Uncovering the Trend-Following Strategy

To help currency managers.

Henrik H. Pedersen and Gerben J. de Zwart

HENRIK H. PEDERSEN is a director of the CitiFX Risk Advisory Group in London.
henrik.pedersen@citigroup.com

GERBEN J. DE ZWART is a senior researcher at Robeco Quantitative Research in Rotterdam, the Netherlands.
g.j.de.zwart@robeco.nl

T rend-following strategies generate currency forecasts by extrapolating past sequences of currency market movements into the future. Research on how these strategies extract value in currency markets includes Levich and Thomas [1993], LeBaron [1999], and Sullivan, Timmerman, and White [1999]. Success is not evident across all markets, however. Lee and Mathur [1995], for example, note that most studies showing favorable results are based on U.S. dollar-denominated currencies. So why is it that trend strategies seem to work well for some currencies, but not for others?

A key criterion for economic success in trend-following seems to be selection of the right currency pairs to trade. Yet despite a wealth of analysis, it is hard to explain the exact statistical characteristics that make trend-following successful within a universal framework. The primary reason is the limited number of liquid and frequently traded currencies available to the analyst, which means that asset selection is typically based on past performance that may not be repeated in the future.

To create a universal framework that can improve this selection process, we build a model that can explain the profitability drivers of trend-following strategies for any given currency pair. We first use a large number of random walk simulations of currency moves. By not relying on historical data, we are able to overcome the information barrier caused by the limited number of past currency paths.

This methodology also allows us to examine the effects of certain specified characteristics for drift (mean), kurtosis, skew, and volatility on trend model performance. Having created the model from simulations, we use empirical data to verify its predictions. This enables us to highlight the statistical attributes that make trend-following successful across a number of frequently traded currency pairs, and to analyze the short-term implications for trading performance.

This research is interesting for three reasons. First, the model explains why trend-following strategies work for certain currencies and not for others within a universal framework. Second, the model identifies the statistical characteristics that influence the performance of trend-following strategies. Finally, the results are applicable to evaluation of trend-following strategies in other asset classes.

TREND MODEL PERFORMANCE

The trend model used in this analysis is a moving-average model using three moving averages (31, 61, and 117 days). This model was proposed by Acar and Lequeux [1998] as a representative benchmark for the various durations followed by active currency traders. Position-taking is simply a function of spot price versus each of the three moving averages (TMA). If the spot is above (or below) all three averages, the active position is 3/3 = 1 (or −1). If the spot is above (or below) two of the moving averages, the position is (2 − 1)/3 = 1/3 (or −1/3). We call this model the TMA model (or strategy).

To create trading signals for the TMA model, we simulate 22,000 individual exchange rate series, each consisting of 5,118 daily rates. On the 118th day, the last of the three moving averages will have a trading signal and thereby activate the model. This allows us to calculate the model's profitability over 5,000 trading days (almost 20 years). The annualized gross profit of the TMA strategy will be referred to as the TMA result.

We do not consider trading costs or interest rate differentials as they are of limited relevance to our purpose.

SIMULATION PARAMETERS

As the standard normal distribution does not capture the skewness and fat-tailed return characteristics experienced in currency markets, we must be able to incorporate the first four moments into the simulated distributions. These are:

1. *Drift* or mean measures the average return on a currency. It is a common assumption in the market that the long-term expected return for currency markets is zero. The actual observed drift will be a function of the time interval that we have chosen to observe. Baz et al. [2001] discuss the zero currency risk premium in more detail.
2. *Volatility* is defined as the normalized standard deviation of currency returns. Visually it describes the width of a distribution so that high volatility provides a wider range of simulated outcomes than low volatility.
3. *Skew* measures the degree of asymmetry of a distribution around its mean, or more simply the observed concentration of positive or negative outcomes. It is therefore natural to assume that skewness could have a positive effect on trend-following strategies.
4. *Kurtosis* characterizes the relative peakedness (and fat-tailedness) over a normal distribution (excess kurtosis). The presence of fat tails and the ability of trend-following models to capture these have also been frequently cited as an explanation of the effectiveness of trend models in currency markets.[1]

Johnson [1949] proposes a set of four normalizing translations and a translation function in order to fit any feasible set of sample values for all four moments. An algorithm to fit the Johnson curves by moments from a normal Monte Carlo simulation was later developed by Hill, Hill, and Holder [1976]. We use this algorithm here in order to create the real-life distribution shapes actually experienced in currency markets.

Exhibit 1 sets out the distribution parameters we use in the simulations. Only one parameter is changed at any one time in order to isolate the factor impact on trend performance. 1,000 exchange rate series are produced for each of the 22 different settings. In each case, we verify that the four moments produced by the algorithm are always equal to the desired settings.

