NHH

Norwegian School of Economics

Bergen, Fall 2020

Risk and Return in Yield Curve Arbitrage

A Survey of the USD and EUR Interest Rate Swap Markets

Brage Ager-Wick and Ngan Luong

Supervisor: Petter Bjerksund

Master's Thesis, Financial Economics

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

We would like to thank Petter Bjerksund for his patient guidance and valuable insights.

The empirical work for this thesis was conducted in R. R-scripts can be shared upon request.

Abstract

This thesis extends the research of Duarte, Longstaff and Yu (2007) by looking at the risk and return characteristics of yield curve arbitrage. Like in Duarte et al., return indexes are created by implementing a particular version of the strategy on historical data. We extend the analysis to include both USD and EUR swap markets. The sample period is from 2006-2020, which is more recent than in Duarte et al. (1988-2004). While the USD strategy produces risk-adjusted excess returns of over five percent per year, the EUR strategy underperforms, which we argue is a result of the term structure model not being well suited to describe the abnormal shape of the EUR swap curve that manifests over much of the sample period. For both USD and EUR, performance is much better over the first half of the sample (2006-2012) than over the second half (2013-2020), which coincides with a fall in swap rate volatility. Still, risk factor exposure is low for both strategies, though it is higher for USD than for EUR. We conclude that there is potential for risk-adjusted excess returns in yield curve arbitrage, but that the strategy suffers when there are structural changes in the shape and volatility of the term structure.

Contents

1 Introduction

Alfred Winslow Jones coined the term “hedged fund” and created the first hedge fund structure in 1949 (Mallaby, 2010). From then on, the popularity of hedge funds has exploded, and their total assets under management has grown from a few million to over $3.3 trillion by 2020 (Barclays Hedge, 2020). Naturally, as hedge funds’ assets under management grew, new families of investment strategies emerged; one of these was fixed-income arbitrage, which is the main subject of our paper. Fixed income arbitrage strategies try to profit from mispricings in fixed income markets, and are often associated with the demise of hedge fund LTCM in the late-1990. In spite of LTCM’s failure, fixed income arbitrage remains popular to this day with more than $750 billion of capital invested in fixed income arbitrage hedge funds at the end of 2020 (Barclays Hedge, 2020).

As detailed by Lowenstein (2001), LTCM quadrupled investors’ money over a period of four years (1994 to 1998) by, at least mainly, engaging in fixed income arbitrage strategies. The fund’s returns were accompanied by low levels of volatility, resulting in a Sharpe ratio of 4.35 (net of fees) before its demise. However, with a leverage ratio of 25, the fund eventually lost all its capital following the 1998 Russian debt default and the ensuing market turbulence, when virtually all their trades failed. Supposed “risk-free” arbitrage strategies turned out to be less “risk-free” than theorized.

In retrospect, LTCM’s failure demonstrated how fixed income arbitrage strategies that were supposed to carry little risk could lead to painful losses (Lowenstein, 2001). The natural question to ask, then, is whether fixed income arbitrage is truly arbitrage, or if the return from these strategies is primarily a reward for being exposed to systematic risk factors. In a widely read paper, Duarte, Longstaff and Yu (2007) find that while some of the strategies are indeed arbitrage-like, other strategies have significant risk factor exposure, particularly the ones that require little skill to implement. They suggest that strategies requiring “intellectual capital” to implement may generate positive, risk-adjusted excess returns; still, even for these strategies, returns are far from “risk-free”. Still, their results, as they argue, indicate that there is more to fixed income arbitrage than “picking up nickels in front of a steamroller”, i.e. earning a small, positive return most of the time only to suffer large losses in times of market stress (like when selling uncovered index puts).

In this thesis we look at one of the fixed income arbitrage strategies from Duarte et al. in greater detail, namely yield curve arbitrage. Following their methodology, we create return indexes for the 2006-2020 period by implementing the strategy on historical data of USD and EUR par swap rates. Our implementation relies on an affine three-factor term structure model to both identify and hedge “cheap” and “rich” maturities on the yield curve. The strategies are implemented on daily data, and every practical detail of the USD and EUR swap markets are taken into account in order to obtain realistic return series. In addition to looking beyond U.S. markets, our thesis extends the research of Duarte et al. by looking at the strategy over a new sample period, which, among other events, includes the financial crisis of 2007-2009 and the unorthodox monetary policy response that followed; we hypothesize that the ensuing effects, like increased central bank control over the yield curve and lower interest rate volatility, may have impacted its profitability.

Over the 2006-2020 sample, with the initial amount of capital set so that the strategies produce a return volatility of 12% per year, the USD strategy produces an excess return of ca. 5.5% per year, implying a Sharpe ratio of 0.462. The excess return of the EUR strategy is not statistically different from zero, which we argue is a result of the three-factor model not being well suited to describe the peculiar shape of the EUR swap curve that manifests over much of the sample period. For both USD and EUR, performance is markedly better over the first half of the sample (2006-2012) than over the later half (2013-2020), which coincides with a fall in par rates volatility. This suggests that the returns to the strategy are positively correlated with the level of interest rate volatility, which is opposite to that of other arbitrage strategies where volatility is usually unwanted.

Seemingly, adjusting for exposure to risk factors does not materially lower the excess returns of the yield curve strategies. The USD strategy produces a risk-adjusted excess return of over five percent per year, which is significant at the 10% level; for EUR, the risk-adjusted excess return is zero. The R^2 is roughly 10% for EUR and roughly 22% for USD, the latter of which being higher due to negative correlation with the stock market and positive correlation with U.S. Treasuries. The included risk factors control for both equity, credit, interest rate and volatility risk factor exposure.

The findings in this thesis are broadly consistent with the results of Duarte et al., who report that yield curve arbitrage produces positive, risk-adjusted excess returns with a limited degree of risk factor exposure. That being said, our results from EUR swap markets illustrate that the strategy is exposed to *model risk* and to the risk of structural changes in the shape and volatility of the term structure. We believe that this is an important finding, and view it as the key contribution of this thesis. In addition, we believe our implementation of the strategy can serve as inspiration for others.

The rest of this thesis is organized as follows. Section 2 reviews key papers related to the risk and return characteristics of arbitrage. Section 3 introduces the necessary theory, with a focus on swap and term structure modeling. Section 4 introduces the concept of yield curve arbitrage and Section 5 looks at our implementation of the strategy. Section 6 presents results. Finally, Section 7 concludes.

2 Literature Review

2.1 Arbitrage in the Literature

Arbitrage is an age-old concept, but it was not until the emergence of arbitrage pricing in the early 1970s that academics began studying the risk and return of arbitrage strategies. In an academic context, arbitrage refers to a trading strategy that is costless at inception but that guarantees a strictly positive return – a "free lunch", so to speak. This is done by exploiting a relative mispricing between a security and its replicating portfolio; if the cheap one is bought and the expensive one is sold, an arbitrageur can pocket the difference and have no remaining obligations since the cash flows, by construction, net out to zero. In reality, the arbitrage term is used in a wider sense to describe trading strategies that exploit mispricings between *similar* securities. As a result, arbitrage strategies do not typically lead to risk-free profits.

In Shleifer and Vishny (1997), the authors define a more realistic view of arbitrage by looking into professional arbitrage and its implications for security pricing. From a traditional point of view, an arbitrageur obtains a risk-free profit by exploiting mispricings between two similar portfolios and the subsequent correction to fair value, and the strategy is theoretically supposed to carry no risk and to require no capital. However, Shleifer et al. argue that interim losses on such strategies can force arbitrageurs to liquidate their positions at a loss in order to preserve their capital. This implies that arbitrage strategies carry risk and require capital, and that arbitrage opportunities may persist if traders are unwilling to take on the risk.

Like Shleifer et al., Pontiff (2005) argues that arbitrage is indeed risky. Pontiff focuses on the idiosyncratic risk of arbitrage strategies, and makes the case that the idiosyncratic risk of such strategies are unhedgeable. As a result, arbitrageurs must trade off the expected profit from an arbitrage trade and the idiosyncratic risk to which the trade exposes them.

Likewise, Patton (2009) – in which the author looks at "market neutral" hedge funds – finds evidence that "market neutral" strategies are indeed often exposed to some risk factors. Patton points out that such strategies – including arbitrage – are often associated with making trades that are neutral with respect to some key market variable, like the general level of rates or the returns on some stock index. Yet, exposure to residual variables still remain.

Mitchell and Pulvino (2012) investigate a specific example of the scenario outlined in Shleifer et al., namely when debt financing was pulled from arbitrage hedge funds during the financial crisis. Instead of forcing prices of similar securities to converge, arbitrageurs had to liquidate existing positions, thereby causing the level of mispricing to increase; this naturally induced losses. Their findings illustrate that, since arbitrage strategies often require leverage, traders run the risk of having to liquidate their positions if debt financing is pulled.

Moving on to particular strategies, Gatev, Goetzmann and Rouwenhorst (1999) study pairs trading over the 1962 to 1997 period. A form of statistical arbitrage, pairs trading entails forming pairs of stocks that tend to move together; when the spread between the two widens, one buys the “loser” and shorts the winner. If history repeats itself, prices will narrow and the arbitrageur will make money. Their results suggest that pairs trading produces positive excess returns with a low degree of correlation to the S&P 500; nevertheless, Gatev et al. realize that the strategies are trading intensive and that the profitability of the strategies depends upon the cost and impact of execution. The authors extended their sample period in Gatev et al. (2006), where they conclude that the excess return is reward for keeping markets efficient.

Mitchell and Pulvino (2001) study the risk and return characteristics of risk arbitrage, which is a strategy that aims to profit from the spread between a target company’s stock price and the offer price. In their paper, Mitchell and Pulvino analyze 4,750 mergers (1963 to 1998) in order to construct a return index. The authors find that the returns to risk arbitrage are very similar to those obtained from writing uncovered index put options and that, adjusted for risk, the excess return is ca. 4% per year. They postulate that this excess return represents a reward paid to risk arbitrageurs for providing liquidity, particularly during market crashes.

Next, Argarwal, Fung, Loon and Naik (2011) study the risk and return characteristics of convertible bond arbitrage. Most commonly, the strategy involves taking a long position in a convertible bond while delta-hedging the equity risk; the rationale of the strategy is that the convertible bond is sometimes priced ineffciently relative to the issuer’s stock (the embedded equity option is often cheap). Their results suggest that convertible arbitrageurs are rewarded for playing an intermediation role of funding issuers while transferring part of the equity risk of the convertibles to the equity market through their hedging of the equity option.

Seemingly, most studies of arbitrage strategies seem to conclude that arbitrage is risky and that arbitrageurs are rewarded for providing some kind of service, be it intermediation or liquidity provision. Most studies also document negatively skewed return distributions; while these trading strategies tend to make money on average, they occasionally suffer large losses, like option selling. Such losses can be a result of the nature of the strategy, but it may also be a result of having to liquidate positions at a loss if debt financing dries up or if interim losses result in margin calls that the arbitrageur cannot meet without selling his holdings.

2.2 Duarte, Longstaff, and Yu

Duarte, Longstaff and Yu (2007) is arguably the most in-depth study of risk and return in arbitrage to date. In their paper, Duarte et al. construct hypothetical return indexes for five popular fixed income arbitrage strategies over the 1988 to 2004 period, and examine both the return and risk factor exposure of the strategies.[1] Since our paper is largely inspired by theirs, we have devoted a full subsection to discussing the paper's most salient results.

Duarte et al. first consider swap spread arbitrage. In its essence, swap spread arbitrage involves receiving fixed in an interest rate swap (receiving *S*) while shorting a Treasury bond of the same maturity as the swap through repo (thus paying *T*). The swap finances at LIBOR, while the short position earns the repo rate *r*; therefore, if *S* - *T* is greater than LIBOR - *r* over the life of the trade, the arbitrageur should make money. Historically, that has often been true, but the trade is quite risky; if LIBOR rates increase relative to risk-free rates *after* initiation of the trade, the strategy fails. Duarte et al. find that the excess return to swap spread arbitrage is roughly 5% per year over the sample; the risk-adjusted excess return, however, is zero, as the strategy has significant exposure to specific risk factors. The authors thus conclude that there is very little "arbitrage" in swap spread arbitrage, and that the positive excess return simply is a reward for being exposed to financial sector events (by paying LIBOR rates).

Duarte et al. thereafter look at yield curve arbitrage, which, of course, is the subject of our thesis. Yield curve arbitrage is more involved than swap spread arbitrage, and the strategy is presented in detail in Sections 4 and 5. In brief, the strategy involves identifying cheap/rich points along the term structure with the help of term structure models. Like in our thesis, they implement the strategy by trading swaps. Duarte et al. find that yield curve arbitrage produces risk-adjusted excess returns of ca. 4% per year with a low degree of exposure to risk factors – in their view, the positive, risk-adjusted excess return is a result of the strategy requiring skill to implement, and is not simply reward for being exposed to some type of risk.

Later in their paper, Duarte et al. also look at credit, mortgage, and volatility arbitrage, the details of which are beyond the scope of this thesis. In any case, both credit and mortgage arbitrage generate positive alphas, which is not the case for volatility arbitrage.

The authors argue that the strategies with positive risk-adjusted excess return succeed because they demand "intellectual capital"; for instance, yield curve arbitrage requires traders to calibrate a multi-factor term structure model, which is not a trivial exercise. In contrast, the

[1] As mentioned in the introduction, Duarte et al. only look at U.S. markets.

more trivial strategies, like swap spread arbitrage, require less skill; in equilibrium, therefore, perfect competition will drive down their risk-adjusted excess returns to zero. That, however, will not be the case for the strategies requiring “skill”, where arbitrageurs can rely on superior financial know-how, information, or modeling to outperform their competitors.

We shall refer to Duarte et al. as “DLY” for the remainder of this thesis.

3 Theory

3.1 Interest Rate Swaps

3.1.1 Par Swaps

An interest rate swap (IRS) is a derivative contract through which two parties agree to exchange interest payments calculated at a fixed rate for interest payments calculated at a rate that changes over time, typically at a short-term rate such as LIBOR. Being an OTC contract, interest rate swaps may be customized to the parties' needs, but the most liquid swaps are the so-called par swaps. In a par swap, the fixed rate is set so that the contract is fair at initiation; in other words, given the expected evolution of the floating rate, paying a fixed rate $C(T)$ for T years has the same present value as paying the floating rate, at some periodicity, over those T years (for example, paying 3-month USD-LIBOR every quarter). Consequently, par swap rate $C(T)$ may be interpreted as the expected, weighted-average floating rate over the swap's term, where weights are given by the discount factors to the floating rates' dates.

Given that the floating rate of the IRS is taken to be the risk-free rate, it can be shown that the present value of the floating leg equals par, or the swap's notional (see Chapter 16 in Tuckman et al. for a justification). As a result, $C(T)$ is chosen so that the present value of the fixed leg also equals par. If the fixed leg accrues semiannually, which is typically the case for USD swaps,

(3.1) $$C(T) = \frac{1 - D(T)}{\sum_{i=1}^{2T} D(\frac{i}{2})}$$

where $D(i)$ is the discount factor to time $t = i$ (here, discount factors are calculated from the projected risk-free rates). If the fixed leg accrues once per year, which is typically the case for EUR swaps,

(3.2) $$C(T) = \frac{1 - D(T)}{\sum_{i=1}^{T} D(i)}.$$

Hence, if one has a function for the discount factor to an arbitrary time, one also has a function for the T-year par swap rate. This result will be used later in the thesis.

3.1.2 Non-Par Swaps

If an investor enters into a 10-year swap at the 10-year par swap rate, the contract has, per the results of the previous subsection, a net present value (NPV) of zero at initiation – the par swap rate is set such that this is the case. However, as time passes, the swap contract is no longer a 10-year swap, and the term structure of par swap rates has likely changed; as a result the swap contract will have a nonzero NPV. For instance, the right to receive fixed at 2.875% on some notional for 10 years will be valuable if the 10-year par swap rate declines to 1.50%. The exact value is found by discounting the contractually fixed payments at a discount curve constructed from the current term structure of par swap rates (how such a curve is constructed is explained in the Appendix); the present value of the fixed leg is then compared to the value of the floating leg to obtain the swap's NPV. In notation,

(3.3) $$NPV = \omega[V_0^{FIXED} - V_0^{FLOAT}]$$

where ω is -1 when paying fixed and +1 when receiving fixed and where

$$V_0^{FIXED} = \sum_{i=1}^{n}[N \times \bar{C} \times \delta(t_{i-1}, t_i) \times D(t_i)] + [N \times D(t_n)]$$

and

$$V_0^{FLOAT} = \{N + [N \times \bar{L} \times \delta(t_0, t_k)]\} \times D(t_k)$$

where N is the notional of the swap, $\bar{C}$ the swap's fixed leg rate, n the number of remaining fixed leg payments, $\delta(t_{i-1}, t_i)$ the time (years) between two fixed leg payments at t_{i-1} and t_i (at the appropriate day count convention), $D(t_i)$ the discount factor to time t_i, $\bar{L}$ the previous floating rate fixing, t_k the time to the next floating rate fixing, and t_0 is "today".

In the previous subsection it was argued that the value of the floating leg equals par at initiation; this is also the case at each reset date. However, between two reset dates, the value of the floating leg can be different from par; now, its value equals par plus the predetermined floating rate payment for the next payment date. Therefore its value will be *higher* than par if $\bar{L}$ is *higher* than the spot rate implied by the discount factor to t_k, and vice versa.

Note that the formulas in (3.3) assume that the swap's notional is exchanged when the swap matures. This is a common convention when calculating the NPV of swaps.

3.2 Term Structure Modeling

3.2.1 PCA of the Yield Curve

Earlier studies have demonstrated that there are predominantly three factors that drive movements in the term structure. Litterman and Scheinkman (1991) estimate common factors through principal component analysis (PCA) of spot rates, derived from U.S. Treasury yields, over the 1984-1988 period; they find that three factors explain 96% of the variance of yields. Through the mechanics of PCA (which is beyond the scope of this thesis) these "factors" are constructed in such a way that they are independent and thereby uncorrelated. Litterman et al. call these factors *level*, *steepness*, and *curvature*; the level factor causes a parallel shift of the spot rate curve, while the steepness factor lowers spot rates up to five years and – at the same time – raises spot rates of longer maturities. Finally, the curvature factor increases the overall curvature of the spot rate curve up to 20 years, which is associated with interest rate volatility. Litterman et al. also find that the first factor, the level factor, is by far the most important; the level factor explains 89.5% of total variance, compared to 8.5% for the steepness factor.

Later studies have confirmed the results of Litterman et al. (see, for example, Baygün, Showers and Cherpelis (2000)). Instead of presenting the results from such studies, which are by now quite old, we present our own results for USD and EUR par swap rates over the 2006 to 2020 period. Here, PCA is applied to weekly changes in par swap rates for maturities from one to 30 years. Results are shown in *Figure 3.1* and *3.2* below; the factor loading shows how a one-standard-deviation change in the factor impacts par swap rates of different terms. These factors are quite similar to the ones shown in Litterman et al. and similar studies, although the shape of the curvature factor is somewhat different. The fraction of total variance explained – which is listed in parentheses – is also quite similar. Quite remarkably, the factors are close to identical for USD and EUR which, to some extent, illustrates the factors' pervasiveness.

