DYNAMIC BETA-NEUTRAL PORTFOLIO BASED ON MACROECONOMIC INDICATORS

by

Vadym Nahorniak

A thesis submitted in partial fulfillment of
the requirements for the degree of

MA in Economic Analysis.

Kyiv School of Economics

2026

Thesis Supervisor: _____ Professor Olesia Verchenko

Approved by _____
Head of the KSE Defense Committee, Professor

_____

_____

_____

Date _____

Kyiv School of Economics

Abstract

DYNAMIC BETA-NEUTRAL PORTFOLIO BASED ON MACROECONOMIC INDICATORS

by Vadym Nahorniak

Thesis Supervisor: Professor Olesia Verchenko

FX markets never stop moving, which makes it genuinely difficult to separate what a portfolio earns on its own from what it picks up just by riding a broad dollar or euro trend. A beta-neutral setup addresses this by offsetting long and short positions until the net market sensitivity is close to zero. What is left after that reflects how the individual pairs actually performed, not which way the market happened to go that day. This study builds and tests exactly that kind of portfolio across ten major USD and EUR currency pairs over the 2020 to 2024 period.

TABLE OF CONTENTS

LIST OF FIGURES iv

LIST OF TABLES v

AI USE DECLARATION vi

LIST OF ABBREVIATIONS vii

Chapter 1. INTRODUCTION 1

Chapter 2. LITERATURE REVIEW 5

Chapter 3. METHODOLOGY 9

3.1 Data transformation 9

3.2 PCA 12

3.3 LDA 13

3.4 Smart Beta 18

3.5 Rebalancing 20

3.6 Trading model 22

Chapter 4. DATA 26

Chapter 5. ESTIMATION RESULTS 36

5.1 Comparison of LDA with and without macroeconomic variables 36

5.2 Constructing a Beta-Neutral Portfolio 42

5.3 Simulated trading results 47

Chapter 6. CONCLUSIONS AND POLICY RECOMMENDATIONS 51

6.1 General conclusions 51

6.2 Contribution and Limitations 51

TABLE OF CONTENTS — Continued

WORKS CITED 53

LIST OF FIGURES

Number *Page*

Figure 1. Daily FX Prices (2010–2024). 27

Figure 2. Daily FX Return (2010–2024). 28

Figure 3. DXY and EUR_PURE Index Dynamics (2010–2024). 30

Figure 4. Correlation Structure of USD and EUR Baskets. 31

Figure 5. U.S. Macroeconomic Indicators (2010–2024). 32

Figure 6. Euro-Area Macroeconomic Indicators (2010–2024). 34

Figure 7. LDA Forecast Accuracy by Year. 37

Figure 8. LDA with Macro Forecast Accuracy by Year. 39

Figure 9. Top-10 USD Macro Variable Combinations. 41

Figure 10. Top-10 EUR Macro Variable Combinations. 42

Figure 11. FX Benchmark vs DXY and EURO_Pure. 43

Figure 12. Distribution of Total Returns. 44

Figure 13. Distribution of Maximum Drawdowns. 45

Figure 14. Correlation matrix by portfolio factors. 46

Figure 15. Cumulative Returns: All Strategies. 47

Figure 16. Final Return Distribution. 48

Figure 17. Skewness Details: Using a Portfolio with the Best Returns. 49

Figure 18. Compare Drawdown and Return. 50

LIST OF TABLES

Number *Page*

Table 1. Descriptive Statistics of FX Pairs (2010–2024). 29

Table 2. Descriptive Statistics of Currency Indices (2010–2024).....30

Table 3. Descriptive Statistics of U.S. Macroeconomic Variables (2010–2024). .33

Table 4. Descriptive Statistics of Euro Macroeconomic Variables (2010–2024).35

AI USE DECLARATION

During the preparation of this thesis, Generative AI (Claude, Grammarly) were used as a supplementary tool to:

support grammar correction and improvement of sentence clarity in Chapters I-VI of this thesis

idea brainstorming on a) LDA classifier implementation, b) portfolio optimization.

Any suggestions produced by the AI tool were critically evaluated and substantially modified by the author before being incorporated into the thesis. These tools were not used to generate research findings, produce literature summaries used directly in the thesis, create citations, or conduct any form of data analysis. All research questions, theoretical framing, methodology, analysis, interpretation of results, and conclusions were developed independently by the author. The author takes full responsibility for the originality, validity, and academic integrity of the work presented in this thesis

LIST OF ABBREVIATIONS

CVaR. Conditional Value at Risk.

DXY. U.S. Dollar Index.

EXY. Euro Currency Index.

EURO_PURE. Beta-Adjusted Euro Currency Index.

FX. Foreign Exchange.

LDA. Linear Discriminant Analysis.

PCA. Principal Component Analysis.

VIX. CBOE Volatility Index.

TNX. 10-Year Treasury Yield.

NAV. Net Asset Value.

IRX. 13-Week Treasury Bill Rate.

HICP. Harmonised Index of Consumer Prices.

Chapter 1

INTRODUCTION

Quick advancement of global financial markets has outlined the crucial role of the right and effective risk management in the investing process in order to maintain stability and constant growth. Among the promising approaches, one is the design and implementation of beta-neutral portfolios. The idea behind a beta-neutral portfolio is to cancel out the sensitivity to the general market movements by taking into account assets with different levels of dependence on the market. This is a very relevant concept in the current volatile situation, especially during a financial crisis, when systemic risks dismantle conventional investment strategies. Therefore, this proposed research tries to look for ways in which beta-neutral strategy can be optimized in its response to macroeconomic fluctuations. FX markets never stop moving, and that creates a specific measurement problem: when a position makes money, it is not always clear how much of that came from the pair itself versus a broad dollar or euro swing that lifted the whole basket. Balancing long and short positions until the net market exposure sits near zero is one way to cut through that noise. This study builds and tests exactly that kind of portfolio across ten major USD and EUR currency pairs over 2020 to 2024, with weights updated daily and risk aversion tied to macro regime signals.

Both the 2008 crisis and COVID-19 showed what happens when correlations shift faster than a static model can adjust (Reinhart and Rogoff, 2009; Baker et al., 2020; Foglia, Recchioni, and Polinesi, 2021). A beta-neutral setup at least ensures that the losses, when they come, reflect pair-specific moves rather than a market bet that was never intentional.

This study important due to the part of macroeconomic indicators take in shaping portfolio performance is evolving. The indicators that have to become a significant part in the risk and return dynamic system, are including interest rates, inflation, GDP growth, and unemployment rates. Foglia, Recchioni, and Polinesi (2021) has showed us how all these indicators affect portfolio stability and how during crisis periods, these factors effects are generally disrupting. For sure, the increase in interest rates normally results into a decline in valuations of those assets that are sensitive to borrowing funds, while real returns are diminished by inflationary pressures. Combination of these factors in combination with increased sensitivity during crisis periods will probably make traditional approaches to portfolio management less effective.

Despite the increasing recognition of these challenges, major drawbacks are associated with the already existing beta-neutral strategies. Main flaw of this approach is static positions. Traders often design them to be under predicted not dynamic conditions. However, in modern fast-paced world, such approach is annihilated by rapid and radical changes that happens all the time. Such approach additionally leaves an asset portfolio vulnerable to external shocks, ruining the concept of beta neutrality. Overcoming this limitation making the approach more dynamic. It adapts it with real-time adjustment of composition, based on changes of current macroeconomic indicators. With implementation of this idea strategy would likely not only keep beta neutrality, but also increase the resilience of the portfolio to volatility, thereby maximizing returns. Resulting in returns maximization.

The main motivation of this portfolio design is to achieve returns rate close to U.S. Treasury bonds but with higher variation of asset distribution and risk management flexibility. Such approach would benefit following market stakeholders as: · Companies with massive foreign currency reserves; · Structured product designers who require stable, low-beta returns as a

foundation. Gaining bond-like performance with dynamic rebalancing, will result portfolio becoming vital alternative to traditional investment strategies. Hopefully, with greater liquidity and adaptation to chaos market conditions.

The main goal of this study is to investigate how macroeconomic indicators affect the performance of different beta-neutral portfolios. Because beta-neutral strategies are designed to reduce market risk, but most existing methods are not quite adaptive to changes in the economy. The problem this study will be addressed by using regression-based beta risk estimations and dynamic rebalancing strategy, which allows for real-time portfolio adjustments and more response efficiency. Hence changes in interest rates, inflation, GDP growth, and unemployment will not have such impact.

The methodological framework in this study, is addressed to these objectives:

1. Primarily, construction of a beta-neutral currency portfolio, that include two major USD and EUR currency pairs. These currencies would represent good candidates for the particular research exercise. Given that their key role in trade and international financial transactions. Sets the case for investigating various influences stemming from macroeconomic perspectives.

2. Using regression models, alike to used in benchmark indices, DXY for USD and EXY for EUR, there are going to be calculated coefficients for each currency. This analysis will help the study identify the sensitivity of according currency pairs to market movements and macroeconomic changes.

3. Rolling beta coefficients are estimated daily for each currency pair against the sector benchmark indices, following Black (1993). Macroeconomic variables are compressed into PCA factors before entering the LDA classifier, where the regime LDA classification

design follows Foglia, Recchioni, and Polinesi (2021). Portfolio weights are then determined by a CVaR optimizer where the risk aversion parameter shifts between bullish and bearish modes according to the LDA signal.

The dollar loading is stripped from the euro index per Lustig, Roussanov, and Verdelhan (2011). A conditional skewness filter from Harvey and Siddique (2000) manages the leverage overlay. The portfolio rebalances when beta drifts past a threshold or when cumulative P&L crosses volatility-scaled boundaries, rather than on a fixed schedule.

To maintain portfolio neutrality without inducing excessive turnover, the model avoids fixing asset weights at the start of each period, relying instead on an event-driven rebalancing mechanism as its core feature. Portfolio adjustments are triggered dynamically whenever beta drifts beyond a specified threshold or cumulative gains and losses cross adaptive, volatility-scaled boundaries.

Chapter 2 reviews the literature, Chapter 3 lays out the methodology, Chapter 4 describes the data, Chapter 5 presents the empirical results, and Chapter 6 concludes the study.

