The Temptation of Emergence Or Don't Rush Into Economic(al) Explanations

Norman Ehrentreich*
Martin-Luther-Universität Halle-Wittenberg

May 31, 2005

Abstract

This paper discusses how agents in the SFI-ASM are able to acquire more wealth than non-classifier agents, even though Ehrentreich (2005) had established that the classifier system does not provide any useful trading information. Besides learning speed, the number of activated rules and the specific selection mechanisms that agents employ in choosing a trading rule to act upon are identified as additional factors that influence wealth accumulation. It is also shown that higher wealth levels in the SFI-ASM are generally a sign of less efficient model behavior, which is counterintuitive to economic intuition. To avoid similar pitfalls, I finally propose as a dictum for agent-based simulations that emergent behavior should only be the last resort when explaining complicated or surprising model behavior.

JEL Classification: G12; G14; D83

Keywords: Emergence; Classifier Systems; Learning; Asset Pricing; Financial Time Series; Genetic Algorithms

*University of Halle-Wittenberg, Department of Banking and Finance, Universitätsplatz 8–9, 06099 Halle/Saale (Germany), E-mail: ehrentreich@wiwi.uni-halle.de., current phone: +1-612-706-7819 (Minneapolis, MN, USA). The market model was programmed using Java and the RePast-library. Comments are welcome.

1 Introduction

As one of the pioneering stock market simulations, the SFI-ASM was able to generate emergent behavior. For fast learning rates of agents, a complex regime emerged which was characterized by higher price volatility, GARCH-behavior, cross-correlation between price and trading volume, and significant levels of technical trading. In their 1997 paper, Arthur et al. asked themselves to what extent the existence of the complex regime is an artifact of design assumptions in their model. They found by "*varying both the model's parameters and the expectational-learning mechanism, that the complex regime and the qualitative phenomena associated with it are robust. These are not an artifact of some deficiency in the model*" (p. 35).

In Ehrentreich (2005) it was shown that these results were based on the design of the mutation operator which introduced an upward bias in the level of set trading bits in the classifier system. When corrected by a bit-neutral mutation operator, NESFI-agents (**N**orman **E**hrentreich's SFI-agents) always find the correct homogeneous rational expectations (hre) equilibrium of no-bit usage, independent of their learning speed. Emergence of technical trading cannot be observed anymore and complex time series behavior arises only for much higher learning speeds than in the original model.

An analysis wealth levels in both model versions was omitted in Ehrentreich (2005). Attentive readers might have noticed, though, that the endogenous abandoning of the classifier system by NESFI-agents is in apparent contradiction to studies by Joshi et al. (1998) and Joshi et al. (2002) who found that technical trading can generate excess profit. The results of the NESFI-ASM, on the other hand, imply that none of the trading bits improves the forecast accuracy of agent's trading rules. Unless a rule's forecast accuracy is a wrong proxy for its profitability, the convergence at the zero-bit solution should imply that classifier-agents, i.e., those who use the information provided by their classifier system, cannot outperform agents who neglect the classifier system altogether. This is in contradiction to studies by Joshi et al. (1998 and 2002) who found that agents with access to technical trading bits acquire more wealth than fundamental traders.

Unfortunately, things are more complicated than that. Wealth differences between different trader types—technical or fundamental traders, fast or slow learning agents, SFI- or NESFI-agents, classifier or non-classifier agents—

most generally arise. While the explanation of these wealth differences had often been sought in elaborate economic reasoning, this paper identifies the number of active trading rules as the main reason for varying wealth levels.

In order to be self-contained, section 2 will briefly recapitulate the design of the SFI-ASM. This will include a short discussion of its mutation operator and its rectification by Ehrentreich (2005). More details, however, are given in the original contributions by Arthur et al. (1997) or LeBaron et al. (1999). Section 3 will briefly introduce two studies by Joshi et al. (1998 and 2002). In section 4, alternative explanations for wealth divergence will be provided. Based on these findings, a final section will draw some general conclusion for agent-based simulations.

2 The Original SFI-ASM

The SFI-ASM is inhabited by N traders who are all initially endowed with one unit of risky stock and $20,000$ units of cash. During each period, traders have to decide how much to invest in risky stock and how much to keep in cash which yields a risk-free rate of return r_f. The stock pays a stochastic dividend per period which is generated by a stationary AR(1)-process

$$d_{t+1} = \bar{d} + \rho(d_t - \bar{d}) + \epsilon_{t+1}, \tag{1}$$

with $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. Traders have identical constant absolute risk-aversion expected utility functions

$$U(W_{t+1}) = -e^{-\lambda W_{t+1}}, \tag{2}$$

with λ being the degree of risk-aversion and W_{t+1} being an agent's expected wealth level in the next period. Given the budget constraint

$$W_{t+1} = x_t(p_{t+1} + d_{t+1}) + (1 + r_f)(W_t - p_t x_t), \tag{3}$$

the optimal amount of stock $\widehat{x}_t$ that an agent desires to hold is

$$\widehat{x}_t = \frac{E_t[p_{t+1} + d_{t+1}] - p_t(1 + r_f)}{\lambda \sigma^2_{t,p+d}}, \tag{4}$$

where $E_t[p_{t+1}+d_{t+1}]$ is the expectation in t about the next period's realization of the stock's price and dividend, and $\sigma^2_{t,p+d}$ is the empirically observed variance of the combined price plus dividend time series. A specialist balances the effective demands with the fixed supply of shares by setting a market clearing price in an iterative process.

