Master's Thesis

Financial Economics

Volume and Volatility in the Icelandic Foreign Exchange Market

Sveinn Þórarinsson

HÁSKÓLI ÍSLANDS

University of Iceland
Department of Economics
Advisor: Helgi Tómasson
September 2009

Preface

This thesis is submitted in partial fulfillment of the requirements for the Degree of Master of Science in the Department of Economics at the University of Iceland, and is 30 ECTS credits. I would like to thank my advisor Helgi Tómasson for his guidance. I would also like to express my gratitude to the Central Bank of Iceland for its support.

Abstract

In this study we examine the relationship between trading volumes and volatility in the Icelandic foreign exchange interbank market using a data set spanning more than seven years. We review the theoretical and empirical literature on the relationship between trading volumes and volatility and discuss the econometric methodology employed. Theory predicts that trading volumes and volatility are positively correlated and a large number of empirical studies from a variety of different market settings have found this to be the case. The results of this study indicate that this is also holds true for the Icelandic interbank market. Furthermore, we find that the relationship is stronger in times of high volatility than low volatility.

Contents

1 Introduction **9**

2 Background & Concepts **11**
- 2.1 Econometric Methodology 11
  - 2.1.1 ARIMA Models 11
  - 2.1.2 GARCH 16
- 2.2 Icelandic Foreign Exchange Interbank Market 17

3 Literature Review **19**
- 3.1 The Mixture of Distributions Hypothesis 19
- 3.2 Empirical Studies 20

4 Analysis **23**
- 4.1 Inspection of Data & Preliminary Analysis 23
  - 4.1.1 The Dataset 23
  - 4.1.2 Exchange Rates, Trading Volumes & Volatility 24
- 4.2 The Relationship Between Volume and Volatility 27
  - 4.2.1 Modeling Trading Volume 27
  - 4.2.2 Modeling Exchange Rate Volatility 29
  - 4.2.3 Regression 29

5 Summary and Conclusions **33**

A Tables **35**

Bibliography **37**

1 Introduction

The goal of this study is to examine the relationship between exchange rate volatility and foreign exchange trading volumes. In particular it examines weather the same stylized facts concerning trading volumes and volatility apply for the foreign exchange interbank market of the very small open economy of Iceland. Furthermore we investigate weather this relationship changes during periods of significant turbulence in the foreign exchange market.

The foreign exchange market, both abroad and in Iceland, has seen tremendous growth over the last two decades. Research on exchange rate economics has also grown, and the field has seen a number of important developments as econometrics progresses and data becomes more available. However there are a number of notoriously elusive puzzles that have yet to be solved and perhaps in part due to that fact a new strand of literature, microstructure, has emerged which approaches these issues from a somewhat different perspective and is built on a different set of assumptions. Microstructure models often assume that market agents have heterogenous expectations and are more concerned with modeling the behavior of market agents and information flows than traditional macroeconomic models.

Such models provide the theoretical basis for the subject we investigate in this study, the relationship between volume and volatility. This subject has received considerable attention in the literature on financial markets from both a theoretical and an empirical point of view and is important for a number of reasons. Volatility is arguably the most important concept in finance and can be thought of as a measure of uncertainty or risk. High exchange rate volatility for a given period corresponds to a high degree of uncertainty regarding exchange rate levels for that period. Since one definition of a liquid market is where large transactions can be executed with a small impact on prices (or exchange rates), the relationship between volume and volatility can therefore be related to market liquidity. Furthermore it is seen as providing insight into the structure of financial markets by relating information arrival to market prices. For the economy we study in this paper, the foreign exchange rate is of special importance due to monetary policy implications and a high degree foreign currency denominated debt, and is as such an interesting topic of study.

The study is organized as follows. In Chapter 2 we discuss the econometric methodology used in this study and describe the institutional features of the market in question, the Icelandic foreign exchange interbank market. In Chapter 3 we review the literature on the relationship between trading volumes and volatility, focusing on the most influential models in the theoretical part as well as reviewing empirical studies. In Chapter 4 we describe the data set and conduct the analysis, inspecting the data before moving on to modeling volumes and volatility, and finally examining the relationship between the two. Chapter 5 concludes with a summary.

2 Background & Concepts

2.1 Econometric Methodology[1]

2.1.1 ARIMA Models

A time series is a collection of random variables $\{r_t\}$ where each random variable represents a value, such as the price of an asset, sampled at some frequency, e.g. daily. An ARIMA-class model is a linear time series model which is composed of an autoregressive part (AR) and an moving average part (MA). Before we move on to the details however, it is necessary to introduce the concepts which these models are based on. We end this section with a brief discussion on model selection.

Stationarity

Intuitively we can think of a strictly stationary process, $\{y_t\}$, as follows. If a given number, k, of points are chosen arbitrarily from the process then their joint probability distribution should stay the same if we shift all these points equally far in time. That is, the joint probability distribution of $(y_{t_1}, y_{t_2}, ..., y_{t_k})$ is the same as the joint probability distribution of $(y_{t_1+c}, y_{t_2+c}, ..., y_{t_k+c})$, for all c. We can state this more formally in the following way.

A Strictly Stationary Process

Denoting the cumulative distribution function by F, a process $\{y_t\}$ is strictly stationary if for any $t_1, t_2, ..., t_T$, T and c we have

$$F_{x_{t_1}, x_{t_2}, ..., x_{t_T}}(x_1, ..., x_T) = F_{x_{t_1+c}, x_{t_2+c}, ..., x_{t_T+c}}(x_1, ..., x_T). \tag{2.1}$$

In practice, the conditions for a strictly stationary series are hard to verify empirically. A slightly weaker condition is therefore often used.

[1]This section is largely based on Hayashi [23] chapter 6, Tsay [32] sections 2.6, 3.4, 3.5 and Brooks [8] chapters 5 and 8.

A Weakly Stationary Process

A process $\{y_t\}$ is weakly stationary if it satisfies

$$E[y_t] = \mu \tag{2.2}$$

$$Var[y_t] = \sigma^2 \tag{2.3}$$

$$Cov[y_t, y_{t-c}] = \gamma_c \tag{2.4}$$

This implies that a plot of the time series y_t as a function of time would show it fluctuating around a constant level, due to condition 2.2, and with a constant variation, due to 2.3. The third condition is also an important one as the auto-covariance function says how y_t relates linearly to, or co-varies with, its previous values. This condition says that for a stationary series the covariance between y_{t_1} and y_{t_2} depends only on the difference between t_1 and t_2, in particular the covariance of e.g. y_4 and y_{11} is equal to the covariance of y_{10} and y_{17}. Strict stationarity implies weak stationarity but not vice versa. If y_t is normally distributed however the two conditions are equivalent.

Autocorrelation Function

Since linear time series models try to capture linear dependence between y_t and its previous values the concept of serial correlations or autocorrelations play an important role in analyzing a stationary time series. Continuing our discussion from above, in order to work with auto-covariances, and to be able to compare the auto-covariances of different time series, it is convenient to make them independent of units of measurement. They are therefore normalized, by dividing by the variances, so that they take values in the range $[-1, 1]$. The result is the autocorrelation function.