Some interpretation of these three factors is in order. The level factor – which causes a near parallel shift – is driven by economic news that are long-lived in nature; a natural way to think about such a shift is that long-run inflation expectations or the long-run real interest rate increases which, in turn, raises swap rates across the curve. In contrast, the steepness factor – referred to as *slope* in the figures – is driven by news about monetary policy; if, for example, the Federal Reserve hints at lower policy rates ahead, short term swap rates will fall and long term swap rates will, at least in many cases, rise as investors anticipate higher inflation ahead on the back of low policy rates. The third factor is associated with interest rate volatility.

Figure 3.1

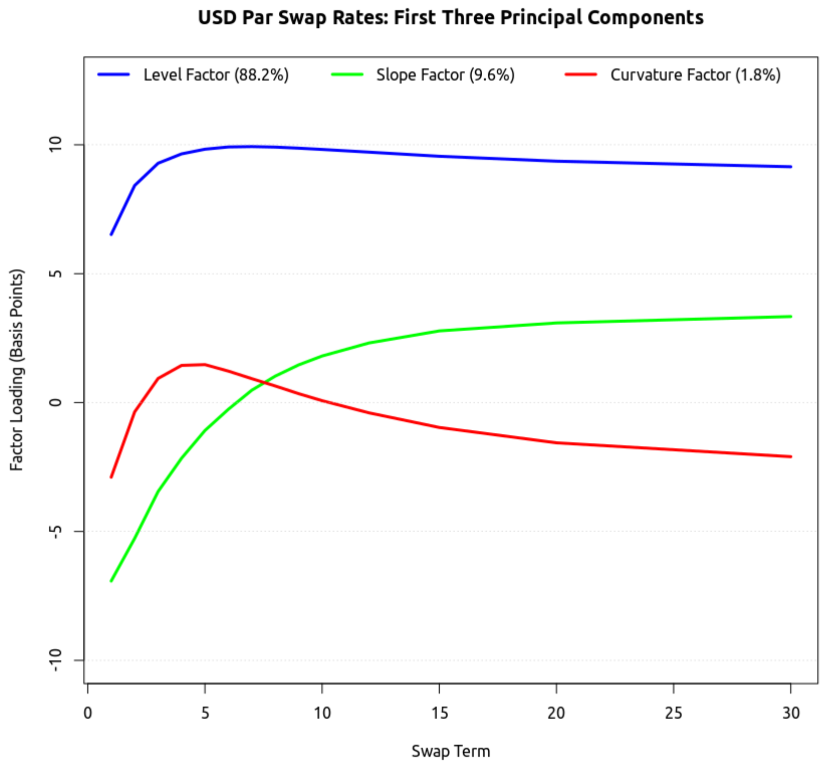

This figure shows the first three principal components of changes in USD swap rates over the 2006-2020 period.

Figure 3.2

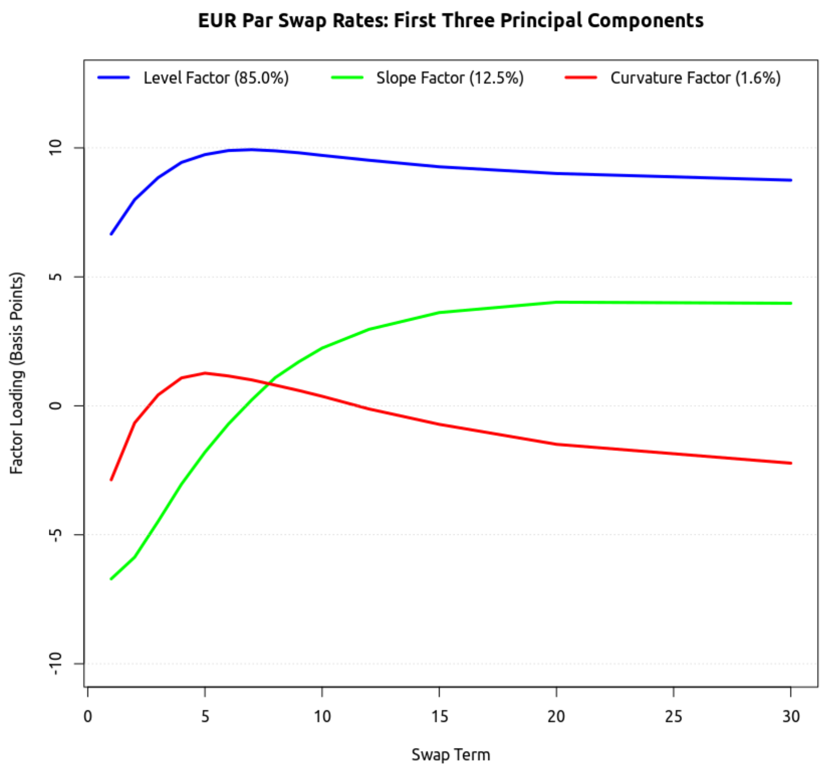

This figure shows the first three principal components of changes in EUR swap rates over the 2006-2020 period.

3.2.2 Short Rate Models

In its essence, a short rate model specifies a stochastic process for the short-term rate, typically denoted r_t. The dynamics of r_t is usually specified as an Itô process of the form

$$dr_t = \mu(t, r_t)dt + \sigma(t, r_t)dW_t \tag{3.4}$$

where dW_t is a standard Wiener process.[1] There are many different specifications of $\mu(t, r_t)$ and $\sigma(t, r_t)$ in the short rate literature, and one popular specification will be introduced in the next subsection. In some cases, these drift and volatility functions are time-dependent so as to allow for a close fit to both the term structure of interest rates and interest rate volatility; still, in other cases, the functions are constants. In either case, the short-term rate r_t is assumed to evolve from its starting value r_0 via dr_t so that a change in the short-term rate is given by the sum of a non-random drift term and a random perturbation proportional to $\sigma(t, r_t)$.

In some cases, short rate models are used to price interest rate options, which involves numerical techniques such as binomial trees and Monte Carlo simulation.[2] In other cases, the interest is not so much on the short-term rate itself, but rather on the term structure of interest rates *implied* by the model. This involves solving

$$\mathbb{E}_0[e^{-\int_0^T r(u)du}] \tag{3.5}$$

to obtain an analytical expression for the discount factor to time T (see Brigo and Mercurio). Since par swap rates can be written as a function of the discount factors over the swap's term, one may also obtain an analytical expression for par swap rates. The model-implied par swap rates are then functions of the short rate model's parameters (and time T). In other words, in the cases where (3.5) has a solution, a user can specify the stochastic process for r_t, solve for par swap rates, and compare model-implied swap rates to swap rates observed in the markets. Of course, for the model to have predictive value, the process for r_t has to be meaningful and sufficiently complex to reflect real-world interest rate dynamics. Taken the other way around, a user may also start with a set of market swap rates and parameterize the short rate model so as to minimize deviations between market-observed and model-implied swap rates.

---

[1] As this thesis is not about stochastic calculus we will refer the reader to Wiersema (2008) for an introduction to Itô calculus and Wiener processes.
[2] Brigo and Mercurio (2006) is a comprehensive reference for short rate models in the context of options pricing.

3.2.3 The Vasicek Model

Short rate modeling began with a seminal paper by Oldrich Vasicek (1977), where the author also proposed a particular specification of equation (3.4) which has come to be known as the Vasicek model. In the Vasicek model, the dynamics of r_t is specified as

(3.6) $$dr_t = \kappa(\theta - r_t)dt + \sigma dW_t$$

where both κ, θ, and σ are constants. The model is a so-called mean-reverting model: when r_t is different from θ, r_t tends toward θ at a speed determined by κ. Hence θ represents the long-run interest rate to which r_t converges from its initial value. These dynamics reflect the empirical observation that interest rates tend to exhibit mean-reversion, which is quite natural when, as in many economies, central banks have an inflation target; if, for example, the target is 2% – and if real GDP growth is, say, 2% – the interest rate should be around 4% over time. Nonetheless, the expected short-term rate can only move monotonically downward or upward in the model, which means that it will not be able to describe situations in which investors see lower short-term rates in the near term but eventually higher short-term rates in the long term. As a consequence, the Vasicek model only allows for simple downward- and upward-sloping term structures; an "inverted" yield curve, for instance, is not possible in this model.

The solution to the stochastic differential equation in (3.6) is

(3.7) $$r_t = r_0 e^{-\kappa t} + \theta(1 - e^{-\kappa t}) + \sigma e^{-\kappa t} \int_0^t e^{\kappa s} dW_s$$

where the integral is an Itô integral.

To solve for discount factors, one substitutes (3.7) into (3.5). The solution – expressed here as a spot rate instead of a discount factor – is given by equation

(3.8) $$\hat{r}(T) = \theta + \frac{1 - e^{-\kappa T}}{\kappa T}(r_0 - \theta) - \frac{\sigma^2}{2\kappa^2}(1 + \frac{1 - e^{-2\kappa T}}{2\kappa T} - 2\frac{1 - e^{-\kappa T}}{\kappa T})$$

where $\hat{r}(T)$ is the T-year spot rate.[3]

To solve for model-implied par swap rates, one uses equation 3.8 iteratively.

[3] Discount factors may then be obtained via the relationship: *DF*(T) = exp[-*r*(T) x T], where *r*(T) is the spot rate.

While the Vasicek model is elegant and intuitive, it is too simplistic for many real-life applications. Firstly, the model has only one factor, namely the short-term rate r_t; from (3.8) it is evident that the initial value of the short-term rate, r_0, is the only variable affecting rates of different terms (the other parameters are constants). In other words, the only way spot rates of different terms can change is through a change in r_0, implying that the interest rate risk of 30-year bond can be hedged with a 1-year bond (by neutralizing the exposure to r_0). Second, the model is not flexible enough to fit most term structures, as can be seen from *Figure 3.3* – in the figure, the value of r_0 is chosen so that the model-implied 10-year swap rate equals the market rate. The model is clearly not flexible enough to fit this particular term structure shape when the constant parameters are set "reasonably" (θ = 4.0%, κ = 0.10, σ = 1.0%).

Figure 3.3

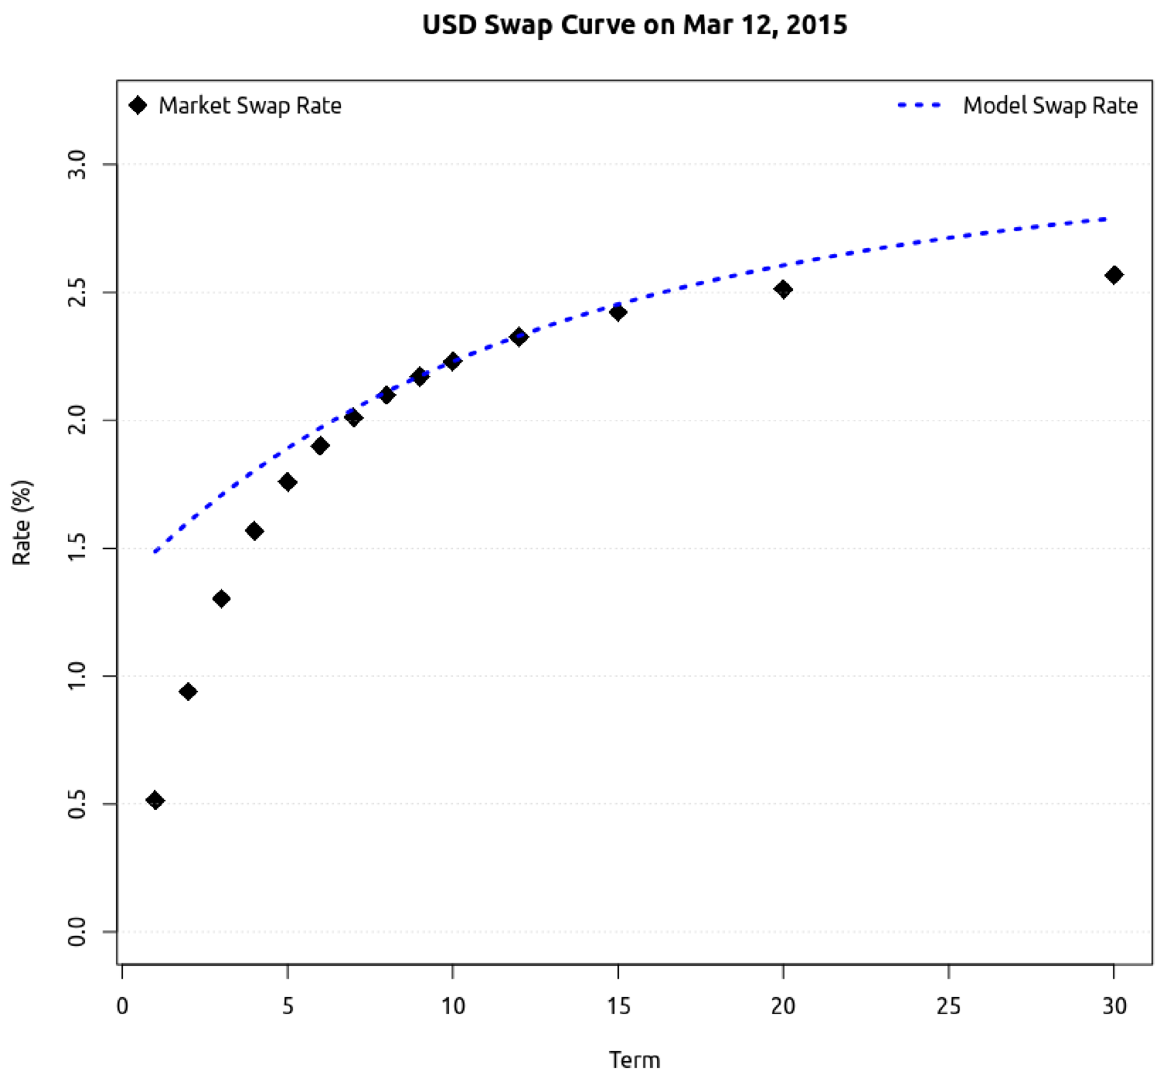

This figure shows the Vasicek model fitted to match the 10-year USD swap rate on 03/12/2015.

3.2.4 Three-Factor Model with Coupled SDEs

The three-factor model from Tuckman and Serrat (2011) is described by the following system of coupled stochastic differential equations (SDEs)

$$
\begin{aligned}
dr_t &= -\alpha_r(r_t - M_t)dt \\
dM_t &= -\alpha_M(M_t - L_t)dt + \sigma_M dW_t^1 \\
dL_t &= -\alpha_L(L_t - \theta)dt + \sigma_L dW_t^2
\end{aligned}
\tag{3.9}
$$

where dW_t is a standard Wiener process and $\mathbb{E}[dW_t^1 dW_t^2] = \rho dt$. In (3.9), short-term rate r_t mean-reverts to M_t at a speed determined by α_r; M_t, on the other hand, mean-reverts to L_t at a speed determined by α_M, while L_t mean-reverts to θ (which is a constant) at a speed set by α_L. However, both M_t and L_t fluctuate about their expected paths with volatility σ_M and σ_L, respectively; r_t itself has no random term, but fluctuates due to its direct tracking of M_t and L_t. In the model, M_t is meant to represent a medium-term interest rate factor while L_t is meant to represent a long-term factor. M_t and L_t are correlated with correlation ρ.

Though complex at first sight, the above system of coupled SDEs has a very intuitive interpretation. Like explained in Tuckman and Serrat, L_t is meant to reflect long-term trends in demographics, productivity, or technology, or other factors that influence the long-run real interest rate; since these things are, by nature, not very volatile, σ_L should be low. Moving on to the next equation, M_t is meant to reflect monetary cycles around that long-term trend, and will hence fluctuate around the value of L_t; since monetary cycles are, by nature, short-lived and volatile, both α_M and σ_M ought to be relatively high. Finally, the process for r_t is meant to reflect the behaviour of a central bank that pegs the short rate at a level consistent with the state of the monetary cycle which, then, implies that α_r should be relatively high. Notice that the process for r_t has no stochastic term, which is consistent with its interpretation.

With the short-rate process from equation 3.9, spot rates in the three-factor model are given by the formula in Appendix 8.1. This formula is a function of θ, α_r, α_M, α_L, σ_M, σ_L, and ρ, as well as the initial values of the model's three factors. Next, since par swap rates can be written as a function of the spot rates over the swap's term, one may derive model-implied par swap rates. These model-implied swap rates can then be compared with market-observed par swap rates. This is typically done by calibrating the initial values of the three factors such that *three* swap rates implied by the model match *three* swap rates in the market; in that way, model rates are broadly consistent with the current term structure of par swap rates.

With the short-rate process specified as above, the model is able to match a variety of term structure shapes only by changing the initial value of the three factors (and *not* the fixed parameters). One example is given in *Figure 3.4* below, where the factors are calibrated to the 1-, 10-, and 30-year USD swap rates on Feb. 24, 2020. The model has no difficulty producing an "inverted" yield curve; in this example, the initial value of r_t is higher than M_t, while M_t is lower than L_t – thus, r_t first falls toward M_t but then increases as M_t increases.

Figure 3.4

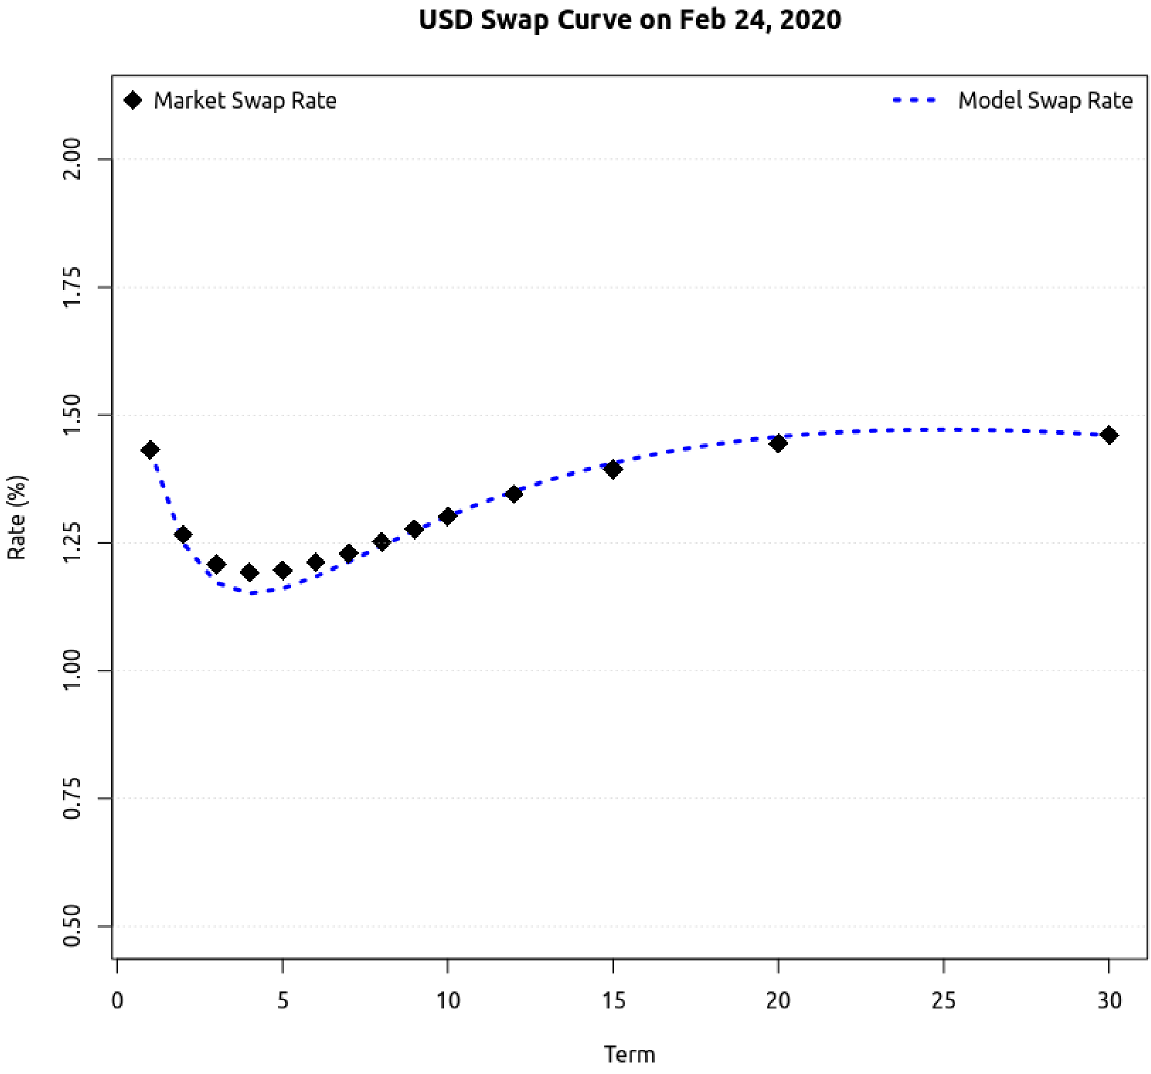

This figure shows the three-factor model fitted to match the 1-, 10-, and 30-year USD swap rates on 02/24/2020.

