Optimization methods for entry and exit points in relative value trading

Introduction

In this article, we explore how optimising entry and exit points can improve the performance of mean reversion strategies, specifically within a purchasing power parity (PPP) framework. We address this optimisation challenge, by exploring two complementary approaches: a theoretical model-based approach and a numerical method. Due to the simplicity of the intuition behind the approach and their consistency across asset classes, mean reversion strategies are often particularly attractive to investors. However, the success of such strategies depends largely on how accurate of the trading bands, i.e. the entry and exit thresholds, are defined. Accurate calibration is of major importance, as setting these bands too tightly can imply excessive transaction costs due to the multitude of executed trades, while setting them too wide misses valuable opportunities. In this article, we apply an approach proposed by Zeng and Lee (2014) that optimizes entry and exit points on the assumption that deviations from PPP follow an Ornstein-Uhlenbeck (OU) process.

A choice of a linear RV model for currencies: Purchasing Power Parity

Purchasing Power Parity (PPP) is a popular approach to monetary theory that posits that in the long run, currency returns should be driven by cross-country inflation differentials. While the theory provides a compelling rationale for long-term currency valuation, short run FX rates often fluctuate heavily and deviate from PPP predictions. Short-run deviations are very common and are often influenced by factors ranging from market frictions to economic policies such as capital controls. In this article, we explore whether these deviations exhibit mean-reverting behaviour and whether such dynamics can be exploited.

We run PPP regressions on 15 currencies against the U.S. Dollar in the following form:

$$r_t = \alpha + \beta \pi_t + \varepsilon_t$$

Where r_t is the monthly rate of appreciation/depreciation of a currency and π_t is its inflation differential with the USD. Since the inflation data we used is monthly, while FX data was provided on a daily basis, we make monthly predictions every time a new inflation data point becomes available, and during the month we measure dislocations in the daily FX rate from that month's prediction. While PPP is conventionally a long-run theory, we chose a window of 1 month to measure returns and inflation differentials because we discovered that longer windows (e.g. 1 year) would slow down reversion speeds and make relative value trading less attractive.

A common approach is to then normalize the residuals by demeaning and standardizing them, before generating signals. We instead chose to create the signal-generating series by transforming the residuals with a sigmoid function (in our case, $\tanh$ to scale the data in the interval $[-1,1]$ and to ensure that extreme values do not overinfluence trading decisions.

Next, we'll explore the approach proposed by Zeng and Lee (2014) and how we implemented it practically.

Analytical determination of optimal Entry and Exit points

Lin and Barziy (2021) rightly point out the two core challenges in developing a pairs trading strategy. The first is identifying which assets to pair in order to form a process exhibiting mean-reverting behavior. The second is determining when to trade – essentially, how to select optimal entry and exit points for the strategy.

Traditionally, these `entry and exit decisions are made by defining thresholds tied to the standard deviation of the spread or dislocation between asset prices. A trade is triggered when the spread hits a multiple of its standard deviation, with the proportionality constant typically found through empirical optimization. This optimization usually targets either the total return of the strategy or its Sharpe ratio.

The approach taken in the Hudson paper by Lin and Barziy (2021) deviates from this empirical framework, aligning instead with the theoretical methods introduced by Bertram (2010) for long-only strategies and extended by Zeng and Lee (2014) to allow for both long and short positions. Rather than relying on backtested parameter fitting, this line of work seeks to analytically optimize the expected return per unit of time.

Assuming the price spread between the paired assets follows an exponential Ornstein-Uhlenbeck process, the problem is reframed in terms of the first-passage time of the process. This makes it possible to derive closed-form expressions for the expected trade duration, its variance, the expected return per unit of time, and the variance of that return. Based on these results, optimal trading thresholds can then be determined by maximizing either the expected return or the Sharpe ratio.

