

Adaptive agent-based modeling in finance: selected applications

Marcello Esposito

Adaptive agent-based modeling in finance

Selected applications

M. Esposito*

Abstract

Since its inception, the Efficient Market Hypothesis (EMH) has faced persistent challenges, as numerous anomalies - such as volatility clustering, excessive trading volumes, and herding behaviour - exposed gaps between theoretical predictions and actual market dynamics. In response, economists developed alternative frameworks that relaxed EMH's strict assumptions, distinguishing between different types of investors (e.g., "chartists" and "fundamentalists") and incorporating bounded rationality, learning, and adaptation. This line of research gave rise to agent-based models, which conceptualize financial markets as adaptive ecosystems and rely on simulations to capture investor interactions and the evolution of trading strategies. This paper reviews central modelling choices - such as the definition of investor heterogeneity, the specification of preferences, the mechanisms of price formation, and the processes of strategy selection - and discusses their implications for balancing realism with the complexity of calibration.

Keywords: adaptive behaviour, agent-based modelling

JEL: G11 - G14 - G17

* School of Economics and Management, Università Carlo Cattaneo - LIUC, Corso Matteotti 22, 21053 Castellanza, Italy. Email: mesposito@liuc.it

“While the economist is prepared to accept that insects may follow some mechanical rule of behaviour which may have evolved so as to be close to optimal ... if the same model is applied to economic agents, he usually requires that the behaviour should be fully and consciously optimizing.”

Kirman (1993) “Ants, Rationality, and Recruitment”

1. Introduction

While the Efficient Market Hypothesis (EMH) is often regarded as one of the greatest intellectual achievements in economics, many researchers have, from the very beginning, highlighted the gaps between the theory’s implications and the realities of financial markets. The most significant of these discrepancies evolved into ‘puzzles’ that successive generations of economists have sought to solve. These puzzles were not limited to statistical anomalies in stock prices or the persistent behavioural ‘errors’ of non-professional investors. They extended to the very structure and functioning of financial markets. In fact, trading activity, active asset management, and nearly all forms of research - not only technical analysis, but also fundamental and macroeconomic analysis - are difficult to reconcile with the principles of EMH.

Broadly speaking, the strategy for addressing these puzzles has often involved relaxing one or more of the core assumptions underlying the EMH. For instance, some approaches distinguish between different types of economic agents, with certain groups following suboptimal or even illusory rules, such as believing that past price patterns are likely to repeat. If this sounds like an ad hoc solution, that’s because, in many ways, it is.

While relying on specific behavioural biases can help explain specific financial anomalies, it risks straying from the scientific method and the falsifiability principle championed by Karl Popper.¹ To address this concern, another branch of economic research developed agent-based models. These models incorporate forms of ‘bounded’ rationality, depicting agents as following adaptive strategies aimed at discovering and adopting the most successful investment rules.

The approach of interpreting financial markets through the lens of ecology succeeds in reproducing - and, to some extent, explaining - certain observed anomalies, such as volatility clustering and herding behaviour, without abandoning the notion of optimizing, though bounded, rational behaviour on the part of economic agents. Furthermore, the use of computer simulations to model interactions among investors provides an

¹ See Popper (1959)

experimental framework, helping to anchor these models in testable scenarios and uphold Karl Popper's Falsification Principle.

However, it is important to note that there is no single, established 'model' of an adaptive financial ecosystem to serve as a reference point. Even when focusing on a highly simplified representation of a financial market, different choices regarding key elements of the experimental design might lead to different models - each with distinct implications for how the simulated markets function.²

Some of these choices are relatively straightforward to grasp: for instance, how investor heterogeneity is expressed (differences in expected returns or in expected volatility?) and how it is formed (which forecasting models or heuristics do they use?), or how individual excess demands are aggregated to determine the market price (via a *walrasian* auctioneer or a market-maker?). Others may appear more technical, yet they too significantly affect both the feasibility of the simulations and their outcomes. To name just a couple: the selection of the utility function to represent agents' preferences, and the mechanism by which agents interact and select strategies (an explicit fitness function or an implicit birth - death process?).

More generally, these design choices create significant challenges when it comes to calibrating such models and using them to generate predictions on aspects beyond the specific phenomenon being studied. Unlike the EMH, even the simplest adaptive agent - based models face a kind of 'curse of dimensionality': the number of parameters to estimate goes far beyond the usual mean and variance of stock returns, and, in most cases, these parameters are not directly observable.

Since there is no single, widely accepted model, it is useful to review the key issues involved in defining a financial ecosystem and to present the solutions we consider most appropriate. In this paper, we adopt the Brock - Hommes (1998) framework as our theoretical foundation, showing how to use it for modelling market dynamics when the world is populated only by two types of investors, chartist and fundamentalist, as in Gaunersdorfer - Hommes (2008). We will then modify some of its assumptions to incorporate more realistic features, at least from an asset management point of view: in our model, the market is populated by traditional asset allocators (mutual funds, pension funds, discretionary mandates) constrained in their risk profile by the mandate received by their clients. Finally, we will see how to simulate high frequency

² See Palmer et al (1994) and LeBaron (2002) for an historical perspective.

interactions, borrowing the birth - death process used by Kirman (1993) to model herding behaviour.

In conclusion, sixty years after the foundational works of Samuelson (1965) and Fama (1965), the EMH still stands as the theoretical benchmark for asset pricing and financial economics. However, many-agent adaptive models of investor interactions are beginning to offer a laboratory-like alternative - an environment where markets are simulated under explicit assumptions about investor heterogeneity.

2. Efficient market hypothesis: anomalies and puzzles

A review of the vast literature on the departures from the Efficient Market Hypothesis (EMH) is beyond the scope of this paper. The first anomalies in financial assets prices were described and published in the same years in which the EMH foundations were laid out by Samuelson (1965) and Fama (1965).³ Many anomalies were of statistical nature, like volatility clustering, and inspired the research activity in fractals, chaos theory, and autoregressive conditional volatility modelling. But some of them were in a certain sense much deeper and structural.