The advantage of this method is that we are not bound by any historical bias of combinations that might have been present in past exchange rates. A drawback is that it is based on the efficient markets assumption. This means that the results could be different from those of a non-random data series. For practical reasons, we have also excluded the effect of interest rate differentials on exchange rate movements from our simulations. First, this is not one of the factors we intend to evaluate, and, second, it would make the

Overview of Four Moment Settings Used in Simulations

	Basic Setting	Variations of Basic Setting						Observed Moments*
Mean	0.00%	0.01%	0.02%	0.03%	0.04%	0.05%	0.10%	0-0.01%
Volatility	0.60%	0.20%	0.40%	0.60%	0.80%	1.00%	-	0.25-0.75%
Skew	0.00	0.25	0.50	0.75	1.00	1.25	-	0.0-0.7
Kurtosis	0	1	2	3	5	10	15	1-14

* Mean and skew are absolute values only. All observed moments are listed in Exhibit 4.

computations significantly more complex.

The settings are chosen so as to include most values that can be observed in real markets. We also include some extreme levels of daily drift despite the long-term zero drift assumption we have noted. This is because in practice, it will always be possible to find a sample period when the average drift is not zero. A frequently discussed example is the 20% fall in the USD/JPY exchange rate in the autumn of 1998.[2]

Note that we consider only the absolute values of drift and skew in this analysis. This is valid because the model allows long as well as short positions; in practice, it is irrelevant to a trend model whether the distribution mean, or skew, is positive or negative.

MODEL TRADING FREQUENCY

The standard four moments describe only the overall distribution of the returns, not the sequence in which they are realized. The sequence, or path, is basically the visual representation that we can observe in a typical line graph. As we can generate a given distribution using an indefinite number of different paths, a path descriptive measure is required in order to describe exactly how a certain outcome is produced.

As the trading frequency generated by the trend model is directly related to the compilation of each individual currency path, we introduce a measure of trading frequency called Tfreq. Tfreq is expressed as a percentage number and is calculated as follows:

$$Tfreq = \frac{\sum_{t=1}^{N} Abs(P_t - P_{t-1})}{N} \qquad (1)$$

where:

P_t is the position size at time t with $-1 \le P \le 1$, and N is the total number of potential trading events in the period we observe.

The purpose of this measure is to provide a simple categorization of the simulated path types leading to certain performance levels. Because we compute nominal position changes, Equation (1) will produce more or less identical results, whatever the number of moving averages used.

This is an advantage, as it makes the path definition universally applicable. In effect, it is an advanced measure of autocorrelation because it is filtered for reverse moves that will not cause a position change.[3] A simulation sample with the basic settings illustrates the importance of the trade frequency.

Exhibit 2 shows the outcome of 1,000 different random walk simulations. We can see the strong correlation between the Tfreq and the TMA result. This result is not surprising but rather intuitive.

On average, the more a trend model trades (in and out of positions in a non-trending market), the less likely it is to be profitable. The long straight fitted line (full simulation) shows the best fit using linear regression, and we see that the R^2 is high at 0.84. This means that 84% of the TMA result can be explained by variation in trade frequency. Also note that the trend strategy produces a positive result for all values below 14% and a negative result for all values above 16%.
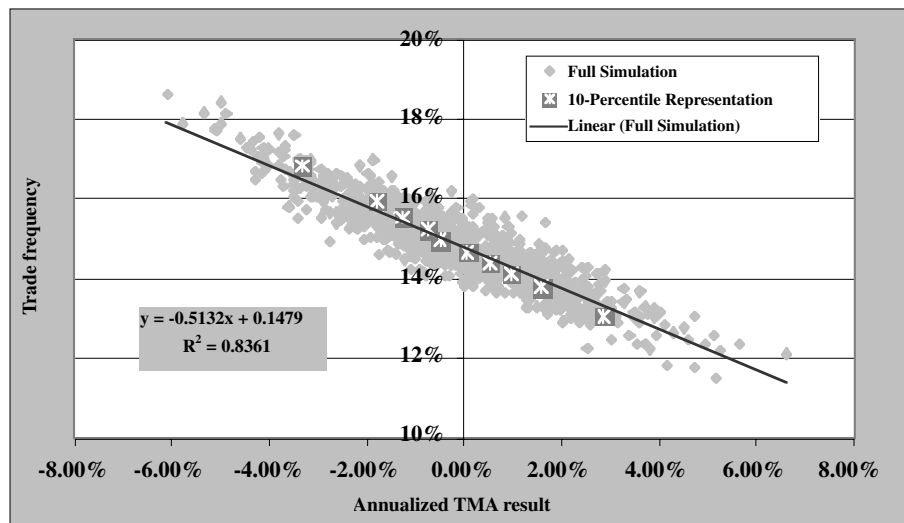
To obtain a manageable number of path categories for the analysis, the results of each simulation setting are divided into 10 percentiles (ranked by trade frequency), and the average results for each set of 100 observations are calculated (illustrated in Exhibit 2 as the 10-percentile representation). We see that this is also a linear representation of the data set, but note that it is slightly steeper than the fitted line, especially at the extreme values of Tfreq.

