Trading Rule Search with Autoregressive Inference Agents

McGregor J. Collie, David L. Dowe and Leigh J. Fitzgibbon
Clayton School of Information Technology
Monash University, Wellington Road, Vic. 3800, Australia
gjc@csse.monash.edu.au

Abstract

The use of the agent-based paradigm in modelling financial markets provides an intuitively natural approach and is a well established technique. In contrast with the assumptions and conclusions of the efficient markets hypothesis (EMH), agent based models provide a refreshing causal approach to understanding the emergence of the general stylized facts of financial markets. In this report we present details of an agent-based stock market simulation in which traders utilise a hybrid mixture of common information criteria based inference procedures, including minimum message length (MML) inference. Traders in our model compete with each other using a range of different inference techniques to infer the parameters and appropriate order of simple autoregressive (AR) models of stock price evolution. We show that in the presence of a noisy AR signal, MML traders significantly outperform their competitors, and in fact do well even in the absence of such a signal.

1 Introduction

The efficient markets hypothesis (EMH) as popularised by Fama [26] and others (e.g., Jensen (1978)[35] and Malkiel (1973)[43]) presents us with the claim that the market is 'efficient with respect to [an] information set ... if it is impossible to make economic profits by trading on the basis of [that] information' [35].

This model ignores the behaviour of individual trading agents in the system, relegating them to an arbitrage role in which more efficient traders exploit less efficient traders to keep the market correctly priced, with respect to the information set they are basing their decisions upon. See, e.g., Mayhew [46] for further discussion. The EMH has been disputed on various conceptual grounds. Farmer and Lo [27] point out that the EMH is 'not a well-poised and empirically refutable hypothesis' (Ibidem, p.2), and that only concepts of *relative* efficiency make any sense. Furthermore, they note that by extending the concept of strict efficiency to scientific research one ends up with the conclusion that any worthwhile research must have already been done, and hence that no progress is possible; e.g., if it were worth working out whether markets are efficient or not, someone would have already done it. Dowe and Korb [24] object to the EMH on various other grounds. One ground upon which they object is computational; the undecidability of Kolmogorov complexity [55, 36, 15] and the relationship between Minimum Message Length (MML) inference and Kolmogorov complexity [66] means that we can rarely, if ever, prove that our inference technique is superior to that of other potential traders in the market, and hence we can never (or rarely) preclude the existence of a superior inference technique to our own. It follows then that we cannot, in general, prove that a non-trivial market is efficient.

Agent based micro-simulation models of individual trading explore the possibility that the emergent behaviour of individual interactions is responsible for observed market behaviours [18]. Tesfatsion [59] conducts a review of agent based approaches in finance, and LeBaron [40] reviews the methodology and construction of agent-based stock market simulations in particular. Within the context of heterogeneous agent-based simulation, one can introduce agents that correspond with rational expectations (RE) type assumptions (see e.g., Chen and Yeh [16]), as well as agents that correspond to behavioural models of individual economic behaviour (for a review of behavioural finance the reader is referred to Shiller [54], and Barberis and Thaler [6]).

In this simulation we introduce a set of trading agents into an artificial single stock trading environment who attempt to model the evolution of the stock price using an autoregressive (AR) model. Furthermore, the AR agents are divided into subsets of agents who use different *information criteria* (IC) to select an autoregressive model order. The method used to generate parameter estimates for the AR models used here is that of conditional least squares, as outlined in [28], modified such that the maximum likelihood estimators of the standard deviation are used for the stationary model variance estimates, as described in Box, Jenkins and Reinsel ([12, pp296-304]).

The use of an information criterion as a means to selecting a parsimonious model to explain observed data is fairly controversial in terms of implementation, if not in principle. Commonly used information criteria which have been used to penalise models with an excessive number of explanatory factors (typically leading to increasingly poor predictive accuracy) include Akaike's AIC [2], Hurvich and Tsai's corrected AIC [33], Schwarz's Bayesian IC (BIC) [53], Rissanen's 1978 mimimum description length (MDL) [51, 52], Hannan and Quinn's IC (HQIC) [31], and C.S. Wallace's Minimum Message Length (MML) criterion ([64] (with Boulton) [69] (with Freeman) [66] (with Dowe) [63]). Based upon the work of Fitzgibbon, Dowe and Vahid [28] in which various inference techniques are compared on various generated AR signals, we implement within this simulation agents that embody a range of different information criteria based inference techniques, and allow them to compete directly with each other[1].

The stock market simulation used here is an extension of the artificial stock market presented in Collie [18], in which agents are selected randomly from a trader pool of fixed total size to appraise the market (consider the sequence of past prices) and potentially submit bids to a double-auction process. Other stock market simulations using an agent-based trading methodology include the pioneering work of the Santa Fe simulation [32, 4], the Genoese simulation [50] and others, e.g. [42, 34, 17], See LeBaron [41] and also his website[2] for further references.

The next section outlines the set-up of the stock market simulation in more detail, and presents a closer examination of the information criterion used by the trading agents. We then discuss the results obtained, and their implications for discussions of market efficiency. Finally we conclude with potential directions for further research.

2 Simulation Design

2.1 Stock Market Design

As mentioned, the stock market simulation used here is an extension of the artificial stock market first presented in Collie [18]. Agents participate in multiple rounds of a continuous double-sided auction of a single tradable asset, submitting buy and sell orders at fixed prices ('at limit' bidding). Unmatched or partially matched orders are submitted to an 'order book', as commonly employed in modern exchanges. New trades are matched against existing orders in the book. Such order books are often publicly visible, and are often used to estimate underlying 'depth' and hidden information in the market for a security [9, 11], but agents in this simulation do not consider this information. Each trading round consists of a random number of traders randomly selected from the trader pool. When an agent is selected from the trader pool they appraise the market according to their characteristic trading strategy or trading rules, and potentially submit an order to the market. If, upon the new order being submitted, a matching or partially matching pre-existing bid is found in the order book, then these bids are matched and the trade is executed at the average of the buy and sell prices. Where two existing orders enter the book at the same price, the earlier order takes precedence. Agents in this simulation trade a single stock with no dividend; the total amount of cash and shares available remains constant throughout the simulation, although total wealth levels may fluctuate with the current trading price of the asset.