With the spot rate specification in Section 8.1, the term structure is a function of three affine factors – not the original r, M, and L. These three factors, $\bar{x}_1$, $\bar{x}_2$, and $\bar{x}_3$, are a result of a "reduced form"-transformation of the model. Even though they no longer have the exact same interpretation as r, M, and L, we will refer to them as such so as to not overwhelm the reader with mathematical detail. In any case, the spot rate function (and, as a result, the swap rate function) can be differentiated with respect to the initial values of the three factors to see how a change in a factor's initial value impacts swap rates of various terms; this is illustrated in *Figure 3.5* below, which is based on a particular parameterization of the three-factor model (we will return to the subject of parameterization). Here, the "long-term"-factor creates a near parallel shift, while the other two factors create a flattening/steepening of the curve. Looking back to Section 3.2.1, the shift caused by the long-term factor is reminiscent of the *level* shift, while the shift caused by the medium-term factor is reminiscent of the *slope* shift (if rotated). Thus, hedging these two rate factors is reminiscent of hedging to the first two PCs.

Figure 3.5

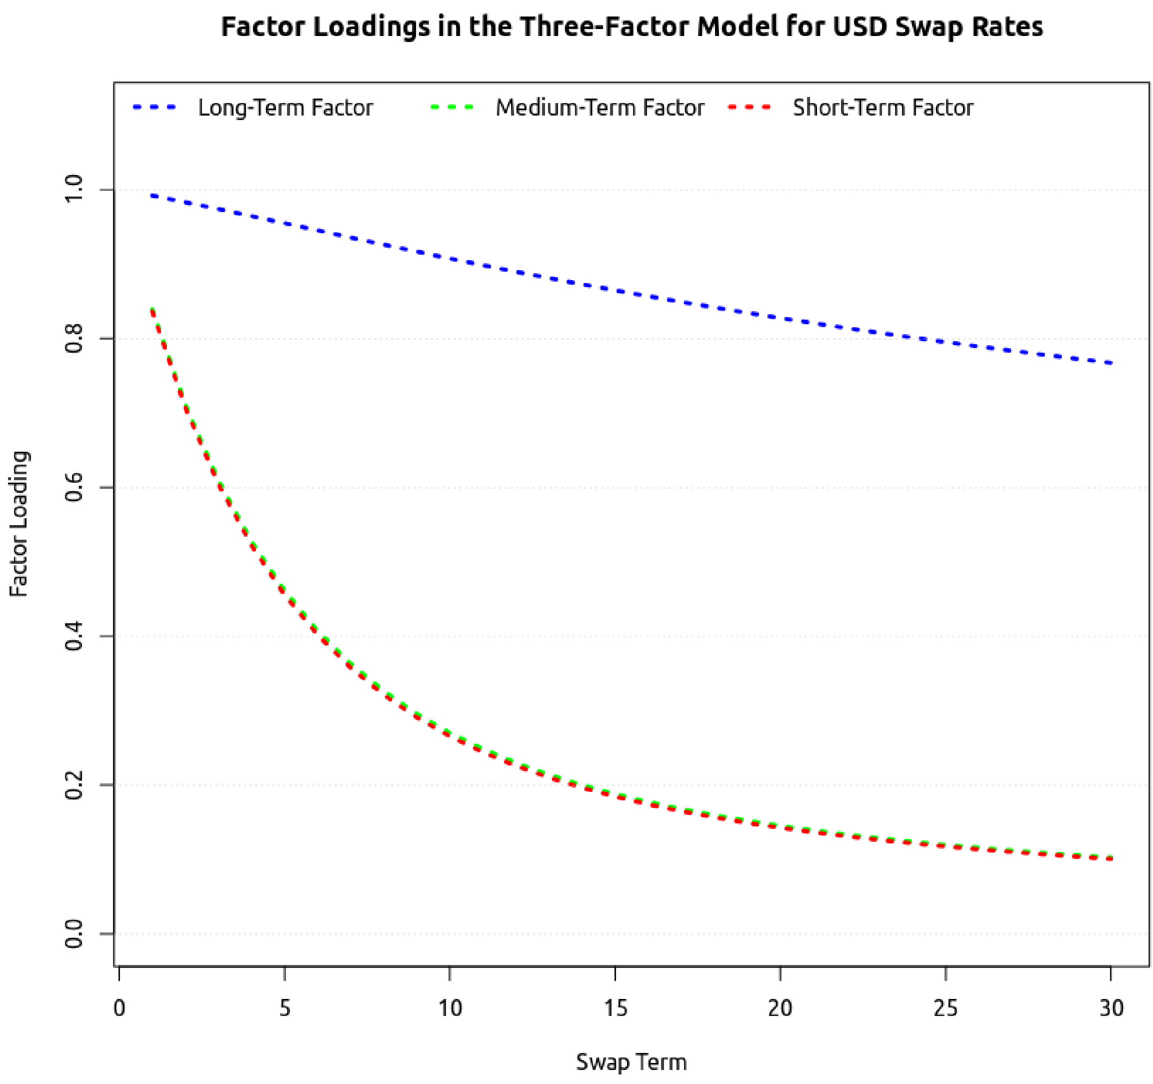

This figure shows the derivatives of swap rates with respect to factor values in the three-factor model (for USD).

4 Yield Curve Arbitrage

There is no general definition of yield curve arbitrage. Rather, the term encompasses a family of strategies in which investors take a view on the relative value of a bond (or a swap) *versus* other bonds (or other swaps). Such strategies may involve the identification of "cheap" and "rich" maturities along the yield curve, be it an on-the-run Treasury bond curve or a swap curve; it may also involve identifying "cheap" and "rich" bonds that are very similar but that differ in price for a variety of reasons, for instance on-/off-the-run features or "specialness" in the repo markets. In any case, the investor would go long the cheap security and short the rich security in such a way that the portfolio is not exposed to general interest rate movements but only to the convergence (or divergence) of the cheap to the rich security. Here, the "investor" would typically be a dealer or market maker who follows the market closely and who, for this reason, will be in a favorable position to identify, and profit from, mispricings; the "investor" could also be a fixed income relative value hedge fund or a proprietary trading desk within an investment bank. In this thesis we will concentrate on the first version of the strategy, namely the identification of cheap and rich points along the term structure; in addition, the thesis will focus exclusively on interest rate swaps since the market has become the *de facto* interest rate benchmark in many currencies over the past couple of decades (PIMCO, 2020).

The identification of cheap and rich points on the swap curve can be based on *a priori* views about the path of future interest rates, or it can be based on a term structure model like one of the ones from the previous section. Since the first approach is inherently too subjective for a quantitative analysis, the thesis will focus on the latter approach. Normally, this involves calibrating a model to match a set of liquid market rates in order to evaluate the cheapness or richness of less liquid market rates, and usually proceeds in the following steps:

1. The initial values of the factors are calibrated in such a way that swap rates implied by the model match the most liquid swap rates in the market, typically the 1-, 2-, 5-, 10-, or 30-year par swap rates.
2. With the calibrated model, swap rates of other maturities are classified as either cheap (i.e. market rates above model rates) or rich (i.e. market rates below model rates).
3. When the cheapness (richness) exceeds a set threshold, a portfolio is formed by going long (short) the relevant swap rate; here, going long implies receiving fixed and going short means paying fixed. Then, offsetting positions in the most liquid swaps are used to hedge interest rate risk.

Looking at *Figure 4.1* – where a particular parameterization of the three-factor model has been fitted to match the 1-, 10-, and 30-year USD swap rates on Mar. 12, 2015 – the 2- to 9-year segment would be classified as cheap while 12-, 15-, and 20-year swap rates would be classified as rich. To take an example of a trade, an investor could decide to receive fixed in a 4-year swap and zero out the resulting factor risk by paying fixed on a portfolio of 1- and 10-year swaps; the trade should prove to be profitable if the 4-year swap rate declines relative to the 1- and 10-year swap rates, i.e. if it converges to its "fair value" in the model.

Figure 4.1

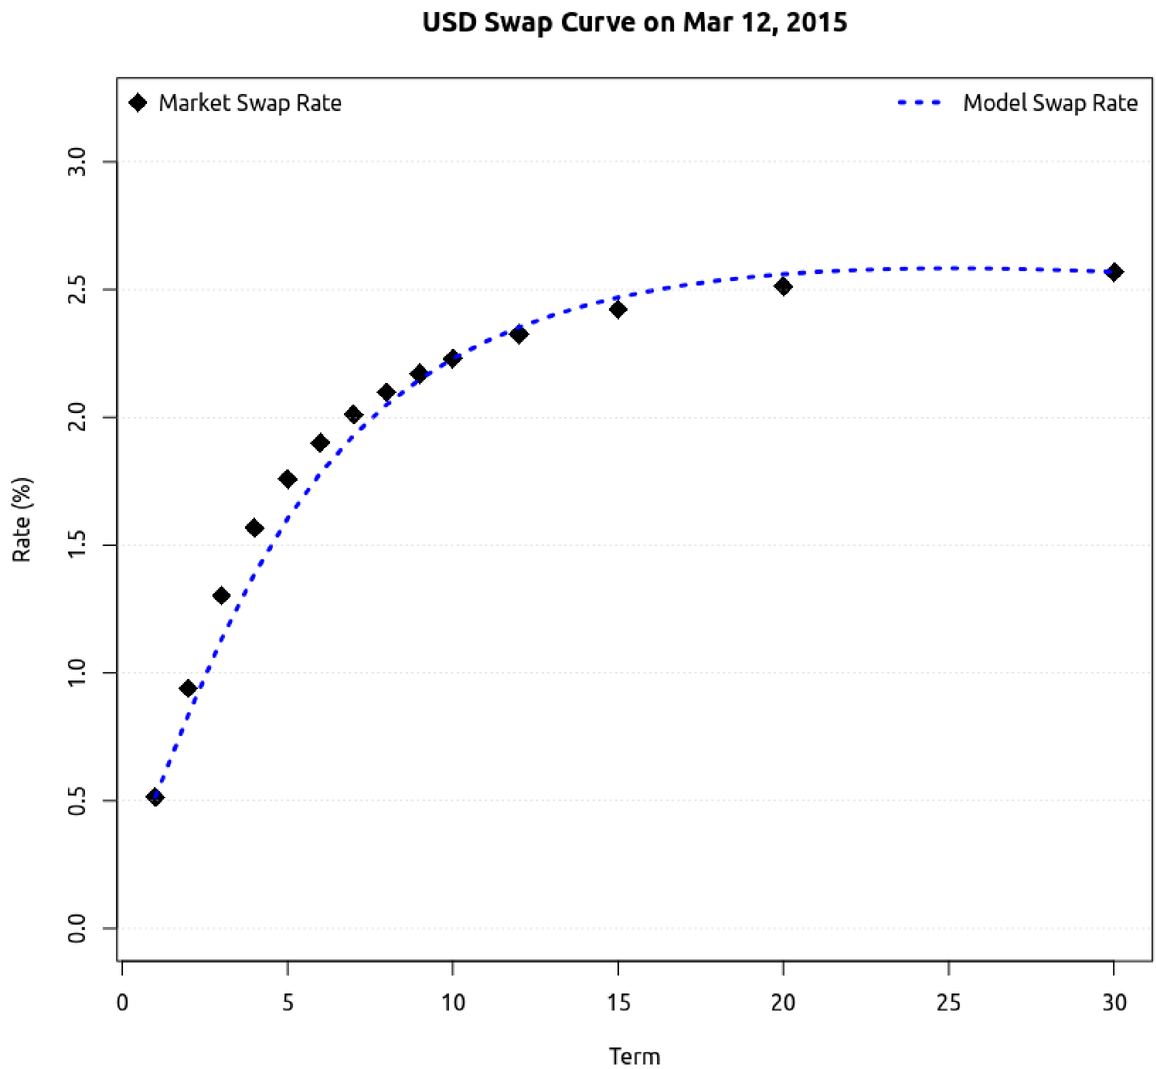

This figure shows the three-factor model fitted to match the 1-, 10-, and 30-year USD swap rates on 03/12/2015.

Positions are held until the relevant swap rate converges to the model's representation of fair value or, if it does not converge, until a certain amount of time has passed. Because the exposure to the interest rate factors in the model is hedged out, the positions are only exposed to the convergence (or divergence) of the cheap or rich swap rates towards “fair rates”; in this way, parallel shifts or a steepening of the curve should not lead to profit nor loss. Still, hedges are only locally effective since they are calculated from derivatives of swap rates with respect to factor values, so large changes in factor values could definitely lead to issues.

Some of the steps outlined above deserves further explanation. By calibrating a model so as to match the most liquid swap rates in the market, one implicitly makes the assumption that these swap rates are fair; in other words, they are neither cheap nor rich in the context of the model. The justification for this assumption is that these swap rates are broadly analyzed, monitored, and traded, and should therefore not deviate far from “fair value”. In this way, the model incorporates the most reliable market information but at the same time allows for other swap rates to be cheap or rich; an alternative approach would be to calibrate the factors so as to minimize deviations between market and model rates over the whole curve, but this would presume that all observed swap rates are equally fair. Since some swap rates are more heavily watched than others, the first method is, in our view, easier to justify. Besides, by matching a set of market rates exactly, an investor can use those swaps to hedge out residual interest rate risk without having to worry about the cheapness or richness of the hedging instruments: they are, from construction, neither cheap nor rich in the context of the fitted model.[4]

The constant parameters of the model, including mean-reversion/volatility parameters, are typically estimated based on historical data or implied from the prices of traded volatility products, like caplets or swaptions.[5] To the extent that the model and its parameters reflect the real-world dynamics of the interest rate process, it should be able to identify which segments are cheap or rich, and hedges should work well. However, if the model’s assumptions are not a suitable description of reality, or if the estimated parameters are inappropriate, its predictive power will be limited; particular segments of the swap curve will always appear cheap or rich and hedges will not work that well. Moreover, even though the model is a suitable description of reality, it may be the case that the estimated parameters are “outdated”; central bank policy may have forever changed the mean-reversion speed of the posited interest rate factors or the level of interest rate volatility, to take but one example. Hence, it is critical to employ a model that is flexible enough to capture real-world term structure movements; nonetheless, it cannot be too flexible because more flexibility invariably leads to too many parameters.

[4] For instance, if the model is fitted to the 1-, 10- and 30-year swap rates those rates are always fair in the model.
[5] One estimation methodology in presented in Section 5.2.

While there are few papers in the finance literature that focus on these types of trading strategies, a review of earlier findings is warranted. The most related analysis is DLY (2004), which was discussed earlier in the paper. In DLY, they implement a yield curve strategy along the lines outlined above using the two-factor Vasicek model and USD swap market data from 1988 to 2004; they find that the strategy produces positive risk-adjusted excess returns with a "low degree of exposure to risk factors". In Fabozzi, Martellini and Priaulet (2005) they use a statistical model, based on a set of economic variables, to forecast changes in the slope of the term structure; trading strategies, such as butterfly swaps, are then implemented based on the model's prediction. Fabozzi et al. show that such strategies perform well, which suggests that the slope of the term structure may be predictable. To our knowledge, these are the only well-known papers that examine the performance of quantitative yield curve strategies, despite the pervasiveness of such strategies amongst relative value hedge funds and traders.

5 Methodology

5.1 Return Series Construction

In order to analyze the risk and return of yield curve arbitrage, we first implement the strategy on historical data. For this, we follow the methodology outlined in DLY, but improve upon their implementation where we feel improvement is necessary. Here, the short-term rate is assumed to evolve according to the three-factor model of Section 3.2.4, where formulas for spot rates (and hence swap rates) are given in Section 8.1. With end-of-day bid and ask dealer quotes (at a daily frequency) from Jan. 2, 2006 to Sep. 18, 2020 for USD and EUR par swaps, we calibrate the initial value of the three factors to match the mid-market 1-, 10-, and 30-year swap rates in each currency.[6] Then, for each currency, we classify the remaining swap rates – in our case the 2-, 3-, 4-, 5-, 6-, 7-, 8-, 9-, 12-, 15-, and 20-year rates – as either cheap or rich. If the cheapness (richness) exceeds a certain threshold, we receive (pay) fixed in the relevant

[6] The swap quotes are from Bloomberg. The fixed leg of USD interest rate swaps accrues on a 30/360 basis with semi-annual payments while the floating leg accrues on an ACT/360 basis with quarterly payments based on the value of 3-month USD-LIBOR. The fixed leg of EUR interest rate swaps accrues on a 30/360 basis with a single payment every year while the floating leg accrues on an ACT/360 basis with semi-annual payments based on the value of 6-month EURIBOR. In both the USD and EUR markets, the floating rate is set two business days *prior* to the start of the accrual period and is paid in arrears. Moreover, both markets follow the "Modified Following" convention, where a payment that falls on a non-business day is pushed to the next business day, unless that day is in a different month (in which case it is brought forward to the first business day *preceding* the payment date).

swap.[7] Maturities from 2 to 9 years are hedged with a portfolio of 1- and 10-year swaps while 12-, 15-, and 20-year swaps are hedged with a portfolio of 10- and 30-year swaps; the hedges are given by the derivatives of swap values with respect to the long- and medium-term factors of the three-factor model. Positions are then held until the relevant swap rate converges to the model's fair value or until a full calendar year has passed. We do not allow for more than one position in the same swap rate at the same time; in other words, if we have taken a position in a 5-year swap, we are not allowed to take a new position in the 5-year swap until the existing position is closed/unwound.[8] Thus, the return index is constructed from the point of view of a single hedge fund who trades the USD and EUR swap markets on a daily basis.

In the implementations we are careful to take into account every practical detail of the USD and EUR swap markets, including day-count conventions, holidays, non-business days, payment frequencies, and spot settlement.[9] In computing floating leg payments we are careful to follow the precise fixing conventions of the markets; daily fixes for 3-month USD-LIBOR and 6-month EURIBOR are obtained from Bloomberg and ICE. Swap positions are valued on a discount curve constructed according to industry standards (see Section 8.2 for more details on the methodology). Hence, since positions are initiated and valued on a daily basis, the PnL should closely resemble that of a hedge fund pursuing the same exact strategy (with the same inputs, models, etc.). In our view, this is an improvement over DLY's implementation, where they ignore most of the above details and where they use monthly observations.

