

An Asset Pricing Model with Loss Aversion and its Stylized Facts

Radu T. Pruna
School of Electronics and
Computer Science
University of Southampton, UK
Email: rp14g11@soton.ac.uk

Maria Polukarov
School of Electronics and
Computer Science
University of Southampton, UK
Email: mp3@soton.ac.uk

Nicholas R. Jennings
Departments of Computing and
Electrical and Electronic Engineering
Imperial College London, UK
Email: n.jennings@imperial.ac.uk

Abstract—A well-defined agent-based model able to match the widely observed properties of financial assets is valuable for testing the implications of various empirically observed heuristics associated with investors behaviour. In this paper, we extend one of the most successful models in capturing the observed behaviour of traders, and present a new behavioural asset pricing model with heterogeneous agents. Specifically, we introduce a new behavioural bias in the model, loss aversion, and show that it causes a major difference in the agents interactions. As we demonstrate, the resulting dynamics achieve one of the major objectives of the field, replicating a rich set of the stylized facts of financial data. In particular, for the first time our model enables us to match the following empirically observed properties: conditional heavy tails of returns, gains/loss asymmetry, volume power-law and long memory and volume-volatility relations.

I. INTRODUCTION

One of the most influential theories in modern financial markets is the Efficient Market Hypothesis (EMH) [1]. It postulates that prices fully reflect all available information, meaning that they should adjust immediately and correctly to incorporate new information. Consequently, a market with this information efficiency leads to a random sequence of price changes, mathematically known as a martingale, which cannot be predicted. However, many theoretical and empirical implications of the EMH have been tested over the years and the debates regarding market efficiency are still unresolved.

In particular, the rationality levels required by the EMH become impossible in a heterogeneous world, where some agents may be able to solve larger problems quicker than others. To this end, in the last two decades a rich literature on agent-based financial models has developed, considering the markets as populations of different groups of trading agents, thus providing an alternative to the rational expectations framework where all agents have full, unbounded rationality. Generally speaking, agent-based models (ABM) are systems in which a number of heterogeneous agents interact with each other and their environment [2], [3]. These models can provide useful insight on the behaviour of individual agents and also on the effects that emerge from their interaction, making the financial markets a very appealing application for agent-based modelling.

Nowadays, researchers do not only try to illustrate the basic mechanisms of the models, but also quantitatively recreate the

statistical properties of financial markets. These characteristics, so robust across different financial instruments and time periods, are known as stylized facts [4]. They target a wide range of aspects in the market and cannot be explained by the EMH. Therefore, the agent-based models attempt to explain these statistical properties endogenously, by considering the interaction of market participants [5], [6]. However, due to the wide variety of trading strategies, motives and rationality in the real financial markets, it is hard to find the origins of particularly interesting events. This has led to the development of various models that aim to understand the links between empirical regularities of the markets and the complexities of the entire economic system.

Interestingly, models that rely on simple trading agents have proven themselves very efficient in generating important dynamics of real financial markets. Specifically, the fundamentalists vs. chartists setting (to be defined), also known as the 2-type design [5], incorporates all the important mechanisms of financial markets, is capable of recreating their key properties and is simple enough for analysis and computation [7], [8]. In these models, fundamentalists act as a stabilizing force in the market. They believe in the existence of a fundamental price and invest relatively to its value. In contrast, chartists (a.k.a. technical analysts), forecast the future prices entirely by modelling historical data. That is, they do not take into consideration the market fundamentals and base their trading strategies and decisions on observed historical patterns in past prices.

One of the central objectives of ABM is to define models able to generate the properties of real-life financial data and then use them to test various behavioural and economic theories. To this end, we built on the widely-used structural stochastic volatility model of [9] (FW) to test the implications of loss aversion, one of the most researched behavioural heuristics observed in real life investors. Loss aversion offers an explanation to investors' unwillingness to realise and cut their losses. In other words, losses loom larger than gains [10]. Indeed, since agent-based models aim to represent real life traders we contend they must capture their behavioural biases. We are particularly interested in loss aversion, since it was shown to effectively explain a number of financial properties such as high mean and volatility of stock returns [11].

In more detail, we present a new behavioural agent-based model that, incorporates the loss aversion in traders behaviour. In doing so, we illustrate one of the essential objectives of a well defined agent-based financial model, that it can be used to test real life behaviour and policies. This new setting produces several interesting results, for the first time seen in an agent-based setting.

Starting with the FW model, we observe that the price series it generates violates one of the core properties of financial time series – its non-stationarity. This is equivalent to having a price distribution with a mean and variance that do not change over time and strongly contradicts the well-known fact that financial price series are non-stationary [12]. We overcome this problem by extending the original model and changing the motion of the fundamental value over time. As a result, we drastically reduce the non-stationarity of the price series generated.

