the autocorrelation ranges between ±0.01, which must be considered too small to have any noticeable effect on our estimated coefficients.

6.3.4 Heteroscedasticity in the models

As mentioned in section 5.1.4, heteroscedasticity also affects our estimates of the standard deviations of the estimated beta values. As opposed to the violation of the other assumptions, we cannot say anything about the direction of the resulting bias.

In Table 6 we have reported the Breusch-Pagan p-value for each model. Note that this value should be as close to 1 as possible. Since we want to accept the null hypothesis of homoscedasticity, a p-value above 0.9 means that we will accept heteroscedasticity in 10% of the cases, which should be sufficient to discard heteroscedasticity as a serious problem for our coefficients. Anything between 0.9 and 0.95 is less worrisome and anything above 0.95 should be considered unproblematic.

In the majority of our models heteroscedasticity is a problem, which we, from our findings in section 5.2.6, believe to be caused by the lagged volatility variables, although with the full and first periods for CHF as well as the second and third for GBP, the second period for JPY and the second and third period for EUR the problem of heteroscedasticity is smallest with Breusch-Pagan p-value ranging between 0.1098 and 0.5069. The third period of JPY and the full and first periods of EUR have a Breusch-Pagan p-value ranging between 0.1 and 0.05. Depending on the size of these models' estimated betas t-values this needs to be included in our evaluation of their results. This is also the case for the third period of CHF, which is somewhat more heteroscedastic at a p-value equal to 0.0124.

In the last five models; namely the second period of CHF and the full and first period of GBP and JPY, we find undisputable evidence of heteroscedasticity where the Breusch-Pagan p-value is less than 0.0001.

6.3.5 Critique

In section 6.3.1 we described how our data is only near normally distributed and how this affects our estimated coefficients. The leptokurtic nature of our data means that a regression model based on this distribution instead of the normal distribution would fit the data better and give more reliable coefficient estimates. We found this to be too comprehensive for this thesis work, wherefore we chose to continue with our OLS models. One of the disadvantages of our OLS models is the heteroscedasticity present in our residuals. We tried different solutions to solve this problem but none of them were sufficient, wherefore we again chose to continue with our OLS model. If research continues on this subject it might be worthwhile considering setting up a proper GLS model that would solve the problem of heteroscedasticity. Last, but not least, we have the problem of endogeneity discussed in section 3.4, which we found in section 5.2.7 that we could not dismiss. Therefore it might be difficult to conclude anything about the true causality between cost and volatility. We will try to solve this problem in section 7.3.

6.3.6 Summary

Real world data rarely lives up to all the assumptions of OLS or other statistical theories, which we also pointed out in section 6.3.5. Therefore, in sections 6.3.1 to 6.3.4, we have described how our models either obey or violate any of those assumptions. As described, some of the models only have mild violations of some of the assumptions, whereas others are more severe in their violations. None of the models had any autocorrelation in the residuals, and the multicollinearity that does exist between some of the variables is quite small and underestimates the t-values. The violations to take notice of are the heteroscedasticity present in the second period of CHF and the full and first period of GBP and JPY. The lack of normality, or more precisely, the presence of high kurtosis in the full period for CHF and GBP, the full and first period of JPY and the full and third period of EUR also needs to be taken into account when evaluating the results of the models. Furthermore, we mentioned the problem of endogeneity, which we have yet to solve but in section 7.3 we will propose a possible solution to the problem.

6.4 Results

First of all, this section evaluates the results of our own regressions. Through sections 6.4.1.1 to 6.4.1.5 we shall look at the beta coefficients of our lagged volatility variables, VIX, dummy and cost, respectively, as well as the coefficient of determination. Section 6.4.2 compares our results with a naïve model, a LN model and a spill-over model and section 6.4.3 compares our result with the results of previous research done by Mulherin (1990), Aliber, Chowdhry & Yan (2003), Lanne & Vesala (2006), Ronan and Weaver (2001) and Hau (2006).

6.4.1 Our model

In this section we will describe the estimated coefficients of our models presented in Table 6. The evaluation of the coefficients will help us to determine which variables have a significant impact on the prediction of volatility and if this is a positive or a negative relationship. It should be pointed out that we have included some insignificant coefficients in the model, such as VIX, cost and the intercept. We have done so, when they did not disturb the outcome of the model, to be able to see how close they are to be significantly different from zero.

6.4.1.1 Lagged volatility

If we look at the estimated beta-coefficients for the different lags of volatility included in the models, we see that all of them are significantly different from zero at a 1% significance level, except for CHF's third period, which is significant at a 5%-level. Even though the t-values for the full period for all the currencies as well as the first and third period for JPY and EUR, respectively, are somewhat underestimated due to a high kurtosis, we see no reason to discard the results of the estimated beta-values for the lagged volatility measures. The presence of heteroscedasticity does affect our t-values but as mentioned in section 5.1.4, we cannot determine the direction of its bias.

For the full period, all the volatility variables are negative for all currencies. We also find that the sizes of the coefficients are relatively close to each other when comparing the different lags between currencies. The first lag has the highest impact on volatility with a magnitude of approximately -0.6, meaning that a change in the first lagged volatility variable of 1 the week before would result in a decrease of -0.6 the week after. Through the second and third lag of volatility the estimated coefficients stay relatively close to each other between the currencies at about -0.3 and -0.2, respectively.

In the first period we find that only the first and second lag of volatility have a significant impact on volatility for all currencies besides EUR, where the first and third lag are different from zero. For all of the estimated coefficients we find that they have a positive influence on volatility with about 0.46 for the first lag and approximately 0.24 for the second and third lag.

The second period resembles the first one except for EUR, where volatility is differenced opposed to the rest of the currencies. This results in that all the coefficients are negative and numerically larger than they are for the remaining three currencies. For the first and second lag for CHF, GBP and JPY the beta-values are estimated at about 0.39 and 0.27, respectively

- where the EUR is estimated at approximately 0.62, 0.41 and 0.18 for the first to third lag; much like the full period.

Looking at the third period we once again find that all of the coefficients are negative just as they are in the full period. During this period, the inclusion of lags in the models is not as consistent as the previous periods. For CHF we find that the first to third lags are significant as well as the first and second ones for JPY. For GBP and EUR we find that only the first lag is significant. The magnitude of the estimated coefficients resembles the ones for the full period, where the first has an estimated beta-value of about -0.59.The second lag is estimated to approximately -0.32 and the third lag to -0.21.

The results give a clear picture of the relationship between volatility in a certain week, and lagged values of volatility in the previous one to three weeks, as well as the relationship between the differenced volatility and the lagged differences. In all the cases where volatility is I(0) the relationship between volatility and its lagged values is positive, and when volatility is I(1) the relationship is negative – even during the second period where EUR is the only model to use differenced volatility measures.

The positive coefficients of the first and second period indicate that the volatility moves in the same direction for at least three weeks, where the negative coefficients for the full and third period show us that the change in volatility fluctuates more rapidly between weeks.

6.4.1.2 VIX

The inclusion of VIX in the models was to see how the currencies' volatility was affected by the general market volatility. For all the models, except the third period for GBP where we chose not to include VIX due to multicollinearity, we find a positive relationship between VIX and volatility. The estimated beta-value for VIX is significantly different from zero at a 1% significance level for all models except for the second period of EUR where it is only significantly different at a 5%-level, and the first and third period for CHF and GBP, respectively, where they are deemed not to be different from zero even at a 10%-level.

The t-values for VIX are affected by the bias in the same cases as explained under volatility. Since VIX is mostly as significant as volatility we apply the same rationale and have decided to trust the beta values.

6.4.1.3 Outlier-dummy

As described before, we have chosen to include a dummy variable in the full model for GBP and JPY. These dummies have a positive and significant effect on volatility. Their estimated coefficients of 0.0024 and 0.0086 for the full period of GBP and JPY, respectively, seem rather small compared to the other beta-values, but one has to remember the relatively large size of the dummy variables' value of 1 compared to the much smaller volatility and cost measures.

6.4.1.4 Cost

Until now we have seen fairly consistent results among the estimated beta-values for the three previously mentioned variables. When including cost in the model, we found that it was only possible to include it once and not possible to include multiple lags simultaneously, as it was with volatility due to multicollinearity. Therefore, we see that the third period for CHF and JPY, as well as the full period for JPY, only includes the first lag of cost, and the third period of GBP only includes the second lag.

The results for the estimated beta-values for cost indicate that a positive relationship between cost and volatility exists during the first period for all currencies and a negative relationship during the second period, also for all currencies. During the first period, both cost and volatility are non-differenced and the beta-values lie around 5 - with CHF being a bit larger at 7.6 and EUR somewhat smaller at 2.8. During the second period, the beta-values are approximately -0.8 except for EUR, which has a beta-value of -2.5.

For the full period we find that CHF and EUR have a negatively estimated beta-value for cost of about -1, and GBP and JPY have positive beta-value of approximately 1.9 and 0.8, respectively. The third period has mixed results, as well as the full period, with GBP and EUR showing a negative relationship between cost and volatility, and CHF and JPY displaying a positive relationship. The size of the coefficients differs quite a lot during this period with GBP having a beta-value of -21.2 and EUR with a beta-value of -3.3. For CHF and JPY the beta-values lie at 2.8 and 4.4, respectively.

Once again, we have to consider the bias that affects the estimated t-values hence the validity of our beta-values. All the currencies' full periods have overestimated t-values due to high kurtosis. Considering that all but the beta-value for cost for GBP during this period is not accepted at the 10%-level, it is difficult to consider them different from zero. GBP's beta-value for the full period of 1.9 is, although significantly different from zero, at a 1%-level. The first period produces rather reliable beta-values, which, at a 1%-level, are all significantly different from zero. We find that the beta-value for CHF and EUR does not suffer from any bias, whereas the t-values for cost for JPY and GBP are overestimated, but here we also find a somewhat high multicollinearity that underestimates the t-values. It is difficult to say anything about their actual effect but considering that the beta-values are significant at a 1%-level we determine them to be reliable. For the second period, we find that all t-values are underestimated, except for EUR, which, on the contrary, has a significant effect at a 10%-level. The t-values for CHF, GBP and JPY are relatively small, and are probably difficult to assign any significant explanatory power at a 10%-level or below. Looking at the third period, we see that only the beta-values for CHF and GBP are significant at a 1 and 10 per cent level, respectively, without being biased in any way. For JPY we find no bias or any significant explanatory power. The same is the case for EUR, which even suffers from a small overestimation. As with the other variables, heteroscedasticity is also a problem here. In the case that the bias works against us, we assume that the beta coefficients for cost deemed significantly different from zero are still reliable due to their rather large t-values.

6.4.1.5 Coefficient of determination

When we look at the coefficient of determination (R^2) for the models, a very general picture is formed. Most of the coefficients are approximately at 0.35 whereas the model for the full period for CHF has the lowest R^2 of 0.30 and the full period for JPY has the highest of 0.43. When estimating the models for the first period for all the currencies, we found that including the intercept resulted in a very high correlation with cost and therefore decided to exclude the intercept from the model. This procedure imposes a mean of 0 for our volatility, which inflates the R^2 of the model as the total sum of squares is based on a deviation from 0 and not the actual mean of the volatility. For the models, there seem to be no systematic tendency as to what period is explained the best across the currencies.

In section 5.2.4 we showed that the coefficient of determination could be biased in the presence of autocorrelation in the error term, although as we saw in section 6.3.3 our models do not contain any autocorrelation wherefore this is not a problem, and thus our R^2-values are not biased.

6.4.1.6 Summary

Through section 6.4.1, we have shown how much significance we should put into the estimated beta-values of our models and hence what variables are significant in helping us explain the variance in the volatility of the exchange rates.

Considering cost, we find that the estimated beta-values are significantly different from zero and positive for GBP during the full period, for all currencies during the first period and for CHF during the third period. Furthermore we find that the estimated beta-values for EUR and GBP during the second and third period, respectively, are also significantly different from zero although negative. In all other models the beta-value is insignificantly different from zero, while still having a negative value except for JPY during the full and third period. For all other variables included in the model be it volatility, VIX or our dummy variable we argue that they are significantly different from zero, even considering the bias they might be affected by. For all of our models we find R^2-values between 30 and 40 per cent, except for the first period models for all currencies, which have no intercept. Considering the possibility of factors affecting real world data we find our models to give a good estimation of the value of volatility.

6.4.2 Relative to other approaches

In this section we will focus on three alternative approaches to the model we have chosen and see how these new models perform compared to our results displayed in Table 6. The results will be compared by the determination coefficient, the signs, size and their level of significance of the beta values. Furthermore, we will compare the test values of heteroscedasticity, multicollinearity, autocorrelation and normality to determine in which direction the results could be biased and to establish which model works better.

6.4.2.1 The naïve model

The naïve model is called so because we only regress volatility on the first lagged value of volatility. By doing this, we can see if volatility this week explains volatility next week, thus we can write the model for GBP as:

$$V_t^{GBP} = \beta_0 + \beta_1 V_{t-1}^{GBP} + u_t^{GBP} \quad (66)$$

The results for GBP in the naïve model can be found in Table 8 and for CHF, JPY and EUR they can be seen in Appendix 8.

Comparing the results from our model with the naïve model we find some similarities. The beta coefficients for the full and third periods are negative at -0.4660 and -0.5352, respectively, and the coefficients for the first and second period are positive at 0.9530 and 0.9731. Furthermore, all the coefficients are extremely significant. This resembles the same relationship with I(0) and I(1) as explained in section 6.4.1.1. In the first and second periods the intercept was correlated with lagged volatility, which is why we have excluded it. This means that because the procedure imposes a mean of 0 we cannot compare the naïve model to our model based on the determination coefficient. In the cases of the full and third period we can compare the determination coefficients and in both cases our models' coefficient of determination exceeds the naïve model by 9.63 and 5.13 percentage points, respectively.

Comparing the test statistic from our model with the naïve model we find that the Anderson-Darling value is a little smaller, although still only nearly exhibits a normal distribution for the naïve model in the first and second period. When looking at skewness and kurtosis for both models, we find that skewness is slightly larger for our model during all periods and that kurtosis for our model performs better in the second and third period compared to the naïve model. The differences are very small, and therefore we believe that this does not have a significant impact when comparing the two models.

In our model autocorrelation is almost non-existent, but when looking at the naïve model we find high negative autocorrelation in the first and second period. This might be due to omitted variables whereby the residuals compensate for this. The high autocorrelation could also explain the very significant t-values for the lagged volatility of 74.54 and 85.95, respectively. Multicollinearity between the variables is not a problem in the naïve model as only one variable, lagged volatility, is included. Thus we find it hard to compare to our model. Furthermore, as described in section 6.3.2, we do not believe that any of our estimated coefficients from our model will be significantly biased due to multicollinearity but this could be a reason why the beta coefficients of the naïve models' first and second periods are higher compared to our model. This is because the effect of lagged volatility is not divided between several correlated lagged volatility variables as in our model.

Looking at the Breusch-Pagan p-values for our model and the naïve model, we find that in the full period the naïve model performs better with a p-value of 0.0158 compared to <.0001. The first periods are alike for the two models but in the second period our model performs much better with a p-value of 0.2411 compared to 0.0096 in the naïve model. In the third period, the p-value for the naïve model exceeds our model by almost three times with a p-value of 0.6776.

After having compared the two models, we find that the signs of the lagged volatility over the different periods are the same compared to our model. Furthermore, we could see that the size of the beta values for the full and third periods were somewhat the same size as in our model. The first and second periods differed in size compared to our model but this is probably due to correlation between lagged volatility variables in our model. Furthermore, we do not find a significant basis for one of the models to perform better than the other when looking at the normality testing or multicollinearity results. It should be noted though that in the naïve model we found a high negative autocorrelation in the first and second period, which is a severe problem. Furthermore, our model performed better when comparing the results of heteroscedasticity and the determination coefficient. Based on this, we can conclude that overall our model performs better than the naïve model, and that cost has an impact when estimating volatility.

6.4.2.2 Spill-over

The purpose with the spill-over model is to find out if either foreign cost or volatility affects domestic volatility. Meaning, does foreign countries' cost and volatility affect the domestic country's volatility? We will investigate this by expanding the models developed in Table 6 to include the foreign variables.

To see if foreign cost should be included in the models we have found the correlation between all the cost time series. The results can be seen in Table 7. The idea is that if the different cost time series are not correlated the traders for a given currency would change to another market to exploit the lower transaction costs if the cost of the first currency increases. From Table 7 we find that costs are highly correlated thereby indicating that the transaction costs are at least partly determined by the market. Including these in the models will only result in multicollinearity.

	Full					**First**			
	CHF	GBP	JPY	EUR		CHF	GBP	JPY	EUR
CHF		62%	57%		CHF		31%	54%	
GBP	62%		64%		GBP	34%		37%	
JPY	57%	64%			JPY	43%	37%		
EUR	28%	72%	62%		EUR	18%	-8%	30%	
	Second					**Third**			
	CHF	GBP	JPY	EUR		CHF	GBP	JPY	EUR
CHF		69%	66%		CHF		28%	26%	36%
GBP	45%		72%	75%	GBP	24%		45%	52%
JPY	62%	69%		63%	JPY	27%	48%		33%
EUR	48%	75%	63%		EUR	37%	59%	33%	

Table 7 - Table of correlations between the cost variables for the different currencies. Left column indicates how much a currency's cost is correlated with a currency in the top rows cost during its own full, first, second or third period.

Looking at the correlations in Table 7 we find that overall costs across currencies are too highly correlated, which is why they should be excluded from the model. So for the GBP model we can write:

$$V_t^{GBP} = \beta_0 + \beta_1 C_{t-i}^{GBP} + \beta_2 V_{t-1}^{GBP} + \beta_3 V_{t-2}^{GBP} + \beta_4 V_{t-3}^{GBP} + \beta_5 VIX_t + \beta_6 Outlier\ dummy_t + \beta_7 V_t^{CHF} + \beta_8 V_t^{JPY} + \beta_9 V_t^{EUR} + u_t^{GBP} \quad (67)$$

Where $i = \{0,1,2\}$ and in the periods where EUR was not yet introduced the variables take the value 0. This is the case in the full and first period for GBP. The results from GBP can be seen in Table 8 and for CHF, JPY and EUR can be seen in Appendix 9.

Comparing the results from our model in Table 6 with the spill-over model, we find no changes in the sign of cost or lagged cost. In the full and second period the changes in the size of the beta values are not extensive, and the significance of the t-values does not change. Especially noticeable are the first and third period wherein the spill-over model and the estimated beta coefficients decrease to 0.52 and -13.13, respectively, and the cost coefficient of the first period changes from being significant at a 1%-level in our model to not being significant at a 10%-level. The significant level of the third period increases in the spill-over model from a 1 to a 5 per cent level compared to our model.

The sign of the beta values and the significance level of all lagged volatility variables do not change. When comparing the models we only find a negligible difference in the size of the beta values. Looking at VIX we find that the variable in the full, first and second period of our model was significant at a 1%-level, but in the spill-over model the variable is not accepted in any of the periods. This could point to that VIX is not the best measure for the market volatility and that instead an average of the different currencies' volatilities would be a better measure. The outlier-dummy in the full period does not change noticeably in either size or significance when comparing the models.

The foreign volatility variables are all significant at a 1%-level except for JPY in the second period and CHF in the third period in the GBP spill-over model. Comparing these results to the spill-over models for CHF, JPY and EUR from Appendix 9 we find that generally in the third period for all the models the relationship is varying and also that the JPY variables more often are insignificant, which can be seen in all the cases of EUR. We believe this could be due to the distance between the countries. We find it logical that EUR, GBP and CHF all affect each other more than JPY since it is centred on another continent.

Looking at the determination coefficient, we find that in all comparable cases the spill-over models perform better. In all cases we see that the determination coefficient of the spill-over model lies approximately 20 percentage points higher than for our model. This could be due to the foreign volatility variables, which give a better market measure than VIX, wherefore the variation in volatility is explained better compared to or model.

Comparing our model to the spill-over model we find no changes in signs of the beta estimates. Furthermore, we find that the cost coefficients in the first and third period decrease noticeably in the spill-over model, which could indicate that the cost coefficients in our model are not that robust in the first and third period. By comparing the test-statistics from our model with the spill-over model we find that the Anderson-Darling value is a little smaller, but still only near a normal distribution for the spill-over model in the full, first and second period. In the cases of skewness and kurtosis we find that our model only performs better in the second period, but the difference in the other periods are small.

Compared to our models the issue of autocorrelation is more severe in the spill-over model. In all periods the correlation between residuals and lagged residuals exceeds ±0.12, which will affect the t-statistics. Looking at the condition index in the spill-over model, we find that in all cases our model performs better. Worrying, though, is the second period of the spill-over model where we find a condition index of 16.33, which means that the size of the beta values of the correlated variables could be incorrectly distributed among them. Looking at the Breusch-Pagan p-values, we find that our model performs better in the second period. The full and the first period for both models are alike, and it is only in the third period that the spill-over model performs better.

After having compared the two models, we find that the spill-over model explains volatility better by looking at the determination coefficient. This is probably due to the better market index created by the foreign volatility variables. Although we have to be careful not to draw our conclusions too quickly based only on this, as the spill-over models performs more poorly overall when looking at the test of multicollinearity, heteroscedasticity and autocorrelation. Based on this, we find it hard to determine which of the two models is better than the other, when comparing the test statistics. Although, if we compare our model with the spill-over model based on the determination coefficient, the spill-over model outperforms our model in all cases. Furthermore, we find that the signs of the cost coefficients have not changed compared to the cost coefficients of our model and they are therefore still in line with previous research.

6.4.2.3 The LN-model

The purpose of comparing an LN-model with our model, is to see if the LN effect, proposed by Hau (2006), will make the model perform better. In his paper, Hau (2006) describes that taking the natural logarithm of the measures of cost and volatility will work as a near Gaussian filter, as we explained in section 3.2.2. This resembles a log-linear model, and it is therefore not possible to compare the sizes of the beta values to our model. The reason for this is that the estimated cost coefficient measures the elasticity of volatility with respect to cost, i.e. the percentage change in volatility for a given percentage change in cost, whereas in our model, we look at changes in percentage points. Furthermore, there is the question of stationarity. As the measures are now changed we might see changes in the stationarity of the periods, but we will assume that these behave in the same way as in our model. We do this so the comparisons of the models are easier to interpret.

When comparing the signs of the costs' beta values from the LN-models and our GBP models, we find that the only change is that in the LN-models' full period the sign is negative. This means that the model has established a negative relationship between volatility and cost, which corresponds to Tobin's (1978) statement, but looking at the associated t-value we find that this beta value is not significant even at a 10%-level. Furthermore, we find that it is only the first period that has a significant cost coefficient at a 1%-level in contrast to the rest of the periods where the null hypothesis of the beta values is accepted, as we found in our model as well. LN-models for CHF, JPY and EUR can be found in Appendix 10.

In the LN-model we find no changes in the signs or the significance level of the lagged volatilities' beta values, compared to our model. Furthermore, do we not see any changes in the outlier-dummy or VIX, except for the second period of the LN-model where the beta value of VIX is now only significant at a 10%-level, compared to our model, which has a significant level of 1%.

Looking at the determination coefficient, we find that our model performs a little better, except for the third period where the LN-model has a determination coefficient of 0.381 compared to ours of 0.337.

Comparing the test statistics of the two different models we find that Hau's (2006) statement, taking the natural logarithm of the measures will feature a near Gaussian distribution, to be partly true. In the full, first and second period the Anderson-darling values improve and in the second period the residuals are normally distributed with a p-value of 0.013. In the third period, however, our model outperforms the LN-model. Here we find an Anderson-darling value of 0.819 compared to the LN-model which has a value of 1.769. Looking at skewness and kurtosis, we find the same pattern as just described; that in general the LN-model performs better. Only in the second period our model performs better with regard to kurtosis and in the third period our model performs better in both kurtosis and skewness.

Looking at autocorrelation, we find that our model performs better in three out of the four cases but the difference between the two models is very small, hence we find the difference insignificant. Comparing the Breusch-Pagan values of the two models, we find that our model performs better in the second and third period. Noticeable is especially the third period in the LN-model where the Breusch-Pagan value is below 0.05. In the tests for multicollinearity our model generally performs better. The first and second period of the LN-model, where the condition index lies between 49.168 and 56.232, are especially troublesome since these numbers indicate an extreme case of multicollinearity in both periods, which leads to accepting the null hypotheses more often than usual. This also means that the estimated beta values of the correlated variables are affected.

After having compared the two models we find that only the first period of the LN-model has a significant positive cost coefficient and that the full, second and third period have negative but insignificant cost coefficients, which correspond to the previous research. That the cost coefficient of the first period does not change in sign or significance level compared to our model could indicate that the coefficient is not that weak as we thought based on section 6.4.2.2. Looking at the test-statistics of the two models, we find that the LN-model generally performs better in the normality tests, even though most of the cases are still only near normally distributed. In the cases of autocorrelation, we believe that the difference is too small to decide which model performs better. In the tests of heteroscedasticity and multicollinearity our models generally perform better, especially when looking at multicollinearity, even though this means that the t-values of the LN-model could be more significant than displayed. Furthermore, we find that the determination coefficients are mostly better in the cases of our model. Based on this and the fact that it complicates the interpretation of the model when taking the natural logarithm to the measures, we conclude that our model is a better model to use when investigating the impact cost has on volatility.