Many activities performed routinely by the most prestigious financial institutions, like for example the research on company's fundamentals or the analysis of macroeconomic dynamics, shouldn't exist (if prices reflect immediately every news, as the EMH implies). And the entire industry of active asset management, comprising traditional mutual funds and the most sophisticated hedge funds, should be a relic of a distant past, like voodoo practices for medical healing. After all, also the weakest form of EMH (i.e. past prices shouldn't tell anything about future stocks' performance) is in contrast with the practice of professional and amateur investing. In 2025, 60 years after the EMH, every financial platform, be it dedicated to professionals (like Bloomberg) or be it free and open to retail investors (like Yahoo Finance), has an impressive graphical array of instruments inspired to technical analysis. And for a good "commercial" reason: before putting an order, a look at the graph of the past price performance is a very hard to die habit, also for the most convinced EMH supporter.

The nature and scale of trading activity in real financial markets is another phenomenon that is largely unexplained by the EMH. According to the EMH, trading is something that should not occur if it were not for those constraints (mostly on

³ See Mandelbrot-Hudson (2004) and Lo (2017) for a critical review of the literature on the Efficient Market Hypothesis departures.

borrowing) and frictions that limit the effective liquidity and the rapidity with which prices should reflect all the relevant information. However, the scale of trading occurring in every financial market, also in those that are populated almost entirely by professional, “informed” agents (like the secondary government bond markets or the forex), is such that there must be a “structural” factor, enshrined in the way in which human beings interact and by interacting form expectations about the future, behind what happens every day in every market around the world.

All these structural departures from the EMH witness that the model of aggregation of individuals’ behaviour behind the dominant “random walk” hypothesis is missing something. In the best case (for the hypothesis), assets’ pricing research and active asset management could be like horoscopes: an irrational anomaly of contemporary capitalist cultures that sometimes in the future will disappear from the rule books of financial advisory. In the worst case, it could be that the “angelic” model of rational behaviour and expectations’ formation is good only to describe a financial market located in Paradise.

2.1. Explaining the puzzles through investors’ heterogeneity

How economic theory tried to explain those puzzles? Varian (1986) stated that trading activity is impossible without assuming substantial differences between investors and he specified also the nature of the differences that can generate trading activity. Differences in endowments or preferences (i.e. risk aversion), for example, are not sufficient because they can generate trading only in the initial phase of the portfolio adjustment toward the optimal composition. According to EMH, changes in information can move the prices, but do not generate trade unless new information arrival is differently interpreted by the investors. The reason why a difference in opinion generates a trading activity, while a difference in information does not, is clearly stated in Varian (1992): *“if one person tells another person something that is perceived as information, the second person will adjust his views to incorporate the additional information. If one person tells another person something that is perceived as ‘just an opinion’ then no adjustment in views will take place.”*

One of the first ideas was then to imagine the existence of different types of agents interpreting and processing differently the new information arrivals. The distinction initially was then made between “smart” and “noise” traders.⁴ As the denominations suggest, the first type is the ideal, rational individual following the classical optimizing

⁴ See Black (1985)

behaviour. The second type is instead an individual that misinterprets noise for signals about the real value of an asset. Typically, a “chartist” that, observing a pattern in the historical timeseries of a stock price, mistakenly believe in its future repetition and bets on it. The chartist can then react to short term up-trends by buying (“positive feedback” traders or “trend chasers” or “momentum” investors) or selling (“negative feedback” traders or “contrarian” investors).⁵ The activity of noise traders creates the excess volatility and auto-correlation in returns, observed, for example, by Fama and French (1988), and can help explaining also the lower implied volatility of prices when markets are closed, documented in French and Roll (1986) when analysing the open-close prices.

Another interesting consequence of this distinction between “informed” and “feedback” traders is that it can generate “herding” behaviour in simulated, artificial markets, and then replicate (if not explain) the episodes of manias and panics that from time to time occur in financial markets. As it was recognized by the first advocates of the EMH, panics and crashes shouldn’t happen unless some unexpected event hits the economy with the strength of a tsunami.

From a theoretical point of view, herding behaviour has been justified in the literature about informational cascades, initiated by Scharfstein and Stein (1990) and by Day and Huang (1990). Shiller and Pound (1986) found, through a careful designed survey, that spreading of interest on specific stocks by word-of-mouth was relevant and was one of the most important factors behind the investment decisions of professional asset managers. This result brought them to conclude that epidemiological models describing the “contagion of interest” among professional investors should be taken into consideration as an alternative to the standard random walk hypothesis.

The herding behaviour of investors, be it determined by a recruitment mechanism or by an imitation game, can generate interesting stock price dynamics with a very limited set of assumptions and without resorting to exogenous shocks to the system. Since the beginnings of modern financial markets, “fads”, “contagion”, “panic”... are commonly used metaphors to describe endogenous fluctuations in asset prices not justified by economic “fundamentals”. After all, there is no way of explaining a financial crisis or a bubble formation in the stock market using the classical, rational paradigm of the *homo economicus* underpinning the EMH.

⁵ See, for example, Shiller (1984), Shiller (1988), DeLong, Shleifer, Summers, and Waldmann (1990), Cutler, Poterba, and Summers (1990), Campbell and Kyle (1993), Kirman (1993), and Shleifer (2000).

