Implied Skewness and the Cross Section of Foreign Exchange Returns

Ákos Török *

December 9, 2022

ABSTRACT

The paper focuses on the relationship between the carry factor and option-implied skewness from a currency portfolio perspective. I show that it is possible to build 'crash-hedged carry' as well as a 'skew-neutral carry' strategies with attractive risk-return profiles and diversification properties, solely by adding option-implied skewness as an additional signal to interest rates. Since these two signals are correlated, I introduce a portfolio-construction technique that solves the problem of double-sorting a low number of assets on correlated signals. Both 'crash-hedged carry' and 'skew-neutral carry' exhibit positive realized skewness, which contradicts skew-based explanations (e.g. crash-risk) of the carry premium. Also, they offer similar realized returns and an improved Sharpe ratio compared to the traditional carry factor.

JEL Codes: F31, G12, G15

*akos@toeroek.org Vienna University of Economics Business & ZZ Vermögensverwaltung; I would like to thank Nóra Bácsfalvi, Georg Cejnek, Stephan Kranner, Semyon Malamud, Andreas Neuhierl, Igor Pozdeev, Daria Prem, Otto Randl, Angie Rigatou, Alina Steshkova, Stefan Vincenz, Christian Wagner, and Josef Zechner for insightful comments and much appreciated guidance.

I. Introduction

Although skewness and interest rates have often been connected in literature, in particular to show that carry returns are associated with crash risk, a simultaneous use in portfolio construction has not been attempted yet. This paper aims to fill this gap. To better illustrate the motivation of this paper, consider Figure 1 which shows a FX-carry strategy and a strategy going long currencies with high implied skewness and going short currencies with low implied skewness[1]. The two performances are strikingly similar (the monthly returns have a correlation coefficient of 0.87), with carry constantly outperforming the skewness-ranked portfolio by a slight margin. I suspect that this is the consequence of a disparity between the quality of the two signals with respect to downside risk. As the two signals are highly, but not perfectly correlated, there exist currencies that are ranked high according to interest rates but at the same time have a comparatively low ranking when it comes to implied-skewness and vice versa. To further illustrate this case I show how both implied and realized skewness relate to mean excess returns using portfolios constructed from carry and option-implied skewness signals (Figure 2). Both sets of portfolios offer higher returns for more negative skewness. However, the carry portfolios does have higher returns for less realized negative skewness and is also less exposed to crash-risk ex-ante (as proxied by the average option-implied skewness of the portfolios). Existing literature has inferred from the negative relationship of excess returns and skewness that carry portfolios compensate crash-risk. The portfolios constructed from option-implied skewness indicate that implied-skewness is a better signal for future realized crash risk than interest rate differences and consequently, that crash-risk is not the risk that is the reason for the risk premium of FX-carry.

In this paper, I show that by combining both interest rates and option-implied skewness as signals for portfolio construction, one can extract a portfolio of high-yielding, low-skew currencies and a portfolio of low-yielding, highly skewed currencies. These two are similar to the 'corner' portfolios in a double-sort. However, since it is hard to double-sort a low number of

[1] See Formula 6 in Section IV. Positive implied skewness is very rare and can be observed only in a few exchange rates like e.g. USDJPY. Thus, when I refer to skewness being high, throughout the paper I mean currencies which are strongly left-skewed, i.e. exhibit high negative skewness.

assets according to two highly correlated characteristics, I introduce a weighting scheme based on rank differences that is better suited to construct a portfolio that isolates the 'difference' of the two characteristics. This weighting scheme can be interpreted as the equivalent of a double-sort for a rank-based weighting scheme. The two portfolios can be combined to form a 'skew-neutral carry' portfolio where the low-yielding, highly skewed currencies are shorted and high-yielding currencies which are hardly skewed, are bought. The short portfolio can also be used as a short in combination with the long-leg of the traditional carry-factor to produce a 'crash-hedged carry' portfolio. Such a 'crash-hedged carry' portfolio not only offers the same return as the traditional carry factor, but also exhibits lower volatility and a positive realized skewness coefficient. This empirical result contradicts crash-risk as an explanation for the carry risk premium. The paper goes on to show that these new portfolios are not exposed to systematic downside risk and when used as factors in a linear pricing model they are priced whereas a traditional carry-factor is not. Lastly, I show that the two new portfolios add value in a portfolio context when added to existing factors, especially, when investors care about skewness. For this, ex-post efficient frontiers are calculated both in a simple mean-variance and also in a three dimensional mean-variance-skewness setting.

II. Literature Review

The foreign exchange market represents one of the biggest asset classes with an estimated turnover of USD 6.6 trillion per day (Schrimpf and Sushko, 2019). Yet, for a long time this asset class was comparably less researched than its bond or stock market counterparts. Early research starting with Hansen and Hodrick (1980), Meese and Rogoff (1983) and subsequently, Fama (1984) established, that there seemed to be very little connection between currency movements and macroeconomic variables, most importantly, interest rates. Macroeconomic theory would predict that currencies depreciate by the difference in their interest rates, but these early studies failed to confirm this prediction empirically, leading to the so-called uncovered interest rate parity puzzle or forward premium puzzle. This apparent puzzle was meanwhile exploited by investors world-wide, who recognized that they could earn positive returns on average by borrowing money in low-interest currencies and investing it in high-yielding currencies.

However, sustained returns to an investment strategy are hard to justify in asset pricing if they do not represent compensation for some sort of risk that is borne by investors. Consequently, asset pricing literature started to focus on finding what the associated risks of the carry-trade are and a handful of solutions have been proposed, such as consumption risk (Lustig and Verdelhan, 2007), peso problems (Burnside et al., 2011), volatility risk (Menkhoff et al., 2012a) exposure to equity downside-risk (Lettau et al., 2014) and crash risk related to sudden changes in liquidity (Brunnermeier et al., 2008). Also, comparatively late, researchers started to think of foreign exchange in a portfolio context, and Lustig et al. (2011) created a carry-factor as well as a dollar-factor, with which they were able to explain the majority of return variation of portfolios sorted on interest rate differences. This paper started a strand of FX-literature taking a portfolio view that proposed new risk-factors in the FX-space, some of which where known from other asset classes, such as value (Menkhoff et al., 2017) or momentum (Menkhoff et al., 2012b).

Still, the most prominent risk-factor which has persistently earned excess returns is carry. I intend to contribute to the literature by connecting carry to what is in my opinion the single best ex-ante measure of downside-risk, namely option-implied skewness. This has been done before. In particular, some papers argue that carry is compensation for disaster risk (e.g. Brunnermeier et al. (2008), Farhi et al. (2009), Burnside et al. (2011)) and show the strong connection between implied skewness and interest rate differentials. Rafferty (2012) even creates a skewness factor (based on realized skewness) and shows that such a factor can price portfolios sorted on interest rate differences as well as currency momentum and value portfolios. However, instead of just showing an empirical connection between carry and option-implied skewness, I will show that by combining interest rates and option-implied skewness as signals, carry strategies that have seemingly no disaster risk can be isolated.

This paper brings together the FX-carry literature with the literature on option-implied information. Option-implied information is exciting to researchers and practitioners alike as it is forward-looking. Thus, it is much more reactive to new information as any news should instantly be priced in options. This is an advantage compared to historical estimates where news only affect the latest of many data points used in estimation. In their seminal paper Breeden and Litzenberger (1978) show how one can derive state-contingent claims from a set of options with a continuum of strikes and how these translate to a risk-neutral distribution.

From this ($\mathbb{Q}$-)distribution moments can be calculated. Risk-neutral distributions combine the market's views on both its best estimate on true probabilities as well as its risk preferences. However, methods proposed to disentangle these two (Ross, 2015) have been found to be based on assumptions that are too strong (Jackwerth and Menner, 2018). A nice overview of applications for option-implied information can be found in Christoffersen et al. (2013). One application that followed risk-neutral densities is the calculation of risk-neutral moments. Other papers that use option-implied information in FX asset pricing are Della Corte et al. (2016), Mueller et al. (2017) and Bang Nielsen (2018). For equities, Schneider et al. (2020) show that coskewness, proxied by option implied skewness can explain low-risk anomalies. I find that this connection between implied-skewness and coskewness also exists in FX-markets but rather than explaining an existing anomaly, this information can be used to create portfolios which contrary to conventional wisdom earn a carry risk premium with no systematic downside risk. Thus, the long-established fact that investors should care about coskewness (Kraus and Litzenberger (1976), Harvey and Siddique (2000)) not the source of the FX carry premium. This observation contradicts existing literature (Dobrynskaya, 2014) which previously connected risk premiums in FX-carry and equity markets through coskewness/systematic disaster risk.

No paper explores portfolios based on disparities between the signals given by implied skewness and interest rate differentials. The closest paper to this one is Jurek (2014) who shows that one can construct an efficiently tail-hedged portfolio out of currency forwards and options that still earns positive returns. His results suggest that either only a small part of carry returns is due to crash risk[2] or FX-options are mispriced. In contrast to Jurek (2014), I show that one can construct both a skew-neutral carry (RDF) as well as a crash-hedged carry (CAR_{hedged}) strategy without the use of options, solely by combining option-implied skewness and interest rates as signals. Other than its simplicity regarding the instruments, my approach also has the advantage of an increased currency universe. For the calculation of option-implied skewness I only need data on options of each currency's exchange rate against the USD, whereas for efficient hedging with options one would need all crosses in the currency universe. This is a bottleneck which effectively constrains the investment universe to just G10 countries.

[2]The author states that at most one third of the excess return of carry is compensation for crash risk. A number that aligns very well with the positions of CAR_{hml} and SKW_{hml} in Figure 2.

III. Data

The basis for my research is data on foreign exchange spot prices as well as derivatives on exchange rates. It is standard in literature to compute currency returns from FX-Forwards as the prices of these derivatives resemble borrowing in a base currency (USD), exchanging it into a foreign currency, and investing it at the foreign interest rate due to no-arbitrage considerations. Consequently, returns calculated as

$$R_{t+n} = \frac{S_{t+n}}{F_{t,t+n}} - 1. \tag{1}$$

are excess returns. In the above formula S_t is the spot exchange rate at time t of US-Dollars per one unit of foreign currency and $F_{t,t+n}$ is the forward exchange rate at time t for $t+n$. The most common forward maturity used in literature is one month, which is the one I use. Traders mostly quote forwards in 'points' which are usually given (some FX-pairs have different conventions) in $\frac{1}{10000}$ units of the currency that is the numerator in a given pair. This means that $F_{t,t+1}$ in the above equation has to be computed by adding these points to the respective spot rates S_t of the currencies[3].