6.4.2.4 Summary

When evaluating the alternative models described in section 6.4.2.1, 6.4.2.2 and 6.4.2.3, we see that in the naïve model the significance and the signs of the lagged volatility coefficients did not change compared to our model. We only found that the size of the first and second periods coefficients increased, but this could be explained by a possible correlation between lagged volatility variables in our model. In the spill-over model we found no change in signs when comparing the cost coefficients to our model. The only thing noticeable was the decrease in size and in the t-values in the first and third periods' cost coefficients, which could indicate that our models' cost coefficients for these two periods were not that robust. In the LN-model we only found the first periods' cost coefficient to be significant and positive like in our model. Otherwise the cost coefficients from the full, second and third period were negative and insignificant. The fact that the first periods' cost coefficient in the LN-model was significant and positive could indicate that it was not as weak as we thought based on our results from the spill-over model.

Based on the comparison of the test-statistics, we found that our model performed better in general even though the differences in most cases were small. This was also the case when looking at the determination coefficient, except in the cases of the spill-over model, but we believe this is due to the fact that the foreign volatility variables are a better measure for the market volatility than VIX and the reason why the variation in volatility is better explained.

Generally we find the alternative models to come to the same results as ours: that either the cost coefficients are positive and significant or negative and insignificant, which correspond to previous research.

	LN GBP			
Period	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1)
Intercept	-0.0043	-	-1.9476	0.0006
	(-0.59)	-	*(-7.33***)*	*(0.02)*
Cost	-0.0110	0.1895	-0.0550	-
	(-0.30)	*(7.83***)*	*(-1.11)*	-
Cost lag2	-	-	-	-0.0235
	-	-	-	*(-0.16)*
Cost lag3	-	-	-	-
	-	-	-	-
Volatility lag1	-0.5979	0.4607	0.3205	-0.6153
	*(-20.23***)*	*(11.22***)*	*(6.90***)*	*(-8.33***)*
Volatility lag2	-0.2790	0.2573	0.3317	-
	*(-8.41***)*	*(6.33***)*	*(7.09***)*	-
Volatility lag3	-0.1550	-	-	-
	*(-5.26***)*	-	-	-
Delta Vix	0.4244	0.5042	0.2256	-
	*(4.96***)*	*(4.05***)*	*(1.80*)*	-
Outlier-dummy	0.2640	-	-	-
	*(4.70***)*	-	-	-
R^2	0.2915		0.3161	0.3807
N	1095	557	418	116
Autocorrelation	-0.0175	-0.0201	-0.0630	-0.0530
Breusch-Pagan	*<.0001*	0.0728	0.2196	0.0279
Condition index	2.1517	49.1684	56.2319	1.0263
Anderson-Darling	3.8691	1.4714	1.0008	1.7689
Skewness	-0.4932	0.4833	-0.0941	-0.8791
Kurtosis	6.0655	1.2996	2.0039	7.6067

	Spill-over GBP			
Period	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1)	I(0) , I(0)	I(1) , I(0)	I(1) , I(1)
Intercept	0.0000	-	0.0006	0.0000
	(-0.94)	-	*(3.11***)*	*(-0.03)*
Cost	1.5514	0.5164	-0.7420	-
	*(-2.85***)*	*(0.83)*	*(-0.86)*	-
Cost lag2	-	-	-	-13.1341
	-	-	-	*(-2.34**)*
Volatility lag1	-0.4098	0.2366	0.3058	-0.2729
	*(-15.95***)*	*(7.02***)*	*(6.63***)*	*(-4.34***)*
Volatility lag2	-0.1706	0.1100	0.1961	-
	*(-6.26***)*	*(3.35***)*	*(4.96***)*	-
Volatility lag3	-0.0936	-	-	-
	*(-3.90***)*	-	-	-
Delta Vix	0.0000	0.0000	0.0000	-
	(-1.49)	*(0.81)*	*(0.63)*	-
Outlier-dummy	*0.0017*	-	-	-
	*(7.23***)*	-	-	-
CHF-Volatility	0.3952	0.5448	0.2646	-0.0013
	*(17.52***)*	*(17.86***)*	*(7.18***)*	*(-0.001)*
JPY-Volatility	0.0848	-0.0571	0.0158	0.2316
	*(4.66***)*	*(-2.97***)*	*(0.54)*	*(3.92***)*
EUR-Volatility	-	-	0.1949	0.4301
	-	-	*(5.22***)*	*(3.92***)*
R^2	0.5296	-	0.5565	0.6460
N	1095	551	418	116
Autocorrelation	-0.1268	0.1539	-0.1332	-0.2507
Breusch-Pagan	*<.0001*	*<.0001*	0.0289	0.1085
Condition index	2.3126	10.9639	16.3348	4.2974
Anderson-Darling	7.9514	5.4774	2.4509	0.6032
Skewness	0.1733	1.0684	0.7517	0.2940
Kurtosis	5.1732	3.0538	1.4966	0.8801

	Naïve GBP			
Period	*Full*	*First*	*Second*	*Third*
Model	I(1)	I(0)	I(0)	I(1)
Intercept	0.0000	-	-	0.0000
	(-0.04)	-	-	*(0.10)*
Volatility lag1	-0.4660	0.9530	0.9731	-0.5352
	*(-17.41***)*	(74.54***)	*(85.95***)*	*(-6.78***)*
R^2	0.2169	-	-	0.2858
N	1097	552	419	117
Autocorrelation	-0.0752	-0.4002	-0.4846	-0.0035
Breusch-Pagan	0.0158	*<.0001*	0.0096	0.6776
Condition index	1.0005	1.0000	1.0000	1.0035
Anderson-Darling	15.8111	5.4598	1.8015	1.4133
Skewness	1.2678	0.4000	0.1076	0.6780
Kurtosis	8.6465	4.4149	2.0131	2.9104

Table 8 - The results of our alternative approaches for GBP over all four periods. T-values are shown in parentheses and the test statistics are displayed beneath each model. One star means that the null hypothesis is rejected at the 10%-level, two stars mean a rejection at the 5%-level and three stars mean a rejection at the 1%-level.

6.4.3 Relative to previous approaches

To compare our results with previous research, we have chosen to look at Mulherin (1990), Aliber, Chowdhry & Yan (2003), Lanne & Vesala (2006), Ronan & Weaver (2001) and Hau (2006). The reason for this is that the five papers all focus on either currencies or that their models are somewhat comparable to ours. Based on this, it is possible to roughly compare the determination coefficient and the signs of the beta values in the different models. As the authors of the different papers have used alternative cost and volatility measures compared to ours, only the signs and not the size of the beta values can be considered when comparing results. Furthermore, we will not describe the alternative measures in detail as these have been loosely described in section 3.2. Not all the papers comment on the test of normality, multicollinearity, autocorrelation and heteroscedasticity, and therefore it is difficult to tell if their coefficients are biased, and if so in what direction. We are aware that the time periods for the five articles all differ from our, which should be taken into consideration when comparing the models, but to our knowledge no one has investigated this relationship based on the same model in recent time. We are aware that the different currencies used in the papers in some cases differ from our four chosen currencies, but as seen in chapter 4, the currencies often behave alike.

Mulherin (1990) looks at the relationship between margin requirements introduced by the Federal Reserve in 1934 and the volatility of the Dow Jones returns on a daily basis from 1934 to 1987. As described in section 3.1, we believe that even though the research is based on the stock market we can compare it to our results, as the relationship between volatility and costs is expected to act the same way. We are aware that Mulherin (1990) uses margin requirements but this is just a different way to measure transaction costs, wherefore we find the models comparable.

The results of Mulherin's (1990) research do not give a clear picture of the relationship between margins and volatility. When he regresses volatility on margins for the full period from 1934 to 1987 he finds that the signs of the cost coefficients are negative, but not significant even at a 10%-level. He then divides the sample into 4 periods of 12 years and the results reveal positive beta values of the costs, but again they are not significant. Comparing this to our models we find that in the majority of the results the relationship between cost and volatility is positive and significant or negative and insignificant. Mulherin's (1990) adjusted determination coefficients for the different regressions all lie between 0.29 and 0.41, which is approximately the same results we find in our models, except for the 12-year period from

1947 to 1958, where he finds a value of 0.08. He concludes that an imposition of a transaction tax corresponding to a decline in volatility is not an obvious result.

Aliber, Chowdhry & Yan (2003) test four major currencies: British Pound (BP), Deutsche Mark (DM), Japanese Yen (JY) and Swiss Franc (SF), where three out of these four consist with our data. The time span they focus on is from January 1977 to December 1999 and the time series consists of monthly data. Compared to our data span, we look at the same period namely between the first week/month of 1990 to the last month of 1999 for GBP, CHF and JPY and the first week/month of 1999 to the last month of 1999 for EUR. It can be argued that the EUR and Deutsche Mark can be compared to some extent but based on the time span we find it more precise to compare their results with our result from only the first period. As we do not have a determination coefficient for all the first periods we have chosen to use the determination coefficient for the full period as a benchmark. This contains the first period and further we found that the R^2 values from our models were alike overall, therefore we find this a good estimate for comparison.

From regressing volatility on costs and lagged volatility, Aliber, Chowdhry & Yan (2003) find a determination coefficient for BP of 0.31 which compared to our GBP full period is almost identical. For JY and SF the values are somewhat lower. For JY the value is 0.19 compared to 0.43 and for SF they find a determination coefficient of 0.22 compared to our coefficient of 0.30. For DM they find a value of 0.21 compared to our EUR of 0.32, but this comparison should be taken with caution, due to the fact that the time periods and the currencies are different.

Looking at the beta values Aliber, Chowdhry & Yan (2003) find a positive relationship between costs and volatility in all cases, which is the same conclusion we find in first period of our models. Only in the case of DM the beta value of costs is not significant, but for BP, JY and SF the beta values of costs are all accepted at a significance level between 1 and 5 per cent. The beta values of lagged volatility are in all cases positive and accepted at a 1%-level of significance - the same results we find.

Lanne & Vesala (2006) test the Deutsche Mark (DM) and Japanese Yen (Yen) on a daily and intradaily basis. The time span is from 10-01-1992 to 09-30-1993, which like in Aliber, Chowdhry & Yan (2003) is a part of our first period. Therefore we will use the same approach when comparing results.

Lanne & Vesala (2006) regress realized variance on transaction costs, lagged variance and up to two seasonal dummies. In both currency cases of the daily data, they find high determination coefficients of 0.48 for DM and 0.43 for Yen. These are both higher when comparing with EUR's and JPY's full period determination coefficients, but we need to keep in mind the much shorter time span of Lanne & Vesala (2006), which is why we might see smaller fluctuations in their data. Additionally, they find a positive relationship between the cost and variance, where both cost's and lagged variance's beta values are significant at a 1 to 5 per cent level. They point out that heteroscedasticity is present, but they seem to solve the problem by calculating the t-statistics based on White's (1980) robust standard errors. In their intradaily regressions they find similar results and based on these they conclude, like Aliber, Chowdhry & Yan (2003), that an increase in transaction costs leads to increased volatility, which is the same result we find for our currencies' first period.

Ronan & Weaver (2001) research if return volatility is directly related to tick size in the American stock exchange (AMEX). The data span is from 04-01-1997 to 06-30-1997. Their main focus is the regime change in tick size that the AMEX stocks underwent in 05-07-1997 and whether this change affected volatility. They regress volatility on the number of trades and two tick size dummy-variables, which represent transaction costs. When comparing our results to theirs we focus on the dummy-variables, and use these as a comparison to our cost variable.

From their results they find that both dummy-variables have negative beta values and are significant at 5 and 1 per cent levels. The reduction in tick sizes gives way to a lower bid-ask spread and therefore lower transaction costs. The determination coefficient has a value of 0.26, which, compared to all of our models, is somewhat lower. They conclude that when following the adoption of smaller tick sizes the volatility is reduced, which is comparable to our results from the first period. It is troubling, however, that the decrease in volatility could be explained by an independent market wide volatility reduction (Hau (2006)).

Hau's (2006) data is based on the stocks in the CAC40 index (the 40 largest and most liquid French stocks), as well as hourly data. The time span is between January 1995 and December 1998 and his goal is to prove that higher transaction costs, in the form of a tick-dummy representing an exogenous transaction cost proxy, increase stock return volatility. In his model, he regresses volatility on the tick-dummy and lagged volatility. The results conclude that both the tick-dummy and the lagged volatility variables are positive and significant at a

1%-level, which is the same picture as in the former papers described in this section. He also tests for normality where he finds that the Wald χ^2 test is significant at a 1%-level, meaning that the data is normally distributed. He thereby concludes, based on the empirical evidence, that the effect of transaction costs on volatility is positive and significant, both statistically and economically and thus he believes that his evidence supports Friedman's (1953) point of view on noise traders.

After looking at the results from the papers of Mulherin (1990), Aliber, Chowdhry & Yan (2003), Lanne & Vesala (2006), Ronan & Weaver (2001) and Hau (2006) and comparing them to our own, we see that whether on the currency or stock market, the relationship between volatility and cost generally is positive and significant or negative and insignificant. This suggests that the theory of Friedman (1953) holds true, as Hau (2006) also stated. We further find that the determination coefficients found in the previous research take approximately the same values as found in our models.

7 Analysis

In this chapter we will interpret our results, which we described in section 6.4.1. The sign and sizes of the estimated coefficients are compared to the theories of Tobin, Keynes and Friedman described in sections 2.1, 2.2.1 and 2.2.2, respectively. Furthermore, we will evaluate the effect that different values of a Tobin tax will have on volatility. In section 7.2 we will relate our estimated coefficients to the conclusions of the game theoretical model from section 2.5 as to be able to comment on the market composition. In the last section of our analysis, section 7.3, we propose that the exchange rate can function as an instrument variable in a two-stage least squares model in helping to solve the endogeneity problem.

7.1 Interpretation of the results

From our linear regressions we can determine how a change in cost and VIX affects volatility. Meaning, that if for instance cost increases by 1 percentage point, volatility will change by the cost's estimated beta value in percentage points. In this section we will furthermore evaluate the estimated coefficients of our models and determine the effect on volatility when a tax of 0.5 and 0.005 per cent is proposed. Moreover, we will discuss if Eichengreen, Wyplosz & Tobin's (1995) proposed tax is a realistic suggestion, based on our empirical findings.

7.1.1 The relation between volatility and cost

Looking at Table 6 we find that in general the cost coefficients are either positive and significant or negative and insignificant, which consists with the findings in the previous research described in section 3.1 and 6.4.3. In this section we will focus mainly on those cases where we find significant cost coefficients, when discussing the signs and sizes of the cost coefficients and the effect these have on the volatility. In section 6.4.1.4 we concluded that the first period's cost coefficients gave the most reliable results and in section 6.4.2.2 and 6.4.2.3 we concluded that this period was fairly robust, wherefore we find it obvious to discuss an introduction of a transaction tax in the first period of the currencies.

For the first period, we found a general trend amongst the currencies. All the cost coefficients are positive, as mentioned in section 6.4.1.4, which is the contrary to the views described by Tobin and Keynes in sections 2.1 and 2.2.1. The largest beta coefficient is found in CHF, and from a 1 percentage point increase in cost, which was Tobin's (1978) original proposal of a transaction tax, we will see that the volatility increases by 7.65 percentage points. If we compare this to the standard deviation of the weekly volatility measure of 0.18%, which must be the expected movement in the volatility in the first period, we find that the increase in the volatility is over 42 times larger. For GBP and JPY the volatility in the exchange rate is affected by 5.47 and 5.28 percentage points, respectively, by a 1 percentage point increase in the cost coefficients. Here we find standard deviations of the weekly volatility measure to be 0.17 and 0.23 per cent, respectively, which corresponds to an increase in volatility more than 32 and 22 timers larger for GBP and JPY. For EUR we see that the impact on volatility when cost changes by 1 percentage point is 2.80 percentage points. The standard deviation for the weekly volatility measure for EUR is the lowest at 0.14%, which is probably due to the shorter time horizon, compared to the three other currencies. In this case we find that if Tobin's (1978) original tax proposal was introduced the increase in volatility would be 20 times larger compared to the standard deviation.

These results point in the direction of what Friedman (1953) describes, cf. section 2.2.2 – that volatility increases because there are less rational traders moving the price toward its fundamental value, which makes the price discovery process less efficient. As the increase in volatility is a result of a less liquid market, the volume or number of trades would have been interesting to investigate, but as described in section 4.5 these variables were not possible to include in our model.

In four other cases we find significant cost coefficients, which compared to the first periods' cost estimates all are differenced. In two of the cases, GBP's full period and CHF's third period, the Δcost coefficients are positive, with values of 1.89 and 2.77, respectively, meaning that if the difference in cost increases by 1 percentage point the Δvolatility will increase by the size of the associated Δcost coefficients in percentage points. Thereby, the higher the coefficient of cost the wider the span becomes between the volatility at time t and t+1, meaning a more volatile exchange rate. In the third period of GBP, we find the largest Δcost coefficient looking at all the four currencies with a negative value of -21.17. Even though we found in section 6.4.1.4 that this coefficient was not biased in any way, we know from section 6.4.2.2 and 6.4.2.3 that this period does not give very robust Δcost coefficients. Furthermore, we find that the Δcost coefficient of EUR's second period has a negative value of -2.47 at a 10% significance level, meaning that in these two cases the statement made by Tobin (1978) is valid, even though the coefficients are not that robust. In the full, second and third period we find that the majority of the results in Table 6 display negative, but insignificant cost coefficients.

7.1.2 Evaluating the size of a Tobin Tax

In section 2.1 we described that James Tobin proposed a transaction tax on currency exchange of 1% in 1978, which is the one we have referred to in section 7.1.1, when describing the signs and sizes of the cost coefficients and their impact on volatility. Almost two decades later in 1995 Tobin proposed a transaction tax of 0.5% in cooperation with Eichengreen and Wyplosz, cf. section 2.1. This was due to his concern that the international market had become too efficient in the matter that available information moved swiftly around the world, decreasing transaction costs and that the investment market had become too liquid (Tobin (1978)). Although we saw from the results in the spill-over model in section 6.4.2.2 that the relationship between JPY's volatility and the other countries' volatility in 50% of the cases was insignificant indicating that JPY's distance to the other countries might have an influence on how much the other markets are affected by it. This could imply that the market is not as efficient as Tobin originally thought, cf. section 2.1. In this section we will see the impact of an introduction of a 0.5% transaction tax on our empirical findings. We will discuss if this is a realistic suggestion and compare this to others' suggestions of a transaction tax.

As described in section 3.2.1 our cost measure is calculated as a percentage of the cost-free price, meaning that we can easily see from our estimated cost coefficients what will happen with volatility when a transaction tax is introduced. In this case we will focus on a tax of 0.5%, which then would have to be multiplied by the cost coefficient to see the impact this will have on volatility. As described in section 6.4.1.4 we believe that the first period of the currencies gives the most realistic cost coefficients, and therefore we will mainly focus on this period when testing the impact of the proposed transaction tax on the volatility. Furthermore, the first period of the currencies, except for EUR, contains the year 1995 where the tax of 0.5% was proposed. In Table 9 we have displayed the cost coefficients from the first period of the currencies, the Tobin Tax of 0.5%, how much it will affect volatility, the average volatility and costs over the first period, the ratio of the effect of volatility to average volatility and the ratio of the proposed tax to the average costs, where both ratios are calculated in percentage.

Currency	Cost coefficients (β_{Cost})	Tobin tax (C) (%)	Affect on volatility (β_{Cost}*C)	Average volatility (%)	Average costs (%)	(β_{Cost}*C)/ Average volatility (%)	C/ Average costs (%)
GBP	5.47	0.5	2.7367	0.4493	0.0279	609.0668	1789.6804
CHF	7.65	0.5	3.8238	0.6229	0.0338	613.8901	1478.2510
JPY	5.28	0.5	2.6403	0.5517	0.0307	478.5970	1630.8311
EUR	2.80	0.5	1.4005	0.5510	0.0286	254.1937	1746.1654

Table 9 - The impact of a 0.5% Tobin Tax on the currencies first period.

From Table 9 we can see that an introduction of a transaction tax of 0.5% will increase the volatility of the four currencies by 1.40 to 3.84 percentage points. From the seventh column we can see how large the effect on volatility is compared to the average volatility in the first period. In the case of GBP we find that the increase in volatility due to the transaction tax is over 6 times larger than the average volatility over the first period. The same applies when looking at CHF, JPY and EUR, although we find smaller ratios for JPY and EUR of 4.79 and 2.54, respectively. When comparing these effects to the associated standard deviations of the weekly volatility measures for the first period found in section 7.1.1, we see that the increase in volatility is on average over 14 times[1] larger than the average standard deviation for the four currencies. As the effect on volatility is positive and the results thereby work contrary to what Tobin (1972) stated, we conclude that if the tax of 0.5% was introduced in 1995 the volatility would have increased to an extreme degree.

In the case of two of our models we found the cost coefficients to be negative and significant within a 10%-level, which corresponded to the proposal made by Tobin in 1972. This was present in GBP's third period and EUR's second period. Even though we do not find these estimated coefficients to be very robust, and thereby not very reliable, we still find it interesting to see what will happen if a transaction tax is introduced. In these two cases both volatility and the cost variables had a unit root, which is why when we interpret the results we look at how much a percentage increase in Δcost will affect Δvolatility in percentage points. In the case of GBP we found the cost coefficient to be -21.17, meaning that if introducing a tax of 0.5% Δvolatility will decrease by 10.58 percentage points, 36 times the standard deviation of the third period, and in the case of EUR we would see a decrease in Δvolatility of 1.24 percentage points, 11 times the standard deviation of the weekly volatility in the second period. If these two negative beta values were to be reliable estimates the introduction of the transaction tax would have worked very well and in line with Tobin's and Keynes' way of thinking.

[1] $\frac{1}{n}\sum(\beta_{Cost}{*}C)/\frac{1}{n}\sum Std_{Weekly\ volatility\ measure}$

After looking at the impact of a transaction tax of 0.5% on volatility, we can evaluate based on our empirical research whether this is a realistic suggestion. We found that the effect the proposed tax would have on volatility in the first period of the currencies was between 2.54 and 6.14 times the average volatility measure, which we believe are extreme increases and definitely the result of an unrealistically high transaction tax. With this in mind, we have compared the average cost measure, with the proposed transaction tax, which can be seen from Table 9 in the last column. Here we find that the proposed transaction tax of 0.5% is between 14.78 and 17.90 times larger than the average cost for the four different currencies over the first period. It seems unrealistic to introduce a tax that, compared to the average costs, is up to 1,690% larger. Rodney Schmidt has also put this scepticism forth in his report for the North-South institute in 2007. Here Schmidt argues that a transaction tax must not exceed the observed level of the transaction costs, which is exactly what we see when comparing the proposed tax with our average costs.

Some might argue that some of our currencies' first period continues until the middle of 2000 or 2002, and that changes in the underlying market could have affected a fall in transaction costs from 1995 until the end of each period, thus a reason why the tax of 0.5% could be too large. However, comparing the average of costs of 1995 with the average transaction costs of the last year in the currencies' first period we do not see a significant change, which is why we conclude based on this, our empirical evidence and the statement made by Schmidt (2007) that the proposed tax simply was estimated too high.

In November 2007 APPG for Debt, Aid and Trade[2] published a report concerning a sterling stamp duty (SSD) on the United Kingdom's FX market (APPG (2007)). In the report they propose a transaction tax of 0.005%, which is 200 times smaller than James Tobin's original suggestion of a 1% transaction tax. Schmidt (2007) proposes a tax of the same size in the North-South institute report from 2007, as a reasonable estimate of a transaction tax. Comparing this to the average cost of the four currencies we see that the tax does not exceed these, which corresponds to the statement of Schmidt (2007) described earlier in this section. Applying this proposed tax to the four currencies' first period we find the results displayed in Table 10.

[2] The All Party Parliamentary Group (APPG) is a grouping in the UK Parliament that is composed of politicians from all political parties discussing topics reflecting the parliamentarians concerns.

Currency	Cost coefficients (β_{Cost})	Tobin tax (C) (%)	Affect on volatility (β_{Cost}*C)	Average volatility (%)	Average costs (%)	(β_{Cost}*C)/ Average volatility (%)	C/ Average costs (%)
GBP	5.47	0.005	0.0274	0.4493	0.0279	6.0907	17.8968
CHF	7.65	0.005	0.0382	0.6229	0.0338	6.1389	14.7825
JPY	5.28	0.005	0.0264	0.5517	0.0307	4.7860	16.3083
EUR	2.80	0.005	0.0140	0.5510	0.0286	2.5419	17.4617

Table 10 - The impact of a 0.005% Transaction Tax on the four currencies first period.

From Table 10 we can see that the effect on volatility compared to Table 9, as well as the last two columns decreases by one-hundredth and now gives much more realistic results. The proposed tax of 0.005% now represents between 14.8 and 17.9 per cent of the four currencies' average costs and we only see an increase in volatility between 3 and 6 per cent of the average volatility.

To give a graphical view of how the two different taxes just described would have affected the exchange rate in the first period, we have done a Monte Carlo simulation. First, we have found the average exchange rate ($\bar{S}$) of the weekly observations and its average variance ($\bar{V}$) based on the weekly volatility measures. The daily exchange rate is defined as:

$$\mathrm{S} = \frac{S_a + S_b}{2} \qquad (68)$$

where S_a and S_b are the ask and bid price, respectively. It is straightforward to show that the exchange rate can be expressed as equation (68) when isolating S in equations (41) and (42).

Second, we have made 100 different Monte Carlo simulations with random numbers equivalent to the number of observations using a normal distribution with $\mu = \bar{S}$ and $\sigma^2 = \bar{V}$. Third, we have taken the average of the 100 simulations for each random number and plotted it against time. This can be seen for GBP in Figure 8, denoted by simulation 1, which represents the simulated variation in the real exchange rate. Graphical view for CHF, JPY and EUR can be found in Appendix 11.

