

Development and calibration of a currency trading strategy using global optimization

**Mustafa Onur Çağlayan & János
D. Pintér**

Journal of Global Optimization

An International Journal Dealing with
Theoretical and Computational Aspects
of Seeking Global Optima and Their
Applications in Science, Management
and Engineering

ISSN 0925-5001

Volume 56

Number 2

J Glob Optim (2013) 56:353-371

DOI 10.1007/s10898-012-9879-2

Volume 56, No. 2 – June 2013

ISSN 0925-5001

Journal of Global Optimization

Special Issue on JOGO MEC EurOPT 2010-Izmir
Guest Editors: Refail Kasimbeyli, Musa Mammadov and
Cemali Dinçer



 Springer

 Springer

Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media, LLC.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Development and calibration of a currency trading strategy using global optimization

Mustafa Onur Çağlayan · János D. Pintér

Received: 13 February 2012 / Accepted: 14 February 2012 / Published online: 28 February 2012
© Springer Science+Business Media, LLC. 2012

Abstract We have developed a new financial indicator—called the Interest Rate Differentials Adjusted for Volatility (IRDAV) measure—to assist investors in currency markets. On a monthly basis, we rank currency pairs according to this measure and then select a basket of pairs with the highest IRDAV values. Under positive market conditions, an IRDAV based investment strategy (buying a currency with high interest rate and simultaneously selling a currency with low interest rate, after adjusting for volatility of the currency pairs in question) can generate significant returns. However, when the markets turn for the worse and crisis situations evolve, investors exit such money-making strategies suddenly, and—as a result—significant losses can occur. In an effort to minimize these potential losses, we also propose an aggregated Risk Metric that estimates the total risk by looking at various financial indicators across different markets. These risk indicators are used to get timely signals of evolving crises and to flip the strategy from *long* to *short* in a timely fashion, to prevent losses and make further gains even during crisis periods. Since our proprietary model is implemented in Excel as a highly nonlinear “black box” computational procedure, we use suitable global optimization methodology and software—the Lipschitz Global Optimizer solver suite linked to Excel—to maximize the performance of the currency basket, based on our selection of key decision variables. After the introduction of the new currency trading model and its implementation, we present numerical results based on actual market data. Our results clearly show the advantages of using global optimization based parameter settings, compared to the typically used “expert estimates” of the key model parameters.

M. O. Çağlayan (✉)
School of Economics and Administrative Sciences, Özyeğin University, Kusbakisi Caddesi, No. 2,
34662 Altunizade, Istanbul, Turkey
e-mail: mustafa.caglayan@ozyegin.edu.tr

J. D. Pintér
School of Engineering, Özyeğin University, Kusbakisi Caddesi, No. 2, 34662 Altunizade, Istanbul, Turkey
e-mail: janos.pinter@ozyegin.edu.tr

J. D. Pintér
Pintér Consulting Services, Inc., Halifax, Canada

Keywords Currency trading model · IRDAV financial indicator · Aggregated risk metric · Excel model implementation · Lipschitz global optimizer (LGO) solver engine · Global optimization by excel-LGO · Numerical results

1 Introduction

In currency markets, the differences in interest rates among countries have always played an important role for investors (financial organizations or individuals) when making buy/sell investment decisions. The pair-wise strategy of buying the currency of a country with the higher interest rate and simultaneously selling the currency of a country with the lower interest rate—known in financial terminology as the *long carry* [trading] strategy—is regarded as one of the most popular trading strategies in currency markets. In addition to interest rate differentials, another key financial factor for investors, however, is the volatility of the currency pair in question. Therefore, a trading strategy that incorporates both the interest rate differentials and the volatility would be a much better solution for an investor who is interested in currency markets.

In this paper we introduce a new financial indicator—called the Interest Rate Differentials Adjusted for Volatility (IRDAV)—which captures both the interest rate differentials and the volatility of the currency pairs. We define and calculate IRDAV as the annual interest rate differential between the high-yielding and low-yielding currencies, divided by the 6-month implied volatility of the exchange rate returns. We rank currency pairs according to this measure on a monthly basis, and then select a basket of pairs with the highest IRDAV values.

Under positive market conditions, an IRDAV based investment strategy—buying a currency with high interest rate and simultaneously selling a currency with low interest rate, after adjusting for volatility of the currency pairs in question—can generate significant returns. However, whenever the markets take a turn for the worse and crisis situations emerge, investors exit such money-making strategies suddenly. As a result, significant losses can occur. In an effort to minimize these potential losses, we also generate an aggregated Risk Metric (RM) that estimates the total risk by looking at various risk indicators across different markets. We use this RM to get timely signals of evolving crises, and to flip the strategy from *long* to *short* (i.e. selling the currency with the high interest rate and simultaneously buying the paired currency with the low interest rate) in a timely fashion, to prevent losses and make further gains, even during crisis periods.

Figure 1 illustrates the performance of a *long-only carry* strategy at different levels of the RM. The negative readings of the RM indicate increasing *risk preference*, while positive readings indicate rising *risk aversion*. The performance of the *long-only carry* strategy at different levels of the RM is measured both in terms of Information Ratio (IR) and the daily success rate. (The performance measures of the currency basket will be explained in Sect. 2.) As seen in the figure, the performance of the currency basket is positive across the board until the RM reaches extreme *risk aversion* levels, upon which returns become severely negative. Therefore, setting the RM to an appropriate threshold level is critical to get timely signals as to when to get out of *long carry* positions, and to revert trades into *short carry* positions.

The trading methodology used in this study seeks to hold *long carry* positions in a currency basket, composed of currency pairs with the highest IRDAV values during normal market conditions, but then to reverse these to *short carry* positions when our aggregated RM indicates the need for increased risk aversion in global markets. We have built this strategy into our Currency Trading Model (CTM). In short, the CTM is a managed version of our

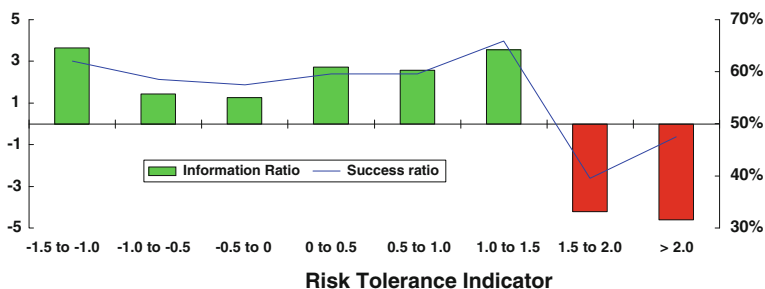


Fig. 1 Performance of a long-only carry strategy at different risk metric levels

IRDAV basket where we use our RM as an overlay over the IRDAV basket, together with other decision variables such as transaction costs, stop-loss levels, and take-profit levels. The CTM is implemented in Excel as a “black box” computational procedure: our numerical tests also indicate that the CTM is highly nonlinear (i.e., multimodal). We then, for the first time in the related financial research, use suitable global optimization methodology—the Lipschitz Global Optimizer (LGO) solver suite linked to Excel—to maximize the performance of the currency basket, based on our selection of key decision variables. LGO integrates several algorithms for constrained global and local optimization: these methods do not require derivatives or other structural model information, except basic model function continuity and local smoothness (to support the theoretical convergence of the embedded algorithms). LGO will be briefly discussed later on, with related references.