After the bid submission and matching process has been cycled through a predetermined number of times, an 'end-of-day' phase is reached, and the order book of outstanding orders is cleared, and the process repeats. Open, close, high and low prices and trade volume are recorded, and the cycle re-initiated for the next 'day', or round, of trading. Simulations are run for an exogenously determined number of trading rounds, or until some other termination condition is reached, such as the cessation of trade by the trading agents. This may occur, for example, if one of the agents has captured the bulk of the available wealth in the market.

Traders are initially allocated equal numbers of shares and an equal value of virtual currency with which to trade. The total amount of shares and currency within the simulation is held constant throughout, but the total amount of wealth available at any one time fluctuates with the current trading value of the asset. There are 100 traders in total in each simulation, of which 40 are random or AR signal generating, with the remaining 60 divided evenly amongst the 6 inference techniques examined[3]. Orders to buy and sell take the form of 'at limit', or fixed price orders. Such order books are often publicly visible, and are often used to estimate underlying 'depth' and hidden information in the market for a security [9, 11], but agents in this simulation do not consider this information.

We present three simulations here, one in which inference traders act in a market with each other and randomly trading 'noise' agents, one in which the random traders are replaced by a set of traders who calculate future price changes as following an exogenously specified noisy AR process, and one in which inference agents interact with each other with noise only from their own trading sig-

[1] We omit pricing approaches like the Capital Assets Pricing Model (CAPM) mean-variance trade-off. Dowe [23] considers it to be largely discredited.

[2] http://people.brandeis.edu/∼blebaron/acf/index.htm

[3] The proportion of the number of random to inference traders is exogenously determined, and is arrived at by attempting to gather enough random traders to provide necessary liquidity, whilst not so many that the emergent properties of the market take too long to appear.

nals. We demonstrate significant performance differences between the inferential traders in the presence of the underlying noisy AR signal, the Gaussian noise traders, and even in the absence of explicitly noisy traders.

Trading agents participating in these simulations are either of the randomly trading variety, or are one of a number of different types of autoregressive inference traders.

2.2 Random Traders

Random trading agents (or 'noise traders') are generally introduced into agent based stock market simulations as a means of providing market liquidity. Collie [18] has shown that simulations involving randomly trading agents with only a simple monetary constraint (finite assets) generate realistic stock price dynamics. One of the earliest models incorporating random traders appears to be that of Kyle in 1985 [38], coming just before Fischer Black's important paper on trading noise in empirical markets [10]. Jack Treynor described a three-trader model of the stock market (trading with a market maker) containing a random ('liquidity-motivated') trader in 1971 [5]. Probably the most well known artificial randomly agent trading model is the 'zero intelligence' model of Gode and Sunder [29], where randomly trading agents subject to a 'budget constraint' achieve high allocative efficiency in a double auction.

The random trading agents used in this model differ slightly from those of the previous random traders in [18]. When a randomly trading agent is selected from the trader pool, they clear any existing, previously unfilled orders remaining in the order book, and choose a uniformly distributed random number from 0 to 1 (random numbers are generated using the Mersenne Twister generator of Matsumoto and Nishimura [45]). This number is then compared to the trader's current ratio of stocks to cash, and (if necessary) an order is submitted to the market to adjust their current position to that of the randomly generated position. The price this order is submitted at is drawn from a Gaussian distribution around the last price change. These generated prices are not bounded below or above by the current price, so that (for example) a randomly generated sell order may be submitted at a price greater than the current price.

2.3 Noisy and Inferential AR Traders

Autoregressive time series processes for a time series $y(t)$ are of the general form

$$y_t = \sum_{i=1}^{p} [\phi_i y_{t-i}] + \epsilon_t \tag{1}$$

where ϵ_t is an $N(\mu_t, \sigma_t^2)$ Gaussian i.i.d. error term with average μ_t and variance σ_t^2.

In our second simulation we introduce both a noisy AR signal and AR signal-detecting inference traders. The AR signal is introduced into the stock price series by replacing the random noise traders of the previous simulation with the noisy AR signal generating traders, who generate trade prices using an autoregressive model like that given above in eq. (1).

The order p of the autoregressive function is varied randomly from one to eight when the simulation is initialised, and the parameters (ϕ_i) of the auto-regression chosen not necessarily to guarantee stationarity. When the order p of the model is greater than one, the individual parameters are initially chosen as $A_i = 1.0 + \epsilon_i$, where ϵ_i is assumed Gaussian i.i.d. with standard deviation of 0.01, and the final parameters are then normalised; $\phi_i = A_i / (\sum_{j=1}^{p} A_j)$. For AR models of order one the stationarity condition is imposed, the parameter ϕ is chosen as $N(0.99, 0.01)$, such that ϕ is less than or equal to one.

Fitzgibbon, Dowe and Vahid in [28] show clear differences in inferential power amongst different information criteria (IC) based inference techniques. We introduce groups of different inference technique based trading agents into the second simulation to potentially exploit the noisy AR signal in the price series, and to compare advantages to using different types of inference techniques. We include agents that assume both a stationary and non-stationary process, and model data accordingly.

We allow inferential agents to construct AR models to forecast future price changes of order $p \leq 8$. In order to determine the order of AR model used to estimate the next price change, traders rely upon their particular IC implementation.

2.4 Competing Inference Traders

The third simulation we undertake involves the removal of all noise traders, with only the different types of inference traders participating in the trading scenario. Liquidity is introduced to the model via the inference traders themselves; when they cannot determine a signal in the past price sequence, due most likely to a failure of the OLS routine, then the traders can submit random bids. Since all inference traders use the same OLS routine to estimate the parameters of their autoregressive models, a failure of the OLS routine does not prefer one kind of trader over another. Traders only have a potential for advantage when all trader types could hypothetically trade on the available past price sequence.

