A Work Project, presented as part of the requirements for the Award of a Master's Degree in Finance from the NOVA School of Business and Economics.

First passage times in portfolio optimization: An analysis of currency carry trade

Pedro Nuno Valadeiro Varela

47210

A PROJECT CARRIED ON THE MASTER'S IN FINANCE PROGRAM UNDER THE SUPERVISION OF:

Professor Paulo M. M. Rodrigues

JANUARY, 2025

First passage times in portfolio optimization: An analysis of currency carry trade [*]

Pedro Nuno Valadeiro Varela

January 2025

Abstract

Currency carry trade strategies have historically delivered notable Sharpe ratios by exploiting the failure of Uncovered Interest Parity (UIP), yet they remain vulnerable to sharp unwinds, crash risk, and high negative skewness. In this paper, we adapt and evaluate the novel first passage times model developed by Zsurkis, Nicolau, and Rodrigues (2024). We argue that this non-parametric approach, which explicitly accounts for intra-horizon risk, offers a potential advancement in both optimization and risk management for currency carry trades. Our empirical findings indicate that this model outperforms traditional long–short strategies and performs at least as well as mean–variance optimization, while reducing negative skewness and kurtosis.

JEL Classification: G11, G15, G17, C14
Keywords: Currency Carry Trade, Portfolio Optimization, Exchange rate dynamics, Intra-horizon risk

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-014-015-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

---

[*] I would like to express my sincere gratitude to Professor Paulo M. M. Rodrigues and Professor Gabriel Zsurkis for their guidance, shared knowledge and availability throughout the duration of my Work Project. Additionally, I extend my heartfelt thanks to my family for their unwavering support.

1 Introduction

Carry trade is a trading strategy widely used in the financial sector, primarily within currency markets, where traders borrow in low-interest-rate currencies, short positions, and invest in higher-interest-rate currencies, long positions (Menkhoff et al., 2012). The idea behind the currency carry trade is to profit from the difference in interest rates between two currencies—also known as the "carry"—while managing the risks associated with their movements. Typically, spot and forward foreign exchange contracts are used to simplify the strategy.

The carry trade strategy became popular in the 1990s and 2000s, especially due to low interest rates in Japan, which encouraged investors to borrow in yen and invest in higher-yielding currencies. This trend was amplified in the early 2000s as global liquidity increased and leverage became more available. This popularity saw another peak before the 2008 financial crisis when risky assets and high-yielding currencies were especially attractive.

The returns generated from this strategy come from the failure of the Uncovered Interest Parity (UIP), documented by Bilson (1981) and Fama (1984) (see also Hansen and Hodrick (1980, 1983), Hodrick and Srivastava (1987), and Engel (1996)). Over the past two decades, the currency carry trade strategy has garnered significant attention due to its potential for generating steady returns. Throughout the period from 1990 to 2012, such strategies delivered Sharpe ratios between 0.40 and 0.55, matching or exceeding those of common equity market factors *(Fama-French/Carhart)* (Jurek, 2014). However, this strategy is subject to multiple risks, namely crash risk, where long-accumulated gains may be erased in a very short time. Papers that study currency carry trade strategies and risks include Burnside et al. (2006, 2007, 2011), Lustig and Verdelhan (2007); Lustig et al. (2011), Brunnermeier et al. (2008), Clarida et al. (2009), Menkhoff et al. (2012), Jurek (2014), and Farhi and Gabaix (2016), which will be further explored in the literature review. The potential for returns to be rapidly erased underscores the importance of selecting currency pairs carefully, accounting for the possibility of sudden devaluations "peso events" that can lead to significant drawdowns. Also, as stated in Burnside et al. (2006), returns from currency carry trade cannot be fully explained by traditional risk factors. Given the non-Gaussian distribution of returns and the difficulty in identifying the risk factors that explain currency returns, it is thus desirable to have a portfolio optimization model tailored for currency carry trade.

Traditional currency carry trade strategies, which involve going long on currencies with high interest rate differentials and short on currencies with negative interest rate differentials relative to the quote currency (see, for example, Burnside et al. (2006) and Lustig and Verdelhan (2007)), often lack robust risk management measures to address crash risk and the inherent volatility of FX markets. Conversely, models that focus exclusively on drawdowns or study only left-tail risk may be overly conservative and lack the flexibility to dynamically manage risk during the investment period (e.g., Basak and Shapiro (2001)). Additionally, classical methods like mean-variance optimization, as proposed by Markowitz (1991), may not sufficiently account for the unique risks and return distribution characteristics of currency markets.

Recently, Zsurkis, Nicolau, and Rodrigues (2024) developed a novel portfolio optimization model based on first passage times using a non-parametric approach. This new framework in portfolio optimization chooses the portfolio that minimizes the risk of returns falling below a specified threshold within a predetermined period, while simultaneously ensuring that the investor's targeted returns are achieved within the defined time frame. By accounting for the dynamic path of returns and not relying on predefined risk factors, this model poses an interesting option to address the challenges posed by the non-linear and non-normal distribution of currency carry trade returns. The objective of this Work Project is to adapt and evaluate the performance of this newly developed portfolio optimization model in constructing portfolios exclusively composed of currency pairs. This study aims to enhance Sharpe ratios while mitigating negative skewness and leptokurtic return distributions, as well as reducing the probability of substantial drawdowns—challenges inherent to this strategy. This approach is compared against the traditional long-short strategy in currency carry trade and the widely used mean-variance optimization method.

The remainder of the Work Project is organized as follows: Section 2 presents a review of the relevant literature on currency carry trade strategies, returns, and associated risks. Section 3 outlines the methodology employed in the analysis. Section 4 reports the empirical results and Section 5 provides the conclusions.

2 Literature Review

Over the past two decades, currency carry trades have garnered significant attention due to their potential for generating steady returns (Burnside et al. (2006), Clarida et al. (2009), Menkhoff et al. (2012)). These returns can be further amplified when incorporating emerging market currencies into the strategy (Burnside et al., 2007). However, carry trades also expose investors to considerable risks, particularly from exchange rate fluctuations, unexpected changes in monetary policy (Brunnermeier et al., 2008), and economic downturns, during which high-yield currencies become more vulnerable, prompting investors to flee to safe havens (Lustig and Verdelhan, 2007; Lustig et al., 2011).

Burnside et al. (2006) investigated the returns from currency speculation and found that carry trades generate significant excess returns that cannot be fully explained by traditional risk factors such as volatility, consumption growth, or other standard risk measures. By regressing accumulated quarterly real excess returns from carry trade strategies on a variety of risk factors—including U.S. per capita consumption, S&P 500 returns, Fama and French (1993) stock market factors, and other macroeconomic indicators—they inferred that no risk factor was significantly correlated with real returns, making traditional risk explanations for carry trade returns empirically implausible. They proposed that price pressures, such as exchange rate changes in response to net order flow and increasing bid-ask spreads with larger order sizes, could explain why the strategy is not fully exploited. The authors achieved annualized Sharpe ratios of 0.69 and 0.35 after accounting for transaction costs, compared to 0.49 for the S&P 500, from 1976 to 2005. They also observed that varying forward horizons had a negligible effect on the Sharpe ratio, which is primarily driven by low volatility rather than high returns. Furthermore, speculators are required to wager very large sums of money in order to profit from this strategy.

Burnside et al. (2007) included emerging market currencies in carry trade portfolios and observed an increase in the Sharpe ratio to 1.92.

Dimic et al. (2016) further explored the effect of political risk on currency carry trades, finding that higher political risk leads to higher carry trade returns, particularly in emerging markets and countries with high-interest-rate differentials. This suggests that political risk is a significant factor contributing to the forward premium puzzle, and investors demand a premium for bearing additional political risk.

Lustig and Verdelhan (2007) approached the carry trade puzzle from a consumption-based asset pricing perspective, challenging the findings of Burnside et al. (2006). They proposed that currency risk premia are linked to global consumption growth risks. By regressing carry trade returns on relevant risk factors, they found a significant and large beta to U.S. consumption growth, indicating that this is a relevant risk factor priced in currency markets. The average interest rate currency tends to depreciate during negative durable consumption growth innovations and appreciate with positive ones. The study revealed that the significant rise in the carry trade risk premium was accompanied by a corresponding increase in the consumption and market betas of carry trade returns. These betas surged dramatically during crisis periods, such as the subprime mortgage crisis, the Russian default in 1998, and the onset of the global financial crisis in 2007. For instance, the market beta HML increased to 1.08 before the Russian default in 1998, implying a depreciation of 1.08% of high-interest-rate currencies relative to low-interest-rate currencies when the stock market declined by 1%.

Their subsequent work (Lustig et al., 2011) demonstrated that currency risk premia are determined by exposure to a single global risk factor HML_{FX}, with interest rates measuring exposure to this factor. High-interest-rate currencies are more exposed to this risk, while low-interest-rate currencies serve as havens during economic downturns. They found that forward discounts determined currencies' exposure to HML_{FX} and that this factor was exacerbated during "bad times".

Brunnermeier et al. (2008) identified 'crash risk' as the primary risk in carry trades, where investment currencies with positive interest rate differentials are linked to negative conditional skewness in exchange rate movements. This implies that long-accumulated gains can be rapidly erased. They also observed that such dramatic exchange rate movements occasionally occur without fundamental news announcements, aligning with the findings of Cutler et al. (1989), who noted that abrupt price movements cannot always be attributed to fundamentals. Brunnermeier et al. (2008) illustrate this with the sharp depreciation of the U.S. dollar against the Japanese yen on October 7 and 8, 1998. They further noted that the unwinding of carry trades may be unrelated to news when speculators approach funding constraints. Consistent with Lustig and Verdelhan (2007), they presented evidence that a higher VIX predicts future higher returns for investment currencies and lower returns for funding currencies. Controlling for the VIX reduces the predictive coefficient for interest rate differentials, aiding in resolving the forward premium puzzle.

Farhi et al. (2009) also explored the existence of crash risk priced in currency markets, finding that it accounted for more than a third of currency risk premia.

Clarida et al. (2009) stated that carry trades resemble FX option strategies that sell out-of-the-money puts on high-interest-rate currencies. Both strategies generate persistent excess returns that unwind sharply when volatility increases.

Burnside et al. (2011) argued that carry trade returns rely on the existence of “peso events”—rare but significant currency devaluations. According to the authors, these events involve high values of the stochastic discount factor rather than very large negative payoffs. They also referred to price pressure as an alternative explanation for returns in emerging markets, meaning the price at which investors can buy or sell currencies depends on the transaction quantity.

Similarly, Farhi and Gabaix (2016) introduced a disaster risk model linking currency crashes to rare but severe economic events. Their model demonstrates how the frequency and severity of these disasters drive currency movements and carry trade returns. They argued that investors bearing disaster risk receive a risk premium of 2.1%, corresponding to 36% of the currency risk.

Jurek (2014) proposed crash-neutral strategies that hedge against significant losses by using deep out-of-the-money currency options to derive currency crash risk. He found that disaster risk explains a substantial portion of carry trade returns, improving risk-adjusted returns without sacrificing profitability.

