DUKASCOPY
RESEARCH PRODUCTS

02/01/2012

Economic Research

Currency Pair Returns: Risks

The Dukascopy Bank SA research department continues its work on statistical properties of currency pair returns. The previous part of the research marked out relationships amongst returns. It was therefore dedicated to a somewhat determinate aspect of their behaviour. In this issue we look into uncertainty of the movements and investigate the parameters that can help reduce it.

The most widely quoted measure of uncertainty in asset movements is volatility. Some sources position it merely as a measure of risk, while others go as far as proposing associated trading strategies. Therefore we set a goal to investigate properties of volatility and see what information it can actually carry.

We base the study on our five top traded currency pairs: EUR/USD, GBP/USD, USD/JPY, USD/CHF and EUR/JPY, - as they seem to be of the greatest interest to our clients. The data used are different frequency exchange rates (2012 ten minute and one hour rates, and 2000-2012 one day rates).

Methodology

Volatility is a measure of how unstable the price of a financial instrument is on a time scale. The more fluctuation it experiences, the higher the volatility is.

Although we mostly use volatility to describe exchange rates per se, the calculations are not based on prices directly. Instead, we derive volatility from returns.

Returns reflect the changes of prices, and asset instability gets captured in their fluctuations. Namely, interchanging positive and negative returns of great magnitude define high instability, while smaller variations indicate calmer movements (see Figure 1).

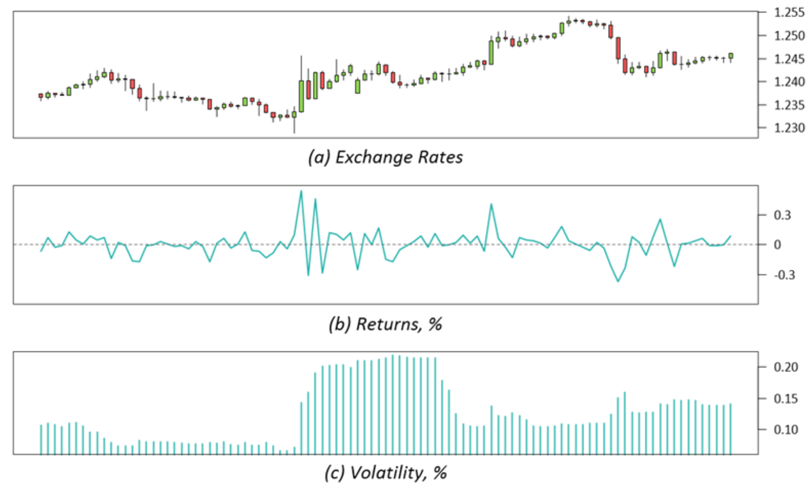

Figure 1: 5-day sample for EUR/USD data

In the following sections we discuss such features as currency pair volatility levels, volatility clustering, periodicity, and effects of scaling volatility.

Volatility Calculation

There are several commonly used formulas for calculating volatility, but they all reflect similar ideas. In fact, the most popular ones – squared returns, return absolute values, and return standard deviations, - are derived from the same notion.

Standard deviation is a statistical measure of data spread around the mean. For currency pair returns it can be estimated as:

$$\sqrt{\frac{\text{sum of squared (return - average return)}}{\text{number of returns - 1}}}$$

Assuming average return to be zero and considering only one value, we get an absolute value measure of volatility. Removing the root from the formula produces a squared return measure.

Figure 2 shows that significant peaks of volatility almost always coincide across the measures. Thus we can expect to receive similar signals whichever formula we use. However, we can also see the differences between the measures, as well as pros and cons of using one or the other. Squared returns and return absolute values appear to be almost as chaotic as returns themselves, giving us little additional information. Besides, they require making an assumption about average return, which automatically makes them less accurate. Standard deviation, on the other hand, is more precise and more consistent, as it averages the value over a certain period. It thus lets us evaluate any new price movement in respect to previous ones, and see whether it really is out of the ordinary. On the downside, the measure is strongly dependent on the said period, and there is no prescribed method for choosing it.

