



Stochastics and Statistics

Dynamic allocations for currency futures under switching regimes signals

Lorenzo Reus^{a,*}, John M. Mulvey^{b,1}^a Science and Engineering Department, Universidad Adolfo Ibañez, Diagonal Las Torres 2640, Santiago, Chile^b Operations Research and Financial Engineering, Sherrerd Hall, Princeton University, Princeton, NJ 08540, United States

ARTICLE INFO

Article history:

Received 29 September 2014

Accepted 12 February 2016

Available online 20 February 2016

Keywords:

Investment analysis

Currency futures

Carry trade

Regime identification

Mean-semivariance portfolio optimization

ABSTRACT

Over the last decades, speculative investors in the FX market have profited in the well known currency carry trade strategy (CT). However, during currencies or global financial crashes, CT produces substantial losses. In this work we present a methodology that enhances CT performance significantly. For our final strategy, constructed backtests show that the mean-semivolatility ratio can be more than doubled with respect to benchmark CT.

To do the latter, we first identify and classify CT returns according to their behavior in different regimes, using a Hidden Markov Model (HMM). The model helps to determine when to open and close positions, depending whether the regime is favorable to CT or not. Finally we employ a mean-semivariance allocation model to improve allocations when positions are opened.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

A currency future (CF) is a popular contract in the FX market. It is an agreement between two parties to exchange one currency for another at some future date at a price fixed on the purchase date. CF allows companies to hedge against currency risk, and it allows speculators to bet on exchange movements with less up-front investments and high leverage.

There are multiple factors that can affect exchange movements. Thus it is hard to derive a model with good prediction powers, especially for the short run. For example, [Cheung, Chinn, and Pascual \(2005\)](#) or [Kilian and Taylor \(2003\)](#) tested several models and found that none could consistently beat random walk forecasts. According to them, U.S. exchange traders think that economic fundamentals are more important at longer horizons, while short-run deviations from the fundamentals are attributed to excess speculation and institutional customer/hedge fund manipulation.

A popular CF strategy among financial investors is the carry trade (CT). CT rests on foreign interest rates. It consists of taking short (long) positions in CF from low (high)-yielding currencies. Hence, this strategy expects that a currency will appreciate/depreciate when the interest is high/low in relation to the other

currencies. The strategy is net zero, that is, the total value in short positions is equal to the total value in long positions. More about CT strategies performance and details can be found in ([Burnside, Eichenbaum, & Rebelo, 2008](#); [Jylha & Suominen, 2009](#)) and ([Galati, Heath, & McGuire, 2007](#)).

In this paper we develop CF strategies that can outperform CT. To achieve this, we first see if CT can be classified in regimes, just like many other markets do. The cyclical behavior of markets facilitates the search for regimes that emerge in different periods of time. The aim and success of this classification method is determined by the heterogeneity of the regimes. In the present case, we look for different behavior of CT returns. As explained below, the regime detection is based on a machine learning process, rather than economic fundamentals. Once regimes are identified, we will be able to create a signal for closing and opening positions depending whether the current regime is conducive to CT strategy or not. The idea is to hold onto a CT strategy when the signal is on and stay out when it is off.

One successful model for regime detection is the Hidden Markov Model (HMM). It assumes data behaves as a Markov process, but states are not observed. In this case, CT returns are the observed data, while regimes are the hidden states. Applying HMM in finance is nothing new. Pioneering work in this field was done by [Hamilton \(1989\)](#), who identified U.S. economy cycles with the gross national product series. [Guidolin and Timmermann \(2007\)](#) utilize a four-state HMM for a series of stock returns. Recently, [Prajogo \(2011\)](#) built an HMM for an agribusiness index.

* Corresponding author. Tel.: +56 961478056.

E-mail addresses: lorenzo.reus@uai.cl (L. Reus), mulvey@princeton.edu (J.M. Mulvey).¹ Tel.: +1 6092585423.

Bae, Kim, and Mulvey (2014) did HMM regime identification on equity, commodities and bond indexes, also using a four-state model. More applications of HMM in Finance can be found in (Mamon, 2014) and (Bhar & Hamori, 2006).

In a second step, we examine ways to allocate long and short positions during periods when positions open, i.e., during periods when carry trade opportunities work in general. CT strategies are based on simple allocation rules; for example, to take equally-weighted long/short positions in the highest/lowest K interest rates currencies (see Deutsche Bank G10 Currency Harvest Fund).

To improve those simple allocation rules, we therefore allocate with a mean-semivariance approach. We use this approach for several reasons. Semivariance can be considered as a downside risk measure because it belongs to the lower partial moment (LPM) of degree two (Fishburn (1977) for more details about LPM). Markowitz (2010) shows that semivariance can measure investor's utility toward risk, and hence suggests to use the geometric mean-semivariance. We prefer to use this risk measure instead of the alternatives, such as the conditional value at risk (CVaR), because although they are widely used in practice, many of them are not considered risk measures because they are LPM of degree one (CVaR) or zero (VaR). In fact, Markowitz (2010) rejects CVaR and VaR as risk measures. For more mean-semivariance applications see Huang (2008), Ballestero (2005), Vasant, Irgolic, Kruger, and Rajaratnam (2014) and Estrada (2007).

The present work is organized as follows. Section 2 analyzes the performance of standard CT. Section 3 shows HMM fitting results for CT returns. Section 4 implements and tests HMM based strategies. Section 5 tests the same HMM-based strategies, but adding mean-semivariance allocation on openings. Finally, Section 6 concludes and points to possible future research.

2. Carry trade performance

When investing in CF at time t , we buy (or sell) foreign exchange at a rate of F_t^T at some date $T > t$. Then at time T , we can earn a profit of $S_T - F_t^T$ ($F_t^T - S_T$) if we sell (buy) that amount on the market. F_t^T is priced according to the covered interest parity. If i and i_c denote the interest rate at time t of home and foreign country respectively, and S_t the current spot rate (units of home currency in one unit of foreign currency), then:

$$F_t^T = S_t \exp((i - i_c)(T - t))$$

The return on CF investment $r_t^T := \frac{S_T}{F_t^T} - 1$ is approximated by

$$r_t^T \approx \log \frac{S_T}{S_t} + (i_c - i)(T - t) \quad (1)$$

To include transactions costs, we denote S_t^{bid} and S_t^{ask} as the bid and ask prices respectively. For long CF, returns in (1) can be expressed as

$$r_t^T = \log \frac{S_T^{bid}}{S_t^{ask}} + (i_c - i)(T - t) \quad (2)$$

Analogously, the return for short CF

$$r_t^T = -\log \frac{S_T^{ask}}{S_t^{bid}} - (i_c - i)(T - t) \quad (3)$$

As previously mentioned, CT policy is based on ranking interest rates among the pool of currencies available. As done in (Brunnermeier, Nagel, & Pedersen, 2008), we buy (sell) equal amounts of CF in the K highest (lowest) interest rates. To test this strategy, we take interest rates differentials (relative to the U.S. interest rate) embedded in forward rates from the G10 currencies. The performance of quarterly rebalancing between January 2000 and June 2015 (62 quarters) is presented in Table 1. The period

Table 1

Benchmark carry trade performance for different values of K . Semi-volatility corresponds to the square root of semivariance. MSV is the ratio between geometric mean and semivolatility. Data are taken from Bloomberg database and was available from Jan 96.