Menkhoff et al. (2012) found that high-interest-rate currencies are negatively related to innovations in global FX volatility, delivering low returns during unexpected high volatility, while low-interest-rate currencies provide a hedge by yielding positive returns. They showed that volatility risk dominates liquidity risk, emphasizing the importance of accounting for volatility in carry trade strategies. Similarly,

Filipe et al. (2023) examined the impact of funding constraints and funding risk on currency carry trades, particularly after the 2008 financial crisis. They found that funding risk significantly affects carry trade activity, returns, and currency correlations. Their theoretical model predicted that when funding constraints bind, both investment and funding currencies crash relative to safe assets like gold, highlighting the importance of funding risk in understanding carry trade dynamics.

Currency carry trade strategies, as evidenced by the literature, exhibit high skewness, significant crash risk, and elevated stochastic discount factor. These characteristics make them particularly vulnerable to sudden drawdowns and volatile return profiles.

The model developed by Zsurkis et al. (2024) offers a promising avenue for advancement within this context. By accounting for intra-horizon risk, it minimizes the probability of breaching a defined threshold within a specified time frame while still achieving the targeted returns over another temporal horizon. This approach addresses the aforementioned challenges of currency carry trades, providing a direct means of incorporating drawdowns and mitigating crash risk. Moreover, employing a non-parametric framework based on Markov chains allows for flexibility in predicting returns without relying on restrictive assumptions such as normally distributed returns.

In contrast to traditional long-short heuristics rooted in interest rate differentials, or even mean-variance (MV) frameworks, this novel approach integrates investor preferences to optimize paths of returns over time. The result is a more robust, accurate framework for managing and mitigating the inherent risks of carry trade strategies.

The literature on currency carry trades highlights both the potential for significant returns and the substantial risks involved. While early studies like Burnside et al. (2006, 2007) emphasized market inefficiencies as the source of returns, subsequent research focused on risk factors such as global consumption growth, crash risk, volatility, and funding constraints. The inclusion of emerging market currencies can enhance returns but introduces additional risks related to liquidity and default. Advancements in modeling and portfolio optimization, such as the approach by Zsurkis et al. (2024), offer a potential pathway for improving currency carry trade strategies by more effectively managing risks associated with drawdowns and crashes with a flexible and adaptive learning framework.

3 Methodology

This section outlines the methodology employed for portfolio optimization. The novel nonparametric approach proposed by Zsurkis, Nicolau, and Rodrigues (2024), which is based on first passage times is adopted, adapting the framework to currency carry trade. Currency carry trades exhibit pronounced negative skewness and significant tail risks.

The nonparametric approach by Zsurkis et al. (2024) can be advantageous to this problem, as it does not require identifying the specific risks or understanding precisely how they influence the strategy, offering a flexible and robust framework for estimating first passage times, particularly within the non-Gaussian return distribution environment typical of carry trades. This approach leverages an optimization model that accounts for intra-horizon risk, making it particularly advantageous during periods of economic uncertainty when expected gains are diminished.

The model minimizes the probability of breaching a downside threshold while ensuring the achievement of at least a required return within a specified time frame. In this Work Project, the model is optimized to secure a minimum return of 3% over 52 weeks (1 year) while reducing the likelihood of a loss of 5% or 10% during the first 13 weeks (1^{st} quarter) of the optimization period. The idea is that an asset manager can change the parameters as he/she whishes, in order to have a better control of how the portfolio will behave. In this context, we consider an asset manager implementing a cost-neutral carry trade strategy using forwards and spots, as described shortly below.

Returns are derived from the Forward Premium Puzzle, which reflects interest rate differentials between currencies, and exploits the failure of the Uncovered Interest Parity (UIP), i.e.,

$$(1+i_t) = (1+i_t^*) * E(S_{t+1})/S_t \tag{1}$$

where i_t represents the domestic interest rate, i_t^* denotes the foreign interest rate, and S_t and $E(S_{t+1})$ are the spot exchange rate and the expected future exchange rate, respectively, expressed as the amount of domestic currency per unit of foreign currency.

In this context, it is expected that currencies with higher interest rates will depreciate over time, until the point that no gains can be extracted from this strategy. However, as discussed in the literature (see, for example, Brunnermeier et al. (2008)), this outcome is inconsistent with empirical evidence, giving rise to the so-called Forward Premium Puzzle. This puzzle, which challenges the foundational Uncovered Interest Parity (UIP) model used in most open-economy international macroeconomic frameworks, stems from a combination of market inefficiencies—such as price frictions (Burnside et al., 2011)—and compensation for risk, including crash risk (Brunnermeier et al., 2008; Lustig et al., 2011). Despite these explanations, there is little academic consensus on the specific risks driving the observed deviations, which highlights the advantage of employing a nonparametric optimization approach, such as this one based on stationary Markov Chains.

In FX markets, forward rates are calculated based on expectations of future spot rates. At maturity, the deviation between the forward rate and the realized spot rate for a currency pair reflects the forward premium puzzle and the returns associated with these carry trade strategies (Fama, 1984; Engel, 1996).

In this Work Project, weekly forward and spot exchange rates are utilized, covering the period from January 8, 1999 to December 13, 2024. Forward rates are calculated using forward points, following the expression below:

$$f_t = s_t + (f_t^p/10000), \tag{2}$$

where f_t is the forward rate at time t, s_t is the spot rate at time t, and f_t^p represents the forward points at time t, expressed in basis points.[1]

Using data from the Bloomberg Terminal, a currency carry trade strategy was developed by selecting currency pairs in the perspective of an US investor where USD serves as the quote currency. The currencies included in this strategy are EUR, AUD, CAD, CZK, DKK, ILS, JPY, KWD, NZD, NOK, PLN, SAR, SGD, ZAR, SEK, CHF, THB, and GBP. The selection of currencies and the data range is based on the larger sample utilized by Lustig and Verdelhan (2007). The sample was narrowed based on data availability and the relevance of the model implementation, focusing on currencies that remain actively traded in the markets.

Additionally, certain spot and forward rates were adjusted to ensure the quote currency was consistently the USD, using the straightforward conversion formula:

$$e_{X/USD} = \frac{1}{e_{USD/X}}. \tag{3}$$

In this strategy the returns from a long position at maturity $t+1$ are defined as:

$$r_{t+1}^{l} = s_{t+1}^{pxlast} - f_t^{pxlast}, \tag{4}$$

where $r_t^l + 1$ is the return at the maturity, $s_t^{pxlast} + 1$ is the spot price at maturity and f_t^{pxlast} is the forward at the time the contract is settled.

These returns come from buying forward the foreign currency at a discount or conversely selling forward the dollar at a premium and selling the foreign currency at the market price at maturity and reversely buying the dollar. Both the spot and the forward are weekly last prices. The returns from a short position can be expressed as:

$$r_{t+1}^{s} = -r_{t+1}^{l}. \tag{5}$$

In theory, the algorithm assigns negative weights to currencies when, during the period of the

---

[1] Forward points for the currency pairs USD/CZK, USD/JPY, USD/KWD, and USD/TBH are expressed not in basis points but in units of 1000, 100, 100 000, and 100, respectively.

sample, the sum of their interest rate differentials against the US dollar is negative.

To perform the optimization, weights for each pair trade currency are initialized so that they sum to 0, i.e.,

$$\sum_{i=1}^{n} w_i = w_1 + w_2 + w_3 + \cdots + w_n = 0, \tag{6}$$

where n represents the number of currencies in the portfolio. Currencies with negative weights $(w < 0)$ represent short positions, while currencies with positive weights $(w > 0)$ represent long positions. During the optimization process, the sum of the positive weights and the absolute sum of the negative weights must each be constrained to not exceed 1, ensuring that the strategy remains cost-neutral.

To proceed with the implementation of the first passage model, consider a vector of weights ω. For $t = 1, \ldots, n$, define the partial sum of returns as,

$$\mathrm{cs}_t(\omega) \;=\; \sum_{i=1}^{t} r_i(\omega), \tag{7}$$

where $r_i(\omega)$ is the return in period i, determined by ω. Let $k_n < 0 < k_p$ be two thresholds such that crossing k_p corresponds to an "up" move of interest, and crossing k_n corresponds to a "down" move. We then define the first passage time as,

$$T^{+}(\omega) \;:=\; \inf\{t \geq 1 : \mathrm{cs}_t(\omega) \geq k_p\}, \qquad T^{-}(\omega) \;:=\; \inf\{t \geq 1 : \mathrm{cs}_t(\omega) \leq k_n\}. \tag{8}$$

In the implementation, whenever $\mathrm{cs}_t(\omega)$ exceeds k_p, the partial sum is reset to zero (i.e., a "take-profit" reset). Consequently, the events "$\mathrm{cs}_t(\omega)$ crosses k_p" and "$\mathrm{cs}_t(\omega)$ crosses k_n" can occur multiple times over an infinite horizon as $n \to \infty$. These recurring threshold-crossing events can therefore be modeled with stationary Markov chains to estimate the distributions of $T^{+}(\omega)$ and $T^{-}(\omega)$. These events are represented by a supplementary process that employs an auxiliary algorithm, which is illustrated below for S^{+} and can be easily transformed for S^{-}.

The auxiliary process algorithm

1. Set $i = 1$, let the initial cumulative sum $cs_0 = 0$, and define $r^+ > 0$ as the desired cumulative growth rate (in percentage).
2. Let $k_p = r^+$
3. For $t = i, i+1, \ldots, n$,

$$S_t^+ := \begin{cases} 1 & \text{if } cs_j < k_p, \text{ for all } j = i, i+1, \ldots, t \\ 0 & \text{otherwise;} \end{cases}$$

4. Run S_t^+ for $t = i, i+1, \ldots, \min(n, t^*)$, where t^* is such that $S_{t^*}^+ = 0$ (indicating a cumulative sum has reached or exceeded k_p). $S_0^+ = 0$ is used for initialization;
5. If $t^* = n$, the procedure is stopped;
6. Set $cs_{t^*} = 0$ and $i = t^* + 1$. Clear the value t^* and return to Step 2.

All the relevant information for the estimation of parameters from hereafter will be based on this auxiliary process S_t^+ (or S_t^-). The estimation approach models S_t^+ and S_t^- as Markov chains with a state space of $\{0, 1\}$, focusing on estimating the transition probabilities between these two states (Zsurkis et al., 2024).

The 1^{st} assumption required for the implementation of the model is the assumption of conditional homogeneity defined below:

$$P(T_{kp} \leq t | cs_t = x) = P(T_{kp} \leq t | cs_t = z), \tag{9}$$

for any values of x and z. This assumption is reasonable in this context of currency carry trade in light of the unpredictability of carry trade returns, which can unwind rapidly. Consequently, we can infer that the time it takes to reach $r\%$ is, in theory, independent of the path it has taken.