Like in DLY, the hedges are designed so as to neutralize the exposure to the factors of the model. However, we only hedge the exposure to the medium-and long-term factors of the three-factor model since, as seen earlier in *Figure 3.5*, hedge ratios from the short-term factor are typically very similar to hedge ratios from the medium-term factor (since α_r and α_M are often close). Mathematically, we solve the below linear system for N_1 and N_2:

$$\begin{bmatrix} \frac{\partial \hat{V}(t_1)}{\partial M} & \frac{\partial \hat{V}(t_2)}{\partial M} \\ \frac{\partial \hat{V}(t_1)}{\partial L} & \frac{\partial \hat{V}(t_2)}{\partial L} \end{bmatrix} \begin{bmatrix} N_1 \\ N_2 \end{bmatrix} = \begin{bmatrix} N_i \frac{\partial \hat{V}(t_i)}{\partial M} \\ N_i \frac{\partial \hat{V}(t_i)}{\partial L} \end{bmatrix},$$

where $\hat{V}(t)$ is the model swap value for term t, M is the medium-term factor, L is the long-

[7] The threshold is 10 basis points for 2-, 3-, 4-, 5-, 6- and 7-year swaps, and 5 basis points for 8-, 9-, 12-, 15- and 20-year swaps. The threshold is lower for the latter group because the deviation between market and model rates is typically much lower for that group, than for the first group.
[8] This only applies to the same swap rate in the same currency.
[9] USD swaps follow the NYC and London holiday calendars. EUR swaps follow the TARGET holiday calendar.

term factor, t_1 is 1 (or 10 years), t_2 is 10 (or 30 years), t_i is the term of the cheap/rich swaps, N_1 is the notional of the t_1-year swap, N_2 is the notional of the t_2-year swap, while N_i is the notional of the t_i-year swap (which we fix to be 100). Of course, the solution is

$$\begin{bmatrix} N_1 \\ N_2 \end{bmatrix} = \begin{bmatrix} \frac{\partial \hat{V}(t_1)}{\partial M} & \frac{\partial \hat{V}(t_2)}{\partial M} \\ \frac{\partial \hat{V}(t_1)}{\partial L} & \frac{\partial \hat{V}(t_2)}{\partial L} \end{bmatrix}^{-1} \begin{bmatrix} N_i \frac{\partial \hat{V}(t_i)}{\partial M} \\ N_i \frac{\partial \hat{V}(t_i)}{\partial L} \end{bmatrix}.$$

Moving on, we first calculate the daily PnL of the strategy; from this daily PnL series we calculate end-of-month PnL and then monthly returns. In both the USD and EUR market, the initial amount of capital is set so that the the strategies produce an annual return volatility of 12% (3.464% per month). In addition, an equal-weight, monthly return series is formed by investing 50% in each market; this return series is scaled to have a volatility of 12% per year by scaling the individual strategies by *x*% (where *x* is greater than 100 since the two strategies are not perfectly correlated). Hence, even though the strategy is implemented on a daily basis we report statistics for monthly returns as that is the normal reporting frequency.

Since we use bid and ask dealer quotes, transaction costs should be realistic. In reality, however, all investors might not be able to transact at those quotes but they should be realistic from the perspective of a large market player; besides, bid-ask spreads for USD and EUR par swaps are very tight, as seen from *Figure 5.1*. For USD swaps, spreads are usually less than a single basis point; for EUR swaps, spreads are usually between 1-4 basis points.

Figure 5.1

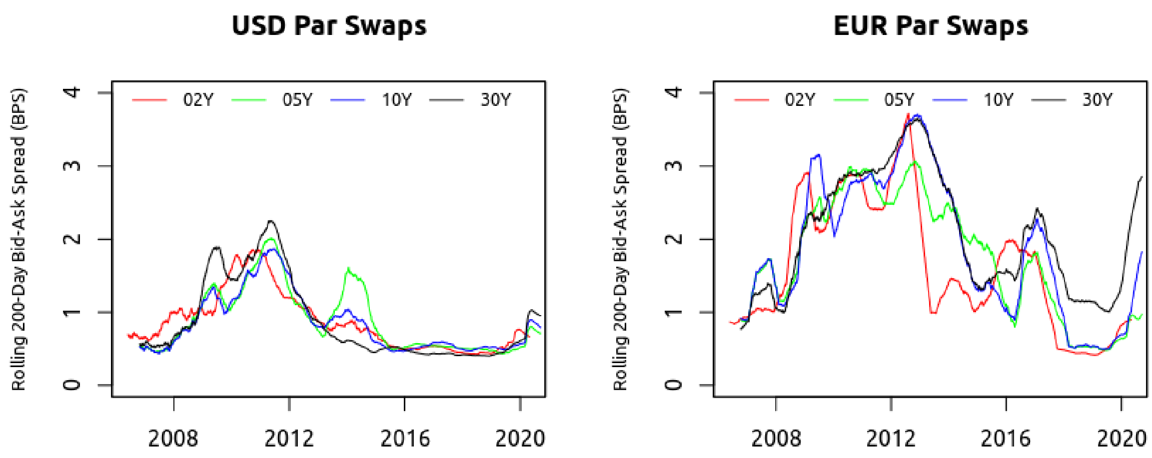

This figure shows the 200-day rolling average bid-ask spread for 2-, 5-, 10-, and 30-year EUR and USD swaps.

5.2 Parameter Calibration

So far, the constant parameters of the three-factor model – θ, α_r, α_M, α_L, σ_M, σ_L, and ρ – have been taken as given. In reality, however, these parameters need to be estimated. For this, we adopt the estimation procedure suggested in DLY (2004), which proceeds in four steps:

1. Choose a trial value of the parameters.
2. For every swap curve in the sample period:
   - 2.1. Calibrate the initial value of the three factors to match the 1-, 10-, and 30-year swap rates. Mathematically,

$$\min_{\bar{x}_1,\bar{x}_2,\bar{x}_3} \sum_t \frac{[\hat{C}(\bar{x}_1,\bar{x}_2,\bar{x}_3,\theta,\alpha_r,\alpha_M,\alpha_L,\sigma_M,\sigma_L,\rho,t) - C(t)]^2}{C(t)}$$

     where $\hat{C}(\ldots,t)$ is the estimated swap rate for term t, $C(t)$ is the market swap rate for term t, and t is in {1, 10, 30}.

   - 2.2. Estimate swap rates for the remaining terms (with $\bar{x}_1$, $\bar{x}_2$, and $\bar{x}_3$ from above).
   - 2.3. Sum the square percentage error between the estimated swap rates and market swap rates. Mathematically,

$$\sum_t \frac{[\hat{C}(\bar{x}_1,\bar{x}_2,\bar{x}_3,\theta,\alpha_r,\alpha_M,\alpha_L,\sigma_M,\sigma_L,\rho,t) - C(t)]^2}{C(t)}$$

     where $\hat{C}(\ldots,t)$ is the estimated swap rate for term t, $C(t)$ is the market swap rate for term t, and t is in {2, 3, 4, 5, 6, 7, 8, 9, 12, 15, 20}.

3. Sum the total error across all swap curves in the sample period.
4. Update the values of θ, α_r, α_M, α_L, σ_M, σ_L, and ρ, and repeat steps 2-4.

The algorithm is repeated until the global minimum is reached. To solve the problem, we rely on the heuristic optimization method of Storn and Price (1997) through the `DEoptim` package in `R`. We further impose the following constraint: $\alpha_r > \alpha_M$, $\alpha_M > \alpha_L$, $\sigma_M > \sigma_L$.[10]

[10] We also impose reasonable bound constraints.

With end-of-month, mid-market USD and EUR par swap rates from Jan. 2006 to Aug. 2020, we estimate a set of parameters for each currency. Like in DLY, the constant parameters are estimated from in-sample data; this may seem like "cheating", since the seven parameters are estimated from data to which the trading strategy is later applied. In reality, however, one would probably reestimate and update these parameters from time to time, particularly after a change of regime; as a result, it would be quite unrealistic to use parameters estimated from a pre-2006 sample over the entire 2006-20 sample period. Of course, the parameters could have been reestimated quarterly or yearly, but that is, in our case, computationally not feasible; it is therefore, in our view, more realistic to estimate the parameters from in-sample data since we are restricted to a single set of parameters.

While the estimation procedure described above is attractive for our purposes, we will end this section with a brief discussion of alternative estimation methods. If the term structure model at hand allows for closed-form solutions to option prices, e.g. European swaptions, one could estimate the parameters by minimizing the sum of squared errors between option prices implied by the model and option prices observed in the market. However, this would involve frequent reestimation, and parameters could be volatile; besides, there need not be a clear link between, on the one hand, mean-reversion and volatility parameters implied by options prices and, on the other hand, mean-reversion and volatility parameters implied by the overall shape of the term structure. Indeed that would be to take the model too literally. Hence, for our task, the estimation procedure suggested in DLY seems more attractive; here, the parameters are in effect estimated to maximize the model fit.

Table 5.1

Estimated Parameters for the Three-Factor Model for USD and EUR Swap Markets

Parameter	USD	EUR
θ	0.036338	0.041238
α_r	0.371104	0.222642
α_M	0.363822	0.045252
α_L	0.018292	0.014048
σ_M	0.058384	0.036982
σ_L	0.006688	0.015962
ρ	-0.138732	-0.938070

6 Results

6.1 Return

In *Table 6.2* (pp. 34), we present summary statistics for the USD and EUR yield curve strategies. Results are presented both for the complete sample period (2006-2020) and for two subperiods, the first from 2006 to 2012 and the second from 2013 to 2020. The USD strategy performs quite well, with a mean excess return of 0.462% per month; since it has been scaled to produce a return volatility of 12% per year (3.464% per month), this implies a Sharpe ratio of 0.462.[11] Over the 2006-2012 sample the mean excess return is higher at 0.907% per month; the return volatility is higher too, but the Sharpe ratio still climbs to 0.677. The excess return is in both cases statistically significant at the 10% level. Over the 2013-2020 sample there is a clear deterioration in performance, with a mean excess return that is statistically insignificant and a Sharpe ratio of just 0.118. Volatility also declines.

Like for the USD strategy, the performance of the EUR strategy is better over the first half of the sample than over the second half. Nevertheless, the EUR strategy performs poorly, with a mean excess return of 0.152% and a Sharpe ratio of 0.152; the mean is not statistically significant, neither over the full sample nor over the two subsamples. Return volatility is also much lower over the second half of the sample, like for the case of USD; for EUR, the return volatility over the second half of the sample (1.402%) is under one third of the volatility over the first half of the sample period (4.819%). The ratio of negative to positive monthly returns is also above 0.5. For USD, that ratio is well below 0.5.

Looking at the other summary statistics, the USD and EUR return series have positive skewness and more kurtosis than would be the case for a normal distribution. In this case, the positive skewness suggests that the strategies are *not* merely designed to earn a small positive return most of the time, only to suffer large losses in times of market turmoil. Indeed, looking at the table, we see that the maximum monthly return for both the USD and EUR strategies is higher than the minimum monthly return (i.e. maximum monthly loss). For instance, the USD strategy has a maximum monthly gain of 16.158% over the entire sample while the maximum monthly loss is -11.994%; the EUR strategy, on the other hand, has a maximum monthly gain of 20.989% and a maximum monthly loss of -11.739%.

[11] We report the annualized Sharpe ratio. When return volatility is 12.0% per year, the annualized Sharpe ratio is equal to the mean monthly return multiplied by 100.

Figure 6.1

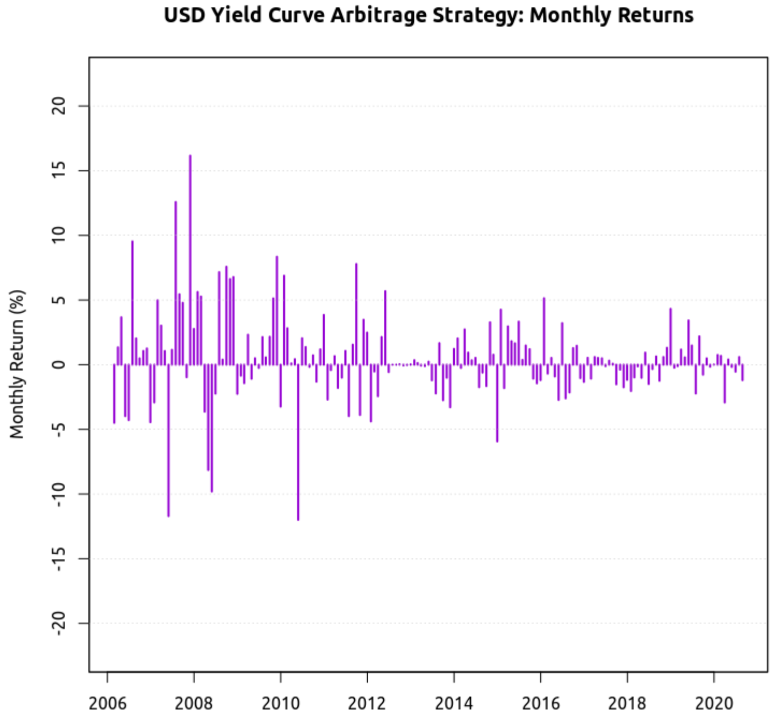

This figure shows the monthly excess returns of the USD yield curve strategy.

Figure 6.2

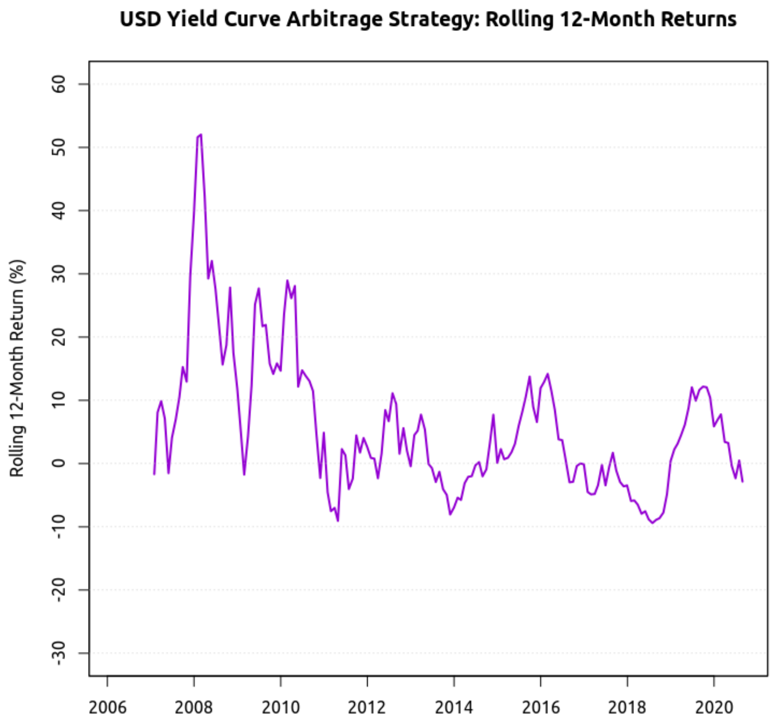

This figure shows the 12-month excess return of the USD yield curve strategy.

Figure 6.3

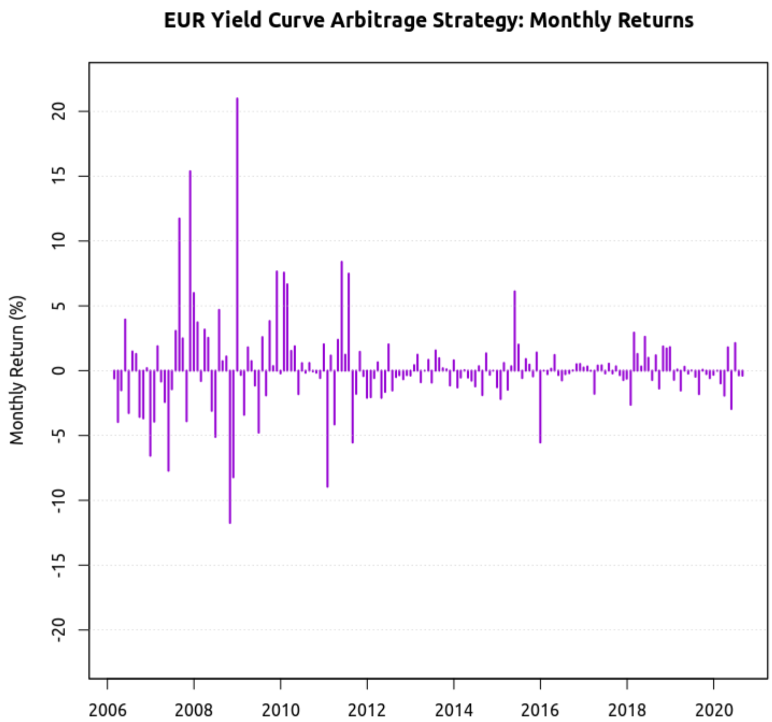

This figure shows the monthly excess returns of the EUR yield curve strategy.

Figure 6.4

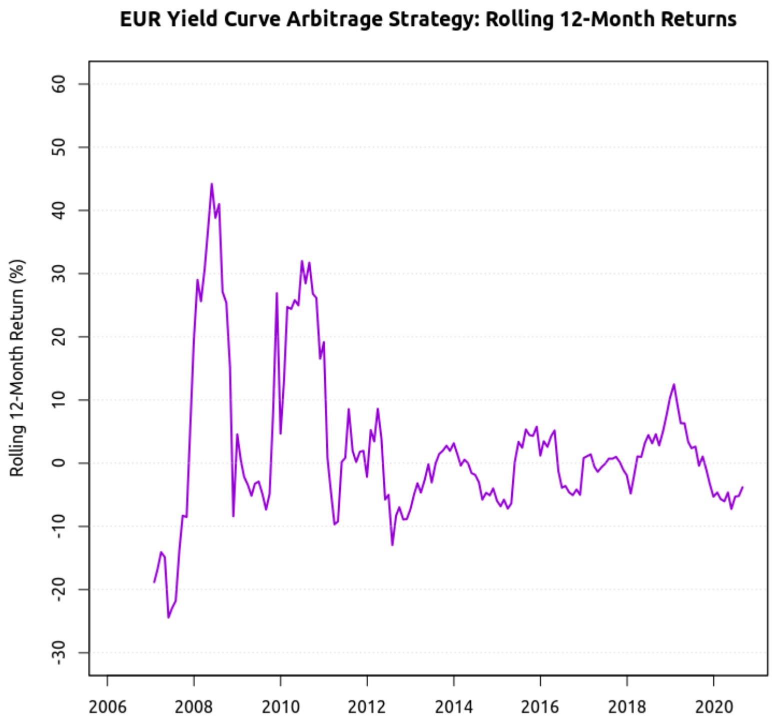

This figure shows the 12-month excess return of the EUR yield curve strategy.