MODEL TO ASSESS TREND PROFITABILITY

Armed with a fifth potential determinant driving trend model performance, we can now proceed toward the ultimate objective. Essentially this means we will seek to predict trend model profitability based on expected (or empirical) values of the mean, volatility, skew, kurtosis, and a path categorization.

**Relationship Between TMA Result and Tfreq
from 1,000 Random Walk Simulations on Basic Settings** (5,000 days)

tion (2) indicates that 99.5% of the simulated TMA result can be explained by the variation in our chosen variables. The standard error of estimation is 0.3%. As we have used the 10th percentile representation to map the relationship, however, these values indicate far too good a fit. This will be corrected when we use the model for prediction and identify a more realistic value for the standard error of estimation.

Interpretation of Results

The mechanics of the model described in Equation (2) can be interpreted as follows:

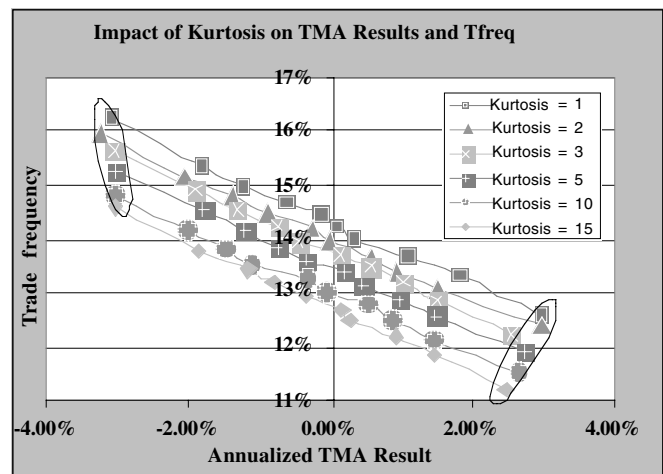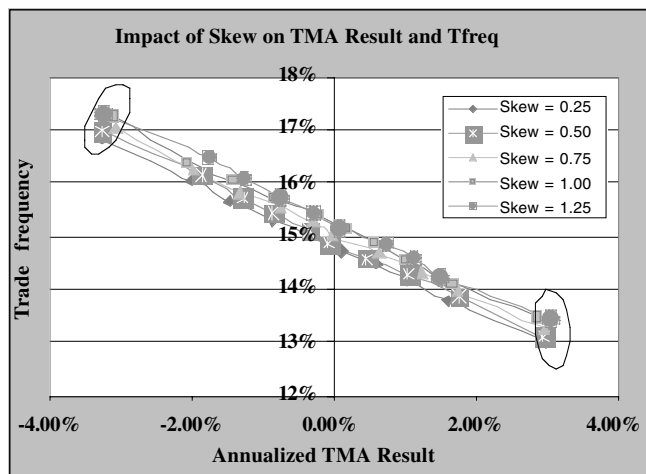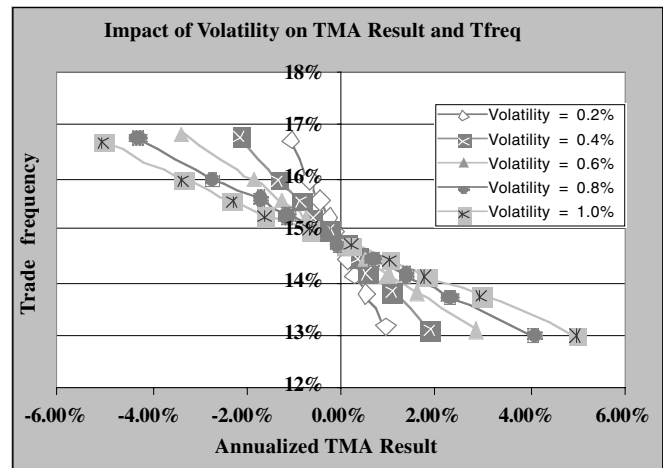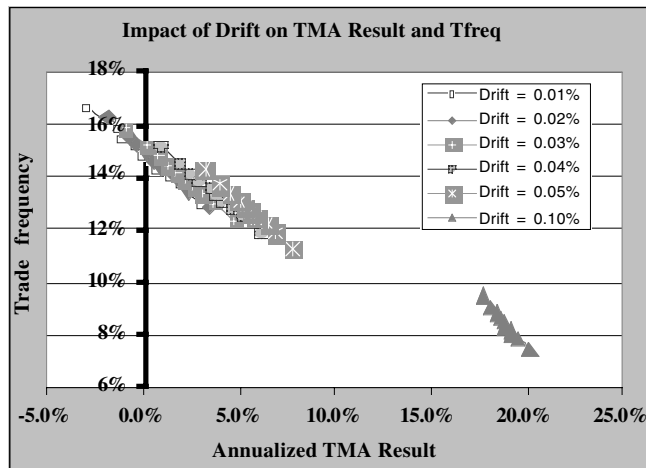Market volatility (38.88 Stdev) determines the profit (or loss) potential of the trend-following strategy. This relationship is direct, so if market volatility doubles, so does the expected TMA result (assuming the other factors remain constant). Accordingly, it is no longer surprising that trend-following models tend to show the best results across the major currency blocks where market volatility is consistently higher than that experienced with anchored or regional currency pairs.