The second type of derivative needed for my analysis are FX-options. As opposed to stock options which are exchange-traded, have a price quoted in currency units, and have strike prices as references, FX options are traded over-the-counter and follow different conventions. They are quoted in Garman and Kohlhagen (1983) implied volatilities and referenced by delta. More specifically, implied volatility data for FX-options can be obtained via direct quotes for at-the-money options but has to be calculated for out-of-the-money options from so-called risk-reversals and butterflies. Risk-reversals quote the difference between implied volatilities of out-of-the-money calls and puts, whereas butterflies quote the average difference between the implied volatilities of puts and calls with a given delta and at-the-money options. Risk-

[3] e.g. a trader might trade a 1-year forward on the USDZAR exchange rate and see a quote of 7000 points. To get to the actual outright exchange rate in one year he has to divide these points by 10000 and add them to the current spot rate. With a spot rate of 15, this would mean a forward exchange rate of $15 + \frac{7000}{10000} = 15.7$ South African Rand per US-Dollar. As we are interested in the FX expressed as USD per ZAR we have to calculate $\frac{1}{15.7}$ to arrive at $F_{t,t+12}$.

reversals and butterflies are usually quoted for deltas of 10% and 25% with which a volatility smile from a put to a call with delta of (-)10% can be spanned. The implied volatilities for (-)10/25% delta puts and calls can be calculated from butterflies and risk-reversals as follows:

$$IV_{call,\delta} = BF_\delta + ATM + \frac{RR_\delta}{2} \tag{2}$$

$$IV_{put,\delta} = BF_\delta + ATM - \frac{RR_\delta}{2} \tag{3}$$

This smile can be transformed to implied volatilities per strike prices for the given deltas which is a necessary step for later calculations. One has to account for various conventions like different premium-adjustments of deltas for different currencies. This is the case since the premiums for different currencies might be paid in domestic or foreign currency depending on the convention for a given exchange rate. Also, conventions might differ as to what is referred to as the at-the-money strike for a given pair.[4]

There are multiple ways to 'connect' the implied volatility dots to obtain a smooth function of implied volatility per strike, the so-called volatility smile. The simplest approach is to use cubic splines to interpolate between the five points. Since the volatility smiles calculated from the data do not span the whole spectrum of strike prices, one has to make assumptions if the smile is to be extended beyond the strikes corresponding to (-)10% delta. The arguably most simple assumption is that the tails of the smile are left flat beyond deltas of (-)10%. Although this approach is not suitable for option-traders, it is common in academic literature as it is easy to replicate and is 'good enough' for most applications. As I intend to do relative comparisons of options smiles (in particular, their skewness) and use the data mostly for ranking purposes, the said approach is appropriate in this context.

Option data are the bottleneck of the empirical work both in regards to time series and cross-section of data. Whereas records for spot and forward data can go back to the 1970s for

[4]The ATM strike price might refer to the current spot rate, forward rate, or the strike price for which a straddle of a put and call option are delta-neutral, depending on the currency pair and maturity of an option. An exhaustive list of conventions and the resulting formulas to retrieve strike prices from deltas can be obtained from text books. For references regarding conventions I can recommend Clark (2011) or Wystup (2017).

various currencies, FX-option data mostly starts in the mid-nineties or even after the financial crisis when it comes to developing markets. Option data is also limited in regards to the cross-section of currencies, in particular, when it comes to more exotic currencies. Generally, emerging markets with more developed capital markets have liquid options-markets (think e.g. Brazil, Russia, South Africa) whereas frontier markets are not covered at all. All in all, these bottlenecks result in a sample of 31 currencies quoted against the USD which can be divided into 9 developed and 22 emerging market currencies. For the calculation of option-implied skewness (as defined in the next section) a whole volatility smile is needed. The availability of full (I require (-10%)-delta option data) volatility smiles constrains the start of my empirical work with the year 2006 (Figure 3). To ensure consistency, Bloomberg (pricing source 'BGNL', 'CMPL' if the former is missing) is used as the single data source for monthly spot, forward and option-implied volatility (in the form of ATM, risk-reversal, and butterfly) data.

IV. Methodology

A. Signals

Two signals have to be extracted from the data. For carry, this is simple as the signal is the domestic interest rate of each country. If we take the view of an US-investor and set the USD as the base currency for all currencies, then we can use the (1-month) interest rate differential priced in the forward $CAR^i_{t,t+1} = \frac{S^i_t}{F^i_{t,t+1}} - 1$ as the signal. If all currencies i are expressed in terms of USD those with the highest interest rates have the highest interest rate differential as the domestic interest cancels out[5].

The second signal is option-implied skewness. For option-implied moments, model-free estimation has become the norm. The first to propose formulas for model-free estimation of option-implied moments were Bakshi et al. (2003). Subsequently, slightly different approaches have been suggested by Kozhan et al. (2013) and Schneider and Trojani (2015). The approaches essentially differ in how options with different moneyness are weighted. I choose the approach

[5]If a currency is quoted as $\frac{USD}{XXX}$, and the interest is higher in the foreign country i than in the US, than $F^i_{t,t+1} < S^i_t$ which means $CAR^i_{t,t+1} > 0$. This is the amount an investor earns if the spot exchange rate stays the same until expiry.

of Schneider and Trojani (2015) which is given by the following formula(s):

$$upperSKEW^{\mathbb{Q}}_{t,t+1} = \frac{6}{p_{t,t+1}} \left(\int_{F_{t,t+1}}^{\infty} \log\left(\frac{K}{F_{t,t+1}}\right) \frac{\sqrt{\frac{K}{F_{t,t+1}}} C_{t,t+1}(K)}{K^2} dK \right) \tag{4}$$

$$lowerSKEW^{\mathbb{Q}}_{t,t+1} = -\frac{6}{p_{t,t+1}} \left(\int_{0}^{F_{t,t+1}} \log\left(\frac{F_{t,t+1}}{K}\right) \frac{\sqrt{\frac{K}{F_{t,t+1}}} P_{t,t+1}(K)}{K^2} dK \right) \tag{5}$$

$$SKEW^{\mathbb{Q}}_{t,t+1} = upperSKEW^{\mathbb{Q}}_{t,t+1} + lowerSKEW^{\mathbb{Q}}_{t,t+1} \tag{6}$$

This method is the most stable and arguably the best suited for real-life circumstances. The reason is that weights for out-of the money options are very small. Thus, how the volatility smile is extrapolated only has a negligible on the estimation and relative ranking of currencies. In any case, since I calculate skewness for the sole purpose of ranking different currencies, my results remain stable when using different approaches. The skewness measure coming out of the above formula is not standardized. Standardized skewness (also referred to as the skewness coefficient) would isolate the 'tailedness' of a distribution from its volatility by dividing by the cube of the standard deviation, whereas the skewness measure used in this paper scales with volatility. This is desirable, as for construction purposes, I need a single measure of (downside-)risk and it makes sense to keep the influence of volatility in the signal. The integral is calculated via a trapezoidal approximation.

B. Portfolio construction

The goal of my work is to isolate a portfolio that trades the 'difference' of two signals, namely carry and skewness. The traditional approach would be to do a double-sort like for example Fama and French (1992) do for size and value. This is called an unconditional double-sort. Quantile-portfolios are constructed across signals, in their case market capitalization and book-to-market. This approach works best for uncorrelated signals and when there are a high number of assets. It is obvious why: One can think about the portfolios by drawing a scatter plot which has the signals along its two axes. For instance, if we do a 3x3 unconditional double-sort, the plane spanned by the x- and the y-axis is divided into 9 rectangles that have to be populated. I draw such a plot for my case in Figure 4. It is comparatively easy to form portfolios from all rectangles if the two signals have a low correlation (size and value) and there are many assets (US equities). This increases the chances of all rectangles being populated. However, it gets harder to populate all rectangles if one has a low number of assets (31 currencies), which are spanned along two signals that are highly correlated, like in the case of implied skewness and carry.

In my case both problems (very high correlation of the signals and low amount of assets) are present. An 3x3 unconditional double-sort results in very concentrated 'corner portfolios' (low skewness, high interest, and high skewness, low interest) which are of particular interest if one wants to construct skew-neutral carry portfolios. The average number of currencies in each corner portfolio is about one. Such concentrated portfolios would result in a very volatile performance. To get corner portfolios with a reasonable amount of currencies (4 to 5, on average) one has to resort to a 2x2 sort. Alternatively, a conditional 2x2 double-sort guarantees that each of the 4 portfolios is populated by a quarter of the assets (see Figure 5). However, these solutions are less than ideal as many of the chosen currencies are very close to the bounds of the corner portfolios. Thus, they do not offer big discrepancies between carry and skewness but still enter into the portfolios with equal weight.

My solution is to focus on the differences in ranks according to the two signals and make the portfolio weights a direct function of these rank differences. I propose the following formula to form orthogonalized ('skew-neutral carry') portfolios. Weights for the rank-difference (RDF)-portfolios are calculated so that each currency i at time t has a weight such that

$$w_{i,t}^{RDF} = \kappa_t \left[\text{rank}(CAR_{t,t+1}^{i}) - \text{rank}(-SKEW_{t,t+1}^{\mathbb{Q},i}) \right] \tag{7}$$

where κ_t is a scaling factor ensuring that both the long and short investments sum to (minus) one. This scheme weights currencies proportional to their rank differences and thus overweights those with high differences in signals. As a result, the strategy should be dominated by high-yielding currencies with limited downside risk on the long side and low-yielding currencies with higher downside risk on the short side. Table I details the weight calculation using the example of March 2020. For a graphical illustration see Figure 6.

Other portfolios, which are constructed from a single signal like those shown in Figure 1 or the ones used in asset pricing tests, are constructed using rank-based weighting as introduced in Asness et al. (2013) and defined by the following formula:

$$w_{i,t} = \kappa_t \left[\text{rank}(z_{i,t}) - \frac{1}{I_t} \sum_{i=1}^{I_t} \text{rank}(z_{i,t}) \right] \tag{8}$$

where κ is again a scaling factor to ensure the sum of weights equal one (-1 for the short positions) and $z_{i,t}$ is a signal for currency i at time t. This method is chosen since it assigns the currencies with the strongest signals the highest (most negative, for shorts) weights. In essence, the weighting scheme introduced in equation 7 can be interpreted as the rank-based equivalent of a double-sort according to quantiles.

Rank-difference based weighting has advantages apart from over-weighting currencies with particularly high discrepancies. Firstly, portfolios are not constrained to a particular rectangle (graphically speaking, see Figure 5), but also currencies that have high discrepancies outside of the (coloured) rectangles can enter into portfolios. This results in more potential assets that can enter the portfolios. All currencies that do not fall directly on the diagonal (Figure 6) enter into either the long or short portfolio. There is one more advantage of the weighting methodology specifically for the crash-hedged carry portfolio: Some currencies that have high-interest rates and therefore enter prominently into the rank-weighted carry trade can be (partially) neutralized by the short-leg based on rank differences. This happens if their skewness ranking is even higher than warranted by the interest rate. The portfolios constructed using rank-difference weighting produce better (co)skewness statistics for skew-neutral and crash-hedged carry strategies than applying a double-sort (as can be seen by comparing the statistics in Tables V and XII). The double-sorted skew-neutral portfolio also improves (co)skewness statistics compared to its carry counterpart but not by as much as its rank-based counterpart. Summary statistics of the weights of various portfolios in IV show that the rank-difference based 'skew-neutral carry' (RDF) is slightly more concentrated and exhibits higher turnover than carry but lower turnover than momentum.

As a comparison, I also look at an alternative weighting methodology: Average-rank-based weighting is used in Fisher et al. (2015) for combining value and momentum stocks. Here average ranks are calculated from two signals and portfolios are constructed from these average weights. Fisher et al. (2015) construct value-weighted quantile portfolios from these average ranks. As value-weighting is not straightforward for FX, I do two versions: a skew-neutral long-short quintile portfolio where currencies are equally weighted and a skew-neutral rank-weighted portfolio based on average ranks. The average ranks are calculated as $\left[\text{rank}(CAR^i_{t,t+1}) + \text{rank}(SKEW^{\mathbb{Q},i}_{t,t+1})\right]/2$ [6]. A graphical illustration of the weighting scheme is shown in Figure 10 and summary statistics of resulting portfolios are shown in Table XIII. Although the resulting portfolios also show reduced skewness, they are inferior to double-sorted and rank-difference-based portfolios when it comes to mean returns.

