RiskMetrics Group™ Research Monthly

Carried away?

Christopher C. Finger
chris.finger@riskmetrics.com

June 2007

Dubbed the "Trade of the Decade" by at least one website, it is difficult to imagine a single trading strategy getting more popular attention than the carry trade has over the last eighteen months. Headlines in early 2006 included "Japan's Boom May Explode Yen-Carry Trade" and "Yen Carry Trade to Unwind—Market Crash Alert". Fears rose again in early 2007: in "What keeps bankers awake at night?" the Economist made the carry trade first on its list. But the fears seem to have subsided, with the Economist acknowledging more recently that the carry trade may have gone "Out With a Whimper."

While carry trade is a generic term, the specific version that has garnered attention lately is the yen carry trade: a trader borrows in Japanese yen at low interest rates, and invests the principal in a higher yielding currency. As long as nothing happens (in particular, as long as the yen does not strengthen), the trader earns the interest rate differential (the carry) and is able to repay the yen loan with the principal from his US investment. The risk is that the yen does in fact strengthen, and the trader winds up using the carry (and possibly some of the principal) to repay the yen loan.

Because the yen carry trade would appear, at least anecdotally, to have been a source for so many profits, its impending death has been the source of just as much speculation. The recipe for disaster is as follows: the Japanese economy strengthens, the Bank of Japan responds by raising interest rates, the higher rates deter borrowing and encourage investing in the yen, and the yen appreciates. The death spiral then commences wherein traders react to the stronger yen by unwinding their positions, which requires purchasing yen, which in turn drives the yen higher for the next trader trying to unwind the trade, and so on. Simply put, the fear is that an outrageous number of traders, not wanting to miss out on the yen carry trade, are positioned similarly, and enough concern about higher yen rates will spark a stampede to get out of the trade, wreaking havoc on a much broader market in what we will call the Great Unwind Scenario.

For a risk manager, the recent chatter raises two questions: one specific and one general. The specific is if I am exposed to the yen carry trade, or at least to those market factors impacted by the Great Unwind, how can I get a sense of what might happen to my positions? The general is to what extent the yen carry trade truly is pervasive, and whether the Great Unwind Scenario really is as great as feared.

Can it be that easy?

The fears over the yen carry trade are the archetypical case for stress tests: the market thus far has not

reflected the fear (neither historical nor implied yen volatility has been elevated), and yet there is the distinct possibility that significant moves could occur. Moreover, history has done us the favor of providing an example of a large unwind of yen carry trades: in the aftermath of the 1998 Long Term Capital Management crisis, the yen appreciated from ¥134 to the dollar to ¥117 during the first week of October, and to below ¥110 by the middle of January 1999. To a large extent, this appreciation is now attributed to the widespread unwinding of the yen carry trade by funds seeking to curtail their leverage. To first order anyway, the investment side of the carry trades in such a period is irrelevant; as long as the market is flooding to pay back yen financing, the yen appreciates relative to most currencies, damaging any trade that is essentially short the yen. This historical period is relevant to any specific version of the yen carry trade.

So the naive view is that we are finished. We simply apply the market moves from a few select periods in late 1998 to our carry trade positions. But what does this really tell us? We must first recall that the events of 1998 were precipitated not by anything specifically Japanese, but rather a widespread desire to decrease leverage. So while these events are reminiscent of a large unwind, they do derive from a different catalyst from what is feared today.

More importantly, we need to consider the behavior of our trader throughout the course of the stress event. The October 1998 move was fast enough for us to consider instantaneous; it is reasonable to assume that we would have been simple victims of the market, with no possibility of significantly managing our exposures. The longer unwinding of the trade, however, appears to have occurred over five months, with the yen appreciating 30% from August 1998 to January 1999. Over this period, there would have been opportunities to pare positions slowly (or erroneously bet on a turnaround), and the decisions made during this time would have impacted our overall loss. Finally, by simply shocking our positions with an instantaneous currency move, we fail to examine to what extent the prevailing carry offsets our currency loss.