Traders have heterogeneous expectations about future prices and dividends $E_t[p_{t+1} + d_{t+1}]$. These expectations are formed through a classifier system by generating linear forecasts

$$E_t[p_{t+1} + d_{t+1}] = a_j(p_t + d_t) + b_j, \tag{5}$$

with a_j and b_j being real valued parameters constituting the predictor part of a chosen trading rule j. Each rule consists of a condition part, a forecast part (predictor), its fitness value, and its forecast accuracy. The condition parts are checked against a Boolean market descriptor D_t which holds current and past price and dividend information. For example, a particular market state could be that the price of the stock is greater than n-times its fundamental value, while at the same time, the 25-period moving average of the stock price is greater than the current price. When a particular predefined condition is met, the corresponding descriptor bit is set to 1, and otherwise to 0. The condition part is coded as a ternary string holding either 1 or 0, depending on whether the corresponding bit in the market descriptor has to be matched or not, or holding the don't care sign # if the rule ignores that particular descriptor bit. Rules with numerous #-signs are quite general, hence, they will be activated more often than more specific rules. The bits of a trading rule may be characterized as either technical or fundamental. Technical bits check only price or trading volume information, while fundamental bits relate the price of a stock to its fundamental value by using dividend information. For example, dividends and prices are checked to determine whether they have increased or decreased, and whether they are above or below certain moving averages. Most importantly, prices are checked against a stock's fundamental value by comparing for each ratio in the brackets to determine whether

$$\text{price x interest rate/dividend} > \left\{\frac{1}{2}, \frac{3}{4}, \frac{7}{8}, 1, \frac{9}{8}, \frac{5}{4}, \frac{3}{2}\right\} \tag{6}$$

is fulfilled.

From the set of 100 individual trading rules that each agent possesses, normally more than one match the market state D_t and are marked as active. From the set of activated rules, agents either choose the rule with the highest fitness value (select best mechanism), or choose a rule with a probability proportionate to its fitness (select roulette mechanism). Only when no rules match the market descriptor, the forecast parameters a and b are determined as a fitness weighted average of all a_j and b_j with $j = 1 \ldots 100$ (select average mechanism). Rule fitness is defined as

$$f_{t,j} = C - \left(\nu_{t,j}^2 + \texttt{bitCost} \times \texttt{specificity}\right), \tag{7}$$

where $\nu_{t,j}^2$ is rule's j forecast accuracy which is measured as a weighted average of previous and current squared forecasting errors

$$\nu_{t,j}^2 = \left(1 - \frac{1}{\theta}\right) \nu_{t-1,j}^2 + \frac{1}{\theta}\Big[(p_t + d_t) - [a_j(p_{t-1} + d_{t-1}) + b_j]\Big]^2. \tag{8}$$

The parameter θ determines the size of the time window that agents consider when estimating a rule's accuracy.[1] The positive constant C in equation 7 ensures positive fitness, `specificity` refers to the number of non-ignored conditions in a rule which are penalized with some associated `bitCost`. The forecast accuracy $\nu_{t,j}^2$ is used as the variance of the combined price plus dividend time series$\sigma_{t,p+d}^2$ in equation 4.

The set of 100 trading rules is, on average, altered every K periods by a genetic algorithm (GA) which replaces the 20 worst rules with new, possibly better ones.[2] The real valued parameters a and b of the prediction part and the bit strings of the condition part are altered by the genetic operators of either *crossover* (with probability $1 - \Pi$) or *mutation* (with probability $\Pi = 0.7$).

For the condition parts, the SFI-ASM uses uniform crossover. Here, an offspring's bit is chosen with equal probability from the corresponding bit positions of either parent. Note that the fraction of bits set in the offspring is an unweighted average of the two parents' bit fraction. Thus, there is no systematic influence on average specificity through the working of the

[1] As LeBaron et al. (1999) have pointed out, the value of θ is a crucial design question since it strongly affects the speed of accuracy adjustment and the resultant learning in the artificial stock market. If $\theta = 1$, the rules would be judged only on the last period's performance and forecast accuracy would be strongly prone to noise. At the other extreme, however, as θ goes to ∞, agents would take all past information into account, implicitly assuming they live in a static world. As in LeBaron et al., a value of 75 is chosen for θ.

[2] Two useful introductions to genetic algorithms, which were originally developed by Holland (1975), are provided in Goldberg (1989) and Mitchell (1996).

crossover operator. As for the real valued parameters, the crossover operator constructs the new parameter values by determining a weighted average of the two parents' values, with $1/\sigma^2_{j,p+d}$ as the weight for each parent. The weights are normalized to sum up to 1.

The real valued parameters of the predictor parts are mutated by adding random numbers to them. The condition parts are mutated by randomly flipping each bit with a small mutation probability of $\pi = 0.03$. Though unsuspicious at first, the specific bit transition probabilities

$$P = \begin{pmatrix} 0 & 1/3 & 2/3 \\ 1/3 & 0 & 2/3 \\ 1/3 & 1/3 & 1/3 \end{pmatrix} \tag{9}$$

with which individual bits are mutated are crucial for the main result of emergent technical trading for faster learning speeds.[3] Even though LeBaron et al. assert that these transition probabilities would, on average, maintain the number of don't care sings #'s in a rule, we find by applying a Markov chain analysis that, in the long run, the fraction of non-# bits converges to one half.

