

CARRY ON

Megan Czasonis, Baykan Pamir, and David Turkington

Megan Czasonis is managing director at State Street Associates in Cambridge, MA.
mczasonis@statestreet.com

Baykan Pamir is assistant vice president at State Street Associates in Cambridge, MA.
bpamir@statestreet.com

David Turkington is senior vice president at State Street Associates in Cambridge, MA.
dturkington@statestreet.com

Forthcoming in the Journal of Alternative Investments

THIS VERSION: September 25, 2018

Abstract

The carry trade in foreign currencies is known for delivering positive returns on average, and for occasionally suffering large losses. While these characteristics prevail on average across time and across currency pairs, we find that interest rate differentials on their own are not sufficient to identify conditions in which currencies exhibit these return and risk attributes reliably. We use three variables – valuation, crowding and volatility – to identify time periods and cross-sections of currencies in which the carry trade performs best. We document a substantial difference in performance between the carry trade applied to high volatility versus low volatility currency pairs. In the full sample from 1984 to 2017, carry in high volatility pairs has consisted of currencies which are undervalued on average, experience greater swings in valuation, and have boom and bust cycles aligned with investor crowding. This finding is consistent with the notion that carry represents a risk premium. Carry in low volatility pairs has the opposite characteristics. Though both strategies performed well prior to the 2008 financial crisis, only carry in high-volatility pairs has worked since.

CARRY ON

For decades, borrowing in countries with low interest rates and investing in countries with high interest rates has been profitable. The return on this simple investment strategy, called the “carry trade” in currency markets, also depends on the exchange rate at which the foreign assets are bought and sold. The trade profits so long as the currency does not depreciate by more than the interest rate differential. On average throughout history and across currency pairs, exchange rates have depreciated by less than the differential and average returns to carry have been positive. Forward contracts offer a convenient way to implement the carry trade. According to Covered Interest Parity (CIP), currency forward rates must reflect the risk-free interest rates of both countries. Otherwise, arbitrageurs can earn riskless profits by borrowing in one country, lending in the other, and hedging away the currency risk. Thus, exchange rates for countries with high interest rates trade at a forward discount to the current spot exchange rate.¹ A long position in discount forward contracts is equivalent to the trade defined earlier: the investor agrees to buy the higher interest rate currency at a set time in the future for a price that is below the current exchange rate, and she will profit by selling it immediately so long as the spot exchange rate does not fall below that level. This effect is called the “forward rate bias” because the forward discounts of higher interest rate currencies systematically over-predict depreciation in their spot rates.

The underlying source of this bias is widely considered a puzzle. Hansen and Hodrick [1980] and Fama [1984] note that the bias allows positive expected profits to foreign exchange

speculation, which stands in contradiction to the efficient market hypothesis. Nonetheless, empirical evidence for the forward rate bias has been persuasive, as documented in surveys by Engel [1996] and others.² Kritzman [1997, 2004] showed that even simple carry trades designed to exploit the bias offered reliably strong returns for decades, even as skeptics have often predicted that widespread knowledge of the forward rate bias may cause it to disappear. Many papers have considered the idea that carry profits should persist in equilibrium as fair compensation for risk (see, for example Hodrick and Srivastava [1984], Brunnermeir [2009] and Burnside [2012]). Even so, this view lacks general consensus because research has struggled to link carry returns to specific economic risk factors. Recent work including Dobrynskaya [2014] and Verdalhan [2018] has advanced the definition of currency risk factors from an asset pricing perspective.

There are also behavioral explanations for carry. Consider, for example, the choice between a forward contract that will profit and one that will lose money if exchange rates hold constant. Kritzman [2004] suggests most investors will prefer the first option. Anchoring to the current exchange rate as a status quo makes the carry trade appear profitable in expectation, and this logic can be self-fulfilling. To the extent other investors also engage in the carry trade, they will move spot rates in a favorable direction which provides ever more reason to invest. Eventually, we should expect speculative bubbles and crashes (see Melvin and Shand [2017] for an historical catalog of carry trade losses). From this perspective risk is a consequence and not necessarily a cause.

In this paper, we begin by documenting that in spite of strong historical performance, the returns of a typical carry strategy in developed market currencies has, from 2013 to 2017,

experienced its longest drawdown since 1984 when our data sample begins.³ Moreover, recent performance does not exhibit the typical boom and bust pattern that has characterized the carry trade in the past. We use this recent sample as a case study to evaluate the drivers of carry performance. First, we decompose carry returns into their interest rate and spot return components, and find that spot rates often moved in favor of carry before the financial crisis, but this trend has since reversed. Second, we show that the carry trade has been overvalued in the recent sample which may partly explain its poor performance. Third, we apply a measure of network centrality to show that investor crowding in the carry trade contributes to both positive returns and subsequent crashes. This result supports the existence of behavioral factors. Fourth, we sort currency pairs based on their level of risk. We find that when the carry trade is applied to a subset of high-volatility currency pairs, its performance has the characteristics we expect for a risk premium. In contrast, the carry trade applied to low-volatility currency pairs does not. Since 2008, the high-volatility carry portfolio has continued to perform well, while the low-volatility carry portfolio has had persistently negative returns. This finding suggests that carry offers a risk premium in some currencies, but not others.