So for stress tests, beyond finding representative events, we must ask ourselves if we are exposed to a position or to a strategy. Particularly if the stress event we identify plays out over months, and not days, then our behavior during the event is an important driver of our exposure.

Positions and strategies

As discussed, the positions involved in the yen carry trade are simple: a short (or borrowing) position in JPY and a long (or investment) position in USD (or another higher yielding currency). Expressing a strategy takes a bit more care. Of course, the possibilities are endless, so we will focus on two examples to illustrate some of the decisions to be made.

We describe a strategy by creating a Net Asset Value (NAV) time series for a hypothetical fund which employs the strategy blindly. The first strategy—"Constant maturity"—begins with an initial amount of capital in a USD cash account, and proceeds as follows:

1. Each week, invest a fixed proportion of the current NAV in a three-month USD deposit. Fund this investment by borrowing in JPY for three months.

Figure 1: Performance of yen carry trade strategies, 1998–2007

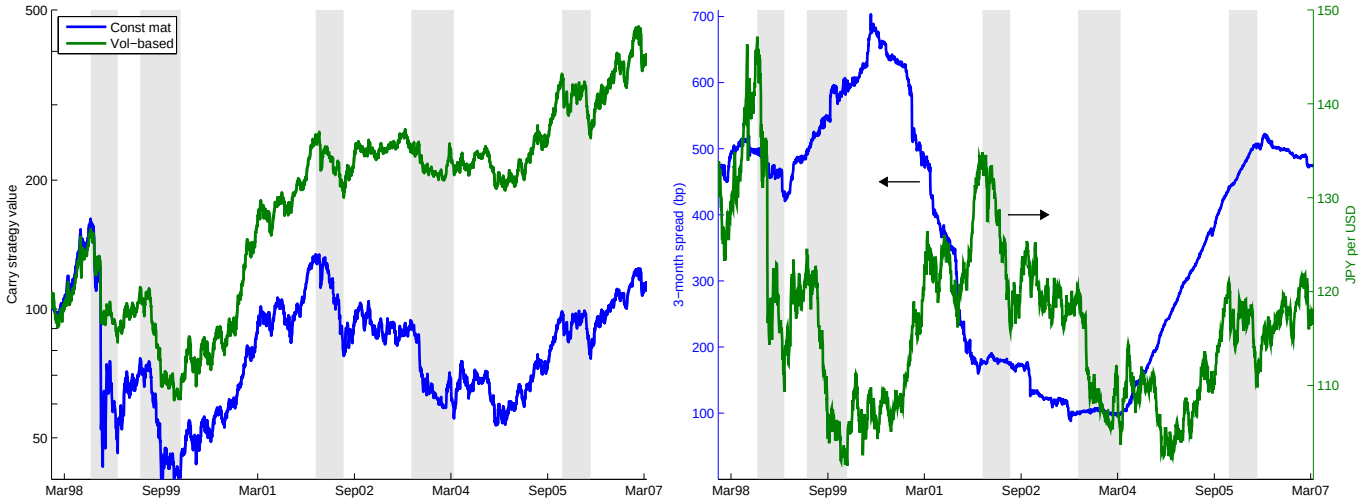

2. Hold each position until it matures, at which point the deposit is redeemed and cash is used to repay the loan. Place any net proceeds into the cash account. If at any point the cash is insufficient to repay the loan, the fund fails.

3. On a daily basis, calculate the fund NAV as the sum of the cash account and the mark-to-market value of the open carry trade positions.

Since all positions are held to maturity, and we invest a constant proportion in these positions, the leverage of the fund will be roughly[1] constant. For our examples, we choose the investment proportion such that the fund leverage is around four.