Because the model usually functions well below the bit-level of one half, the mutation operator introduces an upward tendency in the bit distribution. Invoking the mutation operator more often per time period results in a higher equilibrium bit level, prompting the researchers from the Santa Fe Institute to infer that there is emergent technical trading. Because of the cost they have attached to every non-# bit, they conjectured that emerging trading bits must have, on average, some fitness-based advantages by producing more accurate forecasts.[4] The Markov chain analysis suggests that this emergence

[3]The transition matrix is interpreted as follows: 0- or 1-bits are never left unchanged and are converted into the do not care sign # with a probability of two thirds. After mutation, an initial do not care bit will be either $1, 0$, or # with an equal probability of one third.

[4]Arthur et al. (1997) emphasize that positive bit costs allow them to speak of "emergence" since "*the information represented by a particular bit is used only if agents find it genuinely useful in prediction*" (p. 34). There are several indications that this is a premature conclusion. First, their logic implies that the SFI-model could be forced into a non-bit (or at least a low-bit) usage solution if bit costs are sufficiently high. This, however, is never the case. Second, no significant differences in the level of fundamental and technical trading bits can be detected for various learning speeds or mutation rates. Last, but not least, one cannot assert that agents act upon technical trading bits simply because they

is simply a design artifact of the mutation operator. In Ehrentreich (2005), an alternative bit-neutral mutation operator is suggested, i.e., one that leaves the probability of non-# bits before and after mutation unchanged. The transition matrices, one for fundamental bits

$$P_{fund.} = \begin{pmatrix} 0 & F_{fund.} & 1 - F_{fund.} \\ F_{fund.} & 0 & 1 - F_{fund.} \\ \frac{1}{2}F_{fund.} & \frac{1}{2}F_{fund.} & 1 - F_{fund.} \end{pmatrix} \tag{10}$$

and one for technical bits defined in the same way, are characterized by dynamically adjusting transition probabilities where $F_{fund.}$ refers to the initial fraction of fundamental 0- or 1-bits.

With this new mutation operator, agents endogenously abandon the use of their classifier system. In the long run, the fraction of set trading bits in each agent's rule set always approaches zero. This result confirms the prediction of the hre-equilibrium, given the mean reverting dividend process. Furthermore, the simulated time series are generally closer to the hree-benchmark than in the original SFI-ASM. Only at much faster learning rates than in the original model can the claim of emergent complex price series behavior still be upheld.

3 Wealth Levels in the SFI-ASM: An Economic(al) Explanation

Since agents in the rectified NESFI-ASM endogenously give up the use of their classifier system, it is strongly implied that it does not provide any profitable trading information. If technical trading rules would generate excess profit, technical trading should strive rather than vanish.[5]

The conclusion of a useless classifier system is in direct contradiction to Joshi et al. (1998) who find that agents with access to technical trading

exist. In fact, when counting an agent's rule usage, one realizes that general rules are selected for use much more often than more specific ones.

[5]Strictly speaking, the trading rules are not evaluated according to their wealth generating ability, but on their forecast accuracy. While it is possible that a more accurate trading rule is less successful in terms of generated wealth, Arthur et al. (1997) believed this scenario to be highly unlikely. Their intuition was confirmed by Wilpert (2004) who evaluated trading rules according to their generated profits. His results did not differ much from those of the original SFI-ASM.

bits acquire more wealth than fundamental agents. The wealth differences that they report are not negligible and grow over time, as figure 1 vividly illustrates.

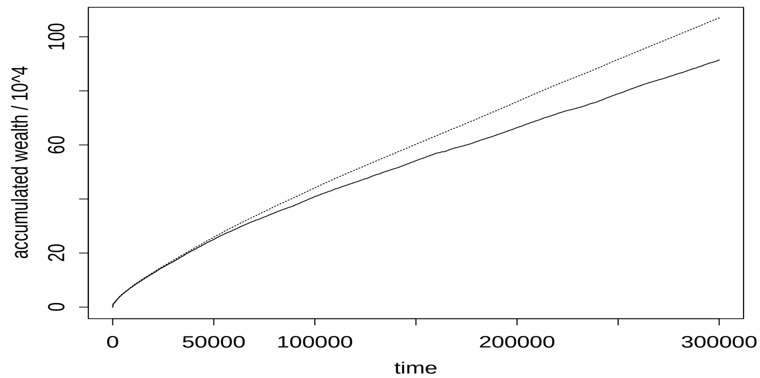

Fig. 1: Wealth levels for the situation when one agent includes technical rules while all others exclude them. Note that the singular agent using technical rules accumulates significantly more wealth than those agents using only fundamental rules almost all through the run, and that this difference grows over time. Source: Joshi et al (1998), p. 11

The observed differences in wealth accumulation between fundamental and technical traders were used by Joshi et al. for a game theoretic analysis. They concluded that technical trading creates a typical prisoner's dilemma. As long as there are fundamental traders in the market, technical traders can acquire more wealth, yet when all agents include technical trading rules, the overall market wealth is less than if everybody excludes them. In a related study, Joshi et al. (2002) vary their focus and determine an agent's optimal learning speed. They find that the unique symmetric Nash-equilibrium implies a fast learning rate that clearly falls into the complex regime. Hence, they conclude that financial markets can operate at sub-optimal equilibria.

The two studies by Joshi et al., however, simply take the observed wealth differences as given and deduce economic interpretations from them. As to where these wealth differences come from, they just offer some vague economic rationalization. They argue, for instance, that if a single agent with access to technical trading bits detects a short term price trend, he might be able to exploit this pattern without dissipating it. If more and more agents detect the price trend, the particular technical trading rules are reinforced through positive feedback, thus making them self-fulfilling prophecies, which can cause bubbles and crashes. Similar arguments on the emergence and profitability of technical trading can be found in Arthur et al. (1997) or Joshi and Bedau (1998).