V. Empirical Results

This section focuses on the empirical properties of 'skew-neutral carry' (RDF) and 'crash-hedged carry' (CAR_{hedged}) portfolios that are constructed using the weighting scheme described in Formula 7 in the previous section. RDF is comprised of long and short legs which are both constructed using rank-difference based weighting, whereas CAR_{hedged} combines the rank-difference based short leg with the long leg of the rank-weighted carry portfolio (for illustrations of weight construction see Figure 6).

Table II shows summary statistics for the weights of the RDF and traditional carry (CAR) strategies. In contrast to CAR, no currency is persistently in the long or short part of the RDF strategy. Only INR has a positive weight more than 90% of the time and it also exhibits by far the highest average weight (14.2%). For the shorts of RDF there are four currencies (CZK, HUF, PLN, SEK) with a negative weight proportion of more than 90%. One can see that most 'classic' carry currencies do not have a clear tendency of being a long or short currency with respect to RDF or exhibit small average weights, suggesting, that the two strategies are unrelated.

Table III presents summary statistics for the returns of the new RDF and CAR_{hedged} strategies as well as other FX-strategies from literature. Panel A summarizes the excess returns of all portfolios and reveals that the period of 2006 until 2022 has not been a great period for most existing FX-strategies. The positive exception is carry which has produced a Sharpe ratio of 0.48. Both the average returns of RDF and CAR_{hedged} are greater than zero with

[6] Note, that unlike in Formula 7 I drop the minus sign in front of $SKEW^{\mathbb{Q},i}_{t,t+1}$. Since currencies with high average rank will be bought, the best would be if these currencies had high interest rates and positive skewness.

2.8% for RDF and 3.2% for CAR_{hedged}. Thus, the return of CAR_{hedged} is about the same as that of the unhedged carry factor. Crucially, CAR_{hedged} delivers on its promise of hedging the crash risk of the carry returns with a positive skewness coefficient. Also, the Sharpe ratio of CAR_{hedged} is 0.12 higher than that of traditional carry due to lower realized volatility. Panel B summarizes the spot returns of the strategies and shows that CAR_{hedged} loses 1.4% less on the spot component than the unhedged CAR portfolio. Conversely, this means that also the ex-ante interest rate difference of the CAR_{hedged} strategy is lower. Nonetheless, it seems to be worth to sacrifice some interest rate difference ex-ante, as with the changed currency composition the spot returns of CAR_{hedged} lose their negative skew. CAR_{hedged} exhibits a positive correlation of 0.79 with the classic CAR due to their shared long leg as can be seen in panel C. On the other hand, RDF is almost completely uncorrelated to CAR. This is particularly interesting as the statistics suggest that similarly to CAR and CAR_{hedged}, RDF also earns money by exploiting interest rate differences, albeit smaller ones than the other two strategies. RDF produced a 0.54 Sharpe ratio over the 16-year period which is also is higher than traditional carry. RDF also exhibits no negative skew and is either negatively or hardly correlated with existing FX-strategies (highest correlation is to value with a coefficient of 0.25).

The returns from taking advantage of discrepancies between skewness and carry are short-term, much like carry itself. This can be seen in Figure 9 which shows the annualized mean return for RDF with different holding periods (implemented with FX-forwards with corresponding maturities). Apart from the 1-month strategy no variant produces significantly positive returns. Also, considering the signal, short-term options (1M) work best for construction purposes although RDF portfolios constructed from longer term options are similar in that they also exhibit almost no realized skew (Table XIV) but their returns are 30-40 bps lower p.a. Furthermore, the RDF strategy is robust to the initial sample of currencies (Figure 11). Varying the currency sample by drawing 27 currencies randomly from the full sample, we notice that the full sample is close to the mode of the distribution of mean returns and skewness coefficients.

Table V takes apart the long and short parts of CAR and RDF (and thus CAR_{hedged} too) and shows summary statistics for each component. RDF_{short} has a more negative skewness coefficient than the regular CAR_{short}. Also, it has a higher (more negative when shorted) correlation with CAR_{long} than CAR_{short}. A similar story applies to various coskewness measures to the Mkt-factor of Fama and French (1992). Coskewness can be fully neutralized in both RDF and CAR_hedged strategies whereas the traditional Carry strategies has negative coskewness with the Mkt-factor.

To more formally show the crash exposure of Carry (CAR), skew-neutral (RDF) and hedged Carry (CAR_{hedged}) portfolios I run three regressions proposed in previous literature. The first one proposed in Lettau et al. (2014) uses only negative observations of the Mkt-factor (as downloaded from Professor French's website). The second one uses the Mkt-factor and its squared error terms e^2_{Mkt} as regressors and is the model used to calculate coskewness in the paper of Kraus and Litzenberger (1976). Thirdly, I regress the portfolio returns on changes in the CBOE VIX Index. Brunnermeier et al. (2008) argue that increases in the VIX are associated with a rise in global risk aversion and worsening funding conditions which leads to unwinds in carry trades. The results are shown in Table VI. All three regressions show similar results: The beta coefficient starts out as significantly positive to (downside) equity market risk for CAR, becomes insignificant for CAR_{hedged} and ends up negative for RDF (insignificant for the downside beta). For ΔVIX the sign changes but the results are the same as in the other two regression specifications: CAR has significant exposure to a form of crash risk whereas $CARhedged$ does not and RDF seems to even be negatively exposed to crash risk. Also, the explanatory power of all regressions is lower for the new factors than for CAR.

Having established that RDF and CAR_{hedged} exhibit better coskewness measures to the stock market than the simple CAR factor, I also investigate if existing FX factors can explain RDF in table VIII. Right-hand side variables are additionally to carry (CAR and the USD-factor (as proposed in Lustig et al. (2011)), momentum (Menkhoff et al., 2012b), value (Asness et al., 2013), and the volatility risk premium (Della Corte et al., 2016). Although all factors have significant beta coefficients to the RDF portfolio, RDF offers significant alpha to all of them individually and collectively. Also, almost half of RDF remains unexplained (in terms of R-squared) by the regression including all factors. This alpha is large as it is more than half of RDF's return and it is higher than the return of all factors except CAR in the sample period. The same exercise is done for the short part of RDF as it is used in the construction of CAR_{hedged}. The results are shown in IX. The short side of RDF is well explained (R-squared

> 0.9) by the other risk factors due to its directional exposure to the USD-factor as well as exposure to the CAR and VRP factor. The short part of RDF also exhibits an alpha which roughly equates to half the size of RDF. Contrary to the long/short RDF's alpha, it is not statistically significant in the regression including all factors but significant when VRP is excluded. A classic asset pricing graph of predicted versus mean returns for all three models is shown in Figure 8.

Now I turn to investigate RDF's and CAR_{hedged}'s use in pricing FX portfolios, I estimate asset pricing models of the form

$$E_t[M_{t+1}R^i_{t+1}] = 0 \tag{9}$$

with a linear stochastic discount factor $M_{t+1} = 1 - b(f_{t+1} - \mu)$, where b is a vector of factor loadings and μ denotes the mean of the factors f. This model implies a beta pricing model where the factor prices of risk λ and the sensitivities to the factors β^i determine a portfolio's expected return $E[R^i] = \lambda' \beta^i$. To obtain the λ, I estimate b so that equation 9 holds using the Generalized Methods of Moments of Hansen (1982) and calculate the λ via the identity $\lambda = \Sigma_{ff} b$[7]. I estimate three models and show the results in Table VII. Firstly, I use USD and CAR as factors. This is the model estimated in Lustig et al. (2011). Secondly, I swap CAR for RDF, and finally, I swap CAR for CAR_{hedged}. As test assets, I use quintile portfolios CAR_i, MOM_i, VAL_i, VRP_i, as well as RDF_{long} and RDF_{short}. All models estimate the price of risk λ of CAR, RDF, and CAR_{hedged} roughly equal their mean returns. However, the estimated standard errors for λ_{RDF} and $\lambda_{CAR_{hedged}}$ are much smaller than that of λ_{CAR}. This results in λ_{RDF} and $\lambda_{CAR_{hedged}}$ being significantly positive while λ_{CAR} is not statistically different from zero. Also, the point estimate of λ_{CAR} is lower than the estimate of Lustig et al. (2011), indicating that the carry risk premium has declined recently. Estimates of b are significant regardless of which of the factors is used in the model. A graph of predicted versus mean returns for the 22 test assets is shown in Figure 8. All three models perform well in explaining the test assets. The third model using USD and CAR_{hedged} as factors offers the best fit both in terms of R^2 and root-mean-square error by a slight margin.

[7] see Cochrane (2009), chapter 13.2 for details

The low correlations of both RDF and CAR_{hedged} with the other factors suggest that these new strategies might be a sensible addition from a portfolio view. To analyze the value of the two new strategies in a portfolio context, I compute the ex-post optimal currency strategy for combinations of FX-risk factors which are detailed in Figure 7 as well as Table X. Specifically, I look at two options: Swapping CAR for CAR_{hedged} and adding RDF to this swap. The new strategies move the efficient frontier resulting in a 0.13 increase in Sharpe ratio (0.01 if only the CAR_{hedged} is used instead of CAR). Thus, in a mean-variance context considerable value is added by the RDF portfolio. However, if skewness is considered as an extra dimension, the new strategies add even more value in terms of Sharpe-ratio. To illustrate this I simulate weights and plot the efficient frontiers with the condition of the realized skewness coefficient being equal or higher than 0.1. This equals slicing the mean-variance-skewness space exactly at the skewness coefficient of 0.1. Swapping CAR for CAR_{hedged} adds 0.13 to the Sharpe-ratio of the mean-variance optimal portfolio and adding RDF adds another 0.02. The details of the skewness-constrained ex-post optimal strategies can be seen in Table XI. The improvement in the skewness coefficient-restricted efficient portfolios is caused by the fact that the CAR_{hedged} strategy has a much better skewness profile than CAR. RDF in turn dominates the weights of the optimal portfolio when added because of its favorable correlation profile with respect to the other factors. Thus, to investors who care about the skewness of their portfolios, adding CAR_{hedged} and RDF to existing factors seems to be worthwhile since they can improve the skewness profile of their portfolios.

VI. Conclusion

The paper finds that although option-implied skewness and forward discounts convey similar information, their less-than-perfect correlation can be used to construct skew-neutral or crash-hedged carry strategies using only FX-Forwards. Both the skew-neutral and crash-hedged carry-trade offer positive skewness coefficients and an increased Sharpe ratio compared to the traditional carry factor while earning money from spot depreciations being smaller than ex-ante interest differences. Additionally, they are unrelated to various measures of systematic downside risk. This strongly suggests that crash risk is not the main source of the carry risk premium. Furthermore, both portfolios offer value when investors care about skewness in their

FX portfolios and in contrast to the carry factor, have a significant price of risk in the sample period from 2006 to 2022. Further efforts need to be made in order to see what risks might be priced in these portfolios.

References

Andrews, D. W. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica: Journal of the Econometric Society*, pages 817–858.

Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013). Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985.

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143.

Bang Nielsen, A. (2018). Forward-looking currency betas.

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of business*, pages 621–651.

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

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

Christoffersen, P., Jacobs, K., and Chang, B. Y. (2013). Forecasting with option-implied information. In *Handbook of economic forecasting*, volume 2, pages 581–656. Elsevier.

Clark, I. J. (2011). *Foreign Exchange Option Pricing: A Practitioner's Guide*, volume 624. John Wiley & Sons.

Cochrane, J. H. (2009). *Asset pricing: Revised edition.* Princeton university press.

Della Corte, P., Ramadorai, T., and Sarno, L. (2016). Volatility risk premia and exchange rate predictability. *Journal of Financial Economics*, 120(1):21–40.

Dobrynskaya, V. (2014). Downside market risk of carry trades. *Review of Finance*, 18(5):1885–1913.