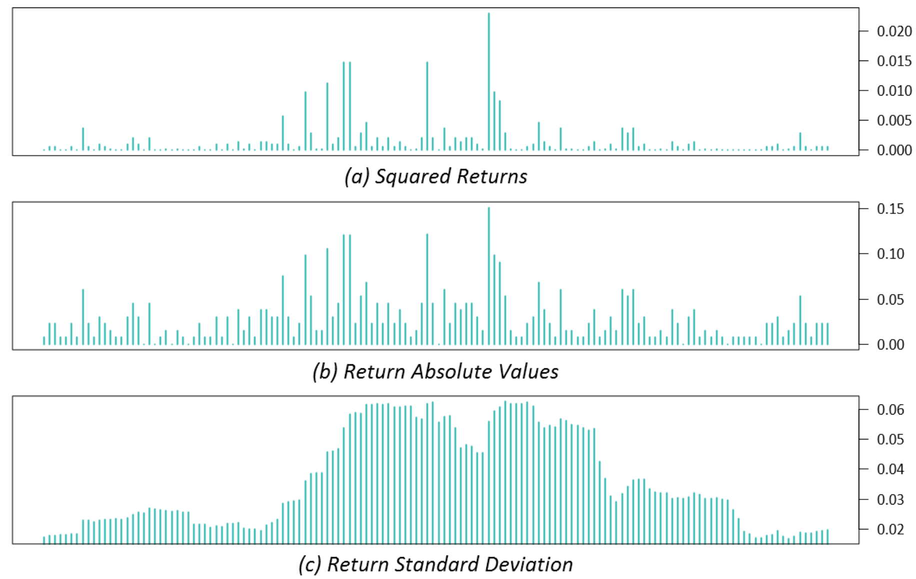

Figure 2: 1-month sample for EUR/USD returns

Another obvious, but nevertheless important, difference is the range of values. Squared returns are typically measured in hundredths of per cent, while absolute values and standard deviation – in tenths. In addition, absolute values tend to have higher peaks than standard deviations. Thus, to interpret the values correctly it is important to know which measure of volatility is used.

Results

1. Each currency pair has a specific volatility range.

We have investigated sets of one-week volatilities for intraday rates and 3-month volatilities for daily data. It appeared that each currency pair has a normal volatility level of a sort - a range of frequently reoccurring values. Furthermore, relative levels hold for all frequencies and time-frames: EUR/JPY seems to be the most volatile, USD/CHF and EUR/USD are in the middle, but USD/JPY and GBP/USD are the calmest. The table below shows approximate ranges that cover half of all observed values. Thus, 50% of volatility values are expected to fall into said intervals, but any numbers outside it can be considered either elevated or low.

Frequency	EURJPY	USDCHF	EURUSD	USDJPY	GBPUSD
10-minute	0.04-0.07	0.03-0.05	0.03-0.05	0.02-0.04	0.02-0.04
1-hour	0.11-0.16	0.08-0.12	0.08-0.12	0.07-0.10	0.06-0.09
1-day (2000-2012)	0.60-0.86	0.59-0.77	0.56-0.70	0.51-0.70	0.48-0.61
1-day (2011-2012)	0.64-0.85	0.60-0.89	0.59-0.73	0.42-0.61	0.43-0.54

Table 1: Normal volatility ranges

The last row shows that ranges may change over time. Here USD/CHF became slightly more volatile in the past two years, while USD/JPY and GBP/USD, on the contrary, calmed even more.

Figure 3 shows EUR/JPY and GBP/USD volatility histograms – diagrams that connect values to the frequencies of their appearance. It serves as a good illustration for the proposed currency pair ranking: not only is EUR/JPY more volatile on average, its values are twice as extreme as GBP/USD. Therefore, it has both a higher general volatility range and further surges out of it.