It is clear from Ehrentreich (2005) that this explanation based on pattern detection and exploitation cannot be upheld. Since NESFI-agents voluntarily forfeit the possible gains from technical trading by giving up their classifier systems, it is implied that the patterns they detected are not profitable in the long run. Thus, the question remains why technical traders accumulate different wealth levels than fundamental traders?[6]

4 Wealth Levels in the (NE)SFI-ASM: Alternative Explanations

Fundamental traders in Joshi et al. (1998) had no access to technical trading bits in their trading rules, yet they were still using the fundamental trading information. NESFI-agents in the corrected stock market version, on the other hand, endogenously give up checking technical and fundamental trading information. It thus seems straightforward to shorten the whole adjustment process and to introduce non-classifier agents, i.e., agents who do not check any environmental information at all and have trading rules that only consist of prediction parts and fitness information. How would these agents compare when they have to compete with old-fashioned SFI-classifier agents?

It was my initial hypothesis was that non-classifier agents should never outperform classifier-agents since only the latter were equipped with an additional tool to analyze and exploit potentially useful information. Figure 2,

[6] Initially, I even doubted the validity of Joshi et al.'s (1998) simulation results. In Ehrentreich (2002), a working paper version of Ehrentreich (2005), I reported that wealth levels of classifier and non-classifier agents rise equally after some initial divergence during the warm-up phase, which is in contradiction to figure 1 with diverging wealth slopes. I then noticed, however, that while the absolute wealth differences were tiny, the classifier-agents did slightly better in 23 out of 25 simulation runs. A paired two sided t-test then revealed that these minuscule wealth differences were significant, sometimes even at the 1% percent level. It was this outperformance of the classifier agents that prompted me to look for alternative explanations how wealth differences in the SFI-ASM come about.

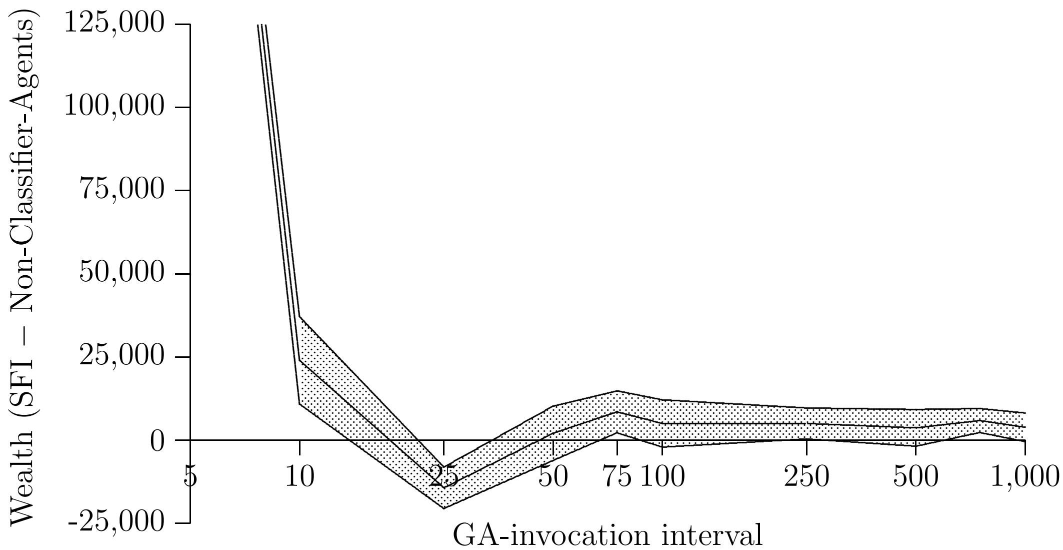

Fig. 2: Wealth differences between SFI-agents and non-classifier agents, both trader types using select best for rule selection, recorded for different GA invocation intervals and averaged over 25 simulation runs. Positive values indicate that SFI-agents outperform the non-classifier agents, the shaded area indicates mean wealth difference ± one standard deviation.

however, tells a different story. While for most GA-invocation intervals the classifier-agents acquire more wealth, there is a significant dip in the curve when non-classifier agents do better than their SFI-counterparts. Better performing non-classifier agents, however, point to reasons other than pattern recognition and exploitation which determine accumulated wealth levels.

4.1 Rule Selection from Activated Rules

First of all, when attempting to explain the wealth differences one stumbles upon an important difference between SFI- and non-classifier agents. For SFI-agents, the mutation operator increases the number of set bits in trading rules. An increase in the learning speed thus results in more specific and more illogical rules, thereby systematically reducing the number of activated rules. The size of the activated rule set for non-classifier agents, on the other hand, is largely independent of the GA-invocation interval. For both agent types, trading rules that had been activated less than "mincount" times since their creation, are not included in the active rule set either. The mincount-parameter was set at a value of five and thus tends to decrease the number of activated rules only for very fast learning speeds.

Secondly, when further investigating wealth performance of different trader types I realized that the model versions used in Joshi et al. (1998 and 2002) differ in one detail from the version used by LeBaron et al. (1999). From the set of activated rules, an agent must choose one for a price and dividend forecast. In LeBaron et al. (1999), always the best, i.e., the rule with the lowest variance estimate from equation 8 is selected. In Joshi et al. (1998 and 2002), however, a rule is randomly chosen from among the active rule set with a probability proportional to it's fitness value. This mechanism is known as roulette wheel selection since a rule's selection probability could be conceptualized as the size of the wedge on a biased roulette wheel.

