Quantitative Analysis of Foreign Exchange Rates

Alexander Becker, Ching-Hao Wang
Boston University, Department of Physics
(Dated: today)

In our class project we have explored foreign exchange data. We analyze daily and hourly returns for the five major currencies, US dollar, Euro, Japanese yen, British pound, and the Swiss franc. The data can be best fitted by a q-Gaussian but we have not been able to model this behavior yet. We start investigating some traits of the distribution, like correlation and time dependence, to open the door for further analysis that hopefully can lead to a better understanding of the FX market in the future.

Keywords: FX markets

INTRODUCTION

We consider the fluctuation in foreign exchange rates. Unlike stock prices, foreign exchange rates are strongly influenced by policies of countries or currency unions. They are not only strongly tied to macroeconomic factors such as inflation and unemployment but also have a direct impact on any other markets because essentially any price is expressed in terms of currencies. Furthermore, a given set of currencies is a closed system in which no overall net gain can be achieved: The rise of one currency compared to a second implies the fall of the second currency compared to the first.

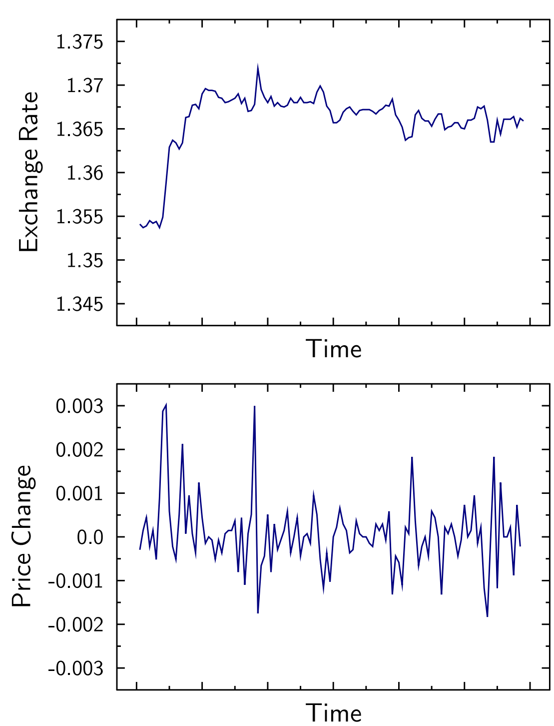

Figure 1: Top: Typical time series for currency exchange rate. Bottom: Returns calculated based on this time series..

Foreign exchange markets are considered to be the most liquid market of all. Given the large amount of research dedicated to equity markets, findings about foreign exchange markets can be used to crosscheck findings about stock prices. For example, the distinct power-law tails of stock price returns are hypothesized to be linked to a lack of liquidity in the market. This means that one should not expect them to be present in returns of currencies. In addition, it has been shown in literature that the power-law distribution is Lévy-unstable since for sufficiently large lag time (i.e. time step dt) it will converge to a Gaussian [3, 4]. Thus, it is of importance to look into the foreign exchange rate distribution to see how it differs from that for the well-identified form of stock market.

We will start our investigation by quantitatively examine the returns on foreign exchange markets. We consider the following currencies: US dollar (USD), Euro (EUR), Japanese yen (JPY), Pound sterling (GBP), and Swiss franc (CHF). These are, with exception of the Australian dollar, the most frequently traded currencies by value, and thus they are most suitable for a quantitative analysis. We will look at the changes in foreign exchange rates over different time horizons from one hour to one day.

Despite the large liquidity present in foreign exchange markets, trade is not evenly distributed across the 24 hours of a day. Instead, the business hours of mostly the European market and the US market appear to be the most liquid periods of the day. We will look into this, investigating the influence of this on the return distribution.

Ultimately, we consider the the relationship between fluctuations of the exchange rate during a certain period and the realized price change within that period.

METHODOLOGY

The basis of analysis is data that spans from Jan 3, 2001, to Mar 28, 2014. For each of the currencies used in this analysis we use hourly data that also contains information about the low and high during each hour. The data has been filtered such that for every hour the data is complete. As a consequence, 67 out of 81,007 trading hours have been neglected. Since the data comprises opening and closing price for each hour, this does not pose a problem for computing the returns.