K	Geometric mean (percent)	Volatility (percent)	Semi volatility (percent)	MSV (percent)	Skew	Kurtosis
1	1.8	13.9	10.6	17	-2.6	61.2
2	2.3	11.4	8.5	27.3	-0.7	9.7
3	1.4	9.3	6.9	20.1	-0.5	10.0
4	0.4	7.6	5.7	7.0	-0.6	7.9
5	0	6.5	4.9	0	-0.6	7.8

January 1996 to December 1999 was assigned for training performance as explained in next sections.

The table shows that CT has been profitable on average. This aligns with the known forward premium puzzle, which basically shows empirically that F_t^T is a poor estimator for S_T . Otherwise, $r \approx 0$ for each currency and hence CT returns would have been close to zero too. The puzzle is analyzed and explained in (Bansal & Dahlquist, 2000; Fama, 1984; Pippenger, 2011) and (Reinert, Rajan, Glass, 2009). If we increase K , we observe diversification effects. The best compromise between profit and downside risk, measured by the mean semivariance ratio (MSV), is when we go long with two currencies and short with another two. The presence of negative skewness means there can be huge losses across periods.

To have some insight of the most active currencies, we look at Fig. 1, which show some metrics about the weights for CT portfolio in time. The New Zealand Dollar (NZD) and the Australian Dollar (AUD) are the main target currencies (where we go long), followed by the Norwegian Krone (NOK). Oppositely, the Swiss Franc (CHF) is always a funding currency. Other funding currencies used are the Japanese Yen (JPY), Swedish Krone (SEK) and Danish Krone (DKK). The rest of the currencies are rarely used and can be either target or funding in time. The best MSV was obtained $K = 2$, meaning that carry trade is mainly constructed with four currencies (NZD, AUD CHF and JPY), followed by three other currencies (SEK, DKK and NOK).

3. Regime construction

The problem with the previous benchmark strategy is that it holds a fixed portfolio for a quarter of a year. No actions are taken when the strategy produces losses. Therefore, we construct an HMM model for CT in order to take early actions. We expect HMM to signal when conditions are unfavorable to CT.

When fitting an HMM, we need to estimate a transition probability matrix P , an initial probability vector and parameters related to the distribution of the data. When the distribution is assumed to follow a normal distribution, the HMM can be calibrated with the Baum Welch Algorithm. For details and explanations about HMM fitting and Baum Welch algorithm see Fraser (2008) and Prajogo (2011).

It's sensible to ask why to fit the HMM with a normal distribution when Table 1 is showing returns distribution is not fitted with such. As emerges below, the strategies proposed in this paper will not depend explicitly on the mean and volatility values of the normal distribution. Rather, they depend on the qualitative characteristic of the regime and the transition probabilities between regimes. The reason to use a normal distribution for the HMM construction is that we have a known algorithm to calibrate the HMM. The assumption for using normal distribution is supported by the discussion in (Mulvey & Zhao, 2011) (pg. 21), about how normal distribution can handle fat-tail distributions properly when applied to HMM on weekly data. Besides, the

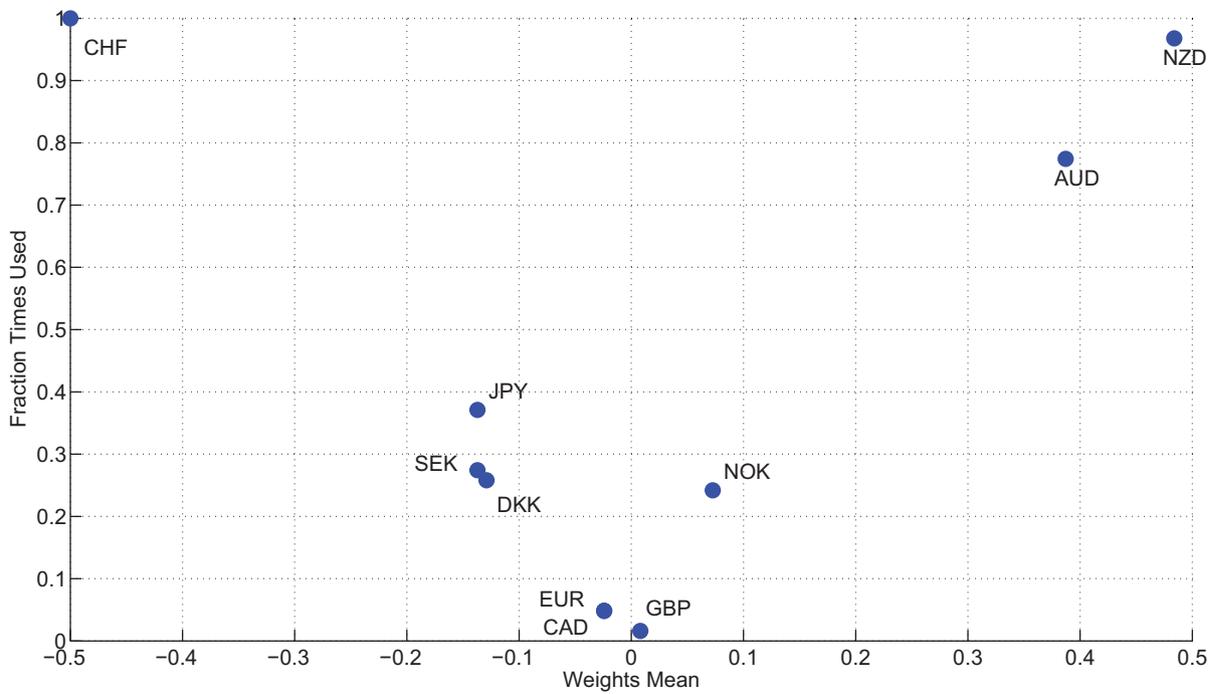


Fig. 1. Allocation composition for the benchmark CT(K = 2) portfolio.

methodology calibrates HMM as new data become available, so it adapts to possible non-stationarity between mean and variance. Whitelaw (1994) points out the problem with the latter. He shows there is non-stationarity between mean and volatility for stock returns, in which case aggregate distribution can be non-normal, even if returns are normal for shorter periods. Finally, any proposed strategy will account for fat tails and downside risk in case returns are not normally distributed.

To fit the HMM, we define first the CT returns so to choose (K). The best performance was with $K = 2$, so we use this strategy. Next, we define the number of states to model the chain. A common criteria used is the Bayesian Information Criteria (BIC), which is a compromise between the likelihood and the number of parameters to be fitted. Another criterion is to look at P structure. For example, it is sensible to be barred from switching between completely opposite regimes when states are more than two. We expect the latter to happen by going through a transition regime. This method is used in (Prajogo, 2011) and Bae et al. (2014).