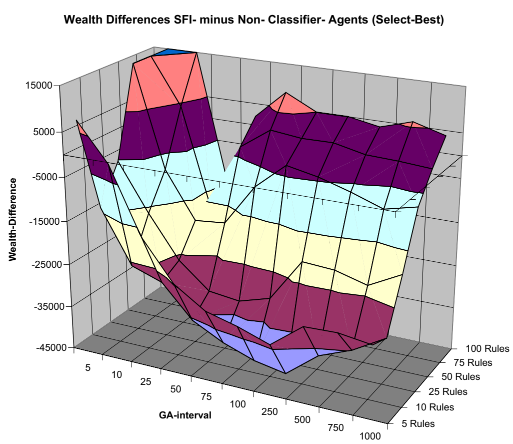

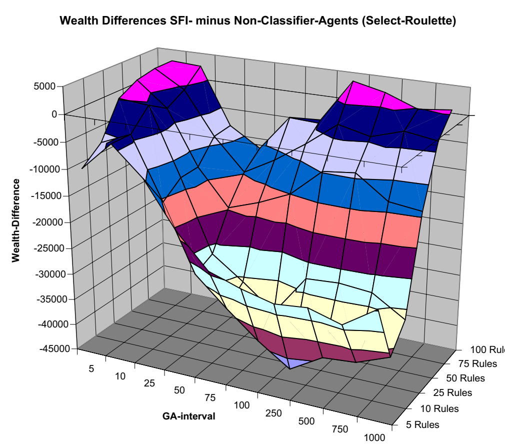

Fig. 3: Wealth differences between SFI- and non-classifier agents as a function of learning speed, number of trading rules they possess, and rule selection mechanism (left: select-best, right: select-roulette). Data were averaged over 25 simulation runs. Note that non-classifier agents outperform the SFI-agents for most parameter combinations (negative wealth differences) and that the peak in the upper left corner for select-best was truncated for better comparability.

Considering the increasing gap between active trading rules for SFI- and non-classifier agents for faster learning speeds, and taking into account the different selection mechanisms used so far, it seemed reasonable to look at wealth levels for different rule set sizes under both selection procedures. Figure 3 thus gives a more complete impression of wealth behavior than figure 2 does. First one notices that wealth differences between SFI- and non-classifier agents do not obey any simple relationship when changing either the learning speed or the rule set size. Both directions have pronounced “wealth valleys”, thus indicating at least two or more overlapping factors that influence an agent’s wealth accumulation. Secondly, while the general shape in both wealth graphs seems similar, there are some noticeable differences such as a peak in the upper left corner for select-best.

While figure 3 does only report wealth differences between SFI- and non-classifier agents, the same analysis had been made for absolute wealth levels in markets with only one agent type (top sections in figures 4 & 5). When only looking at the wealth differences between SFI-and non-classifier agents in figure 3, one is inclined to say that using either select-best or roulette wheel for rule selection does not matter very much. Yet when looking at the absolute wealth levels in figure 4 and 5, a completely different wealth behavior between the two selection mechanisms becomes apparent.

I attribute this difference to a widely known problem with simple roulette wheel selection, i.e., that of scaling invariance. Adding an offset to all fitness values tends to equalize the selection probabilities in roulette wheel selection. The constant C in equation 7 with a value of 100, added to each raw fitness value intended to avoid negative fitness, effectively leads to almost uniform selection probabilities. The fitness information is thus widely neglected, especially in large active rule sets.[7] Given the problem of scaling invariance, select-best seems a better choice for rule selection in the SFI-ASM than simple roulette wheel selection.

The middle and bottom sections in figures $4-5$ point to another possible reason why wealth levels differ for different trader types and learning speeds. Even when told to use select-best or roulette wheel selection for forecast production, SFI-agents may not be able to do so and employ certain fall back methods. By increasing the learning speed for SFI-agents or by reducing the number of rules they possess, the size of their activated rule set shrinks (middle section in figure 4). When there are no active rules at all, SFI-agents derive their forecast parameters as a fitness weighted average of all the rules

[7] Goldberg (1989) and Mitchell (1996) report various approaches to deal with scaling invariance. Sigma scaling, for instance, remedies the problem by taking into account also the mean and standard deviation of fitness values in a population. Other approaches are to use different selection procedures such as ranking-based schemes, select-best being just the simplest to implement.