Table 6.1

Descriptive Statistics for Model-Market Swap Rate Errors

EUR Rate	n	Mean	Median	Standard Deviation	Minimum	Maximum	USD Rate	n	Mean	Median	Standard Deviation	Minimum	Maximum
EUR-02Y	3752	-1.94	-2.06	9.73	-30.12	35.06	USD-02Y	3608	-4.64	-6.36	11.21	-34.82	44.07
EUR-03Y	3752	-4.61	-4.54	12.01	-41.94	25.74	USD-03Y	3608	-7.49	-9.86	13.81	-46.84	33.37
EUR-04Y	3752	-6.22	-5.52	11.45	-41.85	20.45	USD-04Y	3608	-8.23	-10.88	12.77	-44.99	28.72
EUR-05Y	3752	-6.65	-5.78	9.68	-35.27	13.99	USD-05Y	3608	-7.70	-9.59	10.14	-36.07	20.39
EUR-06Y	3752	-6.10	-5.24	7.56	-28.50	11.45	USD-06Y	3608	-6.57	-7.50	7.29	-27.10	15.92
EUR-07Y	3752	-4.87	-4.26	5.42	-20.72	8.88	USD-07Y	3608	-4.97	-5.36	4.80	-18.81	10.65
EUR-08Y	3752	-3.26	-2.87	3.35	-17.10	5.88	USD-08Y	3608	-3.10	-3.24	2.83	-11.17	6.66
EUR-09Y	3752	-1.58	-1.41	1.57	-14.84	4.18	USD-09Y	3608	-1.37	-1.42	1.28	-5.07	3.66
EUR-12Y	3752	2.27	1.98	2.18	-11.03	8.06	USD-12Y	3608	1.68	1.71	1.84	-5.49	6.50
EUR-15Y	3752	4.14	3.76	3.62	-11.20	14.19	USD-15Y	3608	3.20	3.25	2.89	-7.94	10.04
EUR-20Y	3752	4.88	4.55	3.33	-12.06	14.94	USD-20Y	3608	4.38	4.10	2.18	-3.64	10.65

This table reports summary statistics for deviations (in basis points) between swap rates implied by the three-factor model and swap rates observed in the market for 2-, 3-, 4-, 5-, 6-, 7-, 8-, 9-, 12-, 15- and 20-year USD and EUR interest rate par swaps. Deviations are computed at the daily frequency. The sample period is from Jan 2006 to Sep 2020.

Table 6.2

Descriptive Statistics for Yield Curve Arbitrage Strategies

Market	n	Capital	Mean	t-Statistic	Standard Deviation	Minimum	Maximum	Skewness	Kurtosis	Ratio Negative	Serial Correlation	Sharpe Ratio
EUR	175	39.216	0.152	0.58	3.464	-11.739	20.989	1.647	12.973	0.520	0.012	0.152
EUR (2006-2012)	83	-	0.326	0.62	4.819	-11.739	20.989	1.196	7.152	0.542	0.017	0.235
EUR (2013-2020)	92	-	-0.005	-0.03	1.402	-5.535	6.114	0.259	8.006	0.500	-0.072	-0.012
USD	175	25.878	0.462	1.77	3.464	-11.994	16.158	0.402	7.235	0.440	0.094	0.462
USD (2006-2012)	83	-	0.907	1.78	4.636	-11.994	16.158	0.103	4.614	0.386	0.117	0.677
USD (2013-2020)	92	-	0.061	0.33	1.799	-5.932	5.138	0.133	4.040	0.489	-0.149	0.118
EW	175	-	0.386	1.48	3.464	-12.212	19.829	1.213	10.129	0.509	0.093	0.386
EW (2006-2012)	83	-	0.775	1.49	4.753	-12.212	19.829	0.778	5.751	0.482	0.097	0.565
EW (2013-2020)	92	-	0.035	0.22	1.526	-4.532	4.884	-0.004	4.505	0.533	-0.054	0.081

This table reports summary statistics for monthly percentage excess returns of the yield curve strategies. EW is an equally-weighted portfolio of the USD and EUR strategies. Capital is the initial amount required per 100 notional to generate a 12% annualized volatility. "Ratio Negative" is the proportion of negative returns while Sharpe Ratio is the mean excess return divided by standard deviation. The *t*-statistics are corrected for heteroskedasticity and serial correlation. The sample period is from Feb 2006 to Aug 2020.

Moving on, it is worth analyzing why the EUR strategy underperforms compared with the USD strategy, and why, for both strategies, the performance is better in the first half of the sample than in the second half. On the topic of the first question it is difficult to identify what is the exact source of EUR underperformance, but a few issues are definitely discernible if we take a look at *Figure 6.6* on the next page in conjunction with *Table 5.1* (pp. 29). Firstly, over the period following the financial crisis of 2007-2009 up until the mid-2010s, the swap curve has an abnormal shape. In *Figure 6.6*, which shows the EUR swap curve on Jan. 10, 2012, the 20- and 30-year swap rates are low relative to intermediate-term rates, which leads to implied forward rates being steeply downward-sloping from 12 years out, out to 30 years. To be clear, the model has no trouble fitting this swap curve, but the abnormal shape impacts the values of the model's *constant* parameters. Looking back to *Table 5.1*, we see that α_M is low for EUR compared with the USD market; in fact, the half-life implied by the value of α_M in the EUR market is 15.32 years![12] In other words, the abnormal, downward-sloping EUR forward curve necessitates a very low α_M since the influence of the medium-term factor is required beyond the intermediate-term segment in order to generate an inflection point on the forward curve in the 20-year range. This, however, comes at a cost: the medium-term factor is no longer suited to describe the intermediate-term segment of the swap curve and so, as a result, swap rates in the 2- to 9-year segment are more frequently mispriced.

Not only does the EUR model parameterization lead to mispricing, but it also leads to hedges being less effective. Once again, the culprit is α_M – due to its low value, the medium-term factor more or less becomes a long-term factor and, as a result, M and L become more equal. Since positions are hedged by zeroing out the exposures to M and L, portfolios are in this case predominantly protected against parallel shifts. This is depicted in *Figure 6.7* below, which shows the factor loadings in the EUR parameterization of the three-factor model; here, M and L impact swap rates in a rather similar manner. The short-term factor now looks like the medium-term factor in USD, which may suggest that we should hedge out that factor (and not M); this, however, does not yield any improvement. Of course, another possibility would be to zero out the exposure to all three factors of the model with a combination of 1-, 10-, and 30-year swaps for every cheap or rich swap. This method was tested out, but did not improve on the original results; in this case, even more swaps are added to the overall portfolio, which introduces more interest rate risk along the whole curve.

[12] We will not get into the details here, but the half-life is found by solving exp[-α * *t*] = 0.50 for *t*, where α is the mean-reversion speed.

Figure 6.5

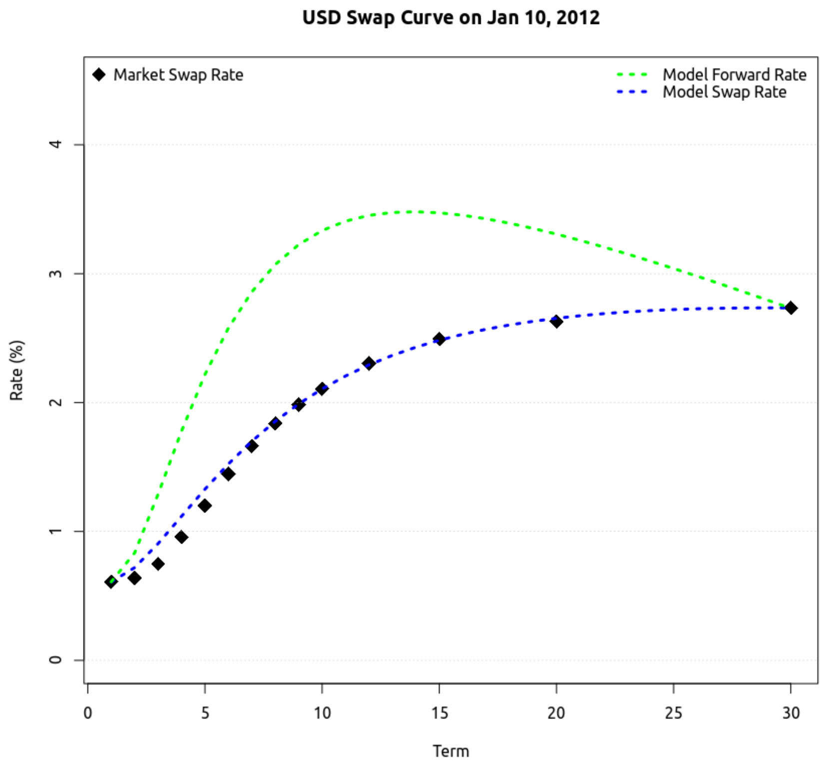

This figure shows the three-factor model fitted to match the 1-, 10-, and 30-year USD swap rates on 01/10/2012.

Figure 6.6

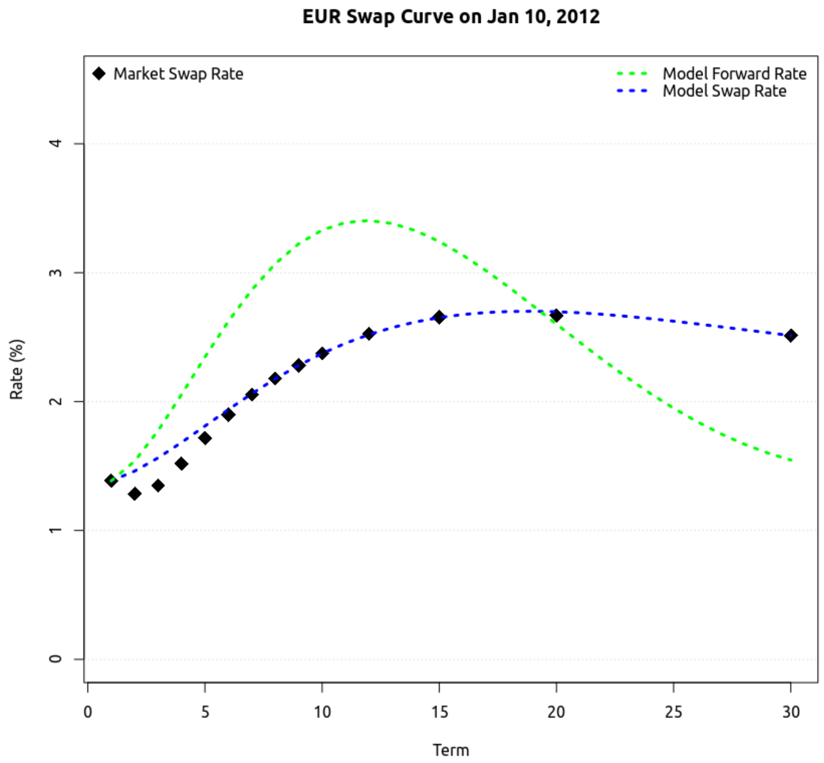

This figure shows the three-factor model fitted to match the 1-, 10-, and 30-year EUR swap rates on 01/10/2012.

Figure 6.7

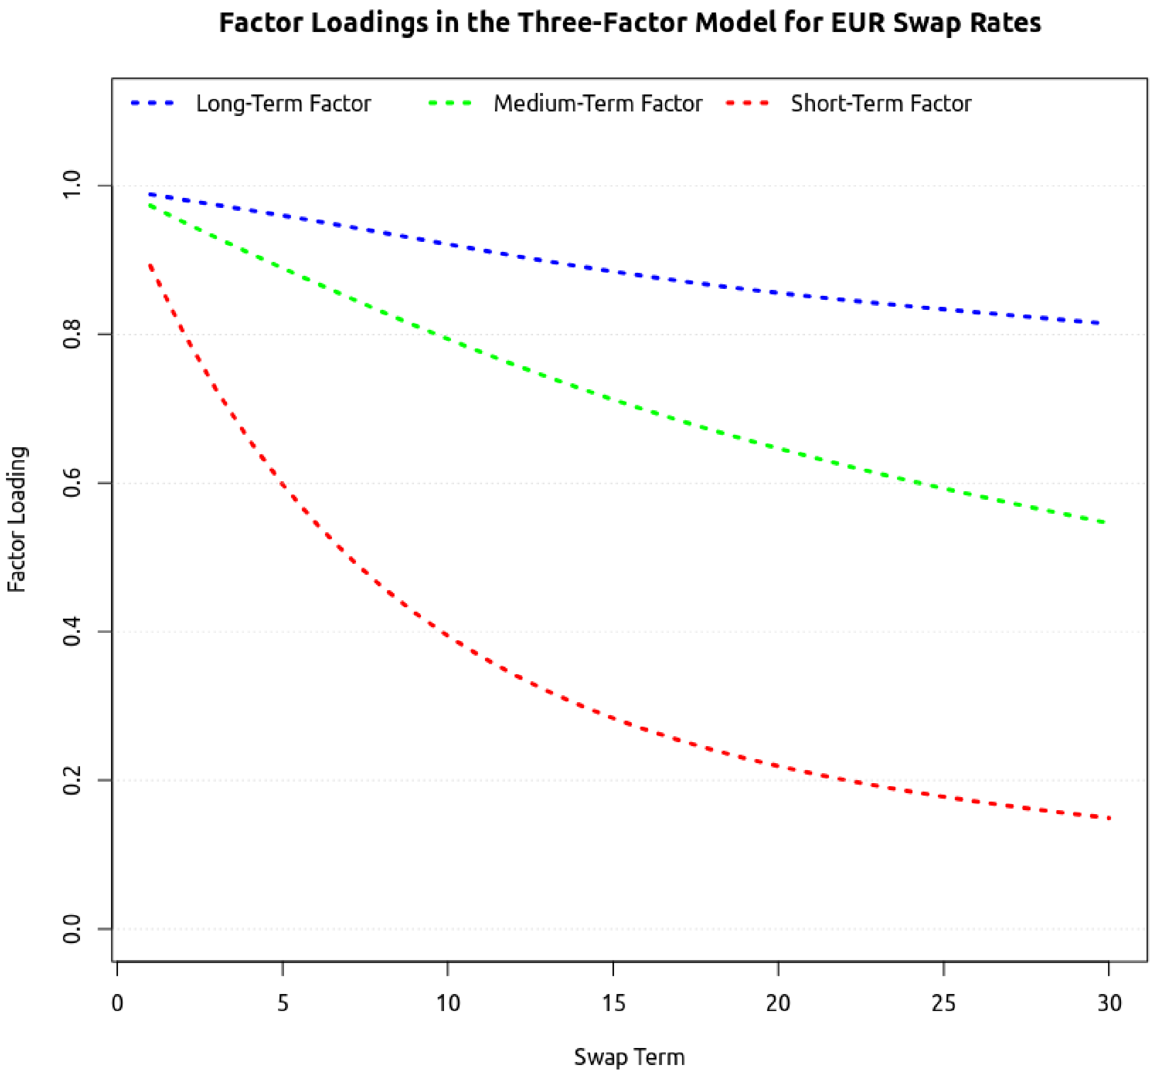

This figure shows the derivatives of swap rates with respect to factor values in the three-factor model (for EUR).

Of course, if the EUR forward curve in *Figure 6.6* is a reflection of the expected value of future short-term rates, the interest rate process of the three-factor model might not be well suited for the EUR market; in fact, an additional interest rate factor would be needed to allow for two inflection points on the swap curve. However, long-term swap rates did *not* decline in the early 2010s based on market participants forecasting lower short-term rates 15 to 30 years into the future, but rather as a result of a decline in the term premium. In Domanski, Shin and Sushko (2017), they argue that this decline was – in addition to the ECB's quantitative easing policies – caused by insurance company and pension fund demand for long-duration assets to match their typical long-duration liabilities. No matter the exact cause, a modeler would need to supplement the three-factor model with a time-varying risk premium in order to capture the peculiar properties of the EUR swap curve; alternatively, a modeler could ignore the long end of the swap curve and only trade the 1- to 10-year segment with the two-factor Vasicek model fitted to match the, say, 1-year and 10-year market rates.

To summarize, the EUR strategy underperforms because the three-factor model is not well suited to describe the peculiar shape of the EUR swap curve that manifests over much of the sample period. Model parameters are forced to take on rather extreme values, which leads to more mispricing and less effective hedge ratios than for the USD case (where the model is better suited to describe the "real" interest rate process).

We end this subsection with a discussion of why, for both markets, the performance is better in the first half of the sample than in the second half. The answer is plainly visible from *Figure 6.8* and *6.9* on the next page, which shows the rolling 250-day volatility of select USD and EUR swap rates (in basis points per year). In both markets swap rate volatility is elevated during the financial crisis of 2007-2009 before it declines and remains low for the larger part of the remaining sample period; the volatility of the 2-year USD rate, for instance, is less than 50 basis points per year from 2012 up to 2020. With a low level of volatility, certain segments of the swap curve may appear cheap or rich for long periods of time and the strategy becomes less profitable; with higher levels of volatility there is much more movement around the fitted swap rates, which essentially increases the probability that certain segments of the curve will move toward "fair value". When volatility is low, however, there is limited movement around those fitted swap rates, which lowers the probability that cheap or rich segments will return to their fair values in the model before a position is closed.

While low levels of swap rate volatility is unfortunate for the strategies, one can make the argument that the dependence on rate volatility is an attractive feature. Like they mention in DLY, fixed income arbitrageurs are sometimes criticized for "picking up nickels in front of a steamroller", meaning that they tend to earn a small positive return most of the time only to suffer huge losses in times of market stress. Thus, from the perspective of a large hedge fund, which probably engages in a lot of different strategies, a yield curve strategy may serve as an important "diversifier"; if the fund's other strategies tend to be negatively correlated with the level of interest rate volatility, a yield curve strategy can help to make the fund's returns more "market neutral" (or at least cushion a loss). The fund could of course remedy the problem by buying fixed income options, e.g. swaptions, but that could lead to steady losses; with a well-implemented yield curve strategy, however, the fund has the potential to earn positive returns even when volatility is declining (like in the USD case).

Figure 6.8

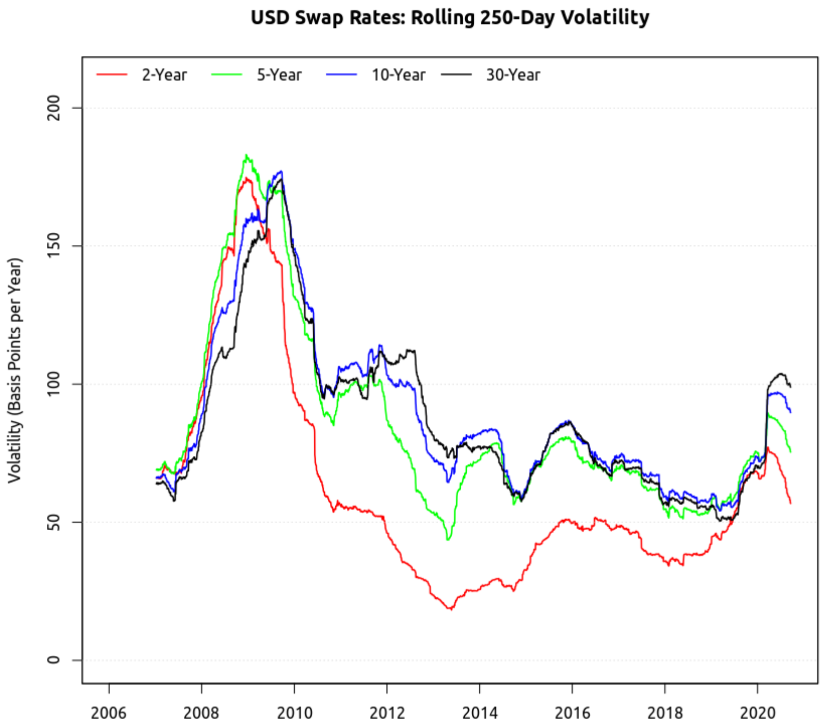

This figure shows the rolling 250-day volatility of (changes in) select USD swap rates (in basis points per year).