Chapter 2

LITERATURE REVIEW

This chapter reviews the literature across three areas relevant to the study: beta-neutral portfolio construction, macroeconomic determinants of risk premia, and regime-based portfolio management. Each section builds toward the gap that the present study addresses.

Traditional asset allocation frameworks remain conceptually centered on Markowitz's (1952) insight that combining low-correlation assets reduces portfolio risk without sacrificing expected returns. However, a major limitation of these models is their failure to distinguish between market-wide systematic risk and asset-specific idiosyncratic risk. It is precisely around this unresolved distinction that beta-neutral strategies are constructed.

Black (1993) quantified market sensitivity and showed that a portfolio where individual betas cancel out is, in theory, insulated from broad shocks. Papageorgiou, Reeves, and Xie (2012) tested this in actual crisis periods and found that correlations move in ways static models do not anticipate, so a portfolio that started neutral can drift quite far from target within weeks. Clarke, de Silva, and Thorley (2020) showed that adjusting weights dynamically handles volatile markets better than any fixed approach. Avramov, Barras, and Kosowski (2013) found that conditioning portfolio decisions on macroeconomic variables keeps performance more stable across different market environments.

The concept of beta-neutrality has been invented in order to keep safe portfolios from the systemic risks inherent in the broader market. In 1993, Fischer Black published research on how beta measures the sensitivity of an

asset to market fluctuations. Portfolios with a near zero betas are theoretically protected from overall market movements. Such portfolios would provide stability during periods of high volatility. Black also figured out that beta-neutral strategies might work worse if there were complex relationships between assets and market conditions, especially during periods of economic downturns or crises. This has resulted in increasing interest in using dynamic approaches to maintain beta neutrality through harsh periods.

The difficulty with market neutrality becomes especially clear during financial crises. Papageorgiou, Reeves, and Xie (2012) studied these challenges. Authors outlined that it is indeed challenging to maintain market neutrality due to asset correlations being very volatile in nature. Most of them change over time unpredictably, which is extremely viewed during crisis periods. Their results highlight that portfolio rebalancing needs to be done with a regular frequency to maintain beta neutrality in chaotic markets. Such idea crosses with the current study idea, so the aims is to devise a framework for dynamic rebalancing and maintaining neutrality during changes in various macroeconomic environments.

When working with currencies, a big problem is the global dollar factor. Lustig, Roussanov, and Verdelhan (2011) showed that if we want to find clean risk premiums, we have to extract these common factors first. If we don't clean out the dollar driver, the individual currency betas get mixed up with global market moves.

Macroeconomic indicators such as interest rates, inflation, GDP growth, and unemployment are key determinants of market dynamics and, therefore, portfolio performance. Lettau and Ludvigson (2001) showing that such risks are very sensitive to these kinds of macroeconomic changes, often in time-varying ways. Their work provided very clearly understanding that

macroeconomic variables such as GDP growth and inflation affect the individual asset return and influence overall risk-return tradeoff. This realization shows us an important insight in portfolio optimization models to use those dynamic relationships.

In the same way, these factors are met during the COVID-19 pandemic, studied by Foglia, Recchioni, and Polinesi (2021). This happened due to the sensitivity of financial markets in most chaotic times in 21st century increases towards changes in macroeconomic indicators same as in various other crisis faced previously. For example, in the case of rate increase, asset estimation fall because of the increased cost of borrowing, and also higher inflation rates has influenced returns on that. This pandemic has created a natural experiment for many of these dynamics and shown the vulnerability of traditional portfolio strategies to adjust to those changes.

Lustig, Roussanov, and Verdelhan (2011) show that risk premia in FX markets are structured around common factors, primarily a global dollar factor and a carry factor**.** Dollar cross rates respond primarily to US conditions. Euro cross rates are formed by euro-area dynamics. This finding motivates the construction of separate classifiers for each basket rather than a single model across the full FX scene. Lettau and Ludvigson (2001) further establish that business cycle variables carry information about time-varying risk premia that unconditional means miss entirely. Supporting the inclusion of GDP growth and unemployment as regime conditioning variables.

The existing literature has established that beta-neutral strategies reduce systematic risk exposure. Macroeconomic variables carry information about time-varying risk premia, and that conditioning on regime signals improves out-of-sample performance. What remains underexplored is the integration of all three within a single FX-specific framework. Most beta-neutral strategies

in the literature target equity markets and do not take into consideration the sector-level structure of currency returns. Most macro-conditioned portfolio frameworks treat the FX universe as homogeneous. Also, they do not build sector-specific regime signals separately for USD and EUR.

The present study attempts to close these gaps. It combines rolling beta estimation, sector-specific LDA and CVaR-based optimization within a fully dynamic rebalancing framework. LDA classifiers are conditioned on separate macro variable sets for the USD and EUR baskets. The out-of-sample evaluation over 2020–2024 covers the COVID-19 shock, the post-pandemic inflation surge, and the most aggressive central bank tightening cycle in decades. It is providing a demanding test of whether the integrated framework holds up when macro conditions shift sharply.

Chapter 3

METHODOLOGY

3.1 Data transformation

The dataset covers daily closing prices for ten currency pairs, five from the USD basket and five from the EUR basket, from January 2010 through December 2024. The sample splits at the end of 2019. Everything up to that point served as the training period for the LDA classifiers. The five years from 2020 through 2024 were held out entirely and used only for out-of-sample evaluation, after all model parameters had been fixed on the training data.

All calculations start from the daily arithmetic return. For each currency pair i, the return on day t is:

$$r_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}} \tag{1}$$

Where $P_{i,t}$ is the closing price of asset i on day t.

Price levels are not stationary, returns are, and arithmetic returns make the ten pairs directly comparable regardless of their absolute price levels.

AUDUSD and GBPUSD sit the wrong way round relative to the other three USD pairs. A rising price in these two means the dollar is getting weaker, not stronger, because the dollar sits in the denominator. To keep the directional convention consistent across the whole basket, their daily returns are sign-reversed before anything else is done with them.

$$r_{i,t} = -r_{i,t} \tag{2}$$

After this adjustment, a positive$r_{i,t}$ uniformly signals dollar appreciation for all five USD pairs, consistent with the direction of the DXY benchmark.

Using DXY and EXY together creates a problem: their correlation is around −0.92, so they essentially cancel each other out. The main reason is that EUR/USD alone makes up about 31.6% of EXY's weight. To get around this, the dollar part of EXY is removed and only the euro-specific movement is kept, which is what EURO_PURE refers to in this study. The benchmark then takes equal parts of DXY and EURO_PURE:

$$r_t^{FX_BENCH} = \theta \cdot r_t^{DXY} + (1-\theta) \cdot r_t^{EUR_PURE}, \qquad \theta = 0{,}5 \tag{3}$$

where r_t^{DXY} denotes the daily return of the US Dollar Index,

$r_t^{EUR_PURE}$ denotes the daily return of the EURO_PURE index,

θ is the benchmark weighting coefficient.

To remove the dollar component from EXY, EXY log-returns are regressed on inverted DXY log-returns using a rolling window, and the residual is taken as EURO_PURE:

$$r_t^{EUR_PURE} = r_t^{EXY} - \beta^{EXY,DXY}(-r_t^{DXY}) \tag{4}$$

Where $\beta^{EXY,DXY}$ is the time-varying sensitivity of EXY to inverted DXY returns.

The beta in this regression is estimated on a rolling 252-day window, which is one full trading year:

$$\beta^{EXY,DXY} = \frac{\mathrm{Cov}(\mathrm{r}_{\mathrm{t}}^{\mathrm{EXY}}, -\mathrm{r}_{\mathrm{t}}^{\mathrm{DXY}})}{Var(-r_t^{DXY})} \tag{5}$$

Once the dollar loading is removed, the residual EURO_PURE series has close to zero correlation with DXY by construction. The benchmark in Equation (3) then combines two factors that move independently of each other, which is what makes it informative about both sides of the FX universe at the same time. Decomposing currency returns into a common dollar component and a currency-specific residual is motivated by Lustig, Roussanov and Verdelhan (2011), who document that a global dollar factor drives common variation across currency pairs.

The LDA classifiers work at monthly frequency, so daily returns need to be compounded up:

$$R_{i,m} = \prod_{t \in m} (1 + r_{i,t}) - 1 \tag{6}$$

where m denotes the calendar month and the product runs over all trading days t within that month.

3.2 PCA

Currency pairs in the same basket tend to move together — not perfectly, but enough that the raw return matrix is highly redundant. Running a regression or a classifier on ten correlated variables simultaneously would be statistically inefficient and prone to overfitting. PCA addresses this by rotating the data into a smaller set of uncorrelated directions that capture most of the variance.

Before decomposing the data, each return series is standardised to have zero mean and unit variance:

$$z_{i,t} = \frac{r_{i,t} - \bar{r}_i}{s_i} \tag{7}$$

where $\bar{r}_i$ is the sample mean return of asset i over the estimation window,

s_i is the sample standard deviation of asset i

The standardised matrix X is then decomposed into factor scores and loadings:

$$X = FW^{\top} \tag{8}$$

where F is constant the scores

W contains the factor loadings (eigenvectors).

Running PCA on the full return matrix gives more components than needed, following the factor extraction approach of Lustig, Roussanov, and Verdelhan (2011). What remains is a cleaner version of the return data with the pair-specific noise filtered out. The reconstructed matrix is:

$$\hat{R} = FW^{\top} \tag{9}$$

Expected returns fed into the optimisation are computed as the rolling mean of the reconstructed series:

$$\mu_{j,t} = \frac{1}{N} \sum_{\tau=1}^{N} \hat{R}_{\tau,j} \tag{10}$$

Where $\hat{R}_{\tau,j}$ denotes the PCA-reconstructed daily return of asset j on day τ,

N denotes the length of the estimation window.

The estimation window is set to 60 days, approximately three calendar months. A shorter window would make the estimates too sensitive to individual daily moves. Additionally, longer window reaction would be slow to shifts in market condition. Sixty days also lines up with the beta estimation window in Section 3.4, so both optimization inputs always reflect the same period of market history. PCA-reconstructed returns, in comparison to raw returns, will reduce the noise that tends to dominate short-window mean estimates in the context of smart beta allocation.