2.2. The co-existence of different type of investors

As one can understand, the issue with noise traders is not so much their existence as their survival. It is true that in specific circumstances the game between informed and noise traders can bring to unexpected “equilibria” where the former can’t exploit their superiority over the latter, but in real world circumstances there is only the “Flaiano Law” that can reasonably avoid the total elimination of noise traders by the informed ones: a periodic and continuous replenishment of the pool of the former.⁶

Given the importance of the issue of survival of the uninformed investors, we close this very short and concise survey with the contribution of Grossman-Stiglitz (1980). They started from the observation that in the informationally efficient markets hypothesis lies a paradox. If all the available information is embodied in the price system, there is no need to make research and collect new information. Then, there is no need to trade on a different view of the course of events and the underlying value of a specific asset because prices would reflect immediately all the available information. Under the perfectly efficient market hypothesis, trading would occur only to adjust the assets holdings to the evolution of the liabilities, and this implies that eventually the markets, as we know them, would collapse and disappear.

They then assumed that markets are not informationally efficient per se and there is the possibility of earning a profit from the research activity because the market price of an asset can diverge from its underlying value. Since research is a costly activity, some traders prefer to perform such activity, while others not. To reach an equilibrium, the informed traders need to observe the real value of an asset with a certain degree of noise that depends negatively on the number of resources dedicated to the research activity. The remuneration derives from the accuracy of the observation of the underlying value and by trading with the uninformed traders on the discrepancy with the market price. In their model, investors that perform “active” asset management co-exist with investors that prefer instead to follow “passive” asset management rules.

However, also in a model where research activity is costly and there is a potential trade-off between being informed and uninformed, one can understand that uninformed traders need to show a very high level of sophistication and smartness: to avoid being wiped out by informed traders, they need to be able (i) to infer from the trading activity

⁶ Flaiano, an Italian writer, once said: “la madre dei cretini è sempre incinta” (“the mother of idiots is always pregnant”)

of the informed traders (and then by the prices' movements) the value of the assets and (ii) to play protective trading strategies.⁷

The possibility for the uninformed traders to devise and follow those trading strategies without making their own research about the macroeconomic environment and each single stocks is a very strong assumption. But, if this assumption is true (and it is a big "if"), an equilibrium can be reached where the profits that the informed traders make cover the costs that they bear on their research activity. Prices remain "informative", also if not in the radical sense of the Efficient Markets Hypothesis, and uninformed traders simply decide to remain uninformed not because they are dumb but because they don't want to sustain those costs.

2.3. Financial markets as ecosystems

After the contributions of the '70s and '80s on investors' heterogeneity, the terrain was ready for building agent-based models, where financial markets dynamics could be obtained through simulations. Differently from other research strands that preferred to focus more on psychological aspects of the human behaviour that are in contrast with the ideal of the *homo economicus*, the adaptive-markets literature preferred to stick to the traditional concept of the economics' tradition (utility functions, budget constraints, optimization...) but injecting forms of "bounded" rationality in the forecasting activity.

The heterogeneity among agents translated in a competition between simple investment strategies, with the agents selecting the most successful among them. In the literature, the selection process has been modelled either explicitly, through a fitness function based on the relative performance of the strategies, or implicitly, through a birth - death stochastic process. In any case, the great majority of the models let the agents migrate from one strategy to another, generating extra-volatility and herding behaviour.

A trading strategy is characterized by the way in which expectations of future prices are formed. Heterogeneity among investors is determined by the choice of the strategy. Depending on the realized profits of the different strategies, investors, characterized by bounded rationality, chose the best strategy and switch among them from time to time. One could then draw a parallel with biology and consider financial markets as

⁷ The trading activity with informed traders must be limited in size and temporally sequenced. For example, if the informed traders knows that tomorrow a pharmaceutical company will pass the FDA last trial, they could propose to the uninformed traders an infinitely leveraged bet on its price jumping. To avoid being wiped out by such types of trades, the uninformed traders accept only limited trades, and those trades must be sequentially differentiated to have the time to interpret the counterparty orders and adjust the price.

ecosystems, populated by agents playing an evolutionary competition game among different strategies and adapting their beliefs to the context.⁸

⁸ See Brock and Hommes (1998), Gaunersdorfer and Hommes (2007) and Chiarella and He (2008).

3. Portfolio choice and wealth dynamics

We will assume that there are only two financial assets in the economy: a risky stock and a risk-free bond. The risky stock is in fixed supply, while the risk-free bond is in perfectly elastic supply. As far as investors are concerned, we will assume that they have identical preferences, represented by a utility function having wealth as its argument.

3.1. Wealth dynamics

Investor's wealth evolves in discrete time according to the following equation:

$$W_{t+1} = W_t \omega_t (1 + r_{s,t+1}) + W_t (1 - \omega_t) (1 + r_f) \quad (1)$$

Where W is the investor's wealth, ω is the portion of the investor wealth invested in stocks, r_f is the risk-free rate, r_s is the (stochastic) return on stocks. Equation (1) can be rewritten as:

$$\begin{aligned} W_{t+1} &= W_t (1 + r_f) + W_t \omega_t (r_{s,t+1} + r_f) \\ &= W_t (1 + r_f) + d_t (r_{s,t+1} - r_f) \end{aligned} \quad (2)$$

where the equity demand is simply defined as:

$$d_t \stackrel{\text{def}}{=} W_t \omega_t \quad (3)$$

If we want to express the equity returns in terms of the stock's price, p , and its dividend, y , we obtain:

$$r_{s,t+1} - r_f = \frac{p_{t+1} + y_{t+1} - p_t}{p_t} - r_f = \frac{p_{t+1} + y_{t+1} - (1 + r_f)p_t}{p_t} \quad (4)$$

Substituting (4) in (2) we obtain:

$$\begin{aligned} W_{t+1} &= W_t (1 + r_f) + D_t \frac{p_{t+1} + y_{t+1} - (1 + r_f)p_t}{p_t} \\ &= W_t (1 + r_f) + z_t R_{t+1} \end{aligned} \quad (5)$$

where the demand for equity shares is represented by the variable z_t :

$$z_t \stackrel{\text{def}}{=} \frac{d_t}{p_t} \quad (6)$$

and the excess return per share is represented by the variable R_t :

$$R_{t+1} = p_{t+1} + y_{t+1} - (1 + r_f)p_t \quad (7)$$

The variables z_t, R_t are particularly relevant in our model and they will be used extensively.