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

Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465.

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

Fisher, G. S., Shah, R., and Titman, S. (2015). Combining value and momentum. *Journal of Investment Management, Forthcoming*.

Garman, M. B. and Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of international Money and Finance*, 2(3):231–237.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054.

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

Harvey, C. R. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of finance*, 55(3):1263–1295.

Jackwerth, J. C. and Menner, M. (2018). Does the ross recovery theorem work empirically? *Available at SSRN 2960733*.

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

Kozhan, R., Neuberger, A., and Schneider, P. (2013). The skew risk premium in the equity index market. *The Review of Financial Studies*, 26(9):2174–2203.

Kraus, A. and Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *The Journal of finance*, 31(4):1085–1100.

Lettau, M., Maggiori, M., and Weber, M. (2014). Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, 114(2):197–225.

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

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

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics*, 14(1-2):3–24.

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

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2012b). Currency momentum strategies. *Journal of Financial Economics*, 106(3):660–684.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2017). Currency value. *The Review of Financial Studies*, 30(2):416–441.

Mueller, P., Stathopoulos, A., and Vedolin, A. (2017). International correlation risk. *Journal of Financial Economics*, 126(2):270–299.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708.

Rafferty, B. (2012). Currency returns, skewness and crash risk. *Skewness and Crash Risk (March 15, 2012)*.

Ross, S. (2015). The recovery theorem. *The Journal of Finance*, 70(2):615–648.

Schneider, P. and Trojani, F. (2015). Fear trading. *Swiss Finance Institute Research Paper*, (15-03).

Schneider, P., Wagner, C., and Zechner, J. (2020). Low-risk anomalies? *The Journal of finance*, 75(5):2673–2718.

Schrimpf, A. and Sushko, V. (2019). Sizing up global foreign exchange markets. *BIS Quarterly Review, December.*

Wystup, U. (2017). *FX options and structured products.* John Wiley & Sons.

VII. Tables

Table I. The table illustrates the weighting scheme according to rank difference given in Formula 7. The data is from March 2020. A graphical illustration can be seen in the middle graph of Figure 6.

FX	CAR	CARrank	SKEW * 10^5	SKEWrank	rankdif	weights
TWD	-7.08	1	-0.11	6	-5	-0.051
CHF	-1.96	2	0.16	2	0	0.000
JPY	-1.93	3	0.83	1	2	0.020
EUR	-1.48	4	-0.05	3	1	0.010
ILS	-1.26	5	-0.50	11	-6	-0.061
KRW	-1.18	6	-1.45	21	-15	-0.152
HUF	-0.96	7	-1.05	20	-13	-0.131
SEK	-0.92	8	-0.68	16	-8	-0.081
CLP	-0.91	9	-0.36	10	-1	-0.010
GBP	-0.84	10	-0.65	14	-4	-0.040
SGD	-0.59	11	-0.08	5	6	0.061
CAD	-0.43	12	-0.56	12	0	0.000
CZK	-0.31	13	-0.96	18	-5	-0.051
NOK	-0.26	14	-3.36	26	-12	-0.121
MYR	-0.21	15	-0.99	19	-4	-0.040
AUD	-0.20	16	-2.43	24	-8	-0.081
THB	0.23	17	-0.26	8	9	0.091
NZD	0.30	18	-2.24	23	-5	-0.051
PLN	0.32	19	-0.81	17	2	0.020
PEN	1.57	20	-0.30	9	11	0.111
CNY	1.95	21	-0.06	4	17	0.172
BRL	2.34	22	-5.02	27	-5	-0.051
COP	2.51	23	-1.56	22	1	0.010
RON	3.37	24	-0.66	15	9	0.091
MXN	5.56	25	-15.75	31	-6	-0.061
ZAR	5.80	26	-3.36	25	1	0.010
RUB	6.51	27	-6.29	29	-2	-0.020
IDR	7.82	28	-5.32	28	0	0.000
PHP	10.42	29	-0.13	7	22	0.222
TRY	11.02	30	-6.69	30	0	0.000
INR	12.67	31	-0.60	13	18	0.182

Table II. The table illustrates summary statistics for the rank difference and carry for various currencies.

FX	RDF mean	RDF sd	RDF % > 0	RDF % < 0	CAR mean	CAR sd	CAR % > 0	CAR % < 0
CZK	-12.0	7.2	6	92	-6.4	4.7	10	87
HUF	-8.1	6.2	2	94	1.8	6.7	60	37
SEK	-8.0	5.8	10	90	-7.5	3.2	0	97
PLN	-7.6	5.6	5	93	0.2	3.4	44	48
COP	-5.1	6.4	12	81	4.7	5.4	84	15
CLP	-4.3	8.1	33	60	2.3	5.6	74	21
KRW	-3.7	8.5	32	63	-2.2	3.4	20	70
NOK	-3.1	6.9	24	70	-2.8	2.8	18	78
RON	-3.0	8.6	37	59	3.4	5.7	74	23
ILS	-2.8	6.2	30	65	-5.6	3.6	6	94
CHF	-2.6	4.8	20	64	-11.4	1.3	0	100
ZAR	-1.8	2.8	21	69	9.5	2.0	100	0
MYR	-1.2	8.8	43	53	-0.5	5.6	49	44
PEN	-1.1	7.5	43	53	2.4	5.3	74	23
TWD	-0.5	6.4	32	65	-9.2	5.0	8	91
EUR	-0.2	5.6	42	45	-8.5	2.6	0	100
BRL	0.4	4.8	53	40	9.6	2.9	99	1
MXN	1.4	5.4	57	32	7.0	3.3	100	0
JPY	2.1	5.1	84	10	-10.2	1.9	0	100
AUD	2.3	6.7	70	27	1.5	4.5	66	26
CAD	2.4	4.4	66	22	-4.0	2.2	1	97
GBP	2.5	7.2	66	26	-4.9	3.2	14	85
IDR	3.1	7.1	68	22	6.7	5.4	87	11
NZD	3.3	4.7	70	21	2.7	4.0	72	23
RUB	3.4	6.5	59	26	8.1	4.8	90	9
THB	3.4	7.1	70	26	0.5	5.1	47	48
TRY	4.5	5.3	68	3	11.6	2.1	100	0
SGD	4.6	6.0	72	20	-4.9	3.5	4	89
PHP	7.8	9.0	79	20	1.4	6.5	69	29
CNY	10.2	12.1	75	20	-0.2	8.7	58	38
INR	14.2	7.1	96	4	8.2	4.1	97	3

Table III. The table illustrates summary statistics for the skew-neutral strategy based on rank differentials (RDF) and hedged carry (CAR_{hedged}) as well as various FX-factor strategies. CAR, SKW, VAL (Asness et al., 2013), MOM (Menkhoff et al., 2012b) and VRP (Della Corte et al., 2016) are constructed with rank-based weights like in Asness et al. (2013) and USD is an equal-weighted portfolio of all currencies against the US-Dollar. The statistics are annualized and include monthly (log-)returns from January 2006 until August 2022.

Panel A: returns all currencies

	RDF	CAR_{hedged}	CAR	SKW	VAL	MOM	VRP	USD
mean	2.76	3.16	3.08	0.76	0.93	0.31	0.42	0.55
sd	5.15	5.30	6.36	7.46	5.36	6.04	4.83	7.56
skew	0.02	0.06	-0.15	-0.17	-0.00	0.16	0.02	-0.12
kurt	-0.20	-0.16	0.07	0.10	0.03	0.20	0.10	-0.16
maxDD	-7.89	-10.33	-10.03	-16.58	-13.80	-22.55	-16.10	-23.84
SR	0.54	0.60	0.48	0.10	0.17	0.05	0.09	0.07

Panel B: spot returns all currencies

	RDF	CAR_{hedged}	CAR	SKW	VAL	MOM	VRP	USD
mean	-0.82	-3.40	-4.81	-4.28	0.80	-0.38	2.16	-1.99
sd	5.19	5.26	6.31	7.43	5.40	6.05	4.87	7.61
skew	-0.04	0.01	-0.17	-0.18	0.00	0.22	0.04	-0.17
kurt	-0.21	-0.15	0.08	0.10	0.04	0.24	0.08	-0.14
SR	-0.16	-0.65	-0.76	-0.58	0.15	-0.06	0.44	-0.26

Panel C: correlations all currencies

	RDF	CAR_{hedged}	CAR	SKW	VAL	MOM	VRP	USD
RDF	1.00	0.57	0.05	-0.52	0.25	-0.39	-0.11	-0.56
CAR_{hedged}	0.57	1.00	0.79	0.36	-0.17	0.03	-0.18	0.06
CAR	0.05	0.79	1.00	0.83	-0.33	0.19	-0.18	0.49
SKW	-0.52	0.36	0.83	1.00	-0.42	0.39	-0.09	0.73
VAL	0.25	-0.17	-0.33	-0.42	1.00	-0.24	0.15	-0.29
MOM	-0.39	0.03	0.19	0.39	-0.24	1.00	-0.01	0.29
VRP	-0.11	-0.18	-0.18	-0.09	0.15	-0.01	1.00	0.24
USD	-0.56	0.06	0.49	0.73	-0.29	0.29	0.24	1.00

Table IV. The table illustrates summary statistics for the weights of various portfolios. Average turnover is calculated as $\frac{1}{t-1}\sum_{t=2}^{T}\sum_{i=1}^{I_t}|w_{i,t}-w_{i,t-1}|$ and the average Herfindahl-Index is defined as $\frac{1}{t}\sum_{t=1}^{T}\sum_{i=1}^{I_t}w_{i,t}^2$.

mean	RDF	CAR_{hedged}	CAR	SKW	VAL	MOM	VRP	USD
... turnover %	112.88	69.18	55.03	57.16	44.96	264.08	71.05	0.03
... Herfindahl-Index	0.23	0.17	0.18	0.18	0.18	0.18	0.18	0.03
... max. weight	19.79	12.72	12.82	12.82	12.82	12.82	12.88	3.23
... min. weight	-19.52	-19.38	-12.82	-12.82	-12.82	-12.82	-12.88	3.23

Table V. The table illustrates summary statistics for the long and short components of the strategy based on rank differentials (RDF) and Carry (CAR). CAR is constructed with rank-based weights like in Asness et al. (2013). The statistics are annualized and include monthly (log-)returns from January 2006 until August 2022. The coskewness measures are calculated relative to the stock market (Mkt-factor taken from Kenneth French's website). The three coskewness measures are (1) taken from Harvey and Siddique (2000), (2) the β_2 from a regression $R_i = \alpha + \beta_1 \cdot Mkt + \beta_2 \cdot Mkt^2$ like in Kraus and Litzenberger (1976), and (3) the covariance of a portfolios excess return on the squared excess return of the US stock market. The bottom part of the table is the correlation matrix.