In an effort to quantify risk aversion in financial markets, a number of financial institutions and organizations have developed and published indices of risk appetite in the last decade. These include the International Monetary Fund’s Risk Appetite Index (IMF RAI 2003) used for market surveillance, JPMorgan’s Liquidity Credit and Volatility Index (JPM LCVI), the UBS FX Risk Index (UBS FX), Westpac’s Risk Appetite Index (WP), Bank of America’s Risk Appetite Monitor (RAM), Merrill Lynch’s Risk Aversion Indicator (ML RAI), Dresdner Kleinwort’s Aggregate Risk Perception Index (ARPI), and Lehman Brothers’ Market Risk Sentiment Index (MARS). While all these indices more or less try to measure the same thing—risk appetite or aversion in markets—their approach and the variables included in the aggregate risk measure vary. Illnig and Aaron (2005) and Gonzalez-Hermosillo (2008) describe and summarize the methodologies and the risk variables utilized in these indices developed by different investment banks, while Coudert and Gex (2007) analyze the predictive power of several risk appetite indices. Kumar and Persaud (2002), by providing empirical evidence from financial markets, focus on contagion and investors’ changing risk appetite as the primary source of crises. Tarashev et al. (2003) study investors’ attitude towards risk by examining stock option prices. Gai and Vause (2004) review the analytical foundation and measurement of investors’ risk appetite and reconcile a number of different approaches with asset pricing theory. Froot and O’Connell (2003), by looking through international investors’ perspective, show that the data on asset holdings and flows across countries can be used to identify changes in risk tolerance. Although these studies approach risk aversion and risk appetite from a more theoretical perspective, in order to test their hypotheses, they do utilize market-oriented measures (indicators developed using financial market data) as we do in our study. In the end, both these previous studies and our study point out crises periods (i.e. rising risk aversion episodes) on approximately the same dates, with only one to 4-day lags.

Our study, although conceptually similar to the various banks' risk appetite indices in terms of the idea of developing a risk measure to gauge the total risk aversion in financial markets, it is unique in the sense that it generates the aggregate risk measure as a weighted sum of various risk measures from different markets, where the weights of each risk measure are obtained by applying global optimization. To our knowledge, this is the first study to conduct global optimization to generate a risk metric, which is then utilized to maximize the performance of a currency trading basket. We believe that our approach is truly innovative, and one that will contribute significantly both to the growing literature on risk appetite and risk aversion and to practical financial decision making.

Our paper is organized as follows. We describe the CTM and the related optimization problem in Sect. 2. In Sect. 3, we provide details regarding the Excel implementation of the CTM. In Sect. 4, we summarize our general approach to the optimization of "black box" system models and describe briefly the Excel-LGO software implementation. In Sect. 5, we present the numerical solution and optimization results to our CTM. Section 6 concludes the paper followed by references.

2 Model description

2.1 IRDAV basket

Our IRDAV measure uses the same concept as the well-known Sharpe Ratio (SR) defined by:

$$SR = \frac{(r_p - r_f)}{\partial r_p}. \quad (1)$$

In (1) r_p is the return on the asset or portfolio p , r_f is the risk-free interest rate and ∂r_p is the standard deviation of returns from p . For further background and applications of the SR, we refer to Sharpe (1966, 1975, 1992), Treynor and Black (1973), Bodie et al. (1993).

For each currency pair included in the analysis, we calculate IRDAV as the annual interest rate differential between the high-yielding and low-yielding currencies, divided by the 6-month implied volatility of the exchange rate returns. This measure is comparable to JPMorgan's *Carry-to-Risk Index* measure, and it can be expressed by:

$$IRDAV = \frac{(i_A - i_B)}{\partial_{FXa/FXb}}. \quad (2)$$

In (2) i_A is the higher interest rate in country A, i_B is the lower interest rate in country B, and $\partial_{FXa/FXb}$ is the 6-month implied volatility of the exchange rate returns between countries A and B. While both our IRDAV measure and JPMorgan's *Carry-to-Risk* measure combine interest rate differentials and volatility in one indicator (as the ratio of former to the latter) and utilize the London Inter-Bank Offer Rates (LIBOR) to measure the size of interest rate differences (*carry*) between countries, the volatility measures utilized in these two indicators differ significantly. In IRDAV, we measure volatility as the 6-month implied volatility of the exchange rate returns, whereas JPMorgan's *Carry-to-Risk* index measures volatility as the annualized standard deviation of the daily exchange returns calculated over a 12 month period. That is, we use implied, as opposed to realized FX volatility, and our term to measure volatility is 6 months as opposed to 1 year. The use of shorter term volatility enables us to have more frequent changes in rankings of IRDAV, which in turn, enables our trading strategy to switch in between currency pairs in a more timely fashion to better reflect the movements in

the currency markets. For further background, we refer to [Çağlayan and Giacomelli \(2003, 2005\)](#), and [Normand et al. \(2004\)](#).

The currencies of the developed countries included in our analysis are USD (US Dollar), EUR (Euro), JPY (Japanese Yen), GBP (British Pound), CHF (Swiss Franc), CAD (Canadian Dollar), SEK (Swedish Krona), NOK (Norwegian Kroner), AUD (Australian Dollar), and NZD (New Zealand Dollar). The specific currency pairs included in our analysis are determined according to the liquidity and transaction volume of the Foreign Exchange (FX) pairs in currency markets. Basically, we have chosen the 18 most liquid FX pairs that have the highest transaction volumes: these are AUD/CAD, AUD/JPY, AUD/USD, EUR/CHF, EUR/GBP, EUR/JPY, EUR/NOK, EUR/SEK, EUR/USD, GBP/CHF, GBP/USD, NOK/SEK, NZD/USD, USD/CAD, USD/CHF, USD/JPY, USD/NOK, and USD/SEK. At the beginning of each month, we rank these currency pairs according to their IRDAV values; then we select the top IRDAV-ranked pairs to compose the actual (monthly) currency basket, where each currency pair in the basket receives an equal portfolio weight. The top ranked FX pairs constitute our *carry* basket for the given month, until new rankings are computed at the beginning of the following month. Within each given month, the *long carry* positions in the basket are held until our RM (which is computed daily) reaches a certain threshold value, whereupon the basket is flipped to *short carry* positions. Note that, we use global optimization to select the best fixed number of currency pairs (out of 18) that should be included in the basket while we maximize the overall performance of the currency portfolio.

2.2 Risk metric

In currency markets, timing the change from the normal markets' "*search for yield*" status to the crisis periods' "*flight to quality*" status is crucial, especially for *carry* models. During these crises, investors suddenly empty their *long carry* positions: thereby a huge depreciation in high-yielding currencies results, together with a significant appreciation of the low-yielding currencies. Consequently, holding *passive* (i.e. *long-only*) *carry* positions in a currency basket at all times does not make sense from a trading perspective, as these positions inevitably get hit very harshly, potentially wiping out long-term positive returns within a matter of weeks. Hence, investors need to be aware of sentiment changes in financial markets to better protect themselves against a sharp decline in the performance of *long carry* trading strategies.