Furthermore, introducing the new behavioural bias into the model not only preserves all the stylized facts presented in [13] but also extends the empirically observed properties matched by agent-based financial models [5]. To the best of our knowledge, the conditional heavy tails in returns are matched for the first time by an ABM and, for the first time in a 2-type design model, we match the volume-volatility relations, gain-loss asymmetry, power-law behaviour and long memory of volume.

The remainder of the paper is organised as follows. We formally define the agent-based model in Sections II-A and II-B, followed by its key implications in Sections II-C and II-D. In Section III, we present our results. We give a thorough econometric analysis of the price series generated by the agent-based model, integrating a wide range of tests and evidence that demonstrate the presence of new stylized facts. Section IV concludes.

II. A BEHAVIOURAL MODEL WITH LOSS AVERSION

A. Price adjustments

The asset price changes are determined by excess demand, as in the FW model, following some of the most prominent examples in the literature [14]. Here, by excess demand we mean the precise positive or negative orders per trading period. The specific demand of each trader type is kept as simple as possible, in the form of the demand per average trader. For fundamentalists, the demand is inversely related to the difference between the current price and the fundamental value. That is, at time t their core demand D_t^f is proportional to the gap $(p_t^f - p_t)$, where p_t is the log price of the asset at time t , while the p_t^f is the fundamental log value at time t . Similarly, the core demand of the chartists' group, D_t^c , is proportional to the latest price changes they have observed, $(p_t - p_{t-1})$, where p_t and p_{t-1} are the log prices at time t and $t - 1$, respectively.

The wide variety of within-group specifications are captured by noise terms added to each of the core demands. These terms encapsulate the within-group heterogeneity and scale with the current size of the group. Specifically, the noise is represented by two normally distributed random variables ϵ_t^f

and ϵ_t^c for fundamentalists and chartists, respectively. Each of the two terms are sampled at every iteration and added to the deterministic part of the demand, leading to the total demand per agent within the corresponding group.

Thus, the net demand of each group for the asset in period t is given by:

$$D_t^f = \phi(p_t^f - p_t) + \epsilon_t^f \quad \epsilon_t^f \sim \mathcal{N}(0, \sigma_f^2) \quad \phi > 0, \epsilon_f > 0, \quad (1)$$

$$D_t^c = \chi(p_t - p_{t-1}) + \epsilon_t^c \quad \epsilon_t^c \sim \mathcal{N}(0, \sigma_c^2) \quad \chi \geq 0, \epsilon_c \geq 0, \quad (2)$$

where ϕ and χ are constants denoting the aggressiveness of traders' demand and σ_t^f and σ_t^c are noise variances.

The agents are allowed to switch strategies at each iteration, so their market fractions fluctuate over time. For simplicity, we fix the agents' population size at $2N$. Let n_t^f and n_t^c be the number of fundamentalists and chartists in the market at time t , respectively. We define the *majority index* of the fundamentalists as:

$$x_t = (n_t^f - n_t^c)/2N, \quad (3)$$

with $x_t \in [-1, 1]$. Thus, a value of $x_t = -1$ (+1) corresponds to a market where all traders are chartists (fundamentalists).

The (scaled) total demand D_t , given by the equation:

$$D_t = n_t^f D_t^f + n_t^c D_t^c = (1 + x_t)D_t^f/2 + (1 - x_t)D_t^c/2 \quad (4)$$

will generally be in disequilibrium. That is, the total demand of the agents will not add up to zero and we will have an excess of either supply or demand. Following the early examples in the literature [14], a market maker is assumed to absorb the excess supply and provide any excess demand. The market maker sets the price by supplying stock out of its inventory and raising the price if there is excess demand, while accumulating stock and lowering the price when there is excess supply. Specifically, the market maker reacts to the imbalance between demand and supply by proportionally adjusting the price with a constant factor $\mu > 0$.

Accordingly, the equation determining the price for the next period $t + 1$ results from Equations 1-4 as:

$$p_{t+1} = p_t + \frac{\mu}{2} [(1 + x_t)\phi(p_t^f - p_t) + (1 - x_t)\chi(p_t - p_{t-1}) + \epsilon_t], \quad (5)$$

$$\epsilon_t \sim \mathcal{N}(0, \sigma_t^2), \quad \sigma_t^2 = [(1 + x_t)^2 \sigma_f^2 + (1 - x_t)^2 \sigma_c^2]/2. \quad (6)$$

The combined variance σ_t depends on the variations of the market fractions of the fundamentalists and chartists. This random time-varying variance is a key feature of the model, being termed the structural stochastic volatility of returns (defined as the log differences in prices) in [9].