	CAR_{long}	CAR_{short}	RDF_{long}	RDF_{short}	CAR	RDF	CAR_{hedged}
mean	1.55	-1.48	1.42	-1.72	3.08	2.76	3.16
sd	9.73	6.56	6.26	9.32	6.36	5.15	5.30
SR	0.16	-0.23	0.23	-0.18	0.48	0.54	0.60
skew	-0.18	-0.04	-0.15	-0.15	-0.15	0.02	0.06
kurt	0.10	0.04	0.18	0.12	0.07	0.05	0.09
HS2000	-0.08	0.01	-0.08	-0.09	-0.05	0.01	0.01
KL1976	-0.25	0.01	-0.18	-0.28	-0.13	0.02	0.02
COV(R_m^2, R_i)	-20.71	10.24	-13.25	-21.01	-8.79	6.41	0.08
CAR_{long}	1.00	0.76	0.90	0.84	0.74	-0.42	0.35
CAR_{short}	0.76	1.00	0.83	0.94	0.13	-0.67	-0.25
RDF_{long}	0.90	0.83	1.00	0.85	0.51	-0.31	0.16
RDF_{short}	0.84	0.94	0.85	1.00	0.32	-0.77	-0.21
CAR	0.74	0.13	0.51	0.32	1.00	0.04	0.79
RDF	-0.42	-0.67	-0.31	-0.77	0.04	1.00	0.57
CAR_{hedged}	0.35	-0.25	0.16	-0.21	0.79	0.57	1.00

Table VI. The table illustrates regressions of Carry (CAR), skew-neutral (RDF) and hedged Carry (CAR_{hedged}) portfolios on different versions of the Mkt-factor (Mkt as downloaded from Kenneth French's website, e^2_{Mkt} being deviations from the mean of Mkt and Mkt_{down} being only negative observations of Mkt. ΔVIX is the change in the CBOE VIX Index. CAR is constructed with rank-based weights like in Asness et al. (2013). The regressions include monthly (log-)returns from January 2006 until August 2022.

	CAR	RDF	CAR_{hedged}	CAR	RDF	CAR_{hedged}	CAR	RDF	CAR_{hedged}
(Intercept)	0.002	0.004	0.004	0.002	0.003**	0.003*	0.003**	0.002**	0.003**
	(0.004)	(0.003)	(0.004)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Mkt_{down}	0.193*	−0.103	0.033						
	(0.080)	(0.065)	(0.080)						
Mkt				0.169***	−0.127***	0.016			
				(0.031)	(0.027)	(0.029)			
e^2_{Mkt}				−0.093	0.009	0.051			
				(0.386)	(0.290)	(0.345)			
ΔVIX							−0.128***	0.091***	−0.019
							(0.019)	(0.019)	(0.017)
R^2	0.127	0.067	0.005	0.181	0.147	0.002	0.157	0.120	0.005
Adj. R^2	0.112	0.052	−0.011	0.172	0.138	−0.009	0.153	0.115	−0.000
Num. obs.	62	62	62	186	186	186	194	194	194

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

Table VII. This table includes results from estimating asset pricing models of the form $E_t[M_{t+1}R^i_{t+1}] = 0$ with a linear stochastic discount factor $M_{t+1} = 1-b(f_{t+1}-\mu)$ via GMM. Test assets are quintile portfolios of CAR_i, MOM_i, VAL_i, VRP_i and the long and short portfolio of RDF. As factors USD, CAR, RDF and CAR_{hedged} are used. Estimates for market prices of risk λ and factor loadings b, as well as adjusted R^2, square-root of mean squared errors, and p-values of χ^2 tests are reported. The models include annualized monthly excess returns from January 2006 until August 2022. Standard errors are adjusted according to Newey and West (1987) using an optimal number of lags according to Andrews (1991).

	λ_{USD}	λ_{CAR}	λ_{RDF}	λ_{CAR_h}	b_{USD}	b_{CAR}	b_{RDF}	b_{CAR_h}	R^2_{adj}	RMSE	χ^2
Model 1	0.23	3.2			-0.31	0.84			51.86	0.76	
	(3.1)	(2.61)			(0.3)	(0.36)					7.41%
Model 2	0.42		2.89		0.6		1.4		49.17	0.78	
	(1.04)		(0.6)		(0.33)		(0.46)				17.16%
Model 3	0.27			3.2	0			0.94	55.57	0.73	
	(1.95)			(1.37)	(0.27)			(0.38)			8.86%

Table VIII. Time-series regression results for RDF-returns with carry, momentum, value, the VRP-factor, and USD-factor as explanatory variables. RDF offers significant alpha compared to some of the factors individually. Also, half the variation in RDF-returns remains unexplained in the full multivariate regression. The best model seems to be the simple CAR/USD model as it explains a good portion of variation with two parameters. In particular, it is interesting to see how CAR on its own is an insignificant explanatory variable but as USD is added it becomes highly significant. The regression includes annualized monthly (log-)returns from January 2006 until August 2022.

	RDF	RDF	RDF	RDF	RDF	RDF	RDF	RDF	RDF
(Intercept)	2.764**	2.654**	2.701**	3.114**	2.813**	2.805***	1.771**	1.821**	1.635*
	(0.961)	(1.073)	(1.032)	(1.046)	(1.068)	(0.777)	(0.710)	(0.752)	(0.738)
CAR		0.036					0.341***	0.383***	0.425***
		(0.064)					(0.065)	(0.056)	(0.055)
MOM			0.207*					0.121**	0.101*
			(0.081)					(0.042)	(0.048)
VAL				−0.375***				−0.241***	−0.236***
				(0.077)				(0.059)	(0.060)
VRP					−0.116				0.156*
					(0.094)				(0.068)
USD						−0.383***	−0.522***	−0.464***	−0.510***
						(0.038)	(0.050)	(0.048)	(0.051)
R^2	0.000	0.002	0.059	0.152	0.012	0.319	0.453	0.539	0.556
Adj. R^2	0.000	−0.003	0.054	0.148	0.007	0.315	0.448	0.530	0.545
Num. obs.	199	199	199	199	199	199	199	199	199

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$

Table IX. Time-series regression results for RDF_{short}-returns with carry, momentum, value, the VRP-factor, and USD-factor as explanatory variables. RDF_{short} offers significant alpha once USD is included as an explanatory variable. The regression includes annualized monthly (log-)returns from January 2006 until August 2022.

	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}	RDF_{short}
(Intercept)	-1.719	-3.172	-1.593	-2.203	-1.904	-1.842^{**}	-1.010^{*}	-1.020^{*}	-0.882
	(2.261)	(2.131)	(2.235)	(2.205)	(2.180)	(0.687)	(0.597)	(0.593)	(0.570)
CAR		0.472^{***}					-0.274^{***}	-0.286^{***}	-0.317^{***}
		(0.128)					(0.045)	(0.040)	(0.039)
MOM			-0.409^{*}					-0.036	-0.021
			(0.184)					(0.029)	(0.032)
VAL				0.519^{**}				0.064	0.060
				(0.189)				(0.042)	(0.044)
VRP					0.447^{*}				-0.115^{*}
					(0.205)				(0.048)
USD						1.167^{***}	1.279^{***}	1.263^{***}	1.297^{***}
						(0.036)	(0.031)	(0.030)	(0.037)
R^2	0.000	0.104	0.070	0.089	0.054	0.903	0.930	0.932	0.935
Adj. R^2	0.000	0.099	0.065	0.085	0.049	0.903	0.929	0.931	0.933
Num. obs.	199	199	199	199	199	199	199	199	199

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table X. This table shows the result of ex-post optimized tangency portfolios calculated from monthly (log)returns from January 2006 until August 2022. The first column shows the weights if CAR is replaced with CAR_{hedged}, the second column shows the mean-variance optimal portfolio weights of existing FX-risk premia, and the third column shows weights for a combination of RDF, CAR_{hedged}, and all other FX-factors except CAR.

	CAR_{hedged} swap	old factors	all factors except CAR
CAR_{hedged}	0.52		0.05
RDF			0.48
CAR		0.58	
MOM	0.14	0.18	0.04
VAL	0.18	0.19	0.22
VRP	0.18	0.30	0.06
USD	-0.03	-0.25	0.15
Sharpe ratio	0.72	0.71	0.84

Table XI. This table shows the result of ex-post optimized tangency portfolios calculated from monthly (log-)returns from January 2006 until August 2022 which are constrained to have a skewness coefficient higher or equal 0.1. The first column shows the weights if CAR is replaced with CAR_{hedged}, the second column shows the mean-variance optimal portfolio weights of existing FX-risk premia, and the third column shows weights for a combination of RDF, CAR_{hedged}, and all other FX-factors except CAR.

	CAR_{hedged} swap	old factors	all factors except CAR
CAR_{hedged}	0.43		0.16
RDF			0.29
CAR		0.54	
MOM	0.18	0.13	0.13
VAL	0.25	0.20	0.32
VRP	0.09	0.35	-0.02
USD	0.04	-0.22	0.12
	0.71	0.58	0.73

VIII. Figures

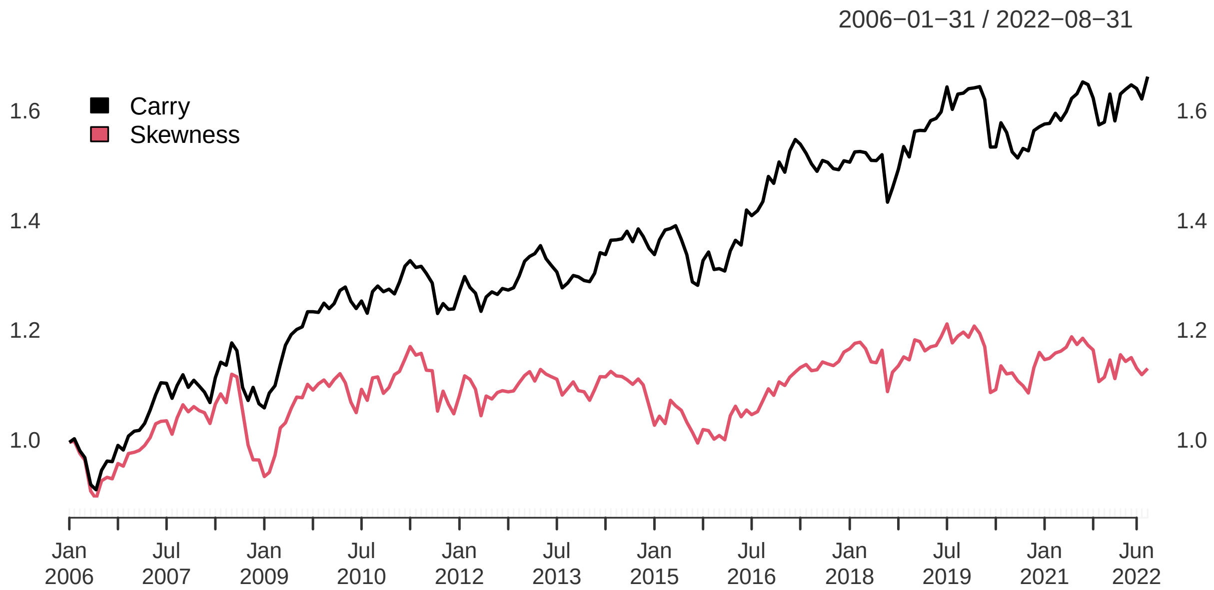

Figure 1. The graph shows the performances of a currency strategy sorted on interest rate differentials (Carry) and a currency strategy sorted on option-implied skewness (Skewness) constructed with rank-based weighting. Option-implied skewness is defined as in Formula 6.

