Survival of the Chartist: An Evolutionary Agent-Based Analysis of Stock Market Trading

Daan Bloembergen[1], Daniel Hennes[2], Simon Parsons[1] and Karl Tuyls[1]
[1]University of Liverpool, United Kingdom
{d.bloembergen, s.d.parsons, k.tuyls}@liverpool.ac.uk
[2]European Space Agency, Noordwijk, The Netherlands
daniel.hennes@esa.int

ABSTRACT

A stock market is a highly complex dynamical system. Stock-price movements are not solely driven by fundamental values but in particular influenced by short term trading behaviour. Chartists use trends to forecast future price directions, whereas fundamentalists estimate stock prices based on dividend payouts or company earnings. Such strategies can similarly be deployed in automatic trading agents, which already account for a large portion of current trading activity. It is therefore vital to understand how these trading strategies behave in different scenarios, and how the interplay of strategies may lead to various market outcomes. In this paper we analyse the dynamics of three different trading strategies: fundamentalist, chartist, and zero-information traders, who base their trading behaviour on the current market price only. We simulate stock markets with various constellations of trading agents, and compare their evolutionary strength. Our results show that it is not straightforward to predict in advance which trading strategy will perform best. Fundamentalists outperform other traders, and drive them out of the market, when information is freely available. If fundamental information is costly, chartists may thrive, potentially destabilising the market.

1. INTRODUCTION

Markets play a central role in today's society and find wide application ranging from stock markets to consumer-to-consumer e-commerce [1, 2]. Achieving accurate market forecasts decides over trading success or losses. There are two main types of trading strategies in todays markets: fundamentalists and chartists [8, 19]. Fundamentalists use a forecasting model that fits the actual economy and correctly identify the fundamental driving forces of the market. Chartists, also called technical analysts, use an autoregressive process to predict future price developments based on recent trends.

One might be tempted to conjecture that fundamentalists eventually drive chartists out of the market. After all, chartists try to exploit an autocorrelation structure in the price series which in turn is mainly a result of their own trading behaviour – not an underlying feature of the market. Rational fundamentalists must surely be superior as they base trading decisions on actual fundamental facts.

However, fundamentalists are not strictly rational. Future fundamental values (e.g. earnings or dividends) of a company are not known at present time and must be predicted using a model. The model must match the economy that drives the market and model parameters must be adjusted accordingly. A mismatch in model choice, or uncertainty in parameter estimates that deviate from the ones that determine the underlying process inevitably cause bounded rationality and thus the risk for false decisions.

In this paper, we study a microscopic market model with traders of three different types: fundamentalists, chartists, and zero-information traders. Zero-information traders use no external information and trade around the current market price. We analyse static markets in which traders follow a pre-set trading strategy, thereby gaining insights into the relative performance of the different strategies, and we check the resulting price series for stylised facts of real financial time series, which indicate the validity of our model. Building on these insights, we then investigate a dynamic market in which traders may switch to a more profitable trading strategy at any time. In particular, we investigate whether chartists can survive in a market and coexist with fundamentalists, or if indeed chartists are eventually driven out of the market.

2. BACKGROUND

Some background information is needed for the remainder of this paper, which we present here. Firstly, we describe auctions and their application in stock markets. Secondly, we discuss evolutionary game theory and in particular the replicator dynamics which form the basis of our analysis. Finally, we introduce heuristic payoff tables, as means of capturing the relative evolutionary strength of high level trading strategies.

2.1 Auctions

Auctions are highly efficient match making mechanisms for trading goods or services. Various sets of rules exists to conduct an auction, yielding different transaction volumes, transaction delays, or allocative market efficiency. Here, we focus on double auctions, which essentially provide a platform for buyers and sellers to meet and exchange a commodity against money. A taxonomy of double auctions especially tailored to automated mechanism design can be found in [15]. Double auctions maintain an open book of bids (offers to buy at a specified price) and asks (offers to sell at a specified price). Two principle forms are the clearing house auction and continuous operation auction. In a clearing house auction, orders are collected for a trading period (e.g., one day) and matched, or cleared, after the trading period is closed. This mode of operation allows for high allocative efficiency, but incurs delays in the transactions.

In contrast, continuous operation immediately establishes a transaction as soon as any trader is willing to buy at the ask price. This mode allows higher transaction rates at the cost of some allocative efficiency. Experiments in this article will use continuous operation mode, since it reflects the day-time operation mode of many stock markets, such as the NYSE [1].

2.2 Evolutionary game theory

Auctions provide a dynamic environment with a lot of traders (agents) that adapt to each other while competing for revenue. Learning in such multi-agent systems is generally complex and poses many challenges that inspire prescriptive, descriptive and normative research [17]. Evolutionary game theory provides a methodology to analyse multi-agent learning, replacing assumptions from game theory like rationality by evolutionary concepts such as pressure of natural selection [22].