Simulation 2 and simulation 2' from Figure 8 represent the scenarios where a tax of 0.005 and 0.5 per cent, respectively, is introduced. In these two cases we have added the effect on volatility (β_{Cost}*C), displayed in Table 9 and Table 10, to the variance of the exchange rate, i.e. $\sigma^2 = \bar{V} + \beta_{\mathrm{Cost}} * \mathrm{C}$. By doing this and comparing the three simulations we can see how the variation in the exchange rate differs with no tax, a low tax and a high tax.

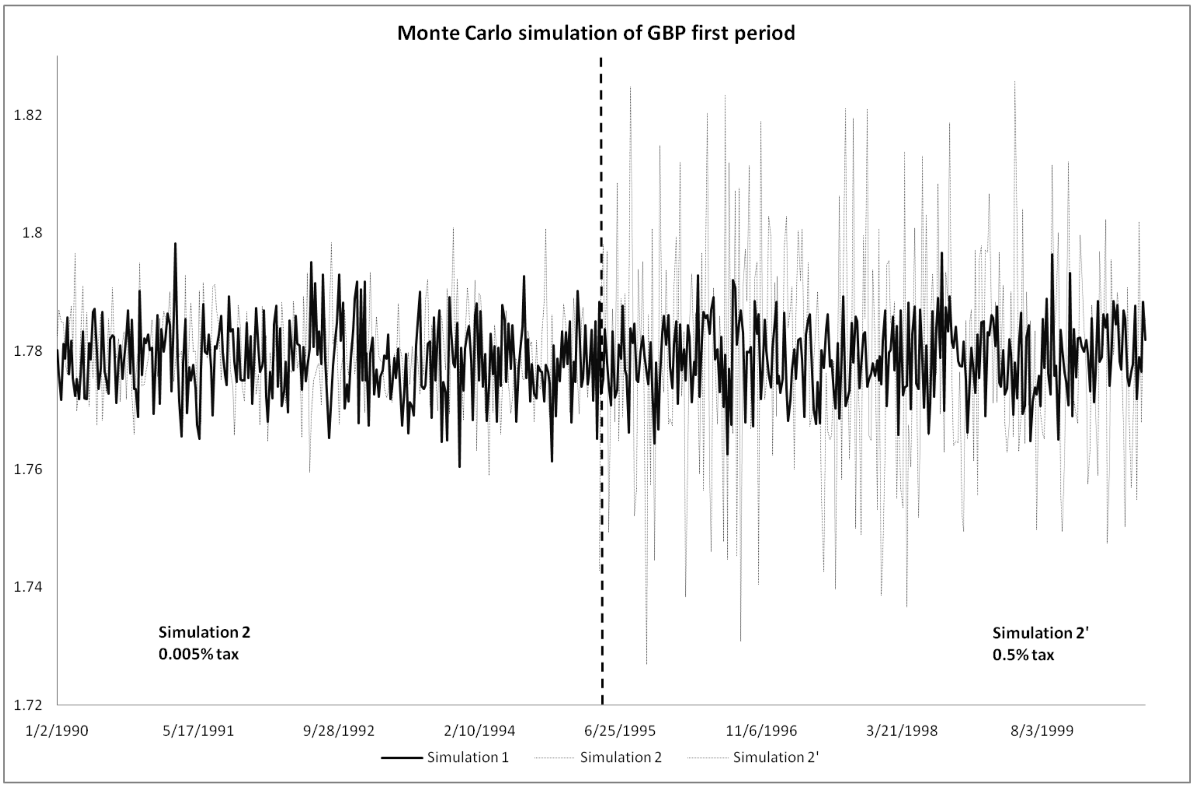

Figure 8 - Monte Carlo simulation of GBP's first period. $\boldsymbol{S_t = \varepsilon_t \sim N(\mu, \sigma^2)}$, where μ = $\boldsymbol{\bar{S}}$ and $\boldsymbol{\sigma^2 = \bar{V} + \beta_{Cost} * C}$.. Simulation 1 represents the real exchange rate wherefore C = 0. For simulation 2 and 2' C takes the value of 0.005% and 0.5%, respectively.

For a better graphical view we have divided Figure 8 into two parts. In the first part, simulation 1 can be compared to simulation 2, which is the scenario where a tax of 0.005% is introduced. In the second part, simulation 1 is compared to simulation 2' – the scenario where a tax of 0.5% is introduced.

It is clear that simulation 2' is much more volatile than simulation 1. This was expected as we found the effect on volatility from a tax introduction of 0.5% to be over 6 times larger than the average volatility for GBP, and the reason why we concluded that the tax was estimated too high. Looking at the simulation 2 compared to simulation 1, we see that the two simulations move close to each other, which indicates that the proposed tax of 0.005% is a much more realistic suggestion than the one proposed by Tobin (1978).

7.1.3 Lagged volatility as a benchmark

In the spill-over models from section 6.4.2.2 we saw that when including foreign volatility variables they became highly significant and VIX, which should resemble the market benchmark, became insignificant. In our models, we have used VIX as the benchmark for the market in general, and in most cases, especially the full, first and second period, VIX is significant at a 1%-level. Although, when looking at our findings just described we believe that an average of the foreign volatility measures would create a better benchmark for the market than VIX does, meaning that this might explain the variation in the volatility better, which could result in a higher determination coefficient in our models as seen in the spill-over models.

7.2 Evaluating the assumptions behind the Tobin Tax

In section 2.1 we described certain assumptions underlying the Tobin tax, which at least had to be upheld for the tax to work. These assumptions were, first, that short term speculation is destabilizing; second, that an increase in transaction costs, such as an added tax, discourages this short term speculation and last, that the added costs will make noise traders base their investments on macroeconomic fundamentals.

In the previous section, section 7.1, we described that an added tax on the transaction costs would drive up the exchange rate volatility in the currencies' first period, which henceforth is the period that we will refer to in this section when commenting on our empirical results. The observed increase in volatility is the opposite result of what Tobin and Keynes describe. It also suggests that the assumptions behind the Tobin tax are not observed in the market. In the following section we will look at what the increase in volatility due to a Tobin tax says about the composition of the market. We will do so by comparing our results with the conclusions of the noise trader model proposed by De Long *et al.* (1990) and Song & Zhang (2005), as we described in section 2.5.

Song & Zhang (2005) concludes that the effect of introducing a transaction tax on volatility warrants both Tobin's view and the contrary one, depending on how dominant the size of noise traders is in the market. According to Song & Zhang (2005), the increase in volatility suggests that our market consists of a larger part of noise traders than informed traders.

We realize that we cannot observe the share of noise traders in the markets we are researching, wherefore we cannot say if our markets correspond to either the equilibrium L or H from Figure 3, although we believe that our results are well in line with the conclusions of Song & Zhang (2005), who themselves would expect to see a positive relationship between a

Tobin tax and volatility in large and often traded exchange rates such as the ones we are analysing in our thesis work.

The increase in volatility also suggests that a larger part of informed traders leave the market compared to noise traders, which increases the share of noise traders to informed traders. This conclusion by Song & Zhang (2005) is based on the assumption that both types of traders' investment horizons are equally long. Although, as we saw in equation (1) in section 2.1, if informed traders had possessed longer investment horizons than noise traders they would have been affected less by the increase in transaction costs. Therefore our results suggest that informed traders' investment horizons are not as long as Keynes might have expected, and because of noise traders' misperception of the prices, informed traders might be trading more often to bet against the noise traders as we discussed in section 2.3.4 (Eichengreen, Tobin & Wyplosz (1995)).

De Long *et al.* (1990) showed that when the share of noise traders approach 0 in the presence of both noise trader risk and fundamental risk, they will always earn higher expected returns than the informed trader, and dominate the market in the long run when the imitation parameter (ζ) is close to 0; indicating that no real imitation of beliefs from period to period exists. If, on the other hand, we assume that the share of noise traders (μ) is increasing it is clear that the expected relative return from equation (29) becomes unstable as μ approaches 1 and equation (29) becomes:

$$E(\Delta \mathrm{R}_{n-i}) = \rho^* - \frac{(\rho^*)^2 + \sigma_\rho^2}{(2\gamma)\left[\sigma_\varepsilon^2 + \frac{\sigma_\rho^2}{(1+r)^2}\right]} \tag{69}$$

It is clear that the expected misperception of the noise traders (ρ*) both influence the expected relative return in a positive and a negative way. It then becomes a question of the size of the two types of risk to determine if the noise traders earn a higher return than the informed traders. If we assume that imitation of beliefs does take place in the market, so traders switch between strategies based on their past performance, we would not expect to see a stable equilibrium of noise traders and informed traders over time, due to the noise traders' uncertain misperception of prices. This instability also points to the fact that our results might have been different had we looked at other time-periods. We do, for example, find an indication of a negative relationship between cost and volatility in the second and third period, as seen in section 6.4.1.4, although it is not significant. Our results correspond with the results of previous research, as we saw in section 3.1, which was also done on major currencies or stock indices. The reason for the increase in volatility might, to a higher degree, be due to the size of the currencies as Song & Zhang (2005) note, than the specific development of the exchange rate during the time period.

Furthermore, our results suggest that the traders, be it noise traders or informed traders, try to maximize their return instead of their utility. In section 2.5.5, we explained how maximizing utility would diminish the share of noise traders in the market, as new traders would mimic the informed traders because of the higher risk associated with noise traders, which results in lower utility. Since the increase in volatility indicates the presence of noise traders in the market it is more likely that all traders base their investment decisions mainly on the basis of past, realized returns and considers the associated risk to a lesser degree as De Long *et al.* (1990) also point out.

If we turn our focus back to the three assumptions of a Tobin tax, we see that the first assumption of short-term speculation being destabilizing is partly true. We have seen indications of noise traders' and informed traders' investment horizons not being too different from each other, which implies that short-term speculation is both stabilizing and destabilizing dependent on the trader. The second assumption: that a Tobin tax will discourage short-term speculation is clear when looking at both equation (1), and the modified equations (34) and (35), although it follows from what we know about assumption 1; namely that the introduction of a Tobin tax will most likely discourage both destabilizing and stabilizing speculation. From section 2.5 we found that the two types of traders' investment horizons are almost identical and that noise traders' demand for the risky asset is less affected by a Tobin tax than that of the informed traders. This is due to noise traders' false misperception of the assets' price and a possibility for higher returns; wherefore they will not necessarily change their investment strategy to be based on the macroeconomic fundamentals of the market. This also suggests that the third assumption does not hold up in practice.

Our results, as well as the results of many of the papers mentioned in section 3.1, indicate that Keynes and Tobin's interpretation of the market structure, as explained in sections 2.1 and 2.2.1, does not correspond with what we see from the market's reaction to an increase in transaction costs. We realize that what our results suggest about the composition of the market, and the investment horizons of noise traders and informed traders, as formulated by

De Long *et al.* (1990) and Song & Zhang (2005), does not resemble the empirical observations of section 2.3. It is difficult to say why this is, but it suggests that it is indeed difficult to measure if traders are either noise traders or informed traders. The difficulty lies in that the fundamental price is cumbersome to observe and therefore it becomes difficult to determine if the actions of the market are essentially following the fundamental development of the currency or somehow trading the currency to become either over- or undervalued depending on the trend in the market, as explained in section 2.3.4.

In our thesis (work) we have first looked at how the market in reality reacts to the introduction of a Tobin tax in section 7.1, and afterwards looked at the conditions that need to be in place for this behaviour to occur. We realize that of course our deductions about the behaviour of the market participants rely on the assumptions of the game theoretical model by De Long *et al.* (1990). One of the most questionable assumptions made by De Long *et al.* (1990) is that the investment horizon of the noise trader and informed trader is the same, as this goes against the beliefs of Keynes who argued that informed traders had long-term investment horizons contrary to noise traders. Fortunately, our empirical results suggest that informed traders' investment horizons actually are close to the investment horizons of noise traders, since our empirical results are line with the game theoretical model.

7.3 Endogeneity

During this section we will deal with the endogeneity problem. We propose that the exchange rate can be used as an instrument variable to determine the true causational relationship between cost and volatility. To test this we set up a two-stage least squares model to determine if the predicted values of cost by the exchange rate can help to predict volatility.

7.3.1 A possible solution to the endogeneity problem

In section 3.4 we discussed the theoretical endogeneity problem between cost and volatility, and explained how the fact that the percentage transactions' cost affects volatility, which again affects the spread creating a self-perpetuating relationship between the two variables. The problem when using endogenous variables is that we cannot be absolutely sure of the statistical causality between cost and volatility which we have discussed in sections 7.1 and 7.2. In the following section, we propose a solution to this problem, which contrary to the proposed solutions by previous research discussed in section 3.4, rely on an instrumental variable that is only correlated with the cost measure and not volatility.

7.3.2 The instrument

Based on a few assumptions, we believe that the exchange rate can work as an instrument in predicting the true relationship between cost and volatility. This realization is based on two assumptions: First, that volatility does not affect the level in the exchange rate, and second that the level in the exchange rate affects the percentage transaction costs.

When the volatility in an asset changes it changes the excess return that a trader demands from his investment but not the level of the asset's price. An excess return of any size can be realized independently from the price level and therefore the currency's price is unaffected by its volatility to any significant degree. This assumption holds true for at least our measure of volatility, cf. equation (54), where we assume volatility to be symmetrical around the true exchange rate (S). In contrast, we do not believe that the exchange rate affects volatility directly. One might assume that the dealer knows how large a percentage transactions cost a trader can tolerate, and therefore widens the spread as the exchange rate appreciates. It is, although, highly unlikely considering the competitive nature of the currency market, wherefore we believe that the bid-ask spread always resembles the minimum required costs. This assumption indicates that there should be no correlation between our instrument and the dependent variable: volatility.

Our second assumption ensures that our instrument has a predictive power over our endogenous variable: cost. Our cost measure is, as explained in section 3.2.1, the spread's percentage share of the exchange rate, wherefore it is partly explained by the exchange rate. This entails that, ceteris paribus, a depreciation or appreciation in the currency will be directly observable in the percentage transactions cost.

We explained in section 3.3 how volatility affects the exchange rate's spread, which is partly what causes the endogeneity problem: The more stable the spread is, the more predictive power we expect to be able to assign our instrument. Therefore, we expect to see the exchange rate for GBP and EUR acting as the strongest instrument, since these currencies' spreads are much more stable than the those of CHF and JPY. We can see this by comparing Figure 6 for GBP with the same figures for the rest of the currencies in Appendix 1.

We believe that our instrument can only be justified to be applicable to the first period where neither cost nor volatility is differenced. Even though volatility should not be affected by the price level of the currency and vice versa, we believe that volatility will have a positive effect on the difference in the exchange rate, as an increase in the required excess return, from an increase in volatility, will affect the difference in the exchange rate proportional to its level. We base this belief on the fact that the differenced exchange rate and volatility are correlated, which violates the assumption of the 2SLS model.

To get a preliminary idea of the relationship between cost, volatility and the exchange rate we have computed their correlations for each currency.

	GBP		CHF	
	Volatility	Exchange rate	Volatility	Exchange rate
Cost	8.93%	-38.19%	26.80%	-4.35%
Volatility		-8.86%		7.90%
	JPY		**EUR**	
	Volatility	Exchange rate	Volatility	Exchange rate
Cost	13.33%	-7.76%	16.99%	-81.70%
Volatility		15.06%		-14.40%

Table 11 – Correlations between the variables for each currency.

As we can see from Table 11 the correlation between the exchange rate and cost only seems significant compared to the correlation between the exchange rate and volatility for GBP and EUR. In the cases of CHF and JPY, the exchange rate's numerical correlation with volatility exceeds its correlation with costs. This is well in line with our presumptions about our data from earlier in this section. We will include all the currencies in our 2SLS model to see if there are any differences in their results.

7.3.3 Two-stage least squares

In order to test if our instrument helps predict the true causational relationship between cost and volatility, denoted by C and V, respectively, we will run a two-stage least squares regression for each currency. Our instrument, the exchange rate (S) is defined by equation (68).

From section 6.2 we have our original OLS regression model:

$$\mathrm{V} = \beta \mathrm{X} + \mathrm{u} \tag{70}$$

where V and u are now two 4x1 vectors defined as:

$$\mathrm{V} = \begin{bmatrix} V_{CHF,first} \\ V_{GBP,first} \\ V_{JPY,first} \\ V_{EUR,first} \end{bmatrix} \tag{71}$$

$$\mathrm{u} = \begin{bmatrix} u_{CHF,first} \\ u_{GBP,first} \\ u_{JPY,first} \\ u_{EUR,first} \end{bmatrix} \tag{72}$$

X is a 4x6 matrix defined as:

$$\mathrm{X} = \begin{bmatrix} 0 & C_{CHF,first} & V_{CHF,first}^{t-1} & V_{CHF,first}^{t-2} & 0 & \Delta VIX_{CHF,first} \\ 0 & C_{GBP,first} & V_{GBP,first}^{t-1} & V_{GBP,first}^{t-2} & 0 & \Delta VIX_{GBP,first} \\ 0 & C_{JPT,first} & V_{JPY,first}^{t-1} & V_{JPY,first}^{t-2} & 0 & \Delta VIX_{JPY,first} \\ 0 & C_{EUR,first} & V_{EUR,first}^{t-1} & 0 & V_{EUR,first}^{t-3} & \Delta VIX_{EUR,first} \end{bmatrix} \tag{73}$$

and β is a 6x1 vector of β's from β_0 to β_5. Equation (70) is the second stage of our 2SLS regression, where we assume that cost (C) is endogenous. On the basis of the discussion about our instrument in section 7.3.2, we can express the first stage of the 2SLS regression as:

$$\mathrm{C} = \delta \mathrm{Z} + \varepsilon \tag{74}$$

where C and ε are two 4x1 vectors defined as:

$$\mathrm{C} = \begin{bmatrix} C_{CHF,first} \\ C_{GBP,first} \\ C_{JPY,first} \\ C_{EUR,first} \end{bmatrix} \tag{75}$$

$$\varepsilon = \begin{bmatrix} \varepsilon_{CHF,first} \\ \varepsilon_{GBP,first} \\ \varepsilon_{JPY,first} \\ \varepsilon_{EUR,first} \end{bmatrix} \tag{76}$$

Z is a 4x6 matrix defined as:

$$\mathrm{Z} = \begin{bmatrix} 0 & V_{CHF,first}^{t-1} & V_{CHF,first}^{t-2} & 0 & \Delta VIX_{CHF,first} & S_{CHF,first} \\ 0 & V_{GBP,first}^{t-1} & V_{GBP,first}^{t-2} & 0 & \Delta VIX_{GBP,first} & S_{GBP,first} \\ 0 & V_{JPY,first}^{t-1} & V_{JPY,first}^{t-2} & 0 & \Delta VIX_{JPY,first} & S_{JPY,first} \\ 0 & V_{EUR,first}^{t-1} & 0 & V_{EUR,first}^{t-3} & \Delta VIX_{EUR,first} & S_{EUR,first} \end{bmatrix} \tag{77}$$

and δ is a 6x1 vector of δ's from δ_0 to δ_5. First of all we need to verify if δ_5 is significant to help explaining Z in equation (74), which corresponds to controlling for the first assumption underlying an instrumental variable as explained in section 3.4. If this is not the case we have no instrument to work with. The second assumption of exogeneity cannot be tested when we have exact identification, i.e. as many instrumental variables as endogenous variables, but we believe that our theoretical reasoning for its validity in section 7.3.2 is sufficient.

7.3.4 Evaluation of our instrument

When we assessed our first results of the models represented by equation (70) and (74) we found that they exhibited quite a large degree of heteroscedasticity. The reason for this being that the level of the exchange rate was not stationary over time. We decided to solve this problem to some degree by adding a time variable to capture the trend in the exchange rate.

The results of our first and second stage regressions are reported in Table 12. From the first stage we can see that we find the delta coefficient relating to the exchange rate to be significant at a 1%-level for all currencies – confirming that our instrument helps to explain the value of cost. Furthermore, we can see that our time variable is significant and that it is primarily the first lag of volatility that has an effect on cost.

When estimating our second stage model we can see that all of our estimated beta coefficients for our predicted costs from the first stage model are positive and significantly different from zero. These results resemble the results from our model represented by equation (62). The results from our 2SLS model gives credence to the belief that the causal relationship between cost and volatility found in our original model actually resembles the variables true causal relationship.

	First stage			
Currency	*GBP*	*CHF*	*JPY*	*EUR*
Intercept	-	-	-	-
	-	-	-	-
Volatility lag1	0.0079	0.0044	0.0028	0.0070
	*(5.63***)*	*(3.96***)*	*(2.21**)*	*(2.20**)*
Volatility lag2	0.0073	0.0005	0.0016	-
	*(5.26***)*	*(0.42)*	*(1.23)*	-
Volatility lag3	-	-	-	0.0053
	-	-	-	*(1.65*)*
Delta Vix	0.0000	0.0000	0.0000	0.0000
	(0.77)	*(0.18)*	*(0.37)*	*(0.40)*
Exchange rate	0.0001	0.0005	0.0345	0.0002
	*(24.96***)*	*(47.40***)*	*(38.72***)*	*(7.96***)*
Time variable	0.0000	0.0000	0.0000	0.0000
	*(3.92***)*	*(-6.23***)*	*(-4.91***)*	*(8.25***)*
R^2	-	-	-	-
N	551	271	566	195
Autocorrelation	0.4335	0.0601	0.2671	0.3047
Breuch-Pagan	0.0044	0.5217	0.1810	0.6540
Condition index	9.6288	8.6600	7.4146	10.4696
Anderson-Darling	1.1843	2.3297	1.2756	3.6570
Skewness	0.2629	0.1411	0.6011	1.4642
Kurtosis	0.7545	3.2681	1.5328	4.0975

	Second stage			
Currency	*GBP*	*CHF*	*JPY*	*EUR*
Intercept	-	-	-	-
	-	-	-	-
Predicted Cost	9.5472	6.3634	5.1051	10.4988
	*(8.14***)*	*(5.64***)*	*(6.18***)*	*(4.31***)*
Volatility lag1	0.3430	0.4305	0.4506	0.4261
	*(7.50***)*	*(6.96***)*	*(10.89***)*	*(5.91***)*
Volatility lag2	0.1650	0.1782	0.2125	-
	*(3.65***)*	*(2.99***)*	*(5.17***)*	-
Volatility lag3	-	-	-	0.1222
	-	-	-	*(1.75*)*
Delta Vix	0.0001	0.0001	0.0002	0.0001
	*(3.35***)*	*(1.75*)*	*(5.07***)*	*(2.16**)*
Time variable	0.0000	0.0000	0.0000	0.0000
	*(-4.08***)*	*(-0.63)*	*(2.14**)*	*(-2.46**)*
R^2	-	-	-	-
N	551	271	566	195
Autocorrelation	0.4494	0.5967	0.5133	0.6842
Breuch-Pagan	<.0001	0.0679	0.0006	0.0824
Condition index	13.7558	8.8799	7.4188	18.5260
Anderson-Darling	8.9676	4.3756	9.7346	1.8423
Skewness	1.0905	1.5747	1.5765	0.8130
Kurtosis	1.4851	5.2040	5.2071	0.8701

Table 12 – Results from the first and second stage of our two-stage least squares models for the first period for all four currencies. T-values are shown in parentheses and the test statistics are displayed beneath each model. One star means that the null hypothesis is rejected at the 10%-level, two stars mean a rejection at the 5%-level and three stars mean a rejection at the 1%-level

7.3.5 Critique of our 2SLS model

The 2SLS model we have used in this section is by no means perfect. It has several drawbacks that need to be brought to light, but none that affect either the first or second stage's estimated coefficients. Therefore, the model still tells us something about the signs of the estimated coefficients for both our instrumental variable in stage 1 and our predicted costs in stage 2. The first noticeable problem is the high autocorrelation in the residuals of both models, which affects our estimated coefficients' t-values. The second problem is less obvious but occurs because we have not been able to compute the corrected error terms from the first stage, which results in larger errors in the second stage. One could also argue that the fact that GBP and EUR's bid-ask spread moves in bands and is therefore very stable, means that the exchange rate causes the variation in our cost measure, wherefore we, to a larger degree, originally tried to predict the volatility via the variation in the exchange rate. We do, although, find similar results in the 2SLS model for GBP and EUR as for CHF and JPY, which have spreads with more variation, wherefore we believe that this is most likely not an issue.

8 Discussion

This chapter aims to discuss the more practical applications regarding the introduction of a Tobin tax. We have divided this chapter into two sections. The first section, section 8.1, focuses on the pros and cons of the Tobin tax from a public point of view and evaluates if such tax will improve welfare. Section 8.2 discusses the Tobin tax's ability to postpone financial crises in the framework of three generations' models of speculative attacks.

8.1 Public welfare benefits

As described in the introduction, chapter 1, we found that after Tobin's (1972) suggestion of a transaction tax the proposal has occasionally been discussed, especially in the context of financial crises in the current era, but it may have gathered its greatest momentum with the outcome of the recent global financial crisis. Over the years institutions such as BIS, G20, AEA, IMF, EU, APEC and more have discussed both the pros and the cons of the Tobin tax. Topics discussed recently in context to the global financial crisis that started in 2007 are amongst others: the impact on employment, the purpose of the tax, implementation of the tax, the size of the tax and thereby if it is feasible and how much money such tax could rise and last, what should the money generated from the tax be spent on?