Autocorrelation Function For A Weakly Stationary Series

Assuming a weakly stationary, or covariance stationary series the lag-c autocorrelation function for y_t is

$$\rho_c = Corr(y_t, y_{t-c}) = \frac{Cov(y_t, y_{t-c})}{\sqrt{Var(y_t)Var(y_{t-c})}} = \frac{\gamma_c}{\gamma_0}. \tag{2.5}$$

In practice, what we have is a finite sample $\{y_t\}_{t=0}^{t=T}$ and therefore it is necessary to define the sample counterpart of Equation 2.5. The lag-c sample autocorrelation function is defined as

$$\hat{\rho}_c = \frac{\sum_{t=c+1}^{T} (y_t - \bar{y})(y_{t-1} - \bar{y})}{\sum_{t=1}^{T} (y_t - \bar{y})^2}, \quad 0 \leq c < T - 1. \tag{2.6}$$

A correlogram, a plot of $\hat{\rho}_c$ as a function of $c = 1, 2, ...$, is used to visualize the sample autocorrelations. Another tool is the partial-autocorrelation function (PACF), which measures the correlation of y_t with a given lagged value y_{t-c} after removing the effects of intermediate values, i.e. $y_{t-1}, ..., y_{t-c+1}$. The PACF can be useful in distinguishing between an AR process and an ARMA process (discussed in more detail below).

A White Noise Process

A time series u_t is called a white noise process if $\{u_t\}$ is a collection of independent and identically distributed random variables. As a consequence of independence, the covariance of any two of the $\{u_t\}$ is zero and therefore the autocorrelation function is zero for all c (except $c = 0$).

AR Models

An autoregressive process of order p, or an AR(p) model, is given by

$$y_t = \mu + \sum_{i=1}^{p} \phi_i y_{t-i} + u_t, \tag{2.7}$$

where $\{u_t\}$ is a white noise process with mean zero and variance σ^2. In order for this to be an empirically workable model, a stationarity condition is assumed. It essentially requires that the lagged values of y_t have a decreasing effect on y_t, i.e. that shocks do not persist indefinitely in the system. To state this compactly[2] we define the Lag-operator as $L^j y_t = y_{t-j}$. We require a finite solution to

$$\phi(L) y_t = u_t, \tag{2.8}$$

where $\phi(L) = 1 - \phi_1 L - ... - \phi_p L^p$. In other words, the stationarity conditions are conditions under which the following solution exists

$$y_t = \phi^{-1}(L) u_t. \tag{2.9}$$

It can be shown[3] that this is satisfied if the roots of the characteristic equation for the polynomial $\phi(z)$, that is

$$\phi(z) = 1 - \phi_1 z - \phi_2 z^2 - ... - \phi_p z^p, \tag{2.10}$$

are all greater than 1 in absolute value.

---

[2] Without loss of generality we have set $\mu = 0$, in Equations 2.9 and 2.10.
[3] See e.g. Proposition 6.3 in Hayashi [23].

MA Models

A moving average process of order q, or an MA(q) model, is given by

$$y_t = \theta_0 + \sum_{i=1}^{q} \theta_i u_{t-i} + u_t, \tag{2.11}$$

where the u_t are from a white noise process. It can be easily verified that this is a weakly stationary process.

ARMA Models

An autoregressive moving average process of order (p,q) is a combination of an AR(p) process and an MA(q) process,

$$y_t = \mu + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{i=1}^{q} \theta_i u_{t-i} + u_t, \tag{2.12}$$

which satisfies the same stationarity condition as the AR(p) process defined in Equation 2.7. This can be written as

$$\phi(L) y_t = \mu + \theta(L) u_t, \tag{2.13}$$

where $\phi(L) = 1 - \phi_1 L - ... - \phi_p L^p$ and $\theta(L) = \theta_0 + \theta_1 L + ... + \theta_q L^q$ and the absolute values of the roots of $\phi(z)$ are all larger than 1.

ARIMA Models

An autoregressive integrated moving average model of order (p,d,q) is essentially an ARMA(p,q) model for y_t differenced d-times i.e. for $\Delta^d y_t$ where Δ is the difference operator, defined as $\Delta y_t = y_t - y_{t-1}$. We can thus write the ARIMA(p,d,q) in its general form as

$$\Delta^d y_t = \mu + \sum_{i=1}^{p} \phi_i \Delta^d y_{t-i} + \sum_{i=1}^{q} \theta_i u_{t-i} + u_t, \tag{2.14}$$

or, expressed in the compact form we defined above,

$$\phi(L) \Delta^d y_t = \mu + \theta(L) u_t. \tag{2.15}$$

As before, u_t are from a white noise process and the stationarity condition, that the roots of $\phi(z)$ are outside the unit circle, has to be satisfied. The qualitative features of the autocorrelation function and partial autocorrelation function for AR, MA and ARMA processes are presented in Table 2.1.

Model Selection

Box and Jenkins [7] suggest the following procedure for estimating ARMA models[4]

1. Identification.

2. Estimation.

3. Diagnostics checking.

The identification stage involves determining the values for (p,q). Box and Jenkins suggest graphical procedures such as plotting the autocorrelation function and partial-autocorrelation function. The resulting patterns can then be compared to the theoretical predictions, outlined in Table 2.1[5].

Table 2.1: Qualitative Features of the ACF and PACF function for AR, MA and ARMA processes.

	AR(p)	MA(q)	ARMA(p,q)
ACF	Tails off	Cuts off after lag q	Tails off
PACF	Cuts off after lag p	Tails off	Tails off

In practice however, these patterns may not be clearly discernible. Another technique is to calculate information criteria for each plausible parameter combination, and select the combination that yields the lowest value. Two of the most common information criteria are Aikaike's information criterion (AIC) and Schwarz's information criterion, defined as

$$AIC = ln(\hat{\sigma}^2) + \frac{2k}{T}, \tag{2.16}$$

$$SBIC = ln(\hat{\sigma}^2) + \frac{k}{T} lnT, \tag{2.17}$$

where $\hat{\sigma}^2$ is the residual variance (the residual sum of squares divided by the number of degrees of freedom, $T-k$), $k = p+q+1$ is the total number of parameters estimated and T the sample size. These expressions therefore give lower values for a better model fit (lower residual variance) while penalizing larger models (higher value of k). The difference between the two criteria is that SBIC has a stronger penalty than AIC.

The second step is to estimate the model that is selected. This can be done with Least Squares estimation or Maximum Likelihood. Box and Jenkins suggest two methods for the third step, the diagnostics checking. One is overfitting, that is fitting a larger model

[4]This also applies to ARIMA models since, as discussed above, they can be viewed as an ARMA model for $\Delta^d y_t$.

[5]This Table is adopted from Table 3.1 in Shumway and Stoffer [30].

than step one requires and then checking weather the extra terms are significant. The other method is to check the residuals for any remaining linear dependance. This can be done by examination of the autocorrelation function and partial-autocorrelation function as well as using the Ljung-Box test. The Ljung-Box test has the null hypothesis that the data are random, and the alternative hypothesis that they are not.

2.1.2 GARCH

GARCH stands for generalized autoregressive conditional heteroskedasticity and is a generalization of Engle's (1982) [16] ARCH model. These models draw their name from the fact that the conditional variance is allowed to change with time, i.e. is heteroskedastic, and is modeled as an autoregressive process of lagged squared error terms. These models try to capture the fact that volatility for financial time series changes over time. In other words, the second moment of the distribution of financial time series, such as exchange rate returns, is serially correlated. Furthermore, these models also capture[6] the fact that volatility often occurs in burst or clusters, i.e. large changes tend to be followed by large changes, and small changes tend to be followed by small changes. A third common property of exchange rate returns that these models reflect is the leptokurtosis, or heavy-tails, of their observed probability distribution. That is, the unconditional kurtosis of the error terms, u_t is larger than that of a random variable from a Gaussian white noise process.