We attempt to gauge the timing of these "sentiment shifts" in the financial markets with our new RM. In order to assess an aggregated risk measure across a wide spectrum of markets, our RM includes six different risk measures, each representing a unique risk for a particular market. Specifically, these six risk measures include: (1) High yield bond spreads over US Treasuries: a risk measure for credit markets; (2) JPMorgan's Emerging Market Bond Index (EMBI+): a risk measure for emerging markets; (3) VIX (implied volatility of the S&P 500 index options): a risk measure for equity markets; (4) Gold-to-Copper price ratio: a risk measure for commodity markets; (5) Swiss Franc trade-weighted (CHF) index: a risk measure for currency markets; and (6) 10-year Swap spreads: a liquidity risk measure. Note that some of these individual risk measures do appear in the literature in the development of risk aversion/appetite indices utilized by different investment banks. For example, High Yield bond spreads over US Treasuries, EMBI+, VIX, and 10-year swap spreads in different forms and/or extensions have been utilized in JPMorgan's LCVI, UBS's FX Risk Index, Westpac's RAI, Bank of America's RAM, and Merrill Lynch's RAI.

Table 1 Performance of carry basket with alternative risk measures utilized as overlays

	Annual return (%)	SD (%)	Information ratio (Return/SD)
Passive long-only strategy	4.29	7.23	0.59
Carry basket with High yield spreads as an overlay	7.65	6.22	1.23
Carry basket with EMBI+ as an overlay	8.52	6.15	1.39
Carry basket with VIX as an overlay	7.03	6.07	1.16
Carry basket with Gold-to-copper price ratio as an overlay	5.21	5.95	0.87
Carry basket with Trade-weighted CHF as an overlay	4.76	5.93	0.80
Carry basket with Swap spreads as an overlay	7.27	5.69	1.28
Carry basket with equally-weighted risk metric as an overlay	10.90	6.01	1.81
Carry basket with optimal weights in risk metric as an overlay	14.98	5.58	2.69

The inclusion of these specific risk measures (in the construction of different banks' risk aversion/appetite indices), which are selected only after running numerous entry tests on various multiple risk measures, shows not only their importance but also their ability to successfully monitor sentiment changes in the markets. Here, in Table 1, we present results from our own empirical analyses and provide evidence that all of these six risk measures (that make up our RM) have proven to point out past periods of sentiment changes successfully. In other words, these six risk measures specified above have given timely signals when to exit *long carry* positions, and to reverse these positions to *short*, thereby making possible further gains even during crisis periods. Table 1 shows during the period 1998–2009 the improvement in the performance of a *passive (long-only)carry* strategy, when the aforementioned individual risk measures are utilized one at a time (separately) as overlays to give signals as to when to flip positions from a *long carry* trade to a *short carry* trade strategy (or vice versa from a *short carry* trade to a *long carry* trade strategy). As can be seen clearly in the table, the use of each individual risk measure separately as an overlay over the simple *long-only* strategy significantly enhances the performance of the *long-only* strategy both in terms of annual returns and information ratios, where IR is measured as the ratio of annual returns to annualized standard deviation of returns. The improvement in the performance of the basket over the *long-only* strategy in each instance shows that all of these six risk measures, as stand-alone indicators, are capable of picking the periods of sentiment changes successfully. For comparison purposes, we also show towards the bottom of Table 1, the performance of the basket where these six risk measures are combined with equal weights to generate an aggregated RM. The performance of the basket where the aggregated RM (with equal weights) is used as an overlay measure over the simple *long-only* strategy is clearly much better compared to the cases where single individual risk measures are utilized as overlay measures, suggesting that these individual risk indicators measure different types of risk in different markets and therefore their combination is more powerful (i.e. successful) in signalling rising risk aversion (or vice versa, improving sentiment) in markets.

We claim, however, that instead of simply giving equal weights to individual risk measures in construction of the RM (as is done by most of the major investment banks), using a global optimization based strategy, which gives the highest weight to a risk indicator that is superior to the other risk measures, could greatly enhance the performance of the *carry* basket. Indeed, as can be seen at the bottom of Table 1, the use of global optimization to set the optimal weights in construction of the RM significantly increases the performance of the

basket both in terms of annual returns and information ratios (optimization-based strategy's 14.98% average annual return and 2.69 IR versus the equal weight strategy's 10.90% average annual return and 1.81 IR). Our innovative use of global optimization to generate the weights of individual risk factors within the RM (to maximize the performance of an *active carry* trade strategy) has never been done in the literature before and we show in Sect. 5, in detail, that using global optimization significantly enhances the performance of our *carry* strategy.

In terms of the construction of the RM, on a daily basis, we first calculate the *Z*-scores (number of standard deviations away from a 6-month average) for each of the six individual risk measures using the formula:

$$Z_i = \frac{(RM_t^i - \mu_t^i)}{\partial RM_t^i}. \quad (3)$$

In (3) RM_t^i is the risk measure i on day t , μ_t^i is the 6-month rolling average of the risk measure i on day t , and ∂RM_t^i is the 6-month standard deviation of the risk measure i on day t . Given the individual scores Z_i , we take their weighted sum to obtain the final daily reading on our RM. Note that, in addition to the weights of each risk measure in the RM, the threshold overall *Z* score level upon which the currency basket positions are flipped back and forth between *long carry* and *short carry* is also determined by the optimization procedure.

One can argue that—due to their apparent connection to the financial sector—the six risk measures used in constructing the aggregate RM are not independent, and may, in fact, be correlated. The optimized weights also serve to take care of this problem (another advantage of using optimization in developing a risk metric), since optimization can potentially give a much lower weight to a risk indicator that is highly correlated with another risk measure—especially to those indicators that do not perform as well as the others. We believe that the global-optimization approach utilized here could also be applied to various investment banks' similar risk metric measures, whereby the effectiveness of banks' risk metrics in pointing out the extreme risk aversion periods can be maximized by optimizing the weights of the individual risk indicators (decision variables) within a global optimization framework.

2.3 Other decision variables and properties of the currency trading model

In addition to IRDAV-ranked currency pairs and aggregated RM, our CTM incorporates all the specifications required for a full-fledged currency trading model, including transaction costs, stop-loss levels, and take-profit levels. The returns reported in this paper are all transaction-cost adjusted. Specifically, the model deducts 0.05% return from each trade in the basket, a realistic assumption given the size of the bid-ask spreads for most of the major currency pairs. Moreover, the model calculates the “*drawdown*” from the peak of returns, and the “*accumulated returns*” from the bottom of returns to use stop-loss and take-profit levels efficiently. We prefer to apply stop-losses and take-profits on the basket rather than on the individual currency pairs so that an extraordinary performance (negative or positive) in one of the currency pairs can be compensated by the performance of other pairs in the basket, without triggering stop-loss and take-profit signals too frequently. Once a stop-loss or take-profit signal is triggered on the basket portfolio; however, the model requires a certain number of days to wait “neutral” before re-entering the trades again. From a financial perspective, since there is no definite “scientific” answer on what a stop-loss or take-profit level should be, and on how many days should one wait “neutral” before re-entering trades after a stop-loss or take-profit, we let the global optimization software determine those key parameters as well, while maximizing the overall performance of the currency basket.