The evolutionary perspective considers a population of individuals, where each individual belongs to one of several species. These species generally relate to atomic strategies, or to information levels within this article. Two core concepts are the *replicator dynamics*, describing how a population evolves, and *evolutionarily stable states*. The replicator dynamics formally define the population change $\dot{\boldsymbol{x}}$ over time, where $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ describes the distribution of n species in the population:

$$\dot{x}_i = x_i \left[f_i(\boldsymbol{x}) - \sum_j x_j f_j(\boldsymbol{x}) \right] \quad (1)$$

The payoff function $f_i(\boldsymbol{x})$ can be interpreted as the Darwinian fitness of each species i. Intuitively, (1) describes how species that do better than average in the population thrive, whereas species that do worse decline. Evolutionarily stable state are such population distributions $\boldsymbol{x}$ that are fixed points of the replicator dynamics, i.e., $\dot{\boldsymbol{x}} = 0$, and where small perturbations $|\hat{\boldsymbol{x}} - \boldsymbol{x}| < \epsilon$ would be driven back to $\boldsymbol{x}$ by selection pressure, i.e., by following the replicator dynamics.

Previous research has demonstrated the viability of evolutionary game theory to analyse meta strategies in simulated auctions [11, 13], and to compare clearing house against continuous double auctions [16]. Here, we follow a similar analysis procedure, building on and extending the market model of Hennes *et al.* [11] to include chartists, as is described in Section 3.

2.3 Heuristic payoff tables

The evolutionary model assumes an infinite population. We cannot compute the payoff for such a population directly, but we can approximate it from evaluations of a finite population. All possible distributions over k trading strategies can be enumerated for a finite population of n individuals. Let N be a matrix, where each row N_i contains one discrete distribution. The matrix will yield $\binom{n+k-1}{n}$ rows. Each distribution over trading strategies can be simulated using the market model, described next in Section 3, returning a vector of average expected relative market revenues $u(N_i)$. Let U be a matrix which captures the revenues corresponding to the rows in N, i.e., $U_i = u(N_i)$. A heuristic payoff table $H = (N, U)$ is proposed by Walsh *et al.* to capture the payoff information for all possible discrete distributions in a finite population [23].

In order to approximate the payoff for an arbitrary mix of strategies in an infinite population distributed over the species according to $\boldsymbol{x}$, n individuals are drawn randomly from the infinite distribution. The probability for selecting a specific row N_i can be computed from $\boldsymbol{x}$ and N_i:

$$P(N_i|x) = \binom{n}{N_{i1}, N_{i2}, \ldots, N_{ik}} \prod_{j=1}^{k} x_j^{N_{ij}}$$

The expected payoff $f_i(\boldsymbol{x})$ is computed as the weighted combination of the payoffs given in all rows:

$$f_i(\boldsymbol{x}) = \frac{\sum_j P(N_j|x) U_{ji}}{1 - (1 - x_i)^k}$$

If a discrete distribution features zero traders of a certain information type, its payoffs cannot be measured and $U_{ji} = 0$. This expected payoff can be used in (1) to compute the evolutionary population change according to the replicator dynamics.

3. MARKET MODEL

Our market model is based on a continuous double auction with open order book, in which all traders can place bids and asks for shares. We closely follow the market model as described by Tóth and Scalas [20], Tóth *et al.* [21], Huber *et al.* [12], and Hennes *et al.* [11] in order to be comparable. In the following we firstly describe the daily operation of the market. We then discuss the value of information, following the dividend discount model, and the trading strategies derived from this model. Finally, we present the different noise and cost functions used in the experiments.

3.1 Market Operation

The current value of a share is inherently determined by the revenue that one is expected to gain from holding the share in the future. In our model, these revenues come from dividends that are paid out regularly based on the number of shares owned at that point in time. The stream of dividends follows a Brownian motion random walk (as in e.g. [20, 12]), given by:

$$D_t = D_{t-1} + \epsilon$$

where D_t denotes the dividend in period t, with $D_0 = 0.2$, and ϵ is a normally distributed random term with $\mu = 0$ and $\sigma = 0.01$, i.e., $\epsilon \sim \mathcal{N}(\mu, \sigma^2)$.

We simulate the market over 30 trading periods, each lasting $10 \cdot n$ time steps, where n is the number of traders present. This ensures that all traders have ample opportunity to trade within each period. All traders start with 1600 units cash and 40 shares, each worth 40 initially. At the beginning of each period, all traders put an initial bid or ask in the book (opening call). Hereafter, at every time step a trader is selected at random who can then either accept an open order, or place a new bid or ask, according to his trading strategy (described below). When an order is accepted, the two traders involved exchange one share (the seller) in return for the asked price (the buyer). At the end of each period, dividend is paid based on the shares owned, and a risk free interest rate (0.1%) is paid over cash. The performance of the traders is measured as their total wealth after the 30 periods, i.e., the sum of their cash and share holdings, where each share is valued according to the discounted future dividends (see below).