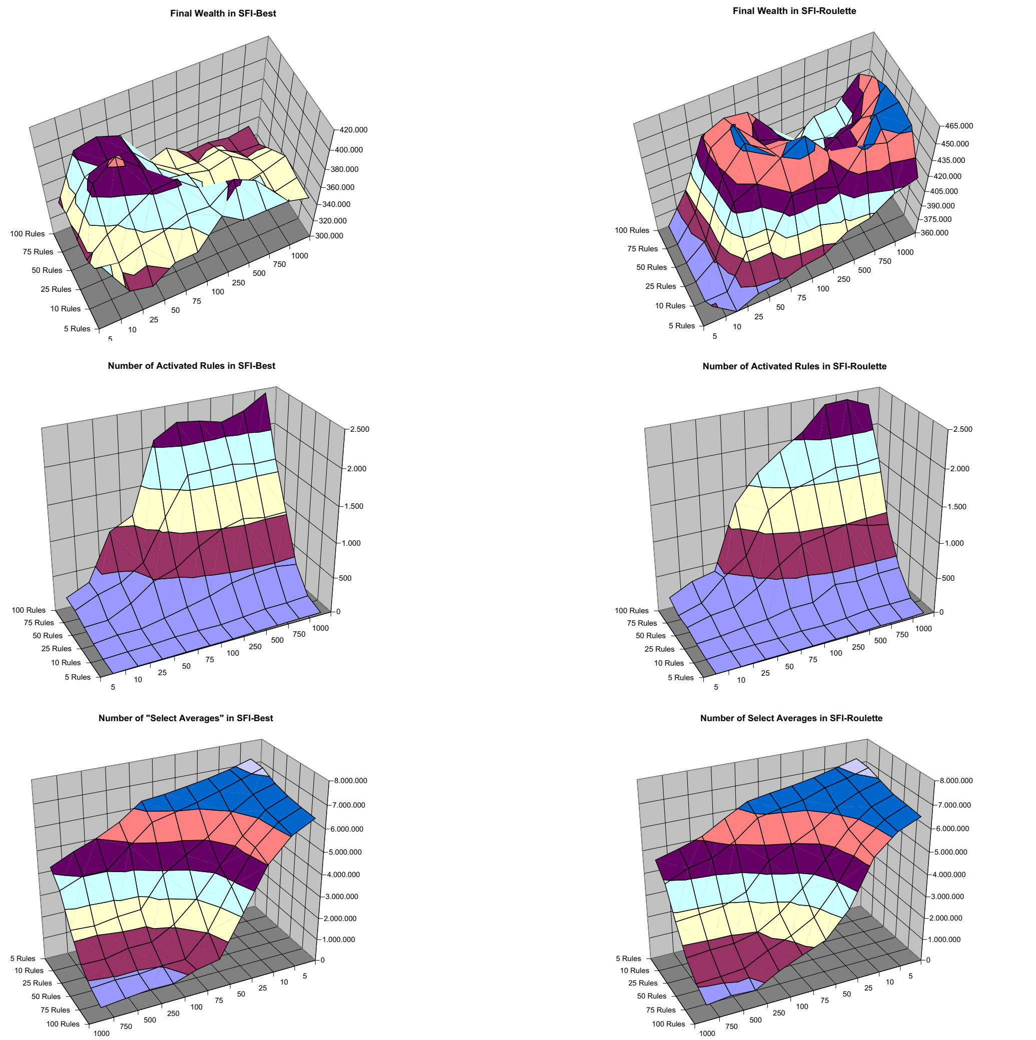

Fig. 4: Top: Final wealth levels for SFI-agents with select-best (left) and roulette-wheel selection (right). Middle: Number of activated rules. Bottom: Number of "select-averages" during simulation. Data were averaged over 10 simulation runs.

that have been activated at least mincount-times in their lifetime. If none of the rules had been activated at least mincount times, SFI-agents resort to use the global average of past prices and dividends as their forecast. Dependent on the parameterization, the relative importance of each method varies greatly, whereby each method alone leads to a different model behavior. In the following, I subsume both fall back mechanism in the SFI-ASM as *select-*

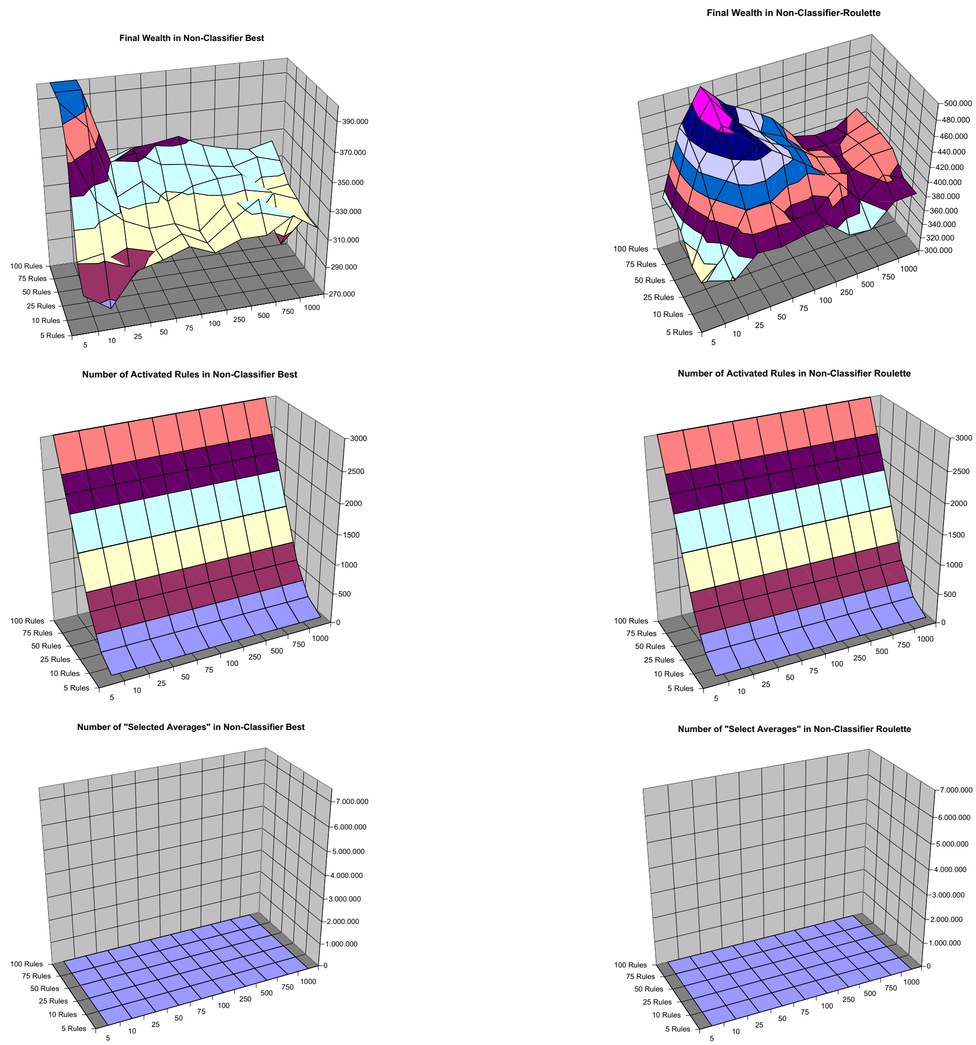

Fig. 5: Top: Final wealth levels for non-classifier agents with select-best (left) and roulette-wheel selection (right). Middle: Number of activated rules. Bottom: Number of "select-averages" during simulation. Data were averaged over 10 simulation runs. Note that the wealth peak in the upper left corner for select best, reaching as high as 1.7 million, was truncated for better visibility.

