Detecting statistical arbitrage opportunities using a combined neural network - GARCH model*

Nikos S. Thomaidis[1**] and Nick Kondakis[1,2]

[1] Dept. of Financial Engineering & Management, University of the Aegean, 8 Michalon Str., GR-821 00, Chios, Greece, tel: +30-2271-0-35454 (35483), fax: +30-2271-0-35499, nthomaid@fme.aegean.gr, URL: http://labs.fme.aegean.gr/decision/

[2] Kepler Asset Management, LLC, 100 North Country Road, Setauket, NY 11733, nick@kondakis.com

Abstract. This paper proposes a hybrid computational intelligent system for the detection of statistical arbitrage opportunities in pairs of assets. The proposed methodology combines nonlinear neural network autoregressive models with GARCH parametrizations of volatility for describing the dynamics of the correction of relative mispricings. First results from this approach seem encouraging; further experimentation on the optimal sampling frequency, the forecasting horizon and the points of entry and exit is conducted, in order to improve the economic value when transaction costs are taken into account.

Keywords: statistical arbitrage, intelligent trading systems, neural networks, GARCH models
JEL classification C14, C22, G11
First version: April 4, 2006; **Current version**: Jan 23, 2012

1 Introduction

In the last few years, a substantial amount of computational intelligent methodologies have been applied to the development of financial forecasting models that attempt to exploit the dynamics of financial markets. A great majority of intelligent approaches employ a network learning technique, such as feedforward, radial basis function or recurrent NN [13, 17], although certain paradigms such as genetically-evolved regression models [5, 8, 11, 14] or inductive fuzzy inference systems [9] are also encountered in the literature. Forecasting experience has shown that predictability in data can increase if modelling is directed to a combination of asset prices rather than (raw) individual time series. Combinations can be seen as a means of improving the signal-to-noise ratio and hence enhancing the predictable component in the data [3].

Statistical arbitrage is an attempt to profit from pricing discrepancies that appear in a group of assets. The detection of mispricings is based upon the identification of a linear combination of assets, or else a "*synthetic*" asset, whose time series is *mean-reverting* and has finite variability. For example, given a set of

* This is an extended working-paper version of the article "An Intelligent Statistical Arbitrage Trading System" appeared in Grigoris Antoniou et al. (ed), ADVANCES IN ARTIFICIAL INTELLIGENCE, Lecture Notes in Computer Science, Volume 3955/2006, pp. 596-599. The original publication of the afore-mentioned paper is available at http://www.springerlink.com/content/l1v8v34t03l22717/.

** This research was financially supported by the Public Benefit Foundation "Alexander S. Onassis" (www.onassis.gr), under the 2003-2004 Scholarships Programme, and by a grant from "Empeirikion" Foundation (www.empirikion.gr).

assets S_1,...,S_k, a statistical mispricing can be considered as a linear combination $\boldsymbol{\beta} = (b_1, b_2, ..., b_k)$ such that the portfolio value

$$\Pi_t \equiv b_1 S_{1t} + b_2 S_{2t} + \cdots + b_k S_{kt}$$

is mean-reverting with zero mean and possibly time-dependent (finite) volatility σ_t^2 (S_{it} is the value of asset i at time t). The vector $\boldsymbol{\beta}$ represents actual quantities of each asset (i.e. number of stocks) to be held in the trading portfolio. The requirement of mean-reversion is to ensure that mispricings eventually "die out" and do not grow indefinitely. If mispricings permanently disrupted the long-run relationship between the assets, it would be impossible to control the risk exposure of the trading portfolio.

The standard approach to identify statistical mispricings is to run a regression of the values of one asset, say S_{1t}, against the others S_{2t}, ..., S_{kt} and test the residuals for mean-reversion. Several tests have been developed for this purpose in the econometric literature, the most famous of which being the Dickey-Fuller and the Phillips-Perron (see e.g. [6]). Note that the residuals of the regression model represent the mispricing at each time t of S_{1t} relative to its fundamental value spanned by the group of fundamentally-related securities $\{S_{1t}$,...,$S_{nt}\}$. The next step is to create a model that describes the *dynamics* of mispricings, i.e. how errors of different magnitude and sign (positive/negative) are corrected over time. Model forecasts are then incorporated into a dynamic trading strategy. An arbitrage trading system identifies the "turning points" of the errors time-series and takes proper positions on the constituent assets when mispricings become *exceptionally* high (i.e. $\boldsymbol{\beta}$ for a positive and $-\boldsymbol{\beta}$ for a negative mispricing). An arbitrage trading strategy, as described above, is not without risk; although profitable in the long run, its instant revenue depends heavily on the speed at which market prices return to the predicted norm within a short period of time. Generally, the weaker the mean-reversion the higher the probability of observing adverse movements of the synthetic.

Several authors have suggested approaches that attempt to take advantage of price discrepancies by taking proper transformations of financial time-series; see e.g. [2, 3, 16] for stocks of FTSE 100, [4, 10] for equity index futures and [12] for exchange rates. Amongst them, [3, 4, 12] employ a neural network model to describe the dynamics of statistical mispricings. In this paper, we propose a new intelligent methodology for the identification of statistical arbitrage opportunities. Our approach deviates from the main trend in that it attempts to detect nonlinearities *both* in the mean and the volatility dynamics of the statistical mispricing. For this purpose, we use a newly proposed class of combined *neural network*-GARCH volatility models. The methodology is applied to the detection of statistical arbitrage opportunities in a pair of two stocks traded in the Mumbai Stock Market (India).

The rest of the paper is organised as follows: In section 2 we describe the application data, including intraday quotes of stock prices. Section 3 presents the methodology for detecting price discrepancies between stocks and section 4 details the NN-GARCH model used to forecast the dynamics of the statistical mispricing. In section 5 we present two arbitrage trading systems based on a high- and low-frequency predictive model. Section 6 concludes the paper and discusses directions for further research.

2 Sample data

For the application and testing of the trading strategy we chose the stocks of Infosys Technologies Ltd and Wipro Ltd, both Application Software companies from the Indian stock market. We did so for two reasons:

1. We plan to further deploy the system onto a larger set of stocks with sector neutrality so we chose two active names from the Technology/Software sector. Choosing stocks from the same industry sector usually results in better mean-reversion behavior. In addition, both companies have active ADR issues in the US which adds some interesting interactions and influences.

2. We are in the process of investigating the extension of statistical arbitrage equity strategies into developing and emerging markets. We are also interested in studying in detail the execution intricacies of various markets and thus we will be paying special attention to trading costs in follow-up work.

Both stocks trade in the National Stock Exchange of India, headquartered in Mumbai, India. The NSEI is a fully automated order-driven market. We are using tick-by-tick data, time and sales as well as best bid and offer and corresponding sizes for the historical period from February 1, 2005 until November 8, 2005. Subsequently the tick data is consolidated into 1-minute bars that include the open, high, low and close price, the total share and tick volume and the volume-weighted average price. We have appropriately adjusted the price and volume data for dividend and split actions. We plan to use the tick data information to work out our trading-cost models in subsequent studies.

3 Identifying statistical mispricings

Figure 1 shows hourly closing prices of Infosys and Winpro from February 2 to November 8[3]. As a first attempt to construct a synthetic asset, we ran a regression of Infosys against Winpro, hence forth S_2 and S_1 respectively, over the first 200 sample observations and we then used the regression coefficients to compute the statistical mispricing. The resulting series is depicted in figure 2. Observe that the estimated combination is weakly mean-reverting especially in the first 600 observations. The Phillips-Perron (PP) test statistic over the whole sample period is -2.0183, which is below (in absolute terms) the 1, 5 and 10% critical levels (-3.88, -3.36 and -3.04 respectively). Hence, the hypothesis of mean reversion cannot be accepted.

In order to control the non-stationarity of the synthetic asset, we adopt an adaptive estimation scheme in which the coefficients of the linear combination are periodically re-calculated. In particular, we define the mispricing as

$$Z_t \equiv S_{2t} - b_0^{t-1} - b_1^{t-1} S_{1t}$$

where b_0^t, b_1^t are estimated on the basis of a window of observations of length L: $\{S_{1j}, S_{2j},\, j = t-L+1, ..., t\}$. Instead of using linear regression, we adopt a slightly different procedure for calculating betas: we define b_1 as the mean price ratio between the two stocks over the specified window and we subsequently choose b_0 so as to minimize the total variation of Z_t within the window, i.e.

$$b_1^t \equiv \text{mean}\,(S_{2j}/S_{1j},\, j = t - L + 1, ..., t)$$

and

$$b_0^t \equiv \text{mean}(S_{2j},\, j = t - L + 1, ..., t) - b_0^t\, \text{mean}(S_{2j},\, j = t - L + 1, ..., t)$$

This procedure has been experimentally found to give more reasonable estimates of the synthetic vector, improving their stability over time. In figure 3 we show the synthetic time series resulting from a different choice of the window length. Observe that the more often the values of betas are updated, the stronger is the mean-reversion of the synthetic and hence the more abrupt are the corrections of mispricings. All depicting series are found mean-reverting; the PP test statistic over the entire sample of observations is -10.09 for $W = 10$, -6.49 for $W = 30$ and -4.56 for $W = 100$, which are above common critical levels. In subsequent experiments, we report results obtained for a synthetic calculated on the basis of a window of 10 observations.

[3] Prices in this diagram are adjusted for splits and dividends.

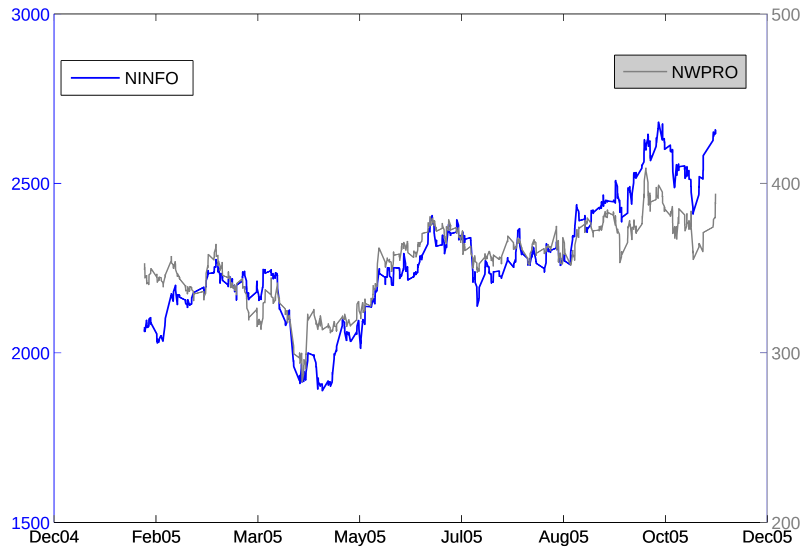

Fig. 1. Hourly data of Infosys and Winpro from 02/02/2005 to 08/11/2005.

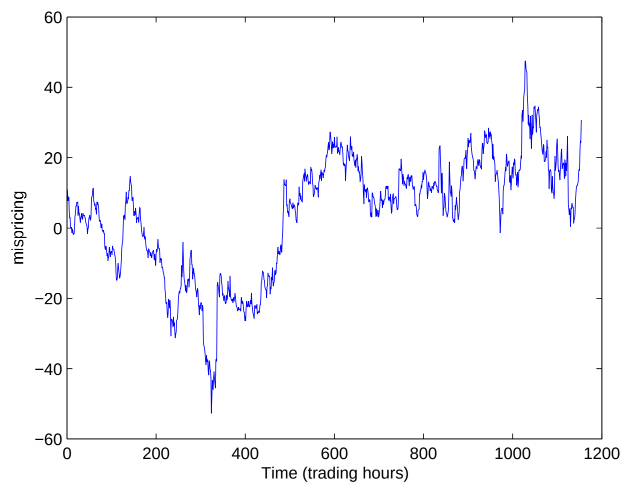

Fig. 2. The synthetic asset constructed from a static regression.

4 Modelling the dynamics of the statistical mispricings

To describe the dynamics of the statistical mispricings we use autoregressive models relating the current level of Z_t to its own history. This gives us an idea of how mispricings of different size and sign (positive/negative)

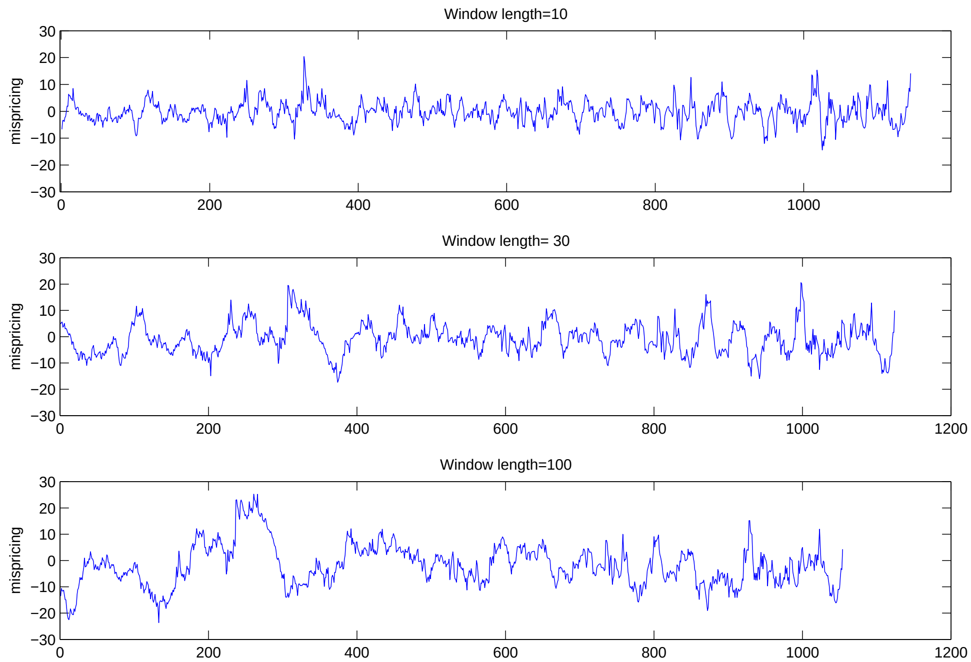

Fig. 3. The synthetic time series obtained from an adaptive estimation scheme for a window length of 10, 30 and 100, respectively.

are corrected over time. We also go one step further to model *both* the mean and the volatility structure of the statistical mispricings. This is because in high sampling frequencies (intra-day data), we find that the volatility of Z_t (i.e. the average uncertainty about the realised value) is not constant over time but strongly depends on the history of Z_t. In particular, large (positive or negative) shocks to Z_t are on average followed up by large shocks of either sign. This "clustering" of volatility, typical in most financial time series, is termed in the literature as *Autoregressive Conditional Heteroskedasticity* (ARCH [7]. Any changes in the short-term volatility level of Z_t deserve special attention from a modelling point of view, as they have important implications for the risk control of the statistical arbitrage. Until today, the most popular models for the volatility dynamics of a time series are the class of GARCH models [1]. A GARCH model effectively shows how a sudden negative or positive mispricing affects the future volatility of mipricings.

In our approach, we attempt to model both nonlinearities in the correction of mispricings as well as volatility clustering effects. For this purpose, we use a recently proposed class of joint neural network-GARCH models[15] that is intended to capture both effects. In this framework, an autoregressive model for the conditional mispricing takes the general form:

$$Z_t = \phi_0 + \boldsymbol{\phi}' \cdot \mathbf{Z}_t + f(\mathbf{Z}_t; \boldsymbol{\theta}) + \epsilon_t \tag{4.1a}$$

$$\epsilon_t | \mathcal{I}_{t-1} \sim N\left(0, \sigma_t^2\right) \tag{4.1b}$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \sigma_{t-i}^2 + \sum_{j=1}^{q} \beta_j \epsilon_{t-j}^2 \tag{4.1c}$$

where $\phi_0 \in \mathbb{R}$, $\mathbf{Z}_t$ is a vector of lags of Z_t ($\mathbf{Z}_t = (Z_{t-1}, Z_{t-2}, ..., Z_{t-n})' \in \mathbb{R}^n$), ϵ_t is the innovation term, $\boldsymbol{\phi} \in \mathbb{R}^n$ and $f(\mathbf{Z}_t; \boldsymbol{\theta})$ is a feedforward neural network with a single hidden layer and l neurons, i.e.

$$f(\mathbf{Z}_t; \theta) = \sum_{j=1}^{l} \lambda_j F(\boldsymbol{w}_j' \cdot \mathbf{Z}_t - c_j)$$

where $F(z) = 1/(1 + e^{-z})$ is the logistic function, $c_j, \lambda_j \in \mathbb{R}$ and $\boldsymbol{w}_j \in \mathbb{R}^n$. With $\mathcal{I}_{t-1}$ we denote the information available up to time t, including the history of mipricings Z_t, shocks ϵ_t and volatilities σ_t^2. Note that in the above specification, ϵ_t is assumed conditionally normally distributed with volatility σ_t^2 that depends on past σ_t^2's as well as (the magnitude of) past *unanticipated* mispricings.

In [15] we propose a complete model-building cycle for NN-GARCH models that comprises the following stages: a) *model specification* (e.g. determining the number of neurons in the hidden layer, the connections from inputs to hidden neurons, etc), b) *parameter estimation* and c) *in-* and *out-of-sample evaluation*. This is a simple-to-complicate model-building approach that starts from a linear model and gradually augments the neural network architecture, if data indicate so. The procedure is briefly described as follows:

1. Estimate a linear model with no GARCH component and choose the optimal set of input variables by means of AIC of SBIC.
2. Test the null hypothesis that the true data-generating process is a linear model against the alternative of a neural network model with a single hidden neuron. If linearity is not rejected at a given confidence level then stop. Otherwise, estimate a NN model with a single neuron and test it against a NN model with an additional neuron. Repeat the above procedure until first acceptance of the null.
3. Once the mean model (4.1a) is specified, test the null hypothesis of no GARCH effects in the volatility of the residuals of the model against the hypothesis that residuals follow a GARCH process of a given order. If null is not rejected then stop. Otherwise, jointly estimate a NN-GARCH model.

There are two important things to note about the above procedure. First, the decision of whether to add an extra neuron is not based on heuristic arguments but on formal statistical tests of "neglected nonlinearity" (see [15] for details). Second, the procedure presented above does not directly estimate a NN-GARCH model but adjusts the specification according to teh complexity of the data set (linear- or nonlinear-in-mean model, with or without a GARCH component). Hence, it aims at producing non-redundant models that are less likely to overfit the data.

5 Application: detecting statistical arbitrage in the pair of Indian stocks

Our methodology for exploiting statistical arbitrage opportunities consists of the following steps:

1. *Constructing a "synthetic asset" and testing for mean-reversion in the price dynamics.* Synthetics are calculated for various sampling frequencies.
2. *Modelling the mispricing-correction mechanism between relative prices.* For this purpose we use the general framework of NN-GARCH models (4.1).

3. *Obtaining 1- and h-step-ahead forecasts for the future value of the mispricing.* Forecasts are given in the form of a *conditional probability density* from which confidence bounds on the future value of the mispricing are derived. The estimation of an h-step-ahead conditional density is based on the simulation of 800 error scenarios. Errors are calculated as $\sigma_t u_t$, where σ_t is the volatility forecast, obtained from equation (4.1c) and u_t are sampled with replacement from the full model's *standardized* residuals (the residuals divided by the estimated sigmas). In this way, we avoid imposing restrictive assumptions on the distribution of the error-generating process.
4. *Implementing a trading system to exploit the predictable component of the mispricing dynamics.* The trading strategy is as follows: buy (long) the synthetic asset if $Z_t < \hat{Z}^{L,\alpha}_{t+h}$ and sell (short) the synthetic asset if $Z_t > \hat{Z}^{H,\alpha}_{t+h}$, where $\hat{Z}^{L,\alpha}_{t+h}$ $\hat{Z}^{H,\alpha}_{t+h}$ denote the $(1-\alpha)\%$ low and high confidence bound on the value of the mispricing h steps ahead[4]. In our approach the confidence interval is a decision variable, which has to be adjusted so that trading results are optimized.

5.1 Experiment A: a high-frequency model

In this section, we give an example of a *microscopic* model, estimated from 1-minute closing prices. The values of the synthetic from August 2 to August 19, a total of 5000 observations approximately, are used for the specification of the mispricing-correction model and the sample observations from August 29 to September 22 for out-of-sample evaluation. At this sampling frequency, we find both nonlinear corrections of mispricings and ARCH volatility effects. The specified NN-GARCH model includes lags 1-10 in the linear part, 1 hidden neuron with variable Z_{t-10} connected to it and a GARCH(1,1) equation.

We report results from a trading system that is based on 5-minute-ahead forecasts. Figure 6 shows the performance of arbitrage trading strategies by varying the confidence level. Observe that as the interval gets narrower ($1-\alpha$ is decreased) the accumulated profit becomes higher, although the number of trades placed on the synthetic is almost exponentially increased. Hence, the average profit per trade gets lower. Figure 7 shows several trading instances of a system with bounds set at 80% confidence. "Circles" represent entry and "crosses" symbolise exit points. Note the effect of the GARCH component of the model, which is to dynamically adjust the confidence bounds, or else the uncertainty about the future realised value, whenever large unexpected mispricings occur. This in general prevents trading in periods of high volatility and risk (see e.g. the first 50 observations of the lower "snapshot" of figure 7). In figure 4, we depict the equity curve for the afore-described trading system for the period September 23 - October 11, 2005. The total number of trades is 823 and the average profit per trade is 78.793/823=0.096 rupees. The profitability of this high-frequency system is severely limited by the fact that positions are necessarily closed at the end of the 1-minute interval[5]. It is important to note that keeping a trade open for a time interval grater than 1 minute is equivalent to not adjusting b_0 and b_1 until the trade is closed. The final outcome of such trades is strongly based upon how well the two synthetic time series, the every 1-min adjusted and that calculated from unchanged estimates of betas, locally resemble each other. It is certain that as the values of the mispricing calculated from unchanged estimates of betas have not been "seen" by the model, the performance of open trades will be unpredictable in the long-run. However, the extend of unpredictability has yet to be evaluated on an experimental basis.

[4] Longing (shorting) the synthetic means buying (selling) 1 stock of S_2 and selling (buying) b_1^{t-1} stocks of S_1.

[5] Recall that the synthetic time series depicted in figure 7 (solid line) assumes periodic recalculation of betas.

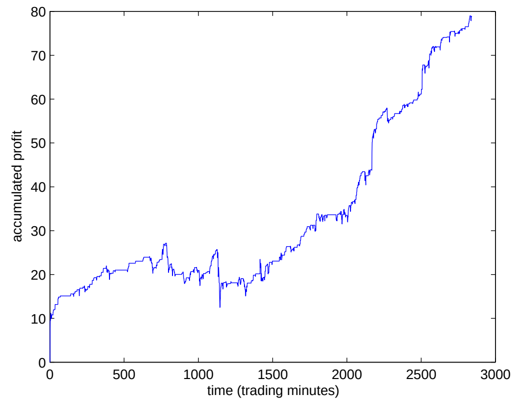

Fig. 4. The profit & loss diagram of the arbitrage trading system described in Experiment A for the sample period September 23 - October 11.

5.2 Experiment B: a lower-frequency model

As a next experiment, we calculate mispricings on hourly averages of stock prices. We use a sample of 804 observations, extending from February 8 to August 19, to specify and estimate the mispricing-correction model and we evaluate the performance of the resulting arbitrage strategy in the period August 29 - October 11, 2005. Noticeably, significant ARCH effects were not detected on hourly synthetic prices, although non-linearities in mean were present. The final error-correction model was thus a pure neural network regression with lags 1-8 in the linear part and 1 hidden neuron with lags 1, 4 and 6-9.

The obvious advantage of low-frequency compared to high-frequency models is that arbitrage trading positions last longer and thus potentially present larger profit opportunities. However, one should bear in mind that as predictions are available not until the next hour, the course of the synthetic within this time interval is largely undetermined. This of course affects the profitability of arbitrage positions. Figure 8 illustrates the point. The solid line shows the mispricing every 1 minute as computed by the hourly estimates of b_0 and b_1 and the dotted lines represent a 30% confidence bound obtained by 2-hour-ahead forecasts. The placing of a trade is based upon the position of the *1-hour ahead* confidence bounds. The encircled areas are typical examples of "heat" cases, where the synthetic moves in a direction adverse to the trading decision. Note that although these trading decisions are correct in the long run, their riskiness is increased when the synthetic experiences high volatility in the first period and hits the upper bound in the second one. Generally, this performance is unavoidable and does not depend on the specific choice of the confidence bounds; it is mainly the result of basing trading decisions on a macroscopic model that overlooks short-term adjustments. In our implementation of trading strategies, we decided to stop a trade whenever it hits the opposite bound in the corresponding sample time period. Of course, this strategy is not globally optimal but it is a way to place a limit on losses due to adverse price movements.

Figure 9 shows the performance of arbitrage trading strategies in the sample period August 29 - September 22, 2005 for varying confidence intervals. Overall, this trading system is more profitable than the one based on 1-minute bars: the number of trades is consistently lower and the average profit per trade is increased at all levels. For wide confidence bounds ($1-\alpha > 0.8$), trading becomes marginally profitable as the average profit exceeds the benchmark transaction cost.

In figure 5, we show the equity curve corresponding to the 80%-confidence trading system for the sample period September 23 - October 11, 2005. The chosen confidence level represents a relatively conservative arbitrage-exploiting policy, which takes a position whenever large mispricings occur. The total number of trades placed by the system is 9.3 and the average profit per trade is 55.60/9=6.18 rupees.

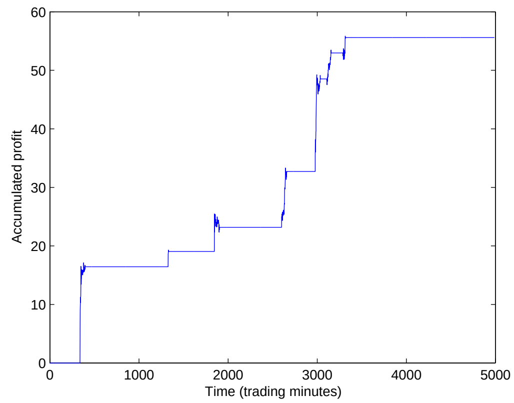

Fig. 5. The evolution of the trading account for a macroscopic arbitrage-detection system with confidence bounds set at 80% (Sample period: September 23 - October 11, 2005).

6 Discussion-Further research

This paper describes a new computational intelligent framework for detecting and exploiting short-term statistical arbitrage opportunities in a group of assets. The innovation of our approach is that it uses conditional density forecasts for the future value of the mispricing to properly design the arbitrage-exploitation strategy. These forecasts are adaptive, in the sense that they explicitly take into account short-term changes in volatility levels.

Experiments presented above show that our models are good detectors of arbitrage opportunities, as the equity curves are statistically sloping upwards. However, the profitability of these trades in a real market environment is still questionable given the various trading costs and market “frictions”. At present, we are conducting further research on the optimal tuning of the parameters of our trading system, stop-loss criteria as well as the possibility of combining forecasts from various sampling-frequency models so as to track both the long- and short-term behaviour of the mispricing dynamics. First results from the latter approach are rather encouraging.

References

1. T. Bollershev, *Generalized autoregressive conditional heteroscedasticity*, Journal of Econometrics **72** (1986), 307–327.

2. N. Burgess, *Statistical arbitrage models of the FTSE 100*, Computational Finance 1999 (Y. Abu-Mostafa, B. LeBaron, A. W Lo, and A. S. Weigend, eds.), The MIT Press, 2000, pp. 297–312.
3. ______, *Cointegration*, Neural Networks and the Financial Markets: predicting, combining and portfolio optimisation (J. Shadbolt and J. G. Taylor, eds.), Springer, 2002, pp. 181–191.
4. N. Burgess and A.N Refenes, *Modelling nonlinear cointegration in international equity index futures*, Neural Networks in Financial Engineering (A. Refenes, Y. Abu-Mostafa, J. Moody, and A. Weigend, eds.), World Scientific, 1996, pp. 50–63.
5. P. Cortez, M. Rocha, and J. Neves, *Genetic and evolutionary algorithms for time series forecasting*, Proceedings of the 14th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems: Engineering of Intelligent Systems, Lecture Notes In Computer Science, Springer-Verlag, 2001, pp. 393 – 402.
6. W. Enders, *Applied econometric analysis*, John Wiley & Sons, 1995.
7. R. F. Engle, *Autoregressive conditional heteroscedasticity with estimates of variance of united kingdom inflation*, Econometrica **50** (1982), 987–1008.
8. A.M. Farley and S. Jones, *Using a genetic algorithm to determine an index of leading economic indicators*, Computational Economics **7** (1996), 163–173.
9. A. Fiordaliso, *A nonlinear forecasts combination method based on Takagi-Sugeno fuzzy systems*, International Journal of Forecasting **14** (1998), 367–379.
10. I. Garrett and N. Taylor, *Intraday and interday basis dynamics: Evidence from the FTSE 100 index futures market*, Studies in Nonlinear Dynamics nd Econometrics **5** (2001), no. 2, 133–152(20).
11. J. R. Koza, *A genetic approach to econometric modeling*, Economics and Cognitive Science (Paul Bourgine and Bernard Walliser, eds.), Pergamon Press, Oxford, UK, 1991, pp. 57–75.
12. E. Steurer and T.H Hann, *Exchange rate forecasting comparison: neural networks, machine learning and linear models*, Neural Networks in Financial Engineering (A. Refenes, Y. Abu-Mostafa, J. Moody, and A. Weigend, eds.), World Scientific, 1996, pp. 113–121.
13. N. R. Swanson and H. White, *Forecasting economic time series using flexible vs fixed specification and linear vs nonlinear econometric models*, International Journal of Forecasting **13** (1997), 439–461.
14. G.G. Szpiro, *Forecasting chaotic time series with genetic algorithms*, Physical Review E **55** (1997), 2557–2568.
15. N.S. Thomaidis and G. Dounias, *A hybrid intelligent system for financial time-series forecasting*, International Journal of Engineering Intelligent Systems **16** (2008), no. 4, 193–213.
16. N. Towers, *Joint optimisation in statistical arbitrage trading*, Neural Networks and the Financial Markets: predicting, combining and portfolio optimisation (J. Shadbolt and J. G. Taylor, eds.), Springer, 2002, pp. 193–201.
17. G. Zhang, B. E. Patuwo, and M. Y. Hu, *Forecasting with artificial neural networks: The state of the art*, International Journal of Forecasting **14** (1998), 35–62.

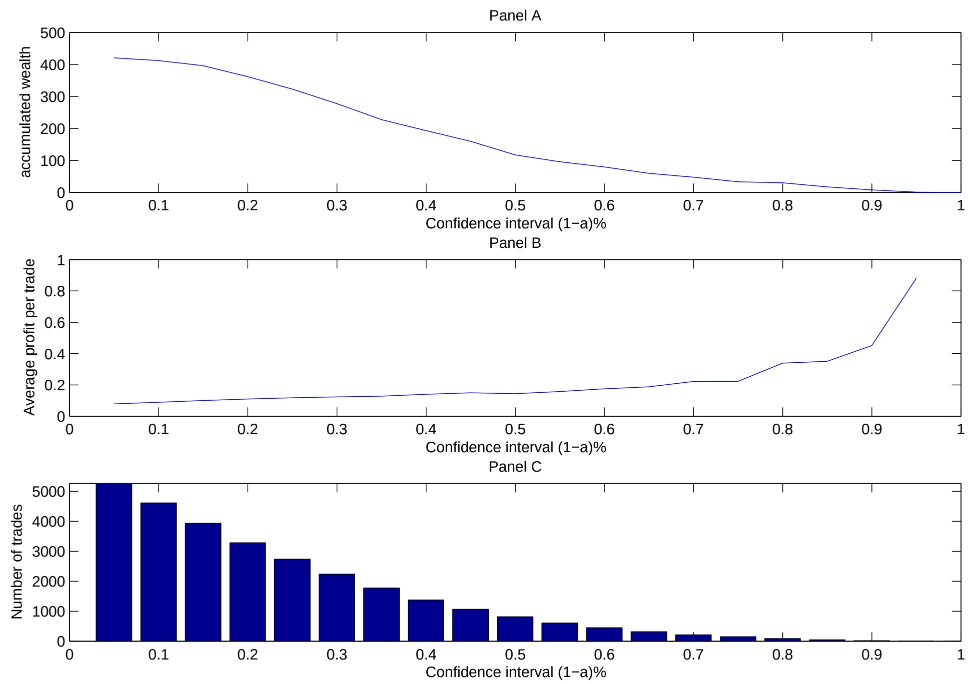

Fig. 6. The performance of the arbitrage trading strategy, corresponding to the microscopic system detailed in Experiment A, as a function of the confidence level (Sample period: August 29 - September 22, 2005). Panel A shows the accumulated profit, panel B the average profit per trade and panel C the total number of executed trades in the designated period.

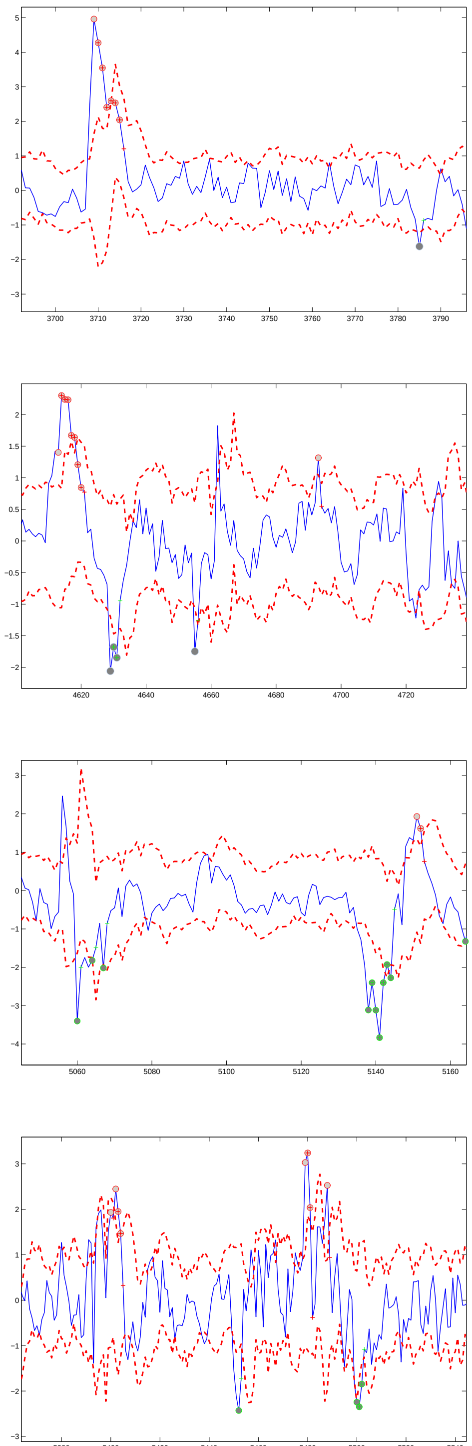

Fig. 7. Trading instances of a microscopic mispricing-detection model that has been estimated on 1-min bars of the mispricings time-series. Dotted lines represent a 80% confidence "envelope" on the value of the mispricing. The decision on opening a trade is taken with respect to 5-minute-ahead forecasts.

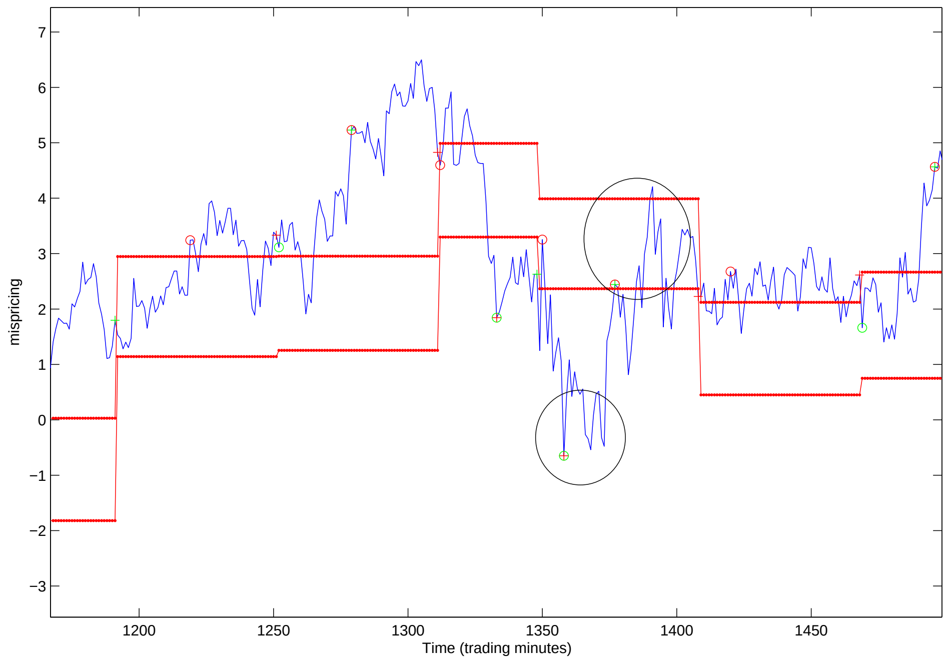

Fig. 8. A trading instance of an arbitrage system based on the macroscopic (1 hour) prediction model of Experiment B. The solid line shows the minute-by-minute evolution of mispricings calculated from hourly estimates of b_0 and b_1 and the dotted lines represent a 30%-confidence envelope based on 2-hour-ahead forecasts. The decision of whether to open a trade is determined by the position of the confidence bounds on the *following* 1-hour time interval.

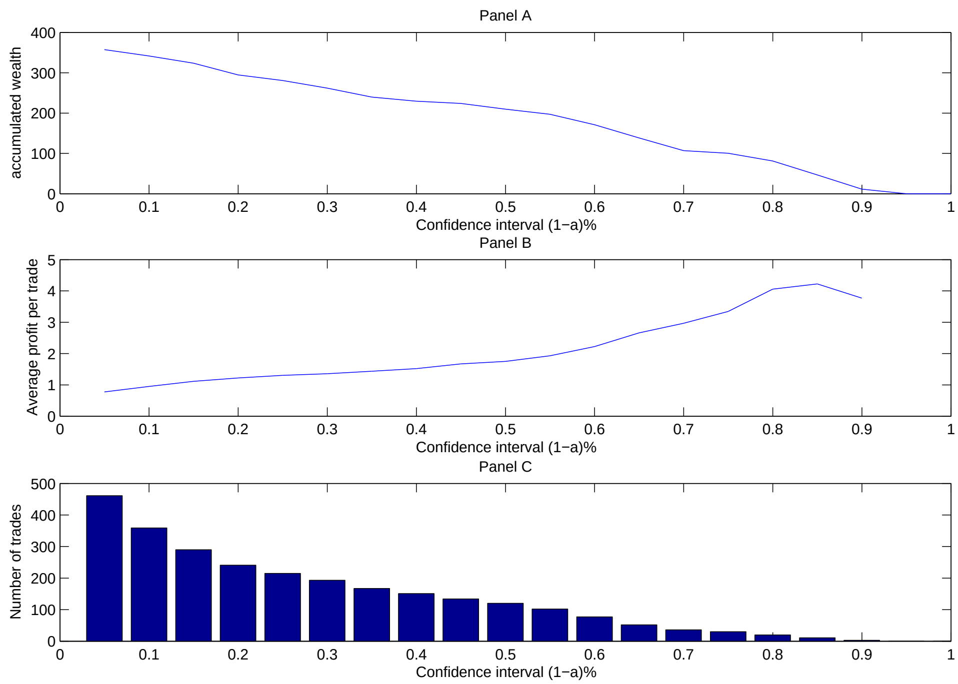

Fig. 9. The performance of the arbitrage trading strategy, corresponding to the macroscopic model detailed in Experiment B, as a function of the confidence level (Sample period: August 29 - September 22, 2005). Panel A shows the accumulated profit, panel B the average profit per trade and panel C the total number of executed trades in the designated period.