Concerns about whether the tax would hurt employment have been expressed by amongst others: APEC (2010) and Schwabish (2005). From section 2.2.2, we have described that an introduction of a transaction tax would lead to a decrease in liquidity, as a result of a reduction in demand i.e. volume. Schwabish (2005) believes that the decrease in demand will lead to unemployment in the financial sector, which then would spread to other sectors thereby affecting the economy as a whole. Contrary to this, you could argue that the tax revenue could be used to create new jobs, but the question is then if the two effects cancel each other out or if one is better than the other? Furthermore, a report from the North-South Institute (Schmidt (2007)) describes that if the tax is small enough, say 0.005% as it suggests, the impact on volume will only lead to a 14% decrease and over time this would not disrupt neither the exchange rate behaviour nor market liquidity. This leads to the next topic discussed – the purpose of the tax. Through this thesis work we have focused on the Tobin tax as an instrument to reduce speculation and thereby an aid to stabilize the exchange rate, but Schmidt (2007) and APPG (2007) do not use the tax as a stabilizing instrument but rather as an opportunity to raise money. APPG (2007) writes:

"*Since the SSD's* [their currency transaction tax] *aim is to raise revenue the intention is for it to skim the market causing the least possible adverse impact to it. The motivation for Tobin's tax was exactly the opposite – to impede the activities of currency traders.*"

With this in mind, this could be the answer to the question of why Tobin suggested, what we believe was an unrealistic estimate of a transaction tax, described in section 7.1.2. If the tax worked intentionally and did not increase volatility, as we have seen from our empirical results and in previous research, cf. Table 9 and section 6.4.3, this would perhaps have worked, as Tobin's intention was that the tax should have a visible effect on the exchange rates. Although we still find that a transaction tax that is up to 17 times larger than the average transaction costs unrealistic. We believe that APPG's (2007), amongst others', suggestions to use the tax purely as a means for revenue might reflect many of the results from previous research that contradicts Tobin's original hypothesis. If the revenue motivated proponents of a transaction tax to base their beliefs on the contradicting results, it is important to introduce a tax so small that it does not affect the market behavior in any significant way, as this would increase volatility, according to previous research.

The next topic is the implementation of the tax. Here financial sectors such as BBA and LIBA (APPG (2007)) have expressed their concern if the tax is only introduced to one market. They believe that this would drive foreign exchange business to other markets with no tax, meaning that the tax would have to be implemented worldwide on all markets, which was also assumed by Eichengreen, Tobin and Wyplosz (1995). An example of their concern could be seen in the mid 1980's when a tax was introduced on the Swedish stock market, described by Umlauf (1993). As a result of the imposed tax, domestic investors started to trade off-shore and foreign investors started trading through London or New York, which is exactly BBA and LIBA's concern. Consequently the tax was abandoned in the early 1990's. In contrast to this, other countries such as the UK, United States, France, Belgium and Switzerland have had more success from introducing a tax on financial transactions, with no apparent negative impact on the market, even though it was not introduced to all markets. A good example is the UK. In 1986 the country introduced the so-called Stamp Duty, which currently taxes transactions regarding stocks, bonds and derivatives. Despite of this imposed tax the UK continues to be one of the leading financial centers and the London Stock Exchange is one of the premier exchanges in the world (APPG (2007)). Based on this example, an introduction of the transaction tax could work, even though it was not introduced to all markets.

The topic concerning the size of the tax depends on the motivation for the tax, as we described earlier in this section – is the target of the tax stability or revenue? Schmidt (2007) gives an estimate of what a tax of 0.005%, the same size of tax we used in Table 10, could generate in revenue if introduced on all sterling foreign transactions worldwide. He points out that the sterling amounts to 15% of the world trade and is the 4th most traded currency after US Dollar, EUR and JPY. The estimate he finds amounts to USD 4.98 billion and is based on the tax of 0.005%, which purpose is to avoid changing fundamental market behavior – the opposite of Tobin's (1972) intention.

The question of what the tax revenue should be used for sees many suggestions. APPG (2007) suggest that the money should be spent on the third world countries. In 2009 the British Prime Minister Gordon Brown and the president of France Nicolas Sarkozy suggested that revenues from the Tobin tax could be spent on the world's fight against climate change, especially in developing countries, which was also discussed at COP15 (FT (2009)). However in November 2009 Prime Minister Gordon Brown and Chancellor Alistair Darling announced that the European Union leaders had agreed to that the IMF should now consider the possibility of introducing an international transaction tax, which would be used to create a fund for bank bailouts and work as insurance for the global taxpayers against future banking crises (BBC (2009)).

BBC (2008) describes the bank rescue packages for the UK and this amount to be approximately USD 850 billion. If we assume that the international transaction tax's only purpose is to generate revenue and therefore should not disrupt the fundamental market behavior, we can assume that the rate will lie around Schmidt's (2007) proposed rate of 0.005%. By doing this, we can compare the USD 850 billion with the estimated yearly revenue of USD 4.98 billion, which would correspond to 0.6% of the total cost. Even if we assume that crises happen every 10 years the amount of money the tax would bring in still is far from the cost of a financial crisis, wherefore we believe that a currency transaction tax would have to be larger than 0.005% and thereby possibly disrupt to the market.

8.2 Postponement of the crisis

Besides the proposed effect of lowering exchange rate volatility, it has been suggested that a Tobin tax could dampen the effect of financial crises. Since the time after the breakdown of the Bretton Woods system, three generations of theories on speculative attacks on currencies have been described. In the following section we will quickly recap these theories as well as discuss how a Tobin tax might help to prevent or postpone these attacks.

The first-generation theories deal with the breakdown of a fixed exchange rate, such as the Bretton Woods system (Krugman (1979)). It is assumed that traders are rational and that a breakdown of the system can only occur because of inconsistent policies with the fixed exchange rate. The fixed exchange rate requires the supply of money to be constant wherefore the time of the collapse can be predicted as a country's currency reserve is depleted and replaced by domestic credit. When traders see that the exchange rate can no longer keep its fixed rate and therefore becomes overvalued relative to its true exchange rate, the traders leave the market resulting in an outflow of foreign money and a depreciation of the domestic currency.

The second-generation theories tone down the assumption that the central bank will not do away with the fixed regime, adding an uncertainty to the possible time of the collapse. As the probability of a collapse increases, traders will increasingly speculate against the currency in the belief that others will do the same (Obstfeld (1986)). These were the sorts of attacks against the European Exchange Rate Mechanism that were believed to be observed from 1992 to 1993.

The third-generation theories are different from the two former generations as they concern themselves with the effects on the value of either corporations' or banks' foreign debt in the case of a depreciation of the currency (Krugman (2001)). This was seen during the Asian financial crisis of 1997 where traders lost confidence in the currencies pegged against the dollar and therefore retracted their investments leading to a depreciation of the currencies and an increase in the value of the foreign debt of the countries' corporations. This led to an increase in bankruptcies, which created a "feed-back loop" of capital outflow as the international companies removed themselves from the countries.

The theories described above can be affected by a Tobin tax in two ways. First of all, the increase in required foreign rate of returns that we described in section 2.1, which is an effect of the added cost by the Tobin tax, and second, how this effect impacts volatility. The first effect expressed by equation (1) and Figure 1, which is the main instrument of a Tobin tax, we believe to be a rather logical assumption. The effect on volatility by an increase in the required foreign rate of return, we have shown to be more questionable.

The assumption of rational traders in the first-generation models implies that traders always know the true exchange rate wherefore the introduction of a Tobin tax will not create more volatility, but only raise the requirement for the foreign rate of return. As the rate of depreciation is known the Tobin tax will delay the point of a speculative attack until the currency has depreciated more than the size of the tax (Jørgensen & Bach (2003)). This will give governments more leeway to stabilize the economy.

The second-generation models' introduction of risk makes the point in time for the collapse uncertain. As the traders are still rational but uncertain about the other traders' perception of the true exchange rate, this makes the traders correspond to Friedman's description of informed traders in section 2.2.2. Jørgensen & Bach (2003) describe how in these models the Tobin tax works as in those of the first-generation by delaying the time of the collapse until the market's expected rate of depreciation exceeds the Tobin tax, although they fail to touch upon the negative effect on volatility that Friedman proposes the tax will have. The increase in volatility that we also found in our thesis work – although our results suggest a somewhat different market composition than Friedman suggested, cf. section 7.2 – would be the result of fewer traders in the market and a less efficient price discovering process. The purpose of a Tobin tax in the two first generations of models is exactly to delay the apparent crises giving the country more time to correct economic factors such as employment, interest rate and inflation. The increase in exchange rate volatility then counteracts this process possibly stimulating the crisis.

The primary way traders gain money on crises is by selling the domestic currency in the hope that it will depreciate by some degree; the larger the depreciation the larger the capital gains. What all the described forms of crises has in common, including the ones like the Asian financial crisis is that a Tobin tax raises the required foreign rate of return, or the required rate of depreciation of the currency. When the rate of depreciation reaches the level of the Tobin tax it stops speculation in the currency as the cost that the tax imposes on each trade results in a loss for the trader. The Tobin tax therefore helps to both slow down speculation against the currency and the outflow of money from the country (Jørgensen & Bach (2003)).

Even though a Tobin tax raises the required rate of depreciation that traders need to gain from a speculative attack on a currency, we find that it most likely also increases the volatility in the exchange rate. This increase in the required rate of depreciation that would push back the point of a breakdown might then be offset by exactly that, an increase in the rate of depreciation, due to the larger required rate of return that traders demand due to the increase in risk, which again will make it more difficult for governments to determine the correct economic policies.

9 Suggestions for further research

In this chapter we will suggest a number of recommendations regarding areas of further research related to some of the problems that we have noted throughout this thesis work as well as new research directions based on some of our conclusions.

- Regarding the data we have thoroughly noted the deviations from normality and that our financial data exhibits leptokurtic structure. Since we have used a model based on the normal distribution to estimate our coefficients we believe that it could be interesting and worthwhile to fit a model based on a leptokurtic distribution opposed to a normal distribution to examine if the results would change.

- In the majority of our models we experience heteroscedasticity which bias the variance of our estimated coefficient. Since we cannot determine the direction of the bias caused by heteroscedasticity we find that there should be spent more time on developing a general least squares model that could correct this problem.

- In this thesis (work) we have suggested an instrument to solve the problem of endogeneity between transaction costs and exchange rate volatility. We have proposed the exchange rate as an instrument and briefly examined its effect, but we feel that more research needs to be done on this subject.

- From the game theoretical model we found that both Tobin's view on the effect of a transaction tax and the opposite one could exist depending on both the pre-existing volatility and the market composition. Since minor currencies, in the model, are associated with low volatility and a small share of noise traders, and we have only examined major currencies, it could be interesting to apply our model to minor currencies to evaluate if the results match.

10 Conclusion

A Tobin tax has often been suggested as an instrument for stabilizing currencies and postponing financial crises. It is based on the conventional wisdom of Keynes and Tobin that a transaction tax would make destabilizing noise traders base their investment decisions on long-term fundamentals similar to informed traders. This is in contrast to the opposing view of Friedman that traders are rational and stabilizing, wherefore a transaction tax would make the price discovering process less efficient.

From our empirical research, we find either a positive and significant or negative and insignificant relationship between transaction costs and exchange rate volatility for the majority of our results, wherefore a Tobin tax would have the exact opposite effect on exchange rate volatility than proposed by James Tobin. This is the same result found in the previous research done on this subject, although we believe to have supported our findings further by correcting for the inherent endogeneity problem between the transaction costs and exchange rate volatility by the use of an instrument variable. We find that using the exchange rate as the instrument for our transaction costs produces the same relationship between transaction costs and exchange rate volatility as documented in our original model. We believe that volatility will have a positive effect on the difference on the exchange rate as these two variables are correlated, wherefore we have only been able to exploit the instrument for the currencies' first period.

Based only on our empirical results one would conclude that Friedman's hypothesis of how a transaction tax is destabilizing is true, which resembles the conclusions of Aliber, Chowdhry & Yan (2003), Lanne & Vesala (2006) and Hau (2006). Although we do not disagree with this conclusion, we propose that the market composition is somewhat different from Friedman's original suggestion in that the market mainly consists of informed traders. We base this proposal on the conclusions of the game theoretical model that an added transaction tax either increases or decreases volatility in major or minor currencies, respectively. Considering that we have investigated four major currencies, the game theoretical model's conclusions suggest that our empirical results resemble a market composition with a larger share of noise traders than informed traders. The game theoretical model suggests that an added transaction tax leads to a larger decrease in informed traders' demand than that of noise traders. This results in an increase in the share of noise traders, which creates more misperception and uncertainty in the price of the exchange rate, i.e. an increase in volatility, which corresponds to our empirical results of the currencies' first periods. Our results are therefore still in line with Friedman's hypothesizes of how the market, with a major currency, would respond to a Tobin tax, although we see that it can occur even when noise traders dominate the market because their risk-sharing effect outweigh the negative effect of the added risk by a marginal noise trader.

We conclude that the proposed tax rate of 0.5% made by Eichengreen, Tobin and Wyplosz (1995) is an unrealistic suggestion when compared to the average transaction costs. Although it satisfies the motivation behind the tax, and would change the market behaviour, it seems impractical to add a transaction tax up to 17 times larger than the average transaction costs. It would also, based on our results from our currencies first periods, increase volatility substantially; up to 6 times the average volatility. The realization that a Tobin tax increases transaction costs and volatility to that degree might be what has spurred the motivation of a smaller transaction tax with the sole purpose of generating revenue. We find that a tax of 0.005% on all GBP transactions worldwide generates revenue of USD 4.98 billion. This revenue has been suggested as an additional aid to either third world countries or to fighting climate changes although in the context of the recent financial crisis it has been suggested as insurance for the global taxpayers against future banking crises. Comparing the billions of dollars spent on bank bailouts in the UK in relation to the financial crisis with the estimated revenue on GBP transactions, we find that it will amount to 0.6% of the total cost of USD 850 billon on a yearly basis. We conclude that this tax rate does not seem reasonable as it will barely cover the cost of a financial crisis even if we assume that crises only occur every 10 years, and increasing the tax to generate more revenue will only result in a change in market behaviour.

11 Bibliography

Umlauf, S. R. (1993). Transaction taxes and the behavior of the Swedish stock market. *Journal of Financial Economies , 33*, 227-240.

Westerfield, R. (1977). The Distribution of Common Stock Price Changes; An Application of Transactions Time and Subordinated Stochastic Models. *Journal of Financial and Quantitative Analysis , 12*, 743-765.

Verbeek, M. (2004). *A Guide to Mordern Econometrics* (2 udg.). John Wiley & Sons, Ltd.

Alizadeh, S., Brandt, M. W., & Diebold, F. X. (2002). Range-Based estimation of Stochastic Volatility Models. *Journal of Finance , 57*, 1047-1092.

Aliber, R. Z., Chowdhry, B., & Yan, S. (2003). Some Evidence that a Tobin Tax on Foreign Exchange Transactions May Increase Volatility. *European Finance Review , 7*, 481-510.

APEC. (2010). *http://www.imf.org/external/np/exr/consult/2009/pdf/Comment91.pdf*.

BBC. (2009). *http://news.bbc.co.uk/2/hi/uk_news/8348653.stm*.

BBC. (2008). *http://news.bbc.co.uk/2/hi/business/7658277.stm*.

Bessembinder, H. (2000). Tick Size, Spreads, and Liquidity: An Analysis of Nasdaq Securities Trading near Ten Dollars. *Journal of Financial Intermediation , 9*, 213-239.

Bessembinder, H. (2003). Trade Execution Costs and Market Quality after Decemalization. *The Journal of Financial and Quantitative Analysis , 38* (4), 747-777.

Bessembinder, H., & Rath, S. (2002). Trading Costs and Return Volatility: Evidence From Exchange Listing.

Boothe, P., & Glassman, D. (1987). The Statistical Distribution of Exchange Rates; Empirical Evidence and Economic Implications. *Journal of International Economics , 22*, 297-319.

Carlson, J. A., & Osler, C. L. (1998). *Rational Speculators and Exchange Raate Volatility.* New York: Federal Reserve Bank of New York.

CBOE. (2011). *http://www.cboe.com/micro/vix/faq.aspx#1*.

Eichengreen, B., Tobin, J., & Wyplosz, C. (1995). Two Cases for Sand in the Wheels of International Finance. *The Economic Journal , 105* (428), 162-172.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy , 98* (4), 703-738.

Dimson, E., & Mussvian, M. (2000). Market efficiency. *The current State of Business Deciplines , 3*, 959-970.

Frankel, J. A. (1996). *How Well Do Foreign Exchange Markets Function: Might a Tobin Tax Help?* Cambridge.

French, K. R., Schwert, W. G., & Stambaugh, R. F. (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics , 19*, 3-29.

Frenkel, J. A., & Levich, R. M. (1979). Covered Interest Arbitrage and Unexploited Profits? Reply. *Journal of Political Economy , 87* (2), 418-422.

Frenkel, J. A., & Levich, R. M. (1975). Covered Interest Arbitrage: Unexploited Profits? *Journal of Political Economy , 83* (2), 325-338.

Frenkel, J. A., & Levich, R. M. (1977). Transaction Costs and Interest Arbitrage: Tranquil versus Turbulent Periods. *The Journal of Political Economy , 85* (6), 1209-1226.

Friedman, M. (1953). *Essays in positive economics.* Chicago: Univercity of Chicago.

FT. (2009). *http://www.ft.com/cms/s/0/aa162054-e65e-11de-bcbe-00144feab49a.html#axzz1F3vAiUU0.*

Gujarati, D. N. (2003). *Basic Econometrics* (4. International Edition udg.). McGraw Hill.

Hau, H. (2006). The Role of Traansaction Costs for Financial Volatility: Evidence from the Paris Bourse. *Journal of the European Economic Assocoation , 4* (4), 862-890.

Haberer, M. (2003). *Some Criticism of the Tobin Tax.* University of Konstanz, Center of Finance and Econometrics, Konstanz.

Harris, L. (1999). *Trading in Pennies: A Survey of the Issues.* Scottsdale, Arizona: Prepared for the NYSE Conference "U.S. Equity Markets in Transition".

Johnston, J., & DiNardo, J. (1997). *Econometric Methods* (4 udg.). (L. Sutton, & C. Berkowitz, Red.) McGraw-Hill Book Co International Editions.

Jones, C. M., & Seguin, P. J. (1997). Transaction Costs and Price Volatility: Evidence from Commission Deregulation. *The American Economic Review , 87* (4), 728-737.

Jørgensen, O., & Bach, C. F. (2003). *Tobin-skatten og modeller for valutakriser.*

Keynes, J. M. (1936). *General Theory of Employment, Interest and Money* (1 udg.). Cambridge: MacMillan and Co., Limited.

Krugman, P. (1979). A Model of Balance-of-Payments Crisis. *Journal of Money, Credit and Banking , 11* (3), 311-325.

Krugman, P. (2001). *Crises: The Next Generation?* Princeton.

Lanne, M., & Vesala, T. (2006). *The effect of a transaction tax on exchange rate volatility.* Hhelsinki, Finland: Bank of Finland Research.

Mulherin, J. H. (1990). Regulation, Trading Volume and Stock Market Volatility. *Revue économique , 41* (5), 923-938.

MCcormick, F. (1979). Covered Interest Arbitrage: Unexploited Profits? Comment. *The Journal of Political Economy , 87* (2), 411-417.

McFarland, J. W., Pettit, R. R., & Sung, K. S. (1982). The Ditribution of Foreign Exchange Price Changes; Trading day Effects and Risk Measurements. *Journal of Finance , 37* (3), 693-715.

Overø, J. E., & Gabrielsen, G. (2005). *Teoretisk Statistik - En Erhvervsøkonomisk tilgang* (3 udg.). København: Rylers I/S.

Obstfeld, M. (1986). Rational and Self-Fulfilling Balance-of-Payments Crises. *The American Economic Review , 76* (1), 72-81.

Summers, L. H., & Summers, V. P. (1989). When Financial Markets Work Too Well: A Cautious Case For a Securities Transaction Tax. *Journal of Financial Research , 3*, 261-286.

Sarno, L., & Taylor, M. P. (2001). The Microstructure of the Foreign-Exchange Market: A selective Survey on the Litterature. *Princeton Studies in International Economies , 89*.

Schulmeister, S. (2010). *Asset Price Fluctuations, Financial Crises and the Stabilizing Effect of a General Transaction Tax.* A joint publication with de Nederlandsche Bank and Rabobank.

Schwabish, J. A. (2005). Estimating Employment Spillover Effects In New York City with an Application to The Stock Transfer Tax. *Public Finance Review , 33*, 663-689.

Schwert, W. G., & Seguin, P. J. (1993). Security Transaction Taxes: An Overview of Costs, Benefits and Unresolved Questions. *Financial Analysts Journal , 49* (5), 27-35.

Schmidt, R. (2007). *The Currency Transaction Tax: rate and Revenue Estimated.* Ottawa, Ontario Canada.

Smith, S. W. (1997). *Digital Signal Processing* (Second ed.). San Diego, California: California Technical Publishing.

Song, F. M., & Zhang, J. (2005). Security Transaction Tax and Market Volatility. *The Economic Journal , 115*, 1103-1120.

Reuters. (2010). *http://thomsonreuters.com/content/press_room/financial/2010_01_18_Morgan_Stanley_FX_Prime*.

Roll, R. (1984). A Simple Implicit Measure of the Effective Bid-Ask Spread in an efficient Market. *The Journal of Finance , 39* (4), 1127-1139.

Ronen, T., & Weaver, d. G. (2001). Teenies' anyone? *Journal of Financial Markets , 4*, 231-260.

Tobin, J. (1978). A Proposal for International Monetary Reform. *Eastern Economic Journal , 4* (3-4), 153-159.