On a market level, it also indicates that trend-following will tend to be less successful in a low-volatility environment. Note that this recognizes the ability of the overall market environment to produce profits for a trend model. This is different from the short-term volatility considerations discussed by, for example, Acar and Lequeux [2001]. They show that adjusting currency positions to reflect one-month market volatility can be profitable.

In our case, this effect is already captured by our path definition. This is because the Tfreq measure identifies market volatility that is non-directional. In practice, it means that we compare paths with similar short-term volatility characteristics, and therefore equal-risk adjusted returns.

The second part of Equation (2) describes the pre-volatility-adjusted impact of each individual element on trend model performance. Perhaps the best way to look at this part is that it incorporates the factors that influence the trading frequency and thereby the actual sequence of return. Using this interpretation we see that, as expected, a high Tfreq will have a negative impact on trend model performance. We can show that the critical value for path prof-

The model is created using linear regression on the 10th percentile representations collected from each of the 22 different simulation settings. This means that the total number of observations used to create the prediction model represents a total of 220 samples from 22,000 individual simulations.

The regression analysis is not straightforward, as we introduce an additional interaction variable to describe the observed impact of the simulated volatility (Stdev) on the remaining input variables, whatever their setting. We also include the quadratic expression for drift. Therefore, the potential list of explanatory variables to determine the TMA result is: Stdev, Stdev × Tfreq, Stdev × Skew, Stdev × Kurtosis, Stdev × Drift, Stdev × Drift^2, and Tfreq.

We intend to include only the coefficients that are statistically significant. This means a few experiments are required in order to derive the best coefficients. In this process, Tfreq is discarded as a single-input variable, as it appears to be insignificant.

The model based on the final regression results is:[4]

$$\text{TMA Result} = 38.88\text{Stdev}(1 - 6.77\text{Tfreq} + 0.0392\text{Skew} - 0.0101\text{Kurtosis} + \text{Effect of Drift}) \quad (2)$$

where: effect of drift is the impact of the daily mean of the distribution computed as Drift(65.65 + 324,600Drift). This shows that the TMA result increases significantly with periodic drift.

The adjusted R^2 for the prediction model in Equa-

EXHIBIT 3
Impact of Factors

Impact of Drift on TMA Result and Tfreq

Legend:
- Drift = 0.01%
- Drift = 0.02%
- Drift = 0.03%
- Drift = 0.04%
- Drift = 0.05%
- Drift = 0.10%

(Y-axis: Trade frequency, 6% to 18%; X-axis: Annualized TMA Result, -5.0% to 25.0%)

Impact of Volatility on TMA Result and Tfreq

Legend:
- Volatility = 0.2%
- Volatility = 0.4%
- Volatility = 0.6%
- Volatility = 0.8%
- Volatility = 1.0%

(Y-axis: Trade frequency, 12% to 18%; X-axis: Annualized TMA Result, -6.00% to 6.00%)

Impact of Skew on TMA Result and Tfreq

Legend:
- Skew = 0.25
- Skew = 0.50
- Skew = 0.75
- Skew = 1.00
- Skew = 1.25

(Y-axis: Trade frequency, 12% to 18%; X-axis: Annualized TMA Result, -4.00% to 4.00%)

Impact of Kurtosis on TMA Results and Tfreq

Legend:
- Kurtosis = 1
- Kurtosis = 2
- Kurtosis = 3
- Kurtosis = 5
- Kurtosis = 10
- Kurtosis = 15

(Y-axis: Trade frequency, 11% to 17%; X-axis: Annualized TMA Result, -4.00% to 4.00%)

itability is 14.8%, no matter the level of market volatility.[5]

Skew will enhance performance, while the opposite is true for kurtosis. Drift will increase the value of the equation and thereby contribute positively to the TMA model result.

Interpretations of all the factors are graphed in Exhibit 3, which depicts the results of the individual simulations by moment. Results with high trend model profitability and low Tfreq are found at the lower right-hand of each graph.

It is of course important to stress that the structure of Equation (2) has been derived from simulations. This means that although the interpretation it offers is in line with our individual experiments and expectations, the complexity of the problem prevents us from verifying this mathematically.

Importance of Coefficients

We can evaluate the relative importance of each of the six input variables in predicting the TMA result by calculating the partial R^2 (coefficient of partial determination) for each x-variable.[6] The partial R^2 measures the mutual relationship between the TMA result and a single variable, when the remaining variables are kept constant. This allows us to identify the proportion of unexplained variation in the TMA result that can be explained by adding a particular variable to the model.