3.3 LDA

To account for changing macro-financial conditions, the study employs a regime-classification framework based on Linear Discriminant Analysis. Two separate LDA classifiers are built, one for each basket, rather than a single model

for everything. One for the USD basket and one for the EUR basket. Each conditioned on its own sector-specific factor structure. This separation is motivated by the finding of Lustig, Roussanov, and Verdelhan (2011) that risk premia in the foreign exchange market are driven by sector-specific macro-financial dynamics: the dollar and the euro respond to different sets of macroeconomic fundamentals and are therefore better modelled independently.

Let $s \in \{USD, EUR\}$ denote the sector. For each sector, the average monthly basket return is computed as the cross-sectional mean of the N_s constituent pairs:

$$\bar{R}_m^s = \frac{1}{N_s} \sum_{i=1}^{N_s} R_{i,m} \tag{11}$$

where N_s denotes the number of currency pairs within sector s.

A month is labelled as a bullish regime if this average return meets or exceeds the in-sample median, and bearish otherwise:

$$y_m^s = 1\{\bar{R}_m^s \geq median(\bar{R}^s)\} \tag{12}$$

where $y_m^s = 1$ denotes a favourable (bullish) market regime,

$y_m^s = 0$ denotes an unfavourable (bearish) market regime.

The median threshold is fixed exclusively on the training period 2010–2019 to prevent look-ahead bias. This binary labelling approach follows the regime-

classification design of Foglia, Recchioni and Polinesi (2021), who demonstrate that a threshold based on the in-sample median of portfolio returns provides a stable and economically interpretable regime indicator.

The feature vector for each monthly observation combines FX PCA factors and macroeconomic PCA factors. For the FX component, we extend the basket return matrix with a sector-level index before applying PCA. For the USD sector, DXY is included as an additional column. With it extracted components capture collective alignment within the five pairs with the broader dollar trend and their co-movement. For the EUR sector, we run two separate specifications. One with EXY and one with EURO_PURE. We want to examine whether removing the dollar component from the euro benchmark results into a more informative regime signal out of sample. Three principal components are extracted from this augmented, lagged matrix:

$$F_m^s = \left[f_{t_m}^1, f_{t_m}^2, f_{t_m}^3\right]^\top \tag{13}$$

For the macroeconomic component, two principal components are extracted from the standardised and lagged macro variables

$$G_m^s = [g_m^1, g_m^2]^\top \tag{14}$$

The combined feature vector that enters the LDA classifier is:

$$Z_m^s = [F_m^s, G_m^s] \tag{15}$$

All features are lagged by one month to prevent look-ahead bias: the classifier predicts the regime for month m using only information available through month m − 1. For the USD basket, the macro variables are: 13-week Treasury bill rate (IRX), 10-year Treasury yield (TNX), inflation expectations (TIP ETF), VIX, US real GDP growth, and US unemployment rate. For the EUR basket, the same market-traded indicators are retained alongside euro-area specific variables. This allows the model to account for both the global dollar factor and macroeconomic conditions specific to the euro area.

The LDA classifier models the posterior probability of a bullish regime given the feature vector:

$$\Pr(y_m^s = 1|Z_m^S) = \frac{\exp(a_1 + \beta_1^\top Z_m^s)}{\sum_{k=0}^{1} \exp(a_k + \beta_k^\top Z_m^s)} \tag{16}$$

where a_k and β_k are the class-specific intercept and coefficient vector for class $k \in \{0, 1\}$,

estimated by maximising the LDA likelihood on the training sample 2010–2019.

The monthly regime signal is obtained by thresholding the posterior probability at 0.5:

$$s_m^s = 1\{\Pr(y_m^s = 1|Z_m^s) \geq 0.5\} \tag{17}$$

This signal is broadcast to daily frequency by assigning each trading day within month m the value s_m^s. Since portfolio optimisation is performed daily, the monthly signals are held constant throughout the month to which they belong.

The daily signals from the two independent sector classifiers are combined via an AND rule: the portfolio enters the bullish regime only when both sectors simultaneously signal a favourable environment:

$$q_t = q_t^{USD} = 1 \ \wedge \ q_t^{EUR} = 1 \tag{18}$$

This strict logic, according to which long signals appear only when these signals are detected in two sectors at the same time is motivated by the need to reduce false positives in regime detection. One sector going bullish could be noise. Both sectors need to agree before the portfolio shifts into return-seeking mode. If only one signals favorably, the strategy stays defensive.

The composite daily signal determines the risk-aversion parameter a_t that governs the portfolio optimisation:

$$a_t = \begin{cases} a_1, q_t = 1 \\ a_0, q_t = 0 \end{cases} \tag{19}$$

In summary, the LDA framework produces a single daily binary signal for each currency basket. One for the USD sector and one for the EUR_PURE sector. Derived from a combination of sector-specific FX dynamics and macroeconomic conditioning. The macroeconomic variables are not used directly in portfolio construction, but are instead distilled into a regime label

through the LDA classifier. This two-step approach, consistent with Foglia, Recchioni and Polinesi (2021), allows the model to translate noisy, low-frequency macro signals into a clean, actionable daily indicator that reflects the broader economic environment in which the portfolio operates. The resulting signals form the foundation for the risk-aversion adjustment in the optimisation step that follows.

3.4 Smart Beta

Foglia, Recchioni, and Polinesi (2021) highlight the benefits of using dynamic weighting in smart beta strategies, particularly during periods of high market volatility. Inspired by their findings, our approach defines portfolio optimization as a balance between maximizing expected returns and minimizing downside risk, where the level of risk aversion adjusts depending on market conditions. The optimization is formulated as:

$$\max_{w}[a_t w^\top \mu_t - (1 - a_t)\, CVaR(w_t)] \tag{20}$$

where: $w \in \mathbb{R}^n$ is vector of asset weights

μ_t is the vector of PCA-based expected returns from Equation (10)

$a_t \in [0,1]$ is a risk aversion parameter that depends on the LDA signal

$CVaR$ is an estimate of the average loss in the 5% worst-case scenario.

CVaR is used instead of variance because it captures what happens in the tail of the distribution, which matters in FX markets where big one-day moves are more common than normal distributions would suggest. The estimation

window matches the one used for expected returns, so both inputs to the optimizer always reflect the same stretch of market history. This follows Rockafellar and Uryasev (2000), extending the VaR-based optimization framework of Foglia, Recchioni, and Polinesi (2021) to CVaR as a more coherent risk measure.

$$\sum_{j} w_j = 1 \tag{21}$$

Short positions are permitted for all assets:

$$w_j \in [-1,1] \tag{22}$$

The bounds $w_j \in [-1,1]$ are necessary to allow the optimiser to achieve beta neutrality after the sign reversal of AUDUSD and GBPUSD in Equation (2): without short-selling, the optimiser cannot offset the positive beta contribution of USD-appreciating pairs with negative-beta positions in the same basket.

Daily portfolio returns are the inner product of weights and individual pair returns:

$$r_{p,t} = \sum_{j=1}^{N} w_{j,t} \cdot r_{j,t} \tag{23}$$

The beta of each individual asset with respect to the FX benchmark is estimated over a rolling window of 60 trading days:

$$\beta_{i,t} = \frac{\mathrm{Cov}\left(r_{it}, r_{\mathrm{t}}^{\mathrm{FX_BENCH}}\right)}{Var\left(r_{\mathrm{t}}^{\mathrm{FX_BENCH}}\right)} \tag{24}$$

This rolling estimation allows betas to adapt to changing market conditions on a daily basis, reflecting the time-varying nature of FX risk exposures. A 60-day window is chosen to match the expected return estimation horizon. Additionally, this will ensure internal consistency between the two key inputs to the optimisation problem. The portfolio beta is the weighted sum of individual asset betas:

$$\beta_{p,t} = \sum_{i=1}^{N} w_{i,t} \cdot \beta_{i,t}^{FX} \tag{25}$$

Together, these components define a fully specified portfolio at each point in time. The weight vector w_t determines both the portfolio return and the portfolio beta. Having an explicit measure of portfolio provides the possibility of controlling its magnitude, which is the mechanism through which the strategy achieves beta-neutrality.

3.5 Rebalancing

The strategy employs event-driven rebalancing rather than fixed calendar rebalancing. Rebalancing is triggered by three independent conditions.

The first condition is a beta breach: the optimisation problem is re-solved whenever the absolute portfolio beta exceeds the threshold δ:

$$\left|\beta_{p,t}\right| > \delta, \qquad \delta \in [0.01, 0.05] \tag{26}$$

The second and third conditions are based on adaptive gain and loss thresholds scaled by rolling portfolio volatility. Portfolio volatility is estimated as the standard deviation of daily portfolio returns over the preceding 21 trading days:

$$\hat{\sigma} = Std(r_t) \tag{27}$$

A positive rebalancing trigger fires when cumulative gains since the last profit/loss event exceed three rolling standard deviations:

$$\frac{\Delta V_t}{V_t} > 3\hat{\sigma} \tag{28}$$

A negative trigger fires when cumulative losses exceed one rolling standard deviation:

$$\frac{\Delta V_t}{V_t} < -\hat{\sigma} \tag{29}$$

where ΔV_t is the change in portfolio value since the last profit/loss rebalancing event.

The gain threshold is three times the loss threshold. Winning positions are given more room to run before locking in, while losses get cut faster. The beta breach trigger runs independently and does not reset the profit/loss reference point.

3.6 Trading model

The beta-neutral portfolio from sections 3.1 to 3.5 keeps market exposure near zero. Returns come from the LDA signal and the CVaR optimizer. The last piece added on top is a leverage overlay. It scales daily exposure up when conditions look asymmetrically positive and cuts back to cash when tail risk starts building.