B. Evolution of the market shares

To complete the model, it remains to set up the motions of the market fractions n_t^f and n_t^c . They are predetermined within each period and change only from one period to the next one. In this section, we modify the setting presented in the FW model to accommodate for our newly introduced behavioural factors. Following the discrete choice

approach (DCA) introduced by [15], the two market shares n_{t+1}^s ($s = f, c$) can be determined using the multinomial logit model. In the basic setting, some payoff indices u_t^c and u_t^f are considered, usually derived from past gains of the two groups. The market fractions can be expressed as $n_{t+1}^s = \exp(\beta u_t^s) / [\exp(\beta u_t^f) + \exp(\beta u_t^c)]$, where β is the intensity of choice. Dividing both numerator and denominator by $\exp(\beta u_t^f)$, the market fraction of fundamentalists is given by $n_{t+1}^f = 1 / \{1 + \exp[-\beta(u_t^f - u_t^c)]\}$.

Note that the difference in any utility variables, $(u_t^f - u_t^c)$, can be viewed as a measure of relative attractiveness of fundamentalist trading. Hence, changing the notation, we make use of the attractiveness level, a_t , defined as the difference $(u_t^f - u_t^c)$. The discrete choice approach is then given by:

$$n_{t+1}^f = \frac{1}{1 + \exp(-\beta a_t)}, \quad n_{t+1}^c = 1 - n_{t+1}^f. \quad (7)$$

We normalize all the demand terms in the price Equation 6 by using a market impact $\mu = 0.01$ and fix the intensity of choice $\beta = 1$. Of course, setting these values is a matter of scaling the market impact on prices and the relative attractiveness level a_t of fundamentalism, respectively.

Note that the market fractions are directly influenced by the attractiveness level. An increase in the index a_t leads to an increase in the market share of the fundamentalists. For this reason, it is extremely important to define the exact mechanism of the attractiveness level and all of its components.

The concept of loss aversion is linked to a level of wealth, from which we can derive the domains of gains and losses [16]. In this context, the first principle that influences the agents' choice of strategies is based on differential profits, incorporating some inertia. In detail, with respect to strategy $s = f, c$, let g_t^s be the short term capital gains that an average agent in this group could realise at time t . This is obtained from the demand formulated at time $t-2$ and executed at the price p_{t-1} of the next day. In general, a utility v_t^s obtained from these gains is defined as $v_t^s = g_t^s + \eta v_{t-1}^s$, where η is a memory coefficient between 0 and 1. Following the same idea, we consider a weighted average of g_t^s and v_{t-1}^s [9] such that:

$$\begin{aligned} g_t^s &= [\exp(p_t) - \exp(p_{t-1})]D_{t-2}^s, \\ w_t^s &= \eta w_{t-1}^s + (1 - \eta)g_t^s. \end{aligned} \quad (8)$$

Note that we replaced the utility symbol v with w to accommodate our interpretation. Furthermore, w_t^s represents the accumulated profits, discounted by the coefficient $\eta < 1$, that would have been earned by an agent who had consistently followed strategy s on a daily basis. In other words, w^s is the accumulated wealth attributed to strategy s .

However, according to prospect theory investors are often risk-averse in the domain of gains but risk-seeking when the changes in wealth are perceived as losses [16]. Similarly, loss aversion influences the agents' overall risk attitude, making them more risk-seeking in the domain of losses. Our implementation of loss aversion is slightly different from prospect theory in the sense that agents are considering a form of




Fig. 1. Simulation results for the network.

accumulated wealth instead of just the latest changes. In order to introduce loss aversion, we modify the agents' wealth function such as:

$$w_t^s = \begin{cases} w_t^s & \text{if } w_t^s \geq 0 \\ -\lambda w_t^s & \text{if } w_t^s < 0, \end{cases} \quad (9)$$

where λ is the loss aversion coefficient, $s = f, c$. Following prospect theory findings, we fix the loss averse coefficient $\lambda = 2.25$ [16]. This means that technical analysts weight losses 2.25 more than gains of same magnitude. Therefore, we introduce the first term in the attractiveness level a^t as being proportional to the difference in accumulated wealth, $(w_t^f - w_t^c)$.

The second principle considered is based on one of the most empirically observed social factors in behavioural literature, known as herding. This is the investors' tendency to follow the strategies of other traders. It is one of the essential irreducible elements that an agent-based model of financial markets should possess [17]. The basic idea of herding is that joining a group becomes more attractive the more adherents it already has. In our model, this principle will be represented by a term proportional to the most recent difference between market fractions, $(n_t^f - n_t^c)$.

The last term of the attractiveness level is based on a predisposition towards one of the two strategies. This is simply expressed by a constant factor, α_0 , which is positive (negative) if there is a priori preference toward fundamentalism (charism). These three elements combined form the attractiveness level, formally defined as:

$$a_t = \alpha_0 + \alpha_n(n_t^f - n_t^c) + \alpha_w(w_t^f - w_t^c), \quad (10)$$

where α_0 is the predisposition parameter, $\alpha_n > 0$ captures the herding parameter and $\alpha_w > 0$ measures the influence of accumulated wealth.