If we measure the realized profits of an investment strategy as the extra wealth earned by the investors with respect to a risk-free investment, we immediately see that:

$$\Pi_{t+1} \stackrel{\text{def}}{=} W_{t+1} - W_t(1 + r_f) = z_t R_{t+1} \quad (8)$$

Note that R_t is outside of the investor's control, whereas z_t is determined by the investor's preferences and by his beliefs about the future. Sometimes it can be useful to measure the (realized) profits in proportion to the original wealth, especially if we need to compare different strategies and we employ logistic functions, involving exponentials. Then, it is sufficient to compute:

$$\pi_{t+1} \stackrel{\text{def}}{=} \frac{\Pi_{t+1}}{W_t} \quad (9)$$

3.2. Portfolio allocation

We assume that preferences are represented by a utility function that reduces the portfolio allocation to the standard mean-variance optimization. The only two types of utility functions that we will consider are the CARA and CRRA types. In Appendix 1, we derive the ω variable as the result of a mean - variance optimization:

$$\omega = \frac{E(r_s) - r_f}{\rho \sigma^2} \quad (10)$$

and we show that demand for equity, d_t , is equal to ω for CARA preferences and is equal to ωW_t for CRRA preferences.

The optimal number of shares, z_t , demanded by the investor is then equal to:

$$z_t = \frac{d_t}{p_t} = \begin{cases} \frac{\omega}{p_t} & \text{CARA} \\ \frac{\omega W_t}{p_t} & \text{CRRA} \end{cases} \quad (11)$$

where p_t is the stock price at time t .

3.3. Stock price determination

But where does p_t come from? In the literature the approaches to price determination are substantially two. In the first approach a *walrasian* auctioneer finds the market clearing price that equates aggregate demand and supply, while in the other approach a market-maker processes the buy-sell orders and adjust the price according to a function depending on the positive or negative excess demand for stocks. In both cases, investors are “price taker” but in the first case the market is always in equilibrium while in the second it is never in equilibrium.

The market-maker approach is very easy to apply but raise some theoretical critiques because it assumes the existence of a market maker whose economic viability should be proved. The auctioneer approach presents less problems from this point of view: there are no inventories to manage and no P&L to worry about. But we can face big “design” issues as far as the stability of the system is concerned. With an arbitrary market-maker reaction function we can, in fact, control the stability of the system and avoid a trading activity that might generate unstable dynamics. This issue is particularly acute with CRRA preferences, where wealth determines the demand for equity and wealth is, in turn, affected by the stock price dynamics.

3.3.A. Walrasian auctioneer

This approach is the one pioneered by Brock-Hommes (1998). In the case of investors’ homogeneity, the nominal aggregate demand for stocks is:

$$D_t = \sum_{i=1}^N d_{i,t} \tag{12}$$

where N is the number of investors. The demand-supply equilibrium determines the stock price, p_t :

$$p_t = \frac{\sum_i^N d_{i,t}}{S_t} \tag{13}$$

Where S_t is the supply of shares.

3.3.B. Market-maker

In this approach a market-maker adjusts the price at which he will buy or sell shares according to a rule.⁹ The rule might consist of a simple linear reaction function where the new price is positively related to the excess demand:

$$p_t = p_{t-1} + b (D_t - S_t) \quad (14)$$

However, when experimenting with simulations, one finds that a linear rule might not be sufficient to avoid unrealistic explosives or implosive price dynamics. We prefer to adopt a generalized logistic function centred around the fundamental price¹⁰ and dependent on the relative excess demand, like this one:

$$p_t = p_t^* - A + \frac{B}{1 + e^{-C \frac{(D_t - S_t)}{S_t}}} \quad (15)$$

If $B=2A$, this formula defines a symmetric band around the fundamental price where the actual price will lie, depending on the relative excess demand.

3.3.C. Other approaches

Alternatives to the first two approaches have been explored in the literature. One consists in simulating an order book like the one of modern electronic stock exchanges and use an algorithm to clear the book.¹¹ Another consists in assuming a random matching of the different investors, like the one that could occur for OTC markets as the FX, that will trade among themselves at the average price, if it is for them convenient to do so.¹²

⁹ See for example Day-Huang (1990)

¹⁰ See Appendix 2.

¹¹ See for example Chiarella-Iori (2002).

¹² See for example Beltratti-Margarita (1993).

4. Investors' heterogeneity

To generate trading activity, the first requirement is to introduce heterogeneity among investors. Differences in wealth alone are not enough. Under CARA preferences, for instance, changes in wealth have no impact on the demand for equities. Even with CRRA preferences, while an increase in one investor's wealth does lead to greater demand for equity (and bonds), that demand must be offset by supply from another investor whose wealth has declined by an equivalent amount. It is highly unrealistic to assume that such wealth-driven adjustments alone could account for the level of trading activity observed in real markets.

An alternative might be to differentiate investors according to their level of risk-aversion. This is apparently more realistic as an assumption, albeit it can justify only an initial trading activity, as Varian (1986) already noted. At the start of the simulations, investors adjust their portfolio according to differences in wealth and in preferences, but afterwards? As in the case of wealth, we need to assume continuous changes in the preferences of the investors, and of opposite sign, to justify significant levels of trading activity.

For these reasons, the economic literature preferred to focus on differences in the reaction to events affecting the probability distribution of future equity returns. The number of candidate events is almost infinite, and we don't need to imagine continuous changes in investors' wealth and preferences to create portfolio's imbalances that needs adjustments by trading on the market. Moreover, their time occurrence is coherent with the frequency at which trading occurs and, what is most important, with the observation of what happens on the markets every day.