Our second strategy—"Volatility-based"—is similar, but with the addition of a bit of intelligence to the choice of both maturity and leverage. As an indicator of the attractiveness of the trade, we utilize the ratio of the interest rate spread to the implied volatility of the exchange rate. When we open a carry trade position, we choose the maturity where this ratio is greatest, and set our leverage proportional to the ratio. We close out each position after one week, and put on a new position. To achieve broadly the same risk profile as the constant maturity strategy, we set the proportionality constant for our leverage such that over our historical period (1998–2007), the average leverage of the fund is four.[2]

We present the performance of the two strategies, as well as the exchange rate and three-month interest rate spread, in Figure 1. Not surprisingly, both strategies exhibit their best performance over periods (2001–2002 and 2005–2006) of sustained yen depreciation. Overall, the volatility-based strategy shows a higher average return (124bp per month to 13bp for the constant maturity strategy) with lower standard deviation of returns (9% per month versus 13% for the constant maturity strategy). The great performance difference between the two funds comes from late 1998, when the volatility-

[1] As a reminder, all leverage measures are rough. See Finger (2006) for further discussion.

[2] Of course, this is something we could not guarantee a priori. To actually implement a strategy of this sort, we could only forecast what future leverage would be, since the actual leverage would fluctuate with the market.

Table 1: Historical stress tests on carry trade strategies

Start	End	Spread change (bp)	JPY appreciation	Instantaneous shock	Const. mat. strategy	Vol-based strategy
11-Aug-98	11-Jan-99	-64	30%	-100%	-72%	-46%
05-Oct-98	09-Oct-98	3	14%	-49%	-58%	-28%
19-May-99	03-Jan-00	101	21%	-74%	-47%	-45%
8-Feb-02	15-Jul-02	-3	15%	-52%	-42%	-28%
04-Aug-03	01-Apr-04	-1	15%	-52%	-41%	-13%
06-Dec-05	17-May-06	65	10%	-33%	-22%	-29%

based strategy employed reduced leverage, and from 2002–2003, when three-month spreads were quite low, and the volatility-based strategy put on positions with longer maturity. The source of the risk in the two funds is almost identical: both funds show a correlation of over 90% to the exchange rate.

Stress testing

To illustrate the impact of the strategies on stress tests, we examine six historical periods with significant yen appreciations, two of which cover the 1998 events. The stress periods are indicated by the gray bars in Figure 1. For each period, we calculate three stress tests. The first is to apply an instantaneous shock, applying the total market move of the period to a three-month carry trade position with leverage of four. The other stress tests involve following the daily market moves throughout the period, applying the strategies of the two funds described above. The results are presented in Table 1.

The first period—the largest yen appreciation—would have ruined the fund as an instantaneous shock, but as this move occurs over five months, the losses on the actual strategies are smaller; still, this period represents the worst loss for all three stress test schemes. Other periods generate notable differences. The one-week shock of October 1998 produces large losses for the instantaneous scheme and the constant maturity strategy, but the volatility-based strategy, by reducing its leverage in the upheaval of August and September, escapes the worst of this move. In fact, the volatility-based strategy suffers a greater loss in the last historical period, which actually comprises the smallest of the yen appreciations.

The lesson here is not that one strategy is better than another. The lesson is that how we assume a trader reacts to market events can have a significant impact on how we interpret stress tests, and in our case even changes which historical periods look the worst or most benign for our carry trade exposures. Particularly as we examine stress scenarios that play out over longer periods, we need to concern ourselves with how our portfolio would react to such a scenario along the way.

This brings us to one last caveat, which is that we have assumed that our strategies can be successfully implemented, even during a stress period. For the constant maturity strategy, this is probably realis-

tic; positions are held to maturity, and new positions (further investment in the carry trade) are in the opposite direction of the market's putative unwinding, where we might expect constrained liquidity. Under the volatility-based strategy, we close carry trade positions each week, leading to reduced exposure as the trade appears less attractive. Since presumably many other market participants would be trying to reduce positions at the same time, depending on the absolute size of our position, we would risk insufficient liquidity to support our strategy. To be more realistic, we might adjust our model strategy by imposing constraints on how quickly we could close losing positions.

How much should we worry?

While we have discussed how we might examine a possible Great Unwind Scenario, we have not addressed the likelihood of such an event. This is an issue on which the media seems to change its mind regularly, and regulators struggle with assessing the size of the market's position in the trade. If an unwind scenario is a sort of financial armageddon, then we should ask how large the positions are that could be unwound. And yet such an answer is elusive.