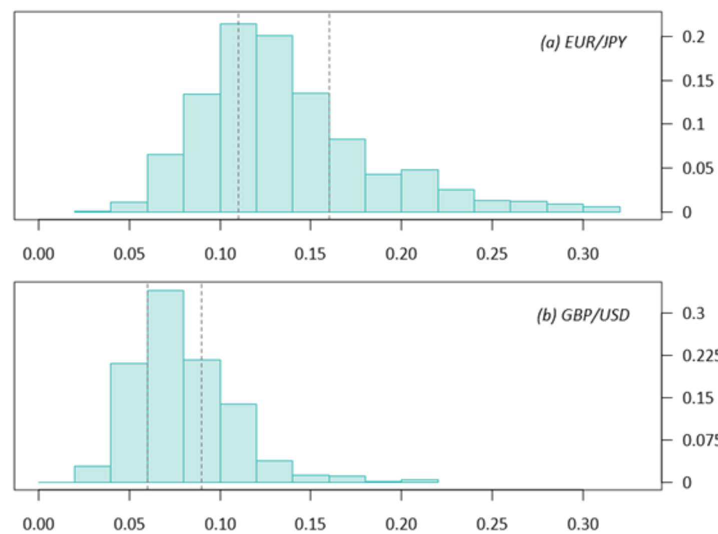

Figure 3: 1-hour volatility histograms with normal volatility ranges

2. Volatility tends to linger on the achieved level.

This phenomenon is known as volatility clustering – high (low) values of volatility are usually followed by high (low) values. This means that, in absence of any disturbing news, currency pair volatility changes gradually, and for some period of time one can expect it to be approximately the same. One way to establish this period is to look into volatility autocorrelation.

Autocorrelation is a measure of dependence between present and past values. It ranges between -1 and 1. A positive autocorrelation indicates that new values tend to follow a previous trend, but negative – that they generally change direction. Lag marks the time difference between present and past: lag 1 is one-period difference, lag 2 – two-period, and so on.

As standard deviation is calculated over several return values, it may seem that positive autocorrelation here is merely due to calculation specifics. However, the maximal lag for which the values stay statistically greater than zero exceeds the period used in the volatility formula. In addition, autocorrelation of currency pair volatility proves to be positive for all definitions. Thus the result indicates a real existing relationship. Furthermore, analysis has shown that autocorrelation values increase with decreasing frequencies.

For 10-minute volatility, autocorrelation goes from 0.3 out of 1 down to 0.01, and maximal lag is 45. Thus the relationship, however weak, holds for up to seven hours. Autocorrelation for hourly volatility keeps above 0.1 for 50 lags, or two days, but for daily values – as long as two months.

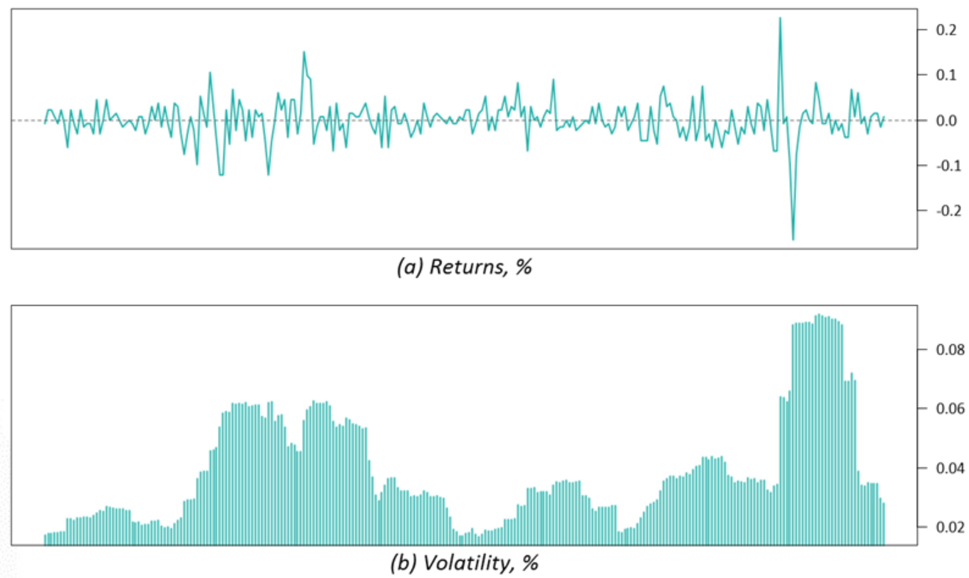

Figure 4: 10-minute EUR/USD returns and volatility

These results, of course, are general, and do not express any strict laws of volatility behaviour. However, they show two things to bear in mind. Firstly, that a single burst in volatility, e.g., in response to a news release, can elevate the instability level for a long period of time after the event. And secondly, that calm and steady periods only last for a limited amount of time, and must not be heavily relied on.

3. Volatility of high frequency rates is periodical.

Periodicity is a quantity's property of taking the same values at regular time intervals.