To train the HMM model, CT returns ($K = 2$) are computed daily to produce sufficient data. Note that (2) (and hence (3)) can be written as the sum of daily returns r_h^{h+1} . If t and T represent two dates in time, then

$$r_h^{h+1} \begin{cases} \log \frac{S_{h+1}^{ask}}{S_h^{ask}} + (i_c - i)h & t \leq h < T - 1 \\ \log \frac{S_T^{buy}}{S_{T-1}^{ask}} + (i_c - i)h & h = T - 1 \end{cases} \quad (4)$$

Fitting results for the HMM model are shown in Figs. 2 and 3. HMM is retrained at every quarter as new data develop. Overall results suggest choosing two regimes until the 2008 financial crisis and then three states.

BIC results are not decisive in determining the number of states, since they are similar in every period. What makes the difference are the values from the transition probabilities, means and volatilities. Before the 2008 crisis, we can clearly see the difference between two regimes. We refer to a “good” (“bad”) state for the state with highest (lowest) mean and lowest (highest) volatility. When using three regimes for periods prior to the 2008 crisis, the third regime looked like a “good” regime.

However, we can identify three different regimes after the 2008 crisis. We label “Transition” the state that moves between the other two in terms of performance. Something not shown in the figures is that for the three state HMM, transition probabilities from “good” to “bad” and vice versa are zero. This is why we call it such, since we need to go through it when going from first to last state and vice versa.

If we look at values of a particular regime in time, there are no major changes except after the 2008 financial crisis. For example, there is a volatility increase for the “bad” regime. In any case, such a change in parameters is expected after such an event. On average, CT strategies had experienced unprecedented losses, the causes of which had been widely discussed. Investors were looking for refuge in more traditional (safe haven) currencies, which in general are used for funding. Therefore, target currencies depreciated in relation to target currencies, leading to losses for CT investment strategies. For more details on safe haven currencies during the 2008 crisis, see McCauley and McGuire (2009), and Rinaldo and Söderlind (2010).

Notice deviations around the mean are considerable, even for the good regime. Therefore, we can still have negative returns for that state. Transition probabilities to the same regime are consistently high, supporting regime existence. When all data are used for calibration, steady state probabilities for the two-states model says that the “bad” regime is visited only 11 percent of the time. For the three-states model, “bad” and “transition” regimes are visited only 34 percent of the time. All these results align with previous benchmark CT results: CT’s are profitable on average, but cause huge losses when they are not.

4. HMM-based carry trade strategy

Regime identification with HMM allows us to design decision rules to close/open CF positions. The rule should follow this idea: We open (close) positions when we have been most likely in the “good” (“bad”) regime during the last V consecutive days. Positions are opened according to the benchmark CT rule.

Strategies with smaller V’s react faster when HMM signals a regime change. That might prevent large drawdowns when CT does

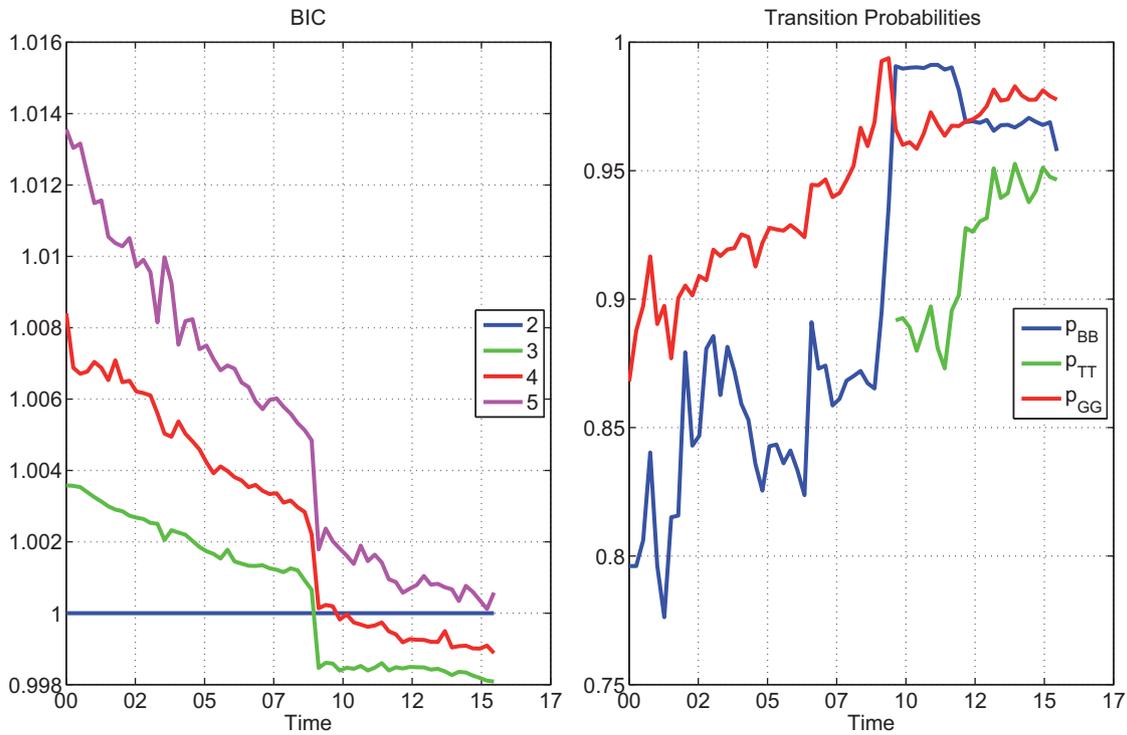


Fig. 2. Quarterly HMM fitting results for daily CT returns ($K = 2$). Left: BIC for different number of states each BIC is standardized (divided) by BIC with 2 states. Right: Transition probabilities with two states.