The key is to consider the entire trading cycle—from entry to exit—as a first-passage event of a mean-reverting process hitting a two-sided boundary. Using the parameters of the OU process, expressions for expected returns and variances are derived and subsequently optimized. The analytical tractability of the problem is enhanced by switching to dimensionless variables, which simplifies the resulting equations. In the end, one obtains a polynomial equation of infinite (but decaying) degree, which must be solved and then transformed back into dimensional terms using the original OU parameters.

When searching for an exit point b for a trade between 0 and its entry point a, with transaction costs c the maximisation problem solved in Zeng and Lee (2014) is the following:

$$\max_{a,b} \frac{a-b-c}{\mathbb{E}(\tau_1) + \mathbb{E}(\tau_2)} = \frac{a-b-c}{\frac{1}{2}\sum_{n=0}^{+\infty} \frac{\left(\sqrt{2}a\right)^{2n+1} - \left(\sqrt{2}b\right)^{2n+1}}{(2n+1)!}\Gamma\left(\frac{2n+1}{2}\right)}$$

$$u.c. \quad 0 \leq b \leq a - c$$

Once they show that $\partial f / \partial b < 0$ for any a, they then set $\partial f / \partial a = 0$ to find a^* through the following equation:

$$\frac{1}{2}\sum_{n=0}^{+\infty} \frac{\left(\sqrt{2}a\right)^{2n+1}}{(2n+1)!}\Gamma\left(\frac{2n+1}{2}\right) = (a-c)\frac{\sqrt{2}}{2}\sum_{n=0}^{+\infty} \frac{\left(\sqrt{2}a\right)^{2n}}{(2n)!}\Gamma\left(\frac{2n+1}{2}\right)$$

They then optimize the same $f(a,b)$ for $-a \leq b \leq 0$.

$$\max_{a,b} f(a,b) = \frac{a-b-c}{E(\tau_1)+E(\tau_2)} = \frac{a-b-c}{\frac{1}{2}\sum_{n=0}^{+\infty}\frac{(\sqrt{2}a)^{2n+1}-(\sqrt{2}b)^{2n+1}}{(2n+1)!}\Gamma\left(\frac{2n+1}{2}\right)}$$

$$s.c. \quad -a \le b \le \min(0, a-c)$$

Where they show that the for the gradient of $f(a,b)$ to be zero, the only solution can only be $a^2 = b^2$, so that $b^* = -a^*$. They then solve the following equation for a^*:

$$\frac{1}{2}\sum_{n=0}^{+\infty}\frac{(\sqrt{2}a)^{2n+1}}{(2n+1)!}\Gamma\left(\frac{2n+1}{2}\right) = \left(a - \frac{c}{2}\right)\frac{\sqrt{2}}{2}\sum_{n=0}^{+\infty}\frac{(\sqrt{2}a)^{2n}}{(2n)!}\Gamma\left(\frac{2n+1}{2}\right)$$

The authors also show that the second maximisation problem always yields greater maximal returns than the first, when transaction costs are present. We implemented a root-finding solver for the largest terminal sum index that we could find that wouldn't result in computing errors.

One of the key findings of Zeng and Lee (2014) is that under the above-mentioned assumptions, when transaction costs are non-zero, the optimal exit points for a long (short) trade should not be set to the long-run mean of the series, but rather to the optimal short (long) entry points, thus allowing for no waiting time between trades.

Estimation of Ornstein-Uhlenbeck parameters

When implementing this particular approach with respect to our PPP strategy, we estimate OU parameters on PPP dislocations from the first half of the data. Specifically, we require μ, θ, σ to adapt the optimized result to the nature of our dislocations. In order to derive these parameters with respect to the dislocations we've observed on our PPP strategy, we can use OLS on a discretized OU process. Alternatively, as other papers have employed in past parameter estimations, MLE is a suitable candidate. However, the computation of this estimation approach is much more computationally expensive, and given standard properties of the log likelihood function, MLE and OLS should theoretically produce equivalent parameter estimates due to the Markov property of OU processes. When performing OLS, we take a discretized version of our OU, particularly of the form:

$$X_t = \alpha + \beta X_{t-1} + \varepsilon_t$$

Discretizing the OU process into, essentially, an AR(1) process is useful because it allows us to use OLS to estimate the α and β terms, which can then be plugged back into the standard OU using the following relations:

$$X_{t+\Delta t} = \theta(1 - e^{-\kappa \Delta t}) + e^{-\kappa \Delta t}X_t + \int_t^{t+\Delta t}\sigma e^{\kappa(s-t)}dW_s$$

$$\mu = -\frac{\ln\beta}{dt}, \qquad \theta = \frac{\beta}{1-\alpha}, \qquad \sigma = \sqrt{\frac{2\kappa \,\text{Var}(\varepsilon)}{1-\beta^2}}$$

Now, how do we apply this method in such a way which avoids forward-looking bias? Simple: we take an expanding sample of the dislocations and re-adjust the parameters for our OU at each particular time-step (in our case, days). You might be wondering: "*Is there a specific reason we prefer an expanding approach to a rolling one in this instance?*" And, in fact, the answer is yes. Using an expanding sample is quite useful when you don't expect the nature of your process' parameters to change dramatically through time, which is different from something like a rolling window, which assumes that older datapoints are irrelevant to future forecasts. Since we believe the former assumption to hold true, the more datapoints, the merrier.

Results and Findings

We run the above-mentioned procedures on 16 years of data from 2008 to 2024 on AUD, BRL, CAD, CLP, CNY, DKK, EUR, GBP, ILS, JPY, KRW, MXN, MYR, NOK and NZD. We calibrate optimal entry/exit points only once at the end of the first half of the data (the train set) and apply those optimal bands to the rest of the data (the test set). More frequent calibration of optimal entry/exit points could be an interesting subject for expansion of this article.
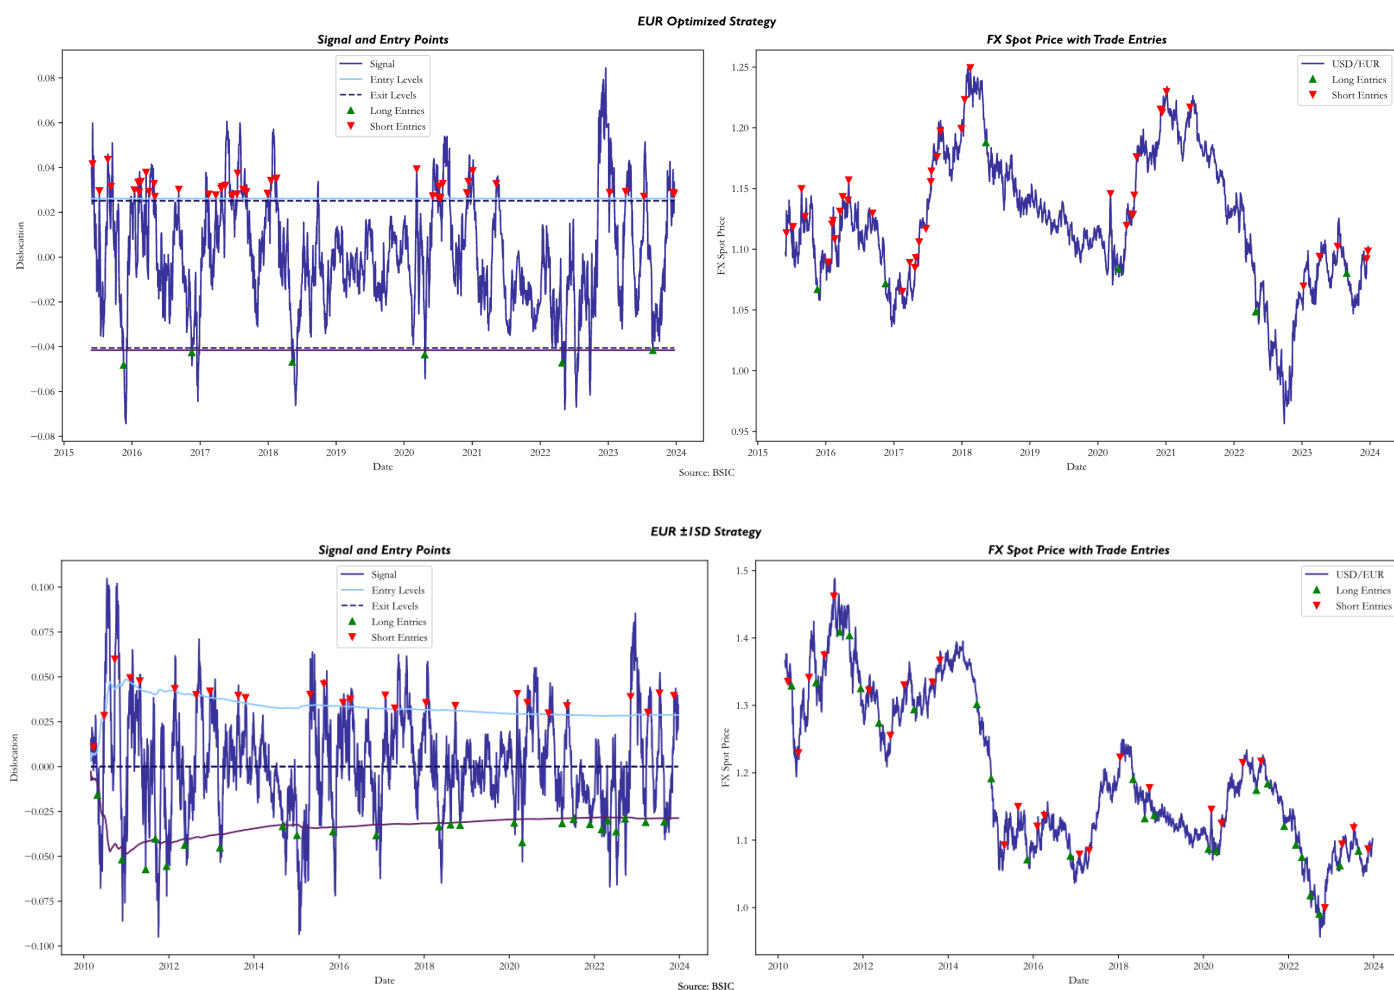
The main subject of interest of the article isn't whether the PPP strategy produces excess returns compared to a benchmark, but rather whether the optimal calibration of entry/exit points produces greater profits compared to an arbitrarily set trading rule. As a benchmark, we chose to display the performance of trading dislocations with a