3.2 Dividend Discount Model

The *dividend discount model* is based on the theory that the intrinsic present value of a share is based on the discounted sum of its future dividend payments. As such, computing the current value of a share relies on a good estimate of these future dividends, which is the core of the fundamentalist trading strategy. The most widely used equation to compute this value is *Gordon's growth model* [10, 9]. The model assumes that we require a certain rate of return $r > 0$ on our investment. For example, if $r = 0.005$, a share must return 0.5% per trading period for it to be a worthwhile investment. This rate r is also called the *discount rate*. Gordon's growth model assumes that dividends grow at a constant rate g. If D_0 is the current dividend payout, the current stock value can be computed as follows:

$$V = \sum_{t=1}^{\infty} D_0 \frac{(1+g)^t}{(1+r)^t} = D_0 \frac{1+g}{r-g}$$

Let us assume the dividends are constant over time, i.e. $\forall i : D_t = D$ and $g = 0$. The stock value simplifies to:

$$V = \sum_{t=1}^{\infty} \frac{D}{(1+r)^t} \tag{2}$$

The infinite series of (2) converges to D/r as $1/|1+r| < 0$ with $r > 0$. For example, a stock that pays a constant dividend of 0.2 per share has a current value of $V = 0.2/0.005 = 40$. This logic underlies the starting price of 40 for each share, given initial dividend value $D_0 = 0.2$.

Differently informed traders can be implemented by varying the amount of foresight knowledge that they have about future dividends. Note that this applies to fundamentalists only; chartists rely on past data only, which is readily available. In trading period $t = k$, we say that fundamentalists of the first information level, F_1, know only the dividend D_k, and in general traders of information level j, labelled F_j, know $D_k, \ldots, D_{k+j-1}$. Therefore, the discounted dividend payoff that is guaranteed for traders with information level F_j is

$$\sum_{i=0}^{j-1} \frac{D_{k+i}}{(1+r)^i}$$

and the future discounted dividends for $t > k + j - 1$ are estimated according to (2) with a constant $D = D_{k+j-1}$:

$$\sum_{t=k+j-1}^{\infty} \frac{D_{k+j-1}}{(1+r)^t} = \frac{D_{k+j-1}}{r} \tag{3}$$

As (3) estimates future discounted dividends from period $t = k+j-1$ on, (3) itself must be discounted by $\frac{1}{(1+r)^{j-1}}$ to adjust payouts to current value prices. The complete stock value estimate for trader F_j is thus:

$$E(V|F_j, k) = \sum_{i=0}^{j-1} \frac{D_{k+i}}{(1+r)^i} + \frac{D_{k+j-1}}{r(1+r)^{j-1}} \tag{4}$$

To put it intuitively, a trader of information level F_j knows j future dividends and assumes dividends stay fixed from that point on. This results in a cumulative information structure, where insiders know at least as much as averagely informed traders.

Algorithm 1 Fundamentalist trading strategy

$pv \leftarrow E(V|\mathrm{F}_j, k)$ according to (4)
if $pv < bestBid$ **then**
    $acceptOrder(bestBid)$
else if $pv > bestAsk$ **then**
    $acceptOrder(bestAsk)$
else
    $\Delta_{ask} = bestAsk - pv$
    $\Delta_{bid} = pv - bestBid$
    if $\Delta_{ask} > \Delta_{bid}$ **then**
        $placeAsk(pv + 0.25 \cdot \Delta_{bid} \cdot \mathcal{N}(0,1))$
    else
        $placeBid(pv + 0.25 \cdot \Delta_{ask} \cdot \mathcal{N}(0,1))$
    end if
end if

3.3 Trading strategies

We use three different trading strategies in our experiments. *Fundamentalist* use their knowledge of future dividends to estimate the current value of the stock and base their trading decision on that estimate. *Chartists* look for trends in the past market prices, and traders without any information use the *zero-information* strategy that only takes the current market price of the shares into account.

3.3.1 Fundamentalists

Fundamentalists completely rely on the information they receive. The fundamentalist strategy is detailed in Algorithm 1 (see also [20]). In essence, the traders compare their estimated present value $pv = E(V|\mathrm{F}_j, k)$, as given by (4), with the current best bid and ask in the book. If they find a bid (ask) with a higher (lower) value than their estimate, they accept the offer. Otherwise, they place a new order between the current best bid and ask prices. Naturally, the trader should own enough shares or cash to accept or place an order.

The cumulative information structure described by the dividend discount model allows to compare fundamentalists with different amounts of foresight knowledge. In the experiments presented in this paper we use two types of fundamental strategies: averagely informed traders with information level 3, and fundamentalists with information level 9. These are chosen as they accurately reflect the dynamics of differently informed traders [11]

3.3.2 Chartists

Chartists analyse past trading prices, and look for trends. If they see an upward trend in the market price, they see this as an opportunity to buy; if the trend goes down, they sell. The algorithm used in this work is summarised in Algorithm 2 (see also [20]). The traders look only at the differences between the four last prices. If each of these differences is positive, the chartist expects an upward trend and is willing to buy at a slightly higher price (lines 2 and 3). If the differences are negative, the expectation is a downward trend, and similarly the chartist will try to sell at a slightly lower price (lines 9 and 10). If no trend can be observed, the chartist places a new order in the book.