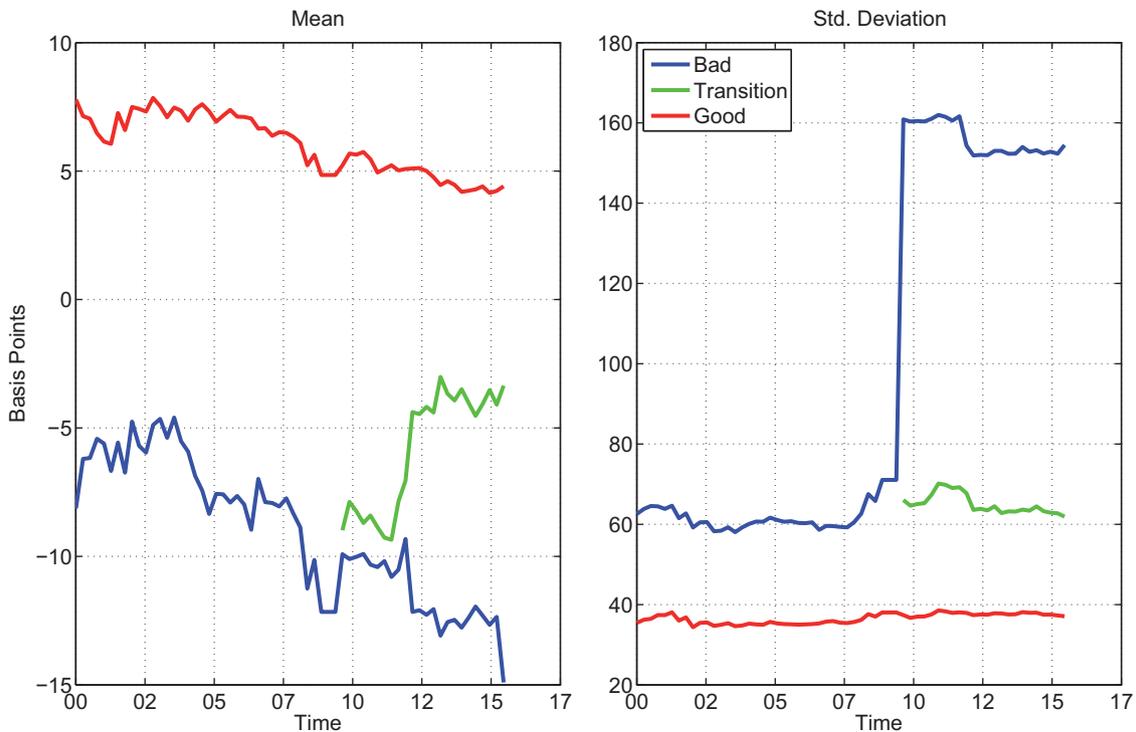


Fig. 3. Quarterly HMM fitting for annualized mean and volatility for daily CT returns ($K = 2$). HMM has two states until 2008 financial crisis and three afterwards.

not work and so avoid losing profit when CT does work. The problem with smaller V 's is that rebalancing frequency increases, which can lead to higher transaction costs. Therefore the right choice of V is not straightforward, and strategies with different V 's should be compared.

To know what is the most likely current regime each day, the Viterbi algorithm is used. This algorithm looks for the most probable sequence of states given the returns observed. The last point of the path corresponds to the current regime. This algorithm is based on a dynamic programming approach and comes with most

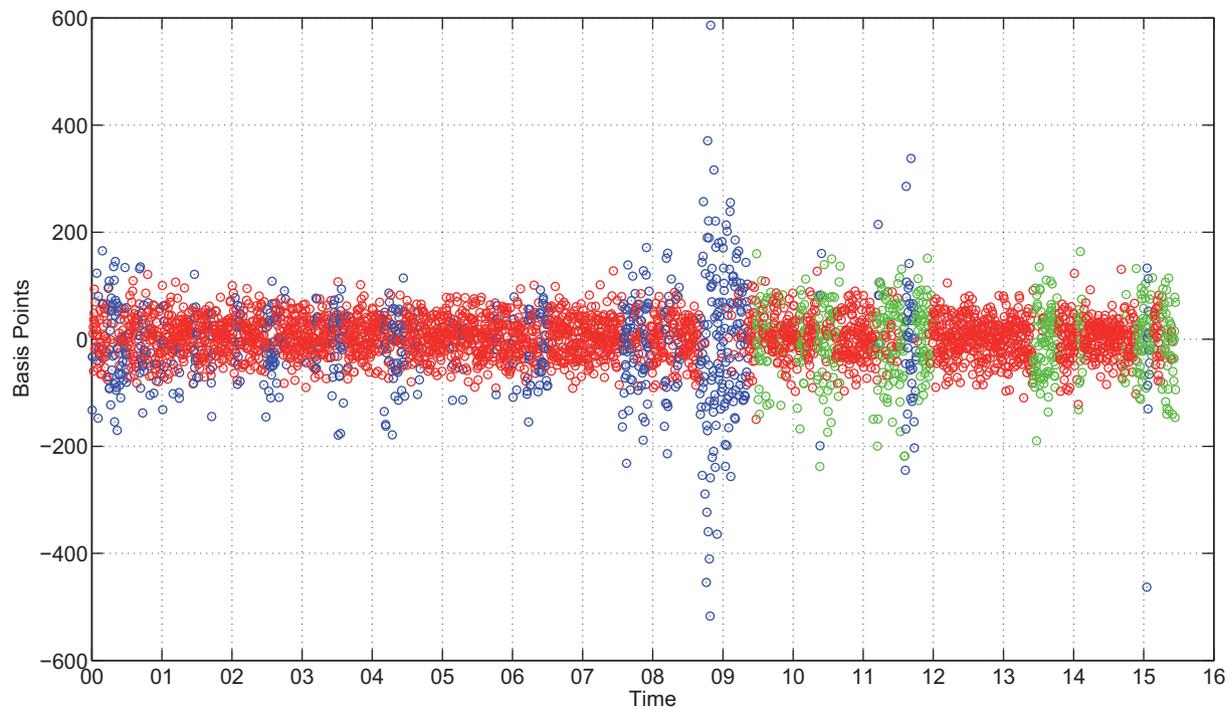


Fig. 4. Daily returns in time, classified by current regime. The blue, green and red marks corresponds to bad, transition and good regime respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

HMM fitting packages. To see how Viterbi works and is used in HMM, see [Viterbi \(1967\)](#), [Fraser \(2008\)](#) and [Prajogo \(2011\)](#).

[Fig. 4](#) shows daily returns (basis points) and the respective current regime determined by Viterbi's algorithm. Almost 70 percent of periods were considered to be in the "good" state and only 17 percent in the "bad" state. Knowing the most likely regime, we can create a signal to open or close positions. We close when the current regime has been the "bad" state for V consecutive years. We open position when the current regime has been the "good" state for V consecutive days. More precisely, a pseudo code for the HMM-based strategy is the following. Denote:

$curr_d$: Regime found by Viterbi algorithm for current day d .

$open_d$: $\begin{cases} 1 & \text{if positions are opened at day } d \\ 0 & \sim \end{cases}$

```

For  $d = 1 : D$ 
   $v \leftarrow 0$ 
  if ( $open_d = 1$ ) {
    if ( $curr_d == \text{"bad"}$ )
      if ( $v = V$ ) Close positions,  $open_{d+1} \leftarrow 0$ 
      else  $v++$ 
    else  $v \leftarrow 0$ 
  } else {
    if ( $curr_d = \text{"good"}$ )
      if ( $v = V$ ) Apply CT( $K = 2$ ) rule,  $open_{d+1} \leftarrow 1$ 
      else  $v++$ 
    else  $v \leftarrow 0$ 
  }
End For

```

We assume there is a full day to open or close positions. To discern the best time to trade during that day is beyond the scope of this work. [Table 2](#) depicts the amount of times positions are opened, as compared to the 62 times we have with quarterly allocations.

Table 2

Number of times positions are opened with different V values.

V	1	2	3	4	5	6	7
# Times	122	78	58	50	38	32	28

Table 3

Signal-based strategies performance for different V values. Period Jan00–Jun15. Measures are annualized.