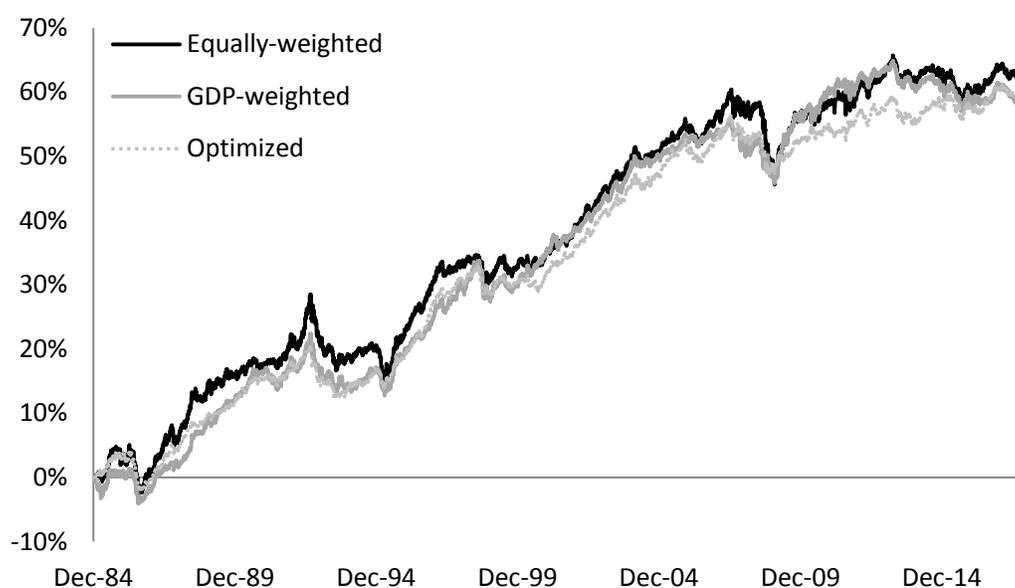
Measuring carry trade performance

We base our analysis on a formulation of the carry trade that is base currency agnostic. Using one currency as a fixed point of reference presents two potential problems. If the amount of long and short positions are allowed to drift freely, the choice of base currency may have an outsized impact on results. On the other hand, if long and short positions are precisely balanced, they will omit information about the base currency altogether. Our approach treats each currency as an equal contributor. We proceed as follows. At the end of each month, we obtain daily spot exchange rates, forward rates, and implied interest rate differentials⁴ for each of 45 currency pairs within the G10 currency universe (Australian dollar, Canadian dollar, Swiss franc, Euro, British pound, Japanese yen, Swedish krona, Norwegian krone, New Zealand dollar, and US dollar).⁵ Next, we align each currency pair such that a long position corresponds to a positive interest rate differential, and we take positions in the 27 pairs with the largest interest rate differentials. We choose 27 pairs to ensure that no single currency comprises more than one third of the long exposures or one third of the short exposures in the portfolio (each currency is represented in exactly nine pairs). This approach also allows for the net exposure of different currencies to vary. It is essentially a more nuanced version of a “top three, bottom three” ranked portfolio.

We apply this methodology to build carry portfolios with equal weights, GDP weights and mean-variance optimal weights.⁶ In each case, we rebalance the portfolios monthly. Exhibit 1 shows the cumulative performance of each strategy from December 1984 to December 2017. All three strategies are highly correlated, and they all perform worse recently than they have in

the past. For simplicity, we focus on the equally-weighted strategy in the remainder of the paper.

Exhibit 1: Cumulative performance of FX carry trade strategies



Notes: All strategies are rescaled to have the same cumulative return at the end of December 2008. The full-sample information ratios equal 0.56 for equally-weighted, 0.63 for GDP-weighted, and 0.67 for optimized.

The impact of interest rate differentials

Following the 2008 financial crisis, interest rates generally declined in developed countries and the spreads between them compressed. All else equal, smaller interest rate differentials imply less opportunity for carry profits. We decompose the total return of carry trade into an interest rate differential component and a spot rate component. Exhibit 2 shows rolling five year spot returns along with the breakeven rate below which the carry trade will result in a net loss.

Although the breakeven rate has compressed over time, the behavior of the spot price suggests an even more dramatic shift. Spot returns were frequently positive from 1984 to 2008, but have been mostly negative in the recent sample. This result suggests that smaller interest rate differentials are not the primary cause of poor carry performance. Next, we turn our attention to factors that explain the behavior of the spot rate.

Exhibit 2: Breakeven spot returns and actual spot returns (5 year rolling window)



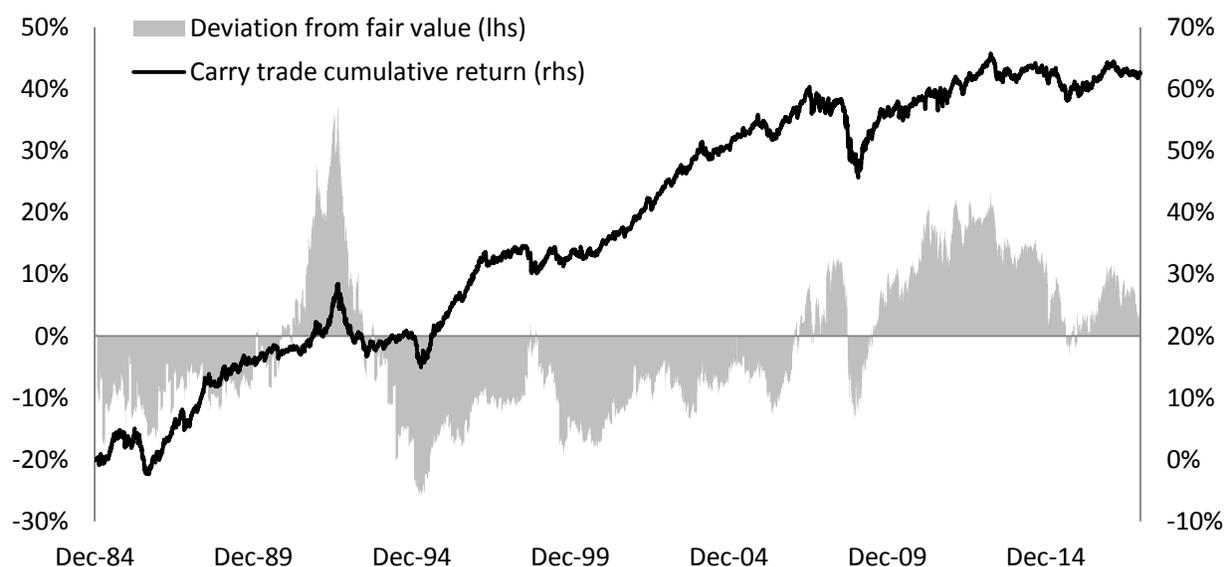
The impact of valuations

Purchasing Power Parity (PPP) suggests a theoretical fair value for one currency versus another. The idea is that over time, exchange rates tend towards equilibrium. If it is cheaper to buy an identical product in Canada than in the United States, for example, Americans will be tempted to buy more goods from Canada. This increases demand for the Canadian dollar and causes it to appreciate. Likewise, countries with expensive products should expect less demand and