C. Motion of the fundamental value

Financial time series are characterised by the lack of predictability, mathematically known as the martingale property. A similar, very common property of financial time series data

TABLE I
BEHAVIOURAL MODEL PARAMETERS ESTIMATED BY MSM

ϕ	χ	σ_f	σ_c	α_0	α_n	α_w	η	μ_p	σ_p
0.996	0.928	0.737	1.736	2.103	1.278	26863	0.987	0.014	0.17

is its non-stationarity. In particular, a non-stationary process is a stochastic process whose joint probability distribution changes when shifted in time. Consequently, its mean and variance change over time. It is commonly assumed that non-stationary data is unpredictable and cannot be forecasted.

Associated with non-stationarity, the unit root states that one is not able to reject the hypothesis that the time series follows a random walk. One of the widely used tests of non-stationarity is the Augmented Dickey Fuller (ADF) unit root test. Following this approach, we apply the classical ADF test (allowing for a constant and trend order) to the FW model, where the fundamental value is kept constant, $p_t^f = 0$. In 10000 simulations, the unit root test was rejected 2736 times with the p-value being less than the critical value at 10%. Moreover the test was rejected 1637 times with the p-value less than the critical value at 5%.

The failure to pass the unit root test is mainly due to the unrealistic assumption of a constant fundamental value. As a solution, we extended the FW model by allowing the fundamental value p_t^f to change over time. Specifically, we set p_t^f to follow a geometric Brownian motion. This is the most widely used model of describing stock price behaviour [18] and is usually applied in quantitative finance. Mathematically, the fundamental price is given by $dp_t^f = \mu_p p_t^f dt + \sigma_p p_t^f dW_t$, where W_t is a Weirner process, μ_p is the percentage drift and σ_p is the percentage volatility. The exact values of the drift and volatility will be estimated together with all the other model's parameters (Table I) using the Method of Simulated Moments (see [13] for more details).

With this change in the model, we run the ADF test once again over 10000 different simulations. Now, the test is rejected only 823 times with the p-value less than the critical value at 10% and 418 times with the p-value less than the critical value at 5%. Thus, we can say that the price series generated by the new behavioural model are non-stationary in more than 93% of the simulations (with 95% confidence interval). In Figure 1 we plot the percentage of the non-stationary price series generated as we increase the number of simulations. The red and blue lines represent the initial and our modified model, respectively. The hard (dotted) lines are the computed ADF results at 10% (5%) critical values. We can observe how the number of non-stationary price series is drastically reduced, from 27% with the initial model (at 10% critical size) to 4% in our modified model (at 5% critical size).

Therefore, we can say that the new model produces more realistic price series, since in real financial markets the prices follow martingales where knowledge of past events does not help predict the mean of the future changes. Unlike the original (FW) model, our improved setting better reflects such fundamental properties of real financial time series as non-

stationarity and unpredictability (our first contribution).

D. Behavioural implications

We now explore the inner workings of the behavioural model and the impact of the loss aversion on the interaction between agents and price changes. First of all, we present a simple run of the behavioural model. Figure 2a illustrates the simulated (log) prices p_t generated by the model, together with the fundamental values p_t^f . There are considerable swings in the price series, similar to the ones we observe in real life price data. The second panel of Figure 2b plots the evolution of the majority index, x_t , over the simulation period. Note how the fraction of fundamentalists and technical analysts changes continuously, varying from periods of fundamentalists dominance to periods of chartists domination. The agents' fractions are usually large and stay at the ends of the interval $[-1, 1]$, mostly having periods when the index is bigger than 0.9 or smaller than -0.9 . This means that an overwhelming majority of agents are either fundamentalists or chartists at different points in time.

In addition, Figure 2c plots the accumulated wealth, w_t , of the two classes of traders. We can observe that when the actual price comes closer to the fundamental value the fundamentalists are more profitable, while the chartists make bigger profits when prices move away from the perceived fundamental value. The volatility or swings in chartists' wealth is expected to be higher than for fundamentalist, since they have a greater variability in demand, $\sigma_c > \sigma_f$.

Finally, Figure 2d demonstrates the implications of the irregular switches in the agents' strategies on returns, specified as percentage points. Since the chartists have a greater variability in demand, by comparing Figures 2b and 2d we observe that the level of returns during a chartism domination exceeds its level in a fundamentalism regime. Therefore, it appears that normal sequences of returns are interrupted by outbursts of increased volatility, when the majority of agents are chartists.