The overlay draws on Harvey and Siddique (2000), who show that conditional skewness is a priced risk factor. When a portfolio's return distribution is right-skewed, the upside asymmetry lowers the effective cost of holding the position and more exposure is warranted. When skewness turns negative, the left tail dominates and the position should be closed rather than held at a premium. To operationalise this, portfolio skewness is estimated as the standardised third central moment of daily returns $r_{p,\tau}$ over a rolling window of $N = 60$ trading days. The estimate uses only returns through day $t-1$, so the leverage decision on day t is formed on the information set Ω_{t-1} with no look-ahead bias:

Skewness is estimated over the previous $N = 60$ daily returns:

$$SK_t = \frac{\frac{1}{N}\Sigma_{p=t-N}^{t-1}\left(r_{p,\tau} - \bar{r}\right)^3}{\left(\frac{1}{N}\Sigma_{p=t-N}^{t-1}\left(r_{p,\tau} - \bar{r}\right)^2\right)^{\frac{3}{2}}} \tag{30}$$

Two fixed thresholds $\kappa^+ > 0 > \kappa^-$ partition each trading day into one of three regimes $\Phi \in \{-1, 0, +1\}$ per Equation (31). Within each regime, the size of the leverage response is calibrated using the second moment of the same distribution. Realised variance over the past $M = 21$ trading days is compared against its own expanding historical mean, which serves as the long-run benchmark:

$$\phi_t = \begin{cases} +1 \;\; if \;\; SK_t > k^+ \\ -1 \;\; if \;\; SK_t > k^- \\ 0 \;\; otherwise \end{cases} \tag{31}$$

When current variance sits below its long-run average, the ratio exceeds one and the strategy takes more exposure. When variance spikes, the base leverage drops automatically. This happens before skewness has time to react, so it works as a second layer of protection. Both moments feed into a single daily multiplier

$$RV_t = \Sigma_{r=t-M}^{t-1} r_{r=t-M}^2 \tag{32}$$

$$\tilde{\sigma}_t^2 = \frac{1}{t-1}\Sigma_s^{t-1} RV_s \tag{33}$$

$$l_t^{base} = \min\left(\frac{\tilde{\sigma}_t^2}{RV_{t-1}}, L_{max}\right) \tag{34}$$

The base leverage and the regime signal combine into the final daily multiplier:

$$l_t = \begin{cases} 0 & if \quad \phi_t = -1 \\ l_t^{base} & if \quad \phi_t = 0 \\ \min\left(l_t^{base}\left(1 + \frac{SK_t}{SK_{t^*}}\right), L_{max}\right) & if \quad \phi_t = +1 \end{cases} \tag{35}$$

Where $SK_{t^*} = \max_{s<t}|SK_s|$ normalises the boost to $(0,1]$.

When skewness turns negative, the position closes to zero entirely. The neutral case is simpler: the portfolio just runs at whatever the variance-scaled base happens to be. If skewness is strongly positive, that base gets scaled up in proportion to how large the skewness signal is, with the total capped at L_{max}.

Leveraged daily return and portfolio value follow directly:

$$r_{p,t}^{l} = l_t \cdot r_{p,t} \tag{36}$$

The overlay is applied after the weight optimisation: w_t from Section 3.5 determines $r_{p,t}$ through Equation (23), and ℓ_t scales the entire position

uniformly. The thresholds $\kappa^{+} = 0.3$ and $\kappa^{-} = -0.3$ are set at approximately one standard deviation of the rolling skewness distribution. $N = 60$ days follows the three-month horizon used in Sections 3.2–3.4; $M = 21$ days follows the one-month horizon of Section 3.5; $L_max = 3$ reflects realistic margin constraints in institutional FX markets.

Chapter 4

DATA

The data chapter covers three areas: the sources of the data, the pre-modelling transformations, and their statistical characteristics. The full sample spans from January 2010 through December 2024. Everything up to 2019 was used to train the LDA classifiers, while the 2020–2024 period was held out exclusively for the evaluation stage.

To capture market-traded indicators, macroeconomic variables were obtained from Yahoo Finance, while lower-frequency releases were extracted from the Federal Reserve Economic Data (FRED) database. These economic metrics complement the financial asset data, which include daily closing values for currency pair prices sourced from Interactive Brokers. Additionally, both the DXY and EXY indices were pulled from TradingView at the same daily frequency.

The portfolio covers ten currency pairs split evenly into two baskets. The USD basket contains AUDUSD, GBPUSD, USDCAD, USDCHF, and USDJPY. The EUR basket contains EURAUD, EURCAD, EURCHF, EURGBP, and EURJPY, with the euro as the base currency throughout.



Figure 1. Daily FX Prices (2010–2024).

Figure 1 traces price levels across the full sample. USDJPY tells the clearest directional story, climbing from around 75 to above 150 as the Fed and the Bank of Japan went in opposite directions on monetary policy. EURCHF is the most dramatic single episode: the Swiss National Bank dropped its exchange rate floor in January 2015 and the franc appreciated by an amount that would normally take years to accumulate, all within one session. GBPUSD dropped sharply after the Brexit result in June 2016. The other pairs moved in tighter ranges through most of the sample, though all of them showed elevated volatility around the COVID-19 shock in early 2020 and again during the 2022 rate hiking cycle.

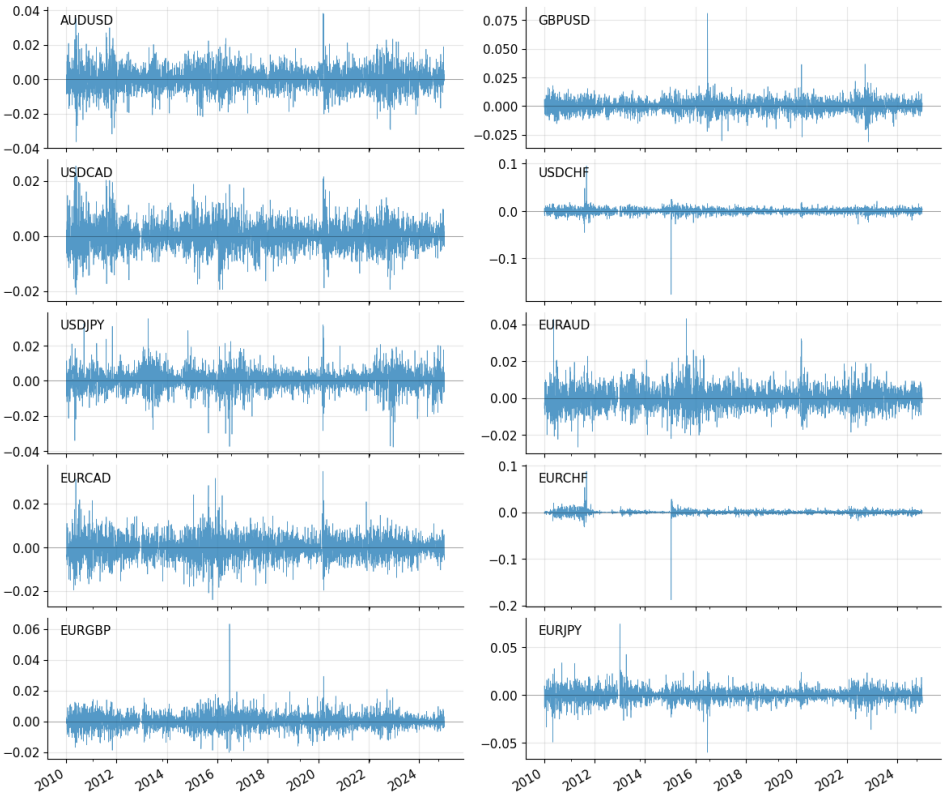


Figure 2. Daily FX Return (2010–2024).

Figure 2 shows the same series in daily return form. Returns sit close to zero on average but have fat tails, and volatility clusters around known stress episodes. The January 2015 EURCHF observation stands out as the single most extreme datapoint in the sample, exceeding fifteen standard deviations on that day.

Table 1. Descriptive Statistics of FX Pairs (2010–2024).

Statistic	N	Mean	St. Dev	Min	Max
AUDUSD	3890	0.000078	0.006675	-0.036335	0.038312
GBPUSD	3890	0.000049	0.005570	-0.031294	0.080886
USDCAD	3890	0.000094	0.004831	-0.021197	0.025340
USDCHF	3890	-0.000012	0.006297	-0.176174	0.095112
USDJPY	3890	0.000153	0.005758	-0.037822	0.035415
EURAUD	3890	0.000032	0.005769	-0.026486	0.042966
EURCAD	3890	0.000011	0.005178	-0.024036	0.034768
EURCHF	3890	-0.000104	0.005140	-0.187528	0.087434
EURGBP	3890	-0.000009	0.004777	-0.020193	0.063168
EURJPY	3890	0.000073	0.006517	-0.060144	0.074479

Table 1 puts numbers on what Figure 1 showed visually. Volatility spreads fairly wide across the basket, from 0.0048 for EURGBP to 0.00668 for AUDUSD. The SNB and Brexit events from the previous section show up clearly in the extremes: EURCHF has the largest negative observation at -0.1875, GBPUSD the largest positive at 0.0809. Means are close to zero throughout, with no pair showing any consistent directional drift over the fifteen years.



Figure 3. DXY and EUR_PURE Index Dynamics (2010–2024).

Figure 3 shows prices and daily returns for DXY and EXY. DXY rose from around 78 in 2011 to above 114 at its 2022 peak, then pulled back to around 100 by end of 2024. EXY moved in the opposite direction for most of the period, falling from above 145 in 2010 to a low near 96 in 2022 before recovering partially.

Table 2. Descriptive Statistics of Currency Indices (2010–2024).

Statistic	N	Mean	St. Dev	Min	Max
DXY	3890	0.000096	0.004407	-0.023722	0.020518
EXY	3890	-0.000071	0.005299	-0.027926	0.030526

Table 2 reports summary statistics for the two currency indices over the full sample period. DXY shows a slight positive mean daily return of 0.0001 with

lower volatility than EXY, while EXY carries a marginally negative mean and wider dispersion, consistent with the inverse relationship between the two indices documented throughout the sample.

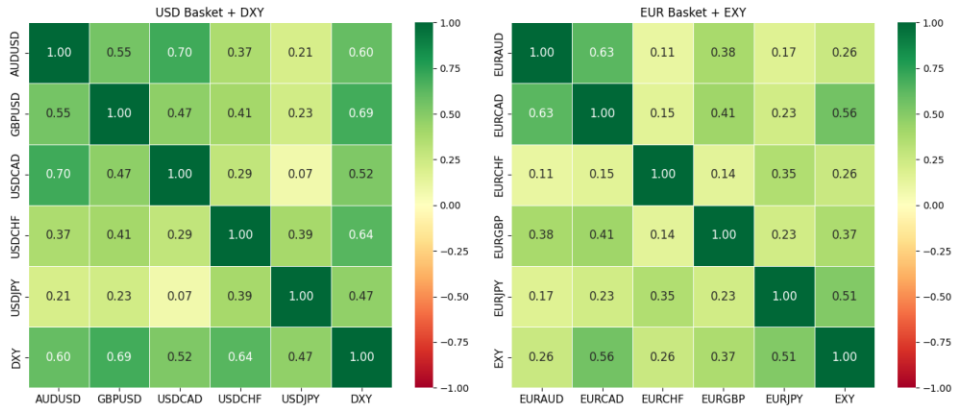


Figure 4. Correlation Structure of USD and EUR Baskets.