3 Inference Technique and Parameter Estimation

Parameter estimation for the AR models is done by individual agents using the techniques and methods outlined in [28], with some implementation differences. The use of IC in model selection has generally been used as a means of augmenting maximum likelihood (ML) techniques so as to identify not simply the model that best fits the data, but rather the model that best *explains* the data; that is, the most parsimonious model. Such models are generally chosen on the basis of minimisation of model complexity [13], or minimum message (hypothesis and data given hypothesis) length [66]. As noted by Hanlon and Forbes [30], these IC in general take the form

$$-n\log(\sigma^2) + Penalty(k,n), \quad (2)$$

where the penalty term is a function of the number of parameters, k, and the sample size, n. The IC above is minimised by the appropriate selection of model order, $p = k-1$.

The four inference trader types used in this simulation that use an IC of the form of eq. (2) use: Akaike's information criteria [3] (hereafter AIC), corrected AIC [33](CAIC), Schwartz's Bayesian IC [53], which is here equivalent to the 1978 implementation of the MDL [51] technique (BIC), and Hannan and Quinn's information criteria [31] (HQIC). Their IC forms are:

$$\text{AIC}(k) = -2\ln(f(\mathbf{y}|\hat{\phi}_k)) + 2k \quad (3)$$

$$\text{CAIC}(k) = -2\ln(f(\mathbf{y}|\hat{\phi}_k)) + \frac{2(k+1)n}{n-k-2} \quad (4)$$

$$\text{BIC}(k) = -2\ln(f(\mathbf{y}|\hat{\phi}_k)) + k\ln n \quad (5)$$

$$\text{HQ}(k) = -2\ln(f(\mathbf{y}|\hat{\phi}_k)) + 2k\ln(\ln(n)) \quad (6)$$

The IC formalism of eq. (2) does not explicitly take into account factors relating to prior probability of potential model choices (see, e.g., [69, p251], [30, 8], [66, pp279-280]).

3.1 MML Inference

The Minimum Message Length (MML) formalism of Wallace et al. [64, 69, 66, 67, 68, 63] differs from the usual informational criteria in that it uses an explicitly Bayesian approach, in that additional terms are included in the information criterion specifying the prior distribution over the model parameters (see, e.g., Baxter and Oliver [8] for details). Furthermore, the inclusion of a term involving the determinant of the expected Fisher information matrix captures further information about the appropriate weighting of observed data from different regions of the data space ([30, section 2]). As outlined in Fitzgibbon, Dowe and Vahid [28], the MML87 IC approximation [69] used in this simulation is described by

$$-\log\left(f(\mathbf{y}|\boldsymbol{\phi})\right) - \log\left(\frac{h(\boldsymbol{\phi})\epsilon^n}{\sqrt{|\mathbf{I}(\boldsymbol{\phi})|}}\right) + \frac{p}{2}\left(1+\log\left(\kappa_p\right)\right) - \log\left(h(p)\right) \quad (7)$$

where $f(\mathbf{y}|\phi))$ is the likelihood function for the model $\boldsymbol{\phi} = (\phi_1, ..., \phi_p, \sigma^2)$ over the n observed data points ($\mathbf{y} = y_1, ..., y_n$), with $p+1 = k$ model parameters. $|\mathbf{I}(\boldsymbol{\phi})|$ is the determinant of the expected Fisher information matrix, κ_p is a space quantising lattice constant, $h(p)$ is a prior over the model order, p, and ϵ is an estimate of data measurement error.

These parameters, along with the estimation of the AR parameters, are described below. They are substituted into equation (7) to give the message length in terms of the autoregressive order, p. The shortest message length for a given string of past prices is then used to determine the appropriate order of AR model used in forecasting the next price, and the estimated standard deviation of this forecast.

3.1.1 Fisher Information

The Fisher information matrix, $\mathbf{I}(\boldsymbol{\phi})$, is obtained from the second partial derivatives of the log-likelihood (10), $E\frac{\partial^2 l}{\partial\phi_i\partial\phi_j}$. Its value is approximately [12, p303][28]

$$\frac{n^{p+1}}{2\sigma^4}|\mathbf{M}_p|^{-1} \quad (8)$$

For the non-stationary case, the partial expected Fisher information is used

$$\mathbf{I}_{y_{p+1},...,y_n}(\boldsymbol{\phi}) \approx \begin{bmatrix} \sigma^{-2}X'X & 0 \\ 0 & \frac{n-p}{2\sigma^4} \end{bmatrix} \quad (9)$$

3.1.2 Bayesian Priors

For a stationary AR time series, the message length given by equation (7) is calculated by substituting in the expected Fisher information (8), the likelihood (10), priors over the number of parameters $h(p)$ (over which we assume a uniform prior)[4], and $h(\boldsymbol{\phi}) \propto \frac{1}{R_p}\frac{1}{\sigma^2}$, the prior over the parameter space R; the stationarity region of the AR series. The stationarity region for $p > 3$ is complicated [47, 7]. As given in [28, equations 21 and 22], we use an (uninformative) uniform prior over the parameter space, since as noted above and also in [28, footnote 4, page 4], the ranges used for p and σ^2 do not affect the model selection in this experiment, although in various circumstances where the absolute

[4] This uniform prior over some parameter range will obviously affect the message length (in terms of an additive constant), but for the purposes of comparing between different message lengths for different AR models, the prior chosen will simply cancel out, and not affect the final message choice. The range of parameters over which the prior is chosen is hence arbitrary, and hence $h(p)$ is set to unity for the calculations.

message length is important it will be necessary to use an alternative, e.g., a Beta distribution [44]. The hypervolume of the parameter space is given by Piccolo [49] in terms of the following recursion:

$$R_1 = 2 \qquad M_1 = 2$$
$$R_p = (M_1 M_3 \times \cdots \times M_{p-1})^2 \quad \text{for } p \text{ even}$$
$$M_{p+1} = \frac{p}{p+1} M_{p-1} \qquad R_{p+1} = R_p M_{p+1}$$

3.1.3 Lattice Constants

As explained in [63, p178, subsection 3.3.4] (also [39, p15-16]), in considering the set of possible intervals that the estimates of ϕ can take within $\boldsymbol{\phi}$, we regard this set as being representable by regular lattice structures, where each such possible set has the same geometry. The most efficient encoding of this structure, that is, the most efficient packing of the lattice of points within the space is called the *optimal quantizing lattice*, describable by the lattice packing constant, κ, which is given exactly here only for the first three dimensions, but is bounded above by $\kappa < ((p/2)!)^{2/p}(2/p)!/\pi$, and from below by $\kappa > ((p/2)!)^{2/p}/((p+2)\pi)$. Conway and Sloane [22, p61] give some of the best known quantizing lattices. The first eight are reproduced in table 1.

p	κ_p
1	$1/12$
2	$5/(36\sqrt{3})$
3	$19/(192 \times 2^{1/3})$
4	0.076603
5	0.075625
6	0.074244
7	0.073116
8	0.071682

Table 1. Quantizing lattice constants, κ_p.