pressure for their currencies to depreciate. However, it is also possible that relative goods prices will adjust rather than exchange rates. Prior research has generally concluded that spot rates tend to move towards fair value, though it may take years for them to converge.⁷

Exhibit 3 shows the relative valuation of carry through time. It is simply the weighted average of valuation per currency for all the currency pairs in the carry portfolio. To compute fair value we use cross-country relative prices from the OECD.⁸ We define relative valuation as the deviation of the actual spot exchange rate from the exchange rate implied by relative prices. Exhibit 3 reveals that strong carry performance often aligns with periods in which carry (as a portfolio) is undervalued. Against this backdrop the trade is doubly compelling: it offers a perceived yield advantage, and there is reason to expect that on average the spot rate should appreciate rather than depreciate. Of course, this alignment of factors will not always occur. In the post-crisis sample we see the opposite: carry currencies have mostly been overvalued and the trade has not performed well.

Exhibit 3: Valuation of the carry trade portfolio over time



The impact of investor crowding

Some positive carry returns have occurred while the underlying positions are overvalued. These periods have been brief and are followed by crashes. They are characterized by positive returns that accrue not only to the interest rate differential, but also to favorable spot returns. Given that the carry trade should be less attractive in terms of valuation in these periods, we consider whether returns are explained by speculative bubbles led by investor crowding. Our conjecture is that following long periods of stable and attractive returns for the carry trade, investors chase performance and increase their allocation. As demand increases, spot rates become overvalued and susceptible to a crash.

We infer crowding in currency factors by applying a measure of network centrality to the covariances of currency returns. Our approach follows the methodology of Kinlaw et al [2012, 2018]. For a given day, we obtain time series of daily spot returns for all 45 currency

pairs over the previous two years. Next, we apply an exponential decay function with a one year half-life to place relatively greater weight on recent observations, and we compute the covariance matrix of weighted returns. With this method, information gradually recedes in relevance over time as the window rolls forward. We extract the top $n = 2$ statistical factors (eigenvectors) from a principal components analysis of the covariance matrix⁹ and compute the centrality for currency pair ij at any given point in time as:

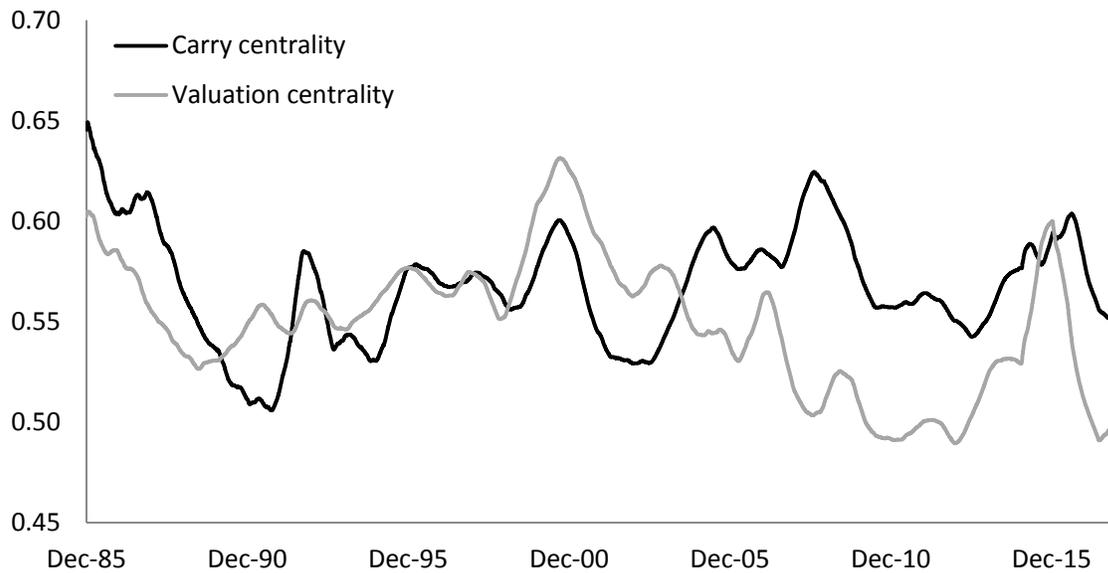
$$Centrality_{ij} = \frac{\sum_{k=1}^n \lambda_k w_{ij;k}}{\sum_{k=1}^n \lambda_k}$$

where $w_{ij;k}$ is the normalized absolute value weight of currency pair ij in eigenvector k , and λ_k is the eigenvalue corresponding to eigenvector k , which represents the variance explained by that eigenvector. Thus, with $n = 2$ the centrality for a given currency pair is equal to the magnitude of its weight in the top principal component multiplied by the total fraction of variance explained by the top component, plus the magnitude of its weight in the second principal component multiplied by the total fraction of variance explained by the second component. Centrality has an intuitive interpretation: in a network graph where nodes represent volatility of currency pairs and correlations represent the strength of the connections between pairs, centrality represents the “importance score” one can assign to each pair such that each node’s importance equals a weighted average of its own price volatility and the volatilities of the other pairs to which it is connected. Centrality scores sum to one across all pairs, so it is easy to see that some are more important than others. All else equal, centrality

will increase if a pair becomes more volatile, if the pairs to which it is connected become more volatile, or if the strength of its correlation to other volatile pairs increases.