The 2^{nd} assumption, Markovian Property and Stationarity, posits that the processes S_t^+ and S_t^- are stationary discrete-time Markov processes of finite order k (Zsurkis et al., 2024). Subsequently, the probability of the threshold T^τ being breached at time t can be written as (Zsurkis et al., 2024, p.3):

$$P(T^\tau = t) = \begin{cases} (1-p_t)\prod_{i=1}^{t-1} p_i, & \text{for } t \leq k, \\ \\ \left((1-p_k)\prod_{i=1}^{k-1} p_i\right) p_k^{t-k}, & \text{for } t > k, \end{cases} \tag{10}$$

where

$$\begin{aligned} p_i &= P(S_i^\tau = 1 \mid S_{i-1}^\tau = 1, S_{i-2}^\tau = 1, \ldots, S_0^\tau = 1) \\ &= P(S_i^\tau = 1 \mid S_{i-1}^\tau = 1, S_{i-2}^\tau = 1, \ldots, S_{i-k}^\tau = 1) \\ &= P(S_k^\tau = 1 \mid S_{k-1}^\tau = 1, S_{k-2}^\tau = 1, \ldots, S_0^\tau = 1) = p_k. \end{aligned} \tag{11}$$

From which it follows that the expected time to reach the threshold is determined by the following expression (Zsurkis et al., 2024, p.3):

$$\begin{aligned} E(T^\tau) &= \sum_{t=1}^{k} t(1-p_t)\prod_{j=1}^{t-1} p_j + (1-p_k)\prod_{j=1}^{k-1} p_j \sum_{t=r+1}^{\infty} t\, p_k^{t-k} \\ &= \sum_{t=1}^{k} t(1-p_t)\prod_{j=1}^{t-1} p_j + \prod_{j=1}^{k-1} p_j\, p_k \frac{1+k-kp_k}{1-p_k}. \end{aligned} \tag{12}$$

For the estimation of the standard errors of $E(T)$ and $P(T=t)$, the variance-covariance matrix of $\hat{p} = (\hat{p}_1, \hat{p}_2, ..., \hat{p}_k)$ will be computed. For that the following multivariate linear model is considered (Zsurkis et al., 2024, p.3),

$$\begin{cases} y_{t1} = p_1 x_{t1} + \varepsilon_1 \\ y_{t2} = p_2 x_{t2} + \varepsilon_2 \\ \vdots \\ y_{tk} = p_k x_{tk} + \varepsilon_k \end{cases}, \tag{13}$$

where $y_{ti} = S_t S_{t-1} \cdots S_{t-(i-1)}$ and $x_{ti} = S_{t-1}S_{t-2}...S_{t-i}$. The system of equations can be expressed in matrix notation as:

$$y_t = X_t p + \varepsilon_t, \tag{14}$$

where $y_t = (y_{t1}, y_{t2}, \ldots, y_{tk})'$, $X_t = \mathrm{diag}(x_{t1}, x_{t2}, \ldots, x_{tk})$, $p = (p_1, p_2, \ldots, p_k)'$ and $\varepsilon_t = (\varepsilon_{t1}, \varepsilon_{t2}, \ldots, \varepsilon_{tk})'$. Building on the above, and following Zsurkis et al. (2024, p. 4), we focus on the OLS estimators,

$$\hat{p} = \left(\sum_t X_t X_t'\right)^{-1} \sum_t X_t y_t. \tag{15}$$

In this model $E(X_i \varepsilon_i) = 0$. As demonstrated by (Zsurkis et al., 2024, p.5), ε_{tk} is serially correlated due to the dependency structure of the Markov chain S_t. Thus, in order to get unbiased estimators we need to compute robust standard errors. For that we use the heteroskedasticity and autocorrelation consistent (HAC) covariance matrix defined below:

$$\sqrt{n}\,(\hat{p} - p = \left(\frac{1}{n}\sum_t^n X_i X_i'\right)^{-1} \frac{1}{\sqrt{n}} \sum_t^n X_i \varepsilon_i \xrightarrow{d} N(0, Q^{-1} S Q^{-1}). \tag{16}$$

where $Q = E(X_i X_i')$ and $S = \lim_{n->\infty} \frac{1}{n} Var(\sum_t^n X_i \varepsilon_i)$.

Regarding the Markov order estimation, we use the new methodology proposed by Zsurkis et al. (2024). This new methodology does not require the construction of multiple models of varying orders, which leads to inefficiency specially when the focus is exclusively on the probabilities p_k, rather than the entire transition matrix (Zhao et al., 2001; Katz, 1981). To estimate the order of the Markov chain efficiently, the following auxiliary regression is used,

$$S_t = \beta_{k-1} x_{t,k} + \beta_k x_{t,k-1} + \varepsilon_t, t = k+1, k+2, ..., n \tag{17}$$

where ε_t is uncorrelated with the regressors $x_{t,k}$ and $x_{t,k-1}$. Both the regressors and the dependent variable are binary, and β_k is defined as:

$$\beta_k = \frac{p_k - p_{k-1}}{1 - p_{k-1}} \tag{18}$$

where $p_k = P(S_t = 1|S_{t-1} = 1, ..., S_{t-k} = 1)$. A $t-test$ is used to evaluate the null hypothesis $H_0 : \beta_k = 0$,

$$t = \frac{\hat{\beta}_k}{\hat{\sigma}\hat{\beta}_k}. \tag{19}$$

If H_0 is not rejected, this implies $p_k = p_{k-1}$, and the true order k^* is smaller than k .

The proposed method from Zsurkis et al. (2024, p.6) determines the Markov chain order k^* sequentially by testing $H_0 : \beta_k = 0$ for increasing values of k. If the null hypothesis is not rejected at a given k, the order k^* is established as $k-1$. Figure 1 illustrates this process for a maximum order $k_{max} = 3$. The maximum order is determined by the Schwert's rule, $k_{max} = [4(n/100)^{(1/4)}]$ (Schwert, 2002; Zsurkis et al., 2024)

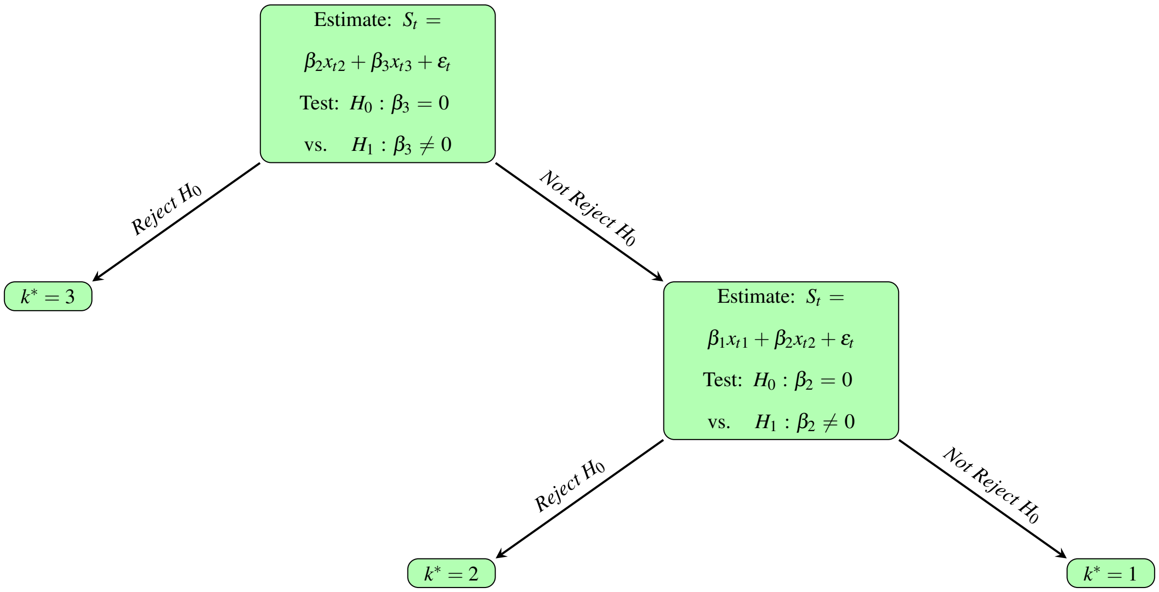

Figure 1: Decision Tree for Markov Chain Order Determination

To solve the proposed portfolio optimization problem, we define the following objective function (Zsurkis et al., 2024, p.6):

$$\Psi(\omega) = P(T^-(\omega) \leq n^-) + \gamma I_E[T^+(\omega)], \tag{20}$$

where

$$I_E[T^+(\omega)] = \begin{cases} 0, & \text{if } \mathbb{E}[T^+(\omega)] \leq n^+, \\ \mathbb{E}[T^+(\omega)], & \text{otherwise}, \end{cases} \tag{21}$$

and γ is a large positive penalty parameter (e.g., $\gamma = 1{,}000{,}000$), ensuring that solutions violating the constraint $\mathbb{E}[T^+(\omega)] \leq n^+$ are heavily penalized. Classical optimization methods are unsuitable for this problem, as evaluating $f(\omega)$ requires the following iterative steps (Zsurkis et al., 2024, p.6):

1. Construct the auxiliary processes S^+ and S^- for first passage times.
2. Estimate the Markov chain orders for S^+ and S^-.
3. Compute the probability functions for T^+ and T^-.
4. Calculate $\mathbb{E}[T^+(\omega)]$ and $P(T^-(\omega) \leq n^-)$.

To overcome these challenges, we employ the Threshold Accepting (TA) algorithm (Dueck and Scheuer, 1990; Zsurkis et al., 2024), a heuristic optimization approach particularly effective for non-smooth objective functions. The steps of the TA algorithm, as outlined further in Appendix A, are as follows:

1. Initialization: Start with a randomly chosen feasible portfolio allocation $\omega^{(0)}$.
2. Neighborhood Exploration: Generate a ”neighbor” solution ω' by perturbing $\omega^{(t)}$ using a predefined neighborhood function.
3. Acceptance Criterion: Accept ω' if $f(\omega') \leq f(\omega^{(t)})$, or if the deviation $f(\omega') - f(\omega^{(t)})$ does not exceed a specified threshold.
4. Threshold Update: Gradually reduce the threshold sequence over iterations.
5. Iteration: Repeat steps 2–4 until convergence criteria are met.

The neighborhood function defines how new solutions are generated, and the threshold sequence determines the acceptance criteria. The algorithms proposed by (Gilli et al., 2019, p. 390–394), which are well-suited for portfolio optimization, are adopted.

To access how this optimization model behaves in terms of relevant metrics, a comparison will be made against a traditional mean-variance optimization, using the function *MVportfolio* from the *NMOF* package in R (Appendix B). To ensure that the portfolio remains cost-neutral, the weights obtained from the mean-variance optimization are normalized as follows:

1. Separate Positive and Negative Weights:

$$\text{Positive Weights: } w_i > 0, \quad \text{Negative Weights: } w_i < 0 \tag{22}$$

2. Normalize Positive Weights:

$$w_i' = \frac{w_i}{\sum_{j=1}^{m} w_j \cdot \mathbb{I}(w_j > 0)} \quad \text{for } w_i > 0 \tag{23}$$

   This ensures that the sum of all positive normalized weights equals 1.