V	Geometric mean (percent)	Volatility (percent)	Semi volatility (percent)	MSV (percent)	Skew	Kurtosis
1	1.1	7.5	5.6	20.1	-0.6	3.6
2	1.2	7.5	5.5	20.8	-0.5	3.5
3	2.1	7.4	5.5	38.8	-0.4	3.3
4	2.0	7.5	5.5	35.9	-0.4	3.5
5	3.1	7.5	5.5	56.4	-0.4	3.6
6	2.5	7.5	5.5	45.9	-0.4	3.7
7	2.9	7.5	5.5	52.2	-0.4	3.6

[Table 3](#) shows the performance of these HMM (signal)-based strategies on the same period tested for benchmark carry trade strategies in [Table 1](#). The best results in terms of MSV are obtained with $V \in \{5, 6, 7\}$, which, outperform all benchmark CT strategies. Improvement comes both from higher geometric mean returns and risk reduction (semi-volatility). Note this risk reduction is equivalent for any of the strategies, since they all have the same semi-volatility of 5.5 percent. Besides, skewness and kurtosis are lower than benchmark returns.

[Fig. 5](#) depicts P&L of best benchmark strategy against the best three HMM-based strategies. Performance is better for these signal-based strategies at almost any starting period and time length chosen. To see this, [Table 5](#) visualizes this finding by comparing the MSV between CT($K = 2$) and HMM ($V = 5$) for different time windows. For example, when the period January 2000 to June 2015 is divided into five equal length sub-periods, the HMM

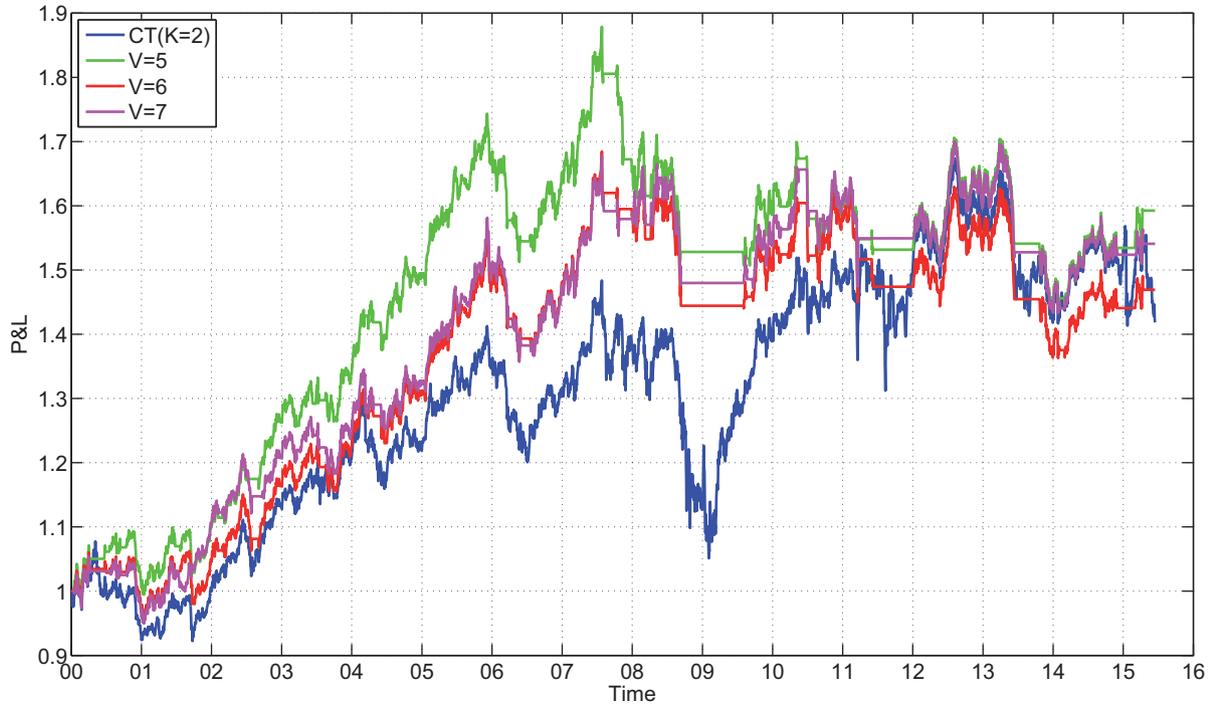


Fig. 5. P&L between CT strategies ($K = 2$) and HMM-based strategies.

strategy outperforms the benchmark at every time period. When the same period is divided in 10 or 15 subperiods this happens 80 percent of times.

As a remark, Fig. 5 helps to see the effects of the Euro–Swiss peg in September 2011. Recalling Fig. 1, CHF is a main funding currency. Prior to this event, CT($K = 2$) had bad performance since CHFUSD increased considerable. However, the peg suppressed the currency and consequently benefited it. The rest of the strategies are on hold during these periods and are not affected by the peg. The peg was abolished unexpectedly in January 2015 and caused a significant increase in CHFUSD which harmed the CT($K = 2$) strategy. Again, the new strategies were not sensitive to this announcement.

5. Mean-semivariance allocation

Although the last results are promising, they still leave room for improvement. Even if previous strategies outperform the benchmark, Fig. 5 shows that at the end we have lost some of the advantage produced in the first eight years. Up to this point, we have applied a very simple allocation rule and we have not done anything in this direction. In this section we show how to improve performance by allocating with other rule when signal turns on.

To put this discussion into context, we are dealing with a currency future portfolio. If allocated at time t , the portfolio return at time T is given by $x'r_t^T$, where r_t^T is the return vector given by (1) and x is the net allocation at each CF. This vector can be decomposed in long x_+ and short x_- positions, since CT is a net zero strategy ($1'x = 0$). Thus define feasible region P :

$$P = \{(x, x_+, x_-) : 1'x_+ = 1, 1'x_- = 1, x_+ \geq 0, x_- \geq 0, x = x_+ - x_-\}$$

The idea is to improve the MSV of the strategy. So if positions are opened at time t , we allocate according to the following problem

$$\max_{x \in P} E\{\log(1 + x'r_t^T)\} \tag{5}$$

$$s.t \ E\{\min\{x'r_t^T - E(x'r_t^T), 0\}^2\} \leq \theta \tag{6}$$

That is, (5) maximizes the geometric mean return and (6) keeps semi-variance below threshold θ . Note vector S_T is random in two ways: Its value is uncertain at T and T is also uncertain itself, since we do not know when we close positions.

We prefer not to use the classical mean-variance allocation because exchange returns on good regime, given by $\log(\frac{S_{t+1}}{S_t})$, are empirically abnormal. (Table 6 holds more information.) Hence, not all risk is captured by volatility. As a remark, Table 7 in the appendix shows how highly correlated exchange returns correlate during a good state. The only one showing low correlation with all others is the JPY.