In this context, we interpret centrality as a proxy for investor crowding. As a currency pair becomes more crowded, order imbalances are likely to increase price impact, and thus volatility, which will be reflected in higher centrality. We also expect investor crowding to cause similar currency pairs to move in tandem, which also increases their centrality. To measure the centrality of valuation and carry factors – rather than individual currency pairs – we compute weighted average centrality scores using cross-sectional percent ranks of the 45 currency pairs sorted by the magnitude of their relative valuation or interest rate differential, respectively. Exhibit 4 plots the one year moving average of these measures.¹⁰

Exhibit 4: Valuation and carry centrality



Next, we investigate these centrality scores in the context of historical carry returns and valuations. We argue that the interaction of these characteristics typically proceeds as a cycle that consists of three phases, as shown in Exhibit 5. First, consider the top panel which spans 1985 to 1995. In the first phase of the cycle, the carry trade is characterized by strong fundamentals: it is undervalued and therefore attractive from both a yield and valuation perspective. We identify the start of this phase – in retrospect – as a low point in valuation. The carry trade generated a return of 3 percent per year on average during this period. The second phase begins when the carry factor becomes more crowded than the valuation factor. We label the start of this phase on the first day that carry centrality rises above valuation centrality. The carry trade generated a positive annualized return of 11 percent during this period, and eventually it became overvalued. The third phase represents a reversion to fair value, often by way of a crash in price. We label the start of this phase on the first day that carry centrality

begins to decline, as this suggests a turning point in investor enthusiasm for carry. In this period, spot rates fell, valuations dropped, and the carry strategy generated losses of 4 percent per year. The subsequent low point in carry valuation marks the end of this cycle and the beginning of a new one. The second cycle in Exhibit 5 proceeded in a similar fashion, spanning 1995 to 2008. Our main point is that a portion of the positive returns to carry aligns with strong fundamentals, while another portion aligns with investor crowding.

Notably, neither strong fundamentals nor investor crowding prevailed from 2009 to 2015, as shown in the third panel of Exhibit 5. The absence of these two factors may help explain the carry trade's poor performance in this period. The fourth and most recent cycle in the data is qualitatively similar to the first two, though it is more muted. Nevertheless, this short period suggests that with attractive valuations, the same dynamics that led to strong carry performance in the past might recur. Exhibit 6 summarizes annualized carry returns in each cycle and the contribution of each phase.

Exhibit 5: Historical FX carry cycles (bars represent valuations on left-hand-side axis; dark line is carry centrality and gray line is valuation centrality, both on right-hand-side axis)