In relation to a stochastic variable like volatility, such a definition is too strict. Rather than taking absolutely identical values, it periodically returns to the same levels, making certain time periods generally more volatile than others. Autocorrelation can be used to examine this property in the same way it was used for clustering.

Figure 5 shows that the relationship for 10-minute volatility tends to strengthen every 12 hours, with interchanging positive and negative values. Thus, for example, a current high volatility level will be replaced by a lower one in the next 12 hours, but will be reached again by the same time the next day. The second part of the figure is a good illustration of this effect. It suggests that midday is the most volatile part of the day, while mornings and evenings are generally significantly calmer.

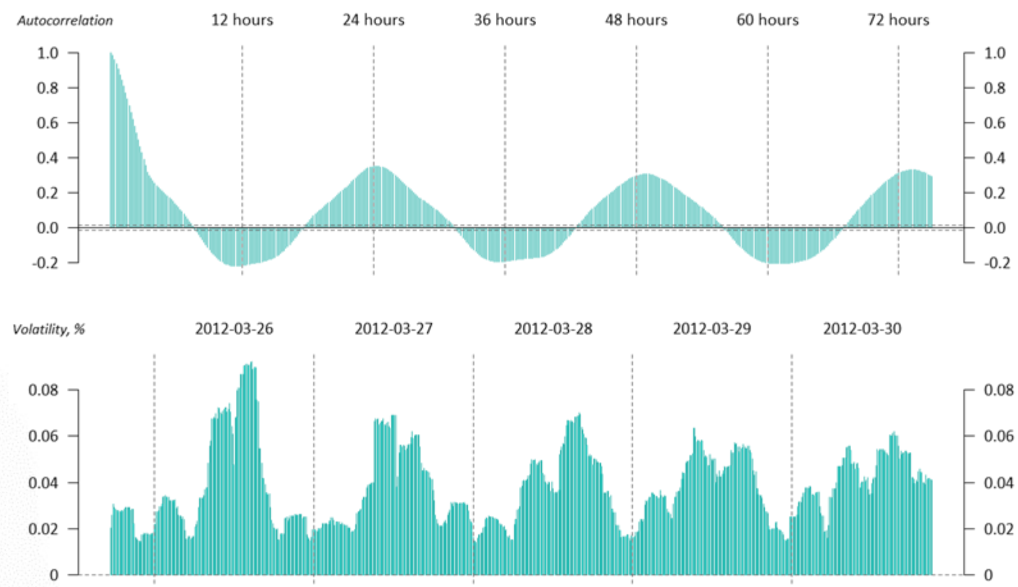

Figure 5 : EUR/USD 10-minute volatility autocorrelation and 1-week sample

The picture stays the same for all examined weeks and almost all pairs. Only the ones with Japanese yen seem to behave differently. EUR/JPY volatility autocorrelations show a little deformed pattern with less negative values, while USD/JPY ones do not have any negative autocorrelations at all. However, periodical strengthening of the relationship is attributable to these pairs as well.

Lower frequency data do not display such a property. Neither hourly, nor daily volatility marks any days, weeks or months as constantly more unstable than others.

4. Volatility scaling often gives significantly inaccurate results.

"Square root of time" scaling is a popular method for converting high frequency volatility into a lower frequency one. It is considered advantageous, because it allows saving both on calculations and data storage. Calculating yearly volatility, for example, would require using yearly exchange rates for the past several years. Not only information like this is hard to find, it is also likely to produce an unreliable outcome, as volatility is too changeable for such broad time frames. Scaling, on the other hand, allows us to calculate daily volatility and use it to get an estimation of a desired value. However, estimations are never precise – they have their confidence intervals and likelihoods. Unfortunately, these are rarely quoted by sources that provide scaled values. That is why it is important to understand the possible degree of inaccuracy, especially in investment planning. Here volatility which is scaled too low can lead to risk underestimation, while one too high can result in overestimating expected profit.

We have calculated scaled volatilities for different frequency data and examined the difference between real and scaled values.

It appears that with the increase of the scale factor, volatility gets mostly underestimated, but rare overestimations become more extreme. Converting volatility from 10 to 30-minute, or from 30-minute to 1-hour, requires scale factors of 3 and 2 respectively. The difference in such cases is close to zero with minor overestimation. Scaling hourly volatility to daily with a factor of 24 shifts the results into underestimation. Real values exceed the scaled ones by 0.05-0.15%, that is, by less than a quarter of the real volatility. Maximal underestimation is about 0.3%. Negative differences, on the other hand, occur less frequently, but reach the level of -0.7%. This means, for example, that when in reality EUR/USD is extremely calm, scaling shows that its volatility is on an average level.

