Int. Fin. Markets, Inst. and Money 37 (2015) 85–98

Contents lists available at ScienceDirect

Journal of International Financial Markets, Institutions & Money

journal homepage: www.elsevier.com/locate/intfin

Return predictability and the 'wisdom of crowds': Genetic Programming trading algorithms, the Marginal Trader Hypothesis and the Hayek Hypothesis

Viktor Manahov[a,*], Robert Hudson[b,1], Hafiz Hoque[a,2]

[a] The University of York, Heslington East, York YO10 5GD, United Kingdom
[b] The University of Hull, Hull HU6 7RX, United Kingdom

ARTICLE INFO

Article history:
Received 2 June 2014
Accepted 24 February 2015
Available online 3 March 2015

Keywords:
Forecasting and simulation
Agent-based modelling
Artificial stock market
Genetic Programming
Marginal Trader Hypothesis
Hayek Hypothesis

ABSTRACT

We develop profitable stock market forecasts for a number of financial instruments and portfolios using a special adaptive form of the Strongly Typed Genetic Programming (STGP)-based trading algorithm. The STGP-based trading algorithm produces one-day-ahead return forecasts for groups of artificial traders with different levels of intelligence and different group sizes. The performance of the algorithm is compared with a number of benchmark forecasts and these comparisons clearly demonstrate the short-term superiority of the STGP-based method in many circumstances. Subsequently we provide detailed analysis of the impact of trader cognitive abilities and trader numbers on the accuracy of forecasting rules which allows us to conduct new experimental tests of the Marginal Trader and the Hayek Hypotheses. We find little support for the Marginal Trader Hypothesis but some evidence for the Hayek Hypothesis.

1. Introduction

A voluminous literature exists on predicting stock market returns using different variables. However, most of the empirical studies suffer a major shortcoming – they find evidence in favour of in-sample return predictability but no presence of out-of-sample profitability. Goyal and Welch (2008) suggest that the predictive ability of nearly all stock return forecasting models does not hold up in out-of-sample validation. Moreover, a number of researchers argue that many studies finding predictability of stock returns might be spurious. This is based on the empirical assumption that the persistence of the forecasting variables and the correlation of their innovations with returns might bias the regression coefficients and negatively affect *t*-statistics (Stambaugh, 1999; Lewellen, 2004). These issues are exacerbated when researchers are data mining, taking into account many variables, and reporting only statistically significant findings (Ferson et al., 2003; Ferreira and Santa-Clara, 2011). Most prior studies investigate the relationship between the available information and the stock returns using simple linear regression. This is quite a restrictive approach as there is no evidence to support the assumption that the relationship between stock returns and the financial and economic variables in these studies is perfectly linear (Enke and Thawornwong, 2005).

* Corresponding author. Tel.: +44 1904325847.
E-mail addresses: viktor.manahov@york.ac.uk (V. Manahov), robert.hudson@hull.ac.uk (R. Hudson), hafiz.hoque@york.ac.uk (H. Hoque).
[1] Tel.: +44 1482463189.
[2] Tel.: +44 190432783.

http://dx.doi.org/10.1016/j.intfin.2015.02.009
1042-4431/

In this paper we use Strongly Typed Genetic Programming (STGP), a new forecasting approach to predict stock market returns, based on intelligent agents interacting in an artificial stock market. Research on financial markets has rarely used STGP models although they offer considerable advantages compared to any other forecasting techniques including the ability to evolve through time and deal with non-linear relationships (see Manahov et al., 2014a,b, for recent papers describing STGP models and applying them in financial markets). We compare the STGP predictions to several appropriate benchmark forecasts on an in-sample and out-of-sample basis using a number of performance measures to investigate whether, to what extent and in which form, stock returns in excess of the risk free rate are indeed predictable and profitable compared to a buy-and-hold strategy with dividends and without dividends. This is done by measuring and quantifying the exact level of generated profit, taking into account the transaction costs of 85 basis points per round trip.

The other important question under investigation in this study relates to how stock markets incorporate information and generate prices. There are two competing hypotheses about how markets work – the Marginal Trader Hypothesis (MTH) and the Hayek Hypothesis. While the MTH postulates that a small fraction of savvy individuals are capable of setting market prices, the Hayek Hypothesis suggests that markets can operate in an efficient manner even when market participants have a limited knowledge of the surrounding environment. Blackwell (2011) and Blackwell and Pickford (2011) argue that the MTH is so intuitive to economists that it seems never to have been tested directly and examined rigorously. In real financial markets it is difficult to determine whether the market price is actually the 'correct' price for any asset because is impossible to know the 'correct' price (Blackwell, 2011). Artificial stock markets do not necessarily suffer from this limitation because the market price of the instrument in the market can reflect some underlying prediction and the accuracy of that prediction can be checked. In addition, in real markets it is clearly not possible to directly control the number and intelligence of the participants. Using our STGP models we conduct experiments which avoid the difficulties in testing the MTH and the Hayek Hypothesis and provide appropriate laboratory conditions for direct examination of both hypotheses. To test the MTH we investigate whether a small fraction of the genetic population, the 'Best Performing Traders', are capable of generating more accurate forecasts. We show that the forecasting process emerging from the collective behaviour (competition and co-evolution) of the entire set of traders in the market is a better predictor than that produced by any small subset of traders, even the 'Best Performing Traders'. We demonstrate that enhanced genetic diversity leads to increased accuracy in forecasts because of more heterogeneous trading rules and behaviour. To examine the validity of the Hayek Hypothesis we populate another model with 'Low Intelligence Traders' in order to investigate their forecasting abilities. We provide evidence that the level of traders' intelligence does not necessarily play a significant role in improving the quality of forecasts (this represents new evidence favourable to the Hayek Hypothesis in a broader environment). The reminder of this paper is organised as follows: Section 2 presents the background to the study and reviews the literature in the field; Section 3 discusses the experimental design; Sections 4 and 5 discuss the simulation results relevant for prediction and assessing the 'wisdom of crowds' respectively and the paper concludes with overall remarks and discusses avenues for further research in Section 6.

2. Background

In this section we outline the background to the study and the relevant prior literature. In Section 2.1 we discuss the forecasting of financial markets and how Genetic Programming has previously been used in this domain. Section 2.2 outlines prior research on the implications of trader cognitive abilities for forecasting accuracy.

2.1. Forecasting and Genetic Programming

Time-series models used for forecasting financial markets have evolved over recent years. The failure of traditional linear models has promoted various innovations. Non-linear models have been increasingly developed (Chan and Ng, 2004; Clements and Smith, 1997). More recently the properties of combined forecasts from different models have been studied in an effort to approximately model the data generation process where forecasters are not able to fully identify the true process (Qian and Rasheed, 2007). Another approach is to implement non-parametric models, such as Genetic Programming (GP) models (Koza, 1995; Iba and Nikolaev, 2000). The advantage of the GP approach, in comparison with traditional models, is that it enables the researcher to be relatively agnostic about the general form of the optimal trading rule and to fully explore the non-differentiable space of trading rules.

A number of papers have applied GP to investigate trading rules in stock markets with varying degrees of success. Allen and Karjalainen (1999) find that out-of-sample the rules they find do not earn excess returns compared to a buy and hold strategy after taking into account transaction costs although they suggest their results were based on a very basic algorithm. Using monthly data Becker and Seshadri (2003) develop trading rules that were able to outperform a buy-and-hold strategy when dividends are excluded from stock returns. Iba and Sasaki (1999) show the GP has superior forecasting accuracy to neural networks in an experiment on the Nikkei 225. Kaboudan (2000) uses GP to produce reasonable one-day-ahead forecasts and develop a single day-trading strategy (SDTS) in which trading decisions are based on GP forecasts of daily highest and lowest prices. Potvin et al. (2004) applied GP to the Canadian TSE 300 index and demonstrated that the trading rules generated by GP are generally beneficial when the market falls or when it is stable. A handful of papers have used GP methods in financial markets other than stock markets. Neely et al. (1997) apply GP techniques to the foreign exchange market to investigate profitable trading rules. They found strong evidence of economically significant *ex ante* excess returns to technical trading rules for six different exchange rates over the period 1981–1995. Chen et al. (1998) and Chidambaran