3. Normalize Negative Weights:

$$w_i' = \frac{w_i}{\sum_{j=1}^{m} |w_j| \cdot \mathbb{I}(w_j < 0)} \quad \text{for } w_i < 0 \tag{24}$$

   This ensures that the sum of all negative normalized weights equals -1.

4. Set Neutral Weights:

$$w_i' = 0 \quad \text{for } w_i = 0 \tag{25}$$

Moreover, a simple currency carry trade strategy, consisting of taking a weekly long position on the $n = 3, 4, 5$ currencies with the highest interest rate differentials against the USD and a short position on the n currencies with the lowest differentials, is also compared. This strategy is similar to that employed by Brunnermeier et al. (2008) and is additionally illustrated in Appendix C.

4 Empirical Results

The in-sample results reported in Table 1 illustrate the differences among the three main strategies: the new first passage times optimization (labeled “Optimization proposed”), the traditional mean-variance approach, and the classical long-short carry trade that, as explained before, goes long on the $n = 3, 4, 5$ highest interest rate currencies and shorts the lowest.

Table 1: Optimal Portfolios In-Sample (Weekly data from January 1999 to December 2024)

Currency	**Optimization Proposed**		**MV**	**Long–Short Portfolio**		
	$r^- = -10\%$	$r^- = -5\%$	$r^+ \geq 3\%$	**3**	**4**	**5**
EUR	-0.264	-0.191	-0.579	-0.283	-0.273	-0.268
AUD	0.156	0.055	0.009	0.376	0.317	0.275
CAD	-0.002	-0.051	-0.048	0.000	0.000	0.000
CZK	0.055	0.116	0.080	0.000	0.000	0.000
DKK	-0.077	-0.096	0.225	0.000	-0.035	-0.034
ILS	0.022	0.171	0.019	0.000	0.000	0.000
JPY	-0.188	-0.336	-0.102	0.000	0.000	0.000
KWD	0.001	0.002	0.038	0.000	0.000	0.000
NZD	0.310	0.168	0.041	0.380	0.320	0.278
NOK	-0.121	-0.146	-0.003	0.000	0.000	0.000
PLN	0.131	0.180	0.060	0.244	0.205	0.178
SAR	0.028	0.005	0.367	0.000	0.000	0.000
SGD	-0.017	-0.011	-0.096	-0.159	-0.153	-0.151
ZAR	0.113	0.074	-0.004	0.000	0.158	0.137
SEK	-0.146	-0.028	-0.116	0.000	0.000	-0.017
CHF	0.034	0.010	0.057	-0.559	-0.539	-0.530
THB	0.149	0.197	0.103	0.000	0.000	0.000
GBP	-0.185	-0.118	-0.052	0.000	0.000	0.130
Et	51	51	112	71	73	84
mean (%)	3.3%	3.4%	1.58%	2.55%	2.51%	2.16%
s.d (%)	6.0%	5.4%	3.38%	8.97%	8.97%	8.22%
Sharpe ratio	0.56	0.63	0.47	0.28	0.28	0.26
Pm						
0.1	0.98%	–	1.7%	9.8%	12.7%	9.7%
0.05	–	2.40%	1.78%	36.3%	35.8%	25.8%
max_dd	-9.63%	-12.58%	-13.4%	-18.0%	-18.2%	-18.5%
skewness	-0.742	-0.624	-0.003	-1.148	-0.979	-1.032
kurtosis	5.823	5.915	0.741	9.429	8.271	8.420
VaR0.95	-9.4%	-8.1%	-5.27%	-13.3%	-13.6%	-12.6%
VaR0.99	-15.2%	-13.8%	-7.74%	-23.6%	-23.5%	-21.5%
CVaR0.95	-14.0%	-12.3%	-6.99%	-20.9%	-20.9%	-19.3%
CVaR0.99	-22.1%	-19.4%	-10.01%	-38.0%	-36.3%	-33.5%

Notes: The expected time **(Et)** refers to the duration required to achieve the defined expected return, which in this case is 3% over 52 weeks. The **mean (%)** and standard deviation **(s.d.%)** are annualized. The **Sharpe ratio** is computed by dividing the mean by the standard deviation. **Pm** represents the probability of breaching the lower threshold that would result in a stop-loss decision. **max_dd** denotes the maximum drawdown. **VaR** and **CVaR** stand for the historical Value at Risk and the historical Conditional Value at Risk, respectively, at confidence levels of 0.95 and 0.99. The long-short portfolio is computed for the top/bottom **n= 3, 4, 5** currencies. (Appendix D)

The “Optimization proposed” determines the portfolio weights driven by the requirement that the annual return stays above 3% while reducing the probability of suffering a drop below either 10% or 5% in the first quarter. This is important because the model will define weights that allow the investor to adopt a more cautious strategy in the following weeks whilst still achieving the predefined return at a longer maturity. Thus, the aim is to reduce the risk of pronounced tails in theory, when market conditions might not be favorable.

This risk control objective explains why some currencies have relatively large or small allocations compared to the other models, reflecting the inherent trade–off between capturing yield differentials and minimizing downside probabilities.

Focusing on the weights in the “Optimization proposed” columns, one notices considerable short positions in EUR, GBP, JPY, and NOK, while simultaneously going long in currencies such as NZD, ILS, THB, and PLN. These allocations reflect not only the relative interest rate differentials (and thus carry potential) but also each currency’s volatility and historical drawdown patterns. NZD, for example, has historically offered higher yields and comparatively stable short-term performance, which explains its positive weight in the new allocation. Conversely, the strategy holds negative weights in classic funding currencies like JPY and EUR, reflecting their tendency toward lower rates but also, in some episodes, safe haven behavior that can generate abrupt reversals-an aspect that the optimization attempts to manage via controlling downside risk. The short positions in NOK, GBP, and other currencies can also stem from their modest carry advantage coupled with occasionally larger drawdowns, so under the first passage perspective, they prove costlier in terms of hitting a lower threshold early in the holding period.

Comparing these allocations with the traditional mean-variance solution (under a target return of at least 3%) reveals how focusing on variance minimization alone can create different trade–offs. Besides achieving a relatively low annualized return of 1.58% and maintaining a modest volatility of 3.38%, the mean-variance approach does not explicitly control the probability of exceeding a specific drawdown threshold. Consequently, it experiences larger drawdowns compared to the proposed optimization. Nonetheless, the in-sample analysis reveals that the VaR and CVaR values are the smallest in absolute terms compared to the other methods evaluated, for instance, the 95% VaR is –5.27%. By contrast, the first passage times portfolios show higher average returns (3.3% and 3.4%), albeit with volatilities of around 5.4%-6.0%. The resulting Sharpe ratios (0.56 and 0.63) are notably above that of the mean-variance solution (0.47), suggesting that, at least in-sample, the new method provides a better balance of risk-adjusted returns once one acknowledges the importance of short-term downside control. This is further evidenced by the probabilities P_m of violating each threshold: for the first passage portfolios, P_m at a -10% loss remains near or below 1%, whereas the mean-variance approach exhibits a higher chance of breaching the lower threshold within thirteen weeks.

Meanwhile, the classic long–short carry portfolios (with 3, 4, and 5 currencies on each side) exhibit annualized returns around 2.5%–2.1% and higher volatility of approximately 8%–9%, leading to Sharpe ratios in the 0.26–0.28 range. This underscores the well documented tradeoff in straightforward carry strategies, which can earn moderate returns but often experience large drawdowns—particularly during "crashes" when risk sentiment reverses.

Indeed, as Brunnermeier et al. (2008) and others have pointed out, these carry trade strategies can generate negatively skewed returns and leptokurtic distributions (fat tails), a feature reflected here by the negative skewness values (around -1.0) and elevated kurtoses (greater than 8). The first passage optimization also shows negative skewness (near -0.6 or -0.7) and moderate excess kurtosis (around 5.8–5.9), but those are less extreme than the long–short carry's tails and, importantly, the average drawdowns remain smaller. These skewness and kurtosis values confirm a characteristic attributed to currency carry trades mentioned in the literature (see, Clarida et al., 2009): their payoffs resemble picking up small gains for extended periods, punctuated by occasional but severe crashes. Despite providing attractive carry, high yield currencies suffer heavily when market participants suddenly rush to unwind positions during "flight to quality" episodes, causing large negative returns. The new first passage times approach tries to mitigate this by focusing explicitly on the probability of hitting a loss threshold in the short run, thus penalizing positions whose historical return paths show deep drawdowns within short windows.

Overall, the in-sample evidence suggests that incorporating an intra horizon downside constraint leads to allocations that can preserve moderate to high returns while keeping a tight rein on tail risk. Consequently, the new optimization strategy typically displays less extreme drawdowns (as seen in the max drawdown figures of around -9.6% to -12.6% versus -18% or worse for the long–short portfolios) and keeps the probability of early losses low.

Although mean-variance methods remain a foundational approach in portfolio design, these results highlight the merit of a first passage times framework for investors concerned about interim losses and tail events—a situation particularly relevant for carry trades, where a few sharp reversals can undo months of accrued gains.

To test the reliability of the model, we examine its performance out-of-sample (Gilli et al., 2019; Zsurkis et al., 2024). Specifically, we utilize a rolling window spanning 19 years to determine the optimal currency weights for the carry trade strategy each week. This approach yields a series of forecasted weights over a period of 362 weeks, covering the time frame from 2018 through 2024.

The out-of-sample performance shown in Table 2 (with plots in Appendix E) provides a clear illustration of how the different strategies fare under the most recent years, including market stresses like the COVID-19 shock in early 2020—the largest turbulence in that window. It is worth noting that differences emerge when comparing the newly proposed first passage times based optimization, the traditional mean-variance approach, and the simple long–short carry strategies.