Figure 6.9

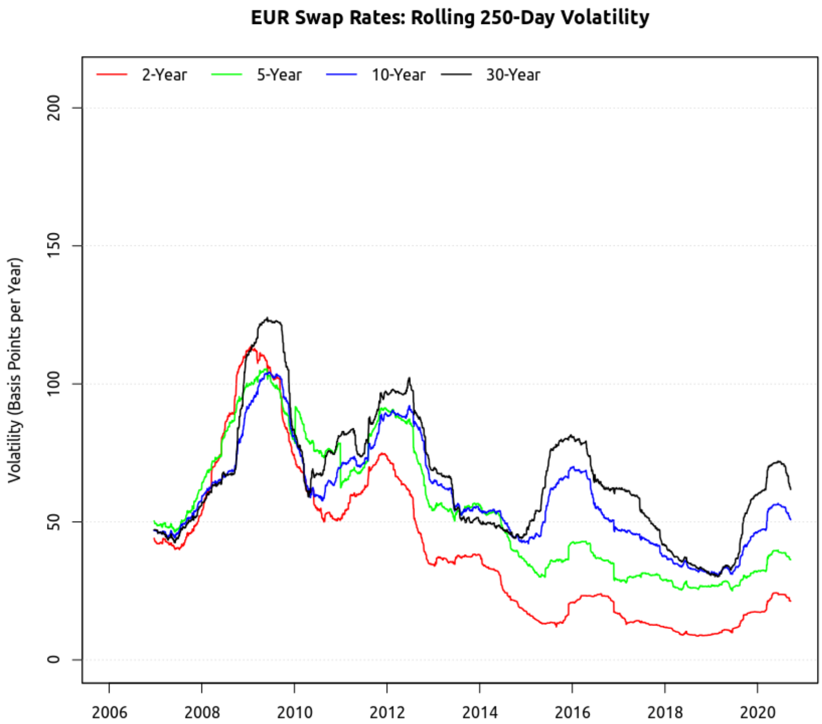

This figure shows the rolling 250-day volatility of (changes in) select EUR swap rates (in basis points per year).

6.2 Risk

To summarize the results of the previous subsection the USD strategy produces excess returns that are positive and statistically significant, with a Sharpe ratio of 0.462; by contrast, the EUR strategy performs poorly with an insignificant mean excess return and a Sharpe ratio of 0.152. Still, it remains to determine whether, especially for the USD case, the excess return is earned as compensation for bearing market risk, or if it represents true “alpha”. As a result, we need to determine if the residual risks of the strategies are correlated with risk factors that have been identified in the literature. To answer the question, we regress the excess returns of the strategies on the excess returns of risk factors that are popular in the literature. We control for equity market risk through the Fama-French U.S. and European market portfolios and the SMB, HML, and UMD portfolios for each market; return data for these portfolios is obtained from Kenneth French’s data library.[13] Here, excess returns from the USD strategy is regressed on the U.S. portfolios, and excess returns from the EUR strategy is regressed on the European portfolios. Moreover, we include the KBW Bank Index *or* the Stoxx Europe 600 Bank Index. To control for default risk, we include the Bloomberg/Barclays U.S. Corporate Bond Index *or* Euro Corporate Bond Index (broad indexes of USD and EUR-denominated corporate bonds). We control for long-term interest rate risk through two indexes from Credit Suisse which – in our case – are meant to serve as proxies for an investment in liquid, long-term U.S. Treasuries and liquid, long-term German bunds. We also control for volatility risk through the returns on CBOE’s VIX Index *or* the Stoxx 50 Volatility Index; since there are futures contracts on these indexes, this risk factor is tradable like the other factors.

The reason for including financial stocks as a risk factor is that financials can serve as a proxy for financial sector risk. Like during the financial crisis of 2007-2009, financials will typically underperform if investors are concerned about the health of bank balance sheets and the financial system as a whole. In such an environment, LIBOR rates – which reflect banks’ borrowing costs – typically increase relative to risk-free rates, like OIS rates. Therefore, if the trading strategy is exposed to, say, the LIBOR/OIS spread, that risk should be captured by the performance of financial stocks. Nevertheless, this link is far from perfect as financial stocks are influenced by many other variables: in Viale, Kolari and Fraser (2009), for example, they show that the return on financial stocks are sensitive to the slope of the yield curve (since the slope of the yield curve affects banks’ interest margins).

---

[13] To be specific, we use the “Fama/French 3 Factors” and “Momentum Factor (Mom)” files for the U.S. market; for Europe, we use the “Fama/French European 3 Factors” and “European Momentum Factor (Mom)” files.

Table 6.3

Regression Results for Yield Curve Arbitrage Strategies

	Intercept		t-Statistics									
Market	α	t-Statistic	R_M	SMB	HML	UMD	R_{BANKS}	R_{BONDS}	R_{GOV}	R_{VOL}	F-Statistic	R^2
EUR	-0.038	-0.11	1.37	-1.50	0.32	-1.88	-1.14	0.27	1.32	1.33	2.43	0.105
EUR (2006-2012)	-0.384	-0.70	2.24	-1.87	0.06	-1.98	-1.78	0.60	1.38	1.15	2.63	0.221
EUR (2013-2020)	0.129	0.87	-1.17	0.44	0.15	-1.14	0.57	0.94	-0.45	-0.07	0.70	0.063
USD	0.426	1.85	-2.35	-1.27	0.72	0.32	1.75	-0.30	3.50	-0.34	6.04	0.226
USD (2006-2012)	0.711	1.68	-1.21	-1.27	-0.15	0.19	1.55	-0.94	3.01	0.38	3.16	0.255
USD (2013-2020)	0.248	1.48	-1.36	0.51	1.41	0.38	-1.13	2.11	0.55	-1.26	4.49	0.302

This table reports summary statistics for the regression of monthly percentage excess returns of the yield curve strategies on the excess returns of various risk factors that are popular in the literature. R_M is the excess return on the Fama-French value-weighted equity market portfolio. SMB, HML, and UMB are the Fama-French equity risk factors. R_{BANKS} is the excess return on the KBW Bank Index (for USD) or the Stoxx Europe 600 Bank Index (for EUR). R_{BONDS} is the excess return on the Bloomberg Barclays U.S. Corporate Bond Index (for USD) or the Bloomberg Barclays Euro Corporate Bond Index (for EUR). R_{GOV} is the excess return on Credit Suisse's 10-Year U.S. Treasury Note Futures Return Index (for USD) or Credit Suisse's 10-Year German Government Bond Futures Return Index (for EUR). R_{VOL} is the excess return on CBOE's VIX Index (for USD) or the Stoxx 50 Volatility Index (for EUR). The *t*-statistics are corrected for heteroskedasticity and serial correlation. The sample period is from Feb 2006 to Aug 2020.

We present regression results in *Table 6.3* on the next page. For the USD strategy, the risk-adjusted mean excess return is 0.426% per month over the entire sample and 0.711% per month over the first half of the sample; in both cases, this risk-adjusted mean excess return is statistically significant at the 10% level. Interestingly, over the second half of the sample, the risk-adjusted mean excess return is *higher* than the mean excess return, but the statistic is not significant at common confidence levels. Over the full sample, the USD strategy is negatively correlated with the market portfolio and positively correlated with long-term Treasury bonds. The positive correlation with long-term Treasury bonds implies that the hedging strategy may not be entirely effective: if the hedging strategy was entirely effective, there would not be any residual interest rate risk remaining on a portfolio level.

While the R^2 of the USD strategy is relatively low, it is still higher than what they find in DLY (where the R^2 is 0.097). Nevertheless, the exposures to the risk factors are negative as well as positive, and the risk-adjusted mean excess return is only slightly lower than the mean excess return from *Table 6.2*; in other words, most of the mean excess return is indeed “skill” – it is not just earned as reward for bearing market risk.

With respect to the EUR strategy, results are a little less informative since the strategy has a low mean excess return to begin with. Adjusting for risk seem to lower the mean excess return slightly, but the intercept is not statistically significant, neither over the full sample nor over any of the two subsamples. Over the full sample, none of the *t*-statistics are significant at the 5% level, and the R^2 is low. Though not significant, the mean excess return is higher over the second half of the sample after adjusting for risk, like for the case of USD; if that is a true reflection of reality it means that the strategy will earn a negative risk premium, which means that negative excess returns can lead to a positive alpha.

With the results from *Table 6.3*, we can conclude that yield curve arbitrage has a fairly low level of exposure to risk factors, though some long-term interest rate risk seem to remain as a result of less-than-perfect hedge ratios. In our view, however, the strategy has significant *model risk*; the strategy is largely based on the assumption that the curvature and shape of the swap curve will be similar in the future as in the past. When this is *not* the case, like for EUR in the second half of the sample, the model breaks down and offers misleading predictions on the cheapness/richness of certain swap rates. Thus the positive, risk-adjusted excess returns of the USD strategy may simply be a reward for bearing model risk. In DLY, they argue that the positive, risk-adjusted excess return is a result of the intellectual capital needed to operate the strategy; be that as it may, but the excess return could also be a reward for relying on a model that may not perform as well in the future as in the past.

7 Conclusion

In this thesis, we examine the risk- and return-characteristics of yield curve arbitrage. Following the approach of Duarte, Longstaff and Yu (2007), we construct hypothetical return indexes for the 2006-2020 period by implementing the strategy on historical data of USD and EUR par swap rates. Our implementation relies on an affine three-factor term structure model to both identify and hedge cheap and rich maturities along the yield curve. Model parameters are estimated through a numerical optimization routine, like in Duarte et al. The strategies are implemented on daily data, and every practical detail of the USD and EUR swap markets are taken into account in order to produce the most realistic return series.

Over the 2006-2020 sample, with the initial amount of capital set so that the strategies produce a return volatility of 12% per year, the USD strategy produces an excess return of ca. 5.5% per year, implying a Sharpe ratio of 0.462. The excess return of the EUR strategy is not statistically different from zero, which we argue is a result of the three-factor model not being well suited to describe the peculiar shape of the EUR swap curve that manifests over much of the sample period. For both USD and EUR, performance is markedly better over the first half of the sample (2006-2012) than over the later half (2013-2020), which coincides with a fall in par rates volatility. This suggests that the returns to the strategy are positively correlated with the level of interest rate volatility, which is opposite to that of other arbitrage strategies – such as risk arbitrage – where heightened volatility usually leads to losses.

Seemingly, adjusting for exposure to risk factors does not materially lower the excess returns of the yield curve strategies. The USD strategy produces a risk-adjusted excess return of over five percent per year, which is significant at the 10% level; for EUR, the risk-adjusted excess return is zero. The R^2 is roughly 10% for EUR and roughly 22% for USD, the latter of which being higher due to negative correlation with the stock market and positive correlation with Treasury bonds. The positive correlation with Treasury bonds illustrates that the hedges are less than perfect and that more effective hedges can lower the R^2.

The findings in this thesis are broadly consistent with the results of Duarte et al., who report that yield curve arbitrage produces positive, risk-adjusted excess returns with a limited degree of exposure to risk factors (over the 1988-2004 period, USD).

The question remains whether yield curve arbitrage produces excess returns and, if so, whether that excess return is “alpha” or merely reward for bearing risk. Our results from USD swap markets suggest that a well-implemented strategy has the potential to generate positive, risk-adjusted excess returns with a limited degree of exposure to risk factors. On the contrary, our results from EUR swap markets illustrate that the strategy is exposed to model risk and to the risk of structural changes in the shape and volatility of the term structure. That being said, the model risk is likely overstated in our purely quantitative implementation; in reality, someone pursuing the strategy would have more flexibility in terms of both revising the model and tuning the model's parameters if performance starts to deteriorate. We therefore conclude that there is potential for risk-adjusted excess returns, but that this potential can only be unlocked through astute modeling and deep knowledge of market dynamics. This conclusion is broadly similar to the one in Duarte et al., where they propose that the returns to yield curve arbitrage are a result of the strategy requiring intellectual capital to implement.

With respect to future profitability, we believe the main risk of the strategy, especially following the monetary policy response to the COVID-19 recession, is increased central bank influence over the yield curve, which may create difficult-to-model term structure shapes and dampen term structure volatility. This is clearly visible from the second half of the sample, in which increased central bank influence and lower interest rate volatility – particularly in EUR markets – coincides with a marked stagnation in strategy profitability.

Finally, we would like to encourage further research in what we regard as a somewhat under-researched, yet still very interesting, area of finance. It would be interesting to see how other implementations of yield curve arbitrage – perhaps using a different model or parameter estimation procedures – perform, and whether risk- and return-characteristics are different in less liquid markets, such as NOK and SEK. Outside this particular realm, we also believe the finance literature would benefit from further research on the risk- and return-characteristics of other popular fixed income arbitrage trades. This will allow investors to make more informed judgements about whether or not to allocate capital to such strategies.

[*This page intentionally left blank*]

8 Appendix

8.1 Closed-Form Solutions for Spot Rates

The continuously compounded spot rate for term T in the three-factor model is given by

$$\hat{r}(T) = \theta + \bar{x}_1 \frac{1 - e^{-\alpha_r T}}{\alpha_r T} + \bar{x}_2 \frac{1 - e^{-\alpha_M T}}{\alpha_M T} + \bar{x}_3 \frac{1 - e^{-\alpha_L T}}{\alpha_L T} + \Upsilon(T)$$

where $\bar{x}_1$, $\bar{x}_2$, and $\bar{x}_3$ are the initial values of the three factors and

$$\begin{aligned}
\Upsilon(T) = & -\frac{1}{2}(G_1^2 + G_2^2) + (\sigma_{11} G_1 + \sigma_{12} G_2) \frac{1 - e^{-\alpha_r T}}{\alpha_r^2 T} \\
& + (\sigma_{21} G_1 - \sigma_{12} G_2) \frac{1 - e^{-\alpha_M T}}{\alpha_M^2 T} - (\sigma_{11} + \sigma_{21}) G_1 \frac{1 - e^{-\alpha_L T}}{\alpha_L^2 T} \\
& - (\sigma_{11}\sigma_{21} - \sigma_{12}^2) \frac{1 - e^{-(\alpha_r + \alpha_M)T}}{\alpha_r \alpha_M (\alpha_r + \alpha_M) T} \\
& + \sigma_{11}(\sigma_{11} + \sigma_{21}) \frac{1 - e^{-(\alpha_r + \alpha_L)T}}{\alpha_r \alpha_L (\alpha_r + \alpha_L) T} \\
& + \sigma_{21}(\sigma_{11} + \sigma_{21}) \frac{1 - e^{-(\alpha_M + \alpha_L)T}}{\alpha_M \alpha_L (\alpha_M + \alpha_L) T} \\
& - (\sigma_{11}^2 + \sigma_{12}^2) \frac{1 - e^{-2\alpha_r T}}{4\alpha_r^3 T} \\
& - (\sigma_{12}^2 + \sigma_{21}^2) \frac{1 - e^{-2\alpha_M T}}{4\alpha_M^3 T} - (\sigma_{11} + \sigma_{21})^2 \frac{1 - e^{-2\alpha_L T}}{4\alpha_L^3 T}
\end{aligned}$$

and

$$\begin{aligned}
G_1 &= \frac{\sigma_{11}}{\alpha_r} + \frac{\sigma_{21}}{\alpha_M} - \frac{\sigma_{11} + \sigma_{21}}{\alpha_L} \\
G_2 &= \frac{\sigma_{12}}{\alpha_r} - \frac{\sigma_{12}}{\alpha_M}
\end{aligned}$$

and

$$\begin{aligned}
\sigma_{11} &= \frac{\alpha_r \alpha_M}{(\alpha_r - \alpha_M)(\alpha_r - \alpha_L)} \sigma_L - \frac{\alpha_r}{\alpha_r - \alpha_M} \rho \sigma_M \\
\sigma_{21} &= \frac{\alpha_r}{\alpha_r - \alpha_M} \rho \sigma_M - \frac{\alpha_r \alpha_M}{(\alpha_r - \alpha_M)(\alpha_M - \alpha_L)} \sigma_L \\
\sigma_{12} &= -\frac{\alpha_r}{\alpha_r - \alpha_M} \sigma_M \sqrt{1 - \rho^2}.
\end{aligned}$$

8.2 Discount Curve Construction

With 1-, 2-, 3-, 4-, 5-, 6-, 7-, 8-, 9-, 10-, 12-, 15-, 20-, and 30-year par swap rates, we first use a custom cubic spline to interpolate par swap rates at semiannual intervals (USD) or at annual intervals (EUR). The starting value of the custom cubic spline is constrained to fit the current value of 3-month USD-LIBOR or 6-month EURIBOR. In notation, we find n splines of form

$$\phi_i(t) = a_i + b_i(t - t_i) + c_i(t - t_i)^2 + d_i(t - t_i)^3, \quad i = 0, ..., n - 1$$

where n is 15, t_i is in {0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 15, 20}, t_n is 30, and

$$\begin{aligned} &\phi_i(t_i) = y_i \text{ and } \phi_i(t_{i+1}) = y_{i+1}, && i = 0, ..., n - 1 \\ &\phi_i'(t_{i+1}) = \phi_{i+1}'(t_{i+1}), && i = 0, ..., n - 2 \\ &\phi_i''(t_{i+1}) = \phi_{i+1}''(t_{i+1}), && i = 0, ..., n - 2 \end{aligned}$$

where y_i is the swap rate for year t_i. We lastly specify that $\phi_{n-1}'(30) = 0$ and $\phi_{n-1}''(30) = 0$.

With semiannual or annual par swap rates out to 30 years, bootstrapping is used to extract the implied discount factors. Then, discount factors for arbitrary dates in the future are found via log-linear interpolation of the corresponding forward rates. In notation,

$$f(t_{i-1}, t_i) = \frac{1}{t_i - t_{i-1}} \ln[\frac{D(t_{i-1})}{D(t_i)}]$$

where $f(t_{i-1}, t_i)$ is the instantaneous forward rate from term t_{i-1} to term t_i, and $D(t_i)$ is the implied discount factor to term t_i. To obtain for the discount factor for an arbitrary date in the interval (t_{i-1}, t_i), say, t_k – where $t_{i-1} < t_k < t_i$ – we use the equation

$$D(t_k) = D(t_{i-1})e^{-f(t_{i-1},t_i)[t_k - t_{i-1}]}.$$

We are aware that many market participants have transitioned away from LIBOR discounting towards OIS discounting following the financial crisis of 2007-2009. An OIS curve, however, is constructed from par swap rates *and* basis swap spreads; since we do not have high-quality data on Fed Funds vs 3-month LIBOR *or* EONIA vs 6-month EURIBOR basis swap spreads, we stick with the traditional methodology (which is still very popular).

8.3 Figures

These figures show the deviation between the swap rates implied by the three-factor model and the market rates.