2.4 Measuring the performance of the currency trading model

Investors, as a principle, demand high returns with minimum volatility from their investments. The IR, defined as the ratio of returns to volatility, includes these two important concepts of finance in one single indicator. That is, the higher the IR from an investment, the more attractive/desirable the investment is for the investor. We measure the overall performance of our CTM with the IR, where IR is defined as:

$$IR = \frac{R_{CTM}}{\partial R_{CTM}}. \tag{4}$$

Here R_{CTM} is the compounded annualized return from the CTM basket during the sample period 1998–2009, and ∂R_{CTM} is the annualized standard deviation of daily returns from the CTM basket during the same period. Recall that the CTM is an optimally managed version of our IRDAV basket where we use our RM as an overlay over the IRDAV basket together with the other aforementioned decision variables (such as stop-loss, take-profit levels) explained in previous sections.

We maximize the IR of the CTM by using global optimization methodology and software to determine the best values for the following key decision variables:

- (1) The (integer) number of currency pairs to be included in the IRDAV basket
- (2) The normalized weights that each risk measure receives in the construction of the RM
- (3) The threshold RM level where positions are switched back and forth between *long carry* and *short carry*
- (4) The stop-loss level of the CTM
- (5) The take-profit level of the CTM
- (6) The (integer) number of days to wait “neutral” after a stop-loss signal is triggered, and
- (7) The (integer) number of days to wait “neutral” after a take-profit signal is triggered

Based on the discussion above, our CTM can be summarized as follows:

Objective function

$$\max IR = \frac{R_{CTM}}{\partial R_{CTM}}$$

Weight factor normalization constraint

$\sum_{i=1}^6 w_i = 1$ here w_i is the weight that risk measure Z_i receives in the aggregated risk metric.

$$0 \leq w_i \leq 1; i = 1, \dots, 6$$

Decision variables, with corresponding given lower and upper bounds

$$ncp_l \leq ncp \leq ncp_u \text{ (an integer variable)}$$

$$rmt_l \leq rmt \leq rmt_u$$

$$sl_l \leq sl \leq sl_u$$

$$tp_l \leq tp \leq tp_u$$

$$nds_l \leq nds \leq nds_u \text{ (an integer variable)}$$

$$ndt_l \leq ndt \leq ndt_u \text{ (an integer variable)}$$

In the CTM, w_i are the weights of our six risk measures that constitute our aggregated RM; ncp is the number of currency pairs to be included in the basket; rmt is the RM threshold level where positions are switched from *long carry* to *short carry*, or vice versa; sl is the stop-loss level; tp is the take-profit level; nds is the number of days to wait neutral after a

stop-loss signal; ndt is the number of days to wait neutral after a take-profit signal. The lower and upper bounds for each decision variables are denoted by attaching the subscript letters l and u to the name of the decision variables.

The respective lower and upper bounds used for each of the decision variables are defined as follows: 0 and 100% for w_i , 2 and 18 for the n_{cp} , 0.5 and 3.0 for the r_{mt} , -4.0 and -1.0% for the sl , 2.0 and 15.0% for the tp , 1 and 20 for the nds , and 1 and 20 for the ndt . The lower and upper bounds for these decision variables were selected so that they make sense from a financial market perspective. Specifically, for the n_{cp} , the lower bound 2 is literally the minimum number of currency pairs to generate a basket and the upper bound 18 is the total number of currency pairs in our sample. For the r_{mt} , the lower bound 0.5 and the upper bound 3.0 are selected so that a reasonable number of days can be labelled as risk-aversion periods. That is, r_{mt} less than 0.5 results in too many days marked as risk-averse periods (almost 40% of our total sample), which does not make sense from a financial market perspective. On the other hand r_{mt} more than 3.0 generates very few days of risk-aversion (less than 1% of our total sample), which again does not make sense from a financial market perspective. For the sl , the lower bound -4.0% is selected so that the trading system does not generate losses beyond 4%, a level that most traders in the financial industry pick as the most that they can tolerate before getting out of positions. On the other hand, -1.0% upper bound for sl is selected so that losses less than 1% do not trigger stop-loss and give a chance for recovery. For the tp , the lower bound 2.0% is selected so that gains less than 2% are not cut short (i.e. positions are not closed too early) before a possible huge rally and the upper bound 15.0% is selected so that the trading system can protect its gains after accumulating a significant 15% return, a rare instance in financial markets, which usually follows a sharp reversal. For nds and ndt , the lower bound 1 is literally the minimum number of days the trading system can wait after a stop-loss or take profit. The upper bound of 20 days for both nds and ndt are selected so that the trading system does not stay neutral beyond 1 month, as there are typically only 22 non-weekend trading days in a month. For comparison purposes, for each of the aforementioned decision variables, we tried broadening the gap between the lower and the upper bounds, but still the global optimization generated the same optimal values for the key decision variables, suggesting that our selection of lower and upper bounds do not impact the outcome of the optimization. We will provide more details regarding the Excel implementation of the CTM in Sect. 3.

Using the available historical data between 1998 and 2009, for each selection of the decision variables listed above, the model calculates the corresponding objective function value (the Information Ratio, IR) based on the entire sample period. Notice that in our model, some of the decision variables are discrete, while others are continuous: hence, we are facing a moderate size mixed integer-continuous nonlinear programming (MINLP) model which in actual practice depends on stochastically varying input. We will address the issue of discrete variables later on in Sect. 4.1.

Based on our extensive set of numerical tests, we have also observed that the CTM is highly nonlinear. Due to its outlined characteristics, the parameterization of the CTM requires global scope derivative-free optimization methodology and software. Our preliminary results showed that the built-in Excel solver could not improve upon a given nominal solution. For this reason, we have chosen to use the LGO software linked to Excel, to find numerically the best possible parameter settings for our CTM.




Fig. 2 The modules of the currency trading model

3 Excel implementation of the currency trading model

The CTM summarized above has been implemented in the Microsoft Excel modelling environment as a stand-alone computational procedure based on the interaction of a set of worksheets. The diagram above shows how the different worksheets of our Excel model interact with each other, in order to calculate the performance of the currency basket, for each setting of the key decision variables.

As depicted in Fig. 2, the CTM Excel spreadsheet is composed of eight worksheets, represented by the individual boxes. Seven of the eight worksheets constitute different components of the CTM and the final worksheet called *CTM Performance* is the worksheet where all aspects of the model are brought together to measure the performance of the currency basket.

The *FX Rates* worksheet is where we store daily exchange rates of the 18 currency pairs included in our analyses. Our sample encompasses a 12-year period from January 1998 to December 2009. The data for the daily exchange rates are obtained from Bloomberg Financial Services, Inc.

In the *FX Volatility* worksheet, we first calculate the daily exchange rate returns on the 18 currency pairs, by using the exchange rates from the *FX Rates* worksheet. Then, using these daily exchange rate returns, we calculate the 6-month FX implied volatility for each of the 18 currency pairs.