Within the USD basket, pairwise correlations are positive throughout, ranging from 0.07 for USDCAD/USDJPY to 0.70 for AUDUSD/USDCAD. DXY loads most strongly on GBPUSD (0.69) and USDCHF (0.64), and least on USDJPY (0.47). Within the EUR basket, EURCHF shows notably weaker co-movement with the other pairs, while EXY correlates most strongly with EURCAD (0.56) and EURJPY (0.51).

For the USD basket, six variables are used: the 13-week T-bill rate (Rates_S), the 10-year Treasury yield (Rates_L), inflation expectations via the TIP ETF (Infl_Exp), VIX, US GDP growth (GDP_YoY), and the unemployment rate (Unemp). The selection follows Foglia, Recchioni, and Polinesi (2021), who condition smart beta allocation on interest rates, volatility, and inflation to capture shifting market regimes. Lettau and Ludvigson (2001) further

motivate the term spread as a predictor of time-varying risk premia. Monthly FRED series are forward-filled to daily frequency.

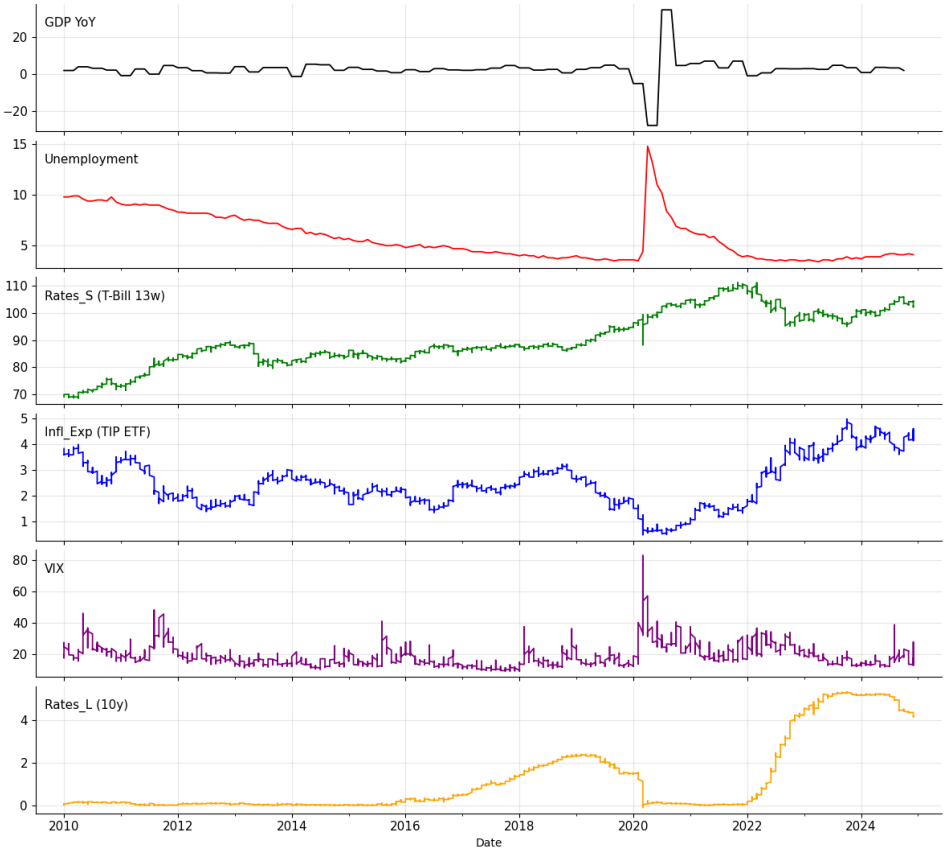


Figure 5. U.S. Macroeconomic Indicators (2010–2024).

Figure 5 shows all six variables across the sample. Short rates stayed near zero through 2015, then rose gradually before spiking sharply during the 2022–2023 tightening cycle. GDP growth dropped hard in 2020 and bounced back quickly, while unemployment hit 14.70% at the peak of the COVID lockdowns.

Table 3. Descriptive Statistics of U.S. Macroeconomic Variables (2010–2024).

Statistic	N	Mean	St. Dev	Min	Max
Rates_S	3773	90.0023	9.9365	68.5333	111.2127
Rates_L	3773	1.1915	1.7194	-0.1050	5.3480
Infl_Exp	3773	2.4920	0.9337	0.4990	4.9880
VIX	3773	18.3768	6.9638	9.1400	82.6900
GDP YoY	125	2.6968	6.6666	-28.0000	34.9000
Unemployment	125	5.7848	2.2404	3.5000	14.8000

Table 3 reports the summary statistics. VIX ranges from 9.14 to 82.69, with the peak in March 2020.

For the EUR basket, the same four market-traded indicators are used alongside five euro-area specific variables from FRED: the unemployment rate, the German 10-year Bund yield, GDP growth, the ECB deposit rate, and HICP inflation.

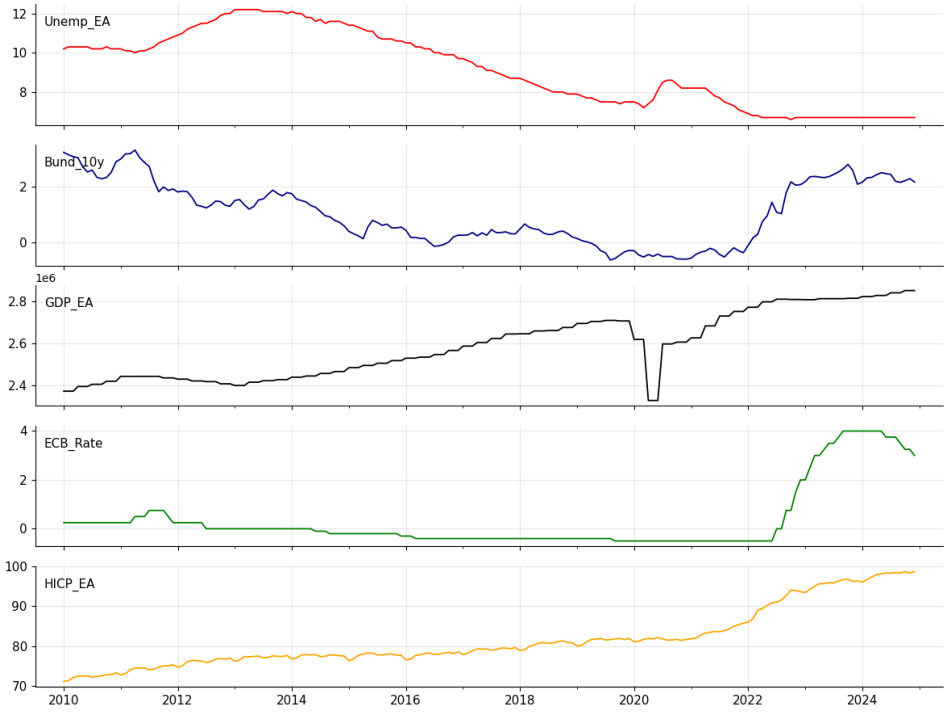


Figure 6. Euro-Area Macroeconomic Indicators (2010–2024).

The ECB deposit rate entered negative territory in 2014 and remained there until mid-2022, a period with no parallel in the US data. Bund yields followed a similar path, turning negative from 2016 through 2021 before reversing sharply as inflation pressures built. HICP inflation was broadly stable through the 2010s before reaching 10.62% in late 2022.

Table 4. Descriptive Statistics of Euro Macroeconomic Variables (2010–2024).

Statistic	N	Mean	St. Dev	Min	Max
Unemp_EA	180	9.1856	1.8949	6.6000	12.2000
Bund_10y	180	1.0428	1.1320	-0.6491	3.3405
GDP_EA	180	2588305	153734.1014	2327598	2851203
ECB_Rate	180	0.3458	1.3214	-0.5000	4.0000
HICP_EA	180	81.6037	7.3507	71.1500	98.7100

Table 4 reports the summary statistics.

Chapter 5

ESTIMATION RESULTS

In building upon the methodological framework and empirical data discussed in the previous chapters, a full-scale model was constructed. The results of its implementation are presented and examined in detail in the following section.

5.1 Comparison of LDA with and without macroeconomic variables

Before proceeding to the results of building a beta-neutral portfolio, it is worth examining the factors that directly influenced the formation of portfolio allocation indicators, namely LDA signals. The LDA classifier produces a binary signal for USD and EUR baskets, where 1 means that the five currency pairs within the USD or EUR basket evenly divided in the next month should show profit, so with a signal of 1, the optimizer focuses on return generation. And when it showed 0, the risk aversion parameter shifts toward minimizing risks with the help of CVaR rather than making a profit.

The study used two separate LDA models. The first model incorporated macroeconomic factors. The second model was based solely on the return data of currency pairs plus DXY for USD currency pairs and EXY, EUR_PURE for EUR basket, without including macroeconomic variables. Comparison of the results of both models allows one to assess the extent to which the addition of macro factors affects forecast accuracy and the efficiency of constructing a beta-neutral portfolio.