In order to assess the impact of the loss aversion to the interaction between agents, we perform a more in-depth analysis of the market fractions. For an accurate comparison, we consider two models that differ just in this respect. Namely, we compare our previously described behavioural model with loss aversion with the same model but without the loss aversion for chartists. In particular, we look at the market factions of chartists for two behavioural models, both having the switching mechanism dependent on wealth (W), herding (H) and a priori predisposition towards one of the two strategies (P). The only difference is that one of the models has the behavioural presence of loss aversion (WHP-LA) for chartists and the other doesn't (WHP).

For a more robust comparison we perform 1000 simulation runs of both WHP and WHP-LA models and look at the median results, plotted in Figure 3. As we have discussed earlier, the fractions stay at the ends of the interval, with the market being dominated by either chartists or fundamentalists. For the WHP model, we have $0.1 < n_f < 0.9$ in 9% of the simulations, while for WHP-LA it is in 7.65% of

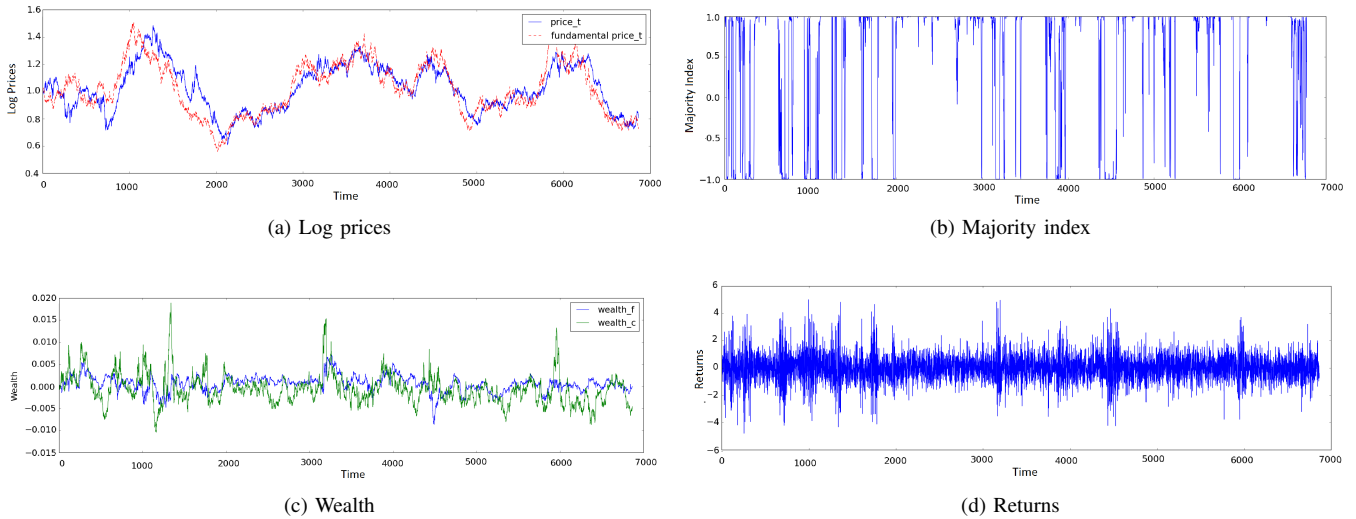


Fig. 2. Simulated time series.

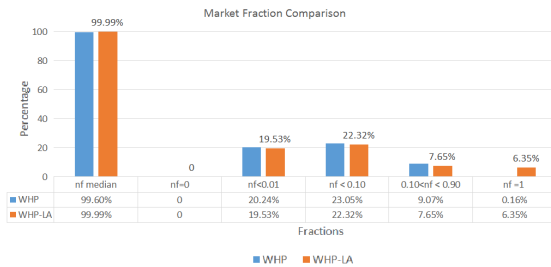


Fig. 3. Market fractions comparison.

the simulations. Furthermore, for both models, the market is heavily dominated by fundamentalists (n_f median = 0.99) with sudden disruptions of chartism, $n_f < 0.01$ in 20% of the simulations. Overall, the fractions appear to be similar for the two models, WHP and WHP-LA.

However, when we look at the ends of the intervals we see major differences between the two settings. While fundamentalists never disappear from the market ($n_f > 0$), we notice that for the WHP-LA model, the chartists are driven out of the market. In particular, the fundamentalists occupy 100% of the market ($n_f = 1$) in quite a few cases. In addition, over a median simulation of $T = 6867$ time periods, the market is fully occupied by fundamentalists 12% of the time. In contrast, for the WHP model, the chartists almost never fully disappear from the market (we obtained $n_f = 1$ in less than 0.2% of the simulations). Since the only difference between the two settings is the loss aversion for chartists, we conclude that it is this behavioural bias that pushes these agents entirely out of the market. This is a significant result since even zero-intelligence strategies, where agents buy or sell randomly, have been proven to survive the competition in agent-based settings [19].