+1/-1 standard deviation rule. The arbitrary trading rule functions as follows: when the transformed dislocation series is at least one standard deviation below (above) its mean, a long (short) signal is generated; positions are cleared when the transformed dislocation series crosses its mean.

Rebased Cumulative P&L from 2016 (GBP)

Rebased Cumulative P&L from 2016 (EUR)
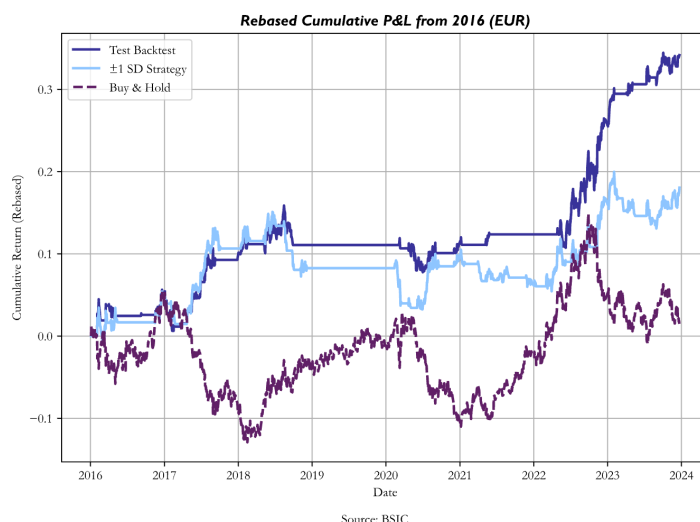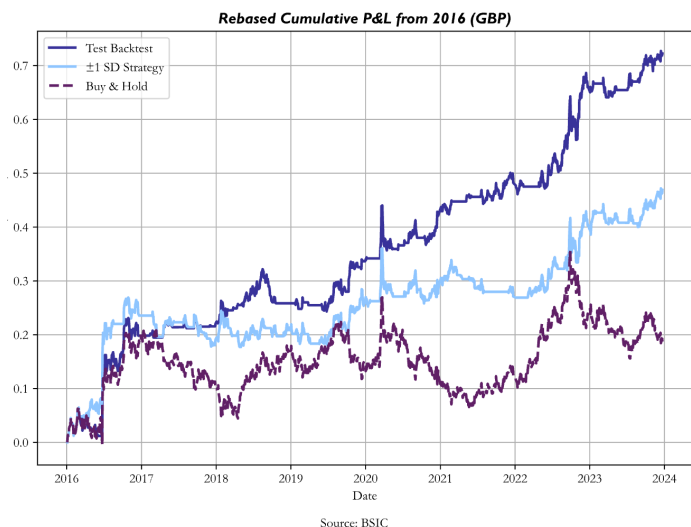
Risk metrics on GBP strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	1,5657	2,1651	1,1802	0,5534	8,35%	8,38%
±1 St. Dev. Strategy	0,9846	1,3810	1,1187	0,5285	8,60%	10,26%
Buy & Hold Benchmark	0,2407	0,3859	1,0502	0,4984	10,10%	20,38%

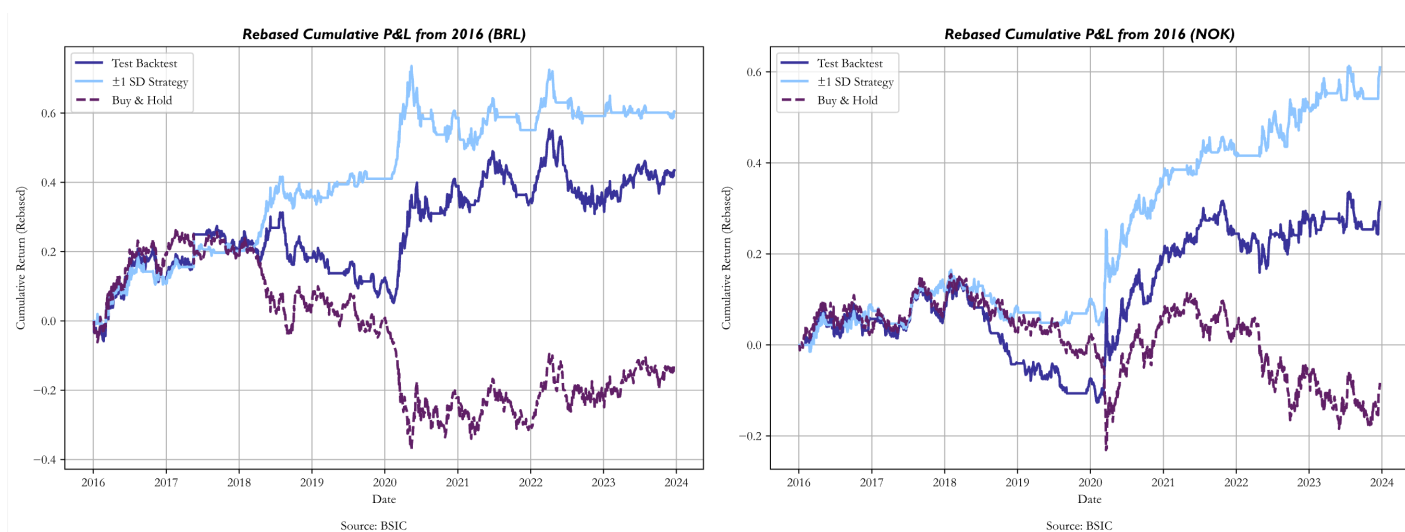| Risk metrics on EUR strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	1,2578	1,3714	1,1475	0,5561	5,40%	8,12%
±1 St. Dev. Strategy	0,6908	0,8217	1,0516	0,5316	5,51%	11,86%
Buy & Hold Benchmark	0,0200	0,0291	0,9799	0,5059	7,78%	18,36%