The results of the first model, which did not use macroeconomic data but only a basket of currencies and their index, were as follows:

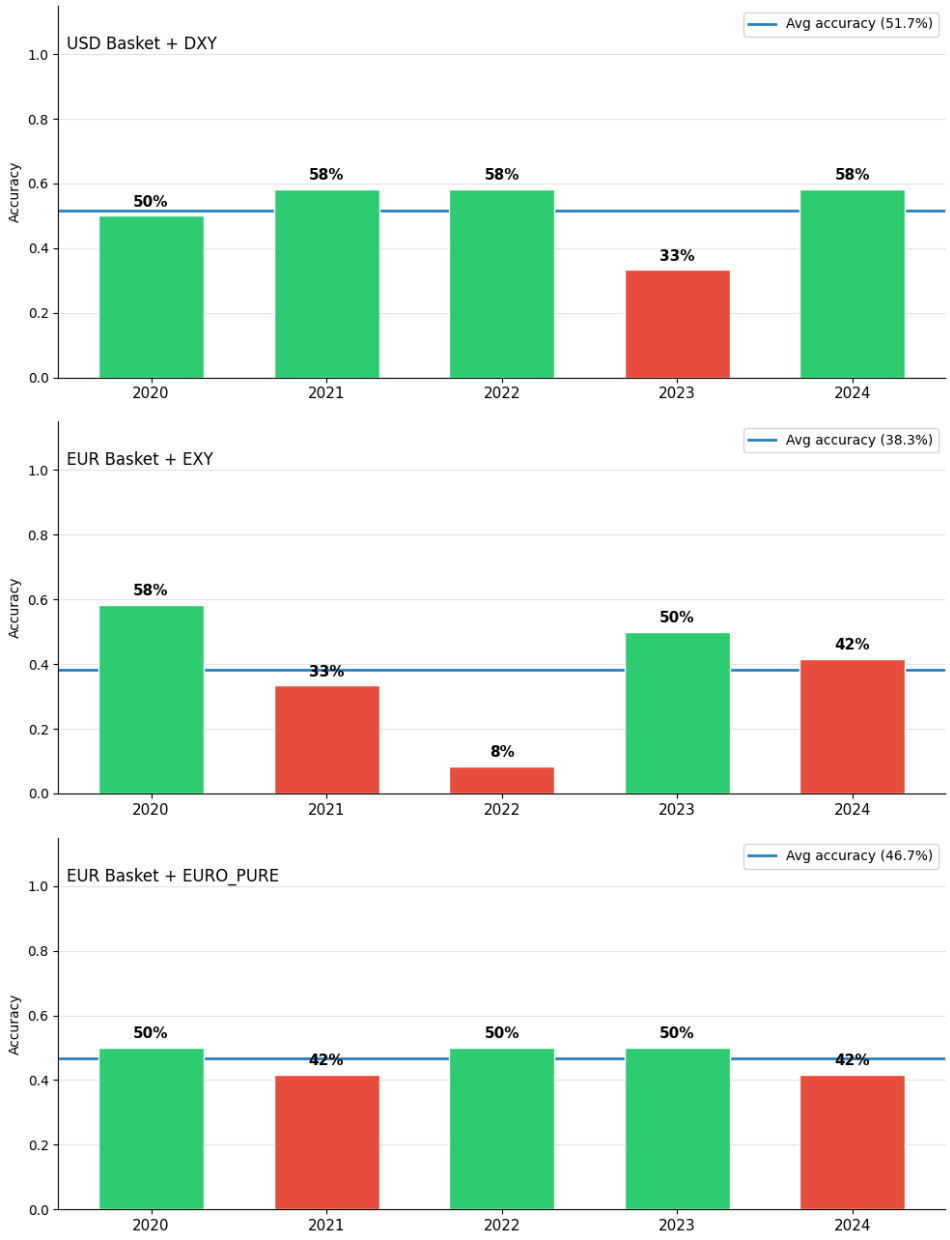


Figure 7. LDA Forecast Accuracy by Year.

The USD basket performs best, with average of 51.7%. It correctly classifies four out of five years. Accuracy level is at 58% in 2021, 2022, and 2024. The only weak year is 2023 at 33%.

The EUR + EXY model average level only 38.3%. There was a collapse to 8% in 2022. The EUR + EURO_PURE model is more stable at 46.7% average level. There are no extreme outliers. It suggests that removing the dollar component from the benchmark produces a cleaner signal. Switching from EXY to EURO_PURE added 8.4 % points to forecast accuracy. With the dollar influence removed, PCA picks up euro-specific patterns instead of global dollar noise. Also, the LDA reads the market regime more accurately as a result.

The forecast using EXY yielded poorer results and less stability in explaining the movements of the EUR basket, so EUR_PURE is used going forward.

The next step is to add all our macro variables to the model and see how much they improve the forecast.



Figure 8. LDA with Macro Forecast Accuracy by Year.

Figure 8 shows the results after adding the best-performing macro combinations to each model.

The USD basket improves from 51.7% to 60% on average, with all five years now above the 50% threshold. The most notable gains are in 2020 and 2021, where accuracy reaches 67%.

The EUR basket with EURO_PURE improves from 46.7% to 68.3% on average, with particularly strong results in 2022 and 2023 at 75% and 83% respectively. Only 2024 falls below threshold at 50%.

Overall, adding macro variables produces a meaningful improvement in both baskets, with the EUR_PURE model benefiting more — jumping 21 percentage points compared to 8 for USD.

For the USD basket, six macroeconomic variables were added: the 13-week T-bill rate (Rates_S), the 10-year Treasury yield (Rates_L), inflation expectations proxied by the TIP ETF (Infl_Exp), market volatility index (VIX), US GDP growth (GDP_YoY), and unemployment (Unemp). For the EUR basket, the same four market-traded indicators were kept and five euro-area specific variables were added on top: unemployment (Unemp_EA), the German 10-year Bund yield (Bund_10y), euro-area GDP growth (GDP_EA), the ECB deposit rate (ECB_Rate), and HICP inflation (HICP_EA).

To identify which macro variables contribute most to forecast accuracy, all subsets of one to three variables were evaluated exhaustively for each basket. The top-10 combinations ranked by out-of-sample accuracy are presented below.

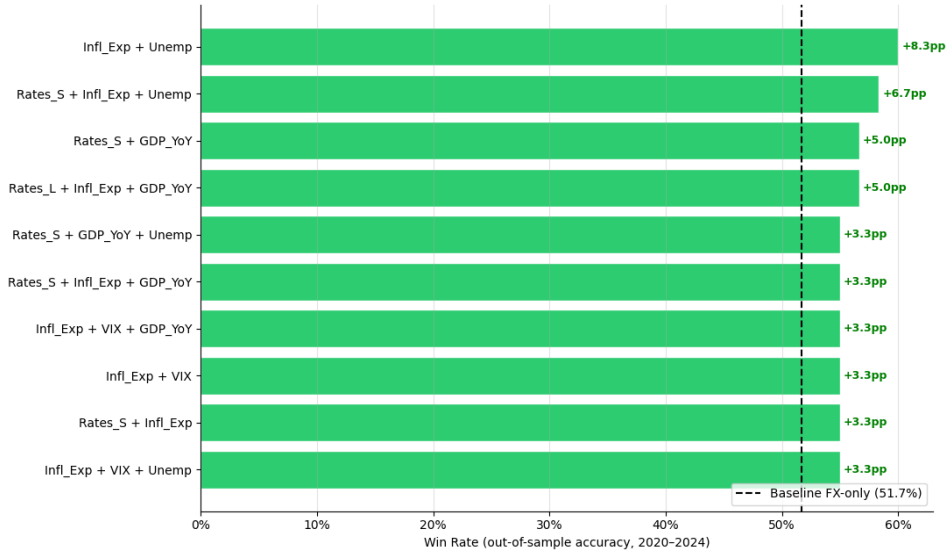


Figure 9. Top-10 USD Macro Variable Combinations.

For the USD basket, the best-performing combination is Infl_Exp + Unemp, which adds 8.3 percentage points over the FX-only baseline of 51.7%. The second-ranked combination, Rates_S + Infl_Exp + Unemp, adds 6.7 percentage points. Inflation expectations appear in eight out of ten top combinations, suggesting it carries the most consistent regime information for the USD basket.

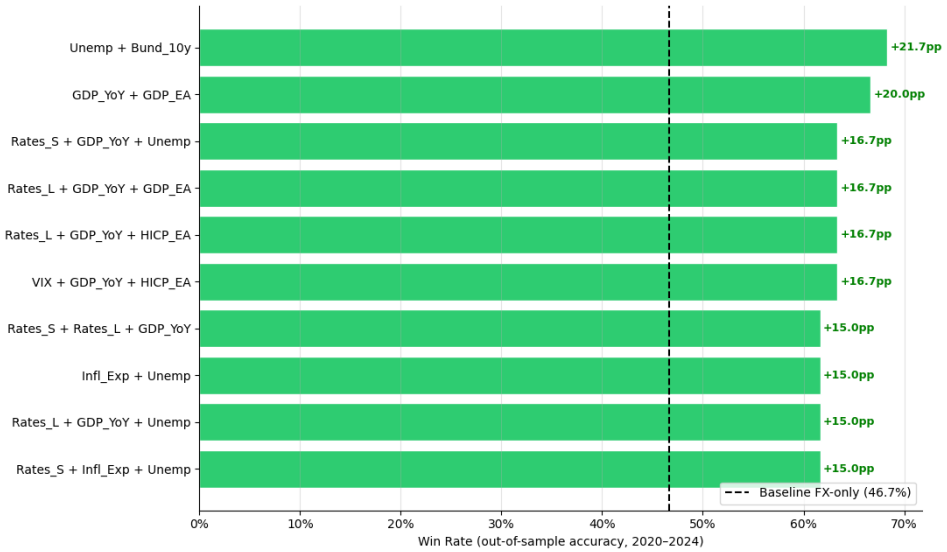


Figure 10. Top-10 EUR Macro Variable Combinations.

For the EUR_PURE basket, the top combination is Unemp + Bund_10y at +21.7pp over the 46.7% baseline. Bund_10y is the only euro-area specific variable in this combination, while Unemp reflects US labor market conditions. GDP_YoY appears in seven of the ten best combinations, confirming that US growth cycle is the dominant macro signal for the EUR basket. Euro-area variables like Bund_10y, GDP_EA, and HICP_EA appear mostly in combination with US indicators rather than standalone, suggesting the EUR basket regimes are driven by an interaction between global and local fundamentals.

5.2 Constructing a Beta-Neutral Portfolio

The first thing we need to demonstrate is the new benchmark we have obtained using equation 3 for the entire set of currency pairs, including the dollar and euro - the FX benchmark.