Engle [16] formulates his ARCH model in the following general terms

$$y_t|\Omega_{t-1} \sim N(x_t\beta, h_t) \tag{2.18}$$
$$h_t = f(\varepsilon_{t-1}, \varepsilon_{t-2}, ..., \varepsilon_{t-p}, ..., \alpha) \tag{2.19}$$
$$\varepsilon = y_t - x_t\beta \tag{2.20}$$

where the conditional distribution of y_t given the information set Ω_{t-1} is Gaussian and the conditional variance h_t is a function of lagged error terms and other exogenous variables. A common way to express the ARCH(p) model is

$$y_t = x_t\beta + u_t, \quad u_t \sim N(0, h_t) \tag{2.21}$$
$$h_t = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2 u_{t-2}^2 + ... + \alpha_q u_{t-p}^2. \tag{2.22}$$

where x_t is a vector of exogenous variables and β a vector of coefficients. Non-negativity constraints, such as $\alpha_j \geq 0, \forall j$, are usually imposed in Equation 2.22 since a negative variance h_t would be meaningless.

[6] As can be seen in Equation 2.22

The generalization of this model, the GARCH model, is much more widely used in the literature however, mainly due to the fact that it is a more parsimonious model (Brooks [8]). The GARCH(p,q) model may be written as

$$y_t = x_t\beta + u_t, \quad u_t \sim N(0, h_t) \tag{2.23}$$

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j}, \tag{2.24}$$

from which we see that the difference between ARCH and GARCH is that in the latter the conditional variance, h_t, is allowed to depend on its own lagged values. This process for h_t is weakly stationary provided that the unconditional variance exists, which is easily shown to be the case when the sum of the coefficients in Equation 2.24 are strictly less than one, i.e. $\sum_{i=1}^{max(p,q)} (\alpha_i + \beta_i) < 1$. When this is not the case or the sum of the coefficients is very close to one, as is often reported in the literature (Sarno and Taylor (2001) [29]), an integrated model is used, or IGARCH.

To see why GARCH is a more parsimonious model than ARCH, suppose we start with a GARCH(1,1) model and substitute the lagged conditional variance, h_{t-1}, for the RHS of Equation 2.24 lagged once. By repeating this process indefinitely we get an infinite order ARCH model. This shows that although the GARCH model implicitly incorporates the influence from a very large number of past squared errors on the current conditional variance, estimation of only a few parameters, three in the GARCH(1,1) case, is required.

For a overview of the large literature on GARCH models see Bollerslev, Chou, and Kroner (1992) [4] and Bollerslev, Engle, and Nelson (1994) [6].

2.2 Icelandic Foreign Exchange Interbank Market[7]

Roots of the Market

The Icelandic interbank market for foreign exchange was established May 28, 1993. From then on the exchange rate was determined by supply and demand of market participants for foreign currency whereas before it had been decided by the Central Bank. In 1995 capital movements between Iceland and other European Economic Area countries were deregulated and two years later the market's opening hours were extended from only a few minutes each day[8] to the hours between 9:15 am and 16:00 pm. At the same time the trading volumes increased significantly, more than doubled between years, and the Central Bank's share of total volume decreased from around 80% in the

[7]This section is based on an article in CBI's Monetary Policy (2001) [9].

[8]So-called "fixing meetings" were held by market participants on each trading day in the Central Bank. Trading was conducted and the exchange rate subsequently fixed.

first four years of the market, down to less then 5% in 1999. From the beginning of the market and until 2001 a flexible fixed exchange rate policy was implemented by the Central Bank. This means the value of the exchange rate index was allowed to deviate from a fixed index level within a certain percentage limit, or deviation band. Should the exchange rate index move close to these limits the Central Bank would intervene - buy or sell currency in the interbank market - to maintain the exchange rate index within the deviation band. The first two years the boundaries were set at $\pm 2.25\%$ then $\pm 6\%$ and finally $\pm 9\%$ until finally on 28th March 2001 the fixed exchange rate regime was abandoned, inflation targeting adopted and the currency allowed to float.

Rules and Market Structure

The Icelandic interbank market is a closed market. A prerequisite for a financial institution to become a member in the market is that it holds a license to act as an intermediary in foreign exchange transactions and that it takes on market making and the associated duties. The duties of a market maker are to be prepared to quote a bid and ask price for foreign currency[9] if requested by another market maker. These quotes are made in a Reuters (or comparable) information system where market participants can observe each other's quotes and are required to be updated at least every 30 seconds. A market maker does not have to present quotes for amounts other than the reference amount of $1.5 million. Trades exceeding $500 thousand are to be reported to the Central Bank within five minutes of taking place. All trades are settled two days after they are made. The spread between bid and ask prices is fixed at 5 cents but can go up to 7 cents. Furthermore, if the exchange rate index changes more than 1.25% from the opening price in a given day the market makers are allowed to increase the spread to 10 cents and 20 cents if the index changes more than 2%. If the market maker has transacted with another market maker he is not required to quote him a binding offer in the next 5 minutes. The official exchange rate of the krona against foreign currencies is fixed between 10:45-11:00 by the Central Bank. It is done by taking an average of the bid and ask prices for the dollar[10] and other foreign currencies against the dollar and using this information the exchange rate of the krona against other foreign currencies is calculated.

At the time of the floating of the currency there were four market participants in the market, Bunadarbanki, Islandsbanki, Kaupthing and Landsbanki. However during the period Kaupthing and Bunadarbanki merged into one bank.

[9]Until December 1, 2006 this currency was the US dollar but was changed then to the euro.
[10]The euro after 1 December, 2006.

3 Literature Review

3.1 The Mixture of Distributions Hypothesis

Theories addressing the volatility-volume relationship date back to at least as early as 1973. Clark [14] proposes the "mixture of distributions hypothesis" (MDH) in which price variability and trading volume are driven by an unobserved common directing variable. This variable reflects the arrival of new information to the market. When there are numerous information shocks to the market, both volume and volatility increase.

Clark's explanation however is secondary to his efforts to explain the empirical fact that the distribution of changes in speculative prices is not Gaussian but rather leptokurtic. This means, in statistical terms, that the price change sample kurtosis is greater than that of the normal distribution. Compared to the normal distribution there are too many small and too many large observations or, put differently, the distribution of price changes has heavier tails and a higher peak around the mean than the normal distribution.

Clark points out that this implies that conditions sufficient for the Central Limit Theorem, which says that the sum of a large number of random variables is approximately normal, are not met by the price change data generating process. He hypothesizes that the number of individual effects added to together to give the price change during a day is random, making the Central Limit Theorem inapplicable. The intuition behind assuming a random number of within-day price changes is the fact that information is available to traders at a varying rate. On days when no new information is available trading is slow and the price process evolves slowly but when new information changes traders' expectations, trading occurs more frequently and the price process evolves faster.