Part of the trouble in assessing the size of the market's carry trade position is that it is not a single trade, but really any position where yen borrowing is used to finance a high yielding investment. Moreover, both the investment and the financing are likely to be over-the-counter positions. And certainly not every loan in yen is to finance a foreign currency position. So how do we size the problem?

One approach is to estimate the quantity of carry trade positions, or at least changes in the quantity, by examining aggregate yen borrowing. Net short yen futures positions on the IMM peaked in January 2007, dropped in April, and peaked again in May. Banking statistics on yen-denominated borrowing reflects some indication of heavy carry trade activity in 2005, but less indication recently.[3] So the picture is mixed, and possibly suggests that some unwinding has already occurred, but we still do not know who owns these positions and why, and whether that would lead to all or some of these positions being candidates for the Great Unwind.

Our search for the likely protagonists in the Great Unwind Scenario takes us naturally to hedge funds. If hedge fund returns are unusually dependent on the carry trade, the logic goes, then there are a greater number of traders apt to rush to unwind should the trade start to appear unprofitable. Our approach, then, is to examine the relationship between hedge fund returns and our carry trade strategies. There is of course a significant logical leap from correlated returns to the prospect of a panic in the yen, but the nature of hedge fund returns is nonetheless a useful complement to the statistics quoted above.

To infer the extent of hedge fund exposure to the carry trade, we perform rolling 24-month regressions of hedge fund indices on a combination of standard fixed income factors[4] and the returns on a carry trade strategy. For the carry trade strategies, we use the aforementioned volatility-based strategy, in which we invest in US interest rates, plus similar strategies where we finance in yen and invest respectively in Australian, Brazilian, and Taiwanese inter-

[3] See Goodman (2007) and McGuire and Upper (2007).

[4] The Lehman Aggregate, Lehman High Yield, and EMBI+ indices

[5] For the AUD strategy, we set an average leverage of four, similar to our USD strategy, but for the BRL and TWD strategies,

Figure 2: Regression coefficients for yen carry trade strategies, rolling 24-month regressions

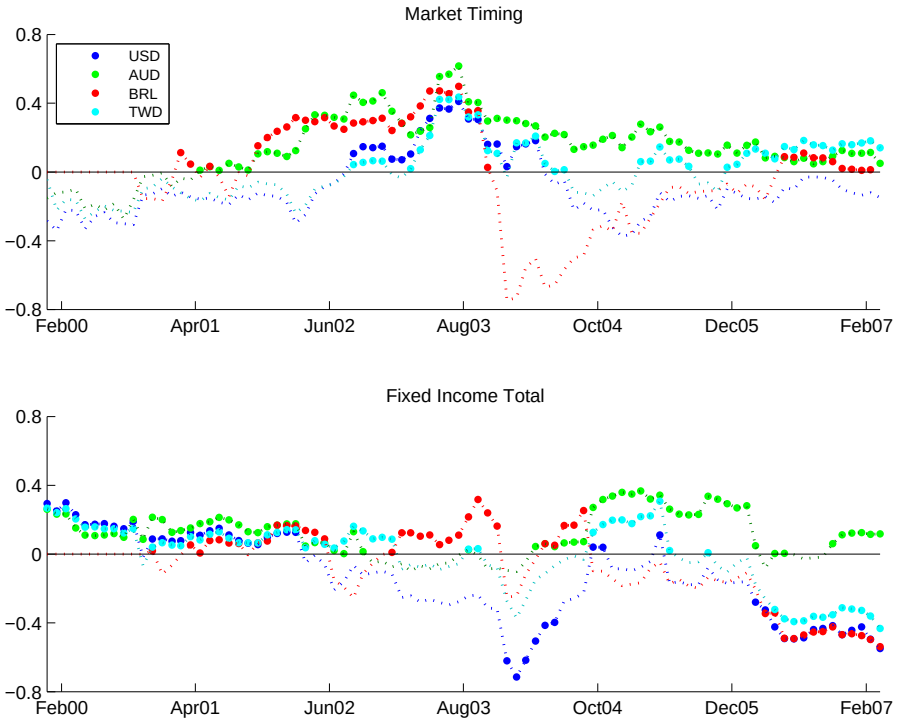

est rates.[5] We do not regress with all of the carry trade series at once, but rather run four separate regression exercises, each using the standard factors plus one of the carry trades.