To build a simulation model, we will assume that the heterogeneity among investors is limited to a "few" types and that it is fully characterized by the way in which investors forecast the risk and return offered by the stocks.

4.1. A "few" types of investors

We will assume that there are H types of investors and a total of N investors. If $N_{h,t}$ is the number of investors belonging to group h , the total number of investors is obviously:

$$\sum_{h=1}^H N_{h,t} = N \tag{16}$$

and the percentage of investors belonging to group h is denoted as $n_{h,t}$:

$$n_{h,t} = \frac{N_{h,t}}{N} \quad (17)$$

The population distribution among different types of investors is in general modelled by letting H and N fixed, while $n_{h,t}$ changes.¹³

To simplify notation, we will assume that the risk preferences are homogeneous inside each group, i.e. the risk aversion coefficient is the same inside each group. Then, the $\omega_{h,t}$ variable is the same inside each group:

$$\omega_{h,t} = \frac{E_{h,t}(r_s) - r_f}{\rho_h \sigma_{h,t}^2} \quad (18)$$

If preferences are of CARA type, this implies that the demand for equity is the same inside each group. In case of CRRA preferences, instead, the individual demand for equity will depend on his wealth, too. And wealth is not the same for each group member either because we can assume a heterogeneous wealth distribution at the beginning of the simulation or because we allow for migration from one group to the other and in this case the history of the investment choices matters for the individual wealth.

The total wealth at time t of the investors belonging to group h is:

$$W_{h,t} = \sum_{i=1}^{N_{h,t}} W_{\{i,h,t\}} = N_{h,t} w_{h,t} \quad (19)$$

where $w_{h,t}$ is the per capita wealth of group h . So that we can obtain a wealth distribution by groups and by individuals belonging to each group. The wealth of group h with respect to total wealth is simply:

$$\frac{W_{h,t}}{W_t} = \frac{N_{h,t} w_{h,t}}{N w_t} = n_{h,t} \frac{w_{h,t}}{w_t} \stackrel{\text{def}}{=} l_{h,t} \quad (20)$$

where W_t is the total wealth at time t and w_t is the population per capita wealth. Note that:

¹³ In Chiarella-He (2001) one can find models with $n_{h,t}$ fixed.

$$\sum_{h=1}^H l_{h,t} = 1 \quad (21)$$

and:

$$n_{h,t} w_{h,t} = w_t l_{h,t} \quad (22)$$

If there were no investors' migration from one type to another, the $l_{h,t}$ variable could measure the wealth created by the investment strategy associated to type h .

4.2. Heterogeneous beliefs

The expected return can be simply expressed as a function of the expected capital gain and the expected dividend:

$$E_{h,t}(r_s) - r_f = E_{h,t} \left[\frac{p_{t+1} + y_{t+1} - p_t}{p_t} \right] - r_f \quad (23)$$

Since:

$$E_{h,t} \left[\frac{p_{t+1} + y_{t+1} - p_t}{p_t} \right] - r_f = \frac{E_{h,t}(p_{t+1} + y_{t+1}) - p_t(1 + r_f)}{p_t} \quad (24)$$

and:

$$\begin{aligned} \sigma_{h,t}^2 &= \text{Var}_{h,t} \left[\frac{p_{t+1} + y_{t+1} - p_t}{p_t} \right] = \frac{\text{Var}_{h,t} [p_{t+1} + y_{t+1} - p_t]}{p_t^2} \\ &= \frac{\text{Var}_{h,t} [p_{t+1} + y_{t+1}]}{p_t^2} \end{aligned} \quad (25)$$

Then, inserting (24) and (25) in (18), we obtain

$$\begin{aligned} \omega_{h,t} &= \frac{1}{\rho_h} \frac{E_{h,t}(p_{t+1} + y_{t+1}) - p_t(1 + r_f)}{p_t} \frac{p_t^2}{\text{Var}_{h,t}[p_{t+1} + y_{t+1}]} \\ &= \frac{p_t}{\rho_h} \frac{E_{h,t}(p_{t+1} + y_{t+1}) - p_t(1 + r_f)}{\text{Var}_{h,t}[p_{t+1} + y_{t+1}]} = \frac{p_t}{\rho_h} \frac{E_{h,t}(R_{t+1})}{\text{Var}_{h,t}[R_{t+1}]} \end{aligned} \quad (26)$$

We will define a new variable to make notation simpler:

$$q_{h,t} \stackrel{\text{def}}{=} \frac{\omega_{h,t}}{p_t} = \frac{E_{h,t}(R_{t+1})}{\rho_h \sigma_{h,t}^2} = \frac{\mu_{h,t} - p_t(1 + r_f)}{\rho_h \lambda_{h,t}^2} \quad (27)$$

Where:

$$\mu_{h,t} \stackrel{\text{def}}{=} E_{h,t}(p_{t+1} + y_{t+1}) \quad (28)$$

$$\lambda_{h,t}^2 \stackrel{\text{def}}{=} \text{Var}_{h,t}(p_{t+1} + y_{t+1}) \quad (29)$$

4.3. Individual and aggregate demands for shares

Note that for CARA preferences the demand for shares, $z_{i,h,t}$, at time t of the i -th individual belonging to group h is exactly equal to $q_{h,t}$. And obviously this is also the average demand, $z_{h,t}$, of the group. The aggregate demand for shares is then very simple to obtain:

$$D_t^{CARA} = \sum_{h=1}^H N_{h,t} q_{h,t} = N \sum_{h=1}^H n_{h,t} q_{h,t} \quad (30)$$

With CRRA preferences, instead, the demand is slightly more complicated because of the presence of the wealth term:

$$z_{i,h,t} = q_{h,t} W_{i,t} \quad (31)$$

If we consider the demand of the average investor of group h , we can rewrite eq. (31) in terms of the average wealth of group h :

$$z_{h,t} = q_{h,t} w_{h,t} \quad (32)$$

If we sum individual demands of each group, we obtain the aggregate demand for equity shares:

$$D_t^{CRRA} = \sum_{h=1}^H N_{h,t} d_{h,t} = N \sum_{h=1}^H n_{h,t} w_{h,t} q_{h,t} = N w_t \sum_{h=1}^H l_{h,t} q_{h,t} \quad (33)$$

We now must determine the price of stocks. We will focus on the Walrasian equilibrium solution, because it is a bit more complex than the market-maker solution¹⁴ and it allows us to determine the “fundamental” stock price which plays an important role in agent - based modelling.