3.2 Parameter Estimation

Given an AR time series (eq. 1) de-trended to insure that the mean is zero ($\mu = 0$) and assuming normality for the ϵ's, we can write the exact log-likelihood as ([12, p299])

$$l(\boldsymbol{\phi}, \sigma_\epsilon | \mathbf{y}) = -\frac{n}{2}\ln(\sigma_\epsilon^2) + \frac{1}{2}\ln|\mathbf{M}_p| - \frac{S(\boldsymbol{\phi})}{2\sigma_\epsilon^2} \qquad (10)$$

where

$$S(\boldsymbol{\phi}) = \mathbf{y}_p'\mathbf{M_p}\mathbf{y}_p + \sum_{t=p+1}^{n}(y_t - \phi_1 y_{t-1} - \ldots - \phi_p y_{t-p})^2 \qquad (11)$$

$(\mathbf{M}_n)^{-1}\sigma_\epsilon^2$ is the $n \times n$ autocovariance matrix of the y's,

$$\begin{bmatrix} \gamma_0 & \gamma_1 & \cdots & \gamma_{n-1} \\ \gamma_1 & \gamma_0 & \cdots & \gamma_{n-2} \\ \vdots & \vdots & & \vdots \\ \gamma_{n-1} & \gamma_{n-2} & \cdots & \gamma_0 \end{bmatrix}$$

Since the maximum likelihood estimates of $\boldsymbol{\phi}$ obtained from $\frac{\partial l}{\partial \phi_j}$ $(j = 1, 2, ..., p)$ generally require analytic solutions which do not always converge [14], we use conditional least squares (OLS) estimates, which do. The likelihood of $y_{p+1}, ..., y_n$ conditional upon the first p values of $\mathbf{y}$ is given by ([12, p297])

$$p(y_{p+1}, ..., y_n | \mathbf{y}_p, \boldsymbol{\phi}, \sigma_\epsilon) = (2\pi\sigma_\epsilon^2)^{-(n-p)/2} \times \exp\left[-\tfrac{1}{2\sigma_\epsilon^2}\textstyle\sum_{t=p+1}^{n}(y_t - \phi_1 y_{t-1} - \cdots - \phi_p y_{t-p})^2\right] \qquad (12)$$

which, if the log is taken, and the partial derivative with respect to σ_ϵ set to zero, yields the conditional least squares estimate for σ_ϵ^2,

$$\hat{\sigma}_\epsilon^2 = \frac{1}{n-p}\sum_{t=p+1}^{n}(y_t - \hat{\phi}_1 y_{t-1} - \cdots - \hat{\phi}_p y_{t-p})^2 \qquad (13)$$

The minimisation of the sum in the above estimate yields the OLS estimates for $\boldsymbol{\phi}$, ([12, p301]) $\widehat{\boldsymbol{\phi}} = \mathbf{D}_p^{-1}\mathbf{d}$, where $D_{ij} = \sum_{t=p+1}^{n} y_{t-i}y_{t-j}$ and $d_i = \sum_{t=p+1}^{n} y_{t-i}y_t$.

4 Results

Figures 1 and 2 show average wealth levels for the different classes of traders across 10 simulations, where the results for each individual simulation are averages across each agent class over 150 trading rounds. Within each trading round there are between 3,600 and 8,200 individual potential trading opportunities for each of the 100 agents in each simulation.

In our first simulation inference agents attempt to model an AR time series from the prices generated by the noise traders and their interaction with them. As one might expect, with little or no signal to detect, the wealth levels of the inference traders don't show much variation; they manage to take most of the noise traders' wealth from them, but their own wealth, affected by the declining price of the stock they are trading, does not increase significantly after an initial, highly volatile trading period.

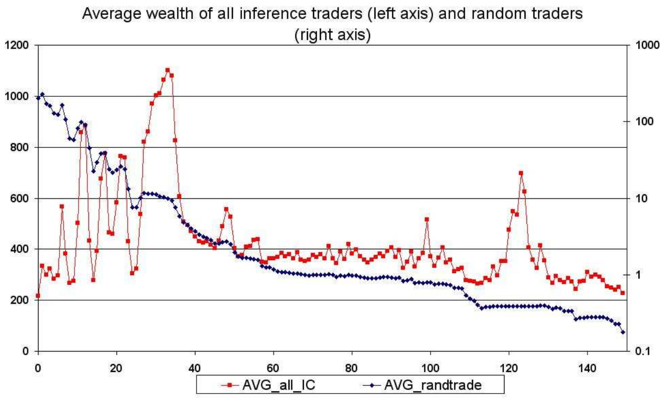

Figure 1. Random noise traders and inference agents average wealth.

Figure 1 shows the average wealth levels for the random (AVG_randtrade) and inference agents (AVG_all_IC), for approximately 150 trading rounds, where wealth levels are averages across 10 simulations, each simulation therein averaging the results for each trader class. Random traders' wealth levels are on the left axis (log scale), and show a rapid decline as the inference agents capture their wealth. Average inference agents' wealth, shown on the right axis, fluctuates significantly in the first 40 rounds of trading, and begins to stabilise when most of the random traders' wealth has been captured. Smaller jumps in average inference trader wealth later in the trading rounds reflect the transfer of wealth into fewer and fewer inferential agents' control. The type of inference trader that captures most of the wealth appears to fluctuate randomly, and did not show any clear outperformer across many simulations.