The results show that the currency path is the most important factor in determining performance (91%). The impact from kurtosis (68%) and drift (56%) is also significant. Skewness is less significant, but still explains 26% of variance on its own. Volatility has no importance at all (0.4%). This might initially come as a surprise, but as illustrated in Equation (2), it is a multiplication variable and

so does not in itself generate trend model profitability (or loss where the path characteristic is unfavorable).

Prediction Accuracy

The prediction accuracy indicated by the regression experiment is significantly exaggerated because the relationship was mapped using (almost) linear representations. To get a more accurate measure, we have to include all the simulations in our considerations. We also want to highlight the different factors that influence the estimation error so that we can use the prediction model universally.

Toward this end, we collect some additional simulation results for a number of subperiods (from 260 to 5,000 days) and volatilities (0.2%, 0.6%, and 1.0%) from our simulation samples using the basic settings described in Exhibit 1.[7] For each subperiod we then calculate the standard error of estimation using Equation (2).

The result of this exercise can be summarized as:

$$\text{Standard Error of Estimate} \cong \frac{113.2 \times Stdev}{N^{0.52}} \quad (3)$$

Equation (3) provides an approximation of the standard error for the predicted TMA result for a given number of potential trading days (N) and market volatility ($Stdev$). It is correct with 99% confidence. It follows that although our model predictions can be applied for all time horizons, they are more accurate for longer time periods. Furthermore, as prediction accuracy declines with increasing volatility, this would indicate that, the higher the market volatility, the wider the range of potential outcomes for any given set of input variables.

EMPIRICAL VERIFICATION OF THE MODEL

To evaluate the model against empirical results from the currency market, we compare the model predictions with actual past performance (from 1994 to 2003) for a number of frequently traded currencies in the marketplace. The predicted results are calculated using Equation (2). The basis of our calculations is the actual daily series from which we have calculated the actual or observed TMA result.

The currency pairs, their empirical moments, and the past performance characteristics used to compute the predictions are given in Exhibit 4. We also show the individual and average model prediction for the TMA result, as well as the contribution from each individual factor to the result.[8]

We see that a low Tfreq (<14.8%) and the presence

of skewness on average produce positive performance, while kurtosis has a negative impact. This is consistent with the factor importance outlined previously. Also note that the impact of drift on performance is close to zero, as there was no significant drift observed for our chosen time interval.

When we examine the performance drivers for each individual currency pair, we see (for example) that the past currency path of EUR/USD has provided the overwhelming contribution to performance, and that kurtosis has a significant negative effect on EUR/NOK.

The last four columns show the predicted TMA result as well as a 95% confidence interval for the prediction. We see that in general our point predictions are close to the observed TMA result. In each case, we show the appropriate standard error for the prediction based on Equation (3) for a sample period of 2,600 days.

If we take a closer look at the 95% confidence intervals, two observations are notable. First, we see that all the observed TMA results from the 14 currency samples fall comfortably within the 95% confidence interval of the model prediction. This is good news, because it provides us statistical verification that our predictions are in line with actual market behavior.

Second, we can see that for EUR/USD, EUR/GBP, USD/JPY, USD/CHF, USD/CAD, and CHF/JPY the lower prediction limit is close to or higher than zero. This means that according to Equation (2) there was a better-than-95% probability of making money on trend-following in this period for these currency pairs. This demonstrates the second use of the model: By mapping the statistical characteristics that influence the performance of trend models, we are also able to identify the currencies most likely to be economically successful.

We conclude that the prediction model can explain trend model performance with statistical significance. This means that we are not only able to explain why some currencies will do better than others when subjected to a trend-following strategy, but we can also identify the statistical distribution characteristics that influence these performance observations.

TMA RESULTS AND CURRENCY MARKET EFFICIENCY

Our model is based on the random walk assumption, but also predicts positive returns for the trend-follower for a large number of currency pairs. This initially seems to contradict the efficient markets hypothesis. This