Source: BSIC

The optimized strategy outperformed (both in terms of Sharpe Ratio and Cumulative Returns) the arbitrary trading rule on the test set on all currencies except CAD, DKK, ILS, KRW, MYR and NOK. A factor that seems to have driven the underperformance of the optimized strategy when trading these 6 currencies against the dollar is linked to non-stationarity of the signal series. In particular, it seems that mean hitting times for trades on these 6 currencies against the USD were somewhat different between train set and test set data. This implies that the OLS parameters that were estimated on the train set caused the optimization process to return relatively wide trading bands, and in a test set with slower mean reversions, a set of relatively tighter trading bands would have been preferrable. Conversely, where mean reversion speeds were greater in the test set than they were in the train set, it could've been preferrable to implement a wider pair of bands, to avoid incurring into excessive transaction costs and benefitting from larger swings.

With certain currencies (especially DKK) we also encountered volatility differences in the signal series between the train set and the test set. When trading bands are fine-tuned to a train set series that is more volatile than the test set one, they tend to be wider, thus decreasing the number of trades as well as increasing position holding times. Conversely, when the signal-generating series tends to be less volatile in the train set data than in the test set data,

more signals are generated, but without capturing larger swings, and incremental transaction costs potentially outweigh incremental profits.

Source: BSIC

Risk metrics on BRL strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,5428	0,7781	1,0380	0,5166	14,26%	26,15%
±1 St. Dev. Strategy	0,7968	0,9630	1,0818	0,5266	12,51%	24,23%
Buy & Hold Benchmark	-0,1112	-0,1572	0,9746	0,5016	15,44%	63,46%

Risk metrics on NOK strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,5924	0,8372	1,1601	0,4928	11,44%	26,89%
±1 St. Dev. Strategy	0,9613	1,2285	1,1618	0,5176	10,43%	12,17%
Buy & Hold Benchmark	-0,0875	-0,1254	1,0406	0,4862	12,73%	38,61%

Source: BSIC

Conclusions and Final Remarks

Our investigations should serve as a first approach to the subject of determining optimal entry and exit points. Further experiments could (and should) be conducted. A first result is that a rule of thumb does not always perform as well as it sounds. The second key result is that the data points in favour of the presented approach being a generally valid one, but with potential for fallacies due to the underlying assumptions on the signal series.

We identified signs that OU parameters may not be stable for very long periods of time (hence volatility and mean reversion speed regime shifts), which could be dealt with by either implementing dynamic drift and diffusion models, or by changing the way in which OLS parameters are estimated (more frequently and/or on shorter timespans). Additionally, since the focus of the article is to present this method as a general approach to RV trading across asset classes, we should also consider that a specific parametric form (as the OU form we imposed) might not be the best fit for all mean-reverting signal series. Experimenting with nonparametric calibration of drift and diffusion coefficients could be a way to tackle the problem. Alternatively, numerical approaches that optimize bands directly on the observed data by maximizing could also be interesting to investigate, if not too computationally expensive.

Lastly, we also suggest that – provided that a solid parameter estimator and model specification be found for the signal series – a Monte Carlo simulation of signal paths could be useful to test the validity of estimated bands, by computing returns on each of those paths using the bands estimated in-sample. This, along with the previously mentioned suggestions on further areas of investigation, provides a decent amount of content to further investigate in the future.