In the second simulation with a noisy AR signal being generated, inference technique traders do significantly better than in the presence of only random traders. Each of the 10 individual simulations has a different randomly selected autoregressive model built in to the noisy AR signal traders within it.

As can be seen in figure 2, wealth levels for AR series modelling inference traders are higher, and more stable, for longer, than in the previous simulation. It is difficult to detect much difference between the performance of inference agents embodying non-MML techniques, which show slight variation based upon the sophistication of the information criteria (IC) they use, but in general the more sophisticated ICs outperform the simpler ones. Using the labels from section 2.3, we have outperformance of BIC (AVG_BIC) over HQ (AVG_HQ) over CAIC (AVG_CAIC) over AIC (AVG_AIC). As the trading rounds increase beyond 80, the MML inference traders split into two, with those MML traders using a model assuming a stationary AR time series (AVG_MMLs) falling off to drift towards the IC traders, whilst the MML inference technique that assumes non-stationarity (AVG_MMLns) significantly outperforms. The collapse of wealth levels towards 200 trading rounds represents the almost complete transfer of wealth to the MML agents.

The third simulation regards only the inference traders' interactions with each other, in the absence of any noise traders. The inference traders act like Gaussian noise traders when they cannot detect a signal in the past price series. In this scenario the non-stationary MML inference trader significantly outperformed the other classes of traders. Figure 3 depicts the averages of the average wealth level of the inference agent class across 10 simulations each containing 100 traders over 200 trading rounds. Once again using the agent class labels from section 2.3, we have the MML non-stationary inference agent (AVGavgMMLns) outperforming all other agent classes. The stationary MML agent (AVGavgMML) came next, capturing the second greatest amount of wealth. The sudden drop in the wealth level of the stationary MML agents near to the 170th round reflects this wealth level existing by this point in only one of the ten aggregated simulations, and this particular simulation halting, or taking too long to continue, at this point. The corrected AIC agent class (AVGavgCAIC) also stood out, capturing the third greatest proportion of wealth before the single simulation remaining with this agent outperforming came to a halt. With this scenario halting, or taking too long to continue, only 4 of the original 10 simulations continued running. With the cessation of the stationary MML out-performance scenario, this dropped to three which completed 200 trading rounds. The only other agent class to show any potential in this scenario was the BIC (MDL) agent (AVGavgBIC), which briefly outperformed all but the non-stationary MML agent in early trading rounds (< 27). From that point forward the BIC agent's declining wealth share was bounded above by the decline of the AIC agent class (AVGavgAIC), and from below by the worst performing agents, the HQ inference agents (AVGavgHQ). It is interesting to note the underperformance of HQ agents in this scenario versus even the simple AIC traders. In seven out of ten of the underlying simulations, the MML non-stationary trader was the best performer at the cessation of execution. The other three to outperform were, as mentioned, CAIC, BIC, and stationary MML. The high rate of halting or time overrun in these simulations, which do not explicitly include noise traders, is due to the greater risk of liquidity loss; under-performing traders tend to go broke.

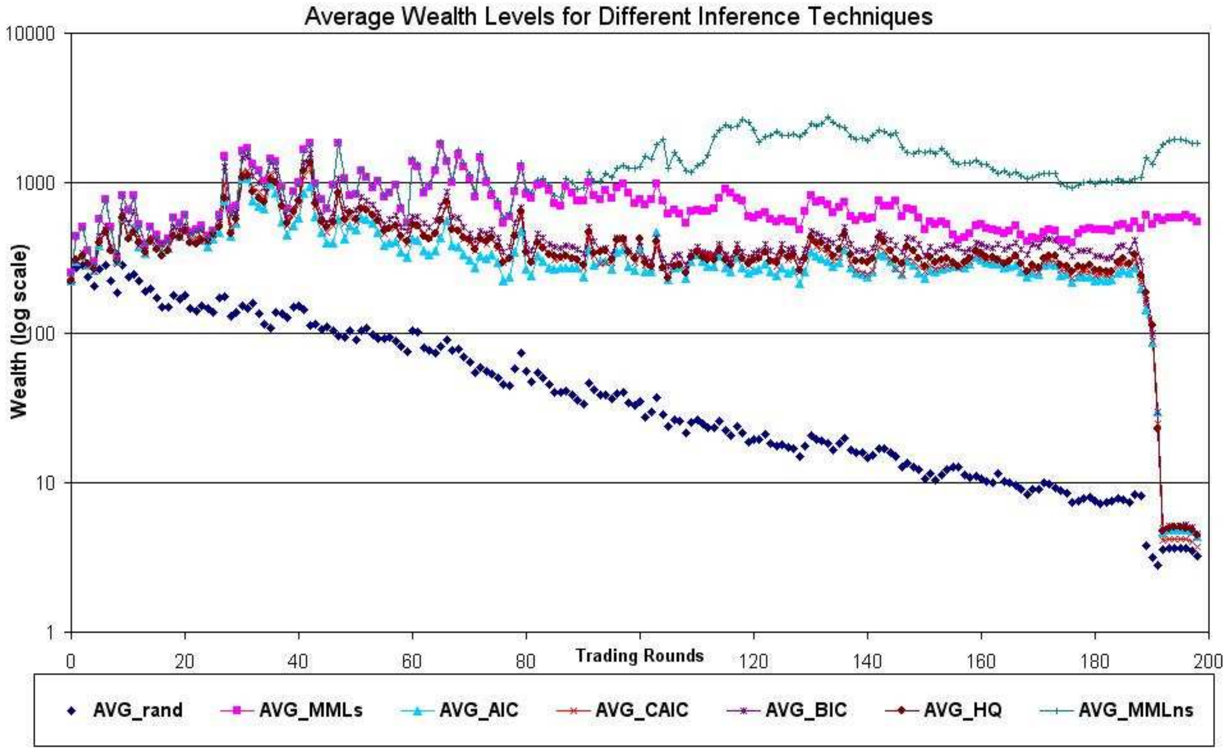

Figure 2. AR noise traders and inference agent wealth levels by type.