E X H I B I T 4
Predictions of TMA Results for Individual Currency Pairs

Currency Pair	Four Empirical Moments and Observed Tfreq.* (calculated from daily returns)					Observed TMA (ex. carry)	Individual Model Factor Contribution (relative to average market volatility)					Predicted		95% Confidence Interval Limits	
	Mean	Stdev	Skew	Kurtosis	Tfreq		Tfreq	Stdev	Skew	Kurtosis	Drift	TMA Model	Approx. Std Error	Lower	Upper
EUR/USD	0.00%	0.61%	0.09	1.60	12.0%	3.80%	4.00%	0.37%	0.07%	-0.35%	0.07%	4.17%	1.16%	1.89%	6.45%
EUR/GBP	0.00%	0.50%	0.19	1.08	13.3%	0.46%	2.21%	-0.24%	0.16%	-0.23%	0.05%	1.94%	0.94%	0.10%	3.79%
EUR/JPY	0.00%	0.74%	0.32	2.80	13.9%	2.58%	1.27%	0.32%	0.27%	-0.61%	0.04%	1.30%	1.40%	-1.45%	4.05%
EUR/CHF	0.00%	0.25%	0.40	4.84	14.9%	-0.63%	-0.11%	0.43%	0.34%	-1.05%	0.04%	-0.35%	0.47%	-1.28%	0.58%
EUR/NOK	0.00%	0.36%	0.53	12.77	13.8%	0.37%	1.46%	0.31%	0.45%	-2.78%	0.00%	-0.56%	0.68%	-1.91%	0.78%
EUR/SEK	0.00%	0.45%	0.24	2.36	14.1%	0.31%	1.01%	-0.13%	0.21%	-0.51%	0.02%	0.59%	0.86%	-1.10%	2.28%
USD/JPY	0.00%	0.72%	0.64	5.02	13.2%	3.74%	2.36%	0.55%	0.54%	-1.09%	0.02%	2.38%	1.37%	-0.32%	5.08%
USD/CHF	0.01%	0.67%	0.18	2.34	13.1%	1.88%	2.50%	0.47%	0.15%	-0.51%	0.13%	2.74%	1.27%	0.24%	5.24%
USD/CAD	0.00%	0.35%	0.09	2.60	13.3%	0.26%	2.13%	-0.60%	0.08%	-0.56%	0.01%	1.05%	0.67%	-0.27%	2.37%
USD/AUD	0.00%	0.62%	0.08	3.80	14.3%	0.52%	0.75%	0.01%	0.07%	-0.83%	0.06%	0.06%	1.17%	-2.25%	2.37%
GBP/USD	0.01%	0.46%	0.09	1.68	15.2%	-1.69%	-0.56%	0.12%	0.08%	-0.37%	0.14%	-0.59%	0.87%	-2.32%	1.13%
GBP/JPY	0.01%	0.75%	0.31	2.86	13.7%	2.06%	1.57%	0.47%	0.26%	-0.62%	0.10%	1.78%	1.43%	-1.03%	4.59%
GBP/CHF	0.00%	0.55%	0.26	1.86	15.4%	-1.41%	-0.90%	0.01%	0.22%	-0.40%	0.01%	-1.07%	1.04%	-3.12%	0.99%
CHF/JPY	0.01%	0.75%	0.22	2.29	12.6%	3.03%	3.22%	1.04%	0.18%	-0.50%	0.10%	4.04%	1.42%	1.24%	6.83%
AVERAGE	0.00%	0.56%	0.26	3.42	13.8%	1.09%	1.49%	0.00%	0.22%	-0.74%	0.01%	0.97%	1.05%	-1.10%	3.05%

Mean and skew are absolute values only.

is because the hypothesis implies that the expected return from betting on market moves (using signals generated from price histories) should be zero.

This is not a contradiction for several reasons. First, the average return—excluding transaction costs—in all our simulations is 0% and in line with expectations. Second, we saw that the random walk simulation was able to show significant results, given certain path and distribution characteristics.

What our model therefore illustrates is that the performance of trend models can be explained by particular market characteristics, or attributes. This observation is in itself not a violation of the efficient markets hypothesis. Rather, it is the persistence of such characteristics in actual market behavior that is.