The returns are calculated the well-established way,

$$r_k(t+\mathrm{d}t) = \ln \frac{S_k(t+\mathrm{d}t)}{S_k(t)}, \tag{1}$$

where $r_k(t)$ describes the return of currency k at time t, $S_k(t)$ is the currency exchange rate for currency k at time t, and $\mathrm{d}t$ is the time step. A typical time series for S and r is shown in figure 1.

The data analysis has been performed with Wolfram Alpha and MatLab.

DATA ANALYSIS

Daily Data

We start our analysis with looking at the daily return of the major currency pairs, in total 10 currency pairs and, thus, 20 return rate time series. Regardless of the currency pair, we calculate the daily return for each of the 24 hours of the day. This yields a list of return values that we bin to investigate the distribution.

However, the binning on the x-axis is not linearly but exponentially spaced; this is to take into account the anticipated result that low changes are much more likely than much larger changes. In praxis, we fix the number of bins that we want to investigate in divide the x-axis into bins that grow exponentially as the price change grows. This allows us to occupy each bin a little more evenly. As for the position on the x-axis of the bin: As long as the bin size is small enough compared to the region of price changes investigated, the midpoint of the interval is a very good approximation.

The result for this analysis is shown in figure 2. In this log-plot we see no clear functional form; instead the curve seems to has dents and inflection points. For very small price changes, the form shows indications of a parabola, only to bend away for higher price changes into what looks like an exponential tail. We will further pursue this analysis in the hourly data, where we look at some of the characteristics of the distribution.

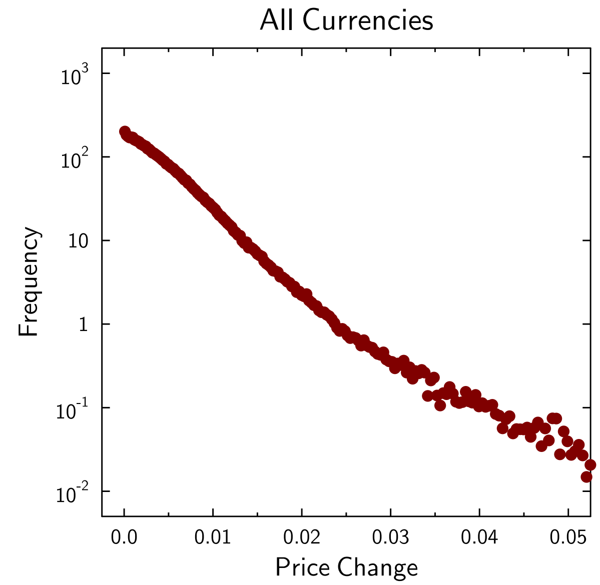

Figure 2: Distribution for the daily foreign exchange rate returns for all currency pairs. This is a log-plot.

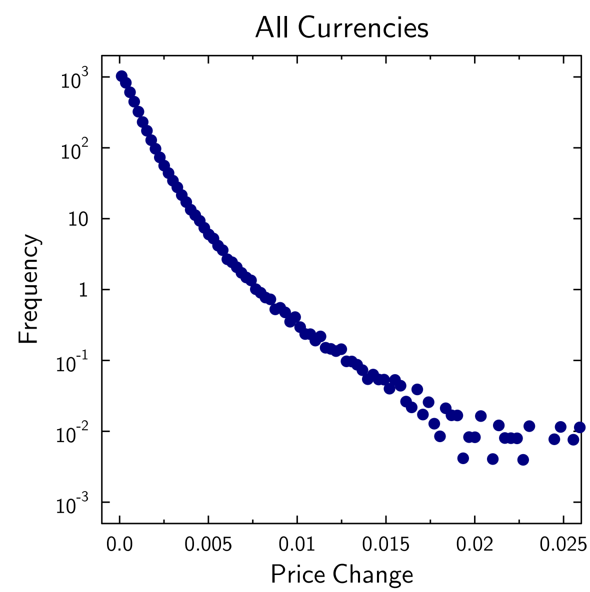

Figure 3: Distribution for the hourly foreign exchange rate returns for all currency pairs. The y-axis is logarithmically scaled.