5 Conclusions

We have demonstrated a simulated stock market in which trading agents embody a class of different AR time series inference techniques, and shown that in the presence of a noisy signal, MML-inference technique based agents significantly outperform other traders using different inference techniques. We have shown such performance advantages to be persistent as long as there exists a noisy signal to exploit, and that a level of outperformance exists amongst inferential agents even in the absence of an explicit noisy signal. Such types of agent-based hybrid inference technique models appear to be a reasonable technique to apply to real markets, and in future work we intend to combine inferential agents with genetic algorithm based search techniques in building superior simulations of empirical market characteristics, and to demonstrate the ability of such models to perform in real trading environments.

The performance of the MML-based inference techniques here is not surprising given the success of MML in earlier applications by Wallace et al. (e.g., [70, 62, 65, 25, 67, 61, 60]) and Dowe et al. (e.g., [48, 56, 57, 1, 28, 20, 21, 58] and references therein) and, we hope, provides further impetus to greater recognition of this methodology and its relevance to practical statistical inference.

In terms of market efficiency, we see here that an agent using a superior inference technique will consistently outperform a lesser one. Furthermore (recalling section 1, and as stated in [24]), since in most markets we can never prove that our inference technique is superior, we can neither in general establish that there does not exist some (as yet unused) trading technique that will outperform, nor in general that a market is efficient[5] [6].

References

[1] Y. Agusta and D. L. Dowe. Unsupervised learning of correlated multivariate Gaussian mixture models using MML. In *Proceedings of the Australian Conference on Artificial Intelligence 2003, Lecture Notes in Artificial Intelligence (LNAI) 2903*, pages 477–489. Springer-Verlag, 2003.

[2] H. Akaike. Information theory and an extension of the maximum likelihood principle. In B. Petrov and F. Csaki, editors, *Second International Symposium on Information Theory*, pages 267–281, Budapest, 1973. Akademiai Kaidó. Reprinted in [37] on pp. 599-624.

[5] This technical report is an expansion of our recent (August 2005) work [19] to appear in November 2005

[6] This research was partially supported by Australian Research Council (ARC) Discovery Grant DP0343650

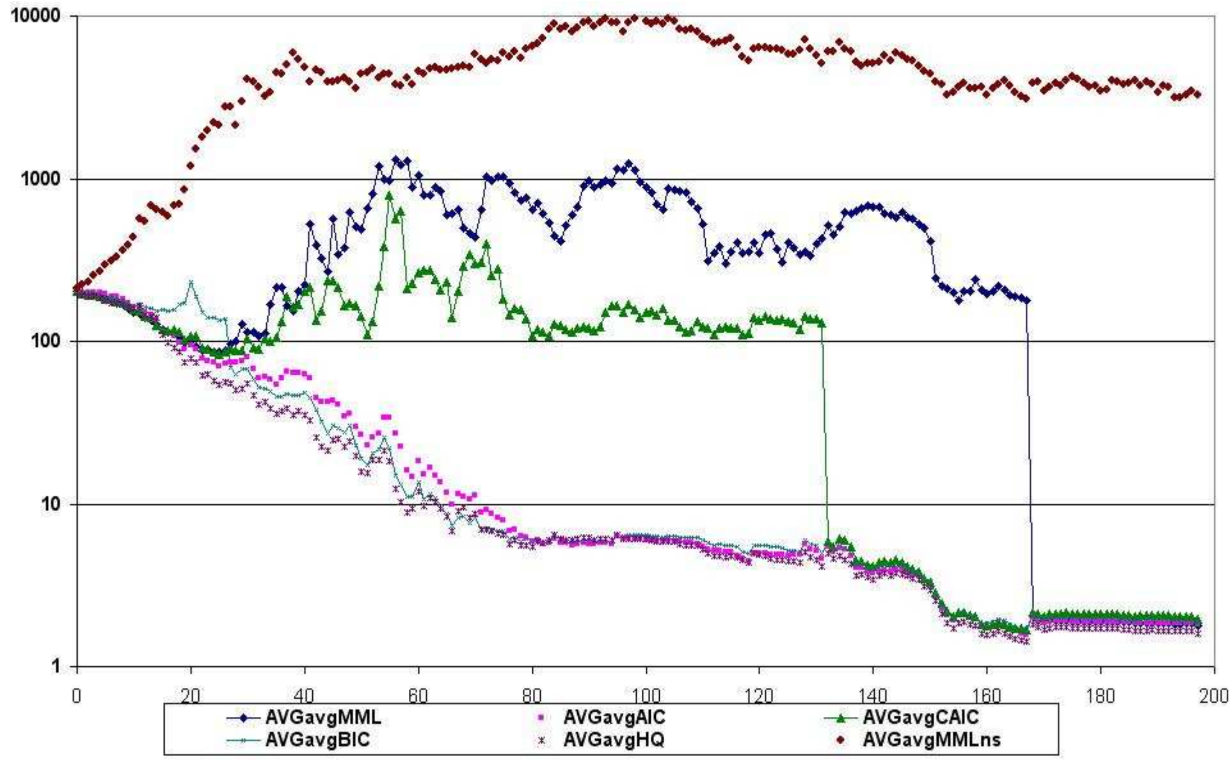

Figure 3. Inference agent wealth levels by type, no noise traders

[3] H. Akaike. A new look at the statistical model identification. *IEEE Transactions On Automatic Control*, 19(6):716–723, December 1974.

[4] W. Arthur, J. Holland, B. LeBaron, R. Palmer, and P. Tayler. Asset pricing under endogenous expectations in an artificial stock market. In W. Arthur, S. Durlauf, and D. Lane, editors, *The economy as an evolving complex system II*, pages 15–44. Addison-Wesley, Reading, MA, 1997.

[5] W. Bagehot. The only game in town. *Financial Analyst Journal*, 22(27):12–14, 1971. (W. Bagehot is a pseudonym for Jack Treynor).