Focusing first on the proposed optimization (with thresholds of -10% and -5% over a 13 week horizon), we see that targeting a 3% annual return while restricting the probability of breaching the early loss threshold yields annualized returns of 1.95% (for the -10% portfolio) and 1.19% (for the -5% one). Their volatilities (4.68% and 3.85%, respectively) are moderate relative to the carry strategies, and their maximum drawdowns (about -8.2% and -6.5%) are materially smaller than those experienced by the long–short portfolios (which range from -10.0% to -11.9%). This difference is especially meaningful given that COVID 19 caused sudden market dislocations, quickly eroding many carry trades. By explicitly emphasizing the avoidance of short-term losses, the first passage method manages to keep larger drawdowns in check, which can be especially beneficial

Table 2: Optimal Portfolios Out-of-Sample (Rolling-window with weekly re-balances from January 2018 to December 2024)

	Optimization Proposed		**MV ($r^+ = 3\%$)**	**Long–Short Portfolio**		
	$r^- = -10\%$	$r^- = -5\%$		**3**	**4**	**5**
mean (%)	1.95%	1.19%	1.86%	2.19%	2.03%	1.38%
s.d (%)	4.68%	3.85%	3.20%	5.74%	4.75%	4.05%
Sharpe ratio	0.42	0.31	0.58	0.38	0.43	0.34
max_dd	-8.2%	-6.5%	-6.4%	-11.9%	-10.0%	-10.9%
skewness	0.051	-0.276	-0.687	-0.332	-0.392	-0.494
kurtosis	0.038	1.097	6.048	0.763	0.914	2.370
VaR0.95	-7.55%	-5.69%	-4.36%	-8.77%	-7.95%	-6.23%
VaR0.99	-9.80%	-10.16%	-7.76%	-13.90%	12.34%	-10.46%
CVaR0.95	-9.12%	-8.53%	-7.10%	-12.48%	-10.90%	-9.23%
CVaR0.99	-11.02%	-12.54%	-13.22%	-18.40%	-15.17%	-14.56%
Top (most positive)						
1st row	NZD	THB	SAR	ZAR	SAR	CZK
2nd row	THB	ILS	DKK	AUD	ZAR	PLN
3rd row	AUD	NZD	THB	CAD	AUD	SAR
Bottom (most negative)						
1st row	GBP	JPY	SEK	GBP	EUR	AUD
2nd row	SEK	SEK	SGD	EUR	CHF	CHF
3rd row	NOK	GBP	EUR	CHF	GBP	GBP

Notes: The **mean (%)** and standard deviation **(s.d.%)** are annualized. The **Sharpe ratio** is computed by dividing the mean by the standard deviation. **Pm** represents the probability of breaching the lower threshold that would result in a stop-loss decision. **max_dd** denotes the maximum drawdown. **VaR** and **CVaR** stand for the historical Value at Risk and the historical Conditional Value at Risk, respectively, at confidence levels of 0.95 and 0.99. The long-short portfolio is computed for the top/bottom **n= 3, 4, 5** currencies. The top and bottom currencies were determined based on the median of weights assigned throughout the sample. See Appendix D for currency details.

during crisis episodes.

Examining the mean-variance optimization strategy (also at a 3% target) reveals a somewhat higher annualized return of 1.86% and a lower volatility of 3.2%, leading to the best Sharpe ratio in the table (0.58). This outcome shows that variance-centric optimization, where smoothing out fluctuations can pay dividends in terms of risk-adjusted performance. However, its maximum drawdown (-6.4%) is similar to that of the more conservative first passage portfolio (-6.5% for the -5% threshold) and does not specifically guard against near term tail events. Indeed, the mean-variance solution also exhibits somewhat large tail risk measures: its $\text{VaR}_{0.99}$ (-7.76%) is not too severe, but its $\text{CVaR}_{0.99}$ extends to -13.22%, indicating that large negative returns still occur if the market experiences a downturn.

Moreover, it presents the largest kurtosis of the strategies and is the most negatively skewed strategy, showing a heavier mass of negative returns relative to other approaches. This arises because currency carry trades do not strictly conform to the geometric Brownian motion (GBM) assumption on which mean-variance optimization relies. Consequently, negative skewness and extreme tail risk may be underrepresented when using a pure variance-based model.

By contrast, the long-short carry portfolios (with 3, 4, or 5 currencies on each side) show higher average returns (2.19%, 2.03%, and 1.38%, respectively) in some cases but also significantly higher volatility (up to 5.74%) and generally deeper maximum drawdowns (-10% to -11.9%). These carry portfolios remain vulnerable to "crash risk," a hallmark of carry trades documented extensively by Brunnermeier and others. The negative skewness (around -0.33 to -0.49) and, in some cases, relatively higher kurtosis (e.g., 2.370 for the 5 currency portfolio) highlight a "fat tailed" return distribution where intermittent but severe losses can wipe out months of accrued gains. Indeed, when market sentiment flipped sharply in early 2020, these strategies suffered disproportionate drawdowns as investors unwound risk positions.

Turning to tail metrics such as VaR and CVaR, the proposed approach with a -10% threshold has $\text{VaR}_{0.95}$ and $\text{VaR}_{0.99}$ around -7.55% and -9.80%, respectively, and the -5% threshold portfolio further tightens risk at -5.69% and -10.16%. In contrast, the long–short carry approach with three currencies sees $\text{VaR}_{0.99}$ around -13.90%, and even deeper tail losses (to -18.4%) in $\text{CVaR}_{0.99}$. This discrepancy reaffirms that while carry strategies can harvest somewhat higher returns in benign markets, they remain exposed to tail events—especially short-term plunges that the first passage methodology was explicitly designed to reduce.

Finally, the currency allocations-shown by the "Top" (most positive) and "Bottom" (most negative) lines-illustrate how each method chooses or avoids certain currencies. For instance, NZD, THB, ILS, and AUD consistently appear in the "Top" for the proposed and mean-variance solutions, consistent with their comparatively favorable yield or relatively stable returns over this period. Meanwhile, GBP, JPY, and SEK tend to land in the "Bottom," reflecting their usage as either funding currencies or as currencies perceived to have had higher crash risk. As an illustration the GBP, relative to the USD has lost close to 15 percent of its value in the weeks after the Brexit decision (Plakandaras et al., 2017), which provoked following devaluations in the subsequent years. The long–short portfolios likewise favor high yielders such as ZAR or SAR on the long side, but they bear correspondingly large exposure when funding conditions swing. In sum, the out-of-sample evidence between 2018 and 2024 confirms that the newly proposed first passage times optimization can meaningfully reduce short-term crash risk while still delivering moderate returns - a particularly valuable trait during market shocks. By contrast, the mean-variance approach maintains a good Sharpe ratio but does not specifically protect against abrupt downward moves, whereas long–short carry continues to exhibit the classic pattern of steady gains punctuated by sudden drawdowns, reflected in its negative skewness and high tail risk. This underlines the practical appeal of first passage optimizations for investors wishing to limit early losses while preserving a target rate of return, offering a balanced alternative that accounts for the pronounced skewness and tail risk inherent in currency markets.

5 Conclusion

This Work Project set out to investigate whether a portfolio optimization framework based on first passage times can improve risk-adjusted performance in currency carry trades relative to traditional approaches. Grounded in the extensive literature on the forward premium puzzle, crash risk, and rare-event dynamics, the model proposed by Zsurkis et al. (2024) was adapted and tested on a set of currencies spanning both developed and emerging markets. By focusing on the intra-horizon probability of hitting a pre-defined drawdown threshold while maintaining a target return, the first passage times optimization addresses many of the shortcomings associated with the negatively skewed, leptokurtic return distributions observed in carry trade strategies.

The empirical analysis provided evidence that including a short-term downside constraint can yield tangible benefits. In both in-sample and out-of-sample tests, first passage times portfolios exhibited fewer large drawdowns, smaller tail risk measures, and improved Sharpe ratios compared to straightforward long–short carry trades. These portfolios also outperformed the classical mean-variance optimization in terms of controlling the probability of early losses—even though the mean-variance approach sometimes achieved a higher risk-adjusted return, it did not specifically mitigate abrupt reversals or the fat-tailed nature of currency returns. In contrast, the newly proposed approach delivered moderate to high average returns with a deliberate cap on short-term downside risk.

Several insights emerge from these findings. First, currency carry strategies remain exposed to "crash risk," often triggered by sudden reversals in risk sentiment or shifts in monetary policy. Second, a focus on longer-horizon averages—common in both classical carry heuristics and mean-variance frameworks—risks overlooking the substantial drawdowns that can accumulate over brief intervals. The first passage times perspective is more aligned with investors who seek to prevent these interim losses from escalating, particularly when higher leverage amplifies downside exposure. Third, incorporating Markov chain-based, non-parametric probability estimates of hitting a drawdown threshold offers the flexibility to capture the path-dependent behavior of returns. This is especially valuable in foreign exchange markets where return distributions often deviate significantly from Gaussian assumptions.

The proposed model does present avenues for future research. First, while the flexibility to set a drawdown threshold (e.g., 5% or 10%) is undoubtedly a strength, the choice of this level remains inherently subjective. Different investors have varying risk appetites and mandates, meaning the model's outputs can shift significantly based on the chosen thresholds. Second, future studies might refine the model further to address additional risk factors such as funding constraints and transaction costs.

Overall, this Work Project contributes to the literature on currency carry trades by presenting a novel approach that explicitly addresses intra-horizon risk. The evidence suggests that this non-parametric approach, based on first passage times optimization, offers a viable alternative for practitioners concerned about the pronounced skewness and tail events endemic to currency strategies. By directly incorporating path-dependent risk considerations, this method can help align currency portfolios more closely with investor tolerance for short-term losses without fully sacrificing long-term carry potential.

References

Basak, S. and A. Shapiro (2001). Value-at-risk-based risk management: optimal policies and asset prices. *The review of financial studies 14*(2), 371–405.

Bilson, J. F. O. (1981, April). Profitability and stability in international currency markets. Working Paper 664, National Bureau of Economic Research.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen (2008). Carry trades and currency crashes. *NBER macroeconomics annual 23*(1), 313–348.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2006, August). The returns to currency speculation. Working Paper 12489, National Bureau of Economic Research.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). Do peso problems explain the returns to the carry trade? *The Review of Financial Studies 24*(3), 853–891.

Burnside, C., M. Eichenbaum, and S. Rebelo (2007). The returns to currency speculation in emerging markets. *American Economic Review 97*(2), 333–338.

Clarida, R., J. Davis, and N. Pedersen (2009). Currency carry trade regimes: Beyond the fama regression. *Journal of International Money and Finance 28*(8), 1375–1389.

Cutler, D. M., J. M. Poterba, and L. H. Summers (1989). What moves the stock market. *Journal of Portfolio Management 15*(3), 4–11.

Dimic, N., V. Orlov, and V. Piljak (2016). The effect of political risk on currency carry trades. *Finance Research Letters 19*, 75–78.

Dueck, G. and T. Scheuer (1990). Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of computational physics 90*(1), 161–175.

Engel, C. (1996). The forward discount anomaly and the risk premium: A survey of recent evidence. *Journal of empirical finance 3*(2), 123–192.

Fama, E. F. (1984). Forward and spot exchange rates. *Journal of monetary economics 14*(3), 319–338.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of financial economics 33*(1), 3–56.

Farhi, E., S. P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009). Crash risk in currency markets. Technical report, National Bureau of Economic Research.

Farhi, E. and X. Gabaix (2016). Rare disasters and exchange rates. *The Quarterly Journal of Economics 131*(1), 1–52.

Filipe, S. F., J. Nissinen, and M. Suominen (2023). Currency carry trades and global funding risk. *Journal of Banking & Finance 149*, 106800.

Gilli, M., D. Maringer, and E. Schumann (2019). *Numerical methods and optimization in finance*. Academic Press.

Hansen, L. P. and R. J. Hodrick (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of political economy 88*(5), 829–853.

Hansen, L. P. and R. J. Hodrick (1983). Risk averse speculation in the forward foreign exchange market: An econometric analysis of linear models. *Exchange rates and international macroeconomics 113*, 152.

Hodrick, R. J. and S. Srivastava (1987). Foreign currency futures. *Journal of International Economics 22*(1-2), 1–24.

Jurek, J. W. (2014). Crash-neutral currency carry trades. *Journal of Financial Economics 113*(3), 325–347.

Katz, R. W. (1981). On some criteria for estimating the order of a markov chain. *Technometrics 23*(3), 243–249.