Appendix: Risk Metrics for other currency pairs

Risk metrics on AUD strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,9294	0,9583	1,1221	0,5341	8,62%	16,16%
±1 St. Dev. Strategy	0,6830	0,7483	1,0831	0,5231	8,64%	14,24%
Buy & Hold Benchmark	0,1312	0,2009	1,0120	0,5027	11,02%	34,99%

Risk metrics on CAD strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,3137	0,4471	1,0634	0,4991	6,58%	20,35%
±1 St. Dev. Strategy	0,3589	0,3899	0,9838	0,5254	5,61%	16,39%
Buy & Hold Benchmark	0,1350	0,2084	1,0413	0,4957	7,18%	17,40%

Risk metrics on CLP strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,6981	0,8889	0,9833	0,5378	10,87%	27,45%
±1 St. Dev. Strategy	0,4320	0,5084	0,9637	0,5318	10,42%	19,11%
Buy & Hold Benchmark	-0,1626	-0,2431	0,9762	0,4992	12,38%	50,94%

Risk metrics on CNY strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	2,2215	3,1099	1,2367	0,5561	4,22%	5,34%
±1 St. Dev. Strategy	1,9516	2,3760	1,2279	0,5568	3,94%	3,91%
Buy & Hold Benchmark	-0,2627	-0,3810	1,0365	0,4787	4,43%	15,51%

Risk metrics on DKK strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,3339	0,5038	1,0824	0,4956	7,11%	17,59%
±1 St. Dev. Strategy	0,7171	0,8520	1,0655	0,5294	5,72%	12,26%
Buy & Hold Benchmark	0,0601	0,0931	1,0200	0,4976	7,89%	25,51%

Risk metrics on JPY strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,9909	1,3706	1,1940	0,5109	8,41%	9,60%
±1 St. Dev. Strategy	1,0023	1,1223	1,1717	0,5290	7,59%	7,82%
Buy & Hold Benchmark	-0,1928	-0,2952	1,0651	0,4755	9,62%	39,16%

Risk metrics on ILS strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,7668	0,8648	1,0394	0,5326	7,05%	10,20%
±1 St. Dev. Strategy	0,3721	0,3722	0,9912	0,5254	6,96%	11,73%
Buy & Hold Benchmark	0,1787	0,2203	0,9280	0,5270	8,40%	27,16%

Risk metrics on KRW strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,3745	0,5353	1,1274	0,4890	6,58%	17,44%
±1 St. Dev. Strategy	1,1827	1,4795	1,2567	0,5105	6,69%	6,76%
Buy & Hold Benchmark	-0,1386	-0,2119	1,0305	0,4866	8,11%	29,50%

Risk metrics on MXN strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,7359	0,9354	0,9809	0,5402	12,06%	20,52%
±1 St. Dev. Strategy	0,5823	0,6033	1,0006	0,5350	10,74%	18,81%
Buy & Hold Benchmark	0,0888	0,1107	0,8505	0,5443	13,13%	34,81%

Risk metrics on MYR strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	0,6499	0,8789	1,1767	0,4941	5,00%	12,78%
±1 St. Dev. Strategy	0,8447	0,8733	1,1153	0,5281	4,64%	10,39%
Buy & Hold Benchmark	-0,1383	-0,2018	1,1220	0,4648	5,73%	21,01%

Risk metrics on NZD strategy returns						
Strategy	Sharpe Ratio	Sortino Ratio	Skew	Hit Ratio	Ann. Volatility of Returns	Max Drawdown
Optimized strategy	1,0051	1,1378	1,0436	0,5481	8,26%	12,03%
±1 St. Dev. Strategy	0,6278	0,7476	0,9774	0,5408	8,57%	10,23%
Buy & Hold Benchmark	0,1398	0,2273	1,0072	0,5040	11,07%	27,57%

Source: BSIC

References

[1] Zeng, Z. & Lee, C. "Pairs trading: optimal thresholds and profitability", 2014

[2] Lin, A. & Barziy, I. "Optimal Trading Thresholds for the O-U Process", 2021

[3] Bertram, W. K. "Analytic solutions for optimal statistical arbitrage trading", 2010

TAGS: Foreign Exchange, Ornstein-Uhlenbeck processes, trading bands, linear regression, PPP