3.3.3 Zero-Information Traders

The zero-information trading strategy only takes the current market price into account when deciding whether to

Algorithm 2 Chartist trading strategy

```
pv ← P_k {current market price}
if P_{k-3} < P_{k-2} ∧ P_{k-2} < P_{k-1} ∧ P_{k-1} < P_k then
    pv ← pv + |N(0,1)|
    if pv > bestAsk then
        acceptOrder(bestAsk)
    else
        placeBid(pv)
    end if
else if P_{k-3} > P_{k-2} ∧ P_{k-2} > P_{k-1} ∧ P_{k-1} > P_k then
    pv ← pv − |N(0,1)|
    if pv < bestBid then
        acceptOrder(bestBid)
    else
        placeAsk(pv)
    end if
else
    Δ_ask = bestAsk − pv
    Δ_bid = pv − bestBid
    if Δ_ask > Δ_bid then
        placeAsk(pv + 0.25 · Δ_bid · N(0,1))
    else
        placeBid(pv + 0.25 · Δ_ask · N(0,1))
    end if
end if
```

accept or place an order. These traders simply trade randomly around the current market price. Specifically, zero-information traders use the fundamentalist strategy given in Algorithm 1, with the difference that line 1 is replaced by $pv \leftarrow P_k$, where P_k is the current market price at time k. The reasoning behind this strategy is based on the *efficient market hypothesis*, which states that the all available information is reflected in the market price [14]. As such, trading around the market price could be a safe choice. The zero-information strategy serves as a base line for both fundamentalists and chartists. Fundamentalists with no foresight (information level 0) trade following this strategy, as do chartists when they do not observe any trend in the market price.

3.4 Cost and noise

So far we have taken the foresight knowledge of fundamentalists as a given. Their information level determines how many trading periods they can look ahead, but other than that the information is reliable and comes for free. In reality, acquiring such information might be costly, and the data itself may be uncertain. For example, limited foresight knowledge might be obtained by reading financial news letters and company statements, whereas a detailed long-term outlook requires hiring experts. Similarly, short-term foresight might be more reliable than long-term estimates. In order to model these effects we introduce various cost and noise functions, which we then use in the experiments in order to investigate their effect on the fundamentalists' performance. Chartists and zero-information traders rely on current and past market prices only, which we assume to be freely available and reliable.

We use three different cost functions in our experiments, as shown in Figure 1a. The fixed cost function assumes that each fundamentalist pays the same fixed amount per trading period, regardless of their information level. We can

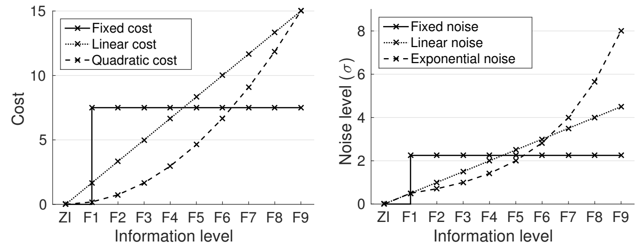

(a) Cost functions **(b)** Noise functions

Figure 1: Different cost and noise functions used in the experiments.

also assume that traders have to pay for each additional bit of information, yielding a linear cost function. Finally, the quadratic cost function is based on the idea that it gets increasingly difficult to obtain more information. In each case, traders pay the required cost at the end of each trading period.

Similarly, we employ three different types of noise functions to model uncertainty in forecasting data, depicted in Figure 1b. Noise is added to each trader's value estimate $E(V)$ when executing Algorithm 1, drawn randomly from a normal distribution. In particular, line 1 of Algorithm 1 is replaced by $pv \leftarrow E(V|F_j, k) + \mathcal{N}(0, \sigma)$, with sigma given as in Figure 1b. In the case of fixed noise, each trader experiences the same level of uncertainty. More realistically, the uncertainty increases with the amount of forecasting, especially when e.g. step-by-step prediction is used [4]. This inspires the exponential noise function. Again, a linear function is used as well as compromise between these two.

4. SIMULATIONS OF A STATIC MARKET

Two types of experiments are conducted to highlight the effect of both cost and noise on the relative return for the different trading algorithms. First, in this section the distribution of trading algorithms in the market is kept fixed, allowing us to investigate the market returns given the various cost and noise scenarios. Moreover, we take a detailed look at some of the characteristics of the resulting price series. Next, in Section 5, traders are allowed to switch their strategy if this is profitable. Evolutionary analysis of the resulting dynamical system indicates which strategy is strongest from a natural selection point of view. Moreover, this analysis shows how the market evolves, and which strategy or set of strategies are economically viable in the long run under different scenarios.

4.1 When is Charting Profitable?

We will now investigate the role of chartists in the market. In particular, we are interested in the viability of the chartist strategy in a market that is essentially dictated by a random process (the dividend stream). After all, chartists rely solely on the presence of trends in the market price. We simulate a market with four types of traders: zero-information (ZI), fundamentalist with information levels 3 and 9 (F_3 and F_9), and chartists (C); 10 traders are used for each strategy. Again, we consider different scenarios using the cost and noise functions of Figures 1a and 1b, and evaluate the relative return for each trading strategy, averaged over 100 sessions of 100 runs each.

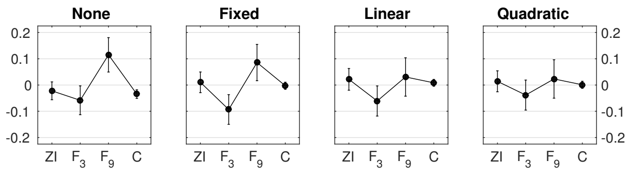

(a) Different cost functions.