Generally speaking, loss aversion makes the chartist weight their losses more than the gains. Hence, when the price is

going toward the fundamental value but fluctuates around, without following any trend, the chartists suffer heavy losses. In the presence of loss aversion, they overweight their losses and make chartism rapidly unattractive. Because of this significant difference in wealth, the switching mechanism then pushes more and more agents to fundamentalism. Moreover, if the prices stay around their fundamental values long enough, the herding mechanism accentuates the movement. This results in the complete disappearance of chartism.

So far, we have defined a new behavioural model with loss aversion and discussed its inner workings. As we have seen, the introduction of loss aversion for one type of the traders determines their disappearance from the market during the simulation period. This is the first time this phenomenon has been observed in an agent-based model of financial markets (our second contribution). In other ABMs, the agents interact with each other and their market fractions are continuously changing but even the most simple strategies survive the competition. However, by introducing a behavioural heuristic to technical analysts, we observe how it leads to their complete extinction from the market. Therefore, in the complex setting of financial markets, where strategies compete with each other, a behavioural bias can be powerful enough to push the agents out from the market. Similarly, behavioural finance literature and psychological experiments argue that people's beliefs lead to instabilities in the market and systematic errors which can be exploited, leading to deviations from market efficiency [20].

III. STYLIZED FACTS

In this section, we explore the statistical properties generated by our model. We show that it matches a rich set of empirically observed stylized facts, demonstrating one of the main objectives of ABM. An in-depth analysis of the time series generated by the model will be performed. First, we demonstrate the existence of the most discussed properties, namely the absence of autocorrelations, heavy tails, volatility




Fig. 4.

clustering and long memory. Moreover, our model is able to match the extensive list of stylized facts discussed in an earlier version of our work [13], but will not be presented for brevity. Next, for the first time in an agent-based model we show the existence of conditional heavy tails in returns. Moreover, we match the gain-loss asymmetry, volume-volatility relations, power-law behaviour and long memory of volume, for the first time with a 2-type design (our third contribution).

A. Absence of autocorrelations, volatility clustering, long memory, heavy tails and conditional heavy tails

One of the most well-known properties of financial data state that price movements do not exhibit any significant autocorrelation. The prices autocorrelation function (ACF) decays very sharply and is usually close to zero [4]. This absence of autocorrelations can be easily observed by plotting the ACF of raw returns at lags from 1 to 100 in Figure 4a, clearly demonstrating that price changes are not autocorrelated.

However, the absence of autocorrelation does not rule out the possibility of non-linear dependencies of returns, since absence of serial correlation does not imply independence. In financial time series, volatility measured as non-linear representations of returns exhibit a much higher positive autocorrelation that persists over time. This phenomenon is a quantitative signature of volatility clustering: large price variations are more likely to be followed by large price variations, as observed in Figure 2d. A common way of confirming the presence of volatility clustering is by considering its autocorrelation function. Even though there are different ways of measuring volatility, the most commonly used ones are the absolute returns [5]. In the ACF of volatility measured as absolute returns (Figure 4a), we observe a positive autocorrelation that persists over time, doubled by its slow decay. This is a clear presence of volatility clustering.


A property closely related to volatility clustering is the decay of the ACF. The long memory effect specifically addresses this decay. Usually, if the decay is slow, similar to a hyperbolic function, we can say that the corresponding process exhibits long memory. One way of observing the decay in the ACF of returns is by fitting a power law of the form A/r^β . We notice that a close fit of the ACF can be obtained, with an exponent $\beta = 0.48$ (Figure 4a), in line with empirical studies [4].

Next, we consider the distribution of returns, another challenging topic in the econometrics literature. In Figure 4b we plot the distribution of returns and immediately observe the deviations from Gaussianity. Our simulated returns, similarly to the ones of many financial assets, are bell shaped but contain more mass in the peak and the tail than normal. Specifically, an excess kurtosis of 2.34 implies a peakiness bigger than normal and a slow asymptotic decay of the probability distribution function. This non-normal decay is the so-called heavy (or fat) tail [21]. Furthermore, we analyse the distribution of absolute returns by computing the well-known Hill index. We obtain a Hill tail index, similar to empirical findings, most studies reporting a tail index higher than two and less than five [4].


Finally, for the first time in an agent-based setting, we discuss a further property observed empirically in financial markets, the so-called conditional heavy tails of returns. It states that even after correcting returns for volatility clustering (e.g. via GARCH-type models [22]), the residual time series still has heavy tails [4]. However, the tails are less heavy than the unconditional distribution. A clear representation of this property can be observed in Figure 4c, where we plot the distribution of the residual returns after correcting for volatility clustering via a GARCH(1,1) model. Compared with the normal distribution superimposed on it, we see that the residuals of returns have a higher peak and longer tails, confirmed by an excess kurtosis of 0.91 and a Hill index of 4.2. In the same time, the tail is less heavy than the one of raw returns.