Denoting Δi_t as the vector of interest rates, differences between each currency and USD at time t , the problem can be written as

$$\max_{x \in P} E\left\{\log\left[1 + x'\left((T - t)\Delta i_t + \log\frac{S_T}{S_t}\right)\right]\right\}$$

$$s.t \ E\{\min\{x'((T - E(T))\Delta i_t + \log S_T - E(\log S_T)), 0\}^2\} \leq \theta$$

5.1. Sample-average approximation

To solve the previous optimization problem, we use the sample-average approximation. That is, we sample different values for time T and S_T to compute expected values as average. Specifically, denote T_n and S_T^n the values of T and S_T in scenario n , we solve:

$$\max_{x \in P} \sum_{n=1}^N \left\{ \log\left[1 + x'\left((T_n - t)\Delta i_t + \log\frac{S_T^n}{S_t}\right)\right]\right\}$$

$$s.t \ \frac{1}{N} \sum_{n=1}^N y_n^2 \leq \theta$$

$$y_n \leq x'\left((T_n - E(T))\Delta i_t + \log S_T^n - \sum_{n=1}^N \log S_T^n\right) \quad \forall n$$

$$y_n \leq 0 \quad \forall n$$

Table 4

Mean-semivolatility allocation performance for period Jan00–Jun15 for different risk tolerance thresholds (θ). Positions are opened according to HMM-based signal $V = 5$. Measures are annualized.

θ (basis points)	Geometric mean (percent)	Volatility (percent)	Semi volatility (percent)	MSV (percent)	Skew	Kurtosis
10	1.0	2.2	1.6	64.9	-0.1	8.4
20	2.3	4.4	3.1	72.4	-0.1	8.5
40	3.4	6.8	4.9	70.0	-0.3	6.5
60	3.6	7.3	5.3	67.4	-0.4	6.1
80	3.5	7.7	5.6	62.0	-0.4	5.4
100	3.6	8.0	5.8	61.7	-0.3	5.0
120	3.6	8.3	6.0	60.8	-0.3	4.5

To construct scenarios, we assume that each currency moves as a geometric Brownian motion with no drift. That is

$$S_{h+1}^c = S_h^c \exp\{-0.5\sigma_c^2 + \sigma_c B_h^1\} \quad \forall c, \quad h : t..T$$

where σ_c is the volatility and B_h^1 is a Brownian motion with unit variance. Note that the Brownian motion is common for every currency, so currency rates are correlated. To estimate σ_c , we use historical currency data under the good regime. Table 6 in the appendix shows estimation values.

To deal with the uncertainty T, note that T is a hitting time of a discrete Markov chain, which is derived from the HMM. T corresponds to the time when we have a non-favorable state (not the last state) for V consecutive days. For the 2 states HMM, denote X_t as the number of consecutive days outside the good regime at day $s \geq t$. $X_t \in \{0..V\}$. The transition matrix structure for this Markov chain is:

$$\begin{bmatrix} p_{GG} & p_{GB} & 0 & \dots & 0 \\ p_{BG} & 0 & p_{BB} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ p_{BG} & 0 & \dots & p_{BB} \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix}$$

$p_{GG}, p_{BG}, p_{GB}, p_{BB}$ are the HMM transition probabilities, where “G” refers to good state and “B” to bad state. T is the hitting time of state V beginning at state 0. For the three HMM states, the Markov chain is a little more complicated, since we need to know if we enter a bad or transition regime once we leave the good regime. But following the same idea, we can get the transition probability for that case.

To get S_T^n , we first obtain the hitting time T and then we choose the vector S_T from the currencies path simulated. Table 8 (appendix) shows that average hitting time is around a quarter of a year.

5.2. Results

Table 4 and Fig. 6 show performance results for this mean-semivariance allocation methodology, applied to the best HMM-based strategy so far ($V = 5$). There is consistent improvement with respect to the same strategy based purely in interest rates differential ($K = 2$) seen in Table 3. That is, the MSV of all strategies are above the 56 percent obtained with previous allocation rule. This improvement is also seen when using other values for V, as Tables 9 and 10 in the appendix show.

Another benefit from this allocation methodology is that it opens more options for investors. By changing θ , each user can choose their preferred risk-return relation. Fig. 7 shows the composition of this allocation method. Main targets and funding currency are the same as with benchmark strategies (see Fig. 1). However, especially for more conservative strategies, allocation tends to be more equally weighted. Not displayed, but important to mention, is that all currencies were used at every period at every strategy shown. Hence, there is more diversification.

6. Conclusion

In this paper, we have shown how to construct a CF portfolio that outperforms standard CT strategies. To this end, we

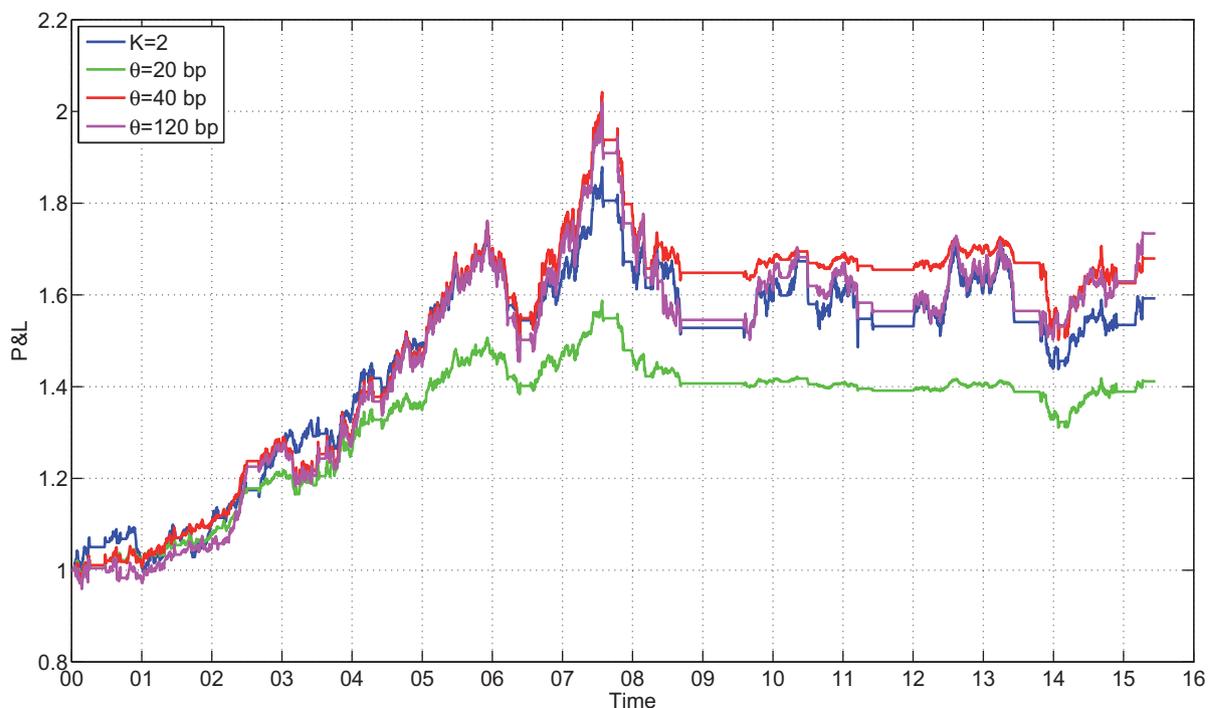


Fig. 6. P&L when $V = 5$ for different allocation settings. $K = 2$ corresponds to the simple benchmark allocation rule, based purely on interest rates. The rest is based on the mean-semivolatility allocation method for different risk tolerance θ .