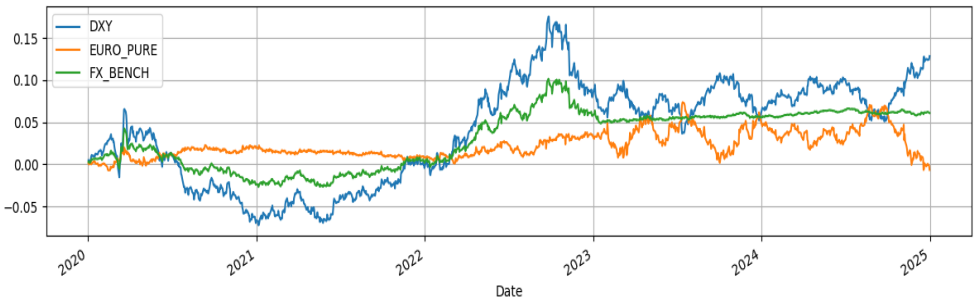


Figure 11. FX Benchmark vs DXY and EURO_Pure.

Figure 11 compares the daily dynamics of the composite FX benchmark against DXY and EURO_PURE over the full sample. The benchmark confirms that the equal-weighted combination captures both the dollar and euro sides of the market without amplifying one of them.

The model was evaluated across all 180 parameter configurations over the out-of-sample period 2020–2024. Three inputs vary across configurations: the bullish risk-aversion parameter α_1, the defensive parameter α_0, and the beta-breach threshold δ.

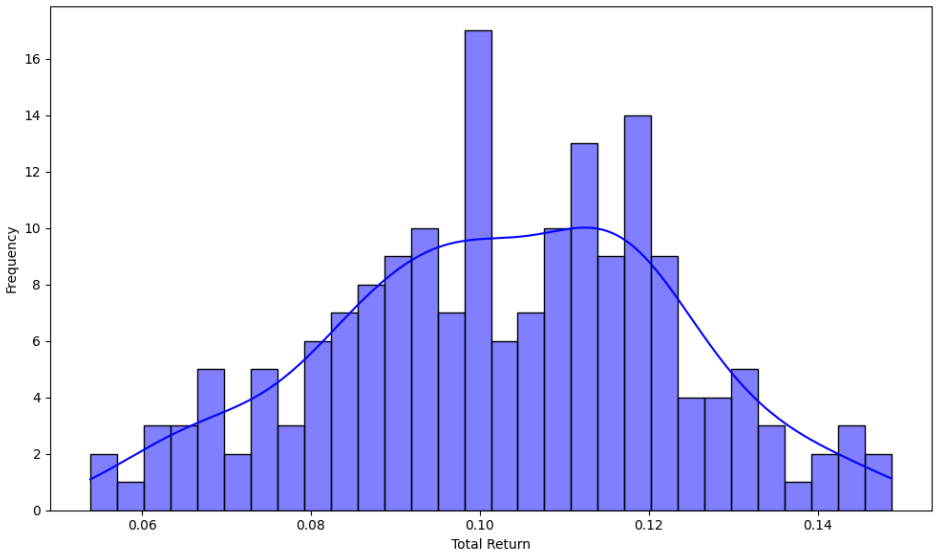


Figure 12. Distribution of Total Returns.

Figure 12 shows the distribution of total returns across all 180 simulations. Returns cluster between 6% and 14%, with a mean of 10.2%. The median is 10.0%. Every configuration finished in positive territory. It reflects the stability of the weight optimisation rather than any single parameter choice. The optimisation procedure proved to be stable. All simulated strategies ended the period with a positive final value. This indicates the effectiveness of the algorithm in forming portfolio weights. It allows for stable generation of positive alpha.

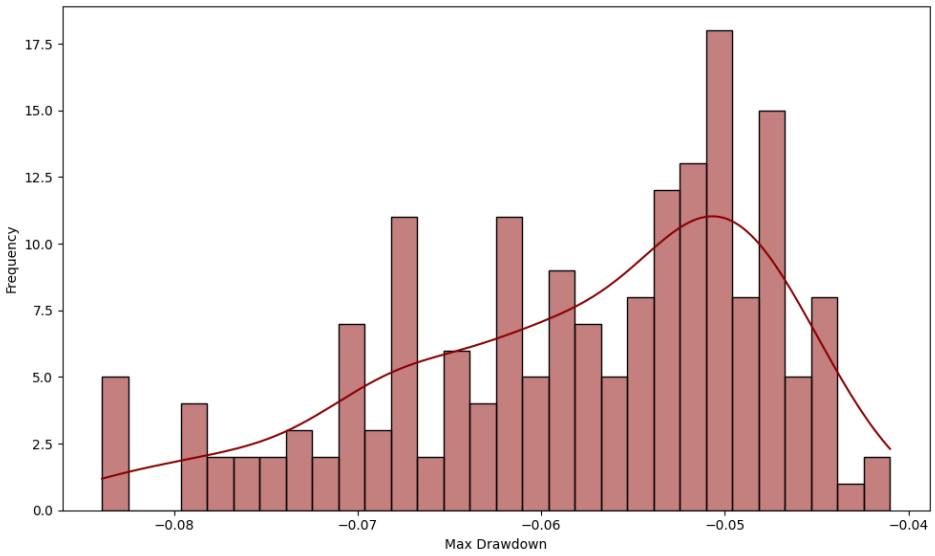


Figure 13. Distribution of Maximum Drawdowns.

Maximum drawdowns across all configurations vary from - 4.5% to - 8.5%. Most of observations are grouped between - 4.5% and - 5.5%. The event-driven rebalancing mechanism keeps losses contained regardless of parameter choice.

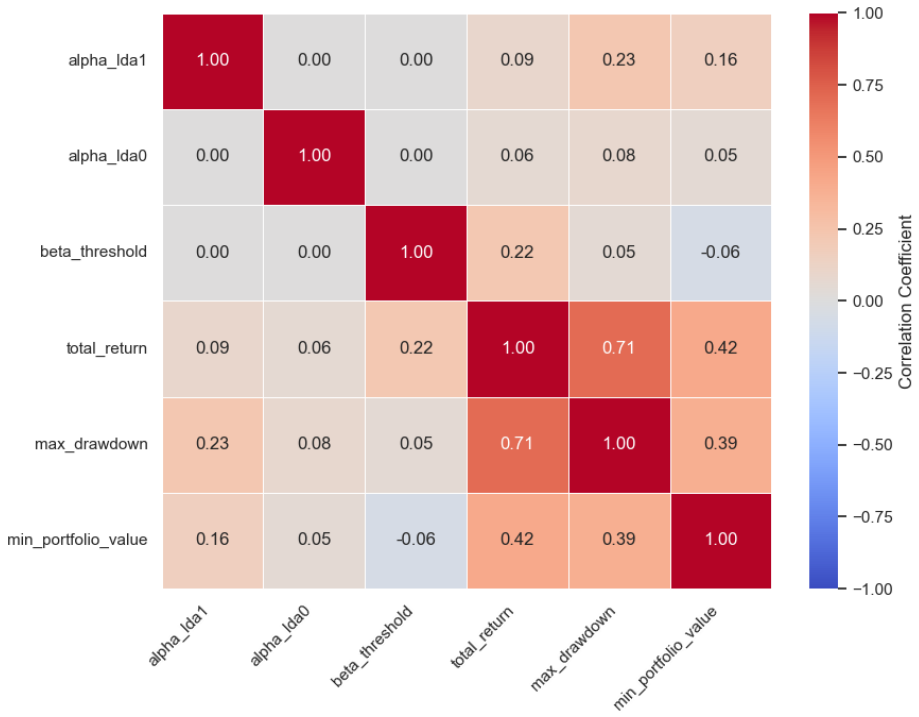


Figure 14. Correlation matrix by portfolio factors.

The graph shows the correlation matrix of all components used in portfolio construction. The beta constraint has a correlation coefficient of 0.22 with returns. This indicates that adjusting this parameter improves the strategy's overall performance. Without significantly increasing the risk of large losses. We can also see that returns correlate strongly with maximum drawdown with a coefficient of 0.71. This shows the classic financial trade-off. The higher the returns a strategy generates in a given simulation, the more the portfolio temporarily declines.

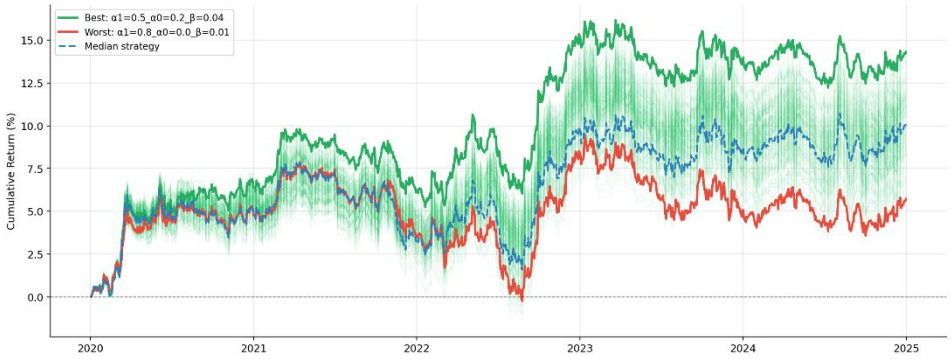


Figure 15. Cumulative Returns: All Strategies.

Figure 15 shows the cumulative return paths of all 180 configurations from 2020 to 2024. The paths diverge after mid-2021, but every configuration recovers and finishes in positive territory.

5.3 Simulated trading results

The return results show that the beta return of the neutral portfolio did not exceed 15 percent over a five-year period under any parameters, which precludes its use as a standalone strategy. Considering that in this study we do not directly account for transaction costs, and in addition to commissions we have a large number of rebalancings, investing will result in losses or negligible returns. We can use the strengths of our model as an indicator for real trading. Our model performs well in the function of optimizing the selection of optimal weights, which in most cases leads to small but stable growth throughout the entire test period. By using leverage and limiting trading under unfavorable conditions, we can ensure rapid growth in profitability and minimize risks. Using our approach as a module for selecting optimal weights within a broader trading model.

The implementation of Harvey and Siddique's (2000) approach in our leverage module acts as an automatic safeguard against large losses.

When the market is calm, the module uses standard leverage to increase profits. However, if the model detects a dangerous change in the returns distribution, leverage is instantly reduced to zero, and the strategy is suspended. Such change could be negative asymmetry. It indicates the risk of a sharp collapse.