12 Appendix

Appendix 1

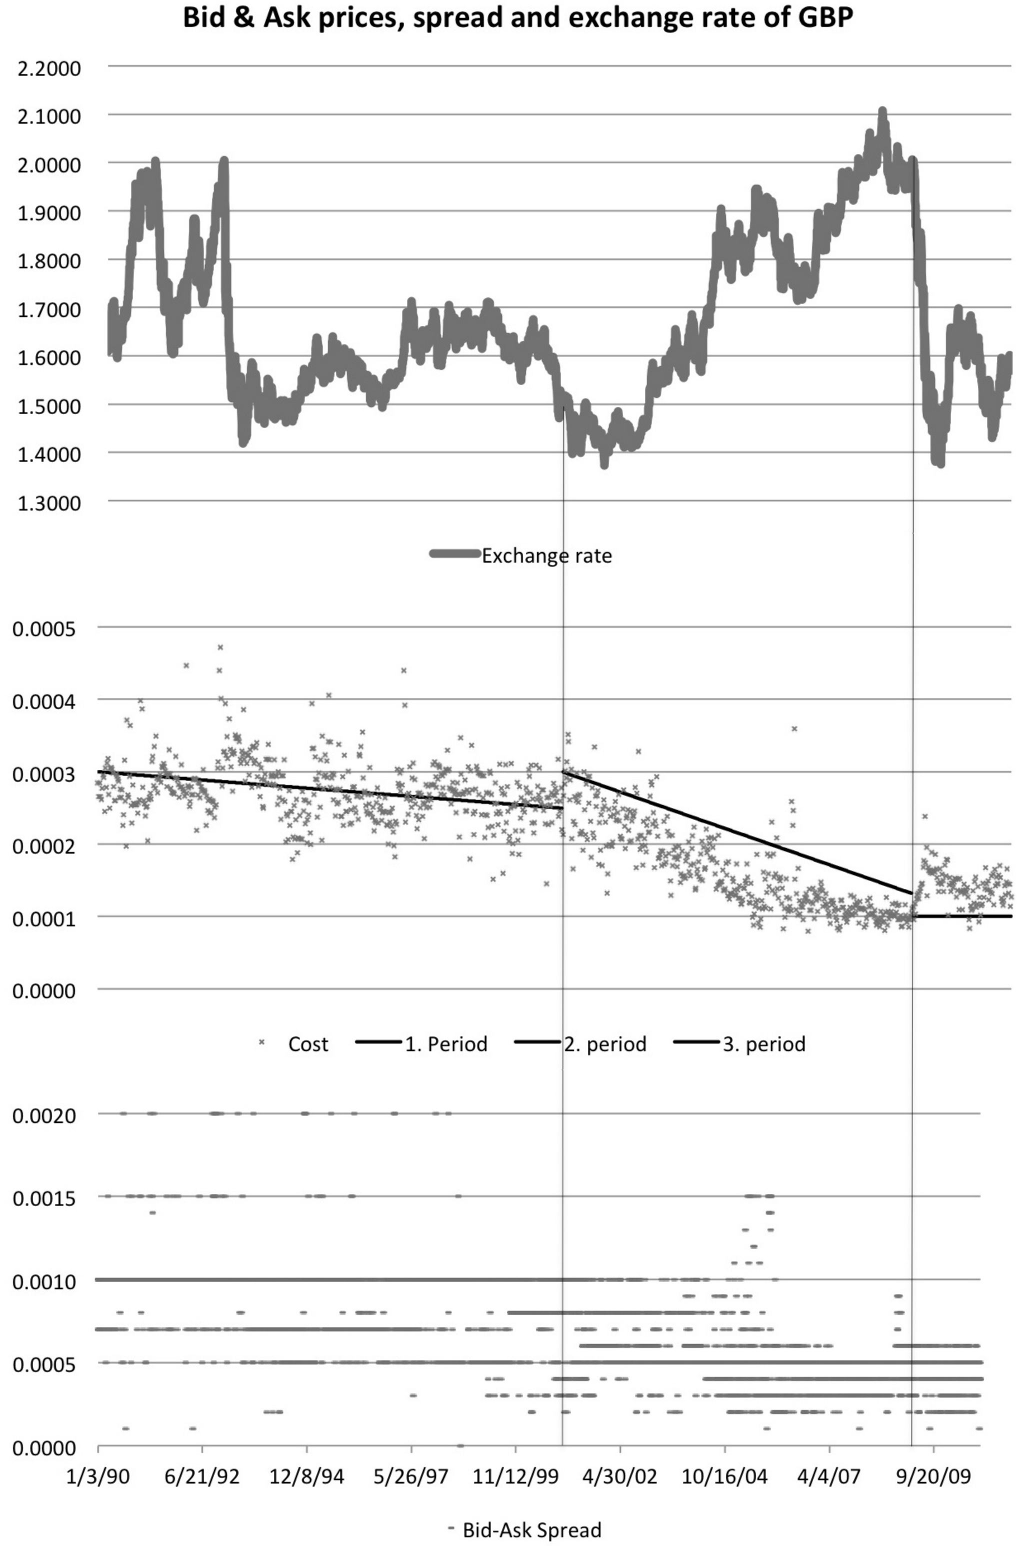

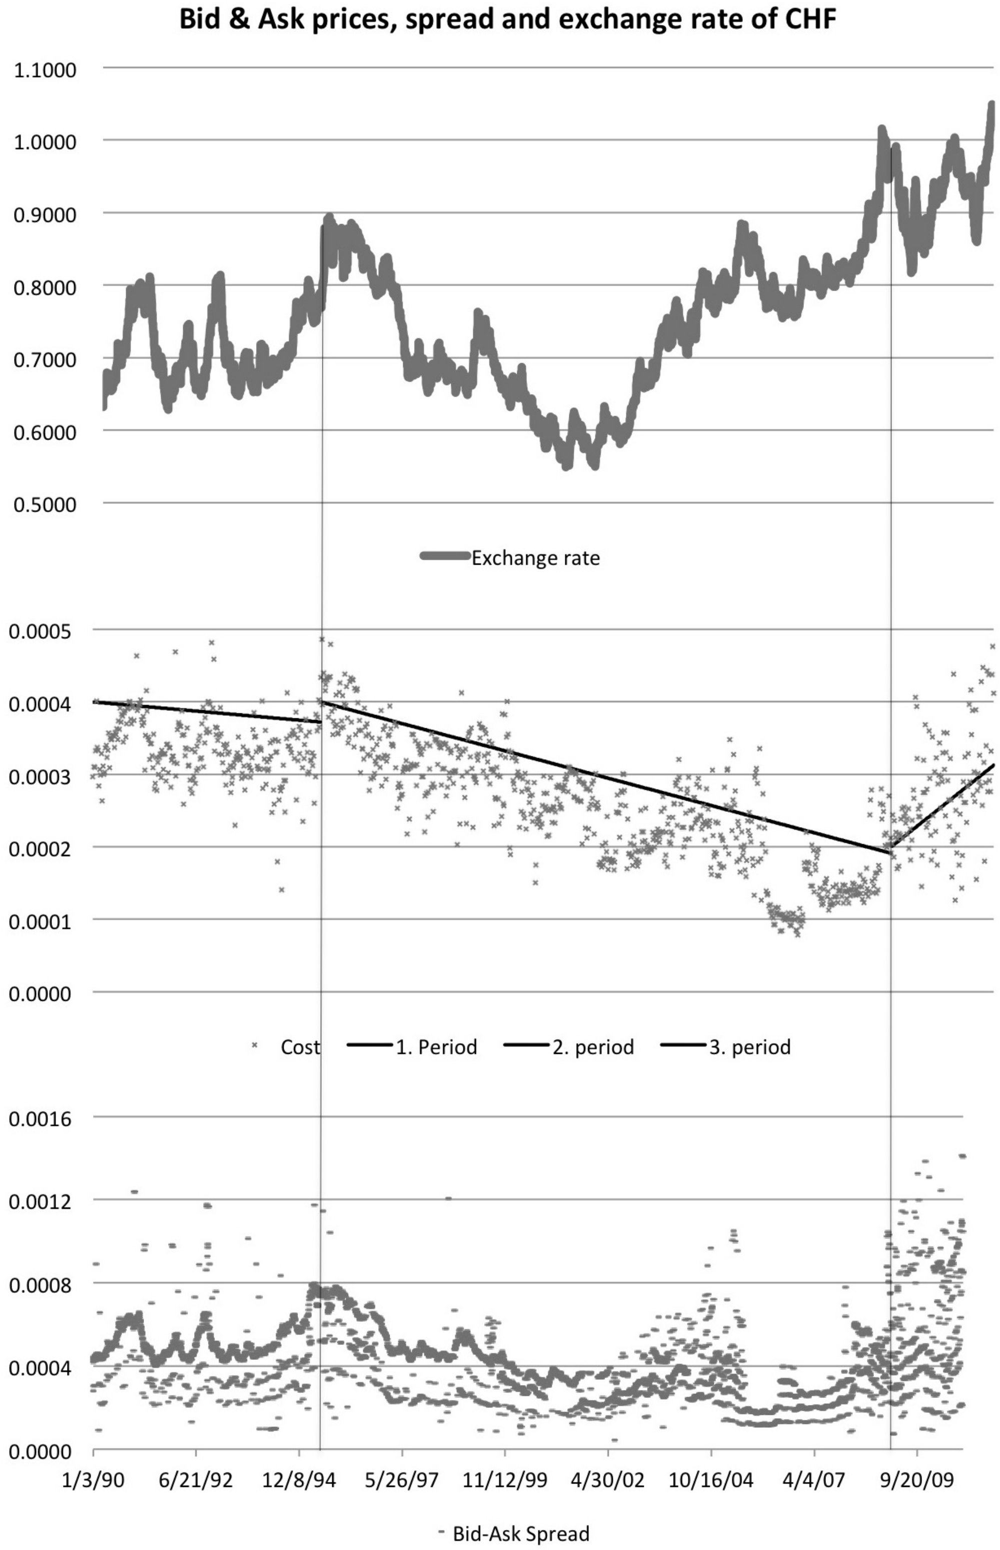
Bid & Ask prices, spread and exchange rate of CHF
1.1000
1.0000
0.9000
0.8000
0.7000
0.6000
0.5000
Exchange rate
0.0005
0.0004
0.0003
0.0002
0.0001
0.0000
Cost
1. Period
2. period
3. period
0.0016
0.0012
0.0008
0.0004
0.0000
1/3/90
6/21/92
12/8/94
5/26/97
11/12/99
4/30/02
10/16/04
4/4/07
9/20/09
Bid-Ask Spread

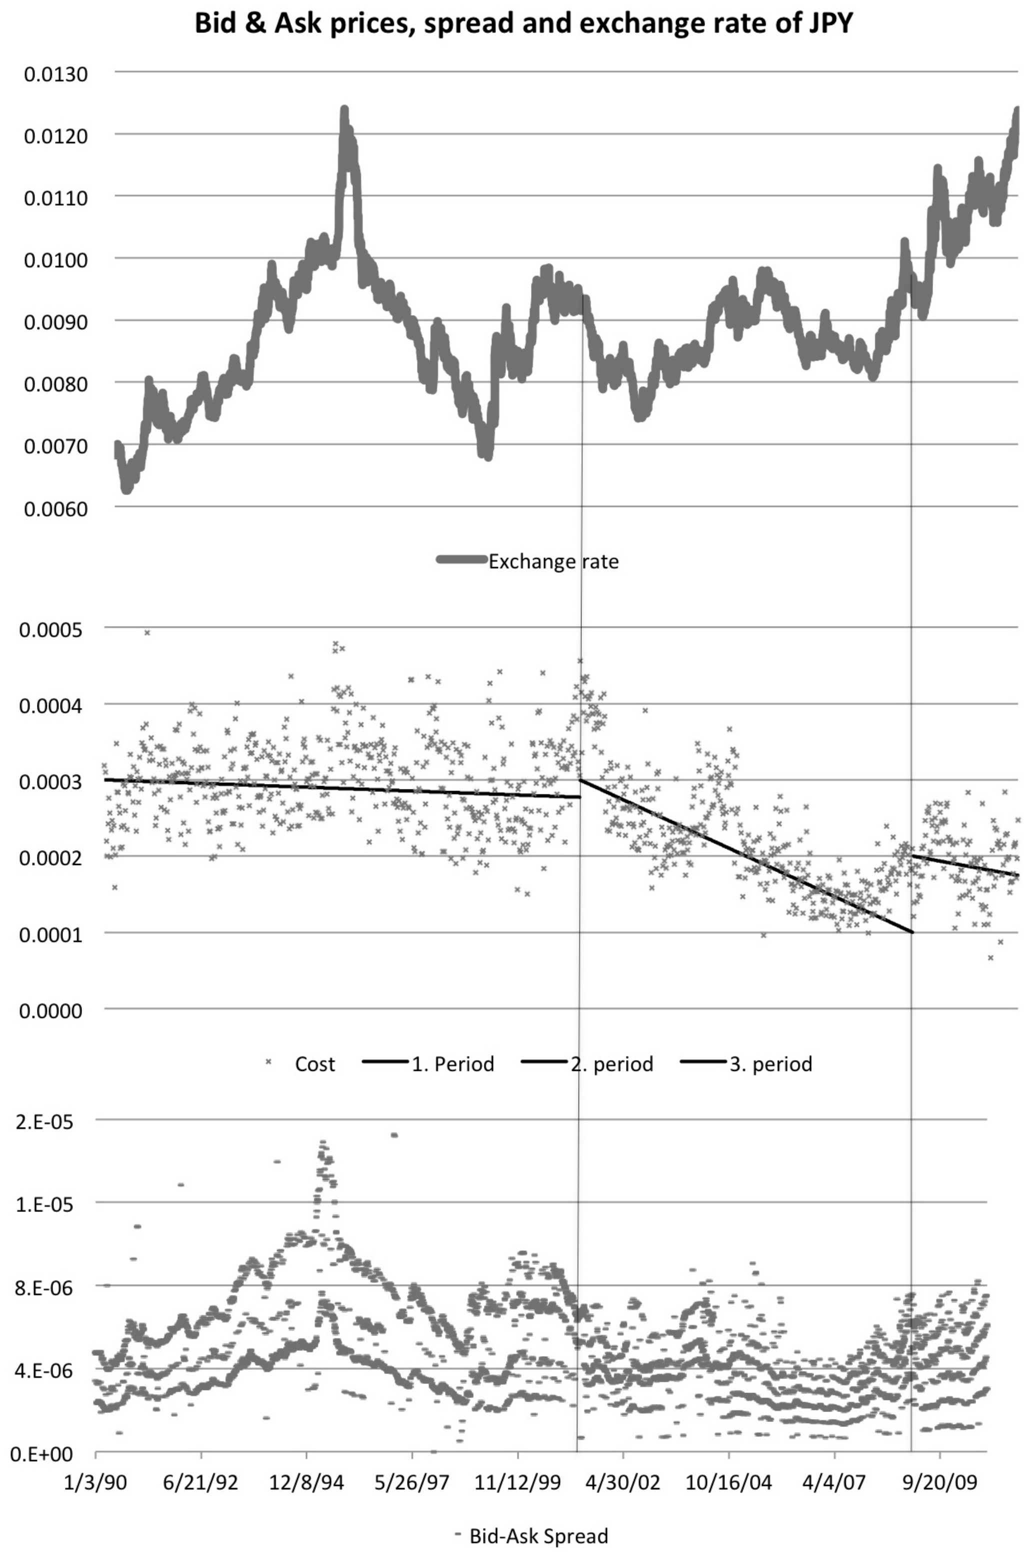
Bid & Ask prices, spread and exchange rate of JPY
0.0130
0.0120
0.0110
0.0100
0.0090
0.0080
0.0070
0.0060
Exchange rate
0.0005
0.0004
0.0003
0.0002
0.0001
0.0000
Cost
1. Period
2. period
3. period
2.E-05
1.E-05
8.E-06
4.E-06
0.E+00
1/3/90
6/21/92
12/8/94
5/26/97
11/12/99
4/30/02
10/16/04
4/4/07
9/20/09
Bid-Ask Spread

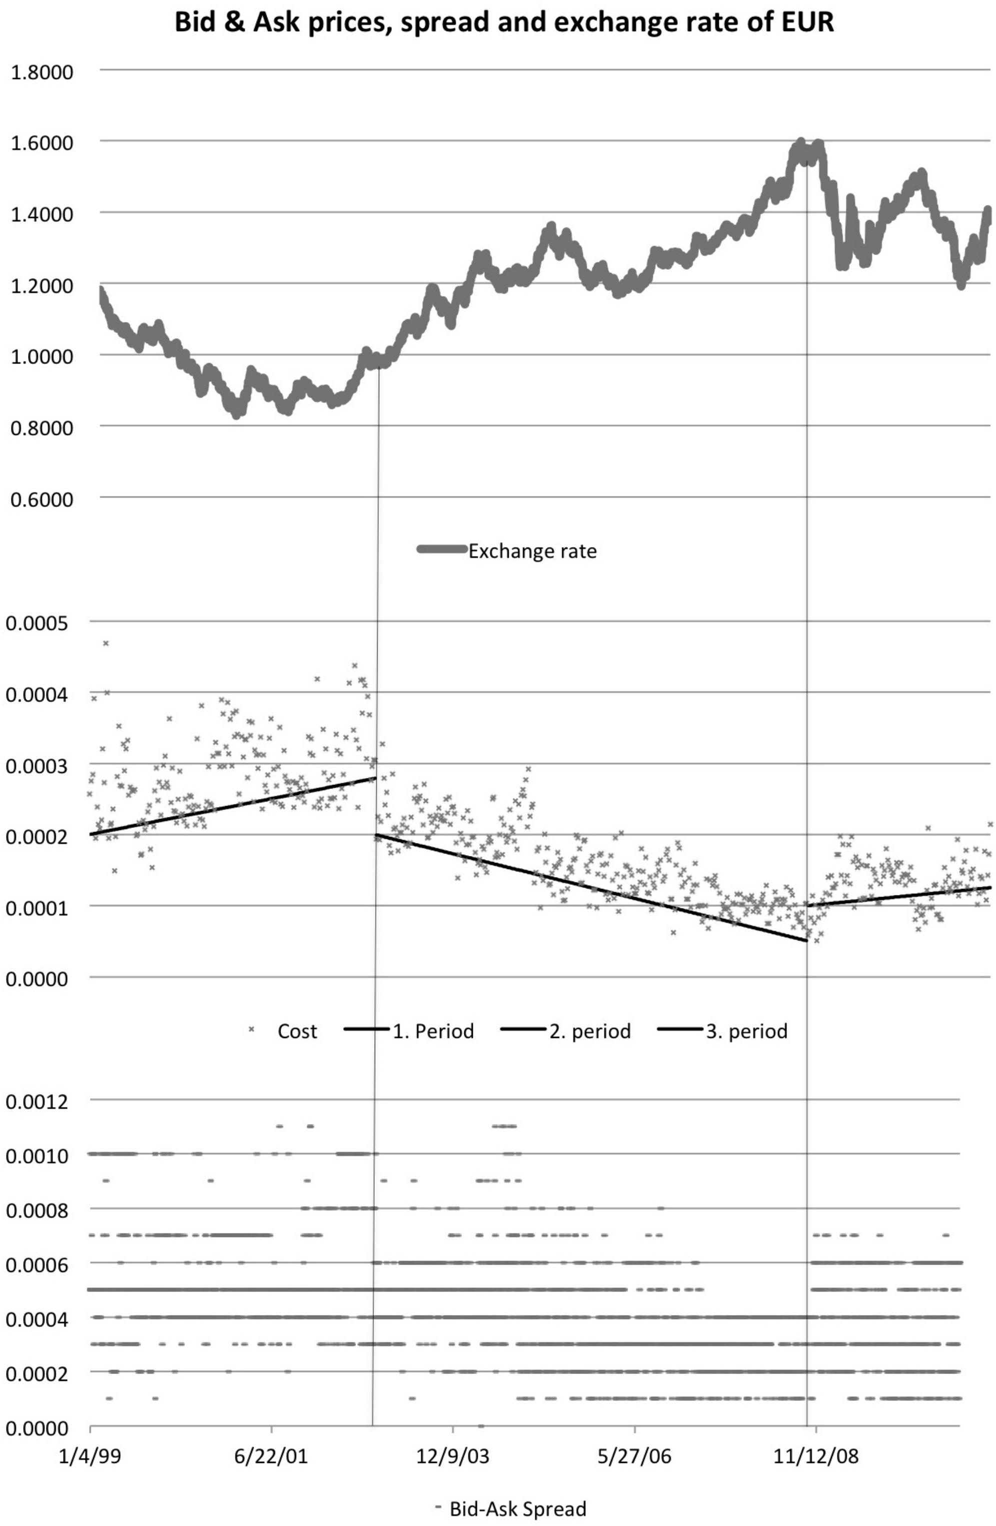
Bid & Ask prices, spread and exchange rate of EUR
1.8000
1.6000
1.4000
1.2000
1.0000
0.8000
0.6000
Exchange rate
0.0005
0.0004
0.0003
0.0002
0.0001
0.0000
Cost
1. Period
2. period
3. period
0.0012
0.0010
0.0008
0.0006
0.0004
0.0002
0.0000
1/4/99
6/22/01
12/9/03
5/27/06
11/12/08
Bid-Ask Spread

Appendix 2

		GBP					CHF					JPY					EUR				
	Period	Daily data	*Full*	*First*	*Second*	*Third*	Daily data	*Full*	*First*	*Second*	*Third*	Daily data	*Full*	*First*	*Second*	*Third*	Daily data	*Full*	*First*	*Second*	*Third*
Cost	Anderson-Darling	81.1862	14.6419	3.1986	9.4406	0.2904	40.5068	4.7723	0.6474	1.3739	1.3196	82.5283	10.6381	2.1236	5.7431	0.2367	43.9210	8.1846	2.7594	3.0007	0.1999
	Skewness	2.7693	1.0947	0.6792	0.7529	0.4625	-0.1005	0.0614	0.3789	0.0979	0.7146	4.7009	4.7702	0.7147	0.8426	-0.2275	1.0228	0.8984	1.2921	0.4270	-0.1108
	Kurtosis	29.6720	6.3046	2.4416	-0.2144	1.7121	11.7237	1.4629	0.8921	0.5056	0.3519	91.8932	71.1964	1.6816	0.3033	-0.1188	1.3798	1.0436	3.5507	-0.6241	-0.2550
Volatility	Anderson-Darling	195.3643	43.9805	14.6127	2.1446	5.5792	120.0808	17.2436	6.2227	5.2737	4.5889	205.4958	31.7556	14.4919	4.9577	4.2805	64.6535	13.2182	3.7357	0.8583	4.3854
	Skewness	2.3445	2.4687	1.9696	0.5926	1.4758	1.7506	1.5918	1.4433	1.0019	2.2172	3.3856	2.6068	2.6945	1.1592	1.7157	1.6934	1.6824	1.0223	0.3250	1.3798
	Kurtosis	10.5587	9.9629	7.4235	0.5759	2.4065	5.6960	6.0887	3.4712	2.1867	11.1100	27.2425	15.9776	17.5785	3.2253	4.5592	5.6222	5.7593	1.2288	-0.3627	2.7620
VIX	Anderson-Darling	137.6379	27.0916	7.8133	6.3372	6.3126	137.6279	27.0916	9.9934	5.7525	7.0962	137.6279	27.0916	7.4187	7.2263	7.4342	137.6279	27.0916	4.9822	10.4691	7.4342
	Skewness	1.9901	1.9418	0.9583	0.7297	1.3218	1.9901	1.9418	1.3592	0.6916	1.3835	1.9901	1.9418	0.9380	0.7702	1.4067	1.9901	1.9418	1.2422	1.0471	1.4067
	Kurtosis	6.9619	6.5538	1.0455	-0.0610	1.0809	6.9619	6.5538	1.5606	0.2740	1.2749	6.9619	6.5538	1.0233	-0.0631	1.3452	6.9619	6.5538	1.6191	0.2688	1.3452

Appendix 3

Removed outliers (yyyyww)

GBP	CHF		JPY
1. period	1. period	2. period	1. period
199052	199051	201000	199033
199152	199052		199052
199237	199238		199100
199251	199414		199552
199544	199419		
199651			

Appendix 4

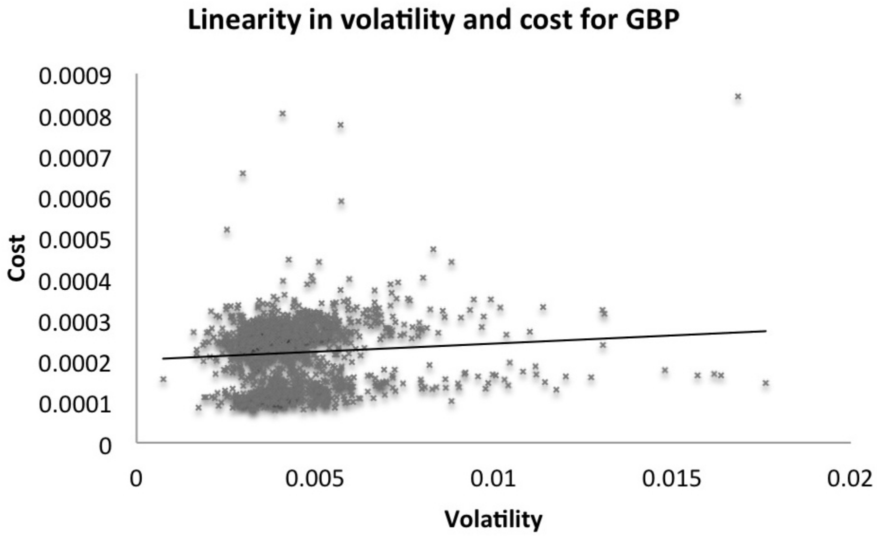
Linearity in volatility and cost for GBP
0.0009
0.0008
0.0007
0.0006
0.0005
0.0004
0.0003
0.0002
0.0001
0
Cost
0
0.005
0.01
0.015
0.02
Volatility

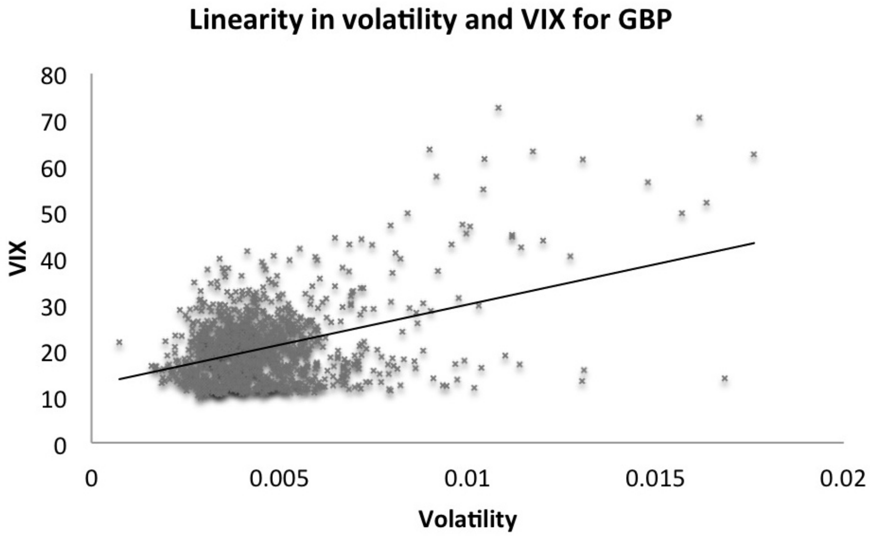
Linearity in volatility and VIX for GBP
80
70
60
50
40
30
20
10
0
VIX
0
0.005
0.01
0.015
0.02
Volatility

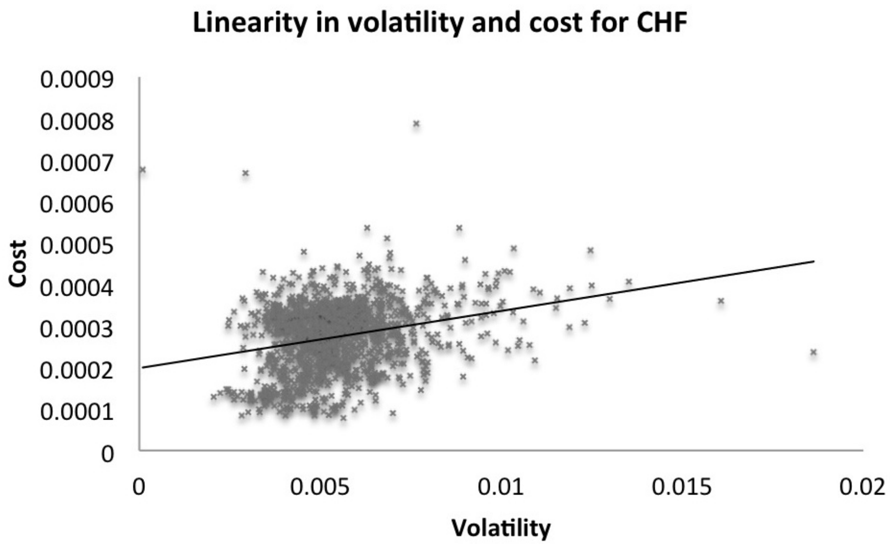
Linearity in volatility and cost for CHF
0.0009
0.0008
0.0007
0.0006
0.0005
0.0004
0.0003
0.0002
0.0001
0
Cost
0
0.005
0.01
0.015
0.02
Volatility

