

Plain noise data: predict this!

In this paper, we examine a case inspired by an online discussion with one of the users on our support forums: www.spiderfinancial.com/forums/financial-time-series/forex/. The case deals with the predictability of the daily closing EUR/USD exchange rate, given its prior history.

To start, we constructed a trading strategy where we buy EUR at market open, and close the position (selling EUR) at market close.

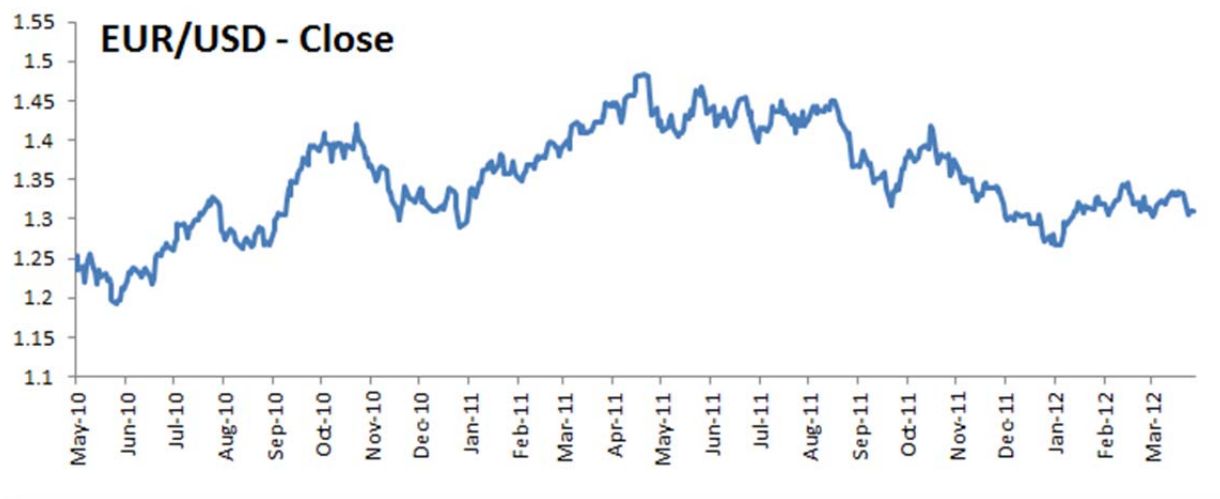
The FX EUR/USD spot market is very efficient as traders constantly monitor and exploit (i.e. arbitrage) any information in the rates time series until it is uneconomical to exploit any further. As a result, it comes as no surprise to us that the daily returns data looks much like a white-noise distribution. This prompts the question: what is next?

Assuming that our main objective for forecasting a next-day EUR/USD closing rate is to generate profits, we took a slightly different approach and devised a strategy with stop-loss settings that not only limit the daily loss, but also generate a positive return.

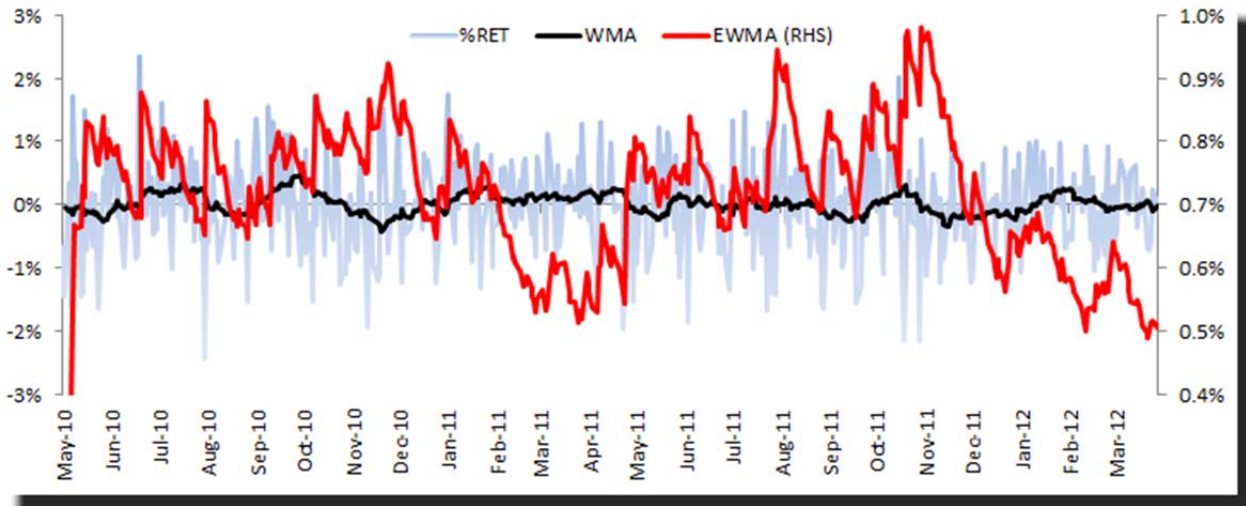
This approach has proven to be more fruitful and will pave the way for many more improvements.