The *Interest Rates* worksheet records the daily 1-month London Inter-Bank Offer Rates (LIBOR) for the currencies of the 10 developed countries (US, Euro area, Japan, United Kingdom, Canada, Switzerland, Sweden, Norway, Australia, and New Zealand) that we include in our study. The daily LIBOR rates are also obtained from Bloomberg Financial Services, Inc.

In the *Interest Rate Differentials* worksheet, using the data from the *Interest Rates* worksheet, we calculate the size of *positive carry* (i.e. the magnitude of the interest rate differentials between high-yielding and low-yielding countries) for each of the 18 currency pairs in question. We also determine which positions (*long* or *short*) in each currency pair would generate that positive carry.

In the *IRDAV Values* worksheet, after computing interest rate differentials and the implied FX volatility separately, we bring these two components together and compute the ratios of interest rate differentials to FX volatility for each of the 18 currency pairs. We calculate IRDAV values at the beginning of each month and rank the 18 pairs from highest to lowest based on their IRDAV values. The IRDAV based rankings help determine which currency pairs are going to be included in the basket for that month.

The *Risk Measures* worksheet contains the daily data used to monitor six different risk indicators (described in Sect. 2.2) from various markets, including credit markets, emerging

markets, equity markets, commodity markets, and currency markets. The daily data for these risk measures are also obtained from Bloomberg Financial Services, Inc.

In the *Risk Metric* worksheet, using the data from *Risk Measures* worksheet, we first calculate the *Z*-scores (number of standard deviations away from a 6-month average) for each of the six individual risk measures on a daily basis. Then, we take the weighted average of these *Z*-scores to get the final daily reading on the RM. Note that the weights of the individual risk measures as well as the overall risk metric threshold level where positions are switched between *long carry* and *short carry* are determined by the global optimization.

Finally, all these components of the model feed into the *CTM Performance* worksheet, where all key decision variables are brought together so that we can run the optimization and do all the calculations to generate the performance statistics (*Information Ratios*, annual returns, and standard deviation of returns) for the currency basket. Specifically, in the *CTM Performance* worksheet we use the data from *FX Rates* and *Interest Rates* worksheet to calculate the daily total returns (movements in the exchange rates plus the daily interest differential earned on the currency position held) on each currency pair included in the basket. In the same *CTM Performance* worksheet, we also use the *IRDAV* values and the rankings of the currency pairs from the *IRDAV Values* worksheet, as an input to determine specifically which currency pairs are going to be included in the currency basket so that we can calculate the performance of the basket with the correct currency pairs. The *CTM Performance* worksheet also takes the daily RM readings from the *Risk Metric* worksheet to signal when the risk metric threshold level is reached so that the positions in the basket can be flipped from *long carry* to *short carry*, or vice versa. Finally, in the *CTM Performance* worksheet, we consider all the specifications that define a general trading model, including transaction costs, stop-loss levels, take-profit levels, and the number of days to wait neutral after a stop-loss or take-profit signal.

It is important to point out that the computational effort related to the optimization is fairly significant. With its given settings (chosen by us), the global optimization run required over 18 hours runtime on a current laptop computer: during this run, over 130,000 optimization iterations (model function evaluations) were executed. These computational requirements are primarily due to two factors. First, the number of decision variables is small but non-trivial. Our preliminary results obtained by local search from several selected starting points clearly indicate that the “black box” model is indeed multimodal. Therefore, based on our experience with LGO and other global optimization software, a global sampling effort in a multi-dimensional decision space with tens or hundreds of thousands of model function evaluations has seemed reasonable. Second, the function evaluations themselves within the CTM are quite complex, and each model evaluation requires about half a second.

To put the cited number of function evaluations in perspective, note that naïve search strategies would require a substantially greater—unaffordable—computational effort. For example, assume that we wish to obtain just 10% relative precision in each (by assumption, continuous) variable by using a uniform grid of sample points. This would then require 10^{12} function evaluations, before (optionally) turning to a local search to refine the numerical solution estimate obtained. Instead, we use about $(1.3) * 10^5$ function evaluations—still a substantial, but manageable computational effort. Given the significance of obtaining a high-quality result in potential real-world applications of our study, we believe the total runtime required should not be seen as excessive either.

Given the fact that the parameterization of the CTM leads to a multimodal optimization problem that comes from the real world of finance, its globally optimal solution is *de facto* unknown. Therefore our objective has been to apply a robust and efficient numerical

optimization procedure to find a high-quality solution which then can be compared to problem-specific “expert solution” estimates.

In the next section we briefly describe the mathematical modelling background and computational tools related to our study. We then provide illustrative numerical results in Sect. 5, and discuss how the optimization procedure has improved the performance of the currency basket.

4 Parameter optimization of “black box” system models using LGO

4.1 Global optimization model statement

The CTM is a special case of the general continuous global optimization (GO) model shown below.

$$\min f(x) \quad \text{subject to } x \in D := \{x : l \leq x \leq u; g_j(x) \leq 0; j = 1, \dots, m\}. \quad (5)$$

In (5) we apply the following notation, and make related assumptions:

$$\begin{aligned} x \in \mathbf{R}^n & \quad n\text{-dimensional real-valued vector of decision variables;} \\ f : \mathbf{R}^n \rightarrow \mathbf{R} & \quad \text{continuous (scalar-valued) objective function;} \\ D \subset \mathbf{R}^n & \quad \text{non-empty set of feasible solutions, a proper subset of } \mathbf{R}^n. \end{aligned}$$

The feasible set D in (5) is defined by

$$\begin{aligned} l \in \mathbf{R}^n & \quad \text{component-wise finite lower bounds on } x; \\ u \in \mathbf{R}^n & \quad \text{component-wise finite upper bounds on } x; \\ g : \mathbf{R}^n \rightarrow \mathbf{R}^m & \quad m\text{-vector of continuous constraint functions.} \end{aligned}$$

Although further technical details are outside the scope of this paper, note that the concise model formulation (5) encompasses many special cases. Specifically, (5) subsumes the general class of mixed integer-continuous (nonlinear) optimization problems, since finitely bounded integers can be represented in the 2-based number system by a sequence of binary variables, and then each binary variable can be described by a pair of trivial bound constraints and a single non-convex (reverse convex) constraint. Thus, our CTM can be considered and numerically treated as an instance of (5).

The key “black box” characteristic of the CTM is due to the fact that its objective function and all embedded functionality needs to be evaluated by a computational procedure implemented in Excel. Although the functions referred to are mostly continuous, some of them are not always so at certain arguments. Similar issues frequently arise in the real world of optimization when the model itself can be a highly complicated and/or a confidential “black box”. The handling of such problems typically requires global scope, robust and efficient derivative-free search methodology, and the capability to cope with the possible occasional discontinuities of the (merely computable) model functions.

To illustrate—as well as to verify—the need for using global optimization, we conducted a number of numerical experiments to explore the performance of the CTM. Our preliminary calculations clearly showed that the basic Excel Solver could not improve upon a given (arbitrarily chosen) nominal solution. In addition, our results also illustrate that different starting points (initial model parameter settings) can lead to corresponding sub-optimal CTM performance, indicating the multimodality of the CTM. These findings altogether indicate the essential requirement to apply global optimization to the CTM parameterization problem.