[6] N. Barberis and R. Thaler. A survey of behavioral finance. In G. Constantinides, M. Harris, and R. Stulz, editors, *Handbook of the Economics of Finance*, chapter 18. Elsevier Science, 2003. also NBER Working Paper 9222 (2002).

[7] O. Barndorff-Neilsen and G. Schou. On the parametrization of autoregressive models by partial autocorrelations. *Journal of Multivariate Analysis*, 3:408–419, 1973.

[8] R. A. Baxter and J. J. Oliver. MDL and MML: Similarities and differences. Tech. Report 207, Department of Computer Science, Monash University, 1994. Amended March 1995.

[9] H. Beltran, J. Grammig, and A. J. Menkveld. Understanding limit order book depth: conditioning on trade informativeness. 2004 Latin American Meetings Series Paper 142, Econometric Society, August 2004. available at http://ideas.repec.org/p/ecm/latm04/142.html.

[10] F. Black. Noise. *The Journal of Finance*, 41(3):529–543, July 1986.

[11] L. Bortoli, A. Frino, E. Jarnecic, and D. Johnstone. Limit order book transparency, execution risk and market depth. Paper delivered in 2005 seminar series, 2004. Available online from http://www.uow.edu.au/commerce/econ/seminarseries.html.

[12] G. Box, G. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice-Hall, New Jersey, 1994. 3rd Edition.

[13] H. Bozdogan. Model selection and Akaike's information criterion (AIC): The general theory and its analytical extensions. *Psychometrika*, 52(3):345–370, September 1987.

[14] P. M. T. Broersen and S. de Waele. Empirical time series analysis and maximum likelihood estimation. In *Proceedings of the 2nd IEEE Benelux Signal Processing Symposium (SPS-2000), Hilvarenbeek, The Netherlands, March 23-24*, 2000.

[15] G. J. Chaitin. On the length of programs for computing binary sequences. *Journal of the ACM*, 13:547–569, 1966.

[16] S. Chen and C. Yeh. On the emergent properties of artificial stock markets: the efficient market hypothesis and the rational expectations hypothesis. *Journal of Economic Behavior and Organization*, 49:217–239, 2002.

[17] C. Chiarella, M. Gallegati, R. Leombruni, and A. Palestrini. Asset price dynamics among heterogeneous interacting agents. *Computational Economics*, 22:213–223, 2003.

[18] M. J. Collie. Stock returns distributions: The emergent behaviour of a multi-agent artificial stock market. In M. Mohammadian, editor, *International Con-*

ference on Intelligent Agents, Web Technologies, and Internet Commerce (IAWTIC'03), Vienna, Austria. University of Canberra, February 2003. Available online at http://www.csse.monash.edu.au/∼gjc/Collie03stockreturndistribs.pdf.

[19] M. J. Collie, D. L. Dowe, and L. J. Fitzgibbon. Stock market simulation and inference technique, 2005. Accepted for inclusion in proceedings of the 5th international conference on Hybrid Intelligent Systems (HIS'05), Rio de Janeiro, Brazil, November 6-9, 2005.

[20] J. W. Comley and D. L. Dowe. General Bayesian networks and asymmetric languages. In *Proceedings of the Hawaii International Conference on Statistics and Related Fields, June 5-8, 2003*, 2003.

[21] J. W. Comley and D. L. Dowe. *Minimum Message Length, MDL and Generalised Bayesian Networks with Asymmetric Languages*, chapter 11, pages 265–294. M.I.T. Press, 2005. [Camera ready copy submitted October 2003].

[22] J. Conway and N. Sloane. *Sphere Packings, Lattices and Groups*. Springer-Verlag, New York, 2 edition, 1993.

[23] D. L. Dowe. The contribution of variance to utility functions. In *Econometric Society Australasian Meeting (ESAM) Conference Proceedings*, volume 3, pages 1–6, July 1997.

[24] D. L. Dowe and K. B. Korb. Conceptual difficulties with the efficient market hypothesis: Towards a naturalized economics. In D. L. Dowe, K. Korb, and J. Oliver, editors, *Proceedings: ISIS (Information, Statistics and Induction in Science)*, pages 212–223. World Scientific, 1996.

[25] D. L. Dowe, J. J. Oliver, and C. S. Wallace. MML estimation of the paramters of the spherical Fisher distribution. In *Proceedings of the 7th International Workshop on Algorithmic Learning Theory (ALT'96), Lecture Notes in Artificial Intelligence (LNAI) 1160, Sydney, Australia, 23-25 October*, pages 213–227, 1996.

[26] E. F. Fama. Efficient capital markets: II. *The Journal of Finance*, 46(5):1575–1617, December 1991.

[27] J. D. Farmer and A. W. Lo. Frontiers of finance: Evolution and efficient markets. *Proceedings of the National Academy of Sciences, USA*, 96:9991–9992, August 1999.

[28] L. J. Fitzgibbon, D. L. Dowe, and F. Vahid. Minimum message length autoregressive model order selection. In M. Palanswami, C. C. Sekhar, G. Venayagamoorthy, S. Mohan, and M. K. Ghantasala, editors, *International Conference on Intelligent Sensing and Information Processing (ICISIP)*, pages 439–444, Chennai, India, January 2004.

[29] K. Gode and S. Sunder. Allocative efficiency of markets with zero intelligence traders: Markets as a partial substitute for individual rationality. *The Journal of Political Economy*, 101(1):119–137, 1993.

[30] B. Hanlon and C. Forbes. Model selection for segmented time series from a Bayesian approach to information compression. Department of Econometrics and Business Statisics Working Paper 8, Monash University, 2002.

[31] E. Hannan and B. Quinn. The determination of the order of an autoregression. *Journal of the Royal Society. Series B (Methodological)*, 41(2):190–195, 1979.

[32] J. Holland and J. Miller. Artificial adaptive agents in economic theory. *American Economic Review: Papers and Proceedings*, 81(2):365–370, 1991.

[33] C. Hurvich and C. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76(2):297–307, June 1989.

[34] G. Iori. A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions. *Journal of Economic Behavior and Organization*, 49:269–285, 2002.