Given the empirical evidence in Exhibit 4, there seems to be a strong case that currency markets are not efficient. Our analysis suggests this is primarily due to the path characteristics. This is in line with Banerjee [1992], who concludes that investors act irrationally and do not trade independently of one another. Central bank intervention may also work against the profit motive assumption of foreign exchange market participants (see Szakmary and Mathur [1997] and LeBaron [1999]).
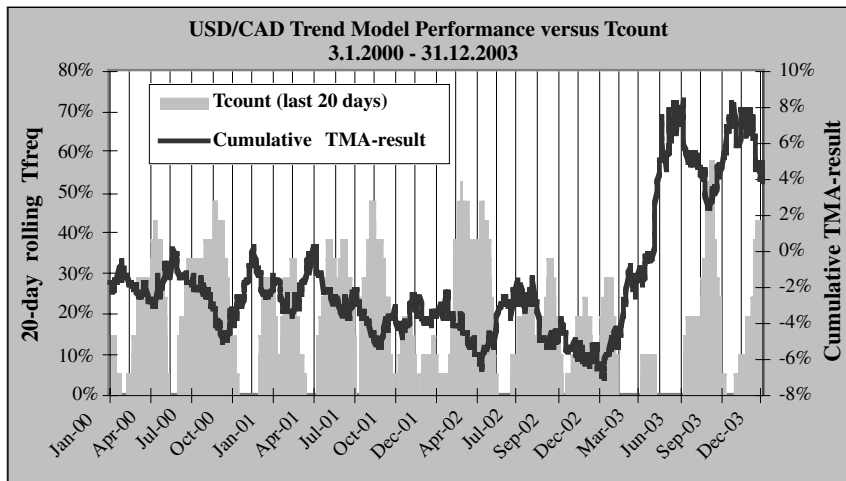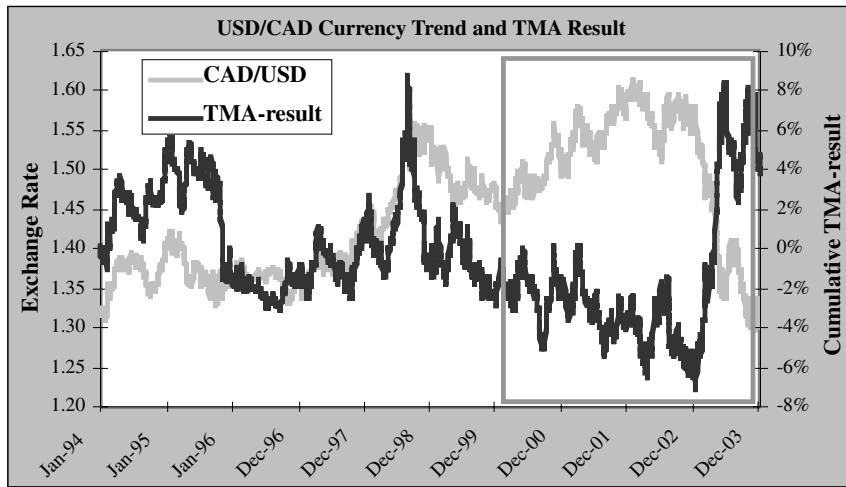
A final real-life consideration that we ignore in the analysis is interest rate differentials (carry). In practice, they will have some impact on model performance because of the need to hold positions for a certain amount of time. When we include the cost (or benefit) of carry in the calculation, we find there seems to be an additional positive impact on TMA model performance to the tune of 0.2% to 0.6% per year.

SHORT-TERM IMPLICATIONS FOR TRADING STRATEGIES

The influence of Tfreq on TMA performance can also be crucial for identifying shorter-term profit opportunities. We choose as a case study USD/CAD because it is generally considered a non-trending currency, according to empirical performance observations.

In effect, this has largely been true, as we can see in Exhibit 5. Except for brief periods of sustained returns in 1994, 1997, and 2003, the Canadian dollar has generally been a difficult currency to trade using a trend model. This is consistent with the statistical properties of USD/CAD identified in Exhibit 4. USD/CAD generally has low volatility (0.3%) and relatively high kurtosis, reducing any profit potential.

EXHIBIT 5
Case Study

USD/CAD Currency Trend and TMA Result

USD/CAD Trend Model Performance versus Tcount
3.1.2000 - 31.12.2003

This evidence shows that USD/CAD does in fact trend, but its statistical characteristics and its past path have significantly reduced the odds of making these trends profitable for the use of model trading. The evidence also shows that the trend/no-trend distinction is in effect much more differentiated than a first look would indicate. Conceptually this means that the key skill required for a trend model user is the ability not only to select the currencies that have a high probability of success, but also to distinguish the times a currency move is profitable and when it is not. We hope the model framework that we have presented here will help currency managers do both.

SUMMARY

We have simulated a total of 22,000 potential currency paths using the random walk assumption in accordance with efficient markets theory. Each path is fitted to the typical return characteristics of the currency market by the use of Johnson distributions. This enables the inclusion of fat tails and skewness in the simulations, and allows us to analyze the effect on trend model performance by independently varying each of the four moments. A fifth measure is introduced to characterize the smoothness of the simulated currency paths and the impact of directional volatility on returns.

The simulations underlie a model to assess trend model profitability using regression analysis. The model predicts trend-following profitability according to the values of mean, volatility, skewness, kurtosis, and path.

On average, kurtosis has a negative impact on trend model results, while the specific path taken by most currency pairs has a positive impact. Equally significantly, we also saw that high market volatility increases the money-making potential of a trend-following strategy, while low market volatility makes it harder to exploit any trends present.

The model is validated using empirical data for the 14 most frequently traded currency pairs. We show that most of the results observed fall within one standard error of the point estimate. Therefore we conclude that the

A key observation, however, is that there seems to be no strong link between currency trends that are visible to the eye and trend model profitability. A good example is the 1999-2001 period, when USD/CAD experienced higher highs and higher lows (the definition of an up-trend), while the TMA model generally produced negative returns.

If we instead focus on the impact of trading frequency on TMA profitability, the picture becomes much clearer. For ease of interpretation, Exhibit 5 shows only the relationship between the profitability of the TMA model and the rolling 20-day Tfreq for USD/CAD for the past three-year period. It is now clear that almost all periods of sustained profitability for the TMA model occur when Tfreq is at zero. We also see that the longer the time without trading, the better the observed model performance.

model offers a universal framework to explain why only some currencies seem to be profitable.