4.2 The LGO software package for nonlinear optimization

Developed since the late 1980s, the Lipschitz Global Optimizer solver engine is currently available for a range of compiler platforms (C, C++, C#, FORTRAN), with links to several prominent optimization modeling languages (AIMMS, AMPL, GAMS, MPL), and to the leading integrated technical computing systems Maple, *Mathematica*, and MATLAB. For algorithmic and implementation details not discussed here, we refer to [Pintér \(1997, 2002, 2007, 2009, 2010a\)](#).

The overall design of LGO is based on the flexible combination of (continuous nonlinear) global and local optimization strategies, with theoretically established convergence properties. The global solvers are continuous branch-and-bound, stochastic single-start and stochastic multi-start algorithm implementations; the local solver is a generalized reduced gradient algorithm implementation. LGO hence can be used for both global and local optimization, under basic continuity conditions. Let us also point out here that Lipschitz-continuity is theoretically required only by one of the global solver options: this option has not been used in the present numerical study.

LGO can also handle MINLP models with a relatively small number of integer variables following the basic reformulation approach outlined above. For example, LGO's current (standalone) MINLP solver capabilities enable the solution of many of the MINLP test problems collected in [Floudas et al. \(1999\)](#); however, we do not make unjustified claims regarding competitive performance for sizeable MINLP models. As our summary results presented in Sect. 5 indicate, LGO successfully handles also the CTM which has only a few integer variables.

LGO's derivative-free design is different from many other nonlinear optimization software packages that require explicit analytical model function information to support model parsing, function decomposition, and other specific operations such as function bounding procedures. Such an analytical approach can support the deterministically guaranteed “precise” solution of a range of GO models which, however, have to be defined in terms of a given set of possible model component functions. In contrast to this approach, the design of LGO also supports the handling of “black box” problems (such as our CTM) that will remain outside of the scope of the analytical optimization approaches referred to above.

In numerical practice, LGO's global search options generate a global solution estimate(s) that is (are) refined by the seamlessly following local search mode(s). Hence, the expected result of using LGO is a global and local search based best feasible solution that meets at least the local optimality criteria. (Note that the latter are theoretically guaranteed only under standard *local* smoothness conditions.) At the same time, one should keep in mind that no global—or other—optimization software will always work satisfactorily, with default option settings and under preset resource limitations related e.g. to the maximum allowed solver execution time or the maximum number of model function evaluation.

Extensive numerical tests and a broad range of practical applications demonstrate that LGO and its platform-specific implementations can find a close numerical approximation of the (known or putative) global solution not only when using “academic” GO test problems, but also in far more complicated and sizeable models. A few illustrative examples of real-world optimization challenges handled by various LGO implementations are laser design ([Isenor et al. 2003](#)), intensity modulated cancer therapy planning ([Tervo et al. 2003](#)), integrated oil and gas production systems operations ([Mason et al. 2007](#)), automotive component design ([Goossens et al. 2007](#)), industrial packing problems ([Castillo et al. 2008](#)), telecommunications networks ([Zukerman et al. 2008](#)), and fuel processing technology analysis ([Pantoleontos](#)

et al. 2009). Note that LGO has also been used previously in the context of other significant financial applications.

4.3 The excel-LGO software implementation

Recently, we developed an LGO solver link to Excel with an easy-to-use worksheet interface to assist users in the model formulation and solution process. Excel-LGO is described in [Pintér 2010b](#); therefore, we will only summarize some essential points here. (It is important to note however, that the Excel-LGO development project and the resulting software product version discussed here is entirely independent from the Excel Solver platform development by Frontline Systems.)

In order to link an Excel based optimization model to LGO, the complete model development spreadsheet is simply extended by an additional worksheet with the name “LGO”. The optimization is then carried out using this new worksheet: therefore all model definitions and dependency links have to be properly imported here to given (standardized) positions. The key static model definitions include the decision variable names, lower bounds, optionally supplied nominal values, and upper bounds, as well as objective and constraint function names, and constraint types. The dependency links serve to evaluate all model function values, based on the algorithmically chosen sequence of decision variable settings. After the model is linked to the “LGO” worksheet, it is ready to solve using Excel-LGO. The outlined model setup structure is self-documenting and transparent, and it supports well the development and solution of complex models such as our CTM.

5 Numerical results and discussion

In an effort to test the potential benefits of a global optimization based parameterization of the CTM, first we divided our 12-year total sample period (1998–2009) into two sub-samples, with 1998 through 2003 being the in-sample (training) period, and 2004 through 2009 being the out-of-sample (verification) period. Next, we solved the model calibration problem using the in-sample period data only, by optimizing the key decision variables introduced in the CTM. Based on the optimal variable values found during the training period, we then checked the performance of the CTM basket during the verification period. For comparison, we also evaluated the out-of-sample period performance of the optimized CTM basket to the performance of a basket where we manually picked the values for each of the key decision variables based on expert intuition and simple “trial and error”.

The numerical results summarized in [Table 2](#) clearly show a significant improvement in the IR based performance of our currency basket after the optimization. The table shows the globally-optimized performance of the CTM basket during the in-sample period as well as the out-of-sample period. The IR of the CTM basket during the training period 1998–2003 (for which we ran the Excel-LGO optimizer) is an eye-opening 3.22. Informally, this means 3.22 units of return for each unit of risk taken, compared to standard financial trading models’ average information ratios that typically lie between 1.00 and 1.50. More importantly, the superior performance of the CTM basket essentially continues during the verification period 2004–2009 (where we use the values of the decision variables derived in the training period, without re-optimization). In fact, during the out-of-sample period, the CTM basket generated 14.3% annual compounded returns, just shy of the 15.7% annual compounded returns calculated during the optimized in-sample period; and the IR of the verification period is still

Table 2 Global optimization based performance of the CTM currency basket

	Annual return (%)	SD (%)	Information ratio
1998	24.49	5.88	4.17
1999	12.47	5.09	2.45
2000	13.58	4.13	3.29
2001	13.38	4.61	2.91
2002	15.32	4.33	3.53
2003	15.36	5.06	3.04
2004	16.29	5.49	2.97
2005	12.74	3.46	3.68
2006	11.84	4.16	2.85
2007	6.33	6.39	0.99
2008	7.28	7.67	0.95
2009	33.13	8.41	3.94
1998–2003 in-sample	15.70	4.87	3.22
2004–2009 out-of-sample	14.27	6.20	2.30
1998–2009 full sample	14.98	5.58	2.69
Key decision variables		Risk metric weights	
Number of pairs in basket	5	High yield	0.245
Risk metric threshold	1.98	EMBI+	0.047
Stop-loss level	–3.66%	VIX	0.231
Take-profit level	9.95%	Gold-to-copper	0.260
Number of days to wait after stop-loss	14	CHF index	0.105
Number of days to wait after take profit	15	Swap spreads	0.112

very high at 2.30. In addition, looking at the yearly performance of the CTM basket during the out-of-sample period in Table 2, we see that the outstanding performance of the basket during 2004–2009 was not due to a couple of outlier “star years”. In fact, the CTM basket generated positive, robust, and consistent returns every single year during 2004–2009, with double-digit returns in 4 out of 6 years.