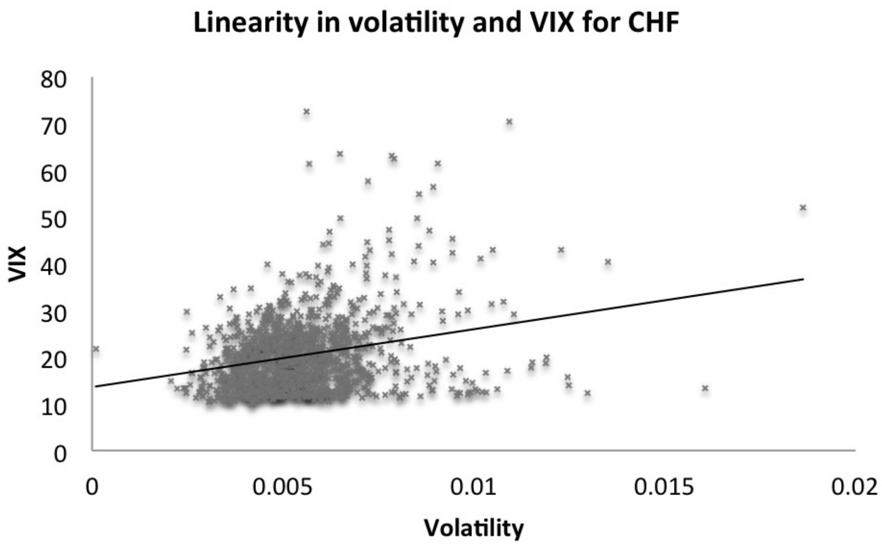
Linearity in volatility and VIX for CHF
80
70
60
50
40
30
20
10
0
VIX
0
0.005
0.01
0.015
0.02
Volatility

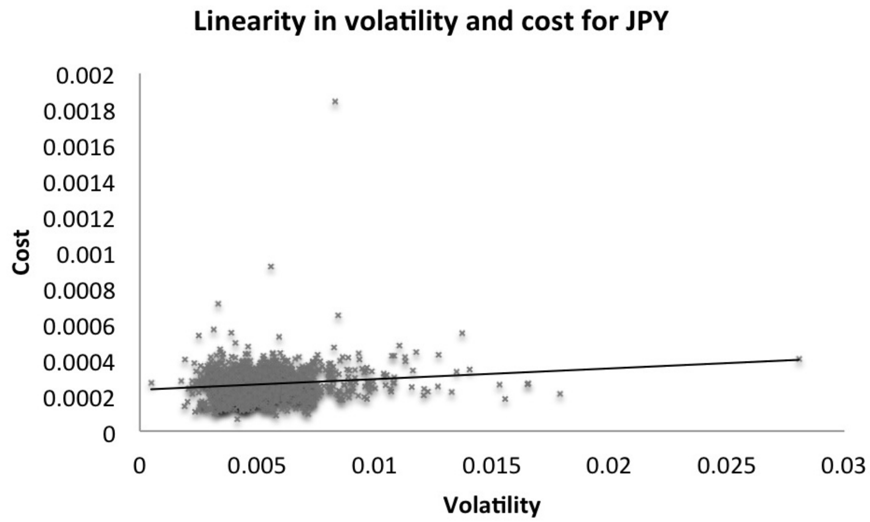
Linearity in volatility and cost for JPY
Cost
0.002
0.0018
0.0016
0.0014
0.0012
0.001
0.0008
0.0006
0.0004
0.0002
0
0
0.005
0.01
0.015
0.02
0.025
0.03
Volatility

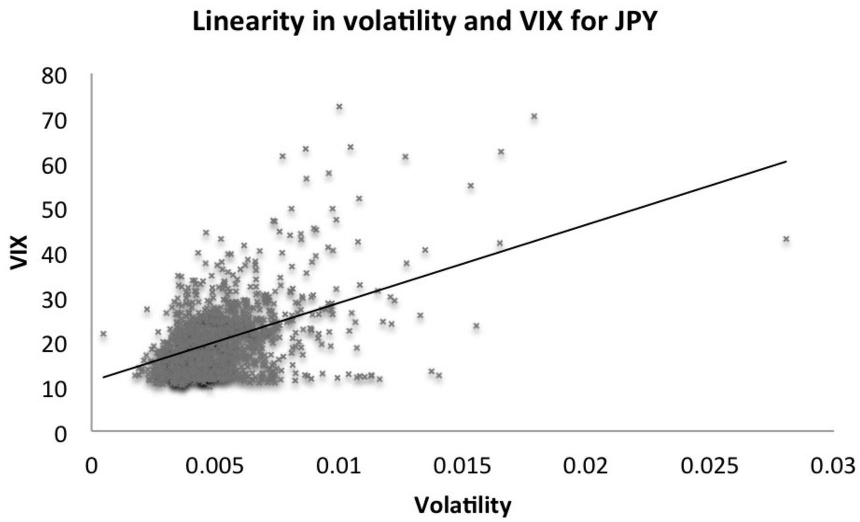
Linearity in volatility and VIX for JPY
VIX
80
70
60
50
40
30
20
10
0
0
0.005
0.01
0.015
0.02
0.025
0.03
Volatility

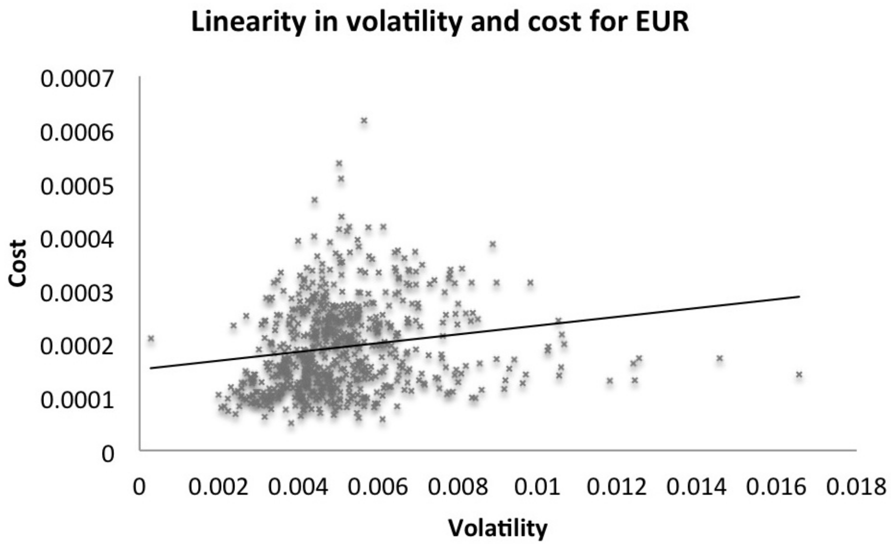
Linearity in volatility and cost for EUR
Cost
0.0007
0.0006
0.0005
0.0004
0.0003
0.0002
0.0001
0
0
0.002
0.004
0.006
0.008
0.01
0.012
0.014
0.016
0.018
Volatility

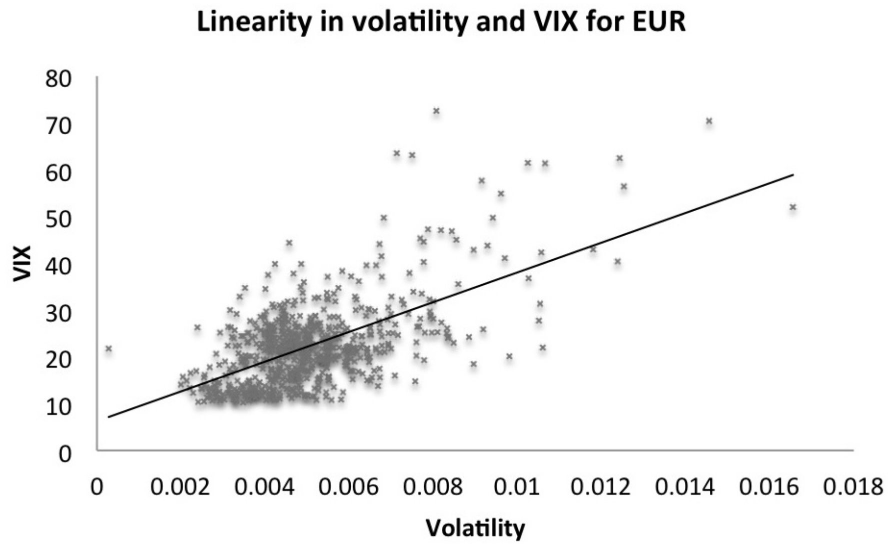
Linearity in volatility and VIX for EUR
VIX
80
70
60
50
40
30
20
10
0
0
0.002
0.004
0.006
0.008
0.01
0.012
0.014
0.016
0.018
Volatility

Appendix 5

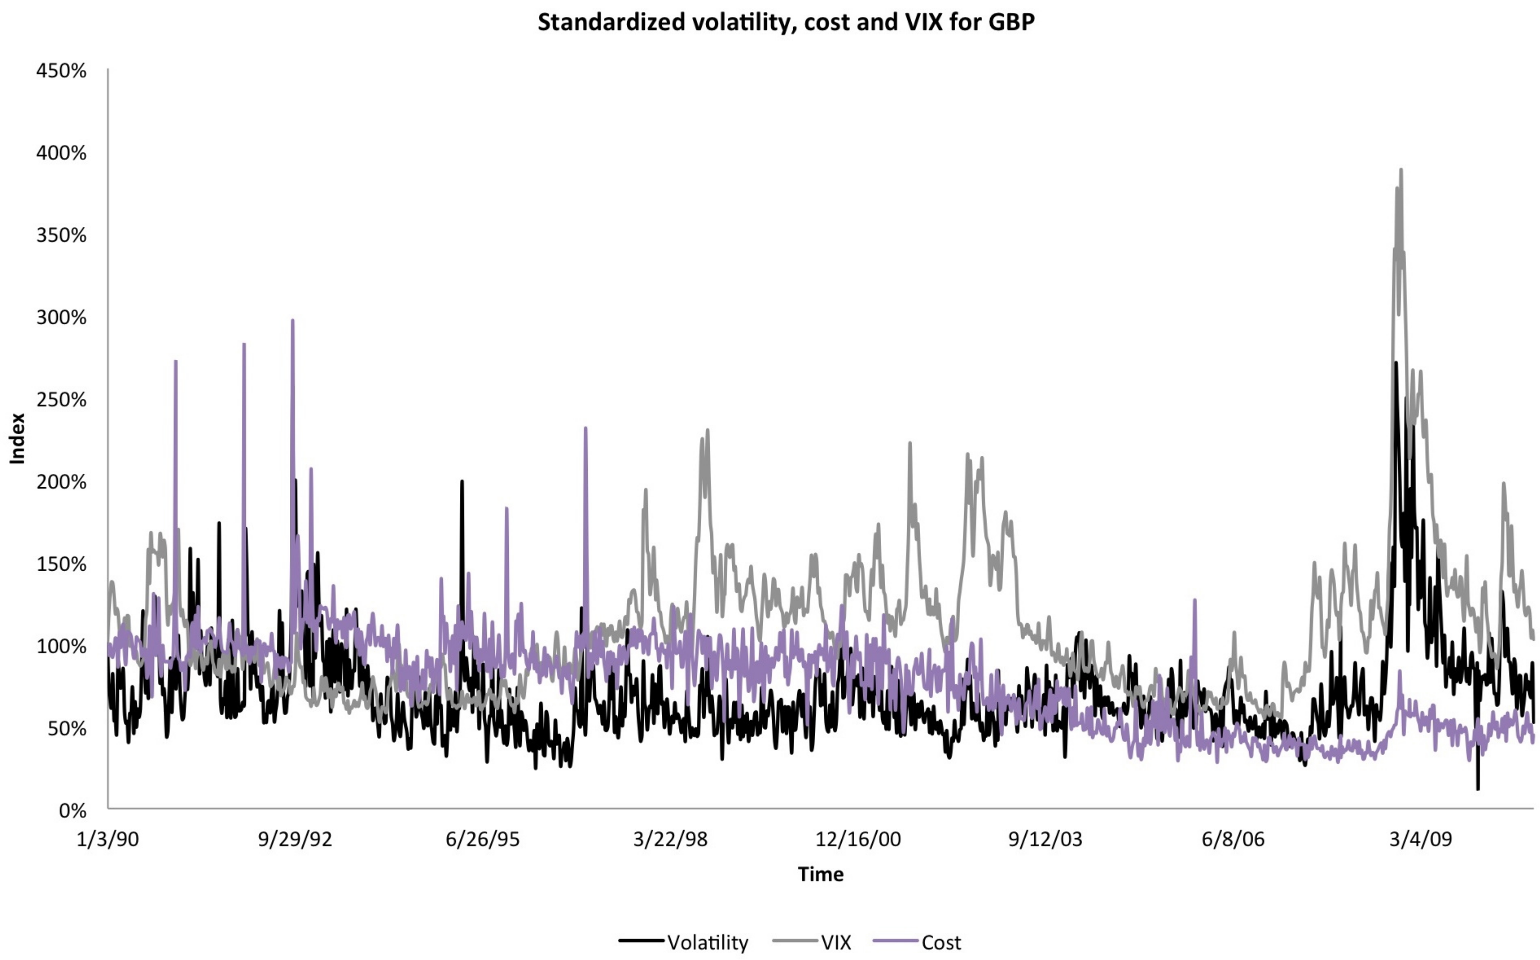
Standardized volatility, cost and VIX for GBP
450%
400%
350%
300%
250%
200%
150%
100%
50%
0%
Index
1/3/90
9/29/92
6/26/95
3/22/98
12/16/00
9/12/03
6/8/06
3/4/09
Time
Volatility
VIX
Cost

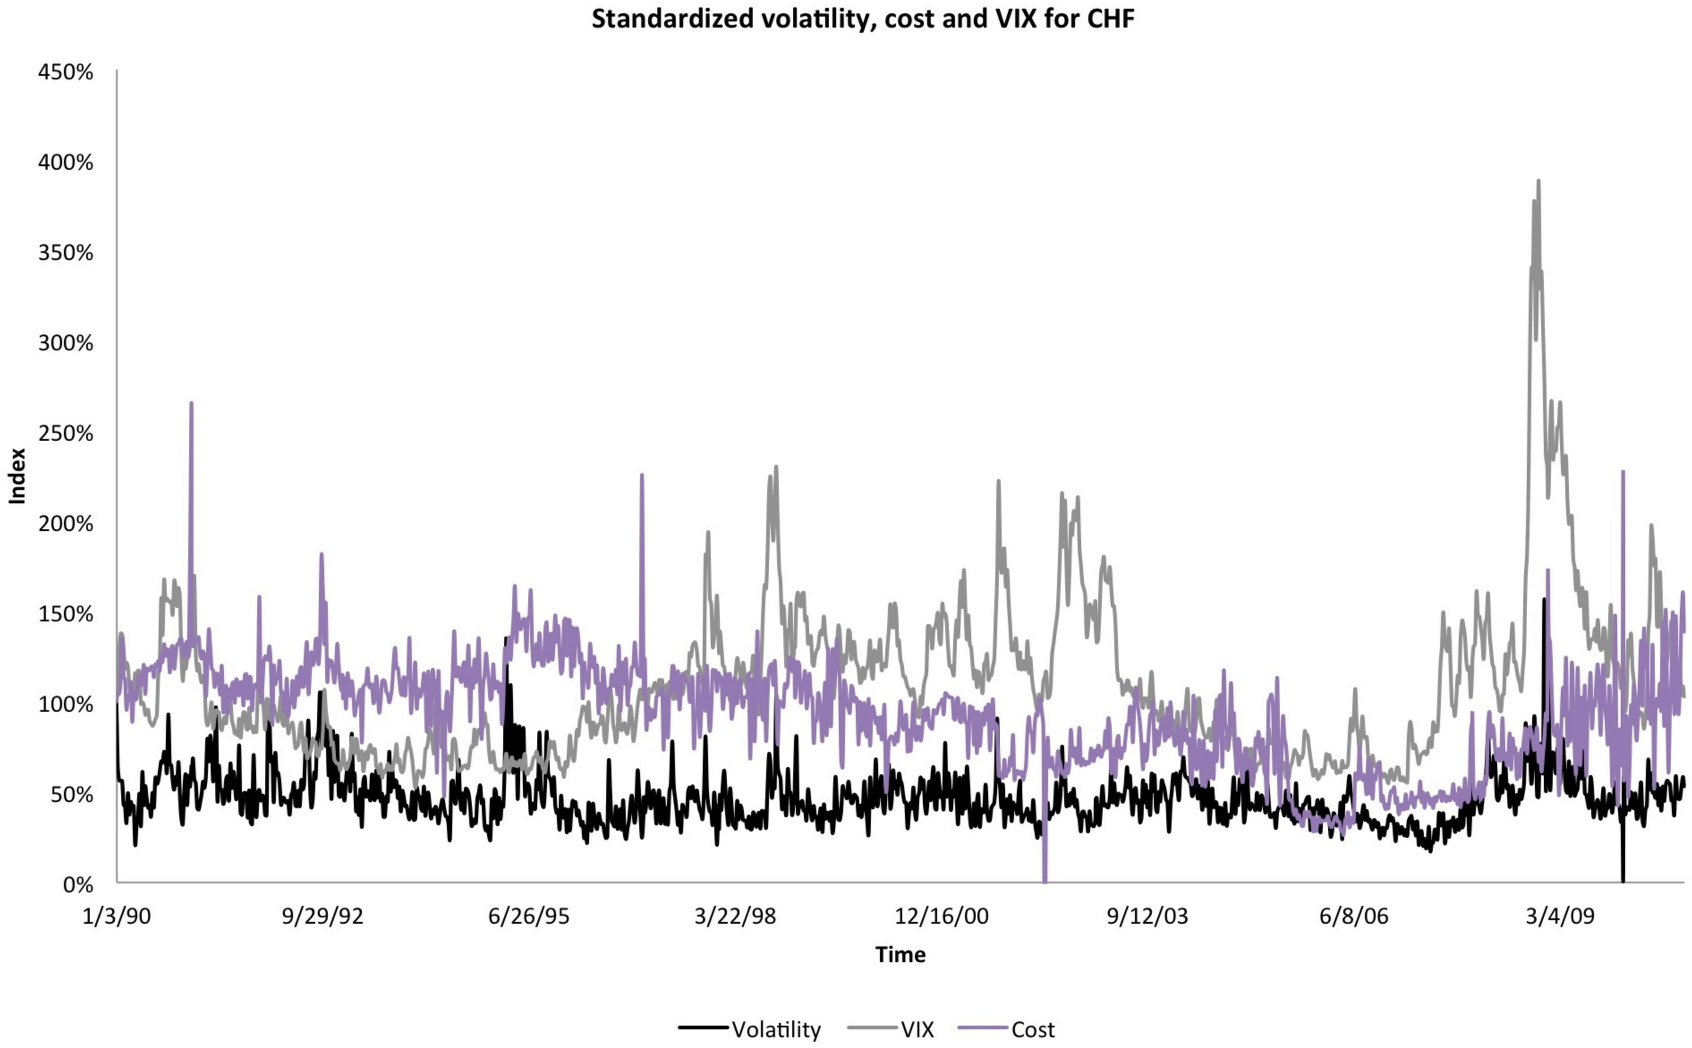
Standardized volatility, cost and VIX for CHF
450%
400%
350%
300%
250%
200%
150%
100%
50%
0%
Index
1/3/90
9/29/92
6/26/95
3/22/98
12/16/00
9/12/03
6/8/06
3/4/09
Time
Volatility
VIX
Cost

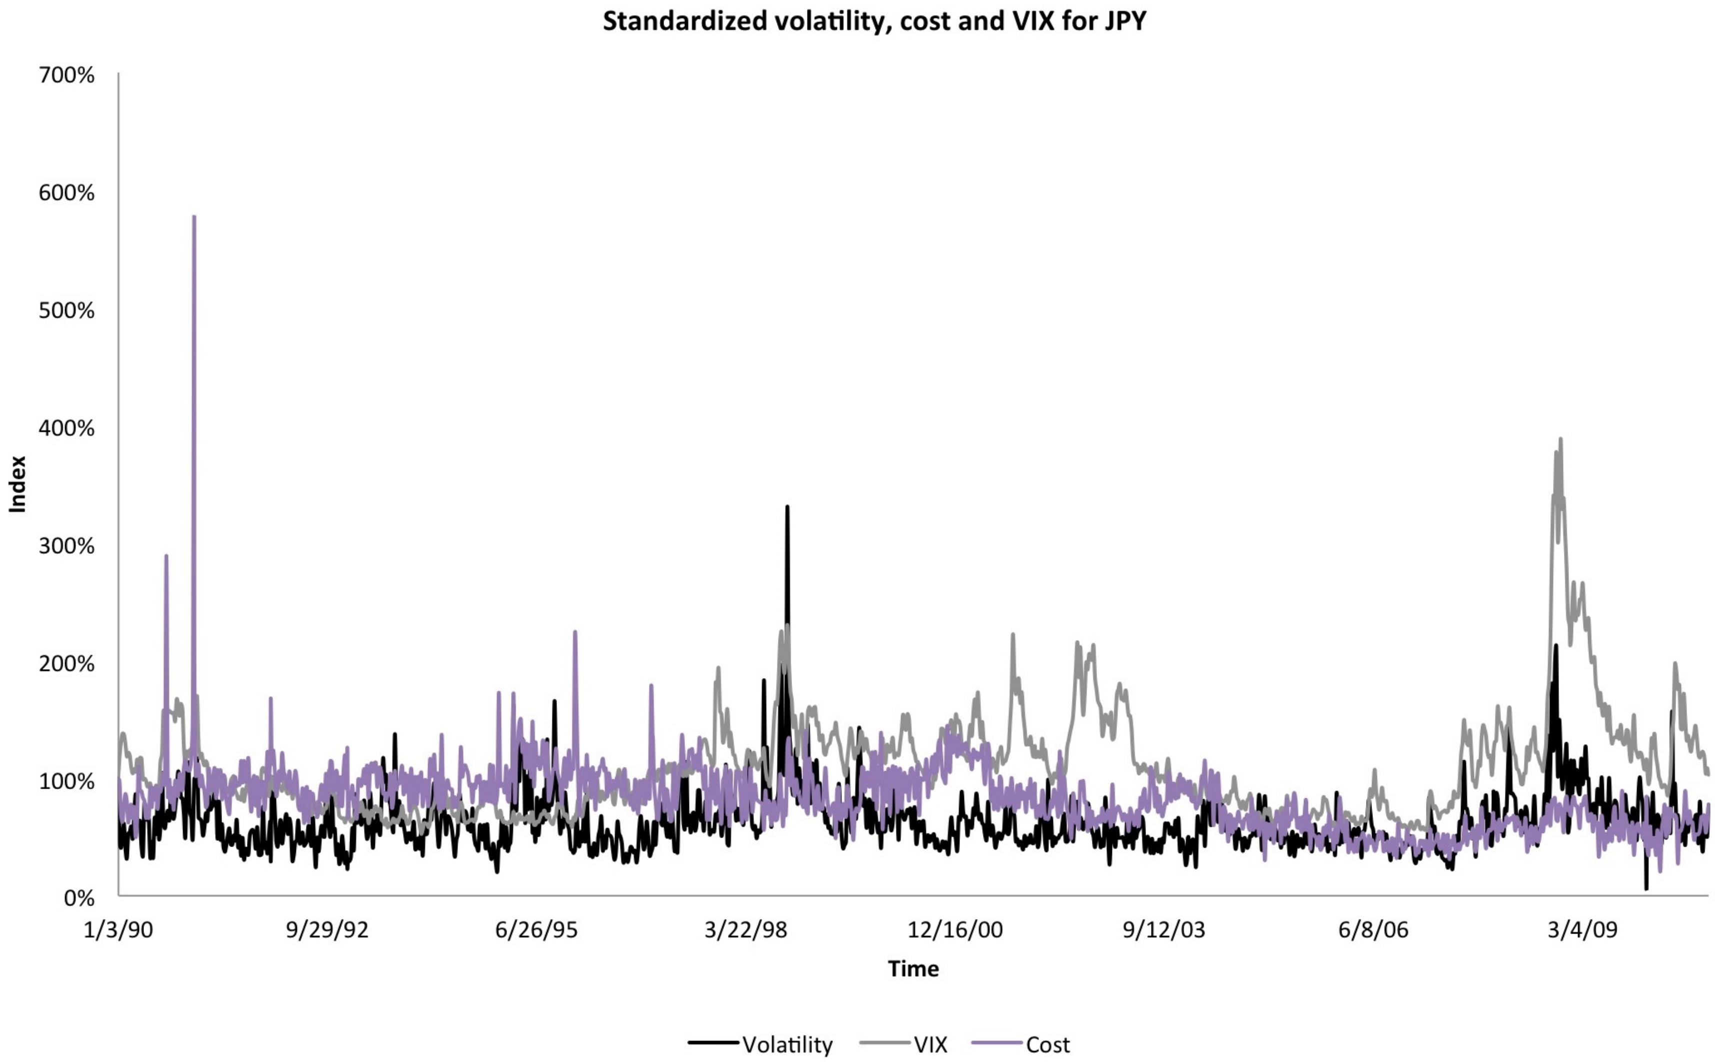
Standardized volatility, cost and VIX for JPY
700%
600%
500%
400%
300%
200%
100%
0%
Index
1/3/90
9/29/92
6/26/95
3/22/98
12/16/00
9/12/03
6/8/06
3/4/09
Time
Volatility
VIX
Cost