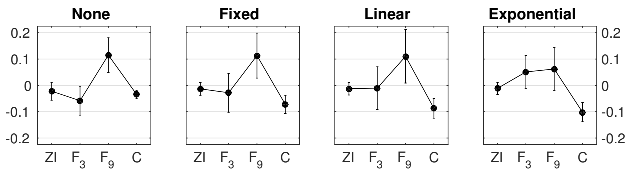

(b) Different noise functions.

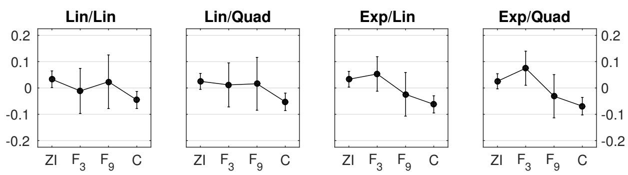

(c) Combination of noise and cost.

Figure 2: Relative returns for a market with a mix of trading strategies; $n = 40$, where 10 traders use each strategy.

The results are presented in Figure 2. Without noise and cost, we see that only highly informed fundamentalists (F_9) perform above market average, whereas in particular the averagely informed traders (F_3) lose out. Zero-information traders and chartists perform around the market average. This effect has been observed before, both in simulation and in human experiments [11, 12, 20]. The addition of cost has an equalising effect on the traders' performance (Figure 2a). As fundamentalists get increasingly charged, zero-information traders and chartists can both profit. The reason is that, as argued before, costs are only subtracted at the end of each trading period, and as such they do not influence the price dynamics of the market. In contrast, noise affects the price dynamics of the market through the fundamentalists' trading behaviour, indirectly influencing the performance of chartists as well (Figure 2b). In this case, increasing noise levels cause the market to be even less predictable, as short-term trends caused by the fundamentalists' foresight knowledge are scrambled. This causes a significant decrease in the chartists' performance. Finally, the combination of cost and noise causes their effects to add up (Figure 2c). The combination of linear noise with either linear or quadratic costs (the left two panels) leads to a situation where zero-information traders achieve the highest performance; in the other cases averagely informed traders do best.

It is clear from these results that both cost and noise can have a large influence on the relative performance of various trading strategies. We analyse these effects in more detail in Section 5, where we apply an evolutionary model to study the dynamics of a market in which traders may switch to more profitable strategies. As such, we can analyse which set of strategies is in equilibrium, and which strategies will die out in the long run under evolutionary pressure.

Table 1: Descriptive statistics for the time series of returns generated by the market model, compared to real-world data of the S&P500 index.

Data	Skewness	Kurtosis
S&P500 index futures	-0.40	15.95
Market without chartists	-0.45	20.47
Market with chartists	-0.44	24.57

4.2 Checking for Stylised Facts

Time series data of many real-world financial assets share a set of common stylised statistical facts [5]. In particular, the distribution of returns is characterised by a heavy tail; there is no significant autocorrelation of returns except for very small (intra-day) time scales; and large price fluctuations tend to be grouped together (volatility clustering). In the following, we check the price series generated by our market model for these stylised facts, focusing in particular on the effect that chartists may have on these measures.

We analyse the price signal resulting from one individual run of the market, both with and without chartists, using the same dividend stream for both scenarios.[1] Specifically, we record the realised prices at every buy and sell action, yielding 4526 price points for the market without chartists, and 5980 price points for the market including chartists. From these data we obtain the log return series r as

$$r(i) = \frac{\log(P_{i+1}) - \log(P_i)}{\Delta t_{i \to i+1}}$$

where P_i is the i^{th} realised price, and $\Delta t_{i \to i+1}$ represents the time difference between two consecutive price realisations. Visual inspection indicates that the distribution of returns resembles a normal distribution, but has a thinner body and bigger tails. In order to verify this, we compute the skewness and kurtosis of the returns. The skewness measures the anti-symmetry in the distribution of a random variable x as

$$\gamma = \frac{E\left[(x-\mu)^3\right]}{\sigma^3}$$

A skewness of 0 means that the data is perfectly symmetrical around its mean. The kurtosis describes the 'peakedness' of the distribution, computed as

$$\kappa = \frac{E\left[(x-\mu)^4\right]}{\sigma^4} - 3$$

where a positive value of κ indicates a heavy tailed distribution; in particular, a normal distribution has $\kappa = 0$. We compare these statistics to those of the S&P500 index futures as reported by [5], in Table 1. The data shows that the distribution of returns generated by our market model is indeed characterised by a heavy tail.

Finally, we look at the autocorrelation of the returns. The lack of significant linear correlations in asset returns has been widely studied for many years [5, 7]. It can be argued that this is a self-correcting property of any market, since the existence of any dependencies in price series would be exploited by some traders, who by that act effectively erase those dependencies [6]. Moreover, a negative autocorrelation of the return series of consecutive transaction prices may be

[1]Similar results are obtained for different dividend streams and simulation trials.

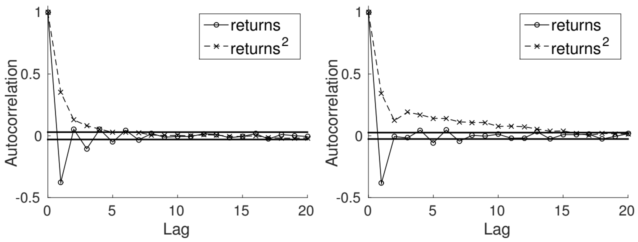

(a) Market without chartists. **(b)** Market with chartists.