B. Gain-loss asymmetry


In order to obtain a deeper understanding of the fluctuations of prices and returns, [23] proposed a different kind of approach involving inverse statistics. In the analysis of financial data, the inverse question can be written as: For a given return on an investment, what is the typical time span needed to obtain this return? This question has been investigated in a series of papers [23], [24], leading to an additional feature of returns. Specifically, the gain/loss asymmetry of returns states that while the maximum of inverse statistics for a positive level of returns occurs at a specific time, the maximum for the same negative level of returns appears earlier. For the first time in



(a) Investment horizon distribution for a return level $|\rho| = 0.05$. The insert is the optimal investment horizon τ_ρ^* for positive and negative levels of returns. (b) The optimal investment horizon τ_ρ^* for positive and negative levels of returns. (c) Volume distribution.



(d) Autocorrelation function of volume.



(e) Volume- volatility realtions. Cross correlation between volatility measured as absolute returns (blue) and squared returns (green) and volume distribution.

Fig. 5.

agent-based modelling, we will match this stylized fact with a 2-type design model.

In more detail, the level of return $\pm\rho$, where the positive or negative sign corresponds to a gain, or a loss respectively, is kept fixed and one looks for the shortest waiting time after t , $\tau_{\pm\rho}(t)$, for which the returns is above or below the predefined $\pm\rho$. The distribution of waiting times for gains $\tau_{+\rho}$, and loses $\tau_{-\rho}$, is denoted $p(\tau_{\pm\rho})$ and corresponds to a fixed return level ρ . Therefore, the maximum of this distribution, τ_ρ^* is the most probable time of production a return ρ , known as the optimal investment horizon.

In Figure 5a we plot the probability distribution function, $p(\tau_{\pm\rho})$ of waiting times $\tau_{\pm\rho}$, with a return level $\rho = 5\%$. We can observe the asymmetry between simulated investment horizon distributions. In particular, for a negative level of return, there is a higher probability to short investment horizons, as compared to what is observed for the positive one. On average, while the optimal investment time for a negative return is 20 days, for the same level of positive returns we obtain an optimal investment time of 23 days.

Similarly, a relevant property for an investor would be the dependence of the optimal horizon τ_ρ^* on the return level ρ . Empirically, the increase of the optimal horizon was found to occur in a systematic fashion, $\tau_\rho^* \propto \rho^\gamma$, with $\gamma \simeq 1.8$. In Figure 5b, we plot the optimal investment horizon τ_ρ^* as a function of ρ . We find a power $\gamma = 1.69$ for negative returns and $\gamma = 1.75$

for positive returns, in line with empirical findings [23].

C. Volume power-law, long memory and volume-volatility relations

We now turn our attention to volume, another highly discussed financial topic. Some of the empirical quantitative properties related to volume include a power-law behaviour, long memory and correlation to volatility. We will demonstrate all these stylized facts with our behavioural model, being for the first time matched by a 2-type model.

We define the volume at time t as the total absolute demand at t of both fundamentalists and chartists. In Figure 5c we can see that the volume distribution has a bigger decay than the exponential distribution superimposed on it, indicating a power-law behaviour. Moreover, an excess kurtosis of 5.47 and a skewness of 1.96 reinforces this behaviour. A Hill index of 3.6 was found.

In Figure 5d we plot the autocorrelation function of volume. We observe a power-law decay with an exponent 0.585. The slow decay is a clear presence of long-memory of volume. Another widely used method for testing the long memory effect is by using the Hurst exponent. Computing the Hurst exponent of the simulated distribution of volume, we obtain the value of 0.68. Most empirical studies report a Hurst exponent in the interval $[0.5, 1]$ for long memory processes.

Therefore, the presence of long memory in volume is clearly demonstrated.

Finally, the dependence between volatility and volume traded has been noticed and documented across different financial instruments at different time scales. Specifically, the relationship between the two market descriptors can be observed via the cross correlation function. In Figure 5e we observe a significantly positive cross correlation between volatility and volumes. This means that a small (large) trading volume is accompanied by a small (large) change in volatility. It is important to note that all measures of volatility are positively correlated with volume and that the dependence remains significant as we increase the time lag.

IV. CONCLUSIONS

Heterogeneous agent-based models that rely on simple automated trading agents have proven themselves very efficient in generating important dynamics of real financial markets. Moreover, a well-defined agent-based model able to match the financial markets' common properties is an important tool for testing various behavioural and economical theories and understanding how they influence the interactions between agents and prices.