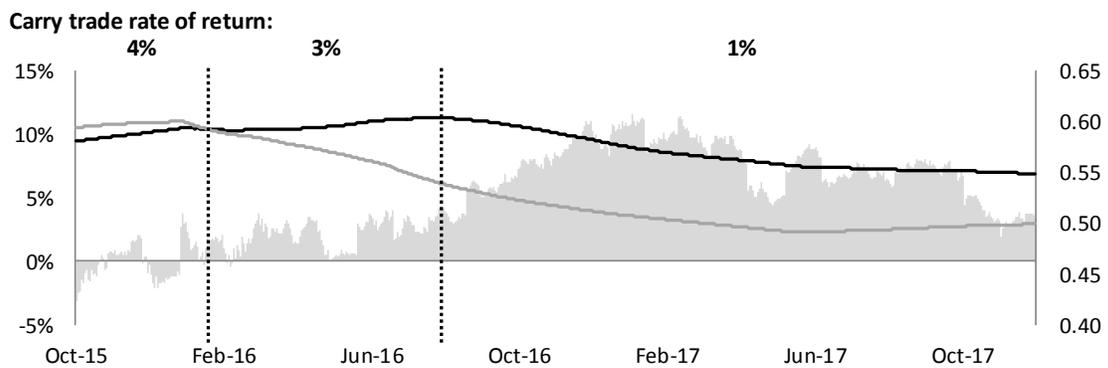
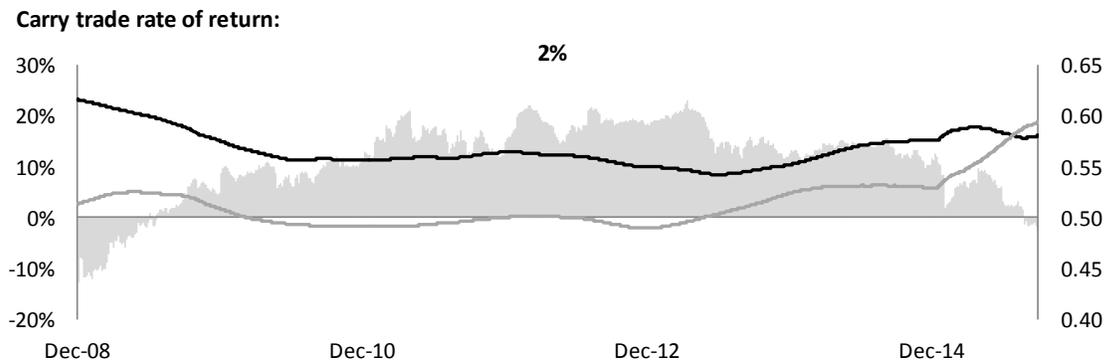
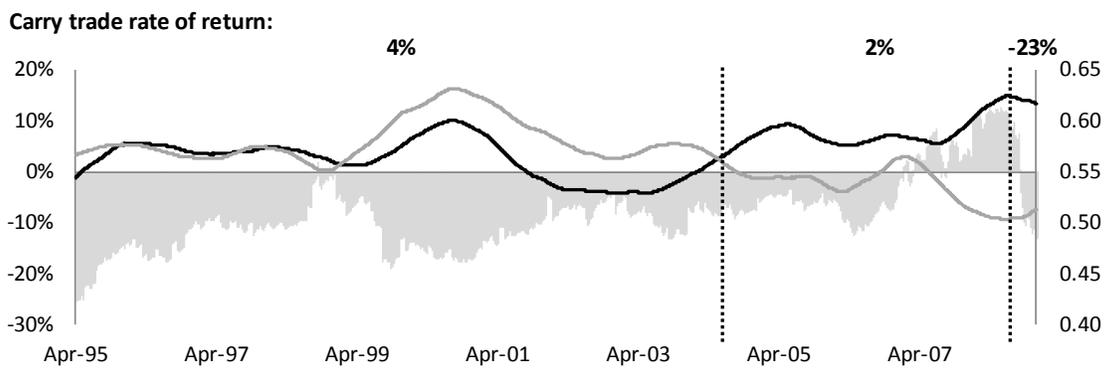
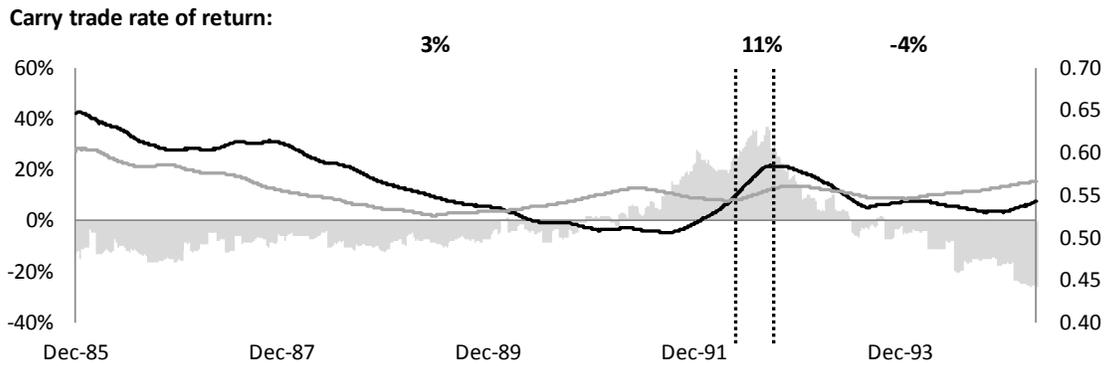
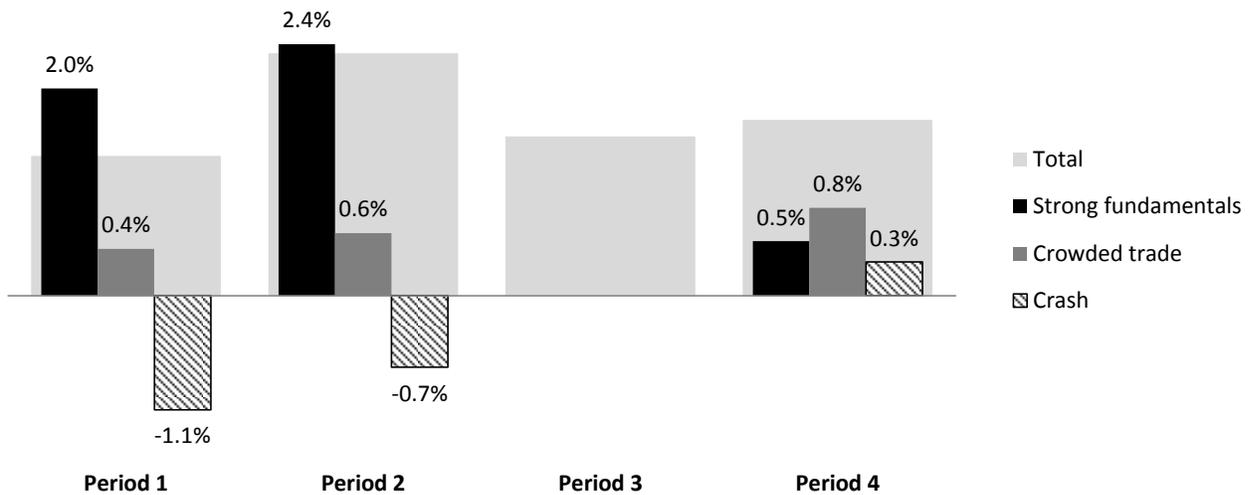


Exhibit 6: Return attribution by phase of cycle



Note: We do not show an attribution for period 3 because its pattern is atypical and does not conform to the same cyclical behavior.

The impact of volatility

Do historical carry returns reflect a risk premium, or an anomaly that is unlikely to sustain in the future? We suggest that both stories may be partly correct. To address this question, we form carry portfolios within subsets of high volatility pairs and low volatility pairs using the previous two years of daily spot returns. Exhibit 7 reveals dramatic differences in the performance characteristics of the two strategies. The carry strategy formed from high volatility pairs is riskier in many ways:

- It has more than twice as much daily volatility.
- Its average deviations from PPP-implied fair value are more than twice as large.
- It has more exposure to systematic global equity risk.
- It is exposed to crash risk following investor crowding.

We also see evidence that the high volatility carry portfolio is compensated in the form of higher returns. In particular:

- It has been undervalued, on average.
- It is associated with larger interest rate differentials, on average.
- Its total return is more than twice as large.
- Its positive returns have persisted post financial crisis.