Figure 3: Autocorrelation of returns (solid line) and squared returns (dashed line). The straight lines indicate the confidence bounds.

observed, caused by the alternation between buy and sell actions close to ask and bid prices, respectively [5]. Additionally, the autocorrelation of squared returns is a quantitative signature of volatility clustering: the idea that large price variations are often followed by large price variations [5].

The autocorrelation of a discrete time series is measured at different lags τ, indicating the time difference between samples for which the correlation is tested [3]. Figure 3 shows the autocorrelation of returns and squared returns for the market with and without chartists. Without chartists, both measures diminish quickly. When chartists are present, on the other hand, the autocorrelation of squared returns diminishes slowly, indicating some volatility clustering in this case. Therefore, the presence of more diverse trading strategies seems to yield more realistic price dynamics. However, note that these findings are based on a single scenario – more extensive analysis is required to make a strong claim. In sum, the market model generates price series that exhibit several stylised facts of real-world financial time series data, supporting the validity of our model.

5. EVOLUTIONARY DYNAMICS

So far, traders have not been able to choose their strategy. Instead, the distribution over trading strategies was kept fixed throughout the simulation. Realistically it can be assumed that traders may be inclined to switch strategy if they are currently performing poorly. In the following we look at the dynamics of a market in which traders are free to change their strategy at any time. Based on the replicator dynamics of evolutionary game theory we visually inspect the dynamics of such markets. This analysis gives insight into the evolutionary strength of various trading strategies, and the fixed points of the dynamics predict the distribution of trading strategies that may be found in a market in equilibrium.

5.1 Visual Inspection

Experiments in Section 4 revealed that chartists only stand a chance when little or no noise is present in the market; in any other case they are outperformed by all other trading strategies. Therefore, we focus our evolutionary analysis on the cost scenarios only, as these provide the most interesting point of analysis. We simulate a market with four trading strategies: zero-information (ZI), fundamentalists of types F_3 and F_9, and chartists (C). We follow the procedure described in Section 2.3 to compute a heuristic payoff table, using 24 traders, distributed over those four strategies. This yields 2925 different discrete permutations in the heuristic payoff table.[2] For each permutation, the relative performance of the involved strategies is estimated by simulating the market for 100 sessions of 10 runs each.

As we have four strategies, visual inspection of the resulting dynamics in the four-dimensional simplex is not straightforward. However, we can get some insights by looking at the different faces of the simplex, which represent those scenarios in which one strategy is absent. Figure 4 shows the four different faces of the simplex, corresponding to the strategy sets {ZI, F_3, F_9}, {ZI, F_3, C}, {ZI, F_9, C}, and {F_3, F_9, C}, for the four different cost scenarios. When no costs are incurred (Figure 4a), both chartists and zero-information traders are consistently outperformed by fundamentalists, in line with findings reported above. Interesting however is the mixed equilibrium that appears where F_3 and F_9 traders coexist. When fixed costs apply (Figure 4b), this mixed equilibrium becomes unstable; instead a stable attractor appears where chartists and highly informed fundamentalists prevail. Averagely informed traders (F_3) incur relatively large costs in this scenario, and are driven out of the market.

The linear and quadratic cost functions yield the most complex dynamics, where all faces of the simplex are qualitatively different, as shown in Figures 4c and 4d. Of particular interest is the striking similarity between the third and fourth panel of both cost scenarios, corresponding to the strategy sets {ZI, F_9, C} and {F_3, F_9, C}. In fact, the face {ZI, F_9, C} is identical under linear and quadratic costs, as the cost for F_9 traders is the same under both functions. Most importantly, chartists can survive in the market in each cost scenario, with the exception of the {ZI, F_3, C} face under quadratic costs, where the averagely informed fundamentalists perform best. In general, however, it can be concluded that the trading behaviour of fundamentalists indirectly reveals their foresight knowledge through the market price, and chartists are able to profit without having to pay the price.

5.2 Numerical Analysis of Stable Attractors

Although informative in many ways, looking only at the faces of the simplex does not reveal the dynamics of the full mix of strategies. The linear and quadratic cost cases in particular warrant further investigation, as each face of their corresponding simplex shows a mixed equilibrium. This raises the question whether a fully mixed internal equilibrium, where each strategy prevails, is present as well. Although visual inspection is difficult, we can locate attracting equilibria numerically by following traces of the dynamical model, starting from different points in the strategy space. This can be done systematically by selecting the starting points from a four-dimensional uniform grid. Specifically, we select each point $\boldsymbol{x} = (x_1, x_2, x_3, x_4)$ such that the individual components $x_i \in \{0.1, 0.2, \ldots, 0.9\}$, constrained by $\sum_i x_i = 1$ to ensure a valid probability distribution. The results are reported in Table 2.

As anticipated, the linear and quadratic cost scenarios indeed give rise to an internal equilibrium where all four trading strategies prevail. Moreover, in the case of quadratic costs the averagely informed traders (F_3) do slightly better than under linear costs, whereas chartists do worse. This

[2]Increasing this number does not change the results significantly – 24 traders are sufficient to capture all relevant dynamics.