4.4. The market equilibrium stock price

The equilibrium condition on the capital market requires that demand for equity shares, D_t , is equal to the fixed supply, S , of equity shares. It is common to assume that the net supply, S , of “external” equity shares, is fixed. Since this is just an annoying factor, in many models, it is set equal to zero.

¹⁴ In the market-maker model, the price is simply obtained by plugging the (35) into the reaction function of the market maker, as represented for example by eq. (17)

In case of CARA preferences, then, we have that the equilibrium on the market for shares requires that:

$$\sum_{h=1}^H n_{h,t} q_{h,t} = 0 \quad (34)$$

Whereas with CRRA preferences, the equilibrium is reached when:

$$\sum_{h=1}^H l_{h,t} q_{h,t} = 0 \quad (35)$$

Recalling the expression for the individual demand for shares, the value of the equity price, p_t , can be obtained by solving the following equation for CRRA preferences:

$$\sum_{h=1}^H l_{h,t} \frac{\mu_{h,t} - p_t(1 + r_f)}{\rho_h \lambda_{h,t}^2} = 0 \quad (36)$$

In many models a further simplification is that individuals share the same preferences, i.e. they have the same risk - aversion parameter ($\rho_h = \rho$). Equation (36) simplifies to:

$$\sum_{h=1}^H l_{h,t} \frac{\mu_{h,t} - p_t(1 + r_f)}{\lambda_{h,t}^2} = 0 \quad (37)$$

4.5. An example of heterogenous investors: fundamentalist VS chartist

The “classical” distinction among investor’s types that has been made in the agent-based financial literature is the one between “fundamentalists” and “chartists”. They share the same preferences and are endowed with the same initial wealth, but they differ in the way in which they form expectations about future stock prices.¹⁵

The first type of investors believes that the future market price will converge to its fundamental value, p^* , with a mean - reversion factor of v :

$$E_{1,t}[p_{t+1}] = p^* + v(p_t - p^*) \quad 0 \leq v \leq 1 \quad (38)$$

This is the reason why they are called “fundamentalists”. Note that, if $v=1$, the fundamentalists adhere to the stricter interpretation of the EMH, because for them the stock price follows a random walk, and the best forecast of future price is the last observed price:

¹⁵ Arthur et al (1996), LeBaron et al (1999).

$$E_{1,t}[p_{t+1}] = p_t \quad \text{if } \nu = 1 \quad (39)$$

The second group of investors, called the “chartist”, forecast future prices based on past price patterns. There are many ways in which chartists can make forecasts using past prices, but to keep things simple, we assume, as in Gaunersdorfer-Hommes (2000), that chartists are “trend followers” following a very naïve rule:

$$E_{2,t}[p_{t+1}] = p_t + \gamma(p_t - p_{t-1}) \quad \gamma > 0 \quad (40)$$

If we assume that preferences are CARA and the estimate of the variance is constant and common to all investors (i.e. $\lambda^2_{h,t} = \lambda^2$), then equation (34) becomes:

$$\sum_{h=1}^H n_{h,t} \mu_{h,t} = p_t (1 + r_f) \quad (41)$$

The stock price today is equal to the discounted weighted average of the future expected price and dividend, where weights depend on numerosity of the different types of investors:

$$p_t = \frac{\sum_{h=1}^H n_{h,t} \mu_{h,t}}{(1 + r_f)} = \frac{\sum_{h=1}^H n_{h,t} E_{h,t}(p_{t+1} + y_{t+1})}{(1 + r_f)} \quad (42)$$

To calculate the demand for equity of the fundamentalist and the chartist, we need simply substitute (38) and (40) in (42) so that we obtain the final expression for the equilibrium price under CARA preferences:

$$p_t = \frac{n_{1,t}[p^* + \nu(p_{t-1} - p^*)] + n_{2,t}[p_{t-1} + \gamma(p_{t-1} - p_{t-2})] + \epsilon_t}{(1 + r_f)} \quad (43)$$

Where a white noise term, $\epsilon_t \sim N(0, \sigma_\epsilon^2)$, is added because, notwithstanding the fact that this equation can generate complex dynamics, in most cases these dynamics do not look as “random” as actual stock prices.

If preferences were of CRRA type the equilibrium price would be:

$$p_t = \frac{l_{1,t}[p^* + \nu(p_{t-1} - p^*)] + l_{2,t}[p_{t-1} + \gamma(p_{t-1} - p_{t-2})] + \epsilon_t}{(1 + r_f)} \quad (44)$$

5. Investors' adaptation

Also if there are examples of models where group's participation does not change, in the great majority of the agent-based models the percentage of individuals belonging to one group or the other is not fixed and changes according to either a hypothetical stochastic process, as in Kirman (1993), or by specifying a switching mechanism based on the relative performance of the strategies followed by the different groups, as in Brock-Hommes (1998). Borrowing the terminology from Ecology, we say that investors follow an adaptation strategy.

The approach based on the evaluation of the relative performance through a so-called "fitness function" is flexible and, among other things, it allows to build a bridge with the machine-learning statistical field: the fitness function can in fact be interpreted as the equivalent of the gain-loss functions used in machine learning. However, this approach introduces a significant level of complexity and, if the goal of the analysis is simply to show how a certain statistical departure from the EMH can be justified by the migration from a strategy to the other, the simpler alternative of a death-birth stochastic process governing this phenomenon is probably a better choice. In any case, we will show how to use the two approaches depending on the analysis' purpose.