Hourly Data

We investigate the hourly returns of the above mentioned currencies, which yields 10 currency pairs and, thus, 20 return rates. Consequently, any distribution will be symmetric around the origin then. Therefore we generally restrict ourselves to plotting only the positive tail of the distribution.

Figure 3 shows the distribution of the return rates in the hourly data. We recognize that it decays slower than an exponential function, which have yielded a straight line. However, figure 4 reveals that we do not find a power law either. For a power law, we would require to be able to fit a straight line to the data in this representation over several orders of magnitude – a task obviously impossible here.

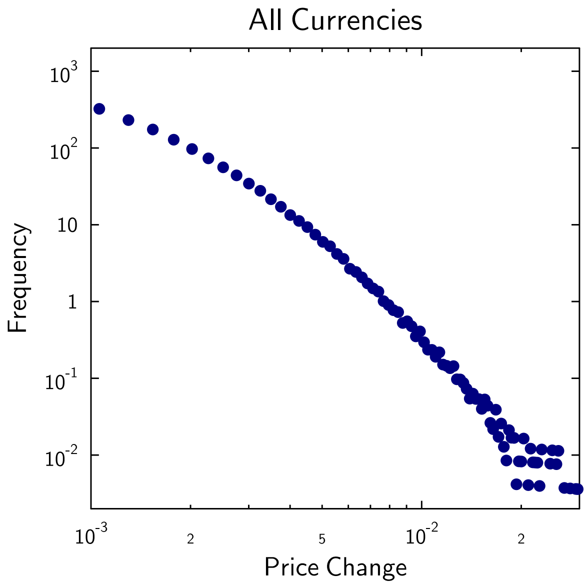

Figure 4: Distribution for the hourly foreign exchange rate returns for all currency pairs. Both axes are logarithmically scaled.

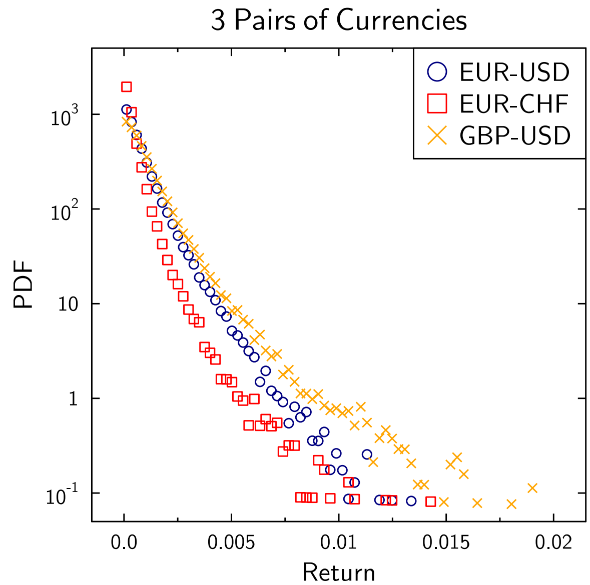

Figure 5: Distribution for the hourly foreign exchange rate returns for all currency pairs. Both axes are logarithmically scaled.

While we expect our result to be more general, it is worthwhile to investigate the distribution from figure 3 further for different currencies. In figure 5 we compare three pairs of currencies. Particularly outstanding is the role of the Swiss franc compared to other currencies, particularly the Euro. As the Swiss government has imposed an exchange rate peg against the Euro, we observe different behavior: Small changes are much more likely than big ones, yielding a much steeper curve.

In general the distribution of foreign exchange is different from Gaussian and to our best knowledge no exact theory is present so far to give it a functional form. It has been argued that the model of a random walk would not be consistent with theoretical framework of foreign exchange rates [2]. However, recently it has been proposed that the empirical data in, say, financial returns in the New York Stock Exchange and NASDAQ, should rather be interpreted as q-gaussian, which, from the physics point of view, stems from the principles of maximizing Tsallis entropy under appropriate constraints [5, 6].