To take account of this fact, Clark generalizes the Central Limit Theorem to allow for a random number of random variables, each representing a within-day price change, and derives the limiting distribution for the sum of those random variables. The variance of this limiting distribution, i.e. the variance of the daily price change, becomes a random variable with a mean proportional to the mean number of daily transactions. He argues that trading volume is related positively to the number of within-day transactions, and thus also to volatility.

Bauwens et al. [1] summarize Clark's MDH in the following way:

$$\Delta P_t = \sum_{i=1}^{N(t)} \Delta P_i, \quad i = 1, ..., N(t), \quad P_0 = P_{N(t-1)}, \tag{3.1}$$

$$\{\Delta P_i\} \quad i.i.d., \qquad \Delta P_i \sim N(0,1), \tag{3.2}$$

$$\frac{E[N(t) \| \nu_t]}{\partial \nu_t} > 0. \tag{3.3}$$

Here the first line (3.1) states that the price increment of period t equals the sum of intra-period increments, line (3.2) is a random walk hypothesis, i.e. that the price increments are independently and identically (normally) distributed, and (3.3) says that the mean of the number of intra-period increments $N(t)$ conditioned on the number of information events ν_t in period t is strictly increasing with ν_t.

A different approach is taken by Epps and Epps [17]. They focus on the microstructure of financial markets and construct a two-parameter portfolio model of traders maximizing utility. In the Epps and Epps model, when new information arrives to the market it affects the reservation prices of different traders in different ways. Consequently, the change in the market price for each transaction is the average of all the changes of traders' reservation prices. A central assumption they make is that there is a positive relationship between the extent to which traders disagree when they revise their reservation prices and the absolute value of the change in the market price. In other words, the more heterogenous beliefs traders have when they respond to information arrival, the more absolute price change occurs.

Tauchen and Pitts [31] extend Clark's MDH in an influential study and derive a more general model of the price change and trading volume on speculative markets than Epps and Epps [17]. Similarly to their model Tauchen and Pitts begin with an equilibrium theory of within-day price determination. Unlike Epps and Epps however they derive an explicit expression for the joint probability distribution of the daily price change and the trading volume. Furthermore, their model allows for the variance of the daily price change and the mean daily trading volume to depend upon three different factors: the average daily rate at which new information flows to the market, the extent to which traders disagree when they respond to new information, and the number of active traders in the market.

3.2 Empirical Studies

A number of empirical studies have examined the relationship between volatility and trading volume. Most of them have found that they exhibit strong contemporanious correlations.

Karpoff [26] provides an overview of nineteen empirical studies from the early literature that examine the volume-volatility relationship. These studies come from a variety of different market settings, such as futures, equity and foreign exchange, although most of the research has focused on stock and future markets due to the fact that data is more easily available for them than for foreign exchange. Indeed, only one of these nineteen studies examines the foreign exchange markets. According to Bauwens et al. [1] the increased availability of data in the nineties changed this to some degree, and in their paper they list ten studies that directly or indirectly investigate the relationship in foreign exchange markets[1].

Galati [18] attributes the lack of data to the fact that, unlike equity markets, foreign exchange markets are for the most part decentralized. He mentions that the most comprehensive source of information in foreign exchange market, the "Central Bank Survey of Foreign Exchange and Derivatives Market Activity", does not provide much time series data for trading volumes. Therefore empirical studies of the foreign exchange volume-volatility relationship have resorted to using various alternative data sources to proxy for foreign exchange trading volumes.

A number of studies have used data on futures contracts to proxy for interbank trading volumes, including Grammatikos and Saunders [21] and Jorion [25]. However, a drawback of using data from the futures market in foreign exchange is that it represents a very small part of the total foreign exchange market and furthermore the futures market has different institutional features than the interbank market (Dumas [15]) .

Bollerslev and Domowitz [5] and Goodhart and Figliuoli [19] use the frequency of indicative quotes as a proxy for trading volume. One possible difficulty with using this data to proxy volumes however is that indicative quotes do not represent actual trades and banks may have programs in place to automatically post quotes at regular time intervals. Frequency of indicative quotes therefore may not represent trading volume accurately (Galati [18]). High frequency data, i.e. data on individual transactions within each day as opposed to a daily aggregate, have also been examined by e.g. Lyons [28] and Goodhart et al [20].

Comprehensive foreign exchange volume data sets are analyzed by Bjonnes et al. [3] and Galati [18]. The data used by Bjonnes et al. is from the Swedish central bank (Riksbank) and covers 90-95% of all daily worldwide trading of the Swedish krona, and Galati examines daily trading volumes of seven currencies from emerging market countries against the dollar. Both studies find a positive correlation between unexpected volume (a measure of information arrival discussed below) and volatility, although Galati finds that correlation between trading volumes and volatility is positive during "normal" periods but turns negative when volatility increases sharply.

[1] Table 1 in their paper

4 Analysis

4.1 Inspection of Data & Preliminary Analysis

4.1.1 The Dataset

We use a publicly available spot trading volumes data set provided by the Central Bank of Iceland from the foreign exchange interbank market (see [12]). The sample period is chosen to begin at the floating of the Icelandic krona, 28/03/2001, and to end when the banking system, and the interbank market, collapsed on the 3/10/2008[1]. This period is selected since it covers the entire time Iceland has had a floating exchange rate. Furthermore it simplifies the analysis somewhat not having to control for different exchange rate regimes[2].

The official exchange rates data for the krona is also publicly available [13]. However, the exchange rates are calculated each day at 10:45-11:00 am as the average of the current bid and ask offers on the market. This presented a potential problem for our analysis since this means there is a temporal mismatch in the measurements of exchange rates and daily spot trading volumes; the volumes are aggregated notional values of spot transactions that occur during the whole day, not just until 11:00 am. This means that should both exchange rates and volumes change significantly after 11:00 am on a given day, change would only be reflected in the volume data for that day, but not the exchange rate. To avoid this problem we decided not to use the official exchange rate, but to use closing prices from Reuters instead. Almost all of the trading until the December 1, 2006 was conducted in USD, but changed to EUR after that. Therefore, data on daily trading volumes corresponds to trading in USD/ISK during the period 28/03/2001-01/12/2006 and trading in EUR/ISK during 01/12/2006-03/10/2008.

To reflect this change we use the USD/ISK exchange rates for the first part of the period and the EUR/ISK for the second part, and combine these two series into one. In order to avoid a large jump in returns on the date when they are conjoined, we scale the USD/ISK exchange rate series by the USD/EUR rate on that date (i.e. USD/ISK

[1]A temporary auction market for foreign exchange was established 15/10/2008 and the foreign exchange interbank market was re-established 04/12/2008 ([10] and [11]).

[2]Bauwens et al. (2005) [1] find however that the volume-volatility relationship is relatively stable across three different exchange rate regimes.

x 1.323365). A similar approach is taken by Bjonnes et al (2005) [3] who construct a single price series from SEK/DEM and SEK/EUR.

All amounts are in billions of Icelandic kronas. The trading volumes are free of double counting issues since, after each trade, only one of the two parties transacting reports the trade to the Central Bank. Statistical analysis in done with R, version 2.9.1.

Data Quality and Error Checking

The data set was examined for possible errors. Due to the fact that the data on trading volumes is publicly available on the Central Bank's website it seems likely that it is of high quality, i.e. it should be mostly free of serious errors such as large outliers due to incorrectly entered data or many missing values. Out of the total of 1868 data points, six missing values were found for the volumes series. The data provider was contacted and subsequently the data set was corrected. A total of eleven days, all occurring in 2001 and 2002, had zero trading volume. This was double checked and turned out to have been the case, no trades were reported to the CBI on those days. To investigate if the exchange rate outliers present in the data were errors we compared the two data sources on exchange rates; daily changes in the Reuters closing prices to the changes in the official exchange rate, and found the differences to be small.