5.1. The fitness function

In this approach, investors compare the different strategies and decide which one to adopt. The evaluation is based on the performance of each strategy, as measured by its realized excess profits at time t , $\Pi_{h,t}$. According to equation (8):

$$\Pi_{h,t} = W_t - W_{t-1}(1 + r_f) = z_{h,t-1}R_t \quad (45)$$

This is the product of the individual number of shares held at time t , $z_{h,t-1}$ (determined in the period before), multiplied by R_t .

It should be noted that eq. (45) is the building block of the original form of the fitness function proposed by Brock-Hommes (1998) and Gaunersdorfer-Hommes (2000). However, this specification, especially under CRRA preferences, can generate explosive computational issues simply because of base effects, when profits are proportional to wealth. For this reason, especially when using CRRA preferences, we

prefer to use the rate of profits instead of profits in eq. (45). In mathematical terms, this implies dividing by p_{t-1} the r.h.s. of eq. (45).

The investors evaluate for each strategy a “fitness” function, which is based on past performances, too:

$$G_{h,t} = \Pi_{h,t} + \eta G_{h,t-1} \quad (46)$$

where the so - called “memory” parameter η measures the use of past fitness values in the calculation of actual fitness. The higher the value of $G_{h,t}$, the higher should be the proportion of traders’ population willing to use the underlying strategy. Acting in this way, investors show a sort of “bounded” rationality.

Fitness might be measured not only by its absolute performance but also by its volatility and then can be measured in a risk-adjusted way. The risk-adjustment can be made perfectly coherent with the preferences of the agents by simply using the utility function as the fitness function. But we don’t think that the higher complexity that these types of risk-adjustment induce is particularly useful, so we will stick to the pure performance interpretation of the fitness function.

The fitness function is the basis to determine the number of investors preferring one strategy to the other. Since the proportion of traders must be a number between 0 and 1, we can resort to a logistic-type functional form as in Brock-Hommes (1998):

$$n_{h,t} = \frac{e^{\beta G_{h,t-1}}}{Z_{t-1}}, \quad h = 1,2 \quad 0 \leq \beta \quad (47)$$

$$Z_{t-1} = \sum_{h=1}^2 e^{\beta G_{h,t-1}} \quad (48)$$

The variable Z is a simple normalization term and the parameter β is called the “choice intensity”. It should be noted that the parameter β plays a very important role in defining the non-linear dynamics of the system. If the choice intensity is zero, the population will be split 50-50 among the two groups and will never change. If the intensity is instead infinite, the population will migrate completely from one winning strategy to the other. Then, the higher the value of the choice intensity and the higher is the probability of observing market overreaction and herd behaviour.

To avoid financial bubbles, a technical correction to the formula used for the chartist proportion is necessary:

$$n_{2,t}^* = n_{2,t} e^{-\frac{(p_{t-1} - p^*)^2}{\alpha}} \quad \alpha > 0 \quad (49)$$

In this way, if chartists drive away too much the price from its fundamental value, the exponential correction term dampens the number of chartists and their influence on the price dynamics.

5.2. Birth-Death process

In a typical birth-death process the state of the system is measured by the number, k , of individuals with a certain characteristic (a disease, a genetic trait ...). The total number of individuals is constant at N and at each time the number k can increase by 1 (“birth”) or decrease by 1 (“death”) or stay put. The resulting Markov process is memoryless: only the current state is relevant for the future states.

It should be clear that the “birth” and “death” terms must be interpreted very loosely. Probably, “mutation” would be more appropriate, if we are inside a genetic application, or “recruitment”, if we are inside a social application. But, in any case, it is important to note that the change of state has not to be irreversible, as the birth-death terms would imply. For example, if we apply those models to epidemiology, the infected individual can change status once immunity has been obtained, or a vaccine has been administered.

Another thing not to forget is that the characteristics that measure the state of the system have not to be “physical” or “genetic”. For example, we can divide the population in two groups according to their preferences, for example those that prefer “beer” and those that prefer “wine”. In this case, the birth - death processes can be helpful to study the oscillations in the preferences inside a certain population through some “recruiting” or “influencing” functions. And the “preferences” have not necessarily to be limited to human individuals. The Kirman (1993) model that we will apply to trading activity, for example, has been originally developed for modelling the foraging and recruiting activities inside ants’ communities.

As in most birth-death processes, Kirman (1993) follows the original “urn” model of Ehrenfest (1907): at regular time intervals one individual is randomly selected and he

can be of type 2 with probability $\frac{N-k}{N}$ and of type 1 with probability $\frac{k}{N}$. Once

selected, the individual might mutate in the other type according to a certain rule and, if this happens, the state of the system changes. To model herding behaviour, Kirman (1993) elaborated a specific birth-death process that is sufficiently flexible to generate many different types of limiting distributions for the state of the process.¹⁶ In his model, the transition probabilities are¹⁷:

$$\text{Prob}(Z_{t+1} = k + 1 | Z_t = k) = \frac{N - k}{N} \left(\frac{k}{N} (1 - \mu) + \epsilon \right) \quad (50)$$

$$\text{Prob}(Z_{t+1} = k - 1 | Z_t = k) = \frac{k}{N} \left(\frac{N - k}{N} (1 - \mu) + \epsilon \right) \quad (51)$$

and

$$\begin{aligned} \text{Prob}(Z_{t+1} = k | Z_t = k) \\ &= 1 - \text{Prob}(Z_{t+1} = k - 1 | Z_t = k) \\ &\quad - \text{Prob}(Z_{t+1} = k + 1 | Z_t = k) \end{aligned} \quad (52)$$

The first term on the right side of the equations is derived from the usual Ehrenfest's sampling. For example, in the first equation it is the probability of selecting a type 2 individual. This individual can meet an individual of type 1 (with probability k/N) and be "recruited" by him with probability $(1 - \mu)$. If this happens, it will change from type 2 to type 1, and the state of the system will change from k to $k+1$.