Lustig, H., N. Roussanov, and A. Verdelhan (2011). Common risk factors in currency markets. *The Review of Financial Studies 24*(11), 3731–3777.

Lustig, H. and A. Verdelhan (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review 97*(1), 89–117.

Markowitz, H. M. (1991). Foundations of portfolio theory. *The journal of finance 46*(2), 469–477.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012). Carry trades and global foreign exchange volatility. *The Journal of Finance 67*(2), 681–718.

Plakandaras, V., R. Gupta, and M. E. Wohar (2017). The depreciation of the pound post-brexit: Could it have been predicted? *Finance Research Letters 21*, 206–213.

Schwert, G. W. (2002). Tests for unit roots: A monte carlo investigation. *Journal of business & economic statistics 20*(1), 5–17.

Zhao, L., C. Dorea, and C. Gonçalves (2001). On determination of the order of a markov chain. *Statistical inference for stochastic processes 4*, 273–282.

Zsurkis, G., J. Nicolau, and P. M. Rodrigues (2024). First passage times in portfolio optimization: A novel nonparametric approach. *European Journal of Operational Research 312*(3), 1074–1085.

6 Appendix

A Auxiliary algorithms for the optimization proposed

Algorithm 1: Neighborhood Algorithm (Zsurkis et al., 2024, p.10)

1. Set ε, which is determined by a draw of a uniformly distributed random variable over $[0, 0.005]$;
2. Randomly select $j_1 \in \{\text{assets with weight} > w_{\text{min}}\}$;
3. Randomly select $j_2 \in \{\text{assets with weight} < w_{\text{max}}\}$;
4. $\omega_{j_1} = \omega_{j_1} - \varepsilon$;
5. $\omega_{j_2} = \omega_{j_2} + \varepsilon$;

where w_{min} and w_{max} are such that $w_{\text{min}} < w_j < w_{\text{max}}$ for $j = 1, \ldots, m$.

Algorithm 2: Threshold Sequence Algorithm (Zsurkis et al., 2024, p.10)

1. Set the number of thresholds n_{rounds} and the number of random steps n_{deltas};
2. **for** $i = 1 : n_{\text{deltas}}$ **do**
    (a) Randomly generate a feasible current solution, say, x_c;

    (b) Generate $x_n \in \mathcal{N}(x_c)$ and compute

    $$i = |(x_n) - (x_c)|,$$

    where $\mathcal{N}(x_c)$ is a neighbor of the current solution defined using the neighborhood algorithm above;

    (c) Set $x_c = x_n$;

    end for;
3. Compute the empirical distribution CDF of i, $i = 1, \ldots, n_{\text{deltas}}$;
4. Compute the threshold sequence

$$\tau_k = \text{CDF}^{-1}\left(\frac{n_{\text{rounds}} - k}{n_{\text{rounds}}}\right), \quad k = 1, \ldots, n_{\text{rounds}},$$

where n_{rounds} are equidistant quantiles.

Algorithm 3: Threshold Accepting (TA) Algorithm (Zsurkis et al., 2024, p.10)

```
1. Use the threshold sequence algorithm to construct the threshold sequence $\tau$;
2. Randomly generate a feasible current solution, say, $x_c$;
3. Set $x^* = x_c$;
4. for $r = 1 : n_{\text{rounds}}$ do
    (a) for $i = 1 : n_{\text{steps}}$ (where $n_{steps}$ is the number of steps per threshold) do
        i. Generate $x_n \in \mathcal{N}(x_c)$ and compute
               $\Delta = (x_n) - (x_c);$
        ii. if $\Delta < \tau_r$ then $x_c = x_n$;
        iii. if $(x_c) \leq (x^*)$ then $x^* = x_c$;
5. end for;
Return $x^*$.
```

B Construction of Mean-Variance Optimized Portfolios Using `MVPortfolio`

This appendix details the methodology employed to construct mean-variance optimized portfolios utilizing the `MVPortfolio` function from the `NMOF` package in `R`.

The mean-variance optimization problem is formulated to identify portfolio weights $\mathbf{w}$ that minimize the portfolio variance for a given expected return. The mathematical representation of this optimization is:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \mathbf{w}^\top \Sigma \mathbf{w} \\ \text{subject to} \quad & \mathbf{w}^\top \mathbf{1} = 1, \\ & \mathbf{w}_{\text{min}} \leq \mathbf{w} \leq \mathbf{w}_{\text{max}}, \\ & \mathbf{w}^\top \mu \geq \mu_{\text{target}}, \end{aligned}$$

where:

- Σ is the covariance matrix of asset returns.
- $\mathbf{w}$ is the vector of portfolio weights.
- $\mathbf{1}$ is a vector of ones, ensuring the weights sum to unity.
- μ is the vector of expected returns for each asset.
- μ_{target} is the target portfolio return.
- $\mathbf{w}_{\text{min}}$ and $\mathbf{w}_{\text{max}}$ are the minimum and maximum allowable weights, respectively.

For the optimization, the following parameters are set:

- $\mu_{\text{target}} = \frac{0.03}{52}$ (assuming an annual expected return of 3%, converted to a weekly basis).
- $\mathbf{w}_{\text{min}} = -1$.
- $\mathbf{w}_{\text{max}} = 1$.

The expected returns vector μ is assumed to be constant at $\frac{0.03}{52}$ for all assets, reflecting a uniform weekly return expectation across the portfolio.

The optimization procedure is executed as follows:

1. Estimate Covariance Matrix: Using historical return data from the in-sample period, compute the covariance matrix Σ:

   $$\Sigma = \text{cov}(\texttt{r}),$$

   where r is the returns matrix by currency.
2. Set Expected Returns: Define the expected returns vector μ with each element set to $\frac{0.03}{52}$.
3. Define Constraints and Execute Optimization: Utilize the `MVPortfolio` function to solve the optimization problem with the specified constraints.
4. Apply Weights to Next Period: The optimized weights **w** are applied to the asset returns in the subsequent period to compute the portfolio return:

   $$\text{StrategyReturn}_t = \mathbf{w}^\top \mathbf{R}_t,$$

   where $\mathbf{R}_t$ is the vector of asset returns at time t.
5. Rolling Optimization (Out-of-Sample): For out-of-sample analysis, the optimization is performed in a rolling window manner. At each time step, the covariance matrix and expected returns are re-estimated using the most recent in-sample data, and new weights are computed to be applied in the subsequent period.

C Construction of Out-of-Sample Long-Short Portfolios

This appendix thoroughly explains the methodology behind the construction of the long-short portfolios. The core idea is to build, at each time t, a portfolio that goes long on the currencies with the highest positive differentials and short on those with the most negative differentials, computed as the spot exchange rate minus the weekly forward at time t. The following steps are applied *period by period* to capture an out-of-sample perspective:

1. Ranking of Differentials. For each date t, we extract the vector of differentials and sort it in descending order.
2. Top- and Bottom-n. After sorting, we pick the top n currencies (largest differentials) to go long and the bottom n currencies (smallest differentials) to go short.
3. Equal Weighting. Assign weights $+\frac{1}{n}$ for each of the top n currencies (long positions) and $-\frac{1}{n}$ for each of the bottom n currencies (short positions). For example, if $n = 3$, then each of the three top differentials receives weight $+\frac{1}{3}$, while the three bottom differentials receive $-\frac{1}{3}$.
4. Repeat Over Time. We repeat the above steps for each row (date) of the differential matrix to obtain a full matrix of weights W, whose rows correspond to time t and columns to the currencies.

Once we have the weight matrix W, the resulting long-short strategy return at time t is simply the dot product of the weight vector at time t with the vector of ex-post returns at time t. In practice,

$$\text{StrategyReturn}_t = \sum_j \left(w_{t,j}\right) \left(\text{Return}_{t,j}\right),$$

where $w_{t,j}$ is the weight for currency j on date t, and $\text{Return}_{t,j}$ is the actual weekly return of currency j on date t. This is computed row-by-row in `R` using `rowSums`.

C.1 In-Sample Variant (Using Averages)

The *in-sample* approach follows the same logic but replaces the period-by-period updating of differentials with a single summary statistic. Instead of computing different weights for every period, one can:

- Compute the mean differential for each currency relative to the USD over the entire in-sample window,
- Rank currencies by these mean differentials,
- Assign a single set of top-n and bottom-n weights using the same logic of $+\frac{1}{n}$ and $-\frac{1}{n}$,
- Apply this same set of weights throughout the period in question.

In other words, for an in-sample strategy, the portfolio is *fixed* based on aggregated statistics over the in-sample horizon rather than updated dynamically as new data arrives.

D Data

Table 3: Tickers for the currency pairs obtained from Bloomberg Terminal

Currency	**Spot Exchange Rate**	**Forward Points**
EUR/USD	EUR BGN CURNCY	EUR1W BGN CURNCY
AUD/USD	AUD BGN CURNCY	AUD1W BGN CURNCY
USD/CAD	CAD BGN CURNCY	CAD1W BGN CURNCY
USD/CZK	CZK BGN CURNCY	CZK1W BGN CURNCY
USD/DKK	DKK BGN CURNCY	DKK1W BGN CURNCY
USD/ILS	ILS BGN CURNCY	ILS1W BGN CURNCY
USD/JPY	JPY BGN CURNCY	JPY1W BGN CURNCY
USD/KWD	KWD BGN CURNCY	KWD1W BGN CURNCY
NZD/USD	NZD BGN CURNCY	NZD1W BGN CURNCY
USD/NOK	NOK BGN CURNCY	NOK1W BGN CURNCY
USD/PLN	PLN BGN CURNCY	PLN1W BGN CURNCY
USD/SAR	SAR BGN CURNCY	SAR1W BGN CURNCY
USD/SGD	SGD BGN CURNCY	SGD1W BGN CURNCY
USD/ZAR	ZAR BGN CURNCY	ZAR1W BGN CURNCY
USD/SEK	SEK BGN CURNCY	SEK1W BGN CURNCY
USD/CHF	CHF BGN CURNCY	CHF1W BGN CURNCY
USD/THB	THB BGN CURNCY	THB1W BGN CURNCY
GBP/USD	GBP BGN CURNCY	GBP1W BGN CURNCY

Notes: The currency pairs were obtained from the Bloomberg Terminal, spanning from January 1999 until December 2024. For currencies with USD as the base currency, the pairs were transformed so that USD became the quote currency using the formula $e_{x/USD} = \frac{1}{e_{USD/x}}$. Forward points were added to the spot exchange rates prior to conversion in units of basis points, except for the currency pairs USD/CZK, USD/JPY, USD/KWD, and USD/TBH, which are expressed in units of 1,000; 100; 100,000; and 100, respectively.