Figure 8.01

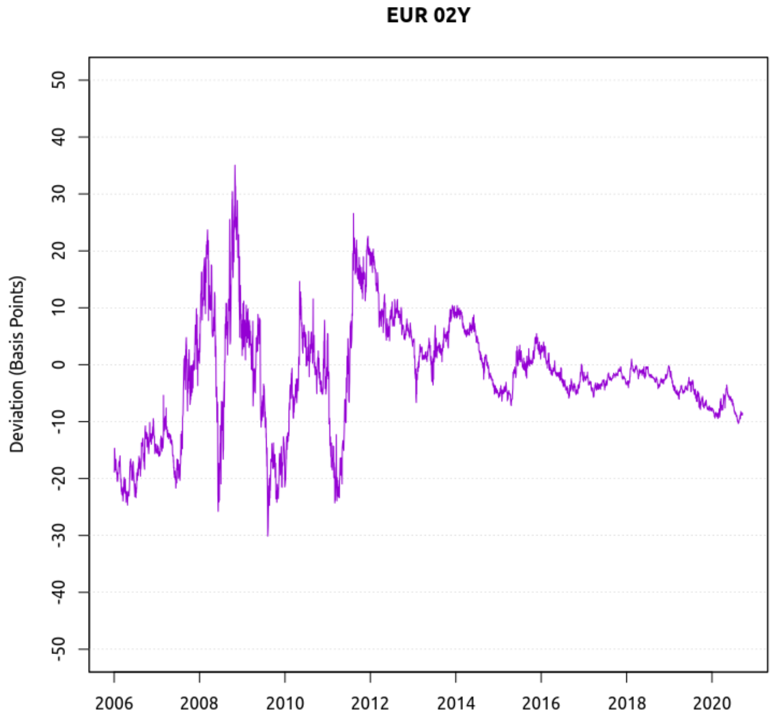

Figure 8.02

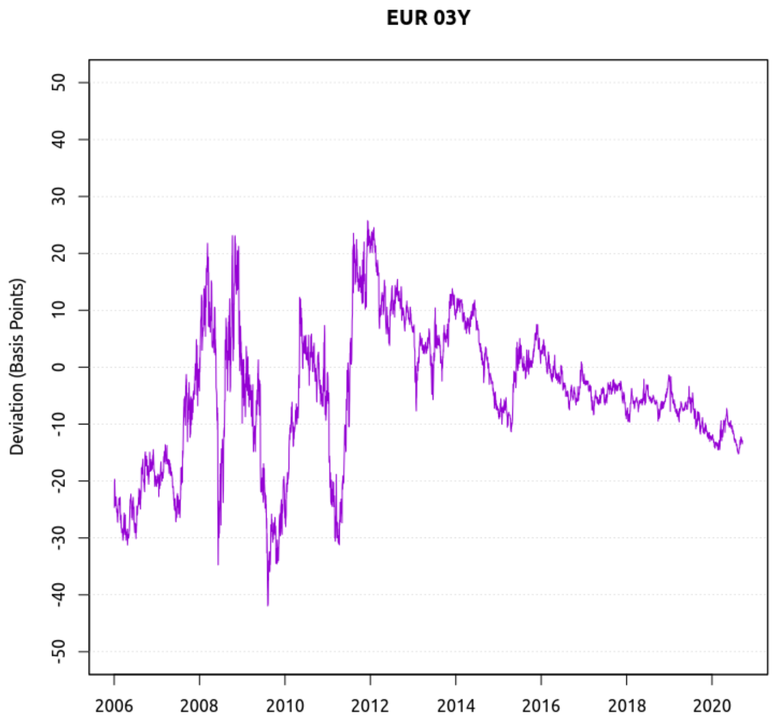

Figure 8.03

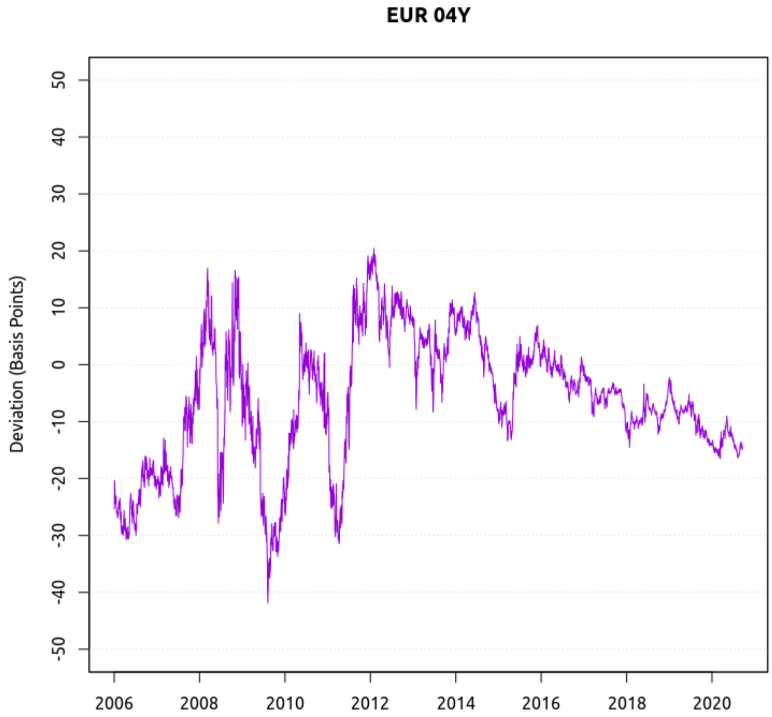

Figure 8.04

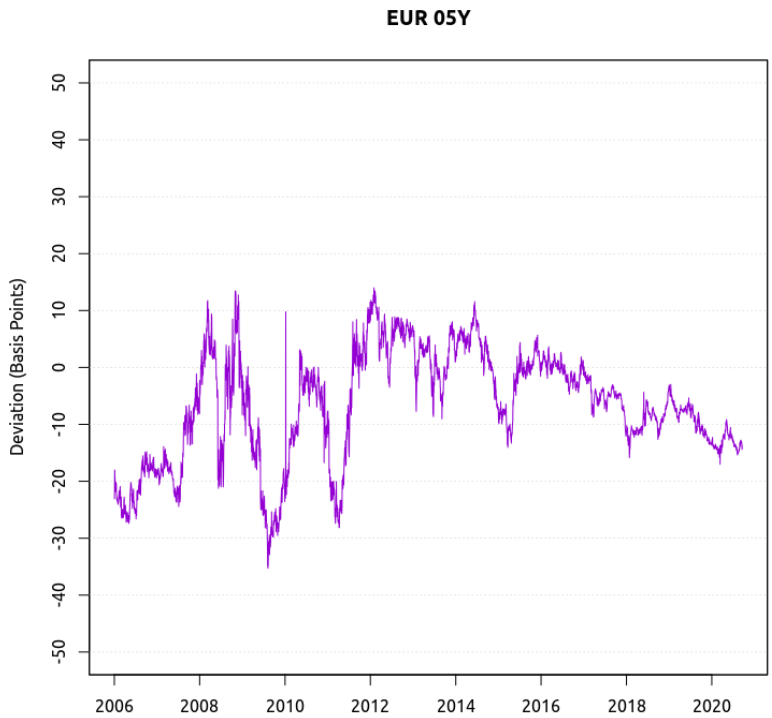

Figure 8.05

Figure 8.06

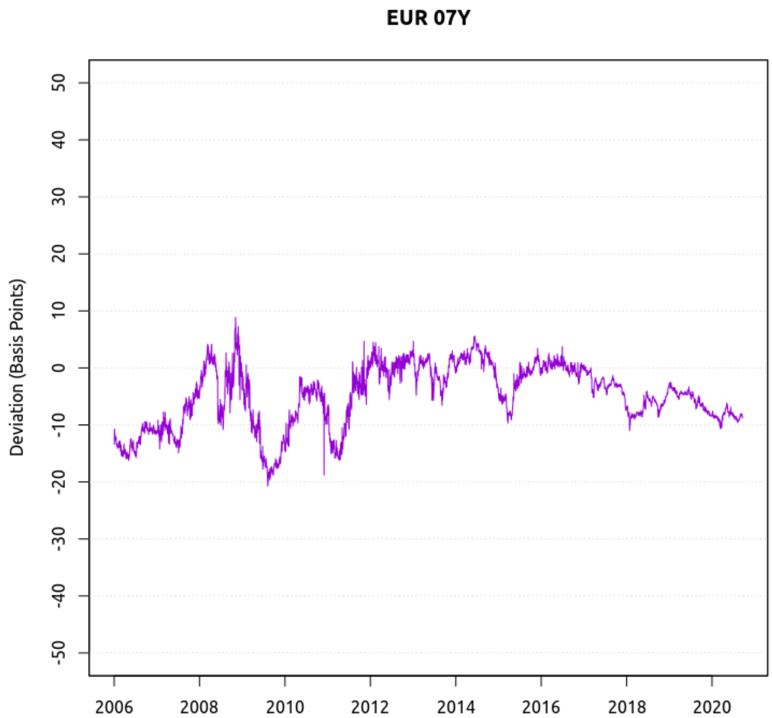

Figure 8.07

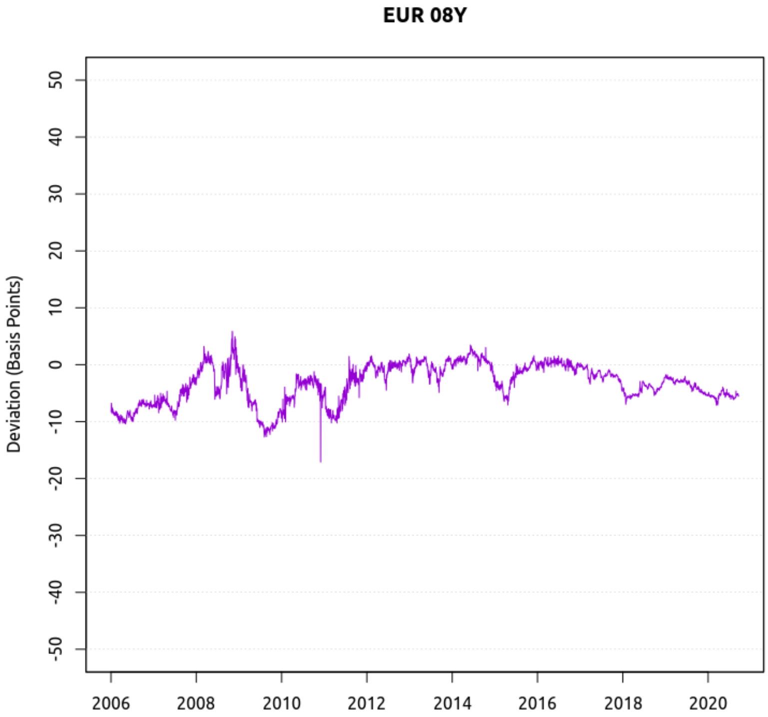

Figure 8.08

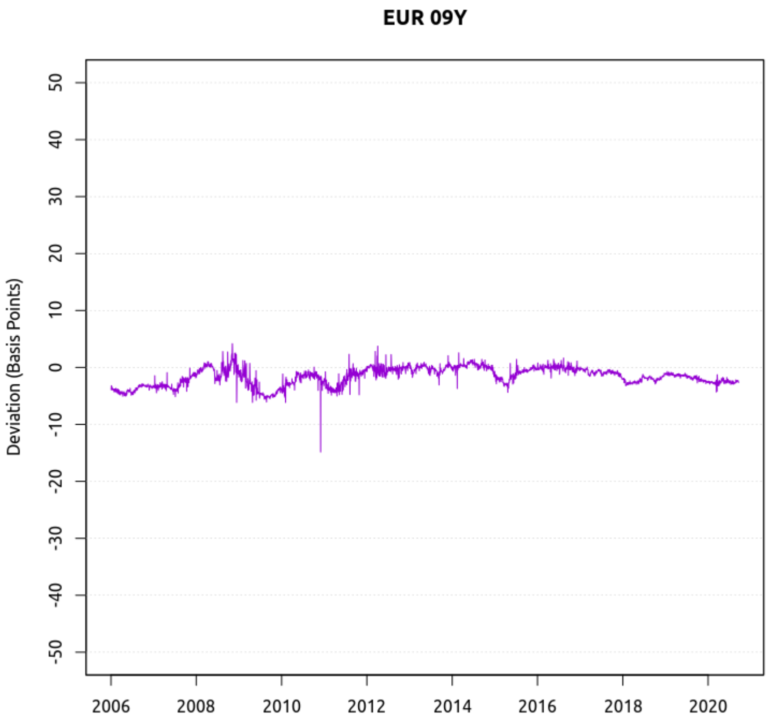

Figure 8.09

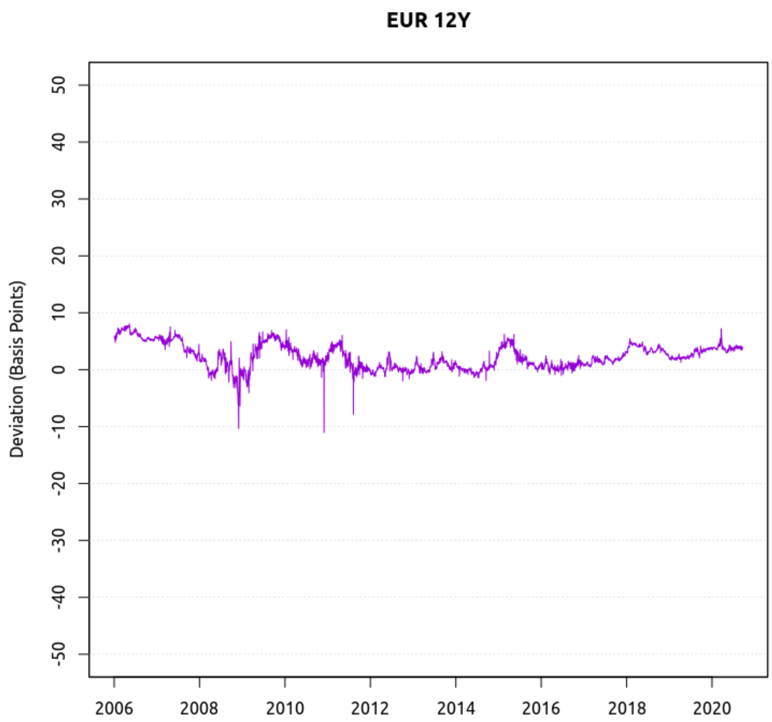

Figure 8.10

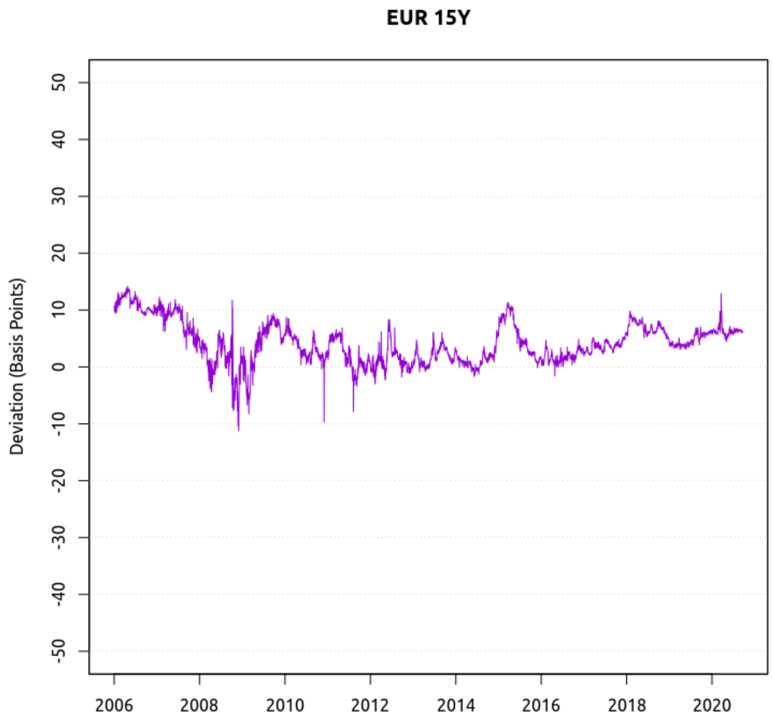

Figure 8.11

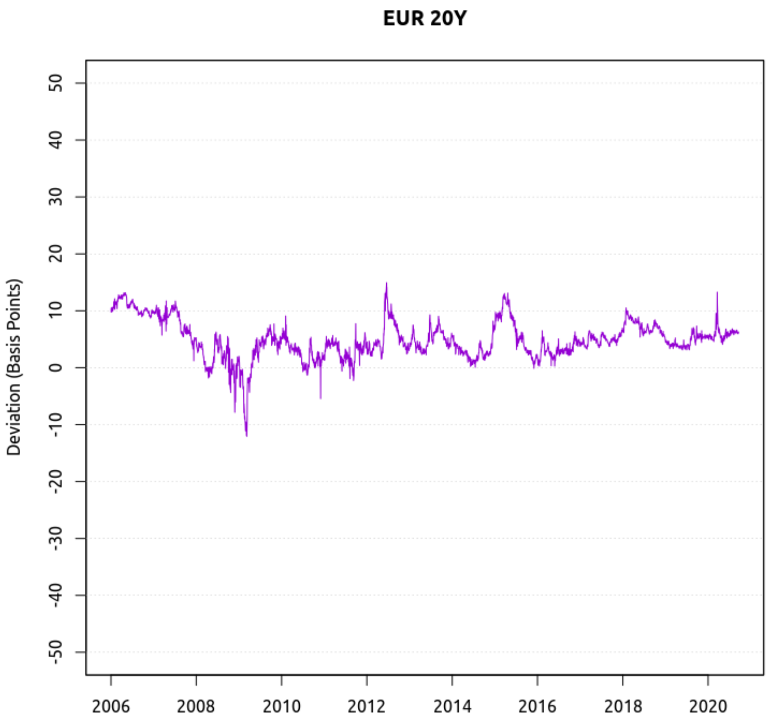

Figure 8.12

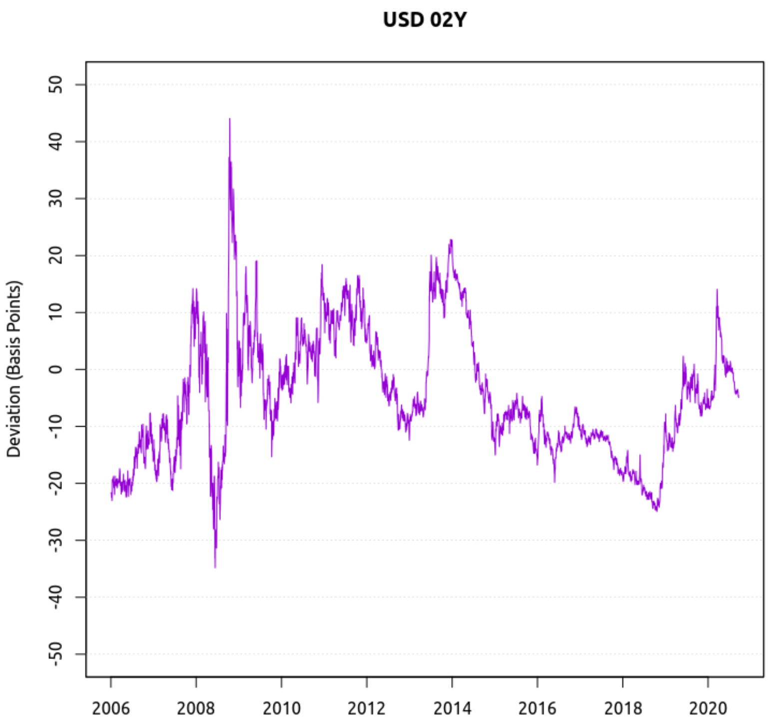

Figure 8.13

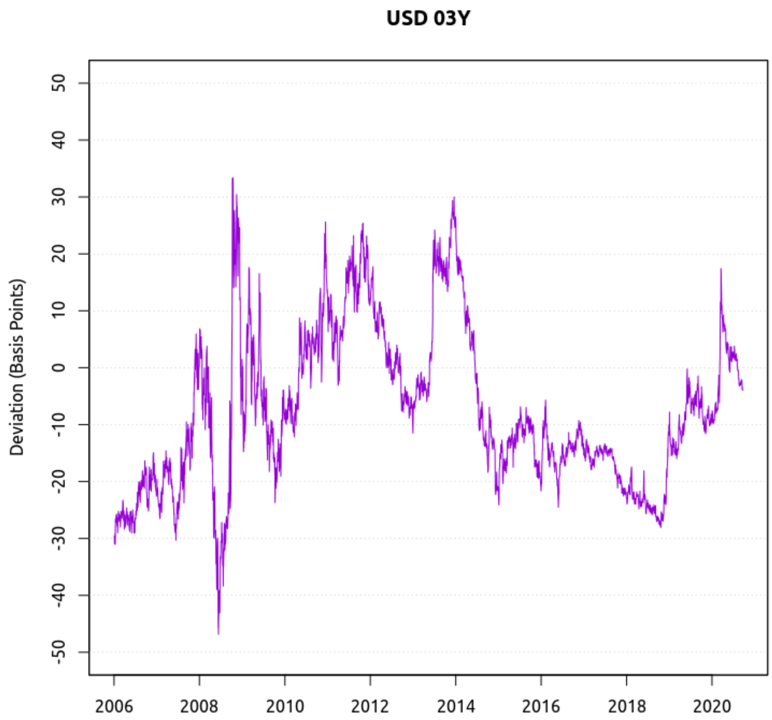

Figure 8.14

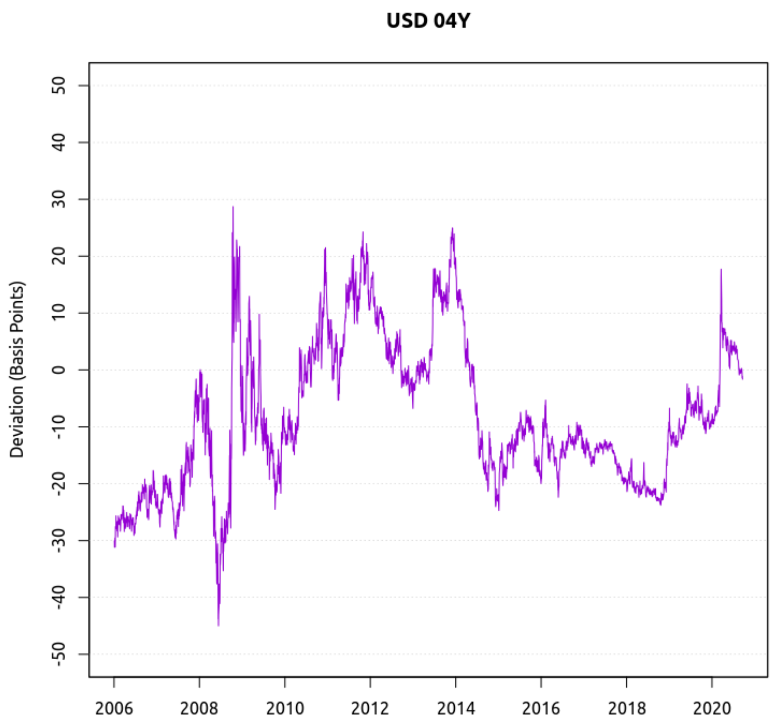

Figure 8.15

Figure 8.16

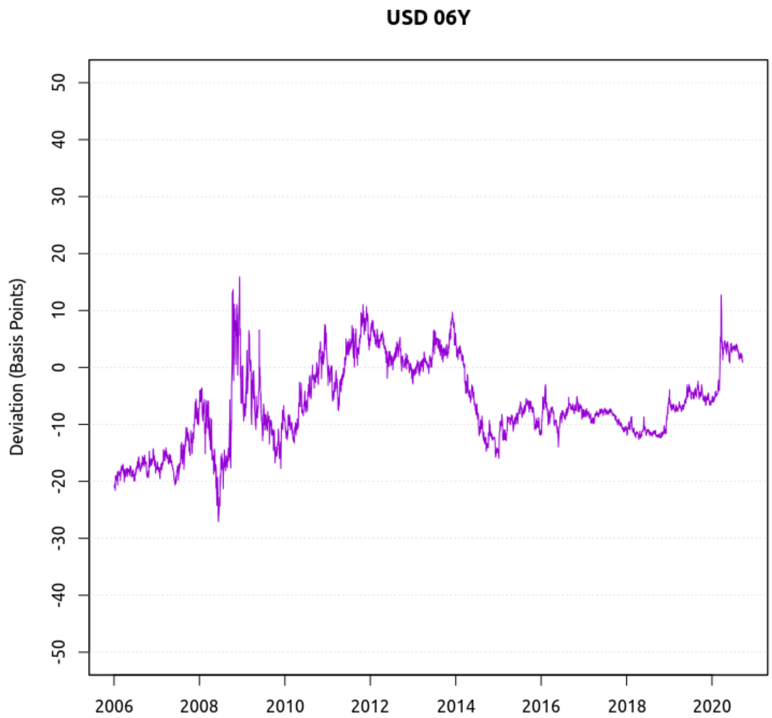

Figure 8.17

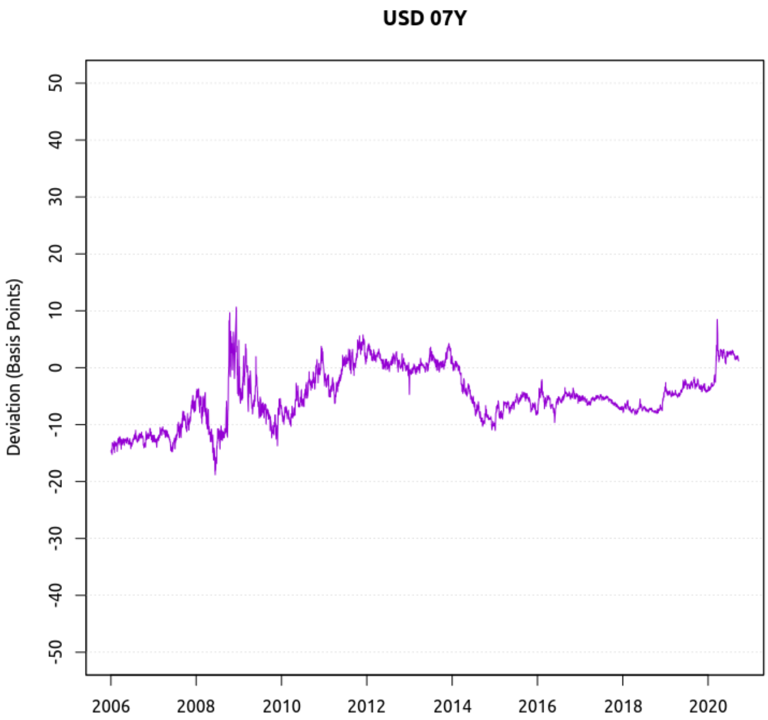

Figure 8.18

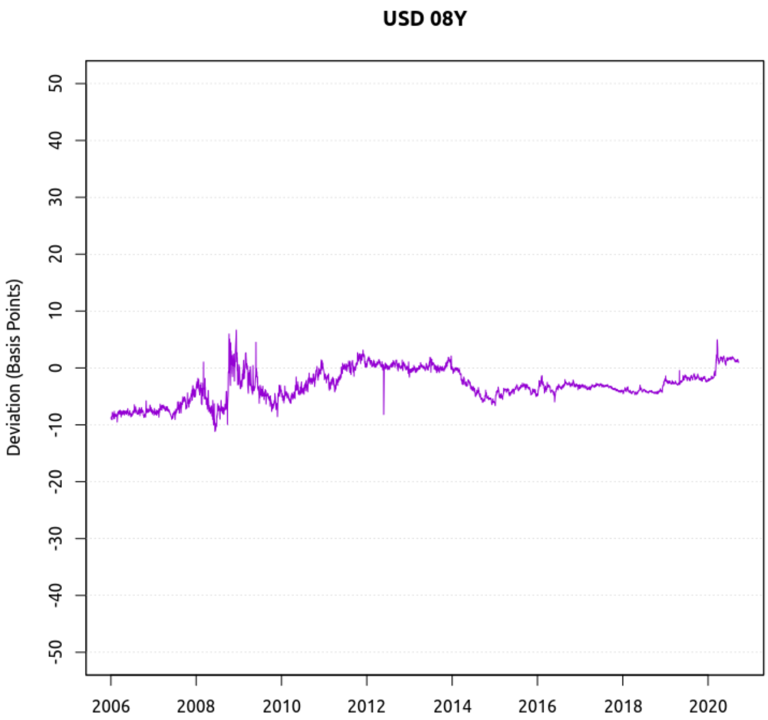

Figure 8.19

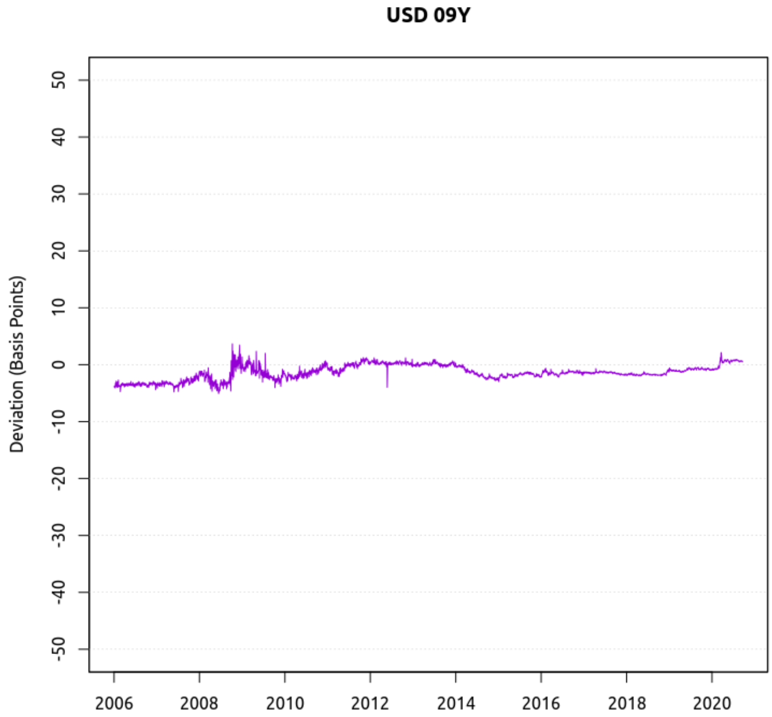

Figure 8.20

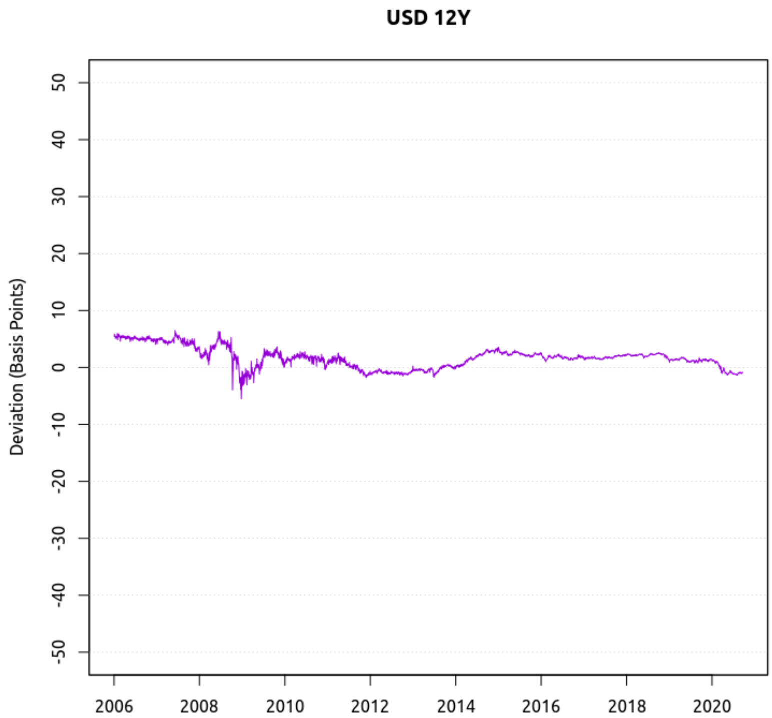

Figure 8.21

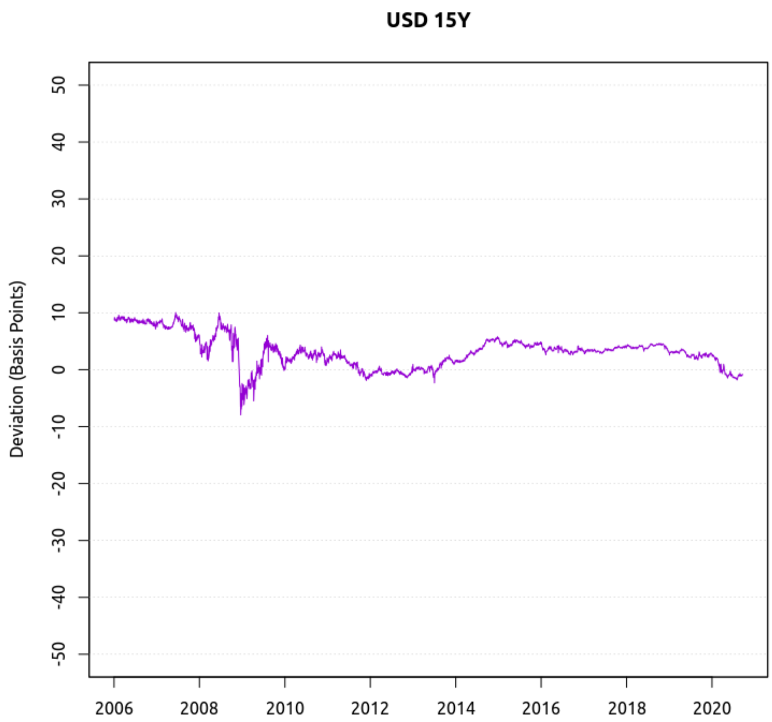

Figure 8.22

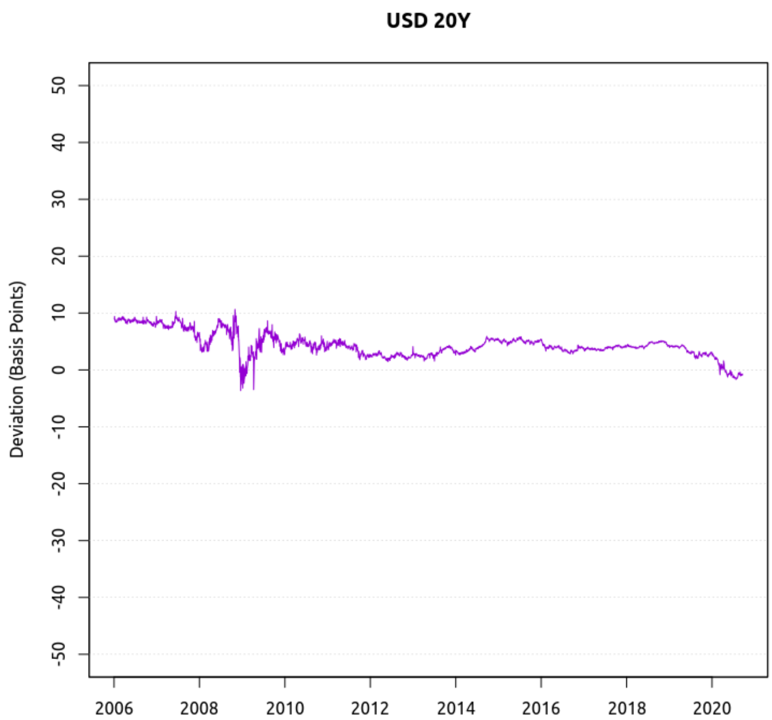

9 References

Agarwal, V., Fung, W., Loon, Y., & Naik, N. (2011). Risk and Return in Convertible Arbitrage. *Journal of Empirical Finance 18*(2), 175-194.

Barclays Hedge (2020). *Hedge Fund Industry Assets Under Management*. Retrieved on Dec. 18, 2020.

Baygün, B., Showers, J., & Cherpelis, G. (2000). Principles of Principal Components. *Salomon Smith Barney, United States Fixed-Income Research*.

Brigo, D. & Mercurio, F. (2006). *Interest Rate Models: Theory and Practice*. Springer.

Domanski, D., Shin, H. S., & Sushko, V. (2017). The Hunt for Duration: Not Waving but Drowning? *BIS Working Paper No. 519*.

Duarte, J., Longstaff, F. & Yu, F. (2007). Risk and Return in Fixed-Income Arbitrage: Nickels in Front of a Steamroller? *The Review of Financial Studies 20*(3), 769-811.