Daily Returns Analysis

For our sample data, we use the FX exchange rate of EUR/USD between March 2010 and April 9th, 2012.



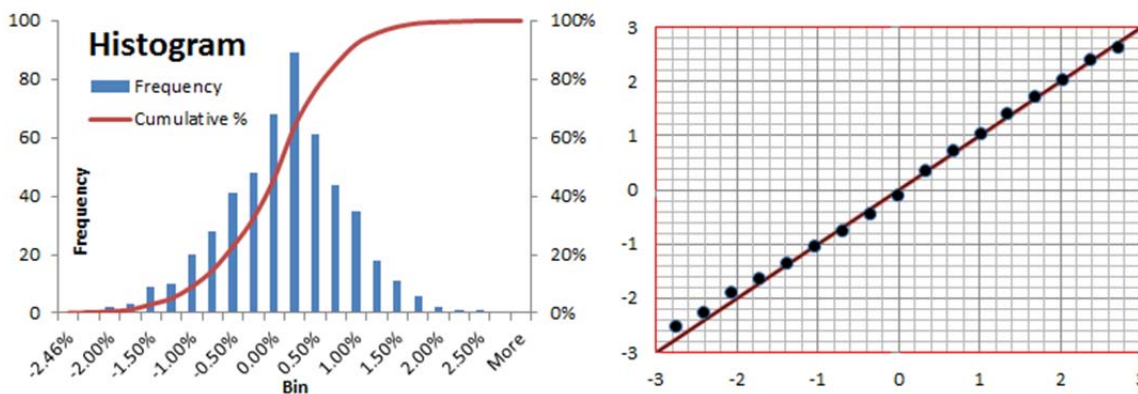
In this paper, we consider a trading strategy that buys EUR (1EUR = FX rate in USD) at market open (MOO) and closes the position (i.e. sells EUR) at market close (MOC).



The daily returns time series shows a steady mean (proxy by 20-day WMA) and a relatively stable volatility (0.5-1.0%). Examining the summary statistics, the average return is not significantly different from zero, with a symmetrical distribution and normal-like tails.

Descriptive Statistics		Significance Test			5.00%	Test	p-value	SIG?
		Target	P-Value	SIG?		White-noise	23.79%	TRUE
AVERAGE:	0.01%	0.000	35.16%	FALSE		Normal Distributed?	13.32%	TRUE
STD DEV:	0.73%					ARCH Effect?	60.83%	FALSE
SKREW:	-0.17	0.000	5.66%	FALSE				
EXCESS-KURTOSIS:	0.29	0.000	10.87%	FALSE				
MEDIAN:	0.06%							
MIN:	-2.40%							
MAX:	2.36%							
Q 1:	-0.44%							
Q 3:	0.46%							

From the table on the bottom right, we can see that daily returns are more like a Gaussian distribution, and thus they possess no information.



Discussion

The daily returns act more like Gaussian noise, so the past returns do not reveal any information that can help us better predict the future points. How about non-linear modeling techniques (e.g. neural networks, genetic algorithms)?

In my opinion, I would rather explore and bring new information to the analysis than spend my time utilizing sophisticated modeling techniques. In our analysis, we try to avoid the over-fitting trap and to build a robust model for reliable forecasting.

Daily High-Low time series

While we restricted our analysis to market open and close, which happen at precise times in the day, the HI and LO represent the Max-Min rates of the intra-day rates. Furthermore, they are impossible to identify for certain, so we can't trade on them.

Nevertheless, the daily HI and LO rates bring new information that is not already in the open-close time series; as a result, we can forecast their relative levels to a certain extent.