average. The bottom section in figure 4 shows how SFI-agents increasingly resort to select-average when the number of activated rules becomes smaller.

Non-classifier agents, on the other hand, are not plagued by the problem of a diminishing active rule set. Their rules are always activated (middle section of figure 5) and hence there is no need for them to resort to any of the fall back methods (bottom section). Even though non-classifier indeed stick to the intended selection mechanism, we still see a wealth peak for GA-intervals around 50 and a valley at about 500 for roulette wheel selection. These remaining wealth differences most likely reflect the working of the GA which changes the forecast parameters, yet other still unidentified influences are possible, too.

Separating and attributing price and wealth effects to each individual factor seems impossible, though. Changing one factor affects the price series which in turn will affect the GA, which in turn will affect the triggering change. Because of these interdependencies, an exhaustive explanation of price and wealth behavior cannot be given. Earlier results in Ehrentreich (2005) have implied that the wealth differences reported by Joshi et al. (1998 and 2002) cannot be attributed to pattern recognition and exploitation. By identifying other factors that influence wealth behavior, the puzzle of diverging wealth levels for different trader types can be at least partly solved.

4.2 Risk-Premium, Taxation, and Wealth Levels

When looking at the absolute wealth levels in figures 4 and 5, an economically trained mind easily spots some potential for optimization. For instance, Joshi et al. (1998) asked themselves whether the observed wealth differences could be helpful in finding an agent's optimal learning speed. Or the designer of an artificial stock market might ask whether there is an optimal number of rules agents should use.

Because of the specific design of the SFI-ASM, however, wealth levels are somewhat counterintuitive to interpret. The hre-equilibrium as the benchmark towards which the model behavior should optimally converge is, for instance, usually characterized by lower wealth levels than when the model is run in non-hree mode. Gulyás et al. (2003) pointed out that in the SFI-ASM, agents usually increase their wealth more or less independent of their actions. In the hre-equilibrium, agents are basically inactive, holding on to their one unit of stock, collecting its dividend, and receiving interest on their cash holdings. To avoid a long-run explosion of wealth levels, wealth in the SFI-ASM is taxed at a rate equal to the exogenous interest rate. When the model is run in hree-mode, it is possible to approximate a theoretical hree-base wealth level W_t^{hree} in period t as

$$W_t^{hree} = C_0 + t\left(\overline{d} - r_f\overline{p^{hree}}\right), \tag{11}$$

r_f being the risk free interest rate, $\overline{d}$ the theoretical dividend mean of the dividend process, and C_0 as agent's initial cash endowment.[8] The constant $\overline{p^{hree}}$ is the theoretical hree-price average which is a linear function of agent's risk aversion, the risk free interest rate, the theoretical dividend mean, speed of mean reversion of the dividend process, and the variance of the dividend noise process.[9] Since risk aversion leads to $\overline{p^{hree}} < \overline{d}/r_f$, equation 11 implies that hree-base wealth W_t^{hree} grows linearly with t.

In a real simulation run, however, stock prices diverge from their theoretical hree-levels. Base-wealth is defined as the wealth of an inactive agent when his one unit of stock is evaluated with the realized prices instead of hree-prices.[10] It is determined as $W_t = C_t + p_t$ with C_t as base cash in t, determined as

$$C_t = C_{t-1} + (d_t - r_f p_t). \tag{12}$$

It is important to realize at this point that base wealth usually exceeds hree-base wealth since the long term average of realized prices p_t is smaller than the theoretical hree-price average. This reflects an additional risk premium that is attached to realized prices since the constant learning of agents introduces some additional noise. Equation 12 also implies that an increase in the risk premium due to higher volatility, reflected in smaller values for p_t, will result in higher wealth levels in the SFI-ASM. It is easy to see that wealth would be maximized stock prices would be zero. An efficient market outcome in the SFI-ASM thus should not be equated with high wealth levels,

[8]The derivation of equation 11 is given in the appendix. The approximation allows us to determine a hree-base wealth level without an actual simulation run. If a particular simulation is run in hree-mode, the exact hree base wealth in t would be $W_t^{hree} = C_0 + t(\overline{d} - r_f\overline{p^{hree}}) + p_t^{hree}$, yet $\overline{d}$ and $\overline{p^{hree}}$ would represent the empirical averages (instead of theoretical averages). Since equation 11 considers only cash, we would have to add the current hree-price of the stock, given the current dividend. In the long run, however, the averages converge to their theoretical means and the stock value becomes negligible in relation to cash holdings.

[9]The exact formula of hree-prices is given in LeBaron et al. (1999) or Ehrentreich (2005). The model parameters in use resulted in a hree-price of 87.92 and a hree-base wealth of 322,096 after 250,000 periods.