E Figures

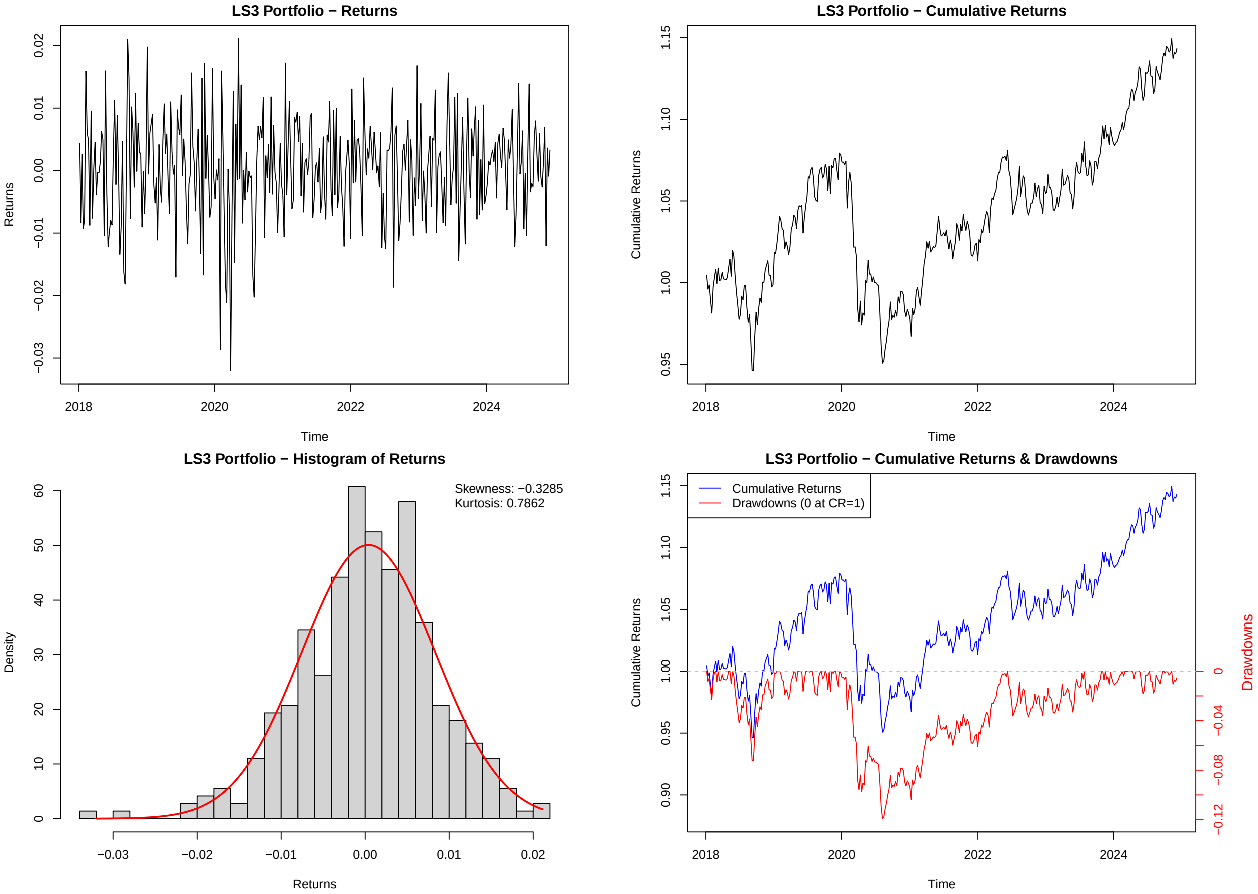

Figure 2: Long-short strategy portfolio, $n = 3$, Out-of-sample

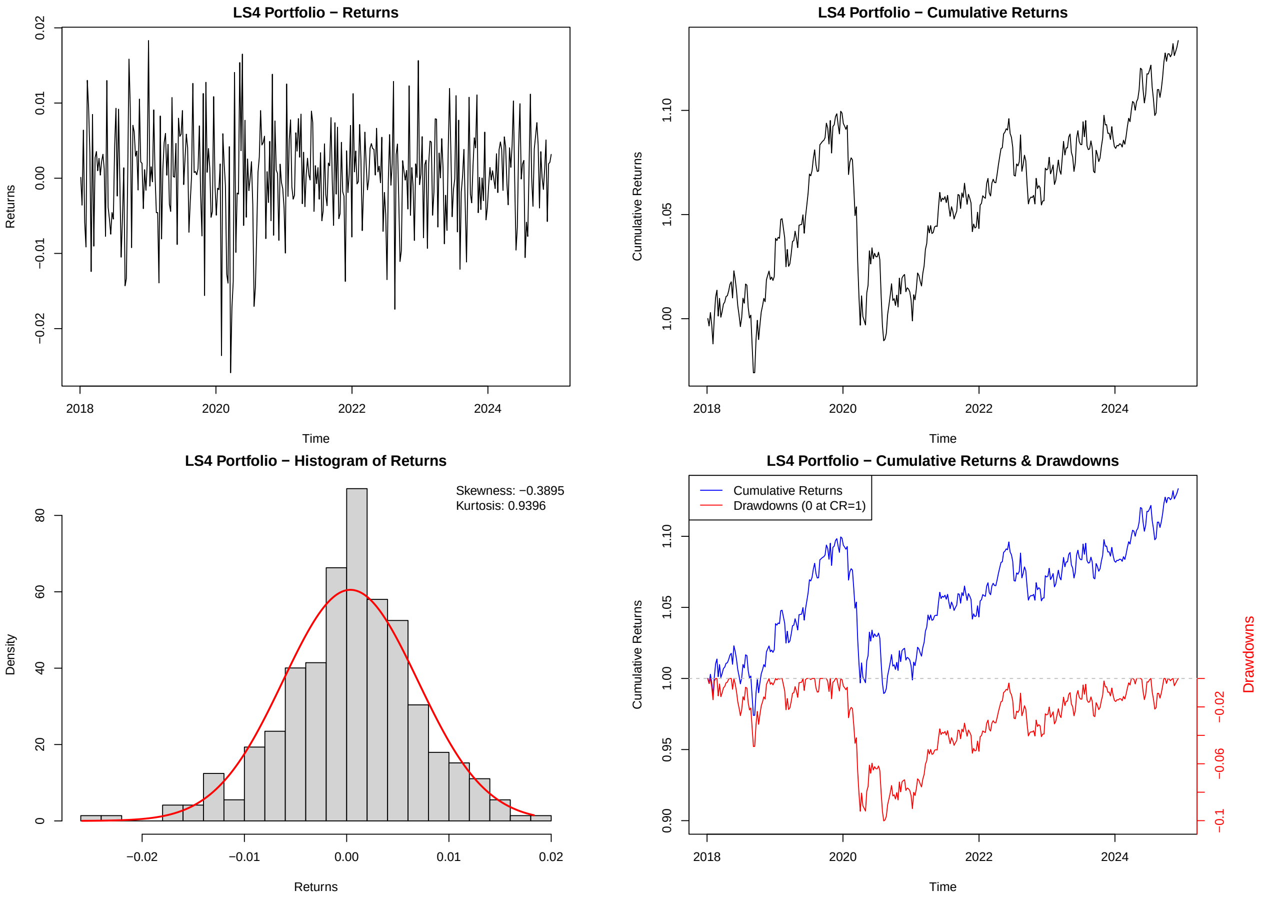

Figure 3: Long-short strategy portfolio, $n = 4$, Out-of-sample

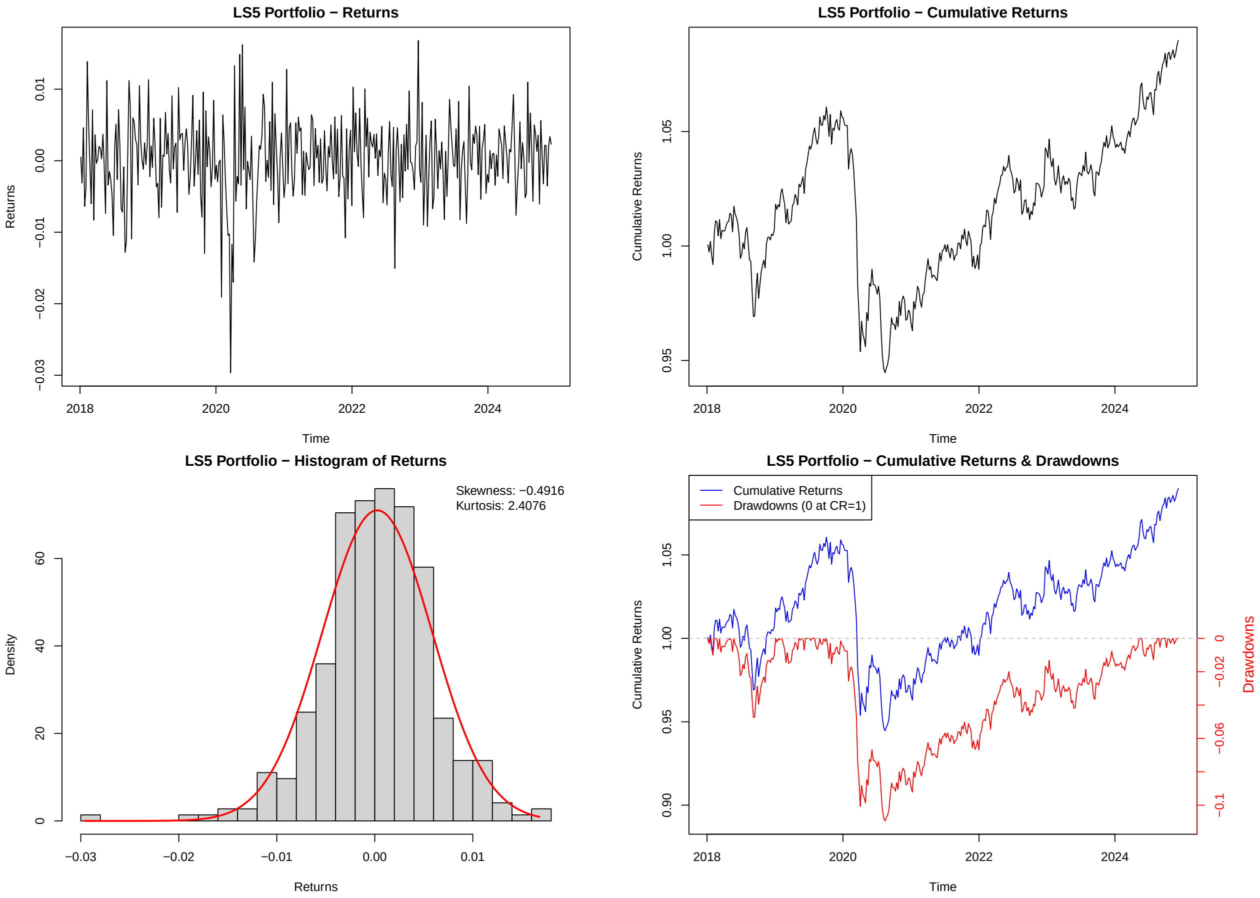

LS5 Portfolio − Returns

LS5 Portfolio − Cumulative Returns

LS5 Portfolio − Histogram of Returns

LS5 Portfolio − Cumulative Returns & Drawdowns