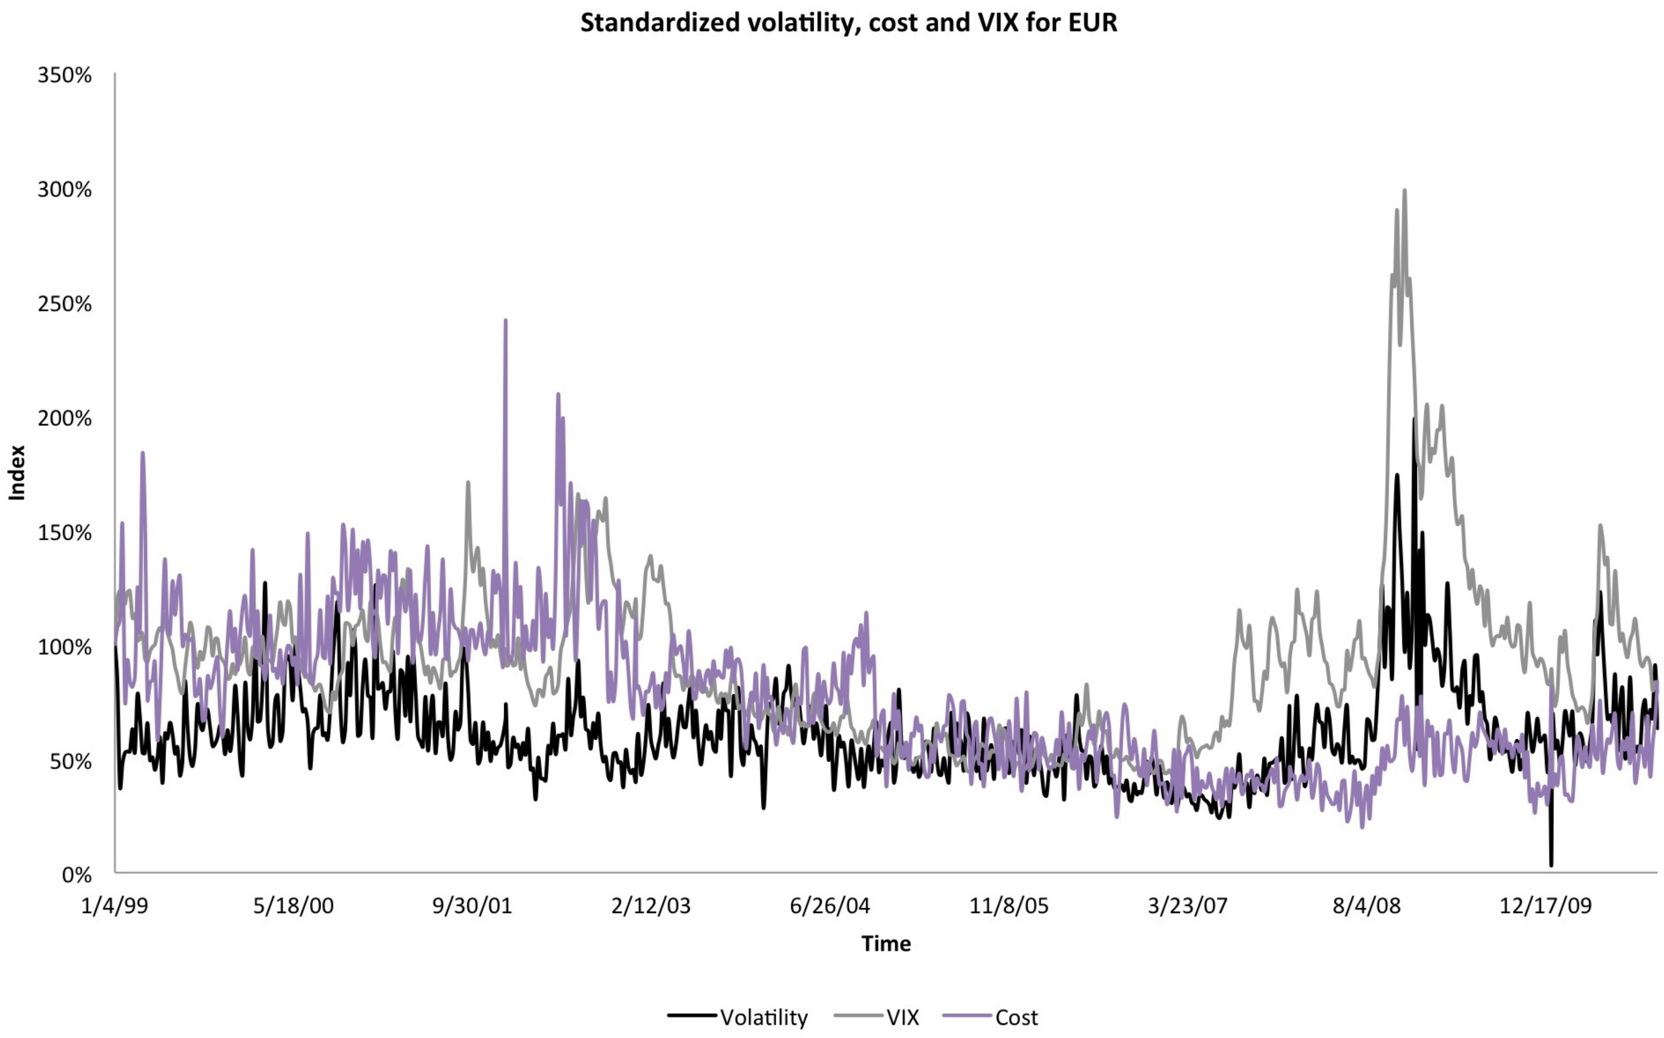
Standardized volatility, cost and VIX for EUR
350%
300%
250%
200%
150%
100%
50%
0%
Index
1/4/99
5/18/00
9/30/01
2/12/03
6/26/04
11/8/05
3/23/07
8/4/08
12/17/09
Time
Volatility
VIX
Cost

Appendix 6

CHF

Period		Scenario	Lag	$\%\Delta_{1-2}$	$\%\Delta_{2-3}$	$\%\Delta_{3-4}$	$\%\Delta_{4-5}$	$\%\Delta_{5-6}$
Full	Volatility	I(1), I(0)	2	9.02%	1.91%	1.13%	0.25%	-0.18%
	Cost	I(1), I(0)	5	10.90%	2.00%	1.90%	2.55%	0.05%
	Volatility	I(1), I(1)	3	5.95%	3.37%	0.69%	0.80%	0.89%
	Cost	I(1), I(1)	5	10.89%	1.98%	2.01%	2.45%	0.54%
First	Volatility	I(0) , I(0)	2	3.10%	-0.10%	-1.77%	-1.57%	-1.45%
	Cost	I(0) , I(0)	4	2.10%	2.99%	0.06%	1.65%	1.27%
	Volatility	I(0) , I(1)	2	3.75%	-1.23%	-1.40%	-0.32%	-0.36%
	Cost	I(0) , I(1)	5	4.51%	5.19%	3.04%	-1.21%	1.10%
Second	Volatility	I(1), I(0)	2	3.83%	2.17%	2.61%	2.05%	-0.28%
	Cost	I(1), I(0)	2	11.38%	0.86%	2.66%	1.20%	-0.17%
Third	Volatility	I(0) , I(1)	3	6.38%	6.06%	-3.10%	-1.54%	1.05%
	Cost	I(0) , I(1)	2	4.14%	-3.44%	-0.92%	-4.19%	0.65%
	Volatility	I(1), I(1)	3	3.39%	5.08%	-3.42%	-1.51%	-0.11%
	Cost	I(1), I(1)	2	4.66%	4.13%	-3.15%	-4.10%	5.34%

JPY

Period		Scenario	Lag	$\%\Delta_{1-2}$	$\%\Delta_{2-3}$	$\%\Delta_{3-4}$	$\%\Delta_{4-5}$	$\%\Delta_{5-6}$
Full	Volatility	I(1), I(0)	3	5.60%	1.76%	1.35%	1.09%	-0.34%
	Cost	I(1), I(0)	3	6.52%	5.89%	2.10%	1.54%	0.52%
	Volatility	I(1), I(1)	3	5.14%	4.10%	2.33%	0.14%	-0.02%
	Cost	I(1), I(1)	4	6.74%	6.05%	2.00%	1.46%	0.50%
First	Volatility	I(0) , I(0)	4	3.64%	1.11%	3.43%	1.74%	-0.35%
	Cost	I(0) , I(0)	2	5.29%	0.71%	1.82%	0.63%	-0.74%
	Volatility	I(0) , I(1)	4	4.36%	7.58%	3.17%	-0.19%	-0.17%
	Cost	I(0) , I(1)	2	5.91%	1.00%	1.63%	0.02%	-0.75%
Second	Volatility	I(1), I(0)	2	3.97%	0.28%	-1.03%	-1.11%	-1.19%
	Cost	I(1), I(0)	3	6.51%	7.36%	1.82%	-0.70%	-0.70%
Third	Volatility	I(0) , I(1)	2	6.85%	-3.89%	0.01%	-0.19%	5.38%
	Cost	I(0) , I(1)	1	0.46%	0.88%	-2.35%	-3.00%	-2.91%
	Volatility	I(1), I(1)	5	6.14%	-0.99%	-0.85%	6.03%	-2.66%
	Cost	I(1), I(1)	4	9.37%	0.55%	-1.07%	-3.61%	-2.03%

GBP

Period		Scenario	Lag	$\%\Delta_{1-2}$	$\%\Delta_{2-3}$	$\%\Delta_{3-4}$	$\%\Delta_{4-5}$	$\%\Delta_{5-6}$
Full	Volatility	I(1), I(0)	3	15.56%	1.34%	0.33%	0.96%	0.19%
	Cost	I(1), I(0)	3	7.70%	4.94%	1.42%	1.20%	2.00%
	Volatility	I(1), I(1)	3	2.33%	1.87%	1.76%	0.82%	0.11%
	Cost	I(1), I(1)	3	7.73%	5.12%	1.22%	1.54%	1.74%
First	Volatility	I(0) , I(0)	2	5.10%	0.00%	-0.35%	-0.17%	-0.60%
	Cost	I(0) , I(0)	2	1.91%	0.70%	-0.12%	-0.64%	-0.25%
	Volatility	I(0) , I(1)	2	3.99%	1.91%	1.24%	0.05%	0.74%
	Cost	I(0) , I(1)	2	2.74%	1.12%	0.80%	-0.27%	0.01%
Second	Volatility	I(1), I(0)	3	9.43%	3.77%	0.35%	0.81%	2.72%
	Cost	I(1), I(0)	2	5.01%	0.07%	1.10%	2.85%	1.17%
	Volatility	I(1), I(1)	3	8.18%	2.09%	0.43%	1.78%	2.64%
	Cost	I(1), I(1)	2	4.97%	0.07%	1.81%	2.15%	1.24%
Third	Volatility	I(1), I(1)	2	2.31%	-3.41%	-0.91%	-1.89%	-1.93%
	Cost	I(1), I(1)	3	1.60%	7.20%	-3.08%	-1.13%	-3.05%

EUR

Period		Scenario	Lag	$\%\Delta_{1-2}$	$\%\Delta_{2-3}$	$\%\Delta_{3-4}$	$\%\Delta_{4-5}$	$\%\Delta_{5-6}$
Full	Volatility	I(1), I(0)	2	15.46%	1.55%	3.31%	0.77%	-0.53%
	Cost	I(1), I(0)	4	7.95%	3.87%	3.83%	0.22%	0.90%
	Volatility	I(1), I(1)	3	5.67%	5.57%	1.43%	-0.03%	0.79%
	Cost	I(1), I(1)	4	8.06%	3.97%	4.12%	-0.12%	1.12%
First	Volatility	I(0) , I(0)	4	-1.94%	2.50%	3.36%	-2.03%	-1.84%
	Cost	I(0) , I(0)	1	-0.92%	-0.56%	0.41%	0.45%	-2.25%
	Volatility	I(0) , I(1)	3	12.20%	9.23%	-1.18%	-1.37%	0.20%
	Cost	I(0) , I(1)	3	-0.81%	2.40%	0.02%	-1.10%	-1.16%
Second	Volatility	I(1), I(0)	3	8.15%	5.21%	0.05%	-1.01%	1.27%
	Cost	I(1), I(0)	4	5.30%	0.27%	5.77%	1.35%	2.02%
Third	Volatility	I(1), I(1)	1	-1.89%	-1.30%	0.50%	-3.46%	-2.96%
	Cost	I(1), I(1)	2	6.87%	0.54%	3.72%	2.12%	-2.42%

Appendix 7

CHF

Period		Scenario	Lag	Skewness	Kurtosis	Anderson-Darling	Auto-correlation	Multi-collinearity	Breusch-Pagan (p)
Full	Volatility	I(1), I(0)	2	1.6956	9.0309	15.8736	(0.0440)	8.7794	<.0001
	Cost	I(1), I(0)	5	0.6804	15.6212	21.9664	(0.0120)	12.6968	0.0709
	Volatility	I(1), I(1)	3	1.2130	7.8954	10.1126	(0.0205)	2.2388	0.0490
	Cost	I(1), I(1)	5	0.6815	15.6975	22.0737	(0.0122)	3.2033	0.0675
First	Volatility	I(0) , I(0)	2	1.5267	5.3834	3.9158	(0.0203)	28.2405	0.0646
	Cost	I(0) , I(0)	4	0.1134	1.5866	1.5631	0.0038	39.6481	0.8391
	Volatility	I(0) , I(1)	2	1.1080	4.1474	2.3630	(0.0112)	21.5021	0.0846
	Cost	I(0) , I(1)	5	(0.0057)	1.5692	1.1907	(0.0102)	32.4375	0.1594
Second	Volatility	I(1), I(0)	2	1.1137	2.9985	6.3583	(0.0219)	9.8055	0.0255
	Cost	I(1), I(0)	2	(0.7821)	18.7497	11.2953	(0.0378)	9.8055	0.6709
Third	Volatility	I(0) , I(1)	3	2.1498	12.0178	2.2802	(0.0151)	12.1200	0.5668
	Cost	I(0) , I(1)	2	0.3255	0.1861	0.5656	(0.0248)	9.9013	0.0309
	Volatility	I(1), I(1)	3	2.3291	12.9676	2.8742	(0.0083)	2.8366	0.8514
	Cost	I(1), I(1)	2	0.1476	0.1441	0.1712	(0.0678)	2.0689	0.0543

GBP

Period		Scenario	Lag	Skewness	Kurtosis	Anderson-Darling	Auto-correlation	Multi-collinearity	Breusch-Pagan (p)
Full	Volatility	I(1), I(0)	3	2.0993	12.2112	22.8484	(0.0098)	9.7140	<.0001
	Cost	I(1), I(0)	3	4.3747	49.3124	49.8483	(0.0277)	9.7140	0.3934
	Volatility	I(1), I(1)	3	1.5541	10.4130	17.6082	(0.0192)	2.1834	0.0008
	Cost	I(1), I(1)	3	4.3405	49.1585	49.9708	(0.0294)	2.1834	0.2828
First	Volatility	I(0) , I(0)	2	1.6924	6.4595	9.8503	(0.0216)	22.0715	<.0001
	Cost	I(0) , I(0)	2	0.1087	3.3991	3.4783	(0.0200)	22.0715	0.4798
	Volatility	I(0) , I(1)	2	0.9840	5.9263	6.7520	(0.0351)	16.6250	<.0001
	Cost	I(0) , I(1)	2	0.6390	3.2781	4.1674	(0.0246)	16.6250	0.4524
Second	Volatility	I(1), I(0)	3	0.7790	0.9955	3.1595	(0.0252)	12.0924	0.5957
	Cost	I(1), I(0)	2	0.1087	3.3991	5.1230	(0.0257)	10.4172	0.3681
	Volatility	I(1), I(1)	3	0.5910	1.1337	2.2547	(0.0223)	2.3747	0.8045
	Cost	I(1), I(1)	2	0.1078	3.4029	5.1381	(0.0258)	1.7415	0.3095
Third	Volatility	I(1), I(1)	2	0.7435	2.8663	0.8768	(0.0162)	1.9943	0.1849
	Cost	I(1), I(1)	3	(0.5206)	1.1596	0.6077	(0.0143)	2.3747	0.3587

JPY

Period		Scenario	Lag	Skewness	Kurtosis	Anderson-Darling	Auto-correlation	Multi-collinearity	Breusch-Pagan (p)
Full	Volatility	I(1), I(0)	3	2.9328	24.5906	29.8462	(0.0148)	8.5959	<.0001
	Cost	I(1), I(0)	3	6.9929	140.6302	47.4648	(0.0381)	8.5959	0.4115
	Volatility	I(1), I(1)	3	1.8478	15.0430	20.5891	(0.0297)	2.0904	0.0788
	Cost	I(1), I(1)	4	7.4419	147.8377	47.0817	(0.0211)	2.5637	0.1727
First	Volatility	I(0) , I(0)	4	2.8765	19.8651	15.7744	(0.0265)	24.2618	<.0001
	Cost	I(0) , I(0)	2	0.7387	2.6388	1.4803	(0.0250)	16.8154	0.1886
	Volatility	I(0) , I(1)	4	2.1363	13.7985	10.7255	(0.0006)	17.9711	0.0398
	Cost	I(0) , I(1)	2	0.7470	2.5911	1.5518	(0.0309)	13.0416	0.1045
Second	Volatility	I(1), I(0)	2	1.2013	3.0546	5.8215	(0.0250)	9.3981	0.0579
	Cost	I(1), I(0)	3	0.0921	1.1678	0.9965	(0.0439)	11.2801	0.7885
Third	Volatility	I(0) , I(1)	2	0.5733	2.9605	1.7333	0.0075	11.7215	0.3304
	Cost	I(0) , I(1)	1	(0.3984)	0.2055	0.5979	(0.0762)	9.1713	0.5147
	Volatility	I(1), I(1)	5	0.5471	2.1908	1.1652	0.0128	3.6143	0.7677
	Cost	I(1), I(1)	4	0.0010	0.5588	0.3352	(0.0082)	3.0696	0.2767

EUR

Period		Scenario	Lag	Skewness	Kurtosis	Anderson-Darling	Auto-correlation	Multi-collinearity	Breusch-Pagan (p)
Full	Volatility	I(1), I(0)	2	1.1531	7.5170	8.9773	(0.0569)	8.9810	<.0001
	Cost	I(1), I(0)	4	1.0937	9.6543	7.3275	(0.0169)	11.9948	0.3642
	Volatility	I(1), I(1)	3	0.3882	6.4273	6.0127	(0.0358)	2.2099	0.0168
	Cost	I(1), I(1)	4	1.0910	9.6437	7.2637	(0.0158)	2.7061	0.3536
First	Volatility	I(0) , I(0)	4	0.9686	1.3715	2.5497	(0.0073)	24.2923	0.0914
	Cost	I(0) , I(0)	1	1.5791	6.5550	2.8333	(0.0546)	12.2946	0.4515
	Volatility	I(0) , I(1)	3	0.5243	0.7396	0.8440	(0.0292)	13.2008	0.6941
	Cost	I(0) , I(1)	3	1.4645	6.2176	2.2979	0.0047	13.2008	0.5448
Second	Volatility	I(1), I(0)	3	0.3575	0.3499	1.0526	(0.0335)	13.4300	0.0375
	Cost	I(1), I(0)	4	(0.1457)	0.7167	0.6196	(0.0302)	15.3052	0.4209
Third	Volatility	I(1), I(1)	1	0.4268	5.3790	2.4466	(0.0832)	1.1384	0.0081
	Cost	I(1), I(1)	2	0.2437	0.9387	0.3749	(0.0495)	2.0875	0.1282

Appendix 8

	Naïve GBP				Naïve CHF				Naïve JPY				Naïve EUR			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1)	I(0)	I(0)	I(1)	I(1)	I(0)	I(0)	I(1)	I(1)	I(0)	I(0)	I(1)	I(1)	I(0)	I(1)	I(1)
Intercept	0.0000	-	-	0.0000	0.0000	-	-	0.0001	0.0000	-	-	0.0000	0.0000	-	0.0000	0.0000
	(-0.04)	-	-	*(0.10)*	*(-0.07)*	-	-	*(0.39)*	*(0.02)*	-	-	*(-0.00)*	*(-0.06)*	-	*(0.15)*	*(0.16)*
Volatility lag1	-0.4660	0.9530	0.9731	-0.5352	-0.4507	0.9585	0.9690	-0.5731	-0.3899	0.9319	0.9682	-0.4449	-0.4668	0.9654	-0.4270	-0.5858
	*(-17.41***)*	(74.54***)	*(85.95***)*	*(-6.78***)*	*(-16.78***)*	(57.25***)	*(103.87***)*	*(-7.64***)*	*(-14.15***)*	(61.79***)	*(75.32***)*	*(-5.52***)*	*(-13.11***)*	(55.23***)	*(-8.06***)*	*(-7.99***)*
R^2	0.2169	-	-	0.2858	0.2045	-	-	0.3293	0.1527	-	-	0.1996	0.2174	-	0.1811	0.3435
N	1097	552	419	117	1097	272	695	121	1097	567	399	124	621	197	296	124
Autocorrelation	-0.0752	-0.4002	-0.4846	-0.0035	-0.1117	-0.4188	-0.3821	-0.0035	-0.0915	-0.3672	-0.3821	-0.1341	-0.1012	-0.2959	-0.1338	-0.0823
Breusch-Pagan	0.0158	*<.0001*	0.0096	0.6776	0.0180	*<.0001*	*<.0001*	0.3391	0.0627	*<.0001*	*<.0001*	0.8778	0.0011	*0.0001*	0.1992	0.8114
Condition index	1.0005	1.0000	1.0000	1.0035	1.0028	1.0000	1.0000	1.0141	1.0020	1.0000	1.0000	1.0056	1.0009	1.0000	1.0018	1.0016
Anderson-Darling	15.8111	5.4598	1.8015	1.4133	9.3460	2.1789	1.8765	3.1676	20.3322	10.6546	2.5472	1.7531	6.1097	0.8105	0.2790	2.5257
Skewness	1.2678	0.4000	0.1076	0.6780	0.9083	0.2137	0.2756	1.6831	1.5733	1.1363	0.2810	0.3656	0.2057	0.0811	0.1714	0.4721
Kurtosis	8.6465	4.4149	2.0131	2.9104	6.5822	2.8081	1.6772	9.1728	16.1582	16.1506	1.7197	3.9799	6.5587	1.6221	0.4232	5.1795

Appendix 9

	Spill-over GBP				**Spill-over CHF**			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1)	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1), I(1)
Intercept	0.0000	-	0.0006	0.0000	0.0000	-	0.0000	0.0000
	(-0.94)	-	(3.11***)	(-0.16)	(-0.03)	-	(0.18)	(0.20)
Cost	1.5514	0.5164	-0.7420	-	0.0733	2.9202	-0.2209	-
	(2.85***)	(0.83)	(-0.86)	-	(0.13)	(3.31***)	(-0.33)	-
Cost lag1	-	-	-	-	-	-	-	-0.7659
	-	-	-	-	-	-	-	(-0.61)
Cost lag2	-	-	-	-7.9411	-	-	-	-
	-	-	-	(-1.42)	-	-	-	-
Volatility lag1	-0.4098	0.2366	0.3058	-0.3147	-0.3920	0.0950	0.2358	-0.3021
	(-15.95***)	(7.02***)	(6.63***)	(-4.66***)	(-15.80***)	(2.23**)	(8.02***)	(-4.49***)
Volatility lag2	-0.1706	0.1100	0.1961	-	-0.2211	0.1081	0.1475	-0.2915
	(-6.26***)	(3.35***)	(4.96***)	-	(-8.32***)	(2.74***)	(5.16***)	(-4.16***)
Volatility lag3	-0.0936	-	-	-	-0.1075	-	-	-0.1157
	(-3.90***)	-	-	-	(-4.59***)	-	-	(-1.96**)
Delta Vix	0.0000	0.0000	0.0000	-	0.0000	0.0001	0.0001	0.0000
	(-1.49)	(0.81)	(0.63)	-	(-1.49)	(1.58)	(4.77***)	(-0.56)
Outlier-dummy	0.0017	-	-	-	-	-	-	-
	(7.23***)	-	-	-	-	-	-	-
GBP-Volatility	-	-	-	-	0.4747	0.6028	0.5539	0.0244
	-	-	-	-	(18.73***)	(15.80***)	(15.24***)	(0.33)
CHF-Volatility	0.3952	0.5448	0.2646	-0.0854	-	-	-	-
	(17.52***)	(17.86***)	(7.18***)	(-0.76)	-	-	-	-
JPY-Volatility	0.0848	-0.0571	0.0158	0.2483	0.1879	0.1661	0.1857	0.0102
	(4.66***)	(-2.97***)	(0.54)	(4.19***)	(10.00***)	(4.45***)	(10.75***)	(0.18)
EUR-Volatility	-	-	0.1949	0.4490	-	-	-	0.7613
	-	-	(5.22***)	(4.12***)	-	-	-	(10.25***)
R^2	0.5296	-	0.5565	0.5262	0.5606	-	0.6163	0.7826
N	1095	551	418	113	1095	271	694	119
Autocorrelation	-0.1268	0.1539	-0.1332	-0.1294	-0.1503	0.1731	0.0902	-0.2233
Breusch-Pagan	<.0001	<.0001	0.0289	0.3178	<.0001	<.0001	<.0001	<.0001
Condition index	2.3126	10.9639	16.3348	3.1501	2.3571	12.6356	13.3058	3.7706
Anderson-Darling	7.9514	5.4774	2.4509	0.5152	4.7797	1.9036	1.5059	1.3797
Skewness	0.1733	1.0684	0.7517	0.3616	0.4839	0.8013	0.1910	0.7418
Kurtosis	5.1732	3.0538	1.4966	0.9541	3.5673	2.2040	0.3264	4.6025