Fabozzi, F., Martellini, L., & Priaulet, P. (2005). Predictability in the Shape of the Term Structure of Interest Rates. *The Journal of Fixed Income 15*(1), 40-53.

Gatev, E., Goetzmann, W., & Rouwenhorst, G. (1999). Pairs Trading: Performance of a Relative-Value Arbitrage Rule. *NBER Working Papers 7032*.

Gatev, E., Goetzmann, W., & Rouwenhorst, G. (2006). Pairs Trading: Performance of a Relative-Value Arbitrage Rule. *The Review of Financial Studies 19*(3), 797-827.

Litterman, R. & Scheinkman, J. (1991). Common Factors Affecting Bond Returns. *The Journal of Fixed Income 1*(1), 54-61.

Lowenstein, R. (2001). *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House.

Mallaby, S. (2010). *More Money Than God: Hedge Funds and the Making of a New Elite*. Penguin Books.

Mitchell, M. & Pulvino, T. (2001). Characteristics of Risk and Return in Risk Arbitrage. *The Journal of Finance 56*(6), 2135-2175.

Mitchell, M. & Pulvino, T. (2012). Arbitrage Crashes and the Speed of Capital. *Journal of Financial Economics 104*(3), 469-490.

Patton, A. (2009). Are “Market Neutral” Hedge Funds Really Market Neutral? *The Review of Financial Studies 22*(7), 2295-2330.

PIMCO (2020). *Understanding Investing – Interest Rate Swaps*. Retrieved on Oct. 10, 2020.

Pontiff, J. (2005). Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics 42*(1-2), 35-52.

Shleifer, A. & Vishny, R. (1997). The Limits of Arbitrage. *The Journal of Finance 52*(1), 35-55.

Storn, R. & Price, K. (1997). Differential Evolution. *Journal of Global Optimization 11*, 341-359.

Tuckman, B. & Serrat, A. (2011). *Fixed Income Securities: Tools for Today’s Markets*. Wiley.

Vasicek, O. (1977). An Equilibrium Characterization of the Term Structure. *Journal of Financial Economics 5*(2), 177-188.

Viale, A., Kolari, J. W., & Fraser, D. R. (2009). Common Risk Factors in Bank Stocks. *Journal of Banking & Finance 33*(3), 464-472.

Wiersema, U. (2008). *Brownian Motion Calculus*. Wiley.