Exhibit 7: Performance characteristics of high versus low volatility carry strategies

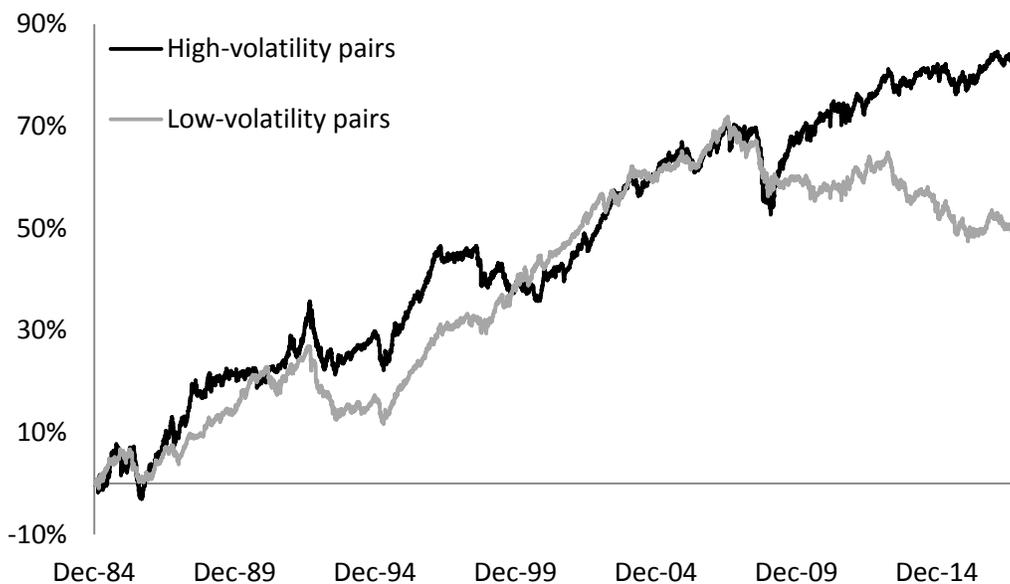
	High volatility pairs	Low volatility pairs
Volatility of daily spot returns (annualized)	4.8%	2.2%
Average size deviation from PPP fair value	15.9%	7.4%
Global equity beta	0.12	0.03
T-statistic of global equity beta	7.51	4.04
Spot return following carry crowding (threshold = 1)	-2.6%	0.3%
Spot return following carry crowding (threshold = 2)	-7.0%	-0.5%
Spot return following carry crowding (threshold = 3)	-17.3%	2.1%
Average valuation relative to PPP fair value	-7.5%	3.9%
Average interest rate differential (annualized)	2.4%	1.5%
Return of carry (annualized)	2.4%	0.9%
Risk-adjusted return (Pre-crisis)	0.48	0.67
Risk-adjusted return (Post-crisis)	0.62	-0.18

Notes: The global equity regression is based on MSCI ACWI monthly returns denominated in local currency. The data spans Jan 1988 through Dec 2017. The regression yields an intercept of 1.7% versus 0.7%, residual risk of 4.3% versus 2.2% and an R-squared of 14% versus 4% for high volatility and low volatility pairs respectively. We conclude that the equity premium only accounts for a portion of the carry trade's performance. The spot return following carry crowding is defined as the difference between the average next day spot return following days with a three-year standardized shift of carry centrality greater than and less than the threshold.

The fact that the low volatility carry strategy embodies less risk suggests that its strong performance pre-crisis may have resulted from temporarily favorable conditions, such as

undervalued spot rates in some pairs and occasional investor enthusiasm for carry during good times. In the absence of these supportive conditions post-crisis, low volatility carry has failed to generate positive returns. In contrast, high volatility carry has continued to perform in line with its historical trend, as shown in Exhibit 8.

Exhibit 8: Cumulative carry returns for high versus low volatility pairs



Notes: Both strategies are rescaled to have the same cumulative return at the end of December 2008.

These results suggest that investors are well advised to accept – and even embrace – the inherent risk of the carry trade. However, we have not yet addressed the possibility that well-informed investors may be able to anticipate and avoid some of the strategy’s occasional losses. There is reason to believe that a dynamic strategy can add value. For example, an investor may use the joint information contained in valuations and centrality to detect bubbles

and avoid the carry trade when it is most vulnerable to loss. Kinlaw et al [2018] illustrate the potential benefits of this approach applied to equity markets. Alternatively, Kritzman and Li [2010] show that a measure of market turbulence is effective at anticipating carry losses. We apply this technique to the carry strategy constructed from high-volatility currency pairs. We compute turbulence from daily currency returns as:

$$Turbulence = \frac{(y - \mu)\Sigma^{-1}(y - \mu)'}{N}$$

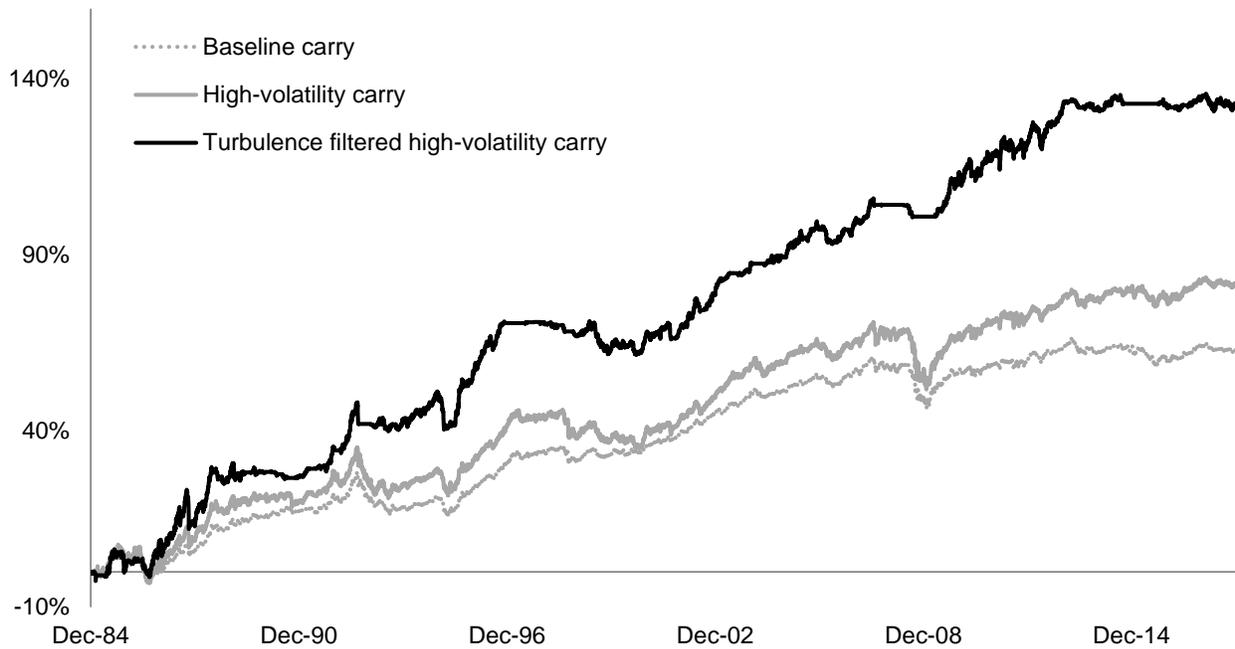
where y is a row vector of cross-sectional daily G10 currency spot returns versus the US dollar, μ is a row vector of daily average spot returns over the prior 3 years, Σ^{-1} is the inverse of the covariance matrix estimated from the prior 3 years of daily returns, and N is the number of currencies included in the calculation.¹¹ Turbulence may rise due to large magnitude returns in particular currency pairs, or because the patterns of returns diverge from their historical correlations. Kinlaw and Turkington [2014] show that correlation surprises of this type are incrementally predictive of future volatility and negative returns to the carry trade. We implement a backtest by taking a rolling 30 day average of turbulence scores, computing the percent rank of this moving average versus its own five year history, and investing notional values of 100, 75, 50, 25 and 0 percent in carry whenever the turbulence percent rank is above 0, 20, 40, 60, and 80, respectively. We include a two day implementation lag. Exhibits 9 and 10 show how the high-volatility carry strategy improves with the turbulence index filter. Once again, we stress that despite the lower information ratio of the high-volatility carry strategy in

the full sample, it performs more reliably than the baseline carry strategy. This is especially true in the period after the financial crisis.

Exhibit 9: Annualized backtest performance

	Baseline Carry	High-volatility Carry	Turbulence Filtered High-volatility Carry
Return	1.84%	2.40%	2.03%
Volatility	3.26%	4.73%	2.47%
Information Ratio	0.57	0.51	0.82
Turnover	0.41	0.51	1.68

Exhibit 10: Cumulative returns



Notes: We rescale the turbulence filtered high-volatility carry to have the same risk as the high-volatility carry. This facilitates comparison while preserving risk-adjusted performance.

Conclusion

The carry trade in foreign currencies is known for delivering positive returns on average, and for occasionally suffering sharp losses. In recent years, though, the typical approach to building carry portfolios has exhibited neither of these defining properties. We investigate features of carry returns that contributed to the strategy's performance in the decades prior to the 2008 global financial crisis, and we evaluate whether these features have changed in the years following the crisis. First, we find that the overall compression of interest rate differentials across currencies does not, on its own, explain the poor performance of the carry trade in the recent sample. Second, we find that much of the positive return in the pre-crisis sample occurred during times when the spot rates of carry positions were fundamentally undervalued. In contrast, the spot rates of carry positions have been overvalued for nearly the entire post-crisis sample. Third, we use a measure of network centrality derived from spot returns to detect crowding in both valuation and carry trades. We document a clear cycle in the pre-crisis sample in which positive carry returns initially align with the valuation factor, followed by periods of crowding in the carry factor and its eventual crash. Fourth, we find that this boom-bust cycle of valuation and carry occurs predominantly in a subset of currency pairs with the largest exchange rate volatility, and not in those with low volatility. The carry trade applied to high volatility pairs has substantially outperformed the carry trade in low volatility pairs from 2009 to 2017, suggesting that the carry trade remains viable if it is implemented with care.

Notes

We thank Jordan Alexiev and Mark Kritzman for helpful comments.

The material presented is for informational purposes only. The views expressed in this material are the views of the authors and are subject to change based on market and other conditions and factors; moreover, they do not necessarily represent the official views of State Street Global Exchange or State Street Corporation and its affiliates.

References

Burnside, Craig. 2012. "Carry Trades and Risk." Ed. Jessica James, Ian W. Marsh and Lucio Sarno. *Handbook of Exchange Rates*. Hoboken: John Wiley & Sons.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2011. "Carry Trade and Momentum in Currency Markets." *Annual Review of Financial Economics* 3 (1).

Brunnermeier, Markus K., Stefan Nagel and Lasse Pedersen. 2009. "Carry Trades and Currency Crashes." Ed. Kenneth Rogoff, Michael Woodford and Daron Acemoglu. *NBER Macroeconomics Annual*.

Daniel, Kent, Robert J. Hodrick and Zhongjin Lu. 2014. "The Carry Trade: Risks and Drawdowns." NBER Working Paper.

Dobrynskaya, Victoria. 2014. "Downside Market Risk of Carry Trades." *Review of Finance* 18 (5): 1885–1913.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan. 2018. "Deviations from Covered Interest Rate Parity." *The Journal of Finance* 73 (3): 915–57.

Engel, Charles. 1996. "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence." *Journal of Empirical Finance* 3.

Fama, Eugene F. 1984. "Forward and Spot Exchange Rates." *Journal of Monetary Economics* 14 (3): 319–38.

Froot, Kenneth A. and Richard T. Thaler. 1990. "Anomalies: Foreign Exchange." *Journal of Economic Perspectives* 4.