	Spill-over JPY				Spill-over EUR			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1), I(1)	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(1), I(1)	I(1) , I(1), I(1)
Intercept	-0.0001	-	0.0005	0.0000	0.0000	-	-0.0015	0.0000
	(-1.24)	-	*(2.26**)*	*(-0.10)*	*(0.06)*	-	*(-7.80***)*	*(0.27)*
Cost	-	0.3580	-1.9876	-	-0.8003	-0.0852	-1.9223	-3.3655
	-	*(0.43)*	*(-2.04**)*	-	*(-1.39)*	*(-0.16)*	*(-1.50)*	*(-1.24)*
Cost lag1	0.8570	-	-	1.4698	-	-	-	-
	(1.81)*	-	-	*(0.48)*	-	-	-	-
Cost lag2	-	-	-	-	-	-	-	-
	-	-	-	-	-	-	-	-
Volatility lag1	-0.4725	0.4051	0.3750	-0.4890	-0.1738	0.1244	-0.6408	-0.1793
	*(-19.20***)*	*(10.72***)*	*(7.83***)*	*(-7.18***)*	*(-6.67***)*	*(3.42***)*	*(-12.59***)*	*(-3.75***)*
Volatility lag2	-0.2854	0.1776	0.1635	-0.3010	-0.0872	-	-0.4166	-
	*(-10.99***)*	*(4.74***)*	*(3.65***)*	*(-4.47***)*	*(-3.31***)*	-	*(-7.44***)*	-
Volatility lag3	-0.1704	-	-	-	-0.0701	-0.0327	-0.2029	-
	*(-7.17***)*	-	-	-	*(-3.07***)*	*(-1.04)*	*(-4.00***)*	-
Delta Vix	0.0002	0.0002	0.0001	0.0002	0.0000	0.0000	0.0001	0.0000
	*(8.04***)*	*(3.75***)*	*(3.03***)*	*(4.89***)*	*(1.85*)*	*(0.18)*	*(2.45**)*	*(0.91)*
Outlier-dummy	0.0073	-	-	-	-	-	-	-
	*(14.87***)*	-	-	-	-	-	-	-
GBP-Volatility	0.0730	-0.1106	0.1076	0.3725	0.2152	0.4523	0.0606	0.2272
	*(1.99**)*	(-1.68*)	(1.75*)	*(3.97***)*	*(7.98***)*	(7.50***)	(0.99)	*(3.75***)*
CHF-Volatility	0.3136	0.4699	0.2303	0.0176	0.6208	0.6502	0.2554	0.5826
	*(9.44***)*	(6.94***)	(4.66***)	*(0.15)*	*(24.49***)*	(13.24***)	(5.30***)	*(10.78***)*
JPY-Volatility	-	-	-	-	0.0326	-0.0407	0.0088	0.0311
	-	-	-	-	*(1.60)*	(-1.55)	(0.23)	*(0.63)*
EUR-Volatility	-	-	0.1380	0.1604	-	-	-	-
	-	-	*(2.92***)*	*(1.18)*	-	-	-	-
R^2	0.5117	-	0.4847	0.5929	0.7735	-	0.4620	0.8057
N	1095	566	393	123	619	193	294	124
Autocorrelation	-0.0782	-0.0782	0.0105	-0.2188	-0.2700	0.2201	0.1921	-0.2616
Breusch-Pagan	<.0001	<.0001	0.1776	0.1889	0.0301	<.0001	0.4244	0.6006
Condition index	2.2824	14.7094	16.8354	4.5228	2.9333	19.0829	15.1475	2.7330
Anderson-Darling	7.7824	11.3068	2.1516	0.2887	7.8271	0.2862	1.1022	1.3177
Skewness	0.5220	2.6480	0.5286	0.1389	-0.1338	0.1893	-0.2766	-0.2778
Kurtosis	2.2577	19.6163	0.9918	1.1375	4.1036	0.6127	0.1462	0.9756

Appendix 10

	LN GBP				LN CHF			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1)	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1)
Intercept	-0.0043	-	-1.9476	0.0006	0.0000	-	-1.6627	-0.0354
	(-0.59)	-	*(-7.33***)*	*(0.02)*	*(-0.01)*	-	*(-9.33***)*	*(-0.88)*
Cost	-0.0110	0.1895	-0.0550	-	-0.1336	0.2414	-0.0443	-
	(-0.30)	*(7.83***)*	*(-1.11)*	-	*(-3.71***)*	*(6.82***)*	*(-1.03)*	-
Cost lag1	-	-	-	-	-	-	-	-0.0792
	-	-	-	-	-	-	-	*(-0.68)*
Cost lag2	-	-	-	-0.0235	-	-	-	-
	-	-	-	*(-0.16)*	-	-	-	-
Volatility lag1	-0.5979	0.4607	0.3205	-0.6153	-0.6203	0.3740	0.4504	-0.1413
	*(-20.23***)*	*(11.22***)*	*(6.90***)*	*(-8.33***)*	*(-21.14***)*	*(6.45***)*	*(12.47***)*	*(-1.45)*
Volatility lag2	-0.2790	0.2573	0.3317	-	-0.3586	0.2519	0.2398	0.0362
	*(-8.41***)*	*(6.33***)*	*(7.09***)*	-	*(-10.90***)*	*(4.47***)*	*(6.58***)*	*(0.38)*
Volatility lag3	-0.1550	-	-	-	-0.2115	-	-	0.0209
	*(-5.26***)*	-	-	-	*(-7.26***)*	-	-	*(0.22)*
Delta Vix	0.4244	0.5042	0.2256	-	0.4985	0.4229	0.5726	0.1412
	*(4.96***)*	*(4.05***)*	*(1.80*)*	-	*(5.51***)*	*(2.70***)*	*(5.89***)*	*(0.36)*
Outlier-dummy	0.2640	-	-	-	-	-	-	-
	*(4.70***)*	-	-	-	-	-	-	-
R^2	0.2915		0.3161	0.3807	0.3200		0.4002	0.0241
N	1095	557	418	116	1095	276	694	117
Autocorrelation	-0.0175	-0.0201	-0.0630	-0.0530	-0.0390	-0.0117	-0.0408	-0.0006
Breusch-Pagan	*<.0001*	0.0728	0.2196	0.0279	*<.0001*	0.9120	0.0415	0.9297
Condition index	2.1517	49.1684	56.2319	1.0263	2.1859	49.4240	54.7118	1.3904
Anderson-Darling	3.8691	1.4714	1.0008	1.7689	8.9532	0.6609	0.3473	7.9958
Skewness	-0.4932	0.4833	-0.0941	-0.8791	-2.3256	0.3947	0.0162	-5.8098
Kurtosis	6.0655	1.2996	2.0039	7.6067	43.6925	1.0110	0.4288	50.6895

	LN JPY				LN EUR			
Period	*Full*	*First*	*Second*	*Third*	*Full*	*First*	*Second*	*Third*
Model	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(0), I(1)	I(1) , I(1)	I(1) , I(1), I(1)	I(0) , I(0), I(1)	I(1) , I(1), I(1)	I(1) , I(1)
Intercept	-0.0043	-	-1.7463	-0.0011	0.0015	-	0.0014	0.0068
	(-0.56)	-	*(-7.11***)*	*(-0.03)*	*(0.16)*	-	*(0.13)*	*(0.20)*
Cost	-	0.1597	-0.0345	-	-0.0870	0.1591	-0.0307	-0.3061
	-	*(7.66***)*	*(-0.77)*	-	*(-2.32**)*	*(3.94***)*	*(-0.69)*	*(-2.64***)*
Cost lag1	0.0804	-	-	0.1374		-	-	-
	*(2.58***)*	-	-	*(1.35)*		-	-	-
Cost lag2	-	-	-	-	-	-	-	-
	-	-	-	-	-	-	-	-
Volatility lag1	-0.5218	0.5227	0.4269	-0.5756	-0.6572	0.5553	-0.6404	-0.5224
	*(-18.24***)*	*(13.08***)*	*(8.85***)*	*(-6.62***)*	*(-16.89***)*	*(9.23***)*	*(-11.12***)*	*(-7.00***)*
Volatility lag2	-0.2975	0.2325	0.2537	-0.3033	0.3966	-	-0.4142	-
	*(-9.63***)*	*(5.80***)*	*(5.23***)*	*(-3.54***)*	*(-9.11***)*	-	*(-6.51***)*	-
Volatility lag3	-0.1677	-	-	-	-0.2344	0.1974	-0.1592	-
	*(-5.90***)*	-	-	-	*(-6.09***)*	*(3.35***)*	*(-2.77***)*	-
Delta Vix	0.5914	0.5294	0.6086	0.7179	0.5294	0.4353	0.3989	0.3464
	*(6.45***)*	*(3.85***)*	*(4.80***)*	*(2.35**)*	*(4.66***)*	*(2.58***)*	*(2.79***)*	*(1.04)*
Outlier-dummy	0.6754	-	-	-	-	-	-	-
	*(7.50***)*	-	-	-	-	-	-	-
R^2	0.2850		0.3922	0.3121	0.3394		0.3175	0.3445
N	1095	570	393	123	619	193	294	124
Autocorrelation	-0.0284	-0.0424	-0.0211	-0.0557	-0.0445	-0.1049	-0.0275	-0.1592
Breusch-Pagan	*0.0020*	0.0276	0.3062	0.1382	*0.0658*	0.8614	0.2351	0.8460
Condition index	1.8851	43.0343	57.1063	1.7138	2.2378	53.5776	2.1723	1.1897
Anderson-Darling	3.1770	1.0591	0.7809	3.3011	5.4157	0.5232	0.2469	4.4856
Skewness	-0.2866	0.2746	-0.1516	-1.5109	-2.0612	0.1814	-0.1908	-1.2889
Kurtosis	6.8623	1.3325	0.3487	15.9246	27.2786	-0.1143	0.7871	15.7813

Appendix 11

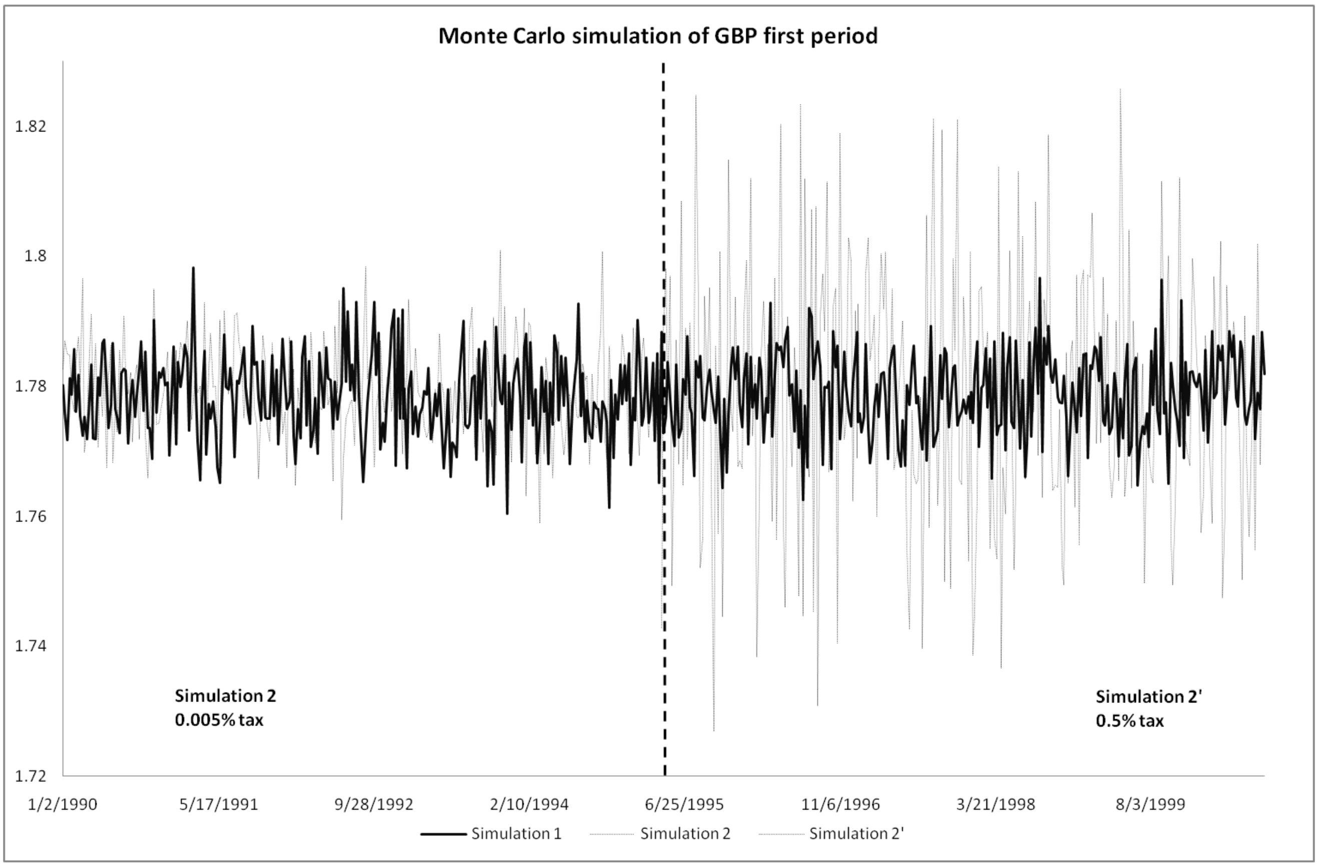
Monte Carlo simulation of GBP first period
1.82
1.8
1.78
1.76
1.74
1.72
Simulation 2
0.005% tax
Simulation 2'
0.5% tax
1/2/1990
5/17/1991
9/28/1992
2/10/1994
6/25/1995
11/6/1996
3/21/1998
8/3/1999
Simulation 1
Simulation 2
Simulation 2'

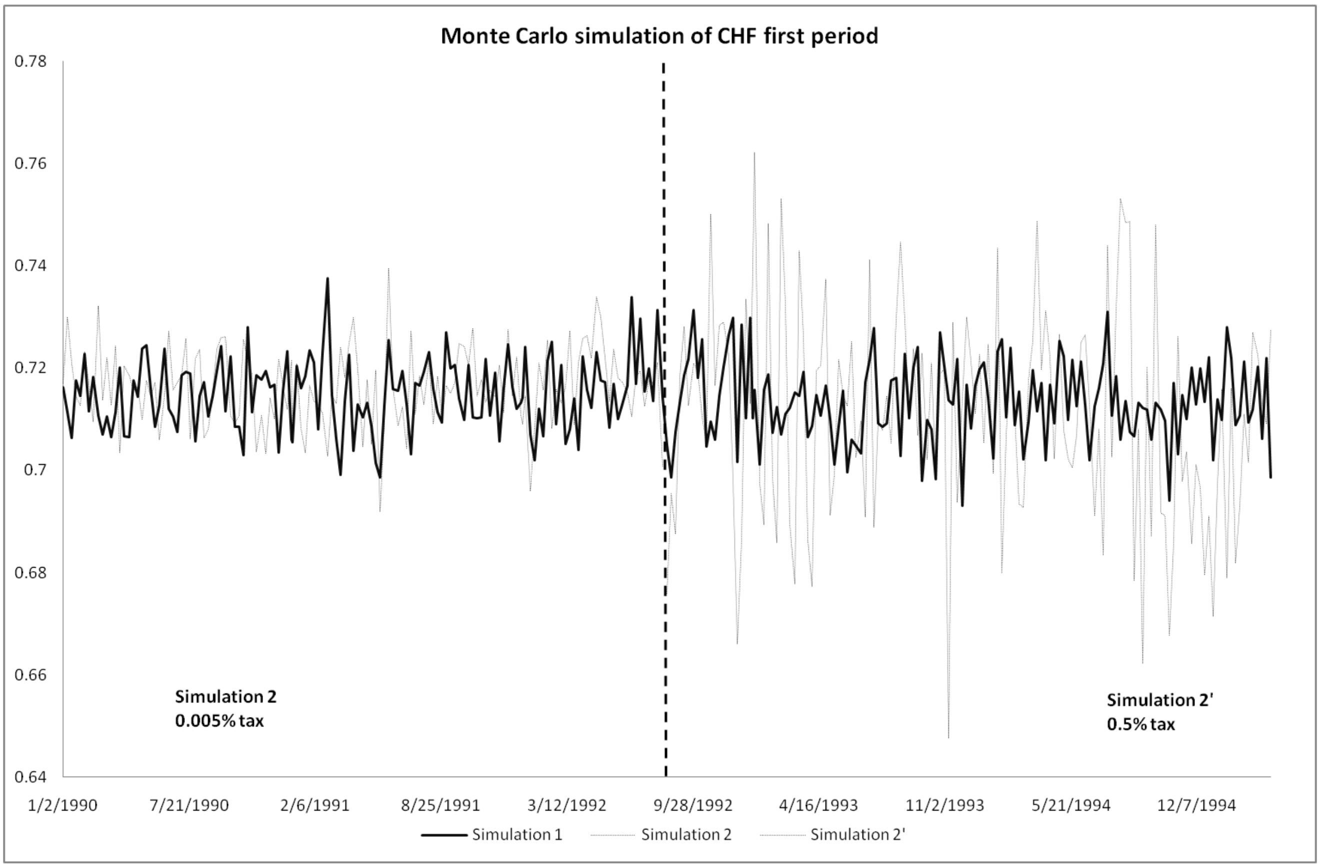
Monte Carlo simulation of CHF first period
0.78
0.76
0.74
0.72
0.7
0.68
0.66
0.64
Simulation 2
0.005% tax
Simulation 2'
0.5% tax
1/2/1990
7/21/1990
2/6/1991
8/25/1991
3/12/1992
9/28/1992
4/16/1993
11/2/1993
5/21/1994
12/7/1994
Simulation 1
Simulation 2
Simulation 2'

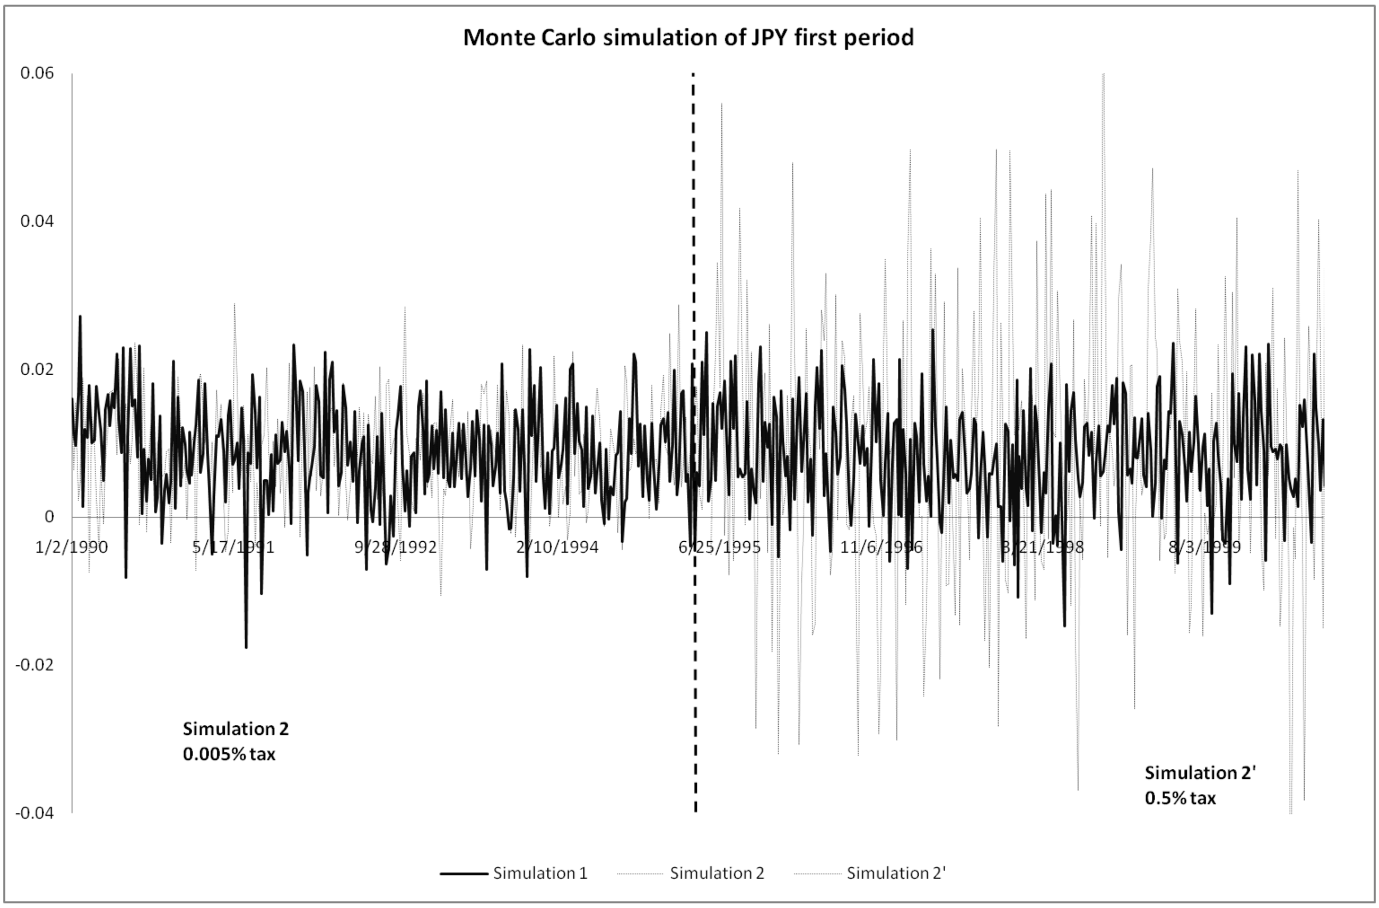
Monte Carlo simulation of JPY first period
0.06
0.04
0.02
0
-0.02
-0.04
1/2/1990
5/17/1991
9/28/1992
2/10/1994
6/25/1995
11/6/1996
3/21/1998
8/3/1999
Simulation 2
0.005% tax
Simulation 2'
0.5% tax
Simulation 1
Simulation 2
Simulation 2'

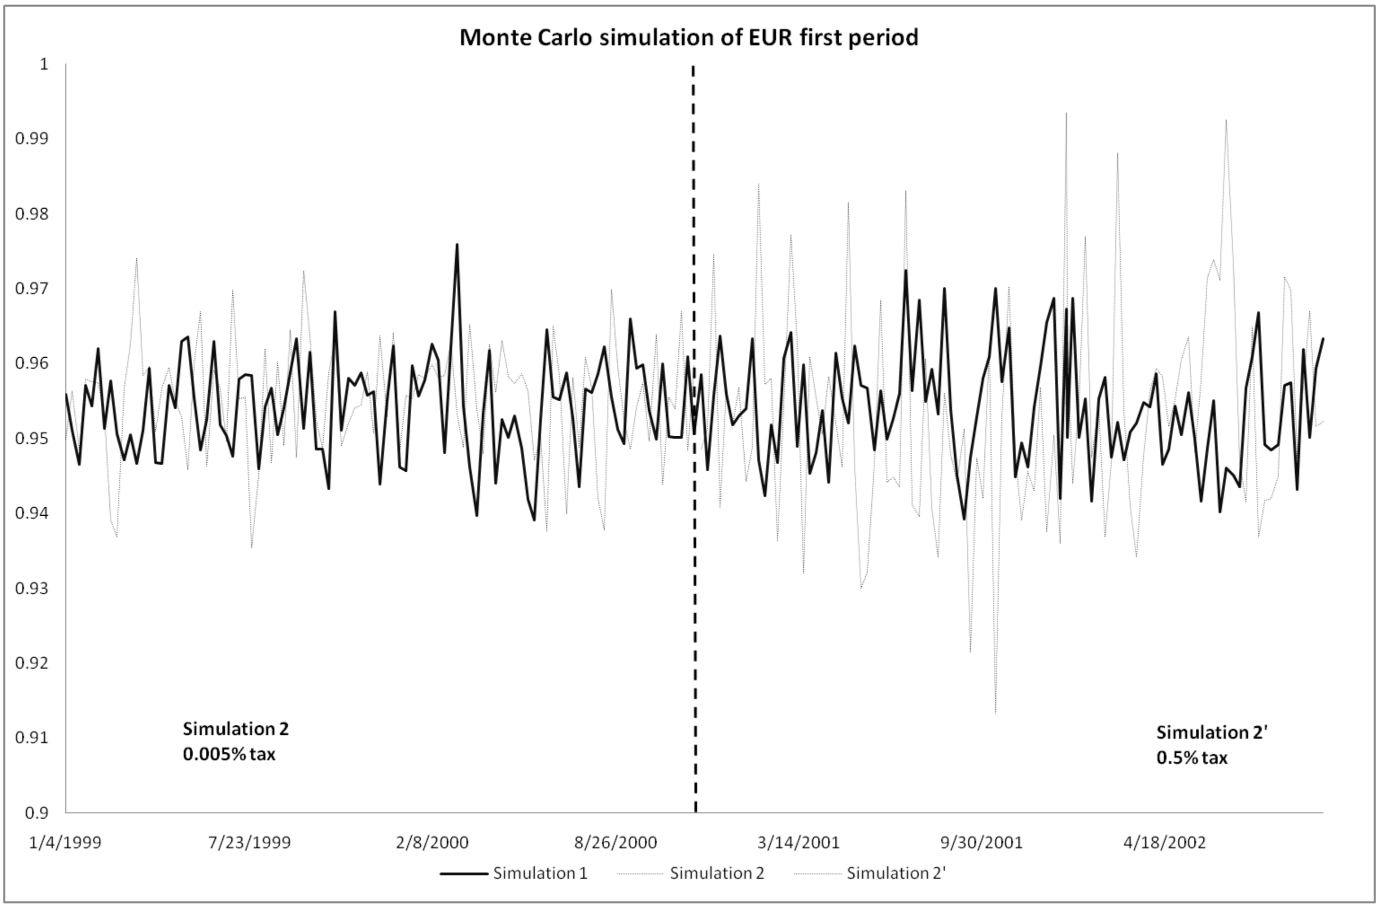
Monte Carlo simulation of EUR first period
1
0.99
0.98
0.97
0.96
0.95
0.94
0.93
0.92
0.91
0.9
1/4/1999
7/23/1999
2/8/2000
8/26/2000
3/14/2001
9/30/2001
4/18/2002
Simulation 2
0.005% tax
Simulation 2'
0.5% tax
Simulation 1
Simulation 2
Simulation 2'