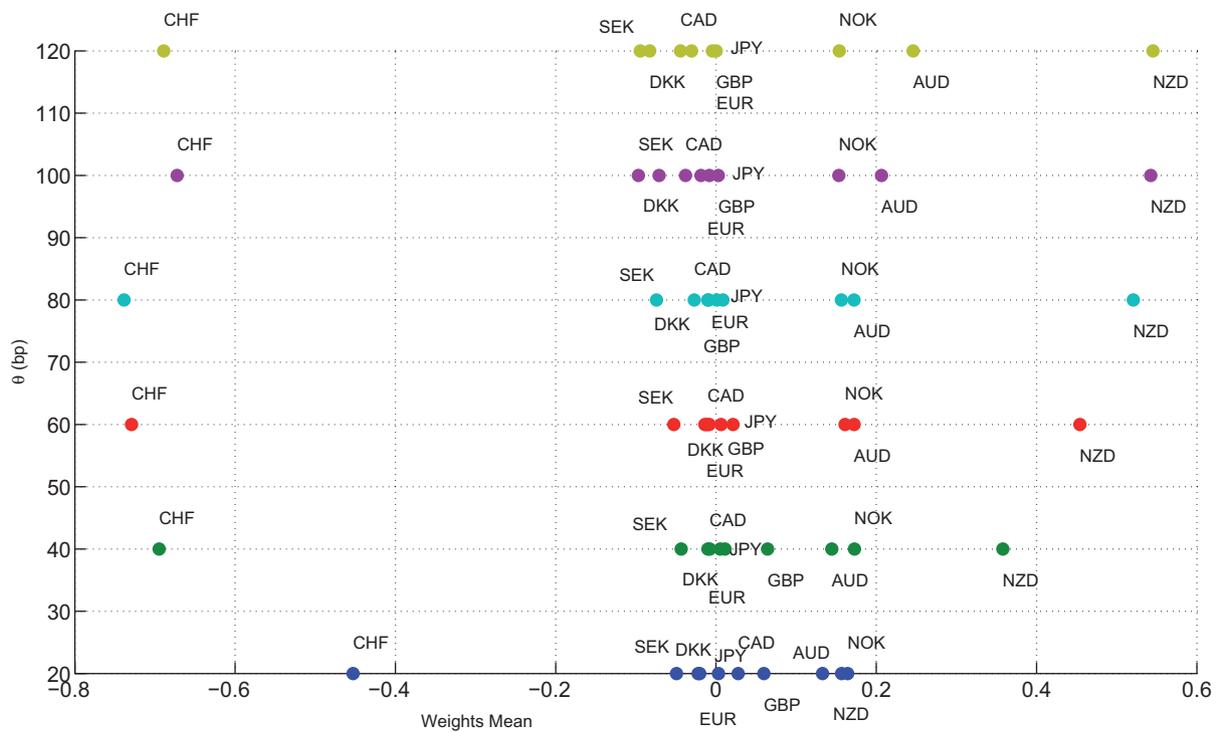


Fig. 7. Allocation composition for mean-semivariance strategies.

implemented a two-step procedure. The first step consisted in constructing a HMM model for some benchmark CT that could classify its returns in different regimes. Fitting results showed that the two-regime model was the most suitable before the 2008 crisis and a three states one on afterwards. Basically, the first regime generated low return/high volatility returns, and the latter one the opposite. With this calibration, we were able to create strategies based on HMM. This is, we open (close) positions when we remain for V consecutive days in the good (bad) regime. Results showed a clear improvement in performance in relation to standard carry trade benchmark strategies, especially for $V \in \{5, 6, 7\}$.

In the second step we improved allocation methodology when positions are opened. To do this, we use a portfolio optimization model that aims for the best compromise between geometric mean return and semi-variance. Backtest results showed that by adjusting risk tolerance (θ), we achieve better and different risk-return results.

Summing up, in terms of performance, our final strategy can increase MSV by 160 percent approx. (from 27.3 percent to 70 percent). Around two thirds of that increase comes just from applying the HMM signal, while the rest come from refining allocations with mean-semivariance approach.

Although we we have achieved good results, more could be tested in order to improve the regime identification. We have not included economic fundamentals to the methodology. It is probable there are economic variables that could be added as observable data (input) to the HMM and improve fitting results, or use this information as a complement to the HMM used in this case, for example, to determine regime changes. Another possibility is to try a different distribution of the data in each regime or use a non-parametric technique.

Regarding data used, a possible extension is to increase the number of currencies in the portfolio. It will be interesting to include currencies from emerging countries and see how much benefit and new opportunities this brings. The main problem of

including more countries comes with data availability, as HMM fitting requires a large amount of data. Additional improvement could be obtained by changing something different when the signal is off. Some strategies still had frequent flat periods that could be exploited. The most intuitive attempt is to reverse CT, i.e., to go long on low interest rates currencies and short on high interest rates currencies. Another possibility is to see what happens when different leverage is used.

Appendix

Table 5

MSV between CT(K = 2) and HMM based strategy (V = 5) within Jan00–Jun15. Both strategies are applied independently over different equal-length time windows. The first two columns shows results for 5 time windows (800 days), next two columns for 10 time windows (400) and last two columns for 15 time windows (260 days). Time windows including financial crisis is in bold. When two negative MSV are obtained, MSV is not the metric to compare. We have to multiply geometric return and semivolatility instead of dividing it. This is the number shown in parenthesis. The one with higher (less negative) value is the best strategy.

Periods	CT	V	CT	V	CT	V
1	77	166	-24	81	-88	14
2	77	154	210	265	161	239
3	-21 (-0.0028)	-40 (-0.0015)	70	164	207	288
4	75	4	85	143	152	162
5	-24	24	24	105	29	197
6			-44 (-0.0097)	-162 (-0.0065)	78	127
7			129	33	-26	2
8			23	-34	0	-28
9			-30 (-0.0013)	-39 (-0.0012)	-48 (-0.0125)	-102 (-0.0027)
10			-17	108	250	148
11					59	-60
12					9	18
13					93	121
14					-97 (-0.0047)	-173 (-0.0045)
15					-32	114

Table 6

Moments measures for daily exchange rates returns $\log(\frac{S_{t+1}}{S_t})$ on good regime state. Period Jan00–Jun15. The standard deviation is local volatility of the geometric brownian motion model for exchange rates. Results are in basis points.