As an alternative strategy to test the potential benefits of a global optimization based parameterization of the CTM, we also ran the optimization on a 1-year rolling-window basis and checked the performance of the CTM in the following year using the values of the decision variables derived from previous year’s optimization. More specifically, we started with the first year of our sample period 1998, and ran the optimization for this year only and obtained optimized values of our key decision variables. Then, using these optimized values for the key decision variables, we generated the out-of-sample period performance of our CTM basket for the year 1999. Next, we ran the optimization for the year 1999 and obtained a new set of optimized values for our key decision variables, and then used this new set of optimized values to generate the out-of-sample period performance of our CTM basket for the year 2000. We followed the same procedure every single year until we reached the end of our sample period in 2009. This strategy generated annual out-of-sample period performance returns for our CTM basket from 1999 through 2009. In Table 3 we report the out-of-sample period performance of our CTM basket from this aforementioned strategy

Table 3 One-year rolling window based optimized performance of the CTM currency basket

	Annual return (%)	SD (%)	Information ratio
1998	NA	NA	NA
1999	10.30	5.09	2.02
2000	11.46	4.13	2.78
2001	10.82	4.60	2.35
2002	11.84	4.32	2.74
2003	10.14	5.04	2.01
2004	15.31	5.48	2.79
2005	11.41	3.46	3.30
2006	9.78	4.15	2.36
2007	8.18	6.39	1.28
2008	9.47	7.68	1.23
2009	30.95	8.40	3.68
1999–2003 sample	10.91	4.65	2.35
2004–2009 sample	13.93	6.19	2.25
1999–2009 sample	12.55	5.54	2.26

both in terms of annual returns and information ratios. Similar to the results obtained earlier, the new 1-year rolling-window optimization strategy also generates very attractive returns. For example, for the full sample period 1999–2009, our CTM generated a very respectable annual compounded return of 12.55%, with an IR of 2.26. More importantly, analyzing the yearly performance of the CTM basket in Table 3, we see that the CTM basket actually generated positive, robust, and consistent returns every single year during 1999–2009, even with double-digit returns in 8 out of 11 years.

Next, for comparison we evaluated the out-of-sample period performances of these optimized baskets to the performance of a basket where we manually selected values for each of the key decision variables based on expert intuition and some trial and error, without using a formal optimization procedure, an approach traditionally used by most financial organizations.

Table 4 shows the performance of the CTM basket with the manually-selected (expert estimates-based, best found) values for the key decision variables. Using the “traditional” model parameter selection approach, we were able to produce information ratio values of the CTM basket for the period 2004–2009 (the first optimization method’s out-of-sample, verification period) up to 1.60. For comparison, as reported in Tables 2 and 3, the information ratios during the same out-of-sample period 2004–2009, after global optimization, are 2.30 and 2.25, respectively. In the world of finance, this is an eye-opening nearly 44 and 41% improvement in the overall performance of the currency basket which is certainly not due to random circumstances and data peculiarities. For further comparison, see also the full sample based IR values in Tables 2 and 3 relative to Table 4. Even this relatively simple comparative exercise shows the significant advantage of using global optimization to select the key parameters of our CTM that reflects real-world complexity.

Finally, in order to further illustrate the benefits of a global optimization based parameterization of our model, we replicated several different investment banks’ risk appetite/aversion indices (constructed as equally-weighted averages of various risk measures) and used these alternative risk indices as an overlay over our IRDAV basket (to flip positions in between *long carry* trade and *short carry* trade during low risk / high risk periods) and compared the

Table 4 Performance of the CTM currency basket without formal optimization

	Annual return (%)	SD (%)	Information ratio
1998	16.01	6.40	2.50
1999	9.18	4.91	1.87
2000	4.88	4.50	1.08
2001	7.59	5.26	1.44
2002	15.62	4.67	3.35
2003	12.95	5.00	2.59
2004	5.55	4.84	1.15
2005	8.51	3.96	2.15
2006	9.47	4.18	2.26
2007	13.28	7.63	1.74
2008	12.95	9.37	1.38
2009	15.52	8.50	1.83
1998–2003 in-sample	10.96	5.14	2.13
2004–2009 out-of-sample	10.83	6.76	1.60
1998–2009 full sample	10.90	6.01	1.81
Key decision variables		Risk metric weights	
Number of pairs in basket	5	High yield	0.1667
Risk metric threshold	1.25	EMBI+	0.1667
Stop-loss level	−2.50%	VIX	0.1667
Take-profit level	8.00%	Gold-to-copper	0.1667
Number of days to wait after stop-loss	5	CHF index	0.1667
Number of days to wait after take profit	1	Swap spreads	0.1667

performances of these baskets to our CTM basket where weights on individual risk measures are assigned after a global optimization. Table 5 reports, for the out-of-sample period 2004–2009, the performance of our optimized CTM basket as well as the performance of alternative FX trading strategies when different investment banks' risk appetite/aversion indices are utilized as overlay measures over our IRDAV basket. As can be seen in Table 5, for the period 2004–2009, applying other investment banks' risk appetite/aversion indices as overlay measures over our IRDAV basket generates annualized compounded returns ranging from 7.58 to 9.41%, with information ratios ranging in between 1.32 and 1.59. On the other hand, applying our aggregated RM with weights derived from global optimization on our IRDAV basket (i.e. our CTM strategy) generates an impressive out-of-sample period annual compounded return of 14.27% with a remarkable IR of 2.30. Note that, if we had not used global optimization in the construction of our RM and instead had given equal weights to each of the six measures that make up our RM, then, during the same time period 2004–2009, our trading strategy would have generated an annual compounded return of only 10.83%, with an IR of 1.60, a performance which is not any different than the performances of other trading baskets tested earlier when alternative banks' equally weighted risk measures are utilized. All in all, these statistics show the significant value added to the performance of a *carry* trade strategy when globally optimized weights, as opposed to equal weights, are assigned to individual risk measures in the development of risk appetite/aversion indices.

Table 5 Performance of IRDAV basket when investment banks' alternative risk measures utilized as overlays

	Annual return (%)	SD (%)	Information ratio
IRDAV basket with JPMorgan's LCVI as an overlay	9.41	6.01	1.57
IRDAV basket with UBS's FX risk index as an overlay	9.25	5.93	1.56
IRDAV basket with Westpac's WP as an overlay	7.58	5.74	1.32
IRDAV basket with Bank of America's RAM as an overlay	9.02	6.09	1.48
IRDAV basket with Merrill Lynch's RAI as an overlay	8.98	5.65	1.59
IRDAV basket with equally-weighted 6 measures of our risk metric as an overlay	10.83	6.76	1.60
IRDAV basket with optimal weights of our risk metric as an overlay (CTM trading basket)	14.27	6.20	2.30

6 Conclusions

In this study we introduce a new financial indicator called the Interest Rate Differentials Adjusted for Volatility (IRDAV). On a monthly basis, we rank currency pairs according to this measure, and then select a basket of pairs with the highest IRDAV values. Under positive market conditions, an IRDAV based investment strategy can generate significant returns. However, when the markets turn for the worse and crisis situations evolve, investors exit such money-making strategies suddenly, and as a result, significant losses can occur. In an effort to minimize these potential losses, we also propose an aggregated RM that estimates the total risk by looking at various financial indicators across different markets. We use this RM to get timely signals of evolving crises and to flip the trading strategy from *long* to *short* in a timely fashion, to prevent losses and make further gains even during crisis periods.