4.1.2 Exchange Rates, Trading Volumes & Volatility

We define the following notation.

$$e_t = (EUR/ISK)_t, \tag{4.1}$$

$$R_t = ln(e_t/e_{t-1}). \tag{4.2}$$

Here e_X denotes the nominal exchange rate for the Icelandic krona against the euro, i.e. $e_t = 100$ means 100 ISK are needed to acquire one EUR on date t. Continuously compounded returns are denoted by R_t. Daily trading volumes, i.e. the aggregate notional value of all foreign exchange spot trades in the interbank market are denoted by V_t and measured in billions of Icelandic kronas.

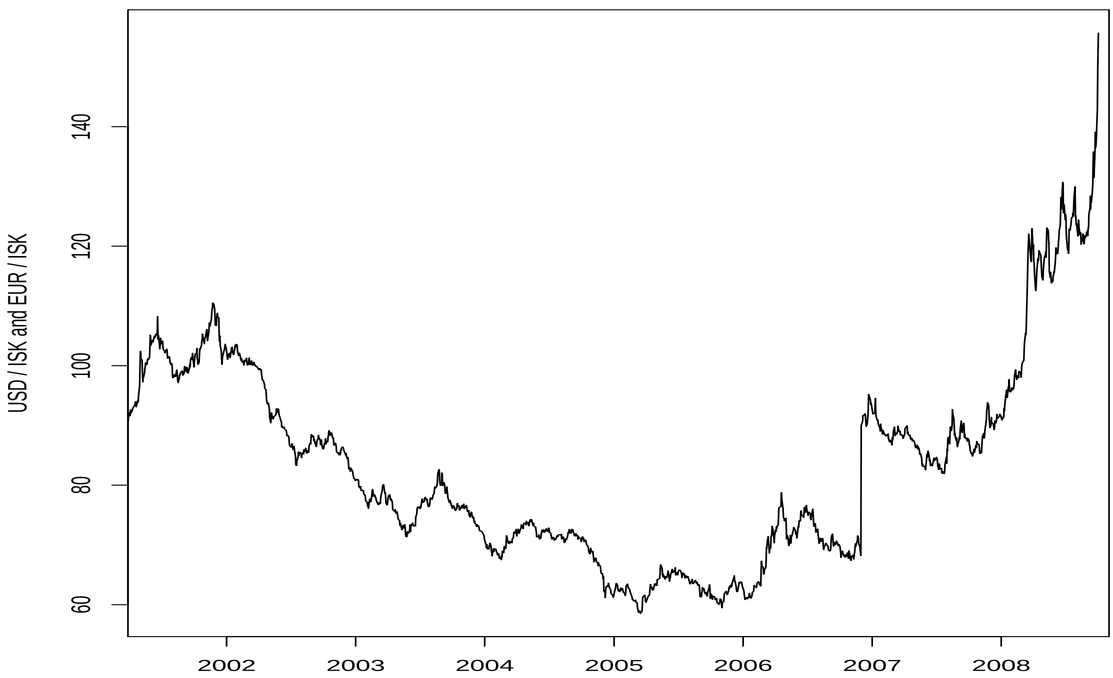

Figure 4.1: The Icelandic krona against the dollar (until 01/12/2006) and the euro (after that). (An increase represents a depreciation in the krona, and decrease an appreciation. Same axis applies for both series)

Figure 4.1 depicts the development of the exchange rate e_t from the floating of the currency until the end of the sample period. In the months following the floating it depreciated significantly, and near the end of November later that year it had lost close to 22% of its value against the dollar. From that point forward however, it began appreciating and in the spring of 2002, the krona had regained its value. A long period of gradual appreciation followed; the next four years, 2002-2006, saw the krona appreciate against the dollar by over 40%. That period of appreciation came to a halt however in the first quarter of 2006, when a number of reports critical of the Icelandic banking system were published, and in less then three months the krona had lost nearly a third of its value. It regained some of its strength over the next year but was hit again at the onset of the global credit crisis in August 2007. Turbulence in foreign exchange markets increased in the first quarter of 2008 and the krona depreciated by 35% in three months. It remained relatively stable throughout the summer but plunged in September around the time when the U.S. bank Lehman Brothers went bankrupt and two weeks later when the Icelandic government announced that it intended to acquire a majority share in one of the three Icelandic banks that formed the foreign exchange interbank market, Glitnir. In the first ten months of 2008, the krona had depreciated by close to 70%.

Summary statistics are presented in Table 4.1.

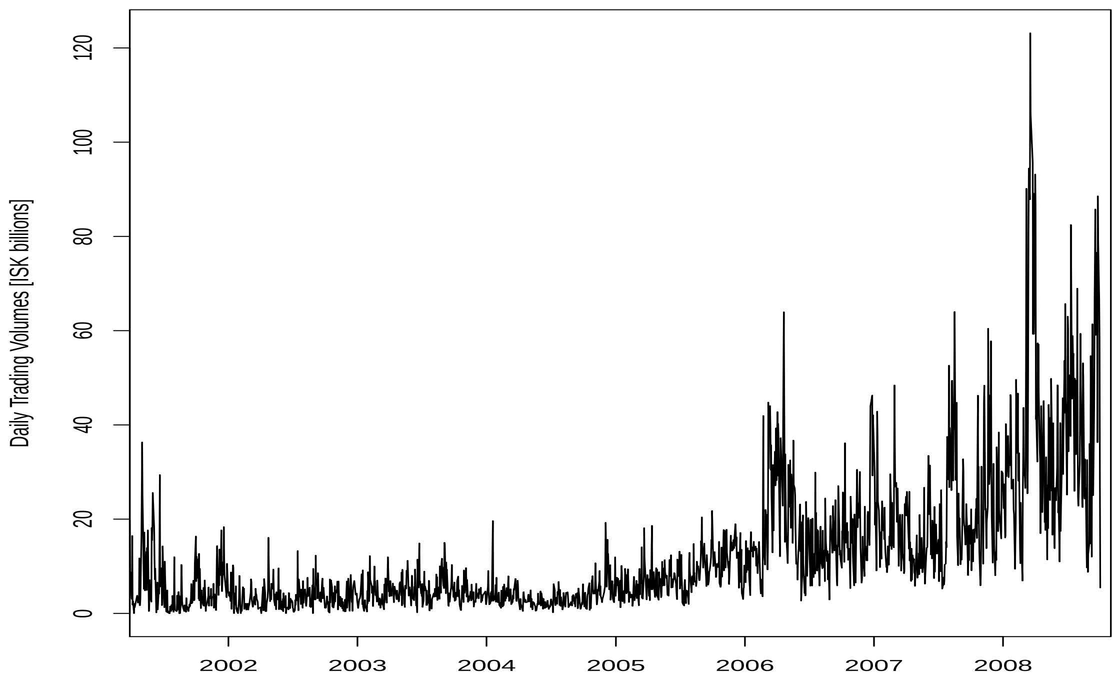

Figure 4.2: Daily spot trading volumes, V_t from the Icelandic interbank foreign exchange market. Period: 28/03/2001-03/10/2008.