A case study of the U.S. and Canadian dollar shows a distinction between currency trends and profitable trends. We saw that sustained past profitability of the trend model was achieved only when the model did not trade. This has potential implications for the trend model user, as this knowledge can be applied in an attempt to avoid unprofitable model activity.

ENDNOTES

The authors thank for helpful comments and suggestions: Emmanuel Acar, Eu-Jin Ang, David Blitz, Harold de Boer, Patrick Houweling, and Antti Ilmanen. The views expressed in this article are the authors' and not necessarily those of their employers.

[1]Most moving-average models captured the USD/JPY fat-tail event in October 1998.

[2]We use the three-letter ISO codes throughout to identify currency pairs. For example, EUR/USD refers to the euro versus the U.S. dollar. The other ISO codes are: GBP = British pound, JPY = Japanese yen, CHF = Swiss franc, NOK = Norwegian krone, SEK = Swedish krona, CAD = Canadian dollar, and AUD = Australian dollar.

[3]Several authors report that the autocorrelation is low in the currency markets. See, for example, Qi and Wu [2001] and Okunev and White [2003].

[4]The original regression result has been transformed from: => TMAresult = −0.0028 + 38.89Stdev −263.13StdevTfreq + 1.5250StdevSkew − 0.3912StdevKurtosis + StdevDrift(2552.83 + 12622941Drift). The constant is set to zero as per our discussion of the efficient markets hypothesis. By definition, there can be no bias (positive or negative), and it is solely a consequence of our 10th-percentile representation. We confirm this by applying the average mean of the full 22,000 simulations, which is zero. We also isolate Stdev and reduce the coefficient values in the equation.

[5]From (2) we can infer that if the TMA result, skew, kurtosis, and drift are equal to zero, then: $0 = 38.88Stdev(1 − 6.77Tfreq) => Tfreq = 3.88 Stdev/263.13 Stdev = 14.8\%$.

[6]For an explanatory variable x_1 and a variable x_2 held constant, the partial correlation coefficient can be computed as follows (see Johnston and DiNardo [1997, p. 77]):

$$r_{y,x_1|x_2} = \frac{r_{y,x_1} - r_{y,x_2}r_{x_1x_2}}{\sqrt{(1 - r_{y,x_2}^2)(1 - r_{x_1,x_2}^2)}}$$

[7]The value of the remaining setting actually does not matter to the calculated standard error of estimation, but it creates a common constraint on our calculations.

[8]To distinguish the contributions, we first measure the impact of Tfreq, skewness, kurtosis, and drift, keeping volatility constant at the sample mean (here 0.56%). This allows us to measure the impact of volatility relative to this average value.

REFERENCES

Acar, E., and P. Lequeux. "A Dynamic Index for Managed Currencies Funds Using CME Currency Contracts." *European Journal of Finance*, 4 (1998), pp. 311-330.

——. "Volatility-Adjusted Exposures." *Currency Issues for Asset Managers—Citigroup*, October 2001.

Banerjee, A.V. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3) (1992), pp. 797-817.

Baz, J., F. Breedon, V. Naik, and J. Peress. "Optimal Portfolios of Foreign Currencies." *The Journal of Portfolio Management*, Fall 2001, pp. 102-111.

Hill, I.D., R. Hill, and L. Holder. "Algorithm AS99: Fitting Johnson Curves by Moments." *Applied Statistics*, 25(2) (1976), pp. 180-189.

Johnson, N.L. "Systems of Frequency Curves Generated by Methods of Translation." *Biometrika*, 36 (1949), pp. 149-176.

Johnston, J., and J. DiNardo. *Econometric Methods*. Singapore: McGraw-Hill, 1997.

LeBaron, B. "Technical Trading Rule Profitability and Foreign Exchange Intervention." *Journal of International Economics*, 49 (1999), pp. 125-143.

Lee, C.I., and I. Mathur. "Trading Rule Profits in European Currency Spot Cross Rates." *Journal of Banking and Finance*, 20 (1995), pp. 949-962.

Levich, R.M., and L.R. Thomas. "The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach." *Journal of International Money and Finance*, Rutgers University, 12 (1993), pp. 451-474.

Okunev, J., and D. White. "Do Momentum Based Strategies Still Work in Foreign Currency Markets?" *Journal of Financial and Quantitative Analysis*, 38(2) (2003), pp. 425-448.

Qi, M., and Y. Wu. "Technical Trading Rule Profitability, Data Snooping, and Reality Check: Evidence from the Foreign Exchange Market." Working paper, Rutgers University, November 2001.

Sullivan, R., A. Timmerman, and H. White. "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *Journal of Finance*, 54 (1999), pp. 1647-1691.

Szakmary, A.C., and I. Mathur. "Central Bank Intervention and Trading Rule Profits in Foreign Exchange Markets." *Journal of International Money and Finance,* 16 (1997), pp. 513-535.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.