[35] M. Jensen. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6:95–102, 1978.

[36] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[37] S. Kotz and N. Johnson, editors. *Breakthroughs in Statistics, Volume 1*. Springer-Verlag, New York, 1992.

[38] A. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1336, November 1985.

[39] A. D. Lanterman. Schwarz, Wallace, and Rissanen: Intertwining themes in theories of model order estimation. *International Statistical Review*, 69(2):185–212, 2001.

[40] B. LeBaron. Building markets with artificial agents: Desired goals, and present techniques. In G. Karakoulas, editor, *Computational Markets*. MIT Press, 1999.

[41] B. LeBaron. A builder's guide to agent-based financial markets. *Quantitative Finance*, 1:254–261, 2001.

[42] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:498–500, February 1999.

[43] B. G. Malkiel. *A Random Walk Down Wall Street*. W.W. Norton and Co., New York, first edition, 1973.

[44] J. M. Marriott and P. Newbold. Bayesian comparison of ARIMA and stationary ARMA models. *International Statistical Review*, 66(3):323–336, December 1998.

[45] M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modelling and Computer Simulation*, 8(1):3–30, January 1998.

[46] S. Mayhew. The impact of derivatives on cash markets: What have we learned? Working paper, Department of Banking and Finance, Terry College of Business, University of Georgia, 2000. available online at www.terry.uga.edu/finance/people/facultyprofiles/mayhew/.

[47] J. F. Monahan. A note on enforcing stationarity in autoregressive-moving average models. *Biometrika*, 71(2):403–404, August 1984.

[48] S. L. Needham and D. L. Dowe. Message length as an effective Ockham's razor in decision tree induction. In *Proceedings of the 8th International Workshop on Artificial Intelligence and Statistics (AI+STATS 2001), Key West, Florida, U.S.A., January 2001*, pages 253–260, 2001.

[49] D. Piccolo. The size of the stationarity and invertibility region of an autoregressive-moving average process. *Journal of Time Series Analysis*, 3(4):245–247, 1982.

[50] M. Raberto, S. Cincotti, S. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:319–327, 2001.

[51] J. Rissanen. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.

[52] J. Rissanen. Hypothesis selection and testing by the MDL principle. *The Computer Journal*, 42(4):260–269, 1999.

[53] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, March 1978.

[54] R. Shiller. From efficient market theory to behavioral finance. Discussion Paper 1385, Cowles Foundation, October 2002.

[55] R. J. Solomonoff. A formal theory of inductive inference, part 1. *Information and Control*, 7(1):1–22, March 1964.

[56] P. J. Tan and D. L. Dowe. MML inference of decision graphs with multi-way joins. In *Proceedings of the 15th Australian Joint Conference on Artificial Intelligence, Canberra, Australia, 2-6 December 2002, published in Lecture Notes in Artificial Intelligence (LNAI) 2557*, pages 131–142. Springer-Verlag, 2002.

[57] P. J. Tan and D. L. Dowe. MML inference of decision graphs with multi-way joins and dynamic attributes. In *Proceedings of the 16th Australian Joint Conference on Artificial Intelligence (AI'03), Perth, Australia, 3-5 December 2003, published in Lecture notes in artificial Intelligence (LNAI) 2903*, pages 269–281. Springer-Verlag, 2003.

[58] P. J. Tan and D. L. Dowe. MML inference of oblique decision trees. In *Proc. 17th Australian Joint Conference on Artificial Intelligence (AI04), Cairns, Qld., Australia*, pages 1082–1088. Springer-Verlag, December 2004.

[59] L. Tesfatsion. Agent-based computational economics. ISU Economics Working Paper 1, Iowa State University, 2003. Revised August 24, 2003.

[60] M. Viswanathan and C. S. Wallace. A note on the comparison of polynomial selection methods. In D. Heckerman and J. Whittaker, editors, *Proceedings of Uncertainty 99: The Seventh International Workshop on Artificial Intelligence and statistics, Fort Lauderdale, Florida, 3-6 January 1999*, pages 169–177, San Francisco, 1999. Morgan Kaufmann.

[61] M. Viswanathan, C. S. Wallace, D. L. Dowe, and K. B. Korb. Finding cutpoints in noisy binary sequences - a revised empirical evaluation. In N. Foo, editor, *Lecture Notes in Artificial Intelligence 1747, Subseries of lecture notes in Computer Science, Advanced Topics in Artificial Intelligence, proceedings of the Twelfth Australian joint Conference on Artificial Intelligence (AI'99), Sydney, Australia, 6-10 December 1999*, pages 405–416. Springer-Verlag, Germany, 1999.

[62] C. S. Wallace. Multiple factor analysis by MML estimation. Technical Report CS 95/218, Department of Computer Science, Monash University, 1995.

[63] C. S. Wallace. *Statistical and Inductive Inference by Minimum Message Length*. Springer, 2005.

[64] C. S. Wallace and D. M. Boulton. An information measure for classification. *The Computer Journal*, 11(2):185–194, August 1968. Available online at http://www3.oup.co.uk/computer_journal/hdb/Volume_11/Issue_02/.

[65] C. S. Wallace and D. L. Dowe. MML estimation of the von Mises concentration parameter. Technical Report CS 93/193, Department of Computer Science, Monash University, 1993.

[66] C. S. Wallace and D. L. Dowe. Minimum message length and Kolmogorov complexity. *Computer Journal*, 42(4):270–283, 1999. [See also replies and rejoinders on pp330-337 and pp345-347].

[67] C. S. Wallace and D. L. Dowe. Refinements of MDL and MML coding. *The Computer Journal*, 42(4):330–337, 1999.

[68] C. S. Wallace and D. L. Dowe. Rejoinder. *The Computer Journal*, 42(4):345–347, 1999.

[69] C. S. Wallace and P. R. Freeman. Estimation and inference by compact coding. *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(3):240–265, 1987.

[70] C. S. Wallace and P. R. Freeman. Single factor analysis by MML estimation. *Journal of the Royal Statistical Society. Series B*, 54(1):195–209, 1992.