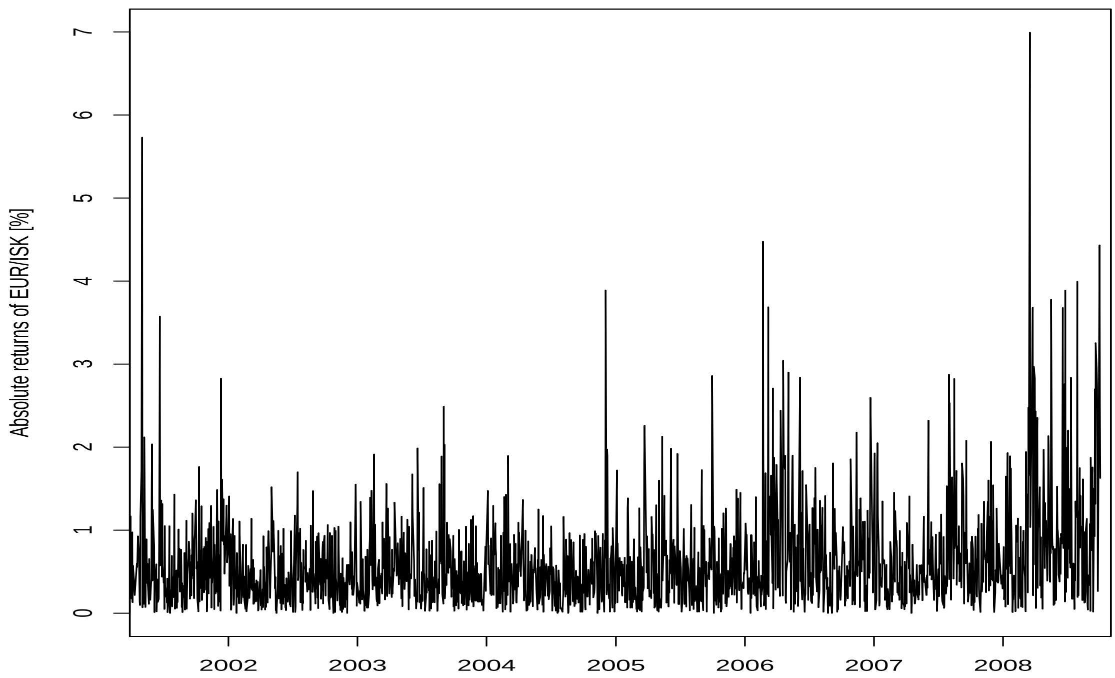

Figure 4.3: Absolute continuously compounded exchange rate returns, $|R_t|$, a measure of volatility.

Figure 4.2 depicts the development of daily spot trading volumes in the Icelandic interbank market for foreign exchange.

Table 4.1: Summary statistics of daily spot volumes, exchange rates, and exchange rate returns. Period: 28/03/2001-03/10/2008.

	Mean	Std.Deviation	Skewness	Excess Kurtosis	Minimum	Maximum
USD	75.72	13.1	0.9	-0.2	58.5	112.7
EUR	89.05	10.9	2.3	6.4	71.3	155.7
Volumes	12.29	14.2	2.6	9.2	0	123.2
Returns	0.01	0.9	0.6	6.2	-4	7

4.2 The Relationship Between Volume and Volatility

An implication of the Tauchen and Pitts model [31] is that volume changes over time for different reasons. In the analyses of Bjonnes et al. [3], Galati [18], Hartmann [22], Jorion [25] and Bessembinder and Seguin [2] a distinction is made between predictable and unpredictable volume. It is assumed that a major part of the information-driven volumes, or activity shocks, come as a surprise to dealers and are unpredictable. Increases in expected activity, or predictable volumes, on the other hand should primarily enhance the liquidity of the market and have little or negative effect on volatility.

The way predictable and unpredictable volumes are distinguished in practice is to identify and model the time series behavior of the volume series. The fitted values of the time series model, one step ahead forecasts, then become the predictable volumes while the residuals are the unpredictable volumes. The studies mentioned above all employ some form of ARIMA models. Bessembinder and Seguin [2] use an ARIMA(0,1,10) model, Jorion [25] chooses ARIMA(2,0,1), Hartmann [22] finds that ARIMA(9,1,1) fits his data best and Bjonnes et al. [3] select an ARIMA(2,0,2) model.

4.2.1 Modeling Trading Volume

Denoting daily trading volume with V_t, we first examine weather the series is stationary. In order to improve the stationarity properties of the series however we work with the logarithm of the volumes, $ln(V_t)$. This has the effect of making the variance of the series more uniform over the period, since the logarithm decreases large values relative to small values. Bessembinder and Seguin [2], Galati [18], Jorion [25] and Hartmann [22] also use log volumes for their analysis.

Augmented Dickey-Fuller and Phillips-Perron tests both reject the null hypothesis of unit-root non-stationarity at 1% significance level. However, Figure 4.2 suggests that the series is non-stationary since there is a significant increase in both mean and variance from February 2006 to the end of the period. Furthermore there are significant and persistent autocorrelation in the series, as the top-left panel in Figure 4.4 shows, which is a characteristic feature of non-stationary series.

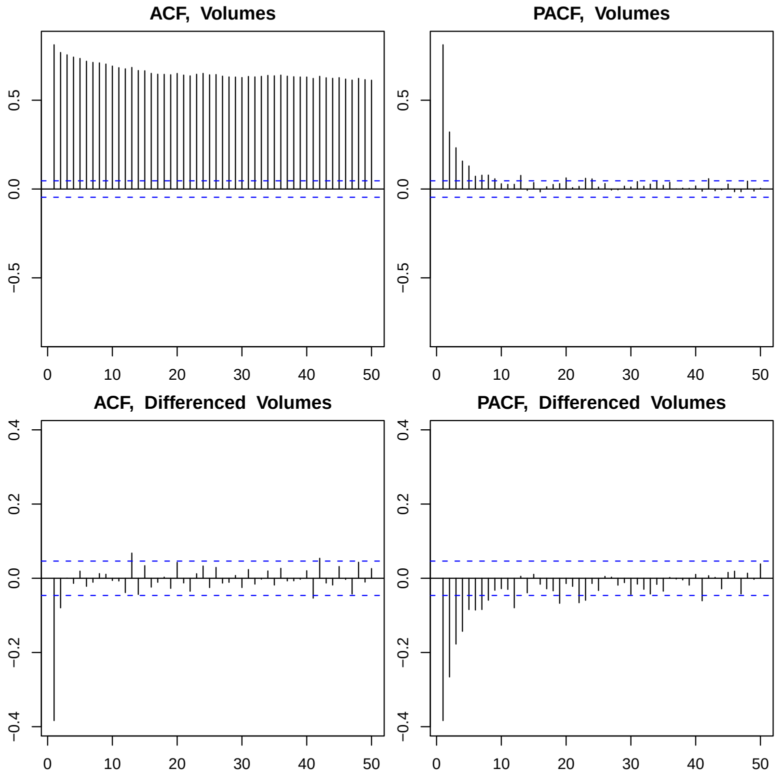

Figure 4.4: Autocorrelation functions and Partial Autocorrelation functions for the Volume series, and first differences of it.