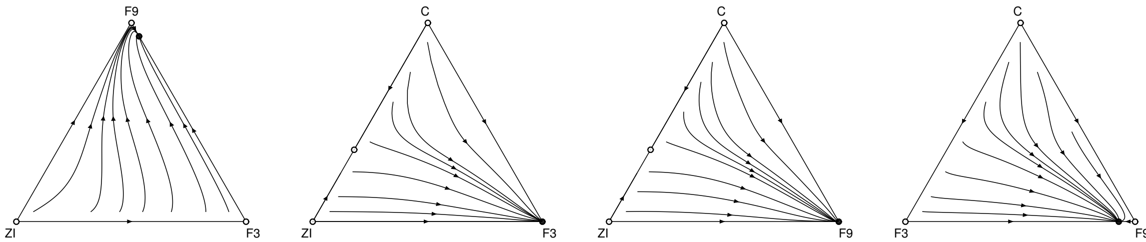

(a) No cost

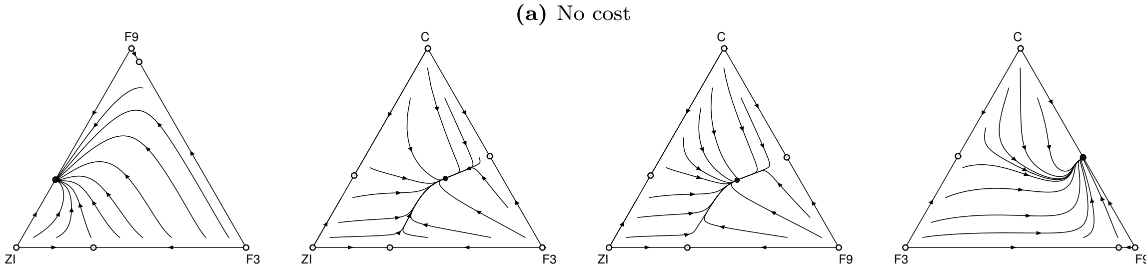

(b) Fixed cost

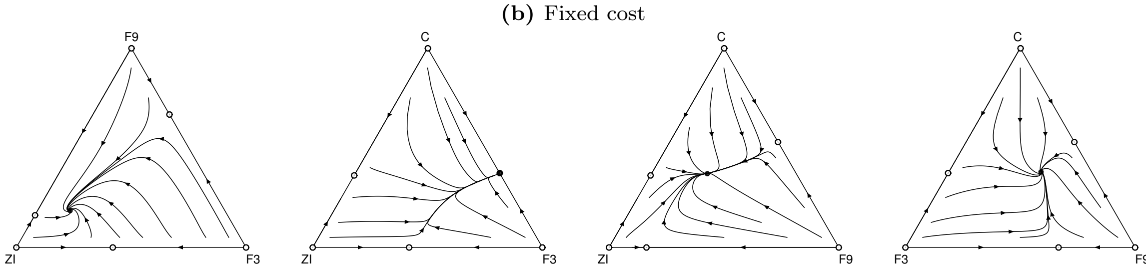

(c) Linear cost

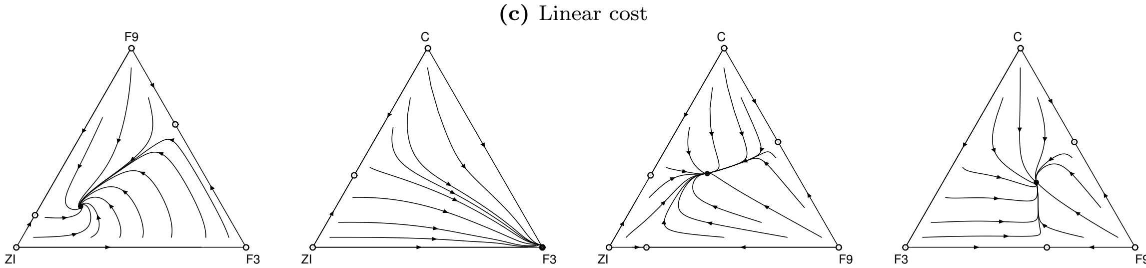

(d) Quadratic cost

Figure 4: Vector fields showing different faces of the four-dimensional simplex for a market with four trading strategies, and different cost functions. Stable attractors are indicated with ● and unstable attractors with ○.

Table 2: Stable equilibria of the four-dimensional simplex.

Cost function	Equilibrium (ZI, F_3, F_9, C)
No cost	(0.00, 0.07, 0.93, 0.00)
Fixed cost	(0.27, 0.00, 0.39, 0.34)
Linear cost	(0.36, 0.10, 0.21, 0.33)
Quadratic cost	(0.35, 0.15, 0.22, 0.28)

finding agrees with the second panel of Figures 4c and 4d, which shows the changing balance between F_3 and C in direct comparison. In the fixed cost scenario, the F_3 strategy dies out in equilibrium, as these traders are disproportionately taxed for their knowledge. Finally, without any cost, both chartists and zero-information traders disappear. Surprisingly, a small fraction of less-informed F_3 traders remains; further research is required to determine whether this is an artifact of the heuristic payoff table, as very little information is available close to the corner points.