This formalism was first introduced with an attempt to explain the non-Gaussian fluctuations in the options of stock prices and indeed it closely fits the empirically observed distribution for many financial time series quite well [7]. However, the q-Gaussian model can only imply how much the distribution deviates from the ideal normal in some of stock market scenarios, a theory that more appropriately describes the foreign exchange market is still lacking. As a result, we would like to pause here and instead comment on the effect of time steps ($\mathrm{d}t$) on the distributions.

As the time step grows, that is the length of the period for which the currency pair has been traded before assessing its new price, we expect the distribution to become flatter. A flatter distribution implies more larger price changes.

To investigate the time effect, we have revisited the data every $\mathrm{d}t$ and calculated the returns, such that between every quota there has exactly $\mathrm{d}t$ passed. In figure 6, the result of this analysis is shown for $\mathrm{d}t = 2, 4, 6, 8, 10, 12$h. As expected, the distribution becomes flatter. Remarkably, there appears to be a cross-over point for the different time scales. To show this cross-over point better, figure 7 provides a close-up of the region.

We have also examined the peak of the distribution to further analyze its scaling behavior. In a first approach, we compare the first point in the binned distribution for each time step with one another. We find that the peak scales as $1/\sqrt{2}$ when doubling the time step length.

These observations open room for a more detailed analysis of the time series' return distributions.

The data set with which we are working not only provides opening and closing prices, but also the highs and lows of that particular hour. As a proxy, we use this to analyze the fluctuation of the price compared to the realized price changes. We introduce the parameter ξ, defined as follows:

$$\xi_k(t) = \frac{\log\left(S_k^{\max}(t)/S_k^{\min}(t)\right)}{r_k(t)}. \tag{2}$$

$S_k^{\max}(t), S_k^{\min}(t)$ are the maximum and minimum value of the currency k in the interval $[t - \mathrm{d}t, t]$, respectively.

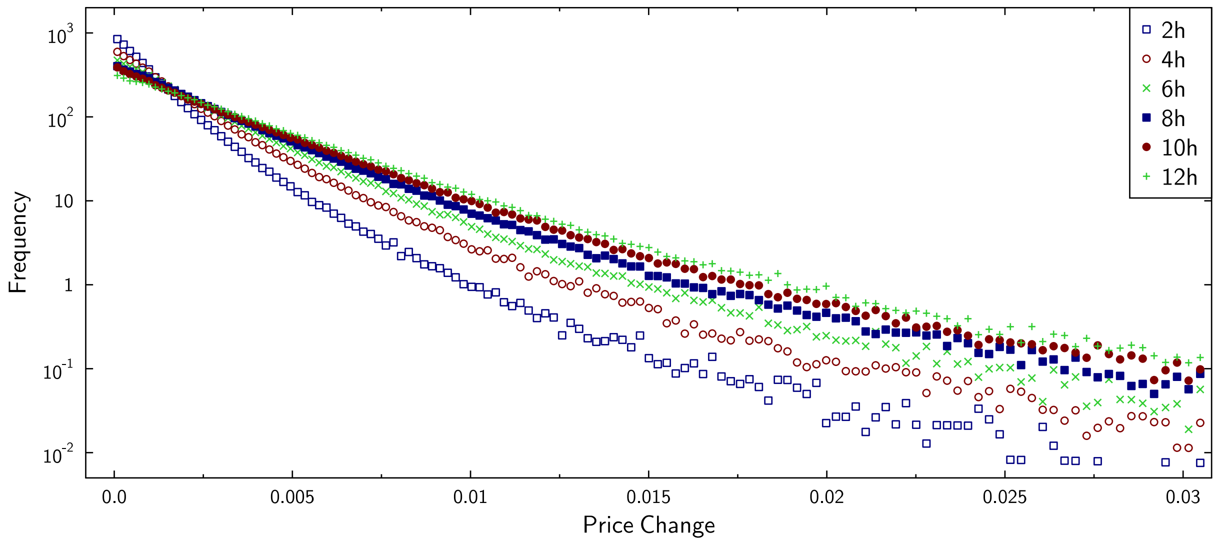

Figure 6: The distribution of price changes for different lengths of the time step dt, ranging from dt = 2h to dt = 12h. We see that the distribution becomes flatter and identify a potential cross-over point for low price changes.