Bjonnes et al. [3] and Hartmann [22] also encounter conflicting evidence for non-stationarity. In his analysis, Hartmann [22] finds that presence of non-stationarity is rejected by augmented Dickey-Fuller and Phillips-Perron tests, but when attempting to fit ARMA models to his volume data the roots of the AR polynomials lie outside the unit circle and the MA(1) parameters are close to -1. He argues that this fact leads to overrejection of the non-stationarity null hypothesis by the standard tests and that the series should be treated as non-stationary.

Brooks [8] recommends confirmatory data analysis when testing for non-stationarity to estimate the robustness of the tests' results. That is, the null hypothesis of non-stationarity and the alternative hypothesis of stationarity should be reversed. We use one such test, due to Kwaitkowski et al. [27], known as the KPSS test. We find that the KPSS test rejects stationarity of the volume series at the 1% level against an alternative hypothesis of a unit-root. We conclude that the volume series is unit-root non-stationary.

To induce stationarity, we construct a series of first differences of the log volumes, $\Delta ln(V_t) = ln(V_t) - ln(V_{t-1})$. The lower left panel in Figure 4.4 indicates that the autocorrelations of the differenced series do not exhibit the same degree of persistence they did for the plain level series. Testing for non-stationarity, it is rejected at 1% significance level by the standard tests like before but this time the KPSS null hypothesis of stationarity cannot be rejected. Thus we conclude that the first difference of the volumes is stationary.

In order to determine what model fits the first differences of volume best we estimate a range of ARIMA(p,1,q) models and calculate the Aikaike and Schwartz Bayesian information criterions for each one. Both information criterions suggest that ARIMA(1,1,3) fit the volume data best. Therefore we estimate the following model:

$$\Delta ln(V_t) = \mu + \phi_1 \Delta ln(V_{t-1}) + \theta_1 u_{t-1} + \theta_2 u_{t-2} + \theta_3 u_{t-3} + u_t, \quad u_t \sim N(0, \sigma^2) \tag{4.3}$$

We conduct a Ljung-Box test for autocorrelation of the residuals of the model fit, u_t, and find that we cannot reject the null hypothesis the data are random, i.e. we conclude that there is no residual autocorrelation, and that the model fits the data sufficiently well.

4.2.2 Modeling Exchange Rate Volatility

We model the continuous exchange rate returns with a GARCH(1,1) model. That is we estimate

$$\begin{aligned} R_t &= \mu + r_t, \quad r_t \sim N(0, h_t) \\ h_t &= \alpha_0 + \alpha_1 r_{t-1}^2 + \beta_1 h_{t-1} \end{aligned} \tag{4.4}$$

Results of the model estimation can be seen in Table 4.2. All coefficients are significant, although the constant for mean return, μ, is only significant at the 5% level.

Table 4.2: GARCH(1,1) model estimated for exchange rate returns, R_t.

	Estimate	Std. Error	t value	Pr(>|t|)
μ	-0.03	0.02	-1.95	0.05
α_0	0.07	0.02	3.10	0.00
α_1	0.17	0.03	4.96	0.00
β_1	0.75	0.06	12.71	0.00

4.2.3 Regression

In order to examine the relationship between volume and volatility we make use of the models we have fit to the volume and volatility data in sections 4.2.1 and 4.2.2. We estimate the following regression

$$|R_t| = \beta_0 + \beta_1 x_t^u + \beta_2 x_t^p + \beta_3 h_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2) \tag{4.5}$$

Here we regress a measure of volatility, the absolute exchange rate returns, on predictable volumes, x_t^p, unpredictable volumes, x_t^u, and predictable volatility, h_t as predicted by the GARCH(1,1) model in Equation (4.4). Predictable volumes are the fitted values, or one step ahead forecasts from Equation (4.3), that is $x_t^p = E_{t-1}[\Delta ln(V_t)]$. Unpredictable volumes are the residuals, or shocks, from the model, that is $x_t^u = \Delta ln(V_t) - E_{t-1}[\Delta ln(V_t)]$. Aside from unpredictable volumes the regressors are determined by information available at time $t-1$, and forecast the next day's volatility. The results of this regression are shown in Table 4.3.

Table 4.3: Volatility estimation, $|R_t|$, with decomposed volumes.

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept, β_0	0.43	0.02	22.87	0.00
Unexpected Volumes, β_1	0.39	0.03	15.03	0.00
Expected Volumes, β_2	-0.11	0.04	-2.77	0.01
GARCH forecast, β_3	0.24	0.02	12.61	0.00

The results indicate that unpredictable volumes have a significant impact on volatility, and moreover a stronger association than expected volatility or GARCH forecast. We find that expected volatilities and expected volumes also have a significant impact although the relationship is negative for expected volumes. This is in line with the Tauchen and Pitts model [31] which can be interpreted as predicting that an increase in expected volumes over time should primarily increase liquidity in the market and decrease volatility.

Our results are broadly consistent with comparable studies. Bjonnes et al. [3] examining trading with the Swedish krona find that unexpected spot volume effects are significant and positive, but that expected volumes are not significantly different from zero. Galati [18] finds unexpected volumes to be significant and positive for five out of seven emerging economies, the Columbian peso, South African rand, Indian rupee, Indonesian rupiah and the Israeli shekel, but expected volumes and GARCH forecasts to be significant and positive for only two out of seven. Jorion [25] examines deutsche mark currency futures and finds that unexpected volumes are positive and significant, as well as the GARCH forecast, h_t.

An interesting question is how stable this relationship is, and in particular weather it is different when there is turbulence in the foreign exchange market, or high volatility, than for periods of low volatility. We now turn to this question.

Periods of Calm, Periods of Stress

The data set we analyze lends itself well to be split up in two sub-periods, one for low volatility and the other for high volatility, since there is a sharp increase in the mean and variance of returns in February 2006 which is sustained throughout the period. We therefore split the data set into Period 1, from 2001-03-30 to 2006-02-20 and Period 2, from 2006-02-21 to 2008-10-02. Tables A.1 and A.2 show summary statistics for these two periods.

The analysis above is then repeated for these two periods. An ARIMA(1,0,2) model is selected for the trading volumes in Period 1 and an ARIMA(1,0,1) for Period 2. For volatilities, GARCH(1,1) is chosen for both periods. Equation 4.5 is then estimated again for the two periods. The results appear in Table 4.4 and Table 4.5.

Table 4.4: Volatility estimation, $|R_t|$, with decomposed volumes. Period: 2001-03-29 to 2006-02-20

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept, β_0	0.46	0.02	18.60	0.00
Unexpected Volumes, β_1	0.27	0.02	10.82	0.00
Expected Volumes, β_2	-0.10	0.04	-2.79	0.01
GARCH forecast, β_3	0.11	0.05	2.36	0.02

Table 4.5: Volatility estimation, $|R_t|$, with decomposed volumes. Period: 2006-02-21 to 2008-10-02

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept,β_0	0.55	0.04	14.37	0.00
Unexpected Volumes,β_1	0.69	0.06	11.30	0.00
Expected Volumes,β_2	-0.17	0.12	-1.45	0.15
GARCH forecast, β_3	0.19	0.02	8.44	0.00

The results indicate that the there is a marked increase in the coefficient estimates for all regressors for the period of turbulence relative to the period of calm. The coefficient on unexpected volume more than doubles. The relationship between volatility and expected volume seems to become more negative, although that indication is not conclusive since the coefficient estimate is not statistically significant from zero anymore. In the context of the MDH, these result imply that the link between information arrival, measured by unexpected volumes, and foreign exchange rate volatility is stronger during turbulent periods. Furthermore that expected volume does not have a significant effect on volatility.