Using real-world data, we demonstrate that by applying suitable global optimization tools we can optimize the key parameters and thereby maximize the performance of our CTM. This innovative and unique trading methodology significantly enhances the performance of a simple *long-only carry* strategy. Our findings are supported by numerical results derived using actual market data that span over 12 years. For completeness, we also briefly review the LGO nonlinear optimization software package and its Excel-LGO implementation.

Our results clearly show the need for, and the validity of, a global optimization based CTM parameterization approach. The combination of the CTM with global optimization leads to a currency trading strategy that is readily applicable, and that leads to superior performance. Our forthcoming research will expand upon the current model, with the objective to make available a robust and efficient tool for real-world financial decision making.

Acknowledgments We wish to acknowledge the valuable contributions of Barış C. Şal and Frank J. Kampas to the Excel-LGO software link development project. We also thank them for technical comments received during our numerical tests related to the CTM presented here.

References

- Bodie, Z., Kane, A., Marcus, A.J.: Investments, 2nd edn. Richard D. Irwin, Homewood (1993)
- Çağlayan, M.O., Giacomelli, D.: Finding lower-risk FX carry trades. JPMorgan Foreign Exchange Research Note, pp. 1–2 (2003)

- Çağlayan, M.O., Giacomelli, D.: Trading JPMorgan's carry-to-risk currency model in G-10 and emerging markets. JPMorgan Foreign Exchange Research Note, pp. 1–16 (2005)
- Castillo, I., Kampas, F.J., Pintér, J.D.: Solving circle packing problems by global optimization: numerical results and industrial applications. *Eur. J. Oper. Res.* **191**, 786–802 (2008)
- Coudert, V., Gex, M.: Does risk aversion drive financial crises? Testing the predictive power of empirical indicators. Centre d'Etudes Prospectives et d'Informations Internationales. Working paper No. 2007–02 (2007)
- Floudas, C.A., Pardalos, P.M., Adjiman, C.S., Esposito, W.R., Gümüs, Z.H., Harding, S.T., Klepeis, J.L., Meyer, C.A., Schweiger, C.A.: *Handbook of Test Problems in Local and Global Optimization*. Kluwer, Dordrecht (1999)
- Froot, K.A., O'Connell, P.G.J.: The risk tolerance of international investors. National Bureau of Economic Research. Working paper No. 10157 (2003)
- Gai, P., Vause, N.: Risk appetite: concept and measurement. Bank of England Financial Stability Review 127–136 (2004)
- Gonzalez-Hermosillo, B.: Investors' risk appetite and global financial market conditions. IMF Working paper series 08/85, 1–75 (2008)
- Goossens, P., McPhee, J., Pintér, J.D., Schmitke, C.: Driving innovation: how mathematical modeling and optimization increase efficiency and productivity in vehicle design. Technical Memorandum, Maplesoft, Waterloo (2007)
- Illnig, M., Aaron, M.: A brief survey of risk-appetite indexes. *Financ. Syst. Rev.* 37–43 (2005)
- International Monetary Fund (IMF): *Global Financial Stability Report*, Chapter 3 (2003)
- Isenor, G., Pintér, J.D., Cada, M.: A global optimization approach to laser design. *Optim. Eng.* **4**, 177–196 (2003)
- Kumar, M., Persaud, A.: Pure contagion and investors' shifting risk appetite: analytical issues and empirical evidence. *Int. Financ.* **5**, 401–436 (2002)
- Mason, T.L., Emelle, C., van Berkel, J., Bagirov, A.M., Kampas, F., Pintér, J.D.: Integrated production system optimization using the Lipschitz Global Optimizer and the Discrete Gradient Method. *J. Ind. Manag. Optim.* **3**(2), 257–277 (2007)
- Normand, J., Çağlayan, M.O., Ko, D., Panigirtzoglou, N., Shen, L.: JPMorgan's FX and commodity barometer. JPMorgan Investment Strategies (10), 1–32 (2004)
- Pantoleontos, G., Basinas, P., Skodras, G., Grammelis, P., Pintér, J.D., Topis, S., Sakellariopoulos, G.P.: A global optimization study on the devolatilisation kinetics of coal, biomass and waste fuels. *Fuel Process. Tech.* **90**, 762–769 (2009)
- Pintér, J.D.: LGO—a program system for continuous and Lipschitz optimization. In: Bomze, I.M., Csendes, T., Horst, R., Pardalos, P.M. (eds.) *Developments in Global Optimization*, pp. 183–197. Kluwer, Dordrecht (1997)
- Pintér, J.D.: *Global optimization: software, test problems, and applications*. In: Pardalos, Romeijn (eds.) *Handbook of Global Optimization*, vol. 2, pp. 515–569. Kluwer, Dordrecht (2002)
- Pintér, J.D.: Nonlinear optimization with GAMS/LGO. *J. Glob. Optim.* **38**, 79–101 (2007)
- Pintér, J.D.: Software development for global optimization. In: Pardalos, P.M., Coleman, T.F. (eds.) *Global Optimization: Methods and Applications*. Fields Institute Communications, vol. 55, pp. 183–204. The American Mathematical Society, Providence (2009)
- Pintér, J.D.: LGO—A Model Development and Solver System for Nonlinear (Global and Local) Optimization. User's Guide. (Current version) Distributed by Pintér Consulting Services, Canada. For professional information summary, see <http://www.pinterconsulting.com> (2010a)
- Pintér, J.D.: Spreadsheet-Based Modeling and Nonlinear Optimization with Excel-LGO—User's Guide. With technical contributions. Şal, B.C., Kampas, F.J. (eds.) Distributed by Pintér Consulting Services, Canada. <http://www.pinterconsulting.com> (2010b)
- Sharpe, W.F.: Mutual fund performance. *J. Bus.* **39**, 119–138 (1966)
- Sharpe, W.F.: Adjusting for risk in portfolio performance measurement. *J. Portfolio Manag.* 29–34 (1975)
- Sharpe, W.F.: Asset allocation: management style and performance measurement, *J. Portfolio Manag.* 7–19 (1992)
- Tarashev, N., Tsatsaronis, K., Karampatos, D.: Investors' attitude towards risk: what can we learn from options? *BIS Q. Rev.* 57–65 (2003)
- Tervo, J., Kolmonen, P., Lyyra-Laitinen, T., Pintér, J.D., Lahtinen, T.: An optimization based approach to the multiple static delivery technique in radiation therapy. *Ann. Oper. Res.* **119**, 205–227 (2003)
- Treynor, J.L., Black, F.: How to use security analysis to improve portfolio selection? *J. Bus.* 66–85 (1973)
- Zukerman, M., Mammadov, M., Tan, L., Ouyevsi, I., Andrew, L.: To be fair or efficient or a bit of both. *Comput. Oper. Res.* **35**, 3787–3806 (2008)