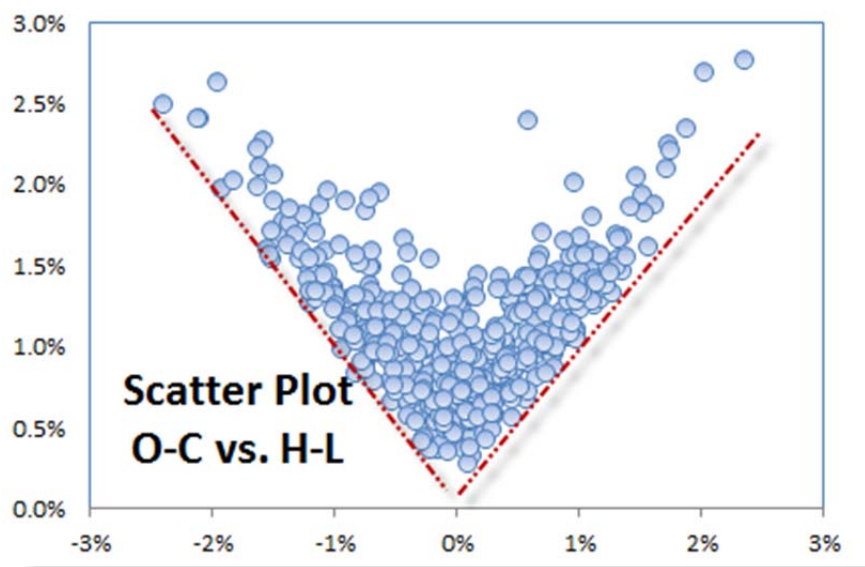


Figure 1: scatter plot for open-close returns vs. HI-LO returns

In the figure above, please note that HI-LO returns are related to the volatility of the open-close returns.

Daily returns with stop-loss


Using the sample data, we implemented a sell stop-loss order with a level set on a daily basis at market open.

Definition: A stop loss is an order to buy or sell an asset (e.g. trading security or commodity) once the price of the asset reaches a specified price, known as the stop price. A Sell stop order is used here to limit a loss on a trading commodity we own.

Using the daily LO rate, we examine whether the stop price has breached for the day; if it breaches, we set the sale price to stop price (no slippage or transaction fees is assumed), otherwise we continue to use the market close price for the sale price.

In this exercise, we examined the following stop price techniques: (1) Fixed percentage below the market open price (e.g. 0.75%), (2) Fixed offset below the market open price.


1. Stop price is set at a fixed percentage below market open



Examining the summary statistics, the data exhibits positive skew and a marginally significant average.

Descriptive Statistics		Significance Test			5.00%	Test	p-value	SIG?
AVERAGE:	0.05%	Target	P-Value	SIG?		White-noise	40.58%	TRUE
STD DEV:	0.64%	0.000	5.47%	FALSE		Normal Distributed?	0.00%	FALSE
SKEW:	0.49	0.000	0.00%	TRUE		ARCH Effect?	30.93%	FALSE
EXCESS-KURTOSIS:	-0.28	0.000	9.24%	FALSE				
MEDIAN:	0.04%							
MIN:	-0.75%							
MAX:	2.36%	Annual						
Q 1:	-0.54%	%RET	11.6%					
Q 3:	0.45%	%STD	10.2%					

To optimize the percentage settings for the stop price, we constructed a data table where we varied the percentage between 0.1% and 2%.



Plotting the mean and volatility of the trading strategy, we picked 0.75% as the best choice. This setting gives the new trading strategy an annual return of 11.6% and a volatility of 10.2%.

Furthermore, the new strategy's returns act like white noise (but not Gaussian), so we have not left any information on the table.

2. Stop price is set at an offset below market open

In an attempt to improve the stop price strategy, we used the EWMA estimate and the WMA of the HI-LO. We followed the same steps: generate strategy returns, use a data table to examine different settings and, finally, pick the best one. The results are comparable to the simple one we illustrated above.

For detailed analysis and illustration, please refer to the Excel worksheet file (available online) for details.

Conclusion

So far, we started from a Gaussian noise data sample, incorporated daily LO rates, and devised a new strategy with stop-loss that not only limits the daily loss, but generates positive returns.

In addition, we can devise another strategy that employs a sell limit order to lock in profits. Using daily HI, we can examine whether the limit price is breached and construct new strategy returns.

How about a strategy that employs both: stop-loss and limit order? Unfortunately, the data set we have doesn't contain information about whether the HI occurs before or after the LO that day.

The takeaway here is that we can start with a data series that looks no more than Gaussian white noise, but with the introduction of new information (e.g. Daily HI and LO, volume, interest-rate parity, G8 summit events, etc.), we can derive a new time series with more favorable characteristics.