5 Summary and Conclusions

In this study we set out to examine the link between trading volume in the interbank foreign exchange market and exchange rate volatility. We have reviewed the theoretical and empirical literature on this subject as well as the econometric methodology employed. The main theory concerning the relationship of trading volumes and volatilities, the Mixture of Distributions Hypothesis, predicts that trading volumes should be positively correlated with volatilities, and a large number of empirical studies in a variety of market settings have found this to be the case. We examined weather this relationship exists in a foreign currency market of a very small open economy, and our results indicate that it is indeed the case. Furthermore we ask the question weather this relationship is different for times of severe market stress. We find that during such periods the link between unexpected trading volumes and volatility becomes stronger.

A Tables

	Mean	Std.Deviation	Skewness	Excess Kurtosis	Minimum	Maximum
USD	78.83	14.3	0.5	-1	58.5	110.5
EUR	85.08	5	-0.7	0.3	71.3	97.4
Volumes	5.35	4.3	1.7	4.3	0	36.3
Returns	-0.03	0.7	0.2	5.8	-3.9	5.7

Table A.1: Summary statistics of daily spot volumes, daily exchange rates levels, and daily exchange rate returns. Period: 2001-03-30 to 2006-02-20

	Mean	Std.Deviation	Skewness	Excess Kurtosis	Minimum	Maximum
USD	69.86	7.4	1.5	4	58.7	112.7
EUR	96.51	14.5	1.4	0.9	77.6	155.7
Volumes	25.26	17.1	1.8	4.5	2.7	123.2
Returns	0.09	1.1	0.6	3.9	-4	7

Table A.2: Summary statistics of daily spot volumes, daily exchange rates levels, and daily exchange rate returns. Period: 2006-02-21 to 2008-10-02

Bibliography

[1] Bauwens, L., Rime, D., Sucarrat, G. (2005). Exchange Rate Volatility And The Mixture of Distribution Hypothesis. *Universite catholique de Louvain, Center for Operations Research and Econometrics*. CORE Discussion Papers.

[2] Bessembinder, H., Seguin, P. J., (1992). Futures-Trading Activity and Stock Price Volatility. *The Journal of Finance*, 47 (5). 2015-2034.

[3] Bjonnes, G. H., Rime, D., Solheim, H. O. Aa., (2005). Volume and volatility in the foreign exchange market: Does it matter who you are? Grauwe, P. (Ed.), *Exchange Rate Economics: Where do we stand? MIT Press*. ISBN-10 0-262-04222-3.

[4] Bollerslev, T., Chou, R. Y., Kroner, K. F., (1992). ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. *Journal of Econometrics*, 52 (1-2), 5-59.

[5] Bollerslev, T., Domowitz, I., (1993). Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance*, 48, 1421-43.

[6] Bollerslev, T., Engle, R. F., Nelson, D. B., (1994). ARCH Model. In R. F. Engle and D. C. McFadden (eds.) *Handbook of Econometrics IV*. 2959-3038. Elsevier Science, Amsterdam.

[7] Box, G. E. P. and Jenkins, G. M., (1976). Time Series Analysis: Forecasting and Control, 2nd edn., *Holden-Day, San Francisco*.

[8] Brooks, C., (2002). Introductory Econometrics For Finance. *Cambridge University Press*. ISBN-10 0-521-79018-2.

[9] Central Bank of Iceland, (2001). The Icelandic foreign exchange market. *Monetary Bulletin*, 3 (3). 57-63.
"http://www.sedlabanki.is/?PageID=234"

[10] Central Bank of Iceland. Temporary foreign exchange arrangements. *Press releases*, 15/10/2008.
"http://www.sedlabanki.is/?PageID=287&NewsID=1913"

[11] Central Bank of Iceland. Interbank Foreign Exchange Market. *Press releases*, 03/12/2008.
"http://www.sedlabanki.is/?PageID=287&NewsID=1988"

[12] Central Bank of Iceland. Data on foreign exchange trading volumes.
"http://www.sedlabanki.is/lisalib/getfile.aspx?itemid=3537"

[13] Central Bank of Iceland. Data on foreign exchange rates.
"http://www.sedlabanki.is/?PageID=37"

[14] Clark, P., (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41 (1), 135-155.

[15] Dumas, B., (1996). Comment. In: Frankel, J., Galli, G., Giovannini, A. (Eds.), *The Microstructure of Foreign Exchange Markets*. University of Chicago Press, Chicago, IL, 37-40.

[16] Engle, R. F., (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50 (4), 987-1007.

[17] Epps, T., Epps, M., (1976). The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. *Econometrica*, 44 (2), 305-321.

[18] Galati, G., (2000). Trading volumes, volatility and spreads in foreign exchange markets: evidence from emerging market countries. Working Paper, nr. 93., *Bank for International Settlements*, Basel, Switzerland.

[19] Goodhart, C., Figliuoli, L., (1991). Every minute counts in financial markets. *Journal of International Money and Finance*, 10, 23-52.

[20] Goodhart, C., Ito, T., Payne, R., (1996). One day in June, 1993: a study of the working of Reuters 2000-2 electronic foreign exchange trading system. In: Frankel, J., Galli, G., Giovannini, A. (Eds.), The Microstructure of Foreign Exchange Markets. *University of Chicago Press*, Chicago, IL. 107-179.

[21] Grammatikos, T., Saunders, A., (1986). Futures price variability: a test of maturity and volume effects. *Journal of Business*, 59, 319-330.

[22] Hartmann, P., (1999). Trading volumes and transaction costs in the foreign exchange market. *Journal of Banking and Finance*, 23, 801-824.

[23] Hayashi, F., (2000). Econometrics. *Princeton University Press*. ISBN 0-691-01018-8.

[24] Jorion, P., (1996). Risk and turnover in the foreign exchange market. Frankel, J., Galli, G., Giovannini, A. (Eds.), *The Microstructure of Foreign Exchange Markets*. University of Chicago Press, Chicago, IL, 19-37.

[25] Karpoff, J., (1987). The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis*, 22 (1), 109-126.

[26] Kwaitkowski, D., Phillips, P. C. B., Schmidt, P. and Shin, Y., (1992). Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root. *Journal of Econometrics* 54. 159-178.

[27] Lyons, R., (1995). Tests of microstructural hypothesis in the foreign exchange market. *Journal of Financial Economics*, 39. 321-351.

[28] Sarno, L., Taylor, M. P. (2001). The Microstructure Of The Foreign-Exchange Market: A Selective Survey Of The Literature. *The International Economics Section, Department of Economics, Princeton University.* ISBN 0-88165-261-X.

[29] Schumway, R., Stoffer, D. S. (2006). Time Series Analysis and Its Applications, 2nd ed. *Springer Science*. ISBN-10 0-378-29317-5.

[30] Tauchen, G., Pitts, M., (1983). The price variability-volume relationship on speculative markets. *Econometrica*, 51, 485-505.

[31] Tsay, R. S., (2005). Analysis of Financial Time Series, 2nd ed. *Wiley-Interscience*. ISBN-10 0-471-69074-0.