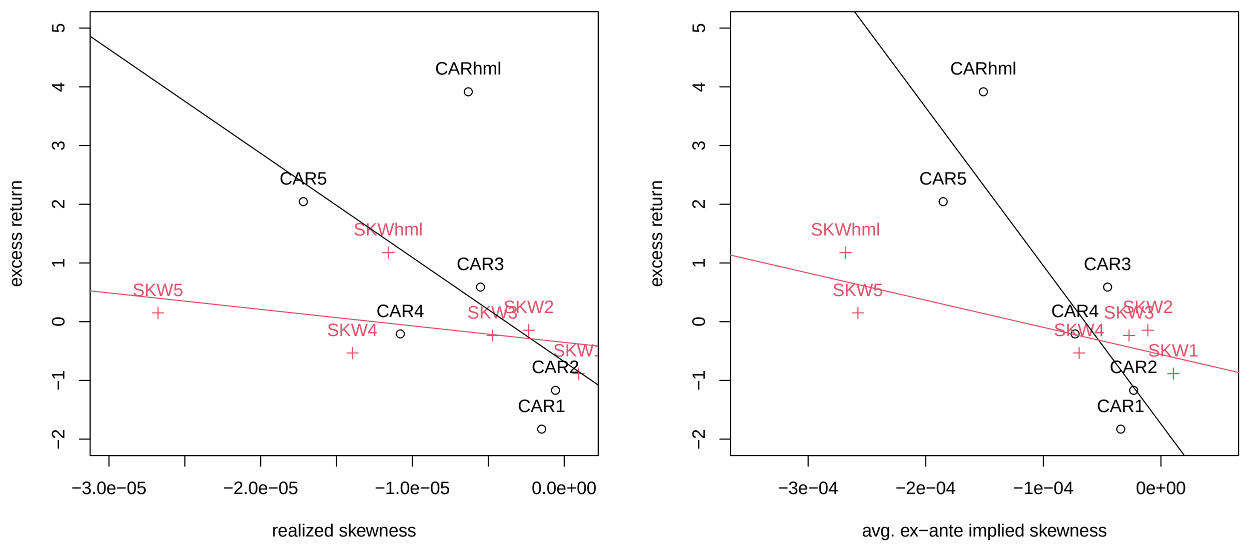

Figure 2. The left graph shows the excess returns and realized skewness of quintile portfolios sorted on interest rate differentials (CAR) and quintile portfolios sorted on option-implied skewness (SKW). The right graph shows the excess returns and average ex-ante implied skewness (Formula 6) of quintile portfolios sorted on interest rate differentials (CAR) and quintile portfolios sorted on option-implied skewness (SKW).

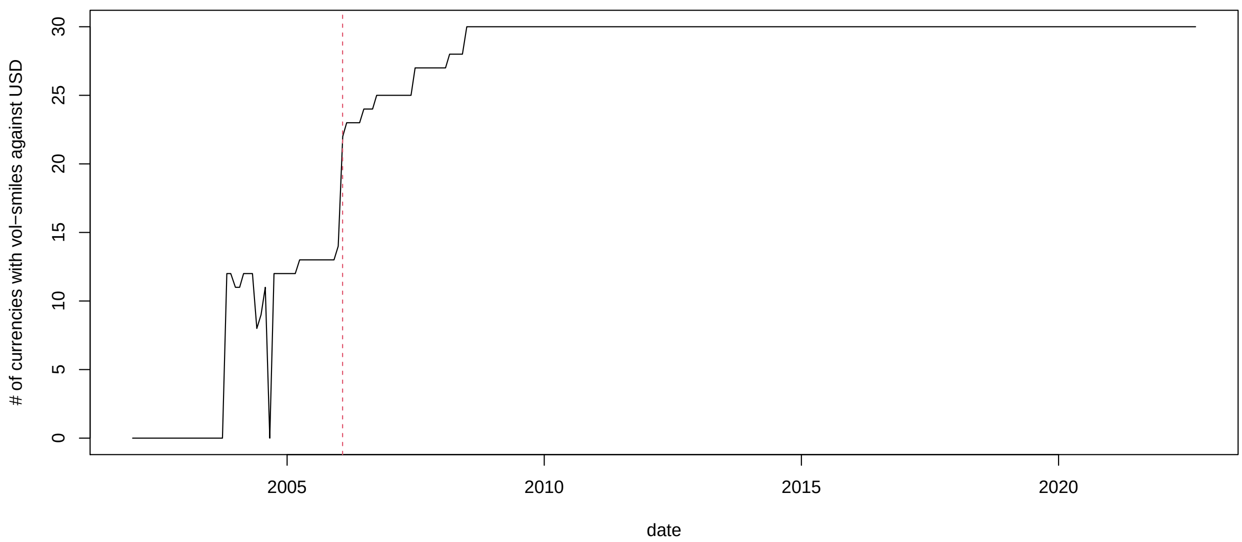

Figure 3. The graph shows the number of available currencies considering option data needed to calculate skewness as given in Formula 6 (10 delta, 25 delta and ATM implied volatilities available). The data number of currencies increases above 20 starting in 2006 and quickly rises to 30.

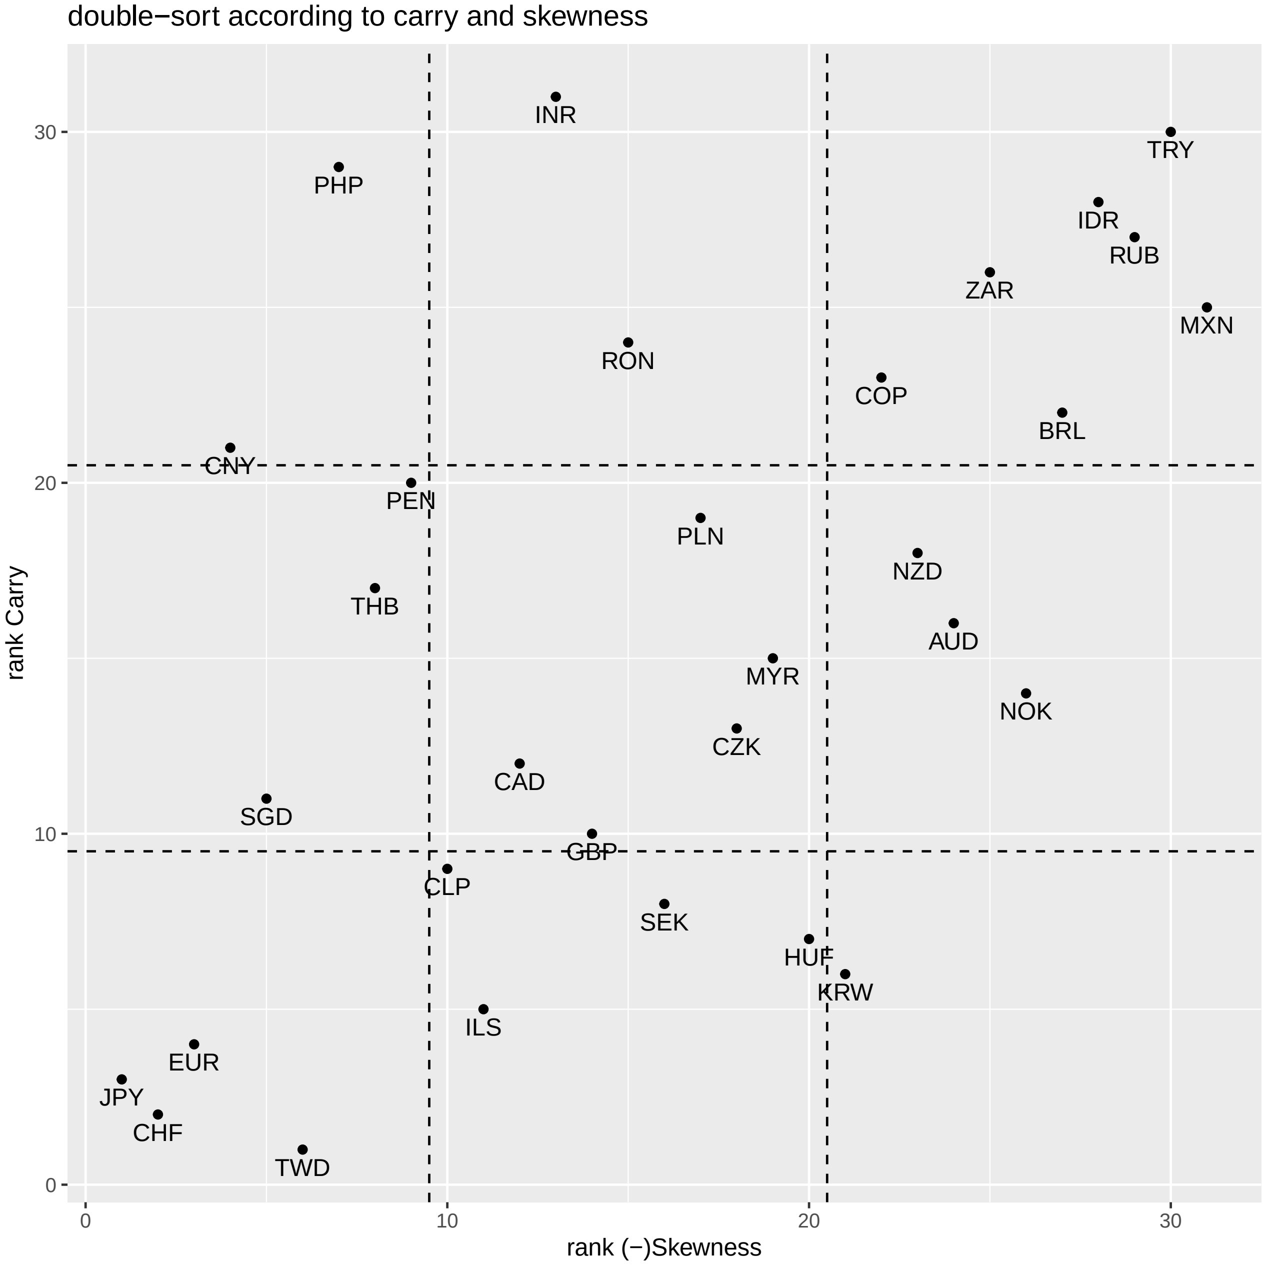

Figure 4. Illustration of an unconditional 3x3 double sort according to carry and option-implied skewness (March 2020).

Figure 5. Illustration of weighting of different portfolios (March 2020). The left graph shows the quartile-based Carry CAR portfolio with equal weights. The middle graph shows the skew-neutral portfolio based an (un)conditional double-sort (unconditional is the colored area, conditional are the coloured dots/currencies), and the right graph illustrates the hedged carry, which is compromised of the long leg of CAR and short leg of the skew-neutral portfolio. Final weights are proportional to vertical lines. Blue dots represent positive weights and red dots negative weights.

Figure 6. Illustration of weighting of different portfolios (March 2020). The left graph shows the Carry (CAR) portfolio with rank-based weights as in Asness et al. (2013). The middle graph shows the skew-neutral portfolio based on rank-differences (RDF) and the right graph illustrates the weights of the hedged carry, which is compromised of the long leg of CAR and short leg of RDF. Final weights are proportional to vertical lines. Blue dots represent positive weights and red dots negative weights.