The ϵ term is introduced to avoid that the system gets stuck in state 0 or N . It can be considered a sort of "self-conversion" factor that operates independently of the recruiting activity of other types of individuals.

¹⁶ Kirman (1993) was inspired by research on herding behaviour in ants' colonies.

¹⁷ In the original formulation of Kirman (1993), the denominator in the fraction contained inside the rightmost parenthesis is $N-1$ and not N . But, for simplicity and ease of comparison with other similar models, we preferred to use N .

6. Simulating adaptation with an explicit fitness function

We will start in section 6.1 with the Gaunersdorfer-Hommes (2000) model, where investors are either chartists or fundamentalist, but they share the same CARA preferences. This application is particularly simple to work with from a computational point of view: wealth is affected by the stock prices dynamics, but under CARA preferences wealth does not affect the demand for stocks and this feature allows to avoid explosive and implosive path for the stock price.

On the contrary, if we assume the more realistic CRRA preferences, the stability of the system is much more difficult to achieve, because demand for stocks is not independent from the wealth level. We will see in section 6.2 how to build models where agents have CRRA preferences. CRRA preferences imply that the portfolio allocated to equity is a fixed proportion of wealth, so they are particularly useful to represent the category of the so-called “asset allocator” investors. It should be noted that almost all forms of collective investment undertakings belong to this category. A mutual fund of whatever nature (active or passive, balanced or specialized, generic or dedicated to retirement savings), for example, will not change the proportion of the portfolio allocated to equities just because there has been a change in the asset under management.

In section 6.3 we will extend the previous model to consider asset allocators pursuing investment strategies with different degrees of riskiness. In this application, the agents are balanced mutual funds with different equity-bond allocation benchmarks. As before, inside each balanced fund category, there could be chartist or fundamentalist asset managers.

In all three applications, the percentage of agents belonging to one group or the other is not fixed and changes according to the relative performance of the strategies followed by the different groups, as in Gaunersdorfer-Hommes (2000). In their model, at each time step, the investors chose the strategy to follow, according to equation (49), and determine the individual demands for equity, according to equation (27). The aggregated stock price is, then, calculated according to equation (43) and the next step follows.

Following previous studies, we take as the time unit the (working) day and we consider 10 years of data. The observations will then be equal to 2520, assuming there are 252 working days inside a calendar year. To calibrate the model, we will choose the period

2015 - 2023.¹⁸ In that period the average annual risk-free interest rate, r_f , has been equal to 1.66% and the average dividend yield, y , has been equal to 1.77%. The stock returns volatility, σ , has been equal to 18.58% while the equity premium has been on average equal to 10.07%.¹⁹ If a representative investor exists and if preferences are of the mean - variance type, we can deduce an implied coefficient of risk aversion equal to 2.91.

6.1. Chartists vs Fundamentalists in a CARA setting: the Gaunersdorfer-Hommes (2000) model

Let's assume CARA preferences and replicate the Gaunersdorfer - Hommes (2000) model. As we said many times before, the CARA assumption is very unrealistic, but it is easier to work with in a simulated environment. It represents then a very good starting point for successive generalizations.

As one can see from the equations previously derived, the dynamics of the system are determined by:²⁰

a deterministic, chaotic component that is endogenous and governs the population migrations from one type to the other through the fitness function.

two random noise terms, affecting the pricing (ϵ_t) and the dividend (δ_t) equations.

To appreciate the contribution of the deterministic and stochastic components to the system's dynamics, let's focus on the first by setting to zero the two random noise variables (ϵ_t, δ_t).

As one can see from Figure 1, certain parametrizations can generate cycles and volatility clustering in stock returns.²¹

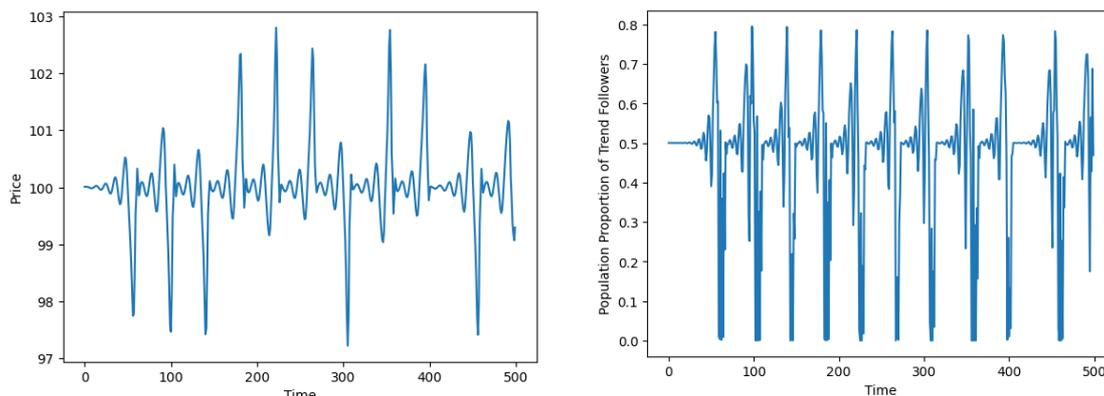
¹⁸ The data are taken from Prof. Damodaran webpage: <https://pages.stern.nyu.edu/~adamodar>. In general, the annual averages are calculated as geometric mean. The stock market is represented by the S&P500 and the risk-free by the 3m TBills.

¹⁹ Federal Reserve Bank of St. Louis (FRED) database, VIXCLS index in <https://fred.stlouisfed.org/series>, starting from 1/1/2000 to 6/12/2025.

²⁰ In appendix 2, we will illustrate the chronological order in which the adaptive system loops.