We note that the choice of the volatility-based strategies over the constant maturity ones in fact has little impact on our results, as the correlation between the two strategies for a given currency is quite high, and we focus simply on whether the carry trades appear to drive the hedge fund returns. Where the choice of strategy does play a role is in benchmarking. The constant in the regression (or alpha) may be interpreted as the excess return of the hedge fund index over the standard factors and carry strategies. Since our two carry trade strategies do show a difference in average return, the alpha estimate varies as well.

We present our results in Figure 2. In each chart, we plot the regression coefficient for each of the carry trade strategies, and emphasize the points in time where the coefficient was statistically significant at 95% confidence. We choose two hedge fund indices—one for market timing strategies and a second for total fixed income strategies—whose results are representative of the behavior we see across the various fixed income-related indices. Our analysis does suggest a general positive exposure to the various carry trades, consistent with our intuition.

The HFRI Market Timing index shows no significant dependence on any of our carry trade strategies until late 2001. From that point forward, the index does appear to derive a significant portion of its returns from the various carry trades; the level of reliance on the trade peaked in 2003, and has

we set an average leverage of two.

been lower, but positive and relatively stable, since 2005. The most consistent dependence seems to be on the AUD investment strategy, though the most recent data exhibits correlations to TWD and BRL as well. The HFRI Fund of Funds Composite and Global Macro indices demonstrate qualitatively similar behavior, with consistent overall positive dependence on the carry trades that persists through the most recent data.

The story for the HFRI Fixed Income Total index is somewhat different. It shows a consistent reliance on the carry trades until 2002, and then peaks of exposure to the BRL in 2003 and the AUD in 2004–2005. Since 2005, there appears to be significant negative exposure to three of our strategies. This might indicate that the funds comprising this index have positioned themselves to profit from a possible unwind scenario; in doing so, however, they would have missed out on what has been a profitable period for the carry trade strategies. Other fixed income HFRI indices (including Fixed Income Arbitrage and Fixed Income Diversified) displayed comparable recent behavior.[6]

Where does that leave us?

Our results, far from definitively settling the issue of the market's exposure to the carry trade, rather add to the body of inconclusive evidence. There are categories of hedge funds that do appear to still be relying on the carry trade, though no more so than they have for the last several years, while there are other categories of hedge funds that appear positioned oppositely. A sharp appreciation in the yen will hurt some, and cause some rush to unwind, but there appear to be plenty of buyers in an unwind scenario. Interesting times, yes, but this does not represent the large systemic exposure that is most feared.

Returning to stress testing, if we must first assess the likelihood of the Great Unwind Scenario, the "yellow light" suggested recently by Bank of America seems most appropriate. The second piece of stress testing—understanding how the scenario might affect us—is subtle. If history is a guide, then the scenario we fear is not one that occurs in one day, but rather takes weeks or months to unfold. How we, or someone managing on our behalf, might react over the course of such a scenario may dramatically alter the outcome, to the point that the scenario we should most be concerned with may actually depend on the behavior we assume. We will never know this behavior in advance, but we can express rules, such as the strategies presented here, that might characterize what we would experience.

Further reading

- Finger, C. (2006). Exposed!, Research Monthly, June.
- Goodman, L. (2007). Switch on yellow lights for the "carry trade"—not red. *Bank of America Forex Focus*, May 30, 3–4.
- McGuire, P. and Upper, C. (2007). Detecting FX carry trades. *BIS Quarterly Review*, March, 8–9.

[6] McGuire and Upper (2007), in a similar exercise, also noted negative dependence of hedge fund indices on some carry trade returns.