In this paper, we extended one of the most recent and successful models in capturing real-life market dynamics. Motivated by the violation of a central property of real financial price series, their non-stationarity, we propose a change in the motion of the fundamental value and overcome it. Next, we presented a new behavioural model of asset pricing where the agents are loss averse. In the new setting, we observed a major difference in the movements of the agents' market fractions. Specifically, the loss averse chartists were driven out of the market at different points in time. Since even the simplest strategies have been proven to survive the competition in an agent-based setting, we can link our finding with the behavioural finance literature, which states that investors' systematic biases lead to unexpected market behaviour. Finally, we provided an in-depth analysis of the time series generated by the model and showed that for the first time it is able to match the empirically observed properties of conditional heavy tails, gain loss asymmetry, volume power law, long memory and volume-volatility relations.

For future work, we will use the model to investigate how other behavioural factors such as gossip or waves of optimism and pessimism influence the way agents interact with each other and test their impact on the market. In this way, we can detect the biases and heuristics that make the prices display some of the anomalies which cannot be explained by traditional financial theory.

REFERENCES

[1] E. F. Fama, "Efficient capital markets: A review of theory and empirical work," *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970.
 [2] L. Tesfatsion, "Agent-based computational economics: Growing economies from the bottom up," *Artificial life*, vol. 8, no. 1, pp. 55–82, 2002.

[3] M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," *The knowledge engineering review*, vol. 10, no. 02, pp. 115–152, 1995.
 [4] R. Cont, "Empirical properties of asset returns: stylized facts and statistical issues," 2001.
 [5] S.-H. Chen, C.-L. Chang, and Y.-R. Du, "Agent-based economic models and econometrics," *The Knowledge Engineering Review*, vol. 27, no. 02, pp. 187–219, 2012.
 [6] B. LeBaron, "Agent-based computational finance," *Handbook of computational economics*, vol. 2, pp. 1187–1233, 2006.
 [7] T. Lux, "Herd behaviour, bubbles and crashes," *The economic journal*, pp. 881–896, 1995.
 [8] W. A. Brock and C. H. Hommes, "Heterogeneous beliefs and routes to chaos in a simple asset pricing model," *Journal of Economic dynamics and Control*, vol. 22, no. 8, pp. 1235–1274, 1998.
 [9] R. Franke and F. Westerhoff, "Structural stochastic volatility in asset pricing dynamics: Estimation and model contest," *Journal of Economic Dynamics and Control*, vol. 36, no. 8, pp. 1193–1211, 2012.
 [10] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica: Journal of the econometric society*, pp. 263–291, 1979.
 [11] N. Barberis, M. Huang, and T. Santos, "Prospect theory and asset prices," National bureau of economic research, Tech. Rep., 1999.
 [12] A. Pagan, "The econometrics of financial markets," *Journal of empirical finance*, vol. 3, no. 1, pp. 15–102, 1996.
 [13] R. T. Pruna, M. Polukarov, and N. R. Jennings, "A new structural stochastic volatility model of asset pricing and its stylized facts," *arXiv preprint arXiv:1604.08824*, 2016.
 [14] J. D. Farmer and S. Joshi, "The price dynamics of common trading strategies," *Journal of Economic Behavior & Organization*, vol. 49, no. 2, pp. 149–171, 2002.
 [15] W. A. Brock and C. H. Hommes, "A rational route to randomness," *Econometrica: Journal of the Econometric Society*, pp. 1059–1095, 1997.
 [16] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.
 [17] V. Alfi, M. Cristelli, L. Pietronero, and A. Zaccaria, "Minimal agent based model for financial markets i," *The European Physical Journal B*, vol. 67, no. 3, pp. 385–397, 2009.
 [18] J. C. Hull, *Options, futures, and other derivatives*. Pearson Education India, 2006.
 [19] D. Ladley, "Zero intelligence in economics and finance," *The Knowledge Engineering Review*, vol. 27, no. 02, pp. 273–286, 2012.
 [20] R. J. Shiller, *Irrational exuberance*. Princeton university press, 2015.
 [21] B. LeBaron and R. Samanta, "Extreme value theory and fat tails in equity markets," *Available at SSRN 873656*, 2005.
 [22] L. Bauwens, S. Laurent, and J. V. Rombouts, "Multivariate garch models: a survey," *Journal of applied econometrics*, vol. 21, no. 1, pp. 79–109, 2006.
 [23] I. Simonsen, M. H. Jensen, and A. Johansen, "Optimal investment horizons," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 27, no. 4, pp. 583–586, 2002.
 [24] M. H. Jensen, A. Johansen, and I. Simonsen, "Inverse statistics in economics: the gain-loss asymmetry," *Physica A: Statistical Mechanics and its Applications*, vol. 324, no. 1, pp. 338–343, 2003.