[10]I owe this idea of base wealth to Gulyás et al. (2003).

in the contrary, the hree-benchmark equilibrium could almost be considered as a lower bound for non-hree wealth levels. Even though this seems quite counterintuitive, it is merely a result of taxation of last period's wealth in the model, while dividends are collected tax-free.

Wealth levels much higher than the corresponding hree-base wealth are thus usually accompanied by rather extreme model behavior, high volatility, and realized prices far away from equilibrium prices. Remember that non-classifier agents with select-best, with many trading rules, and extremely high learning speed acquired excessively high wealth levels (figure 5). The price series accompanied by this, however, do not resemble any real stock prices at all. Because of its specific design, an isolated look at wealth levels in the SFI-ASM can lead researchers easily astray.

5 Conclusion

This paper has dealt with the puzzle how SFI-agents are able to acquire more wealth than non-classifier agents, even though Ehrentreich (2005) had established that the classifier system does not provide any useful trading information. It has been shown that wealth levels in the SFI-ASM are not only tied to learning speed, but also to the number of activated rules and to the specific selection mechanisms that agents employ in choosing a trading rule to act upon. These last factors represent in my opinion more programming technicalities than any economic content. It has also been established that efficient market outcomes close to the hre-equilibrium benchmark are characterized by lower wealth levels than when the market exhibits complex price behavior.

It seems to me that the lesson to be learned from this discussion could be of general interest to the agent-based simulation community. On the one hand, this paper highlights the importance of seemingly inconspicuous programming details such as the appropriate number of genetic individuals or the choice of a specific selection operator. On a more general level, it cautions from premature declarations of emergent behavior or the necessity to produce intuitive economic explanations of model behavior. Even as researchers, it is hard not to be fooled by our confirmation and observer biases. The SFI-ASM was designed to allow for emergent technical trading, and the first signs of it were accepted as proof. The reasons for different wealth performances were immediately sought in elaborate economic arguments such as exploited technical trading patterns. To avoid similar pitfalls, I would therefore like to close with positing a dictum for agent-based models: *Try to link complicated or surprising model behavior to the known model structure. Only if all attempts fail, resort to emergence as a possible explanation.*

Appendix

The cash and wealth positions of agents in the SFI-ASM are determined according to

$$\begin{aligned} W_t &= C_{t-1} + p_t x_t && (13) \\ C_t &= C_{t-1} + r_f C_{t-1} + x_t d_t && (14) \\ C_t &= C_{t-1} - r_f W_t, && (15) \end{aligned}$$

x_t being the amount of stock held by them. Equation 13 denotes the previous wealth before adjustment, equation 14 takes interest and dividend earnings into account, while equation 15 lowers cash through tax payments at a tax rate equal to the risk free interest rate. Taxation is incorporated in the model to avoid an explosive wealth behavior in the long run through compound interest effects. Equations $13 - 15$ can be summarized as

$$C_t = C_{t-1} + x_t \left(d_t - r_f p_t \right). \quad (16)$$

Since base wealth was defined as the wealth of inactive agents, their stock position x_t remain at their initial endowments of one unit. Equation 16 simplifies to

$$C_t = C_{t-1} + (d_t - r_f p_t) \quad (17)$$

and base wealth is thus $W_t = C_t + p_t$. An approximation for hree-base wealth can be obtained by substituting the simulated prices and dividends in period t by their theoretical averages $\overline{p^{hree}}$ and $\overline{d}$. Since these are constants, the recursive relationship of equation 17 can be written as a function of initial cash endowment

$$C_t = C_0 + t \left(\overline{d} - r_f \overline{p^{hree}} \right). \quad (18)$$

References

Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R. and Tayler, P. (1997). Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. N. Durlauf and D. Lane (eds), *The Economy as an Evolving Complex System II*, Vol. 27 of *SFI Studies in the Sciences of Complexity*, Santa Fe Institute, Addison-Wesley, pp. 15–44.

Ehrentreich, N. (2002). The Santa Fe Artificial Stock Market reexamined: Suggested corrections, *Betriebswirtschaftliche Diskussionsbeiträge 45/02*, Martin Luther University Halle-Wittenberg.

Ehrentreich, N. (2005). A corrected version of the Santa Fe Institute Artificial Stock Market model, *Journal of Economic Behavior and Organization*, (forthcoming).

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA.

Gulyás, L., Adamcsek, B. and Kiss, A. (2003). An early agent-based stock market: Replication and participation, *Proceedings 4th Soft-Computing for Economics and Finance Meeting (NEU 2003), May 29, 2003, University Ca' Foscari, Venice, Italy.*
*http://omega.ailab.sztaki.hu/%7Egulya/papers/NEU2003.zip

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, University of Michigan Press, Ann Arbor, MI.

Joshi, S. and Bedau, M. A. (1998). An explanation of generic behavior in an evolving financial market, *Santa Fe Institute Working Paper* (98-12-114).
*http://www.santafe.edu/research/publications/workingpapers/98-12-114.pdf

Joshi, S., Parker, J. and Bedau, M. A. (1998). Technical trading creates a prisoner's dilemma: results from an agent-based model, *Santa Fe Institute Working Paper* (98-12-115).

Joshi, S., Parker, J. and Bedau, M. A. (2002). Financial markets can be at sub-optimal equilibria, *Computational Economics* **19**: 5–23.

LeBaron, B., Arthur, W. B. and Palmer, R. (1999). Time series properties of an artificial stock market, *Journal of Economic Dynamics and Control* **23**: 1487–1516.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*, MIT Press, Cambridge, MA.

Wilpert, M. (2004). *Künstliche Aktienmarktmodelle auf Basis von Classifier-Systems*, Knapp, Frankfurt/M.