

Are foreign currency markets interdependent? Evidence from data mining technologies

A. G. Malliaris*

Mary Malliaris**

Fecha de recepción: 22 de junio de 2011

Fecha de aceptación: 28 de noviembre de 2011

* Loyola University Chicago,
tmallia@luc.edu

** Loyola University Chicago,
mmallia@luc.edu

RESUMEN

En este trabajo se utilizan dos metodologías de minería de datos: Clasificación y Árboles de Regresión (C&RT) y la Regla Generalizada de Inducción (GRI) para descubrir patrones subyacentes en los precios en efectivo diarios de ocho mercados de divisas. Se utiliza información del 2000 al 2009, sin considerar el último año, por que éste se usa para probar la robustez de las reglas encontradas en los nueve años previos. Se comparan los resultados de ambas metodologías. Se discuten varias de las reglas con buen desempeño, aplicadas tanto en los primeros años como en el año de prueba, como evidencia empírica de la interdependencia entre los mercados de divisas. Las reglas mecánicas utilizadas en este artículo pueden usarse de manera complementaria con otros tipos de modelos financieros aplicados a los mercados de divisas.

Palabras clave: minería de datos, mercados de divisas, árboles de clasificación y regresión, regla de inducción generalizada.

Clasificación JEL: C02, C40, C60.

ABSTRACT

This study uses two data mining methodologies: Classification and Regression Trees (C&RT) and Generalized Rule Induction (GRI) to uncover patterns among daily cash closing prices of eight currency markets. Data from 2000 through 2009 is used, with the last year held out to test the robustness of the rules found in the previous nine years. Results from the two methodologies are contrasted. A number of rules which perform well in both the training and testing years are discussed as empirical evidence of interdependence among foreign currency markets. The mechanical rules identified in this paper can usefully supplement other types of financial modeling of foreign currencies.

Key words: *Data mining, foreign currency markets, Classification and Regression Trees, Generalized Rule Induction.*

JEL Classification: C02, C40, C60.

1. Introduction

Currency trading predates both bond and stock trading as a financial innovation. However, the past 50 years of globalization have seen a marked increase in the amount of currency trading. Today, the daily volume of currency transactions in currency futures, forwards, swaps and options dominates all other types of derivatives trading. This volume is driven by globalization that includes both trade and foreign direct investments, by portfolio diversification and by hedging and speculation, among other factors. For a discussion of financial globalization in the past few years, see Lane and Milesi-Ferretti (2001, 2006), Devereux and Sutherland (2007), and Campbell, Medeiros and Viceira (2007).

Financial globalization is a process driven both by global trade and global finance. Global trade existed throughout history, but has been emphasized since the second world war. The primary reason for such an emphasis was the political motivation to employ global trade as an engine for global economic growth. Countries were encouraged to promote their comparative advantage and produce goods and services that could be traded for products that they were not able to produce as efficiently. This process encouraged global specialization and strengthened the interdependence of nations. Simultaneously, capital moved to finance investments in countries that exhibited such comparative advantages and as a result, global banking was accelerated. Global trade and global finance generated a large volume of currency trading with national central banks coordinating their policies to reduce potential risks of currency fluctuations. Furthermore, sovereign nations extended their borrowing from domestic markets to global markets and thus individuals, firms, and sovereign governments all participated in the supply and demand of global currencies. Thus, both global trade and finance intensified the interrelationships among foreign currencies. These fluctuations need to be explained both in the short run and in the long run.

In contrast to econometric approaches, a data-driven modeling approach used in data mining makes no prior assumptions about data distributions or types of relationships. There are no parameters to be estimated nor is there an

assumed model form. Instead, this non-parametric approach searches a large data set to see whether any patterns are exhibited in that set. If the patterns found meet certain minimum requirements, then the pattern is recorded for further inspection. The usefulness of the methodology is judged by looking at new data sets to see whether or not these patterns also occur here in about the same percentage. If so, we say that the data mining model is robust and has found a pattern that holds over time. This study uses two data mining methodologies: a decision tree methodology named Classification and Regression Trees (C&RT), and an association analysis methodology called Generalized Rule Induction (GRI) to uncover patterns among daily cash closing prices of eight currency markets. Data from 2000 through 2009 is used, with the last year held out to test the robustness of the rules found in the previous nine years. Results from the two methodologies are contrasted. A number of rules which perform well in both the training and testing years are discussed as empirical evidence of interdependence among foreign currency markets.

In section 2 of this paper we give a review of various dimensions of currency interdependencies and formulate our fundamental hypothesis. Then in section 3 we describe the novel methodological approach used in this paper to search for interdependencies, discuss the data used in section 4 and the results obtained that confirm the hypothesis stated in section 5. Conclusions are presented in the last section. In particular, we find that to optimize correctly directional movements in currency market prediction, one would follow association analysis rules for the European markets, and decision tree rules for the American ones. For example, GRI indicates that when Australia and Japan move in the same direction on one day, then the European markets also move in that direction. These rules affect the euro, the British Pound and the Swiss franc and are correct between 76 and 87 percent of the time. For the American markets, specific mixed combinations of movement in Asia and Europe are precursors of movement in the Americas. For example, the C&RT decision tree indicates that when Australia and the British Pound are both down, but the Yen is up, then the Mexican Peso is down over 83% of the time.

2. Literature review and hypothesis

Interrelationships between and among foreign currencies have been investigated along four main dimensions. The first dimension involves the covered interest rate parity. This approach argues that a trader has at any moment two

alternative investment strategies involving two foreign currencies. For clarity of argument suppose that we consider the euro and the U.S. dollar. One strategy is to invest for one year, say \$1, in the U.S. and earn the U.S. interest rate, or to convert the \$1 into Euros using today's exchange rate and invest the exchanged Euros at the European Central Bank interest rate and concurrently buy a one year euro forward euro contract so the investor could convert at the end of the year the invested euro and its earned interest into U.S. dollars. In equilibrium these two strategies should yield the same return. Numerous papers test and extend this theory. For example, Coffey, Hrung and Sarkar (2009) show that the theory works well during certain times but not at other times. They document that during the early months of the global financial crisis of 2007, the interest rate parity did not hold and interrelationships among global currencies were interrupted. The interest rate parity is thus used to explain interrelationships among foreign currencies in the short run. This theory is discussed in detail in most international finance and international economics textbooks such as Krugman and Obstfeld (2009) and recently expounded in Lee and Malliaris (2011). This theory associates changes in currencies to changes in the corresponding short-term interest rates of the national central banks and sets the foundations for foreign currency interdependencies attributed to interest rates. Naturally, economists are curious to know if there are additional reasons that may explain currency interdependencies.

The second dimension attributes currency interrelationships to macroeconomic variables such as prices, GDP. Schneller and Vanstone (2010) analyze how the release of various macroeconomic news affects the exchange rate behavior. Interestingly enough they perform numerous tests and conclude that "no exploitable trading patterns were found".

The third dimension of currency interrelationships covers the much broader area of portfolio management. Within this category, aspects of portfolio diversification, hedging and speculation are investigated. Kroencke, Schindler and Schrimpf (2011) offer a comprehensive analysis of international diversification benefits that result from various foreign exchange investment styles. They find that there are significant diversification benefits that can be obtained from foreign currencies. Campbell, Medeiros and Viceira (2007) and more recently Schmittmann (2010) examine the benefits from hedging the currency exposure of international investments in both single and multi-country portfolios.

The last approach involves the extension of behavioral finance to foreign currencies. For example Menkhoff, *et al.* (2011) apply currency momentum strategies to foreign currencies and find significant cross-sectional spreads between past winner and loser currencies.

In the long run, however, currency interrelationships are influenced by each countries' economic fundamentals. These fundamentals include the growth of domestic GDP, rate of inflation, technological advances, labor resources, institutional and political frameworks. Global trade and finance, with their impact on foreign currency markets, have been researched extensively using a variety of methodologies. For example, Elyasiani and Kocagil (2001) study the interdependence and dynamics in the currency futures markets. A different approach to studying the behavior of currency markets is used in Elyasiani, Kocagil, and Mansur (2007) who employ a generalized variance decomposition analysis. A large segment of currency markets research utilizes time series methodologies such as Orlov (2009). In addition, authors such as Nikkin, Sahlstrom, and Vahamaa (2006) employ derivative markets methodologies to explain links among major currencies.

Considering the existing literature that searches for interdependencies among foreign currencies using primarily time series methods and financial reasoning we hypothesize that foreign currency data contain such relationships and employ data mining methodologies described in the next section to discover exact interrelationships.

3. Data mining methodologies

Changes in information and communication technologies have accelerated the processing and transmission of data and ideas to a level far beyond our capabilities of a decade or two ago. This ability to collect considerable amounts of data quickly and easily has led to the rise of a set of investigative techniques designed to analyze large, often sparse, data sets efficiently. These collective techniques are known as data mining. Data mining is a set of exploratory methodologies designed to uncover patterns in large data sets. Though a researcher may have some question driving the search for patterns, the process does not make any assumptions about the distribution of the data, nor are hypotheses stated and tested. Rather, it looks at a large body of data and simply attempts to find patterns. The most interesting and useful patterns are those that remain stable on new data from later sets. Thus one common method for testing any data mining methodology is to feed a new data set, not used for

building the model, through the trained model to test whether the discovered patterns are still valid.

In this study, we begin with data from eight currency markets that reflect the price of 1 US dollar in each of the currencies at the closing time of the local market. The question driving the data mining is a broad one of “does movement in one market influence movement in another market?” We apply two data mining techniques to uncover common patterns in various markets. These patterns are tested against a validation set comprised of a year’s worth of data following the time of the training set.

This study uses two data mining methodologies: Classification and Regression Trees (C&RT) and Generalized Rule Induction (GRI). Unlike many methods from statistics, C&RT did not exist before machine learning methods were available. C&RT is a decision tree methodology that uses recursive partitioning to divide the training data set into groups with the same value of a target variable (Loh, 2011). It begins with all the target values in one large bucket, called the root of the decision tree. The methodology then looks at each input variable one at a time. A measure of impurity is calculated for each variable. Impurity refers to the amount of non-equal target values that would result in the two new child buckets if we were to divide the original target set on the basis of an input variable’s values. The input variable that yields the minimum amount of impurity is used to split the original data set. The data set is said to branch on this variable and the branching results in two buckets. Now, the same technique is applied to each of these two child buckets and another split may result in a further decrease of impurity. The split at each child bucket does not have to be on the same variable. This process continues until no further possible split on an input variable value will decrease the impurity. In C&RT, all splits are binary (unlike some other decision tree methodologies). A bucket, or node, is labeled as “pure” if it contains only one value of the target variable. The overall goal of any decision tree methodology is to minimize impurity and maximize purity. The final result of C&RT is often represented graphically in a tree structure. This visual representation is easy to read. One simply follows the path from the root of the tree to a branches’ end, using the decision made at each split. New data can be fed through a trained tree by following the set of step-wise binary splits to a final leaf of the tree. The forecasted value will be the predominating target value in this leaf.

There are simpler data mining techniques, however, which can be used to forecast in more limited areas with non-numeric data. One of these tech-

niques is association analysis. Association analysis, also called market basket analysis, was designed to answer the question “what occurs together?” and has been used by marketing specialists to discover which products are likely to be purchased by the same customer. It was developed as a technique that would discover interesting rules on large sparse data sets. The data typically used for association analysis is non-numeric and dichotomous. It does not focus on the amount of product purchases, only whether or not some of that product appears in the shopping basket.

GRI is an association analysis technique that was created by Smyth and Goodman (1992) as an algorithm that could be used for the induction of rules from a large set of examples. Rather than using an expert (common, but labor intensive, at that time) to obtain rules, they wanted to create an algorithm that could automatically acquire rules from data, where that data existed. They also wanted their technique to generate rules relating not only two columns of data, but possibly multiple columns. Their focus was thus on finding a set of rules within that data. These rules would be of the form “If A then B” where A could be multiple products. Unlike decision tree analysis, in which any row of data fits only one path in a tree, GRI may generate many accurate rules on a single row. For example, “if bread then milk”, if “diapers then milk”, and “if strawberries and champagne then chocolate” may all apply to a single basket with many items. Each rule is accompanied by two measures, support and confidence. Support is the percentage of times the “if” part of the rule occurs in the data set. When the “if” part does occur, confidence is the percentage of times that the “then” part of the rule is also true. Acceptable levels of support and confidence are set by the researcher before running the model. Lower values will allow more rules to be generated. After a set of GRI rules is created, one then filters the rules for a specific target variable (say, milk), and uses the most accurate rules of the set. For a more detailed discussion of association analysis techniques, see Hand *et al* (2001) or Berry and Linoff (2004).


In today’s digital environment, collecting clean data is much easier than when these techniques were first put into practice. Thus the number of rule induction methods and problems for their application has increased. These methods continue to be popular for approaching problems in finance. See, for example, Batyrshin *et al* (2005, 2007) for applications in time series databases, Albanis and Batchelor (2007) for an example with stock selection, Wang *et al* (2009) applies rule induction to forecasting time series, Tseng (2007) discusses co-movement in international stocks and Bossomaier *et al.* (2010) details a simulation of trust in wealth management.

4. Data set

We began with daily cash closing prices for the Australian Dollar, British Pound, Brazilian Real, Canadian Dollar, Euro, Japanese Yen, Mexican Peso, and Swiss Franc with respect to the US Dollar. That is, the data reflects the amount of each foreign currency that could be purchased with 1 US dollar that day. Though the foreign exchange market is considered to be a 24-hour market, closing prices can be quoted for individual markets in pairs. Thus, when the market in Tokyo closes, the value of the Yen to the dollar can be established for that day. All values in this data set are in these units of the foreign currency per US dollar. The data sample covers the time period from January 2000 through July 2009 and was downloaded from Bloomberg.

The relative movement in these currencies can be seen in Figure 1. In order to view them all in a similar scale, the Mexican Peso has been multiplied by 10 and the Japanese Yen by 100 for the graph. There are a total of 2,491 observations for prices for each of the eight daily closing prices. These prices were split into two disjoint sets for training and validation. Data from January 1 2000 to June 30 2008 was used as the training set (2215 rows), with the remainder, from July 1 2008 to July 21 2009, used as the validation set (276 rows) to judge the effectiveness of the data mining methods.

Figure 1. Currency prices in units of 1 US dollar, peso & yen scaled.



To study the simultaneous directional market movements with non-numeric data, all data was transformed into “Up” or “Down” by comparing the value of the currency at time t with its value on the previous day. The currency markets in this set group into three areas: Asia (Aus. Dollar and Yen), Europe (Pound, Franc, and Euro), and Americas (Real, Can. Dollar, Peso). Data from Asia was used to forecast movement in Europe. Data from Asia and Europe was used to forecast movement in the Americas.

The decision tree tries to develop paths to branches with a pure value of Up or Down on a single target. One decision tree was created for each target in the data set. The rules generated by this methodology give a path of Up and/or Down movements from various currencies to the target.

For the association analysis technique, using a slight twist on the question of what is bought together, we can ask which major currency directional movements occur together, relative to the dollar, on the same day. That is, when the price of 1 US dollar in some foreign currency goes down or up, does the price of this dollar in other currencies move up or down on the same day? This question does not look at the amount of change, only the directions in two or more markets. The GRI methodology creates separate rules for up movement and down movement. It can handle multiple targets at the same time. The GRI methodology was run twice. The first time, the Asian markets were possible antecedents, and the European markets possible consequents. In the second run, the Asian and European markets were possible antecedents and the Americas were possible consequents.

5. Results

Each model developed a large number of rules. In this section, we display one Up rule and one Down rule for each of the target markets. Rules selected were those that did well not only on the training set, but also on the validation set. Accuracy on the validation set is an indication that the pattern will be applicable on future data.

Table 1 shows results from the C&RT decision tree methodology and the best rules for each target market. The first column of this table lists the target variable. Column two displays the inputs that were used by C&RT in the rule, in the order that the methodology selected them. The variables selected for a path to a target are those that the C&RT methodology found useful in splitting the data into more pure groups. Variables not listed in a path were deemed as

not useful by the C&RT methodology for the specific path to the target value. For each target, there may have been several paths to a given value of the target. For example, an Up value of a target may have had a half-dozen different paths that ended with Up. The single path displayed in the table was the one with the greatest stability from training to validation set. For the training and validation sets, columns three and five display the percentage of rows in which the inputs path occurred in the data set; columns four and six display the percentage of time the target value matched the one listed in the row where the input path is as stated. The results are listed with the European targets first, followed by the American targets.

Table 1. Selected rules generated by C&RT.

Best Decision Tree Rules		Training set		Validation set	
Target	Inputs	%Rows	%Correct	%Rows	%Correct
Euro = Up	Australia = Up	45.92	69.8	47.13	78.05
Euro = Down	Australia = Down	54.08	70.5	52.87	70.29
Britain = Up	Australia = Up	45.92	66.11	47.13	76.42
Britain = Down	Australia = Down	54.08	67.3	52.87	66.67
Swiss = Up	Japan = Up	50.67	66.64	46.74	63.11
Swiss = Down	Japan = Down and Australia = Down	32.42	80.22	21.84	80.70
Brazil = Up	Australia = Up and Japan = Down	16.91	58.62	31.42	74.39
Brazil = Down	Australia = Down and Euro = Down and Swiss = Down and Britain = Down	27.8	61.45	23.75	80.64
Mexico = Up	Australia = Up and Swiss = Down and Euro = Up	3.77	64.28	5.36	71.43
Mexico = Down	Australia = Down and Japan = Up and Britain = Down	11.53	63.03	18.77	83.67
Canada = Up	Australia = Up	45.92	65.72	47.13	78.05
Canada = Down	Australia = Down and Euro = Down	38.12	72.24	37.16	83.51

For the six European target rules, only the Swiss Down rule used more than one input. The rules for the Euro and the British Pound use same direction information from the Australian Dollar. The Swiss Franc is more sensitive to same-direction movement in the Japanese Yen, followed by the Australian Dollar only when the movement is down. The percentage of the data sets with inputs matching those shown in each row of the table is similar from training to validation set, with the exception of days when both the Australian Dollar and the Japanese Yen moved down. There was a drop of over 10% in the occurrence of this input pattern. However, of more interest is whether or not the target value is correct when the input path is as stated. The European targets show a nice robustness in these numbers. Two of the up targets, the Euro and the British Pound, show even stronger performance on the validation set, while the others remain about the same.

For the targets in the Americas, we see that the input paths are generally more complex, and thus occur less frequently. However, all paths remain strong in performance on the validation set. In fact, the percentage of time the target value is correct has increased in every validation set example.

The results from the association analysis methodology, GRI, are shown in Table 2. The order and interpretation of the columns is the same as in Table 1. Of interest is the fact that, for most markets, the best rules generated by GRI do not exactly match those generated by C&RT. In GRI, there is no ordering in a path as there is in C&RT, but the combination of inputs matches those selected by C&RT in only three rules: Swiss = Down, Canada = Up, and Canada = Down. The other nine rules are either more complex, or select different market movements as inputs. For all the European markets, directional movement depends on both the Australian Dollar and the Japanese Yen. All market movement is identical to that of Australia and Japan when both move the same way. Because the input combination is more restrictive, we see that the percentage of rows to which the rules apply is smaller. The amount of time the rules are correct is essentially the same for the training and validation sets for down movements of the British Pound, but show an increase in the validation set over the training set for all other European market movements.

In the markets of the Americas, only up movements in Mexico ignore what happened in the Australian Dollar that day. All other American market movements require the Australian Dollar to move the same direction. However, this movement must be in combination with other markets. Up movements in the Brazilian Real and the Mexican Peso look for mixed directions in other mar-

Table 2. Selected rules generated by GRI.

Best Generalized Rule Induction Rules		Training Set		Validation Set	
Target	Inputs	%Row	%Correct	%Rows	%Correct
Euro = Up	Australia = Up and Japan = Up	28.94	78	17.39	85.42
Euro = Down	Australia = Down and Japan = Down	32.42	81.45	21.84	85.71
Britain = Up	Australia = Up and Japan = Up	28.94	72.07	17.39	77.08
Britain = Down	Australia = Down and Japan = Down	32.42	76.71	21.84	76.19
Swiss = Up	Australia = Up and Japan = Up	28.94	76.44	17.39	87.50
Swiss = Down	Australia= Down and Japan = Down	32.42	80.22	21.84	80.70
Brazil = Up	Australia = Up and Swiss = Down and Japan = Down	8.08	64.8	11.23	70.97
Brazil = Down	Australia= Down and Euro= Down	38.06	61.09	37.68	73.08
Mexico = Up	Euro = Up and Britain = Up and Swiss = Down	3.02	64.18	5.8	87.50
Mexico = Down	Australia = Down and Euro = Down and Swiss = Up	4.24	64.89	6.16	76.47
Canada = Up	Australia = Up	45.92	65.72	47.13	78.05
Canada = Down	Australia = Down and Euro = Down	38.12	72.24	37.16	83.51

kets, especially a Down movement in the Swiss Franc. Any of these market movements that take into account the Swiss Franc moves in an opposite way from the Franc that day. We see that some of these rules have a small percent of data rows to which they apply, most notably those of the Mexican Peso. However, when the rule does apply, the validation set shows a high percentage of correctness for the rule.

Table 3 shows, for each of the markets, the model that had the highest percent of correctness on the training set and on the validation set. If we were basing our model choice on only the training set, we would use the GRI model in all but two cases. These two, the Real down and the Peso up, each have a complex path to the target value. However, in data mining, decisions are most often based on results from the validation set, an entirely new set of data not available to the methodology during training. In the validation set, we see that, for the European group of targets, the GRI correctness is higher than or equal to the decision tree values in each rule. For the Americas, however, the C&RT validation set correctness is higher than or equal to that of the GRI in

5 out of 6 rules. In only one rule, for the Mexican Peso moving up, is GRI the best rule to follow. Thus, to optimize being correct on directional movements, one would follow the GRI rules for the European markets, and C&RT for the American ones.

Table 3. Highest training and validation sets correctness per market.

Target	Best training model	% Correct	Best validation model	% Correct
Euro = Up	GRI	78.00	GRI	85.42
Euro = Down	GRI	81.45	GRI	85.71
Britain = Up	GRI	72.07	GRI	77.08
Britain = Down	GRI	76.71	GRI	76.19
Swiss = Up	GRI	76.44	GRI	87.50
Swiss = Down	Equal	80.22	Equal	80.70
Brazil = Up	GRI	64.80	C&RT	74.39
Brazil = Down	C&RT	61.45	C&RT	80.64
Mexico = Up	C&RT	64.28	GRI	87.50
Mexico = Down	GRI	64.89	C&RT	83.67
Canada = Up	Equal	65.72	Equal	78.05
Canada = Down	Equal	72.24	Equal	83.51

6. Conclusions

This paper illustrates the application of two machine learning techniques, decision trees and association analysis, to ten years of currency movement data. These two exploratory methods yield encouraging results for a trader wishing to diversify a basket of currency products. Both methods are trained on nine years of data and validated on the tenth year. The rules resulting from the models show robustness in their application to the validation set, an indication that the trading rules uncovered would be stable over time.

The robustness of the rules obtained from techniques that are unrelated to financial reasoning has useful implications for traders, portfolio managers and hedgers. Currency traders have been very active during the past ten years in carry trades (Jain (2010), Jorda and Taylor (2009)). These trades involve

selecting an appreciating currency that is bought and a depreciating currency that is sold to finance the long position. Such positions are influenced also by domestic interest rates and growth prospect in the two countries. These trades have been successful for several years before the global financial crisis that reversed positive and negative momentum of several currencies. The mechanical rules identified in this paper can usefully supplement financial modeling of carry trades.

Well diversified portfolios often invest a small portion of their capital in foreign currencies. The selection of currencies and the amounts invested in each are guided by modern portfolio selection techniques. These techniques can usefully be supplemented by the results of the methods applied in this paper.

Finally, global trade and its financing involve several risks with currency risks being on the top of the list. Hedging such risks can be performed using the traditional currency hedging techniques of computing optimal hedge ratios for currencies trading in futures markets. For currencies with lower trading volume and no futures contracts or options on these currencies, cross hedging techniques are used that can benefit from the methodologies applied in this paper.

References

- Albanis, G. and Batchelor, R. (2007). "Combining Heterogeneous Classifiers For Stock Selection". *Intelligent Systems in Accounting, Finance & Management*, 15: 1–21. doi: 10.1002/isaf.282
- Batyrshin, I., Herrera-Avelar, R., Sheremetov, L, and Panova, A. (2005). "Association Networks In Time Series Data Mining, Proceedings of the Fuzzy Information Processing Society, pags 754 – 759.
- , (2007). Moving Approximation Transform and Local Trend Associations in Time Series Data Bases. Perception-based Data Mining and Decision Making in Economics and Finance, Studies in Computational Intelligence, Springer, Berlin / Heidelberg.
- Berry, M. and Linoff, G. (2004). *Data Mining Techniques*, Second Edition, Wiley Publishing Inc.
- Bossomaier, T., Standish, R., Harre, M. (2010). Simulation of Trust in Client-Wealth Management Adviser Relationships, *International Journal of Simulation and Process Modelling*, 6(1), 40 – 49.

- Campbell, J.Y., Medeiros, K.S., and Viceira, L.M. (2007). Global Currency Hedging, NBER Working Paper, No. W13088.
- Coffey, Niall, Hrung, Warren B. and Sarkar, Asani, (October 29, 2009). Capital Constraints, Counterparty Risk, and Deviations from Covered Interest Rate Parity. FRB of New York Staff Report No. 393. Available at SSRN: <http://ssrn.com/abstract=1473377>
- Devereux, M.B., and Sutherland, A. (2007). "Financial Globalization and Monetary Policy", *IMF Working Paper*.
- Elyasiani, Elyas, and Kocagil, Ahmet (2001). "Interdependence and Dynamics in Currency futures Markets: A Multivariate Analysis of Intraday Data". *Journal of Banking & Finance*, 25: 1161 – 1186.
- Elyasiani, Elyas, Kocagil, Ahmet and Mansur, iqbal (2007). "Information Transmission and Spillover in Currency Markets: A Generalized Variance Decomposition Analysis". *The Quarterly Review of Economics and Finance*, 47: 312-330.
- Gadiraju, Pavan, (July 19, 2009). A Statistical Arbitrage FX Trading System Based on Short Term FX Volatility Swings Forecasting with Institutional Data on JPY Based Investment Flows Into US Markets. Available at SSRN: <http://ssrn.com/abstract=1436209>
- Hand, D., Mannila, H., and Smyth, P. (2001). "Principles of Data Mining", *The MIT Press*.
- Jain, Apurv, (2010). Are Carry Trade Risks Systematic Risks Now? An Analysis of the Dynamics of Carry Trade Risks. Available at SSRN: <http://ssrn.com/abstract=1546637>
- Jordá, Óscar Alan M. Taylor, (2009). The Carry Trade And Fundamentals: Nothing To Fear But Fear Itself, NBER Working Paper No. 15518.
- Kroencke, Tim-Alexander, Schindler, Felix and Schrimpf, Andreas, (March 20, 2011). International Diversification Benefits with Foreign Exchange Investment Styles CREATES Research Paper No. 2011-10. Available at SSRN: <http://ssrn.com/abstract=1790764>
- Krugman, Paul R., and Maurice Obstfeld, (2009). *International Economics Theory & Policy*, 8th edition. Boston, MA: Addison Wesley.
- Lane, P., and Milesi-Ferretti, G.M. (2001). "The External Wealth of Nations: Measures of Foreign Assets and Liabilities for Industrial and Developing Countries", *Journal of International Economics*, 55, 263-94.
- , (2006). "The External Wealth of Nations Mark II", *IMF Working Paper*, No 06-69.

- Lee, Suk Hun and Malliaris, A. G. (2011). *Currency Markets and International Interest Rate Parity in Survey of International Finance*, Kent Baker and Leigh Riddick, eds., Oxford University Press,. Available at SSRN: <http://ssrn.com/abstract=1747111>
- Loh, W.-Y. (2011). Classification and regression trees. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1: 14–23. doi: 10.1002/widm.8
- Menkhoff, Lukas, Sarno, Lucio, Schmeling, Maik and Schrimpf, Andreas, (May 12, 2011). Currency Momentum Strategies. Available at SSRN: <http://ssrn.com/abstract=1809776>
- Nikkin, Jussi, Sahlstrom, Petri, and Vahamaa, Sami (2006). "Implied Volatility Linkages Among Major European Currencies". *Journal of International Financial Markets, Institutions and Money*, 16:87-103.
- Orlov, Alexei (2009). "A Cospectral Analysis of Exchange Rate Comovements During Asian Financial Crisis". *Journal of International Financial Market, Institutions & Money*, 19: 742-758.
- Schmittmann, Jochen, (June 2010). "Currency Hedging for International Portfolios" *IMF Working Papers*, Vol., pp. 1-44, 2010. Available at SSRN: <http://ssrn.com/abstract=1641006>
- Schneller, Warwick and Vanstone, Bruce James, (August 22, 2010). Predictable Responses in Currency Markets to Macroeconomic News: A Trading System Approach. 23rd Australasian Finance and Banking Conference 2010 Paper. Available at SSRN: <http://ssrn.com/abstract=1663429>
- Smyth P, Goodman RM. (1992). Information Theoretic Approach to Rule Induction From Databases. *IEEE Transactions on Knowledge and Data Engineering*. 4(4):301–316. doi: 10.1109/69.149926
- Tseng, C. (2007). "Data Driven Modeling Of Co-Movement Among International Stock Market", *Journal of Modeling in Management*, 2007, 2(3), 195 – 207.
- Wang, X., Smith-Miles, K., Hyndman, R. (2009). Rule Induction For Forecasting Method Selection: Meta-Learning The Characteristics Of Univariate Time Series, *Neurocomputing*, 72(10-12), 2581-2594.