Currency	Mean	Std. deviation	Skewness	Excess kurtosis
SEK	0.2	71	−0.2	1.5
NOK	0.4	69.2	−0.1	1.3
JPY	−0.7	58.0	0.0	2.3
DKK	0.8	59.1	0.0	0.9
CHF	1.5	62.9	0.1	0.9
CAD	0.6	49.6	−0.2	1.5
NZD	0.9	74.4	−0.2	0.9
GRP	0.8	50.1	0.0	0.9
EUR	0.8	59.0	0.0	0.9
AUD	0.5	67.9	−0.2	2.2

Table 7

Correlation for exchange rates returns $\log(\frac{S_{t+1}}{S_t})$ on good regime state. Period Jan00–Jun15.

Currency	NOK	JPY	DKK	CHF	CAD	NZD	GRP	EUR	AUD
SEK	0.8	0.2	0.8	0.7	0.5	0.5	0.6	0.8	0.6
NOK		0.2	0.8	0.7	0.4	0.5	0.6	0.8	0.6
JPY			0.3	0.4	0.1	0.2	0.3	0.3	0.2
DKK				0.9	0.4	0.5	0.7	1.0	0.6
CHF					0.4	0.5	0.6	0.9	0.5
CAD						0.5	0.4	0.4	0.6
NZD							0.5	0.5	0.8
GRP								0.7	0.5
EUR									0.6

Table 8

Simulation results of hitting time (days) in the last opening day.

V	Mean	Std. dev	Std. error	Quantile 5 percent	Quantile 95 percent
5	65	63	2	7	191
6	70	66	2	9	215
7	74	70	2	10	224

Table 9

Mean-semivolatility allocation performance for period Jan00–Jun15 and different risk tolerance thresholds (θ). Positions are opened according to HMM-based signal $V = 6$. Measures are annualized.

θ (basis points)	Geometric mean (percent)	Volatility (percent)	Semi volatility (percent)	MSV (percent)	Skew	Kurtosis
10	0.8	1.9	1.4	57.1	−0.3	7.0
20	1.7	3.8	2.7	61.9	−0.3	7.3
40	2.8	6.5	4.7	59.9	−0.5	6.4
60	3.2	7.2	5.3	60.1	−0.5	5.9
80	3.4	7.5	5.5	62.2	−0.4	5.3
100	3.5	7.8	5.7	61.5	−0.4	4.8
120	3.5	8.1	5.9	59.0	−0.4	4.4

Table 10

Mean-semivolatility allocation performance for period Jan00–Jun15 for different risk tolerance thresholds (θ). Positions are opened according to HMM-based signal $V = 7$. Measures are annualized.

θ (basis points)	Geometric mean (percent)	Volatility (percent)	Semi volatility (percent)	MSV (percent)	Skew	Kurtosis
10	0.8	1.7	1.2	64.2	−0.6	8.5
20	1.5	3.3	2.4	60.7	−0.6	9.6
40	2.9	5.9	4.3	66.2	−0.6	8.2
60	3.1	7.3	5.4	58.0	−0.6	7.0
80	3.0	7.6	5.6	54.6	−0.6	6.3
100	3.1	7.8	5.7	54.5	−0.6	5.7
120	3.4	8.1	5.9	57.4	−0.5	5.2

References

Bae, G. I., Kim, W. C., & Mulvey, J. M. (2014). Dynamic asset allocation for varied financial markets under regime switching framework. *European Journal of Operational Research*, 234(2), 450–458.

Ballester, E. (2005). Mean-semivariance efficient frontier: A downside risk model for portfolio selection. *Applied Mathematical Finance*, 12(1), 1–15.

Bansal, R., & Dahlquist, M. (2000). The forward premium puzzle: different tales from developed and emerging economies. *Journal of International Economics*, 51(1), 115–144.

Bhar, R., & Hamori, S. (2006). *Hidden Markov models: Applications to financial economics*: Vol. 40. Springer Science & Business Media.

Brunnermeier, M. K., Nagel, S., & Pedersen, L. H. (2008). Carry trades and currency crashes. *Technical Report*. National Bureau of Economic Research.

Burnside, C., Eichenbaum, M., & Rebelo, S. (2008). Carry trade: The gains of diversification. *Journal of the European Economic Association*, 6(2–3), 581–588.

Cheung, Y.-W., Chinn, M. D., & Pascual, A. G. (2005). Empirical exchange rate models of the nineties: Are any fit to survive? *Journal of International Money and Finance*, 24(7), 1150–1175.

Estrada, J. (2007). Mean-semivariance behavior: Downside risk and capital asset pricing. *International Review of Economics & Finance*, 16(2), 169–185.

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

Fishburn, P. C. (1977). Mean-risk analysis with risk associated with below-target returns. *The American Economic Review*, 67(2), 116–126.

Fraser, A. M. (2008). *Hidden Markov models and dynamical systems*. SIAM.

Galati, G., Heath, A., & McGuire, P. (2007). Evidence of carry trade activity. *BIS Quarterly Review*, 3, 27–41.

Guidolin, M., & Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11), 3503–3544.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.

Huang, X. (2008). Mean-semivariance models for fuzzy portfolio selection. *Journal of computational and applied mathematics*, 217(1), 1–8.

Jylha, P., & Suominen, M. (2009). Arbitrage capital and currency carry trade returns. Available at SSRN 1343439.

Kilian, L., & Taylor, M. P. (2003). Why is it so difficult to beat the random walk forecast of exchange rates? *Journal of International Economics*, 60(1), 85–107.

Mamon, R. S. (2014). *Hidden Markov models in finance: Further developments and applications*: Vol. 209. Springer.

Markowitz, H. M. (2010). Portfolio theory: As i still see it. *Annual Review of Financial Economics*, 2(1), 1–23.

McCauley, R. N., & McGuire, P. (2009). Dollar appreciation in 2008: safe haven, carry trades, dollar shortage and overhedging. *BIS Quarterly Review*, 85–91. http://www.bis.org/publ/qtrpdf/r_qt0912.pdf

Mulvey, J. M., & Zhao, Y. (2011). An investment model via regime-switching economic indicators. *Technical Report*. Princeton: Research Report, Bendheim Center for Finance and Operations Research and Financial Engineering Department. Princeton University.

Pippenger, J. (2011). The solution to the forward-bias puzzle. *Journal of International Financial Markets, Institutions and Money*, 21(2), 296–304.

Prajogo, A. U. (2011). *Analyzing patterns in the equity market: ETF investor sentiment and corporate cash holding*. Princeton University.

Reinert, K. A., Rajan, R. S., & Glass, A. J. (2009). The Princeton Encyclopedia of the World Economy.(Two volume set). *Princeton University Press*, 2, 495–497.

Ranaldo, A., & Söderlind, P. (2010). Safe haven currencies*. *Review of Finance*, 14(3), 385–407.

Vasant, J., Irgolic, L., Kruger, R., & Rajaratnam, K. (2014). A comparison of mean-variance and mean-semivariance optimisation on the JSE. *Journal of Applied Business Research (JABR)*, 30(6), 1587–1596.

Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2), 260–269.

Whitelaw, R. F. (1994). Time variations and covariations in the expectation and volatility of stock market returns. *The Journal of Finance*, 49(2), 515–541. doi:10.1111/j.1540-6261.1994.tb05150.x.