6. CONCLUSIONS

In this paper we have employed the evolutionary model of replicator dynamics to analyse the complex strategic interactions of stock market trading. We use an agent-based market model, which includes three different trading strategies: zero-information traders, fundamentalists, and chartists. We investigate the effectiveness of these strategies in our market model, and find that fundamentalists consistently beat the market when information is freely available. More realistically, acquiring fundamental information takes time and effort, which should be taken into account. We find that when information comes at a price, zero-information traders and chartists can survive, leading to a market equilibrium where each trading strategy is present. In such scenarios, chartists can profit from the information that is present in the market price due to fundamentalists' trading behaviour, without having to pay the price. We check the resulting market prices for stylised facts of real-world financial time series, and find that our market model indeed yields realistic price series, which exhibit a fat tailed returns distribution, lack of autocorrelation of returns, and volatility clustering.

Our findings highlight the limits of fundamental trading, and show that it is possible for chartists to survive in the market. We observe a variety of outcomes depending on whether or not cost and noise are a driving factor of the market. This shows that a good understanding of the underlying dynamics are vital if any reasonable predictions about market outcomes are to be made.

Many interesting directions for future research can be identified. Taking a microscopic look at the trading actions in the market equilibrium may give insights to the strengths and weaknesses of the different trading strategies. For example, one might investigate when traders make or lose most of their money, by comparing limit orders and market orders for the different trading strategies, as done by e.g. Stöckl and Kirchler [18] for a market with only fundamental traders. Additionally, the market model can be extended to include various assets, which may yield more complex dynamics. Also, markets do not typically consist of a fixed set of traders. Instead, traders may continuously enter and exit the market, potentially shifting the equilibrium. This may similarly yield more complex dynamics, which may further help to explain the diverse set of traders usually found in real-world stock markets. Finally, comparing the evolutionary dynamics of strategy change against the behaviour of real traders would further validate our model.

REFERENCES

[1] J. J. Angel. Market mechanics: A guide to U.S. stock markets. Technical report, The Nasdaq Stock Market Educational Foundation, Inc., 2002.

[2] P. Bajari and A. Hortacsu. The winner's curse, reserve prices, and endogenous entry: empirical insights from ebay auctions. *RAND Journal of Economics*, pages 329–355, 2003.

[3] G. E. Box, G. M. Jenkins, and G. C. Reinsel. *Time series analysis: forecasting and control*. Prentice Hall, 3rd edition, 1994.

[4] H. Cheng, P.-N. Tan, J. Gao, and J. Scripps. Multistep-ahead time series prediction. In *Proc. of the 10th Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining*, pages 765–774, 2006.

[5] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[6] E. F. Fama. The behavior of stock-market prices. *Journal of Business*, pages 34–105, 1965.

[7] E. F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.

[8] T. Gehrig and L. Menkhoff. Extended evidence on the use of technical analysis in foreign exchange. *International Journal of Finance & Economics*, 11(4):327–338, 2006.

[9] M. Gordon. *The Investment, Financing, and Valuation of the Corporation*. Irwin Series in Economics. Greenwood Press, 1962.

[10] M. J. Gordon. Dividends, earnings, and stock prices. *The Review of Economics and Statistics*, pages 99–105, 1959.

[11] D. Hennes, D. Bloembergen, M. Kaisers, K. Tuyls, and S. Parsons. Evolutionary advantage of foresight in markets. In *Proc. of the Genetic and Evolutionary Computation Conference (GECCO 2012)*, pages 943–950, 2012.

[12] J. Huber, M. Kirchler, and M. Sutter. Is more information always better? Experimental financial markets with cumulative information. *Journal of Economic Behavior and Organization*, 65(1):86 – 104, 2008.

[13] M. Kaisers, K. Tuyls, S. Parsons, and F. Thuijsman. An evolutionary model of multi-agent learning with a varying exploration rate. In *Proc. of The 8th Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 1255–1256. International Foundation for Autonomous Agents and Multiagent Systems, 2009.

[14] B. Malkiel. The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1):59–82, 2003.

[15] J. Niu, K. Cai, S. Parsons, M. Fasli, and X. Yao. A grey-box approach to automated mechanism design. *Electronic Commerce Research and Applications*, 11(1):24–35, 2012.

[16] S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double-auction market designs. *Agent-Mediated Electronic Commerce VI. Theories for and Engineering of Distributed Mechanisms and Systems*, pages 101–114, 2005.

[17] Y. Shoham, R. Powers, and T. Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.

[18] T. Stöckl and M. Kirchler. Trading behavior and profits in experimental asset markets with asymmetric information. *Journal of Behavioral and Experimental Finance*, 2:18–30, 2014.

[19] M. P. Taylor and H. Allen. The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance*, 11(3):304 – 314, 1992.

[20] B. Tóth and E. Scalas. The value of information in financial markets: An agent-based simulation. In J. Huber and M. Hanke, editors, *Information, Interaction, and (In)Efficiency in Financial Markets*. Linde Verlag, 2007.

[21] B. Tóth, E. Scalas, J. Huber, and M. Kirchler. The value of information in a multi-agent market model. *Eur. Phys. J. B*, 55:115–120, 2007.

[22] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416, 2007.

[23] W. Walsh, R. Das, G. Tesauro, and J. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002.