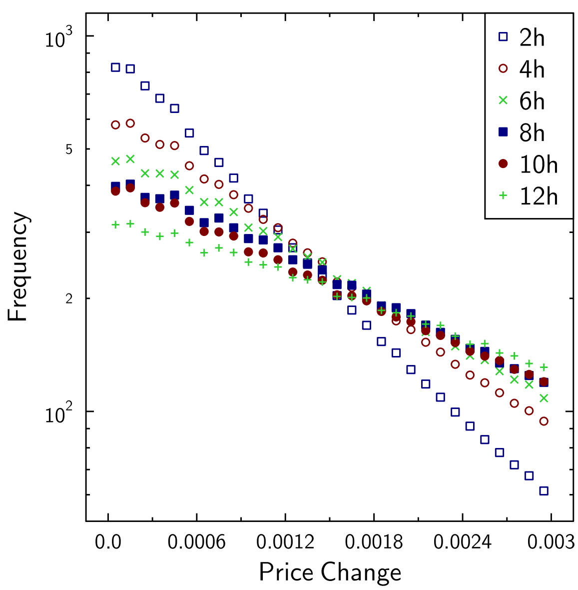

Figure 7: Zooming in on the potential cross-over point in the distribution for different time-steps.

$r_k(t)$ is the return in this same time interval.

The larger ξ, the more the price has been fluctuating compared to its final return. A value of $\xi = 1$ implies that the price has been fluctuation between opening and closing price only. The parameter, however, does not let us infer any information on the size of the fluctuations themselves. One could further analyze how ξ scales with the actual price change for this.

Figure 8 shows the distribution of the parameter ξ. This can be very well fitted to a power law with an exponential cut-off. The log-representation nicely illustrates the exponential tail for large fluctuations.

We have started comparing this parameter to the typical behavior for stock prices. First analyses hint that there is no exponential cut-off in stock-prices. Before we interpret this, however, we will have to extend the data pool for stock price data to see if this holds if we improve our statistics.

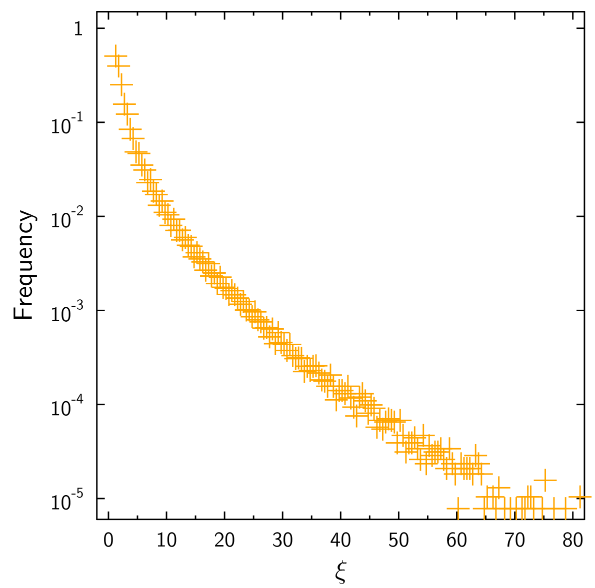

Figure 8: Distribution of the parameter ξ, which describes the fluctuation around the realized return rate.

[1] Y.-F. Chen *et al.*, Phys. Rev. Lett. **107**, 217401 (2011)
[2] F. Alvarez, A. Atkeson, and P. J. Kehoe, AEA Papers and Proceedings (May 2007)
[3] R. Rak, S. Drożdż and J. Kwapień, Physica A **374**, 315 (2007)
[4] S. Drożdż, J. Kwapień , P. Oświecimka and R. Rak, Europhys. Lett. **88**, 60003 (2009)
[5] L. Borland, Phys. Rev. Lett. **89**, 098701 (2002)
[6] Constantino Tsallis, Celia Anteneodo, Lisa Borland, Roberto Osorio, Physica A **324**, 89 (2003)
[7] S. Drożdż, *et al*, New J. Phys. **12**, 105003 (2010)