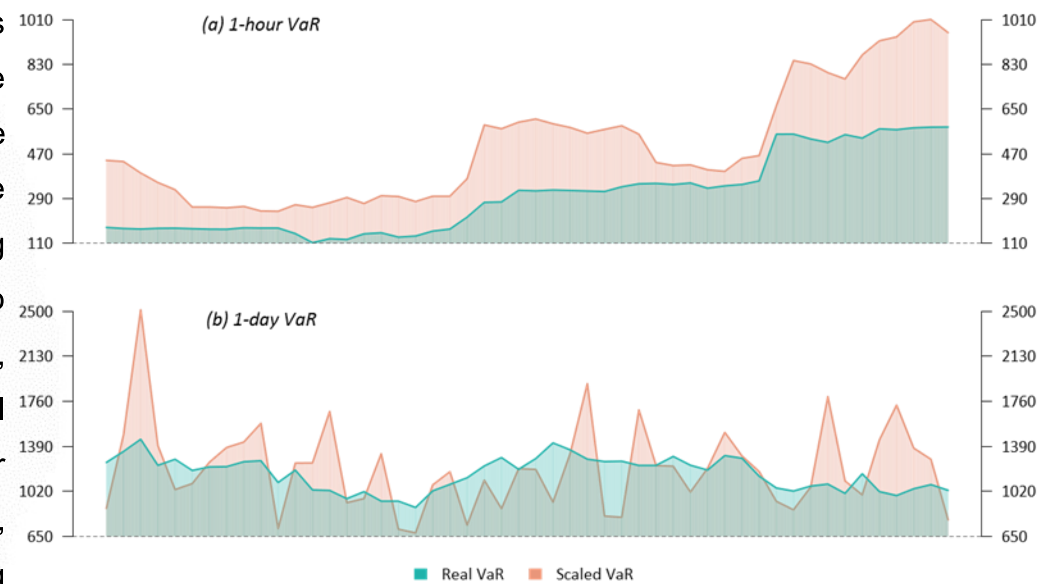

Figure 6: EUR/USD real and scaled Value-at-Risk sample

This effect can be seen in Figure 6. Here we converted volatility into Value-at-Risk for a $100,000 investment. In addition to showing the difference between real and scaled values, the plots highlight how scaling magnifies instability in general.

Conclusion

In this research we have looked into instability of exchange rates and possibilities to reduce its effects. It appeared that under certain circumstances some prognoses are possible. Their accuracy, however, is limited to expecting one or the other level of volatility.

We have found that the magnitude of exchange rate fluctuations changes gradually, and that less stable times generally last longer than calm ones. In addition, high frequency rates seem to be more volatile in the middle of the day, while evenings and mornings offer a calmer picture. This might be useful for both risk-averse and risk tolerant investors while choosing a more preferable trading time.

Scaling appeared to be a rather inaccurate way to convert volatility. It tends to under or overestimate risks, depending on the magnitude of the scaling factor. However, sometimes it is impossible to manage without conversion. In such cases it is important to estimate not only the value, but also its possible error.

Finally, we have established that each currency pair has a typical volatility level, and that some pairs are more unstable than others. Thus, it is possible to balance risks and profits by taking into account this “instability ranking” of currency pairs. However, we have also seen that the levels can change over time. Therefore, to try to use volatility, one must monitor it closely to capture any changes in behaviour.

Newest releases and archive:

- Fundamental Analysis
- Technical Analysis
- Press Review
- Market Research
- Expert Commentary
- Dukascopy Sentiment Index
- Trade Pattern Ideas
- Global Stock Market Review
- Commodity Overview
- Economic Research
- Quarterly Report
- Aggregate Technical Indicator

Additional information:

- Dukascopy Group Home Page
- Market News & Research
- FXSpider
- Live Webinars
- Dukascopy TV
- Daily Pivot Point Levels
- Economic Calendar
- Daily Highs/Lows
- SWFX Sentiment Index
- Movers & Shakers FX
- Forex Calculators
- Currency Converter
- Currency Index
- CoT Charts

Social networks: