Expert Systems With Applications 55 (2016) 70–86

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Improving risk-adjusted performance in high frequency trading using interval type-2 fuzzy logic

Vince Vella*, Wing Lon Ng

Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

ARTICLE INFO

Keywords:
High-frequency trading
ANFIS
Type-2 fuzzy logic
ANFIS/T2

ABSTRACT

In this paper, we investigate the ability of higher order fuzzy systems to handle increased uncertainty, mostly induced by the market microstructure noise inherent in a high frequency trading (HFT) scenario. Whilst many former studies comparing type-1 and type-2 Fuzzy Logic Systems (FLSs) focus on error reduction or market direction accuracy, our interest is predominantly risk-adjusted performance and more in line with both trading practitioners and upcoming regulatory regimes. We propose an innovative approach to design an interval type-2 model which is based on a generalisation of the popular type-1 ANFIS model. The significance of this work stems from the contributions as a result of introducing type-2 fuzzy sets in intelligent trading algorithms, with the objective to improve the risk-adjusted performance with minimal increase in the design and computational complexity. Overall, the proposed ANFIS/T2 model scores significant performance improvements when compared to both standard ANFIS and Buy-and-Hold methods. As a further step, we identify a relationship between the increased trading performance benefits of the proposed type-2 model and higher levels of microstructure noise. The results resolve a desirable need for practitioners, researchers and regulators in the design of expert and intelligent systems for better management of risk in the field of HFT.

1. Introduction

Most transactions in modern, highly computerised, financial markets are being greatly controlled by specialised algorithms which incessantly sift through masses of data and take split second trading decisions. According to a recent study by Brogaard, Hendershott, and Riordan (2014), between 2008 and 2010 high frequency trading (HFT) algorithms accounted for 70% of dollar volume on the NASDAQ exchange. This tends to defy the long standing Efficient Market Hypothesis (EMH) (Fama, 1965, 1970) that states that current prices incorporate all relevant information with no possibility of predictability or excess returns. A number of authors (e.g. Schulmeister, 2009; Zhang, 2010; Rechenthin & Street, 2013; Holmberg, Lönnbark, & Lundström, 2013; Brogaard et al., 2014) insist that the presence of efficient pricing becomes more questionable when investigating short-lived (milli-seconds to few minutes) trades. However, Kearns, Kulesza, and Nevmyvaka (2010) validate the EMH in their study and argue that generating profits from aggressive high-frequency trading (HFT) is next to impossible. These debates keep this domain a very active area of research.

According to Johnson et al. (2013) this new machine dominated reality highlights the need for new theories in support for sub-second financial phenomena during which the human traders lose the ability to react in real time. Due to the non-stationary characteristics of financial time series (see Fama, 1965), applying machine learning techniques to infer predictions is a challenging task and prone to increased error variance. Complexity is heightened given the level of noise in high frequency stock price movements. Incidents like the “flash crash” of 6 May 2010 stress the importance of risk management. As a result, in recent years HFT and algorithmic trading have been the subject of increasing global regulatory attention. As an example, in October 2011, the EU Commission proposed a new version of the Markets in Financial Instruments Directive (MiFID2). MiFID2 will apply from January 2017 and will introduce a new regulatory regime for firms which engage in algorithmic or high-frequency trading on European venues or who provide investment management services directly to clients in the EU. The new regulations are intended to ensure that trading systems are adequately designed and tested to mitigate the risks to which they are exposed.

This acts as a reminder for model or algorithm designers that both investors and regulators are more concerned about risk-adjusted performance rather than directional accuracies, error measures or just profitability. Unfortunately, surveys show that the

* Corresponding author. Tel.: +35699432217.
E-mail addresses: vvella@essex.ac.uk (V. Vella), wlng@bracil.net (W. Lon Ng).

http://dx.doi.org/10.1016/j.eswa.2016.01.056
0957-4174/

great majority of computational finance research disregards the essential requirement that an investor always measures investment returns in line with risk measures in order to compare relative trading performance (Krollner, Vanstone, & Finnie, 2010; Tsai & Wang, 2009). A higher return trading strategy does not necessarily outperform other strategies if the associated risk is also substantially higher. This is the reason why in this paper we focus on risk-adjusted performance.

Zadeh (1975) proposed that increased systems complexity calls for approaches that are significantly different from the traditional methods which are highly effective when applied to mechanistic systems. Fuzzy sets and systems are attributed as an excellent method to deal with situations where the element of uncertainty and imprecision is high, typically prevalent in complex environments (see Wagner & Hagras, 2010). A number of surveys (Krollner et al., 2010; Sahin, Tolun, & Hassanpour, 2012; Tsang, 2009) place Artificial Neural Networks (ANNs) amongst the most popular learning techniques in AI-based financial applications and hybrid models. ANNs, especially in conjunction with fuzzy logic, were found to provide better forecasts.

A frequently cited technique in non-stationary and chaotic time series prediction, which combines ANN and type-1 (T1) fuzzy logic, is the Adaptive Neuro-Fuzzy Inference System (ANFIS) by Jang (1993). Successful application and active continuous research in improving ANFIS based techniques in trading applications is demonstrated by numerous publications (Boyacioglu & Avci, 2010; Chang, Wei, & Cheng, 2011; Chen, 2013; Gradojevic, 2007; Kablan & Ng, 2011; Tan, Quek, & Cheng, 2011; Vella & Ng, 2014b; Wei, Cheng, & Wu, 2014). Moreover Vella and Ng (2014b) showed the increased stability of ANFIS in terms of risk-adjusted performance when compared to ANN alone. Recently, type-2 (T2) fuzzy logic have gained significant academic attention (see review in Melin & Castillo, 2014) and as of today it remains a primary area of research in the fuzzy logic domain (Mendel, Hagras, Tan, Melek, & Ying, 2014). To our best knowledge, the use of higher order fuzzy logic systems (FLSs) in a high frequency trading environment has not been addressed in the literature before. Our intention in this paper is to investigate the possible improvement that can be obtained by generalising ANFIS to interval T2 (IT2) FLS. However, in line with Wu and Mendel (2014), we argue that although T2 FLSs provide the researcher with extensive freedom in design options, the increased computational and design complexity can possibly hinder the wider applications of such systems. This challenge was a key consideration that inspired our innovative and practical IT2 approach that we present in this paper.

The investigation of possible improvements using T2 in HFT is appealing in view of increased uncertainties which are inherent in high frequency data. Albeit the concepts of risk and uncertainty have often been used interchangeably, economists have long distinguished between the two (e.g. Knight, 1921) and also in recent literature (e.g. Nelson & Katzenstein, 2014; Heal & Millner, 2014). Our view is that overall risk can be divided into measurable risk (e.g. flip of a fair coin), and uncertainty, which we categorise as the risk of events to which it is difficult to attach a probability distribution. Our aim is to further reduce the trading uncertainty by utilising T2 FLSs. We have not identified any existing literature that investigates the level of noise (uncertainty), indirectly reflecting the trading frequency, that would warrant the (feasible) use of T2 over T1 fuzzy logic methods for algorithmic trading purposes.

With respect to the above literature review and identified gaps, in this study the objectives can be summarised as follows:

1. To identify practical methods of how the popular ANFIS model can be generalised to an interval T2 Takagi-Sugeno-Kang (IT2 TSK) fuzzy system. We aim to address this with minimal increase in design and computational complexity.
2. To investigate the ability of higher order fuzzy systems to handle increased uncertainties inherent in a HFT scenario.
3. To identify *if* T2 FLSs provide a viable alternative for trading purposes in view of improving risk-adjusted performance.
4. To explore *when* can T2 models offer a more viable approach than T1 alternatives. We analyse this from the perspective of different levels of trading frequencies.

This paper aims to convey a number of contributions. As a first contribution we propose an innovative, but at the same time a more accessible, way of how to design a T2 FLS from an optimised T1 Neuro-Fuzzy FLS (ANFIS/T2). With IT2 there are many, sometimes overwhelming (Wu & Mendel, 2014), design choices to be made which includes the shape of membership functions, number of membership functions, type of fuzzifier, kind of rules, type of i-norm, method to compute the output, and methods for tuning the parameters. We address this from a number of aspects. Firstly, we make use of a fuzzy clustering algorithm for rule identification in order to reduce the number of rules and hence simplify the model. We apply simple rules where antecedents are T2 fuzzy sets and consequents are crisp numbers (A2-C0). Secondly, as our base structure for the T2 model we use the popular ANFIS as a solid benchmark model as it is computationally fast and also has been successfully applied in high frequency trading (Kablan & Ng, 2011; Vella & Ng, 2014b). Thirdly, we reduce the training complexity by reducing the number of tuning parameters, limiting this to varying sizes of the Footprint of Uncertainty (FOU). Our parsimonious approach also reduces the possibility of overfitting and spurious results (see Bailey, Borwein, de Prado, & Zhu, 2014). Finally, we apply an efficient closed form type reduction method.

As a second contribution we shed more light on the theoretical market efficiency debate in HFT. Schulmeister (2009) points towards possible market inefficiencies and profitability of technical trading rules at higher frequencies, this being driven by faster algorithmic trading. Recently, Rechenthin and Street (2013) claimed that when price shocks break the bid-ask spread, which was identified to happen anywhere in between 5 and 10 s, price movements can be predicted for up to one minute. Beyond this point prediction probabilities remained significant for about the next 5 min, dying out completely beyond 30 min. In our case we make use of HFT trade data from a set of stocks listed on the London Stock Exchange and investigate a combination of technical rules on 2 min returns with holding periods ranging between 2 and 10 min. Contrary to findings in AI surveys (Krollner et al., 2010; Tsai & Wang, 2009), we align ourselves with the priorities of investors and regulators and focus on comparing the proposed models using risk-adjusted performance (Choey & Weigend, 1997; Vanstone & Finnie, 2010; Xufre Casqueiro & Rodrigues, 2006). We are not aware of any previous studies which investigate the link between higher order fuzzy systems and risk-adjusted performance.

Finally, as our third contribution we try to answer an important question which explores, from a trading performance perspective, *when* it is viable to apply T2 models rather than T1. Although previous literature found that T2 models can perform better under increasing uncertainties (Aladi, Wagner, & Garibaldi, 2014; Sepulveda, Melin, Díaz, Mancilla, & Montiel, 2006), it is however not clear at which uncertainty level this would reflect in a reasonable improvement in risk-adjusted trading performance. Birkin and Garibaldi (2009) even showed that if the level of noise is too low, T2 models show no significant improvement on T1. A number of authors (Gençay, 1996; Holmberg et al., 2013; Vanstone & Finnie, 2009, 2010) suggest the use of a threshold on the predicted signals below which a trading action is not taken into consideration. This is done to reduce the effect of the underlying noise, however, at the cost of reduced trades and hence possible return. We propose an innovative experiment approach by extending this technique to