Hansen, Lars Peter, and Robert J. Hodrick. 1980. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy* 88 (5): 829–53.

Hodrick, Robert J. 1989. "Risk, Uncertainty and Exchange Rates." *Journal of Monetary Economics* 23: 433–459.

Hodrick, Robert J. and Sanjay Srivastava. 1984. "An Investigation of Risk and Return in Forward Foreign Exchange." *Journal of International Money and Finance* 3: 5–29.

Kinlaw, Will, Mark Kritzman, and David Turkington. 2012. "Toward Determining Systemic Importance." *The Journal of Portfolio Management* 38 (4): 100–111.

Kinlaw, Will, and David Turkington. 2013. "Correlation Surprise." *Journal of Asset Management* 14 (6): 385–99.

Kinlaw, Will, Mark Kritzman, and David Turkington. 2018. "Crowded Trades: Implications for Sector Rotation and Factor Timing." Working Paper.

Kritzman, Mark. 1997. "The Forward Rate Bias." *Economics and Portfolio Strategy*.

Kritzman, Mark. 2004. "The Forward Rate Bias: An Update." *Economics and Portfolio Strategy*.

Kritzman, Mark, and Yuanzhen Li. 2010. "Skulls, Financial Turbulence, and Risk Management." *Financial Analysts Journal* 66 (5): 30–41.

Kritzman, Mark, Yuanzhen Li, Sébastien Page, and Roberto Rigobon. 2011. "Principal Components as a Measure of Systemic Risk." *The Journal of Portfolio Management* 37 (4): 112–26.

Lustig, Hanno, Nikolai Roussanov and Adrien Verdelhan. 2011. "Common Risk Factors in Currency Markets." *Review of Financial Studies*, 24 (11).

Melvin, Michael, and Duncan Shand. 2017. "When Carry Goes Bad: The Magnitude, Causes, and Duration of Currency Carry Unwinds." *Financial Analysts Journal* 73 (1): 121–44.

Taylor, Alan, and Mark Taylor. 2004. "The Purchasing Power Parity Debate." *Journal of Economic Perspectives* 18 (4).

Verdelhan, Adrien. 2018. "The Share of Systematic Variation in Bilateral Exchange Rates." *Journal of Finance* 73 (2).

¹ For simplicity, we describe this relationship in terms of the country with a higher interest rate, which sells at a forward discount. This notion does not sacrifice generality, as each forward contract pertains to a currency pair and can be described equivalently in terms of either currency. Selling currency A forward versus currency B is equivalent to buying currency B forward against currency A.

² See also Hodrick [1989], Froot and Thaler [1990], and Burnside [2012] for surveys of the extensive literature in this area.

³ As of December 29, 2017, the drawdown is still ongoing.

⁴ We imply interest rate differentials assuming the covered interest parity condition holds. Even if this condition does not hold, the implied differentials are appropriate to inform an investment strategy because they represent the yield differential that will be earned on a forward contract, assuming the spot rate remains constant. Recent research from Du et al. (Forthcoming) documents some meaningful divergences from covered interest parity since 2008. These findings have more direct application to understanding the divergences between carry implemented with forward contracts and carry implemented on interest rates directly.

⁵ The exchange rates are quoted for major pairs, most often versus the US dollar. We use these rates to estimate cross rates between any two currencies. We use these cross-rates to arrive at desired portfolio weights, acknowledging that these implied cross rates may not represent actual rates available to trade in the market. After we compute portfolio weights, we aggregate currency positions and model trades at the rates quoted for liquid pairs versus the US dollar. Thus, the backtested returns represent realistic investment strategies. The spot rates, one-month and two-month forward rates are reported by WM/Reuters (London 4pm fix) and Barclays Bank PLC on Datastream. We use the German mark to proxy for the Euro prior to its introduction in January 1999.

⁶ For GDP-weighted carry, we derive GDP-relative weights from annual GDP data (expenditure approach) reported by the OECD and apply these weights to the aggregate currency positions. For optimized carry, we use mean-variance optimization to determine optimal weights across the nine currencies versus the US dollar. We use expected returns equal to each currency's one-month interest rate differential (versus the US dollar) and a covariance matrix derived from three years of trailing spot returns (versus the US dollar). We target an annualized volatility level of 3% and cap the absolute exposure to any currency at one sixth of the total portfolio (which is consistent with the equally-weighted strategy). All strategies are rebalanced at the end of each month. In Exhibit 1, the returns of all three approaches are rescaled to have the same cumulative return at the end of December 2008.

⁷ See Taylor and Taylor [2004] for a useful summary of perspectives on this topic.

⁸ We obtain annual PPPs (for GDP) from the OECD. These represent relative prices between a given country and the US. We use these to estimate relative prices between any two countries. Then, for each currency pair, we calculate its real exchange rate (RER) as the ratio between its nominal spot rate and its PPP measure. To generate a valuation signal, we simply subtract one from each pair's RER. This represents the degree of misvaluation relative to an RER equal to one. It may be possible to improve upon these simple measures of fair value by incorporating additional information.

⁹ Data for pairwise returns contains a degree of redundancy, as some currency pairs are linear combinations of others. This fact does not pose a problem for Principal Components Analysis (PCA), because the technique is well-suited to extracting information from matrices which are not full-rank. In other words, PCA allows us to retain the most important sources of variation while discarding variation that is more likely to reflect noise or that reflects redundancy.

¹⁰ We take the moving average of centrality to visualize trends more easily and to mitigate some noise that is present in daily rates during the 1980s.

¹¹ The turbulence calculation is not sensitive to the choice of base currency, because the multi-variate formula inherently accounts for volatility of foreign currencies against the chosen base currency as well as correlations between each pair of foreign currencies. Therefore, it is appropriate to estimate turbulence from returns denominated in any base currency and the results will apply to the currency universe overall.