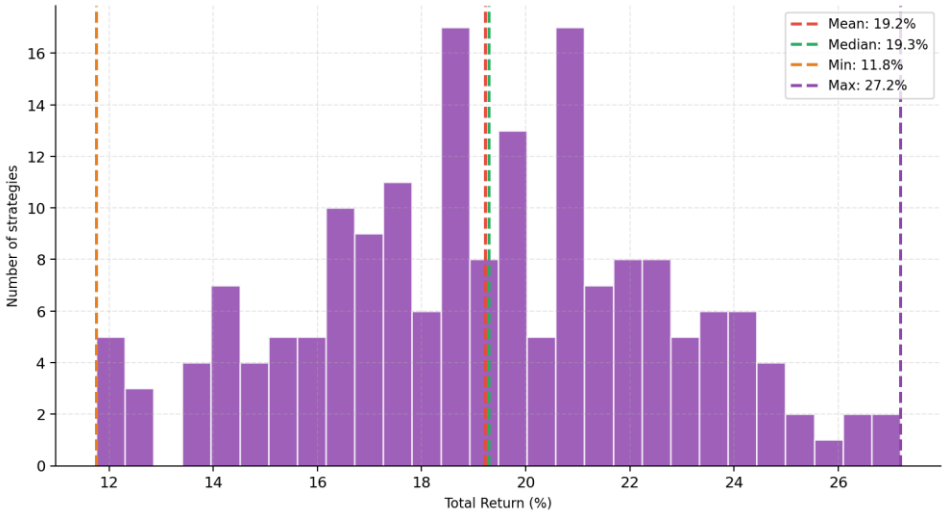


Figure 16. Final Return Distribution.

Figure 16 shows the return distribution after applying the leverage overlay. The median total return increases from 10.0% to 19.3%. The best configuration achieves 27.2%. The worst configuration is at 11.8%. The 5.5% worst unlevered outcome. The mean return increases by 9 percentage points. Even the worst-case scenario ended in the black with a return of 11.8%.

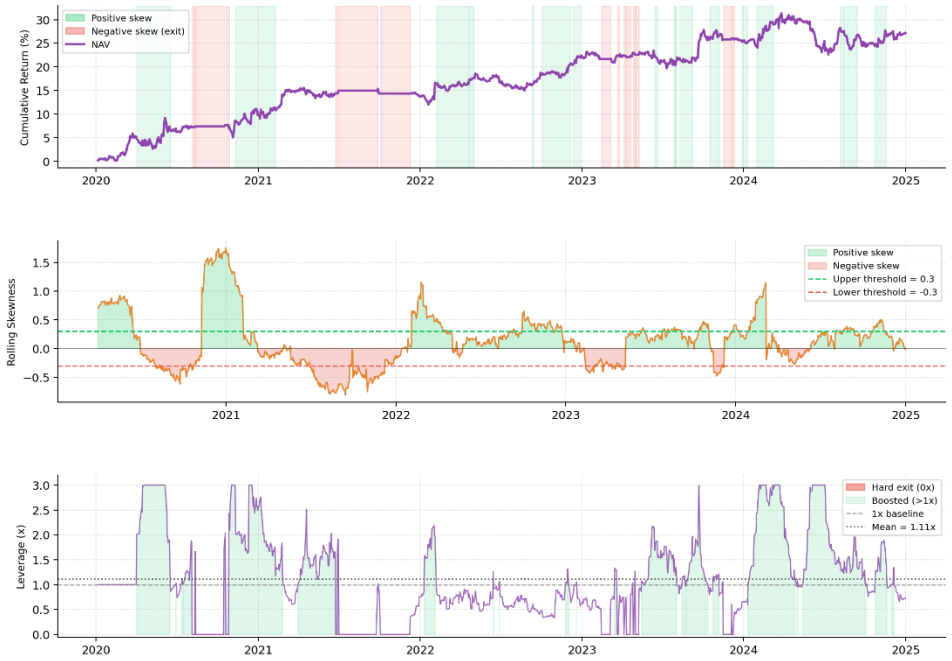


Figure 17. Skewness Details: Using a Portfolio with the Best Returns.

Figure 17 shows the mechanism in detail for the best-performing configuration. Rolling 60-day skewness (middle panel) vary between -0.5 and +1.6 across the sample. Leverage (lower panel) rises 3-times during positive-skew stretches and drops to zero when skewness crosses the lower threshold. The flattenings in the NAV path (upper panel) during early 2021 and mid-2022 occurs simultaneously with hard-exit periods. That confirms the protection mechanism is working as intended.

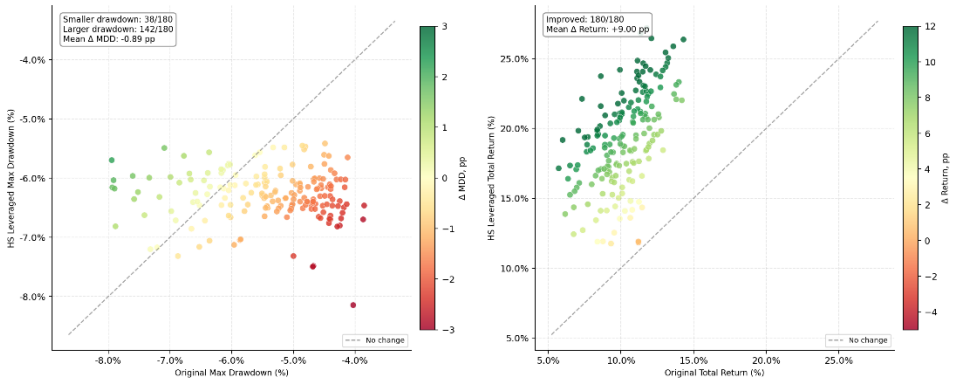


Figure 18. Compare Drawdown and Return.

The chart compares the results of the base model and the modified strategy using the Harvey-Siddique leverage. The right-hand panel clearly demonstrates a dramatic improvement in performance: all 180 simulations showed higher returns, with the average return increasing by 9 percentage points. At the same time, the left panel shows that this increase in profit did not lead to a catastrophic rise in risk: the maximum drawdown changed by an average of just 0.89 percentage points, consistently remaining within safe limits of -5% to -7%.

Chapter 6

CONCLUSIONS AND POLICY RECOMMENDATIONS

6.1 General conclusions

The main contribution of this study is the integration of several existing quantitative tools into a single, FX-specific pipeline. While LDA regime signals, macroeconomic conditioning, CVaR-based weight optimization, and event-driven rebalancing have each appeared separately in the literature, combining them within a unified framework tailored to currency markets represents the core novelty of this research. Empirically, the macro-augmented classifier outperformed the FX-only baseline in out-of-sample accuracy over the 2020–2024 period, allowing the portfolio to achieve positive returns across most parameter combinations.

6.2 Contribution and Limitations

Three primary limitations should be noted. First, transitioning to higher-frequency data would require a complete rebuild of the technical framework rather than a simple modification of the existing model. This is because the current system operates exclusively on daily closing prices, which leaves all intraday dynamics completely hidden.

Second, PCA is applied uniformly across all macroeconomic variable combinations in the grid search, including single-variable cases where the transformation is redundant. While this approach kept accuracy scores strictly comparable across all subsets, it represents a simplification relative to the LDA methodology of Foglia, Recchioni, and Polinesi (2021).

Third, because the rebalancing frequency fluctuates significantly across different configurations, incorporating a single realistic cost assumption was impractical. Consequently, transaction costs were omitted entirely from the performance metrics, though accounting for them would be an absolute necessity in any live implementation.

Finally, the framework itself is highly modular, meaning each component operates within its own layer and can be independently modified or replaced without disrupting the rest of the system.

WORKS CITED

Avramov, Doron, Laurent Barras, and Robert Kosowski. "Hedge Fund Return Predictability Under the Magnifying Glass." *Journal of Financial and Quantitative Analysis* 48, no. 4 (2013): 1057–1083. https://www.jstor.org/stable/43303833

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. "COVID-Induced Economic Uncertainty." *NBER Working Paper* No. 26983, 2020. https://www.nber.org/papers/w26983

Black, Fischer. "Beta and Return." *Journal of Portfolio Management* 20, no. 1 (1993): 8–18. https://www.pm-research.com/content/iijpormgmt/20/1/8

Clarke, Roger, Harindra de Silva, and Steven Thorley. "Risk Management and the Optimal Combination of Equity Market Factors." *Financial Analysts Journal* 76, no. 3 (2020): 57–79. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3486035

Foglia, Matteo, Maria Cristina Recchioni, and Gloria Polinesi. "Smart Beta Allocation and Macroeconomic Variables: The Impact of COVID-19." *Risks* 9, no. 2 (2021): 34. https://www.mdpi.com/2227-9091/9/2/34

Lettau, Martin, and Sydney Ludvigson. "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying." *Journal of Political Economy 109*, no. 6 (2001): 1238–1287. https://www.jstor.org/stable/10.1086/323282

Markowitz, Harry. "Portfolio Selection." *Journal of Finance* 7, no. 1 (1952): 77–91. https://www.jstor.org/stable/2975974

Papageorgiou, Nicolas, Jonathan J. Reeves, and Xuan Xie. "Betas and the Myth of Market Neutrality." *SSRN* Working Paper, 2012. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1974370

Harvey, Campbell R., and Akhtar Siddique. "Conditional Skewness in Asset Pricing Tests." *Journal of Finance* 55, no. 3 (2000): 1263–1295. https://doi.org/10.1111/0022-1082.00247

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan. "Common Risk Factors in Currency Markets." *Review of Financial Studies* 24, no. 11 (2011): 3731–3777. https://doi.org/10.1093/rfs/hhr068

Rockafellar, R. T., and Stanislav Uryasev. "Optimization of Conditional Value-at-Risk." *Journal of Risk* 2, no. 3 (2000): 21–41. https://doi.org/10.21314/JOR.2000.038

Reinhart, Carmen M., and Kenneth S. Rogoff. "This Time Is Different: Eight Centuries of Financial Folly." *Princeton University Press*, 2009. https://www.researchgate.net/publication/227486731_This_Time_Is_Different_Eight_Centuries_of_Financial_Folly