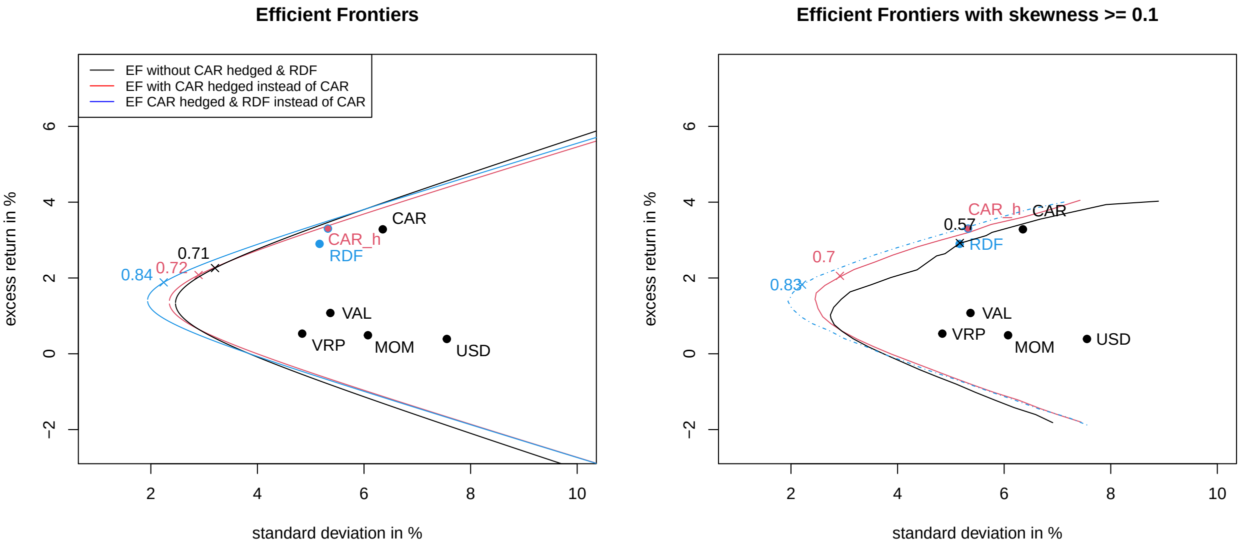

Figure 7. The graphs show the ex-post mean-variance efficient frontiers for various combinations of currency factors. The basis are monthly returns from January 2006 until August 2022.

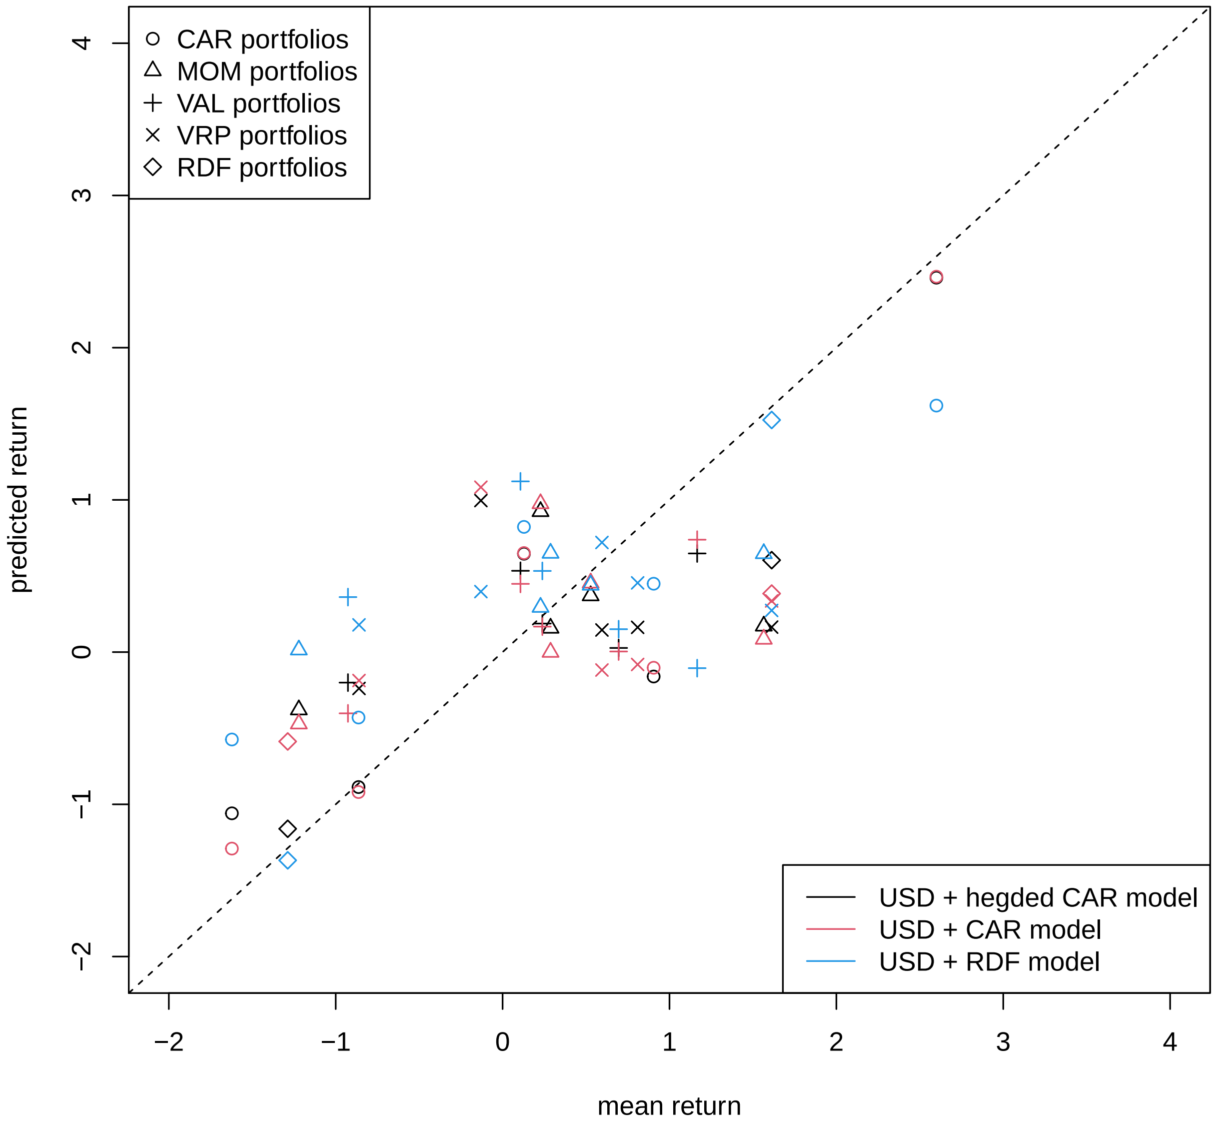

Figure 8. The graphs shows the classic asset pricing chart of predicted return vs. mean returns for 22 different FX portfolios and the three linear models summarized in Table VII. The basis are monthly returns from January 2006 until August 2022.

IX. Appendix

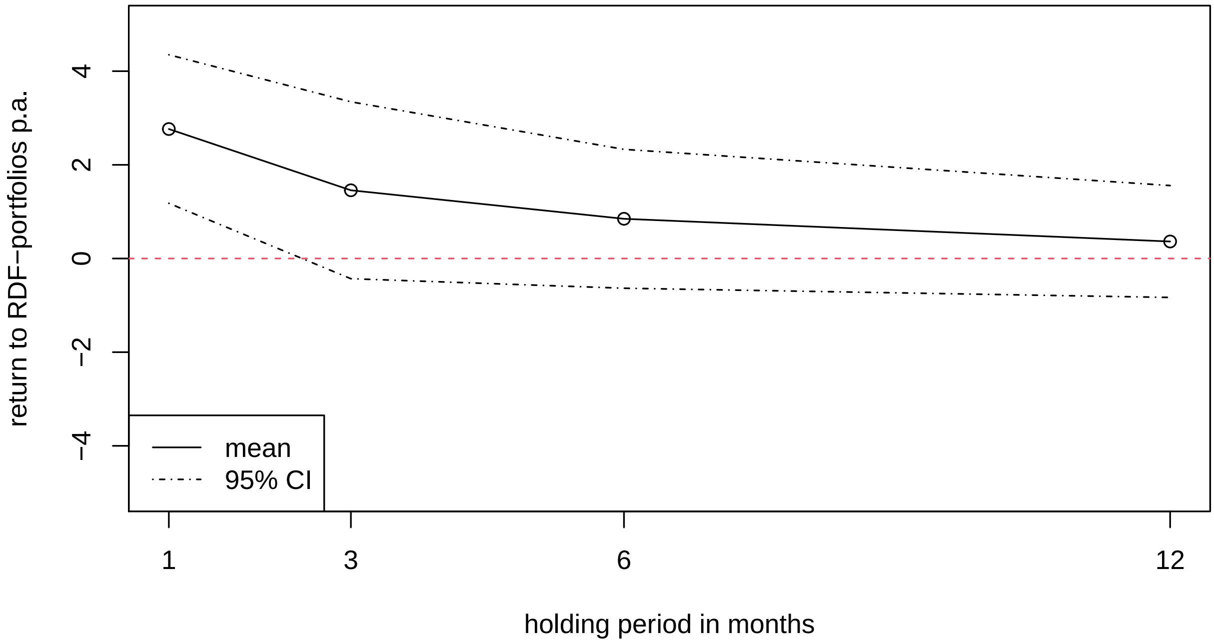

Figure 9. The figure shows the annualized mean return for *RDF* with different holding periods (implemented with FX-forwards with corresponding maturities). The basis are monthly returns from January 2006 until August 2022. The returns (except for the holding period of 1-month) are overlapping and the standard errors are corrected for auto-correlation using the methodology of Newey and West (1987). Returns for holding periods greater than 1 month are not significantly different from 0.