Figure 4: Long-short strategy portfolio, $n = 5$, Out-of-sample

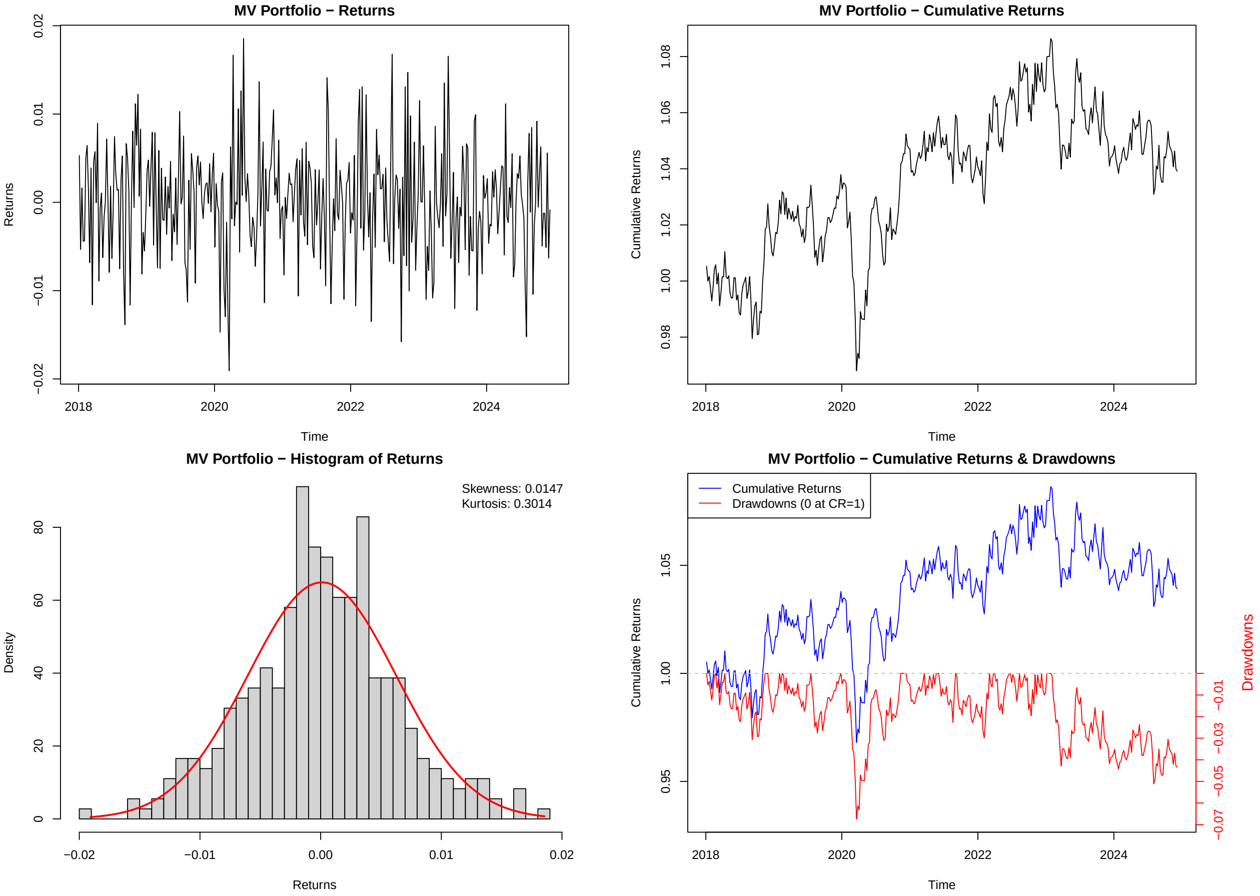

Figure 5: Mean-variance optimization, Out-of-sample

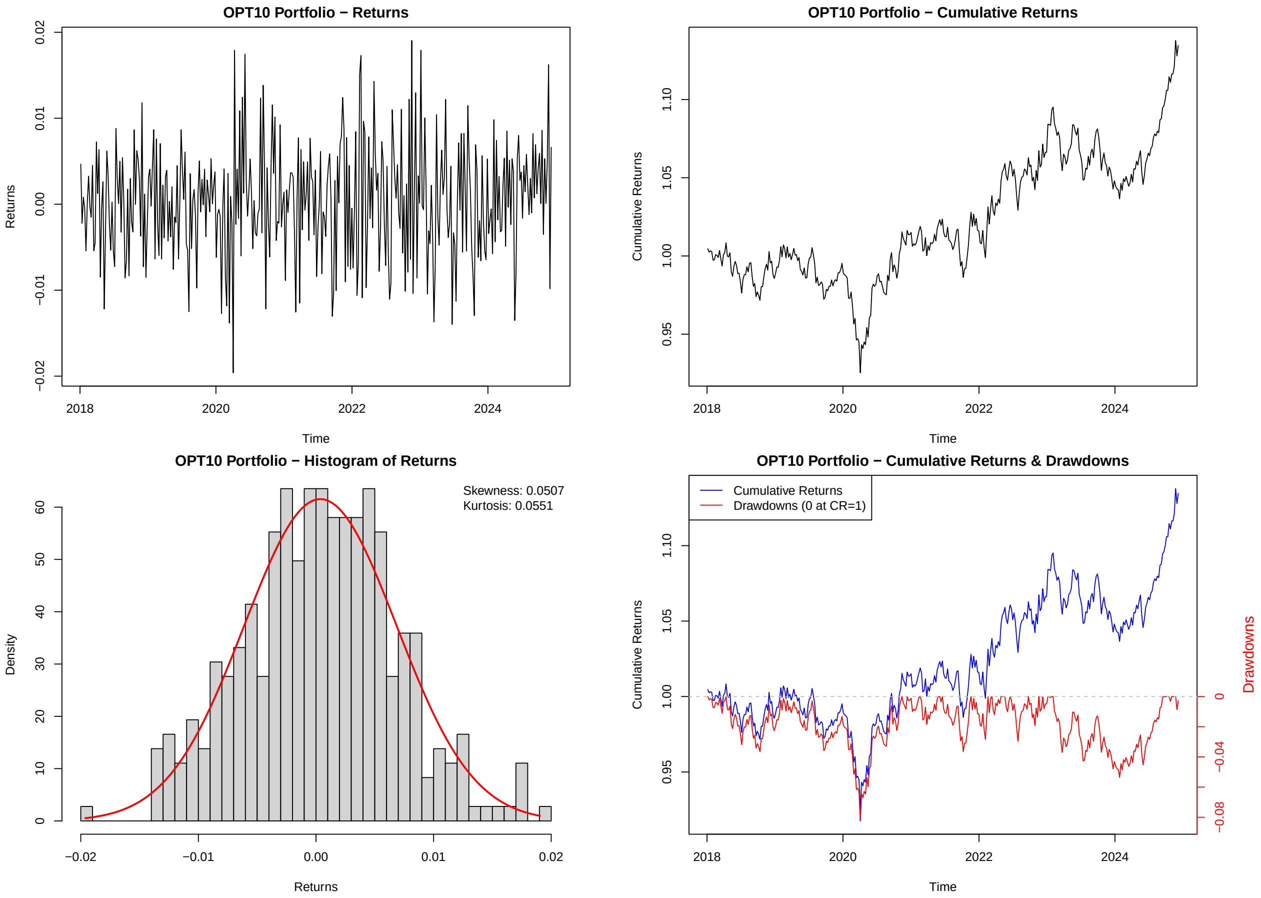

Figure 6: Optimization Proposed, $r^{-} = 0.10$ (lower-threshold), Out-of-sample

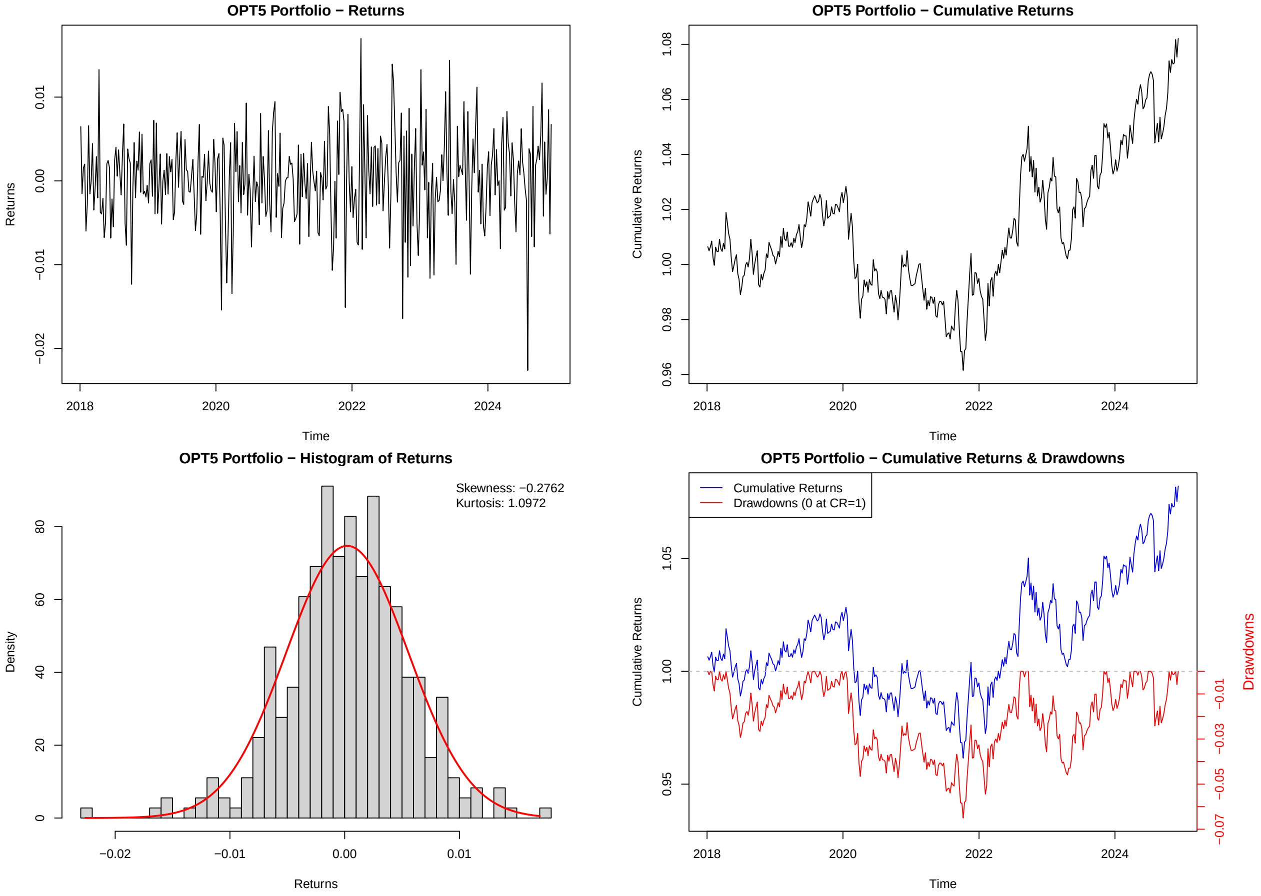

Figure 7: Optimization Proposed, $r^{-} = 0.05$ (lower-threshold), Out-of-sample