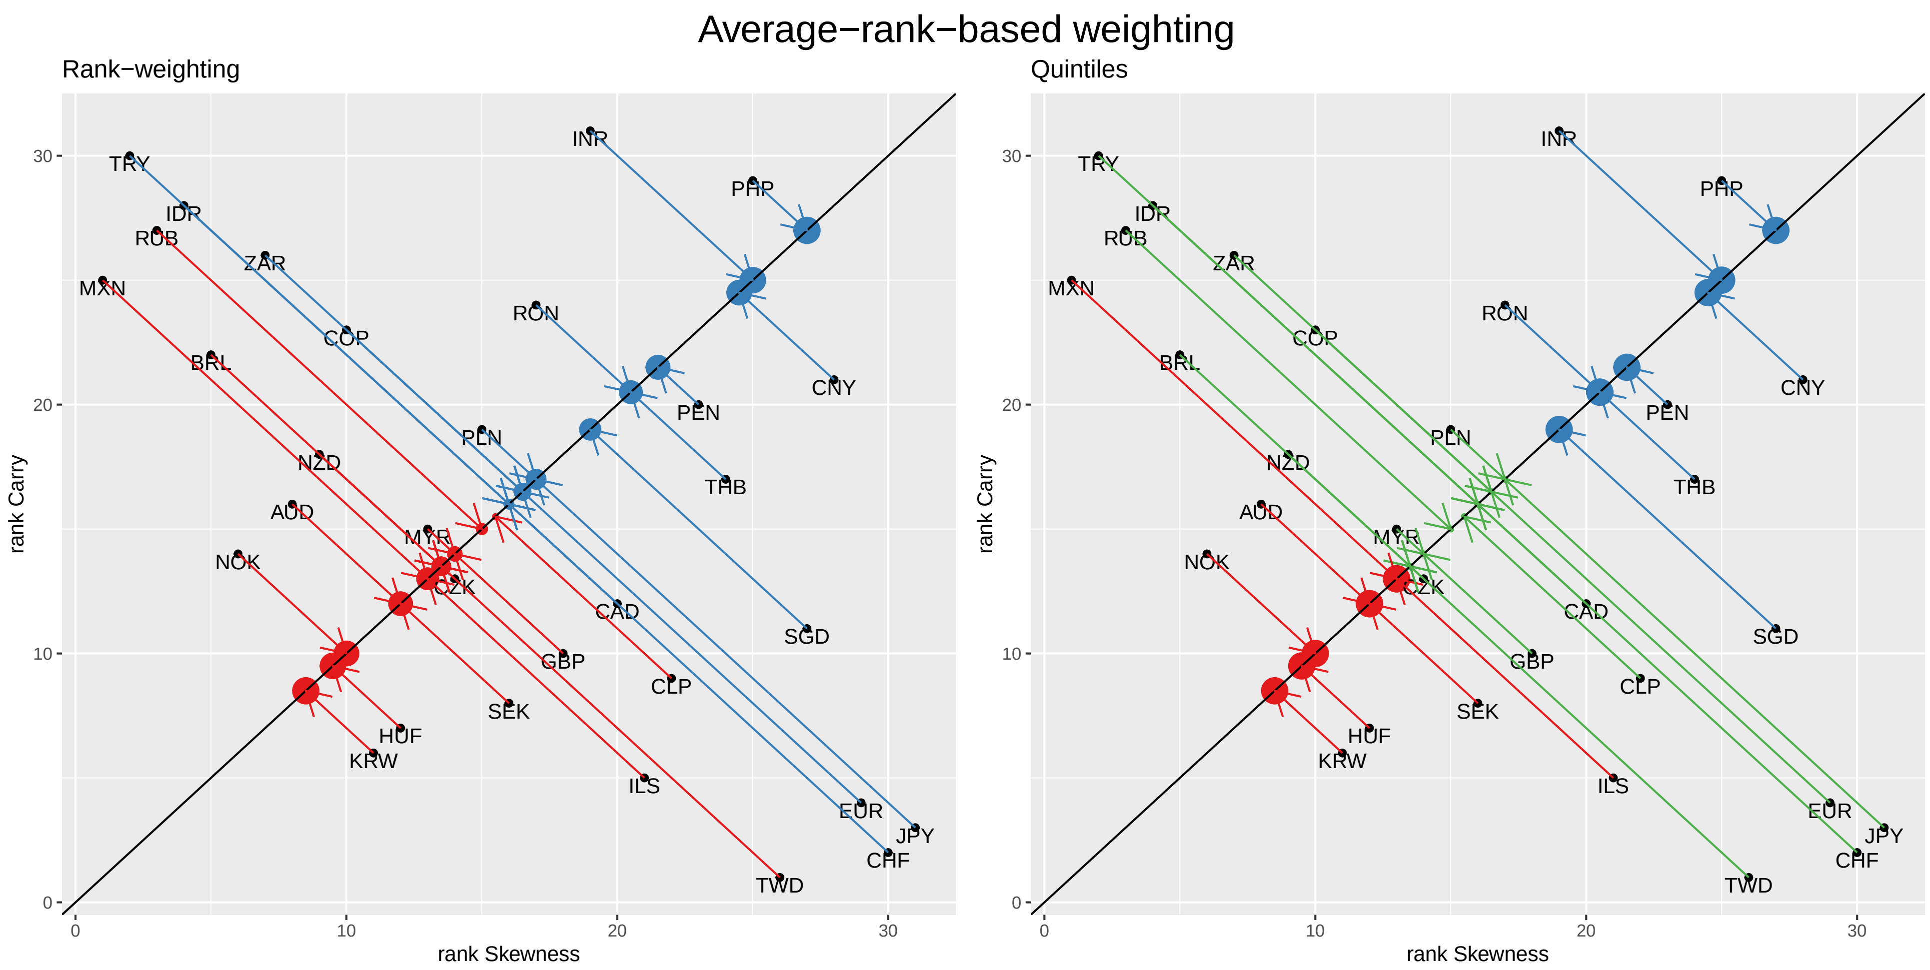

Figure 10. Illustration of average-rank-based weighting on March 2020. First average ranks are calculated and then these are used to create portfolios. Fisher et al. (2015) propose to calculate average ranks and value-weight quantile portfolios to more efficiently combine value and momentum stocks. The right graph illustrates the weights of a skew-neutral long-short quintile portfolio where currencies are equally weighted. The left graph illustrates a skew-neutral rank-weighted portfolio based on average ranks.

Table XII. The table illustrates summary statistics for the long and short components of the skew-neutral strategy based on a 2x2 unconditional double-sort (SN) and Carry (CAR). CAR is constructed from the highest and lowest quartile of currencies according to interest rates. The statistics are annualized and include monthly (log-)returns from January 2006 until August 2022. The bottom part of the table is the correlation matrix.

	CAR_{long}	CAR_{short}	SN_{long}	SN_{short}	CAR	SN	CAR_{hedged}
mean	1.87	-1.55	1.74	-1.62	3.47	2.84	3.26
sd	10.20	6.51	6.61	10.10	7.22	6.68	6.93
SR	0.18	-0.24	0.26	-0.16	0.48	0.43	0.47
skew	-0.17	-0.04	-0.10	-0.14	-0.14	0.01	-0.05
kurt	0.10	0.03	0.19	0.13	0.08	0.02	0.02
HS2000	-0.12	-0.08	-0.10	-0.09	-0.03	-0.02	-0.04
KL1976	-0.34	-0.18	-0.24	-0.31	-0.11	-0.04	-0.08
COV(R_m^2, R_i)	-19.65	-11.83	-14.04	-21.86	-10.52	6.06	1.48
CAR_{long}	1.00	0.71	0.80	0.76	0.77	-0.35	0.36
CAR_{short}	0.71	1.00	0.74	0.90	0.10	-0.61	-0.26
SN_{long}	0.80	0.74	1.00	0.75	0.46	-0.14	0.09
SN_{short}	0.76	0.90	0.75	1.00	0.27	-0.76	-0.33
CAR	0.77	0.10	0.46	0.27	1.00	0.05	0.74
SN	-0.35	-0.61	-0.14	-0.76	0.05	1.00	0.58
CAR_{hedged}	0.36	-0.26	0.09	-0.33	0.74	0.58	1.00

Table XIII. The table illustrates summary statistics for average-rank-based portfolios. RDF_q is a skew-neutral long-short quintile portfolio where currencies are equally weighted and $CAR_{hedged,q}$ uses the short leg of the former to combine it with the long quintile of a classic carry portfolio. RDF_{rw} is a skew-neutral rank-weighted portfolio based on average ranks and $CAR_{hedged,rw}$ combines a rank-weighted long carry portfolio with the short leg of RDF_{rw}. The statistics are annualized and include monthly (log-)returns from January 2006 until August 2022.

	RDF_q	$CAR_{hedged,q}$	RDF_{rw}	$CAR_{hedged,r}$
mean	1.63	1.89	1.94	2.99
sd	5.53	4.47	4.54	4.72
skew	0.01	-0.03	-0.03	-0.01
kurt	0.04	-0.21	0.02	-0.18
SR	0.30	0.42	0.43	0.63

Table XIV. The table shows summary statistics of *RDF* portfolios constructed using four different option maturities to calculate the skewness signal. Although 1-month options are the best, longer-dated options seem to convey similar information and result in similar portfolios with lower returns.

	1-month	3-month	6-month	1-year
mean	2.764	2.397	2.385	2.445
sd	5.151	5.238	5.251	5.292
Sharpe Ratio	0.537	0.458	0.454	0.462
skew	0.016	0.034	0.018	0.033
kurt	0.306	0.323	0.327	0.396

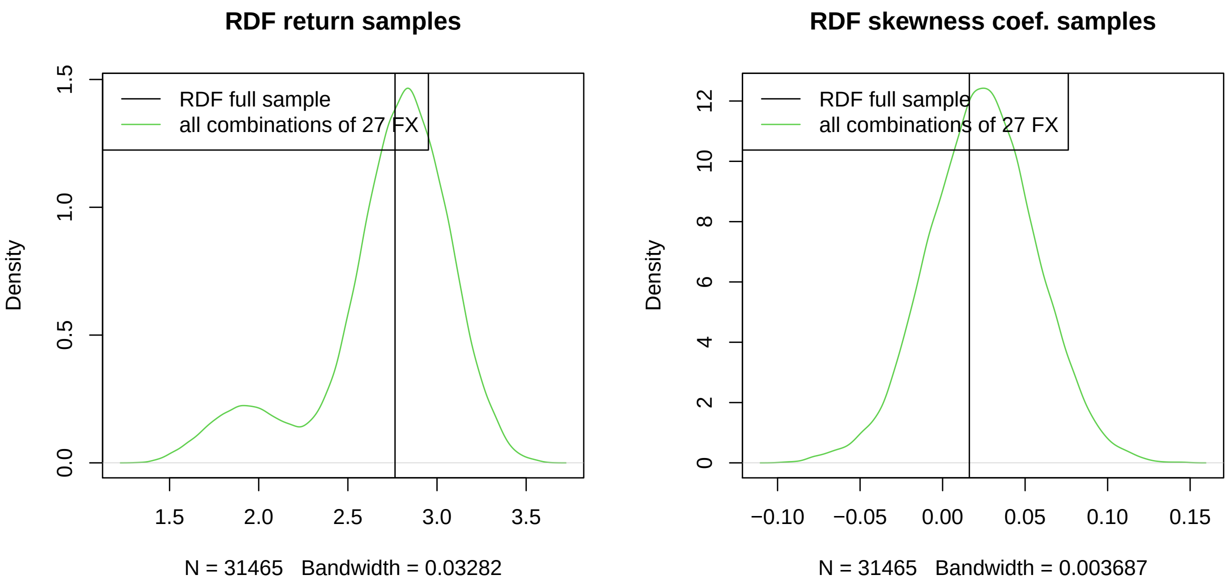

Figure 11. This graph shows the distribution of mean returns and the skewness coefficients of *RDF* for all combinations of 27 currencies selected from the full sample of 31 currencies.

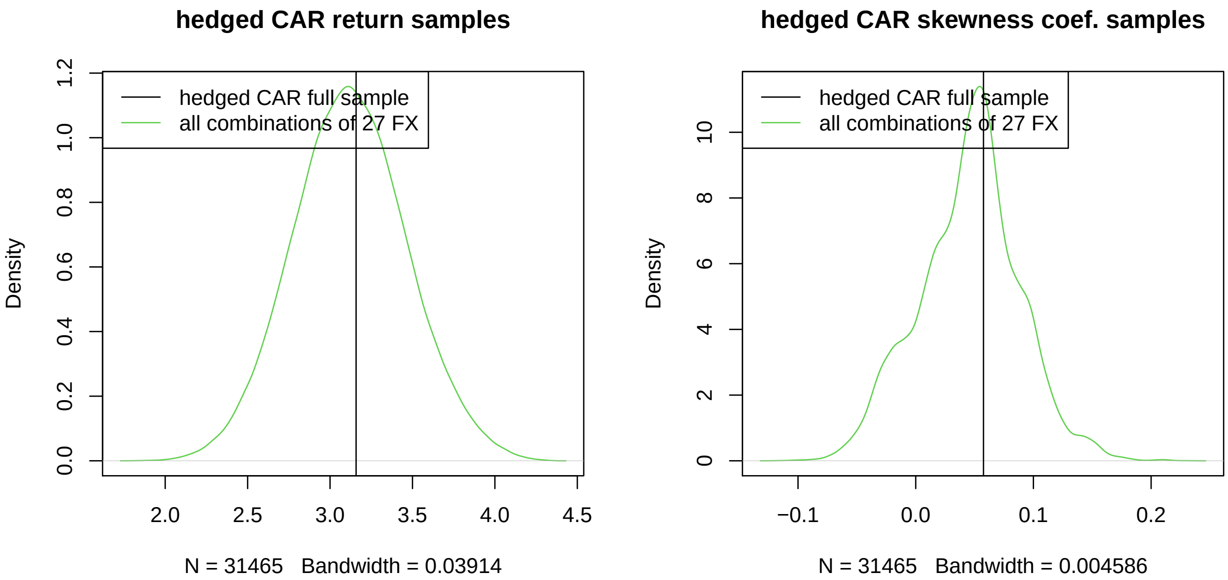

Figure 12. This graph shows the distribution of mean returns and the skewness coefficients of CAR_{hedged} for all combinations of 27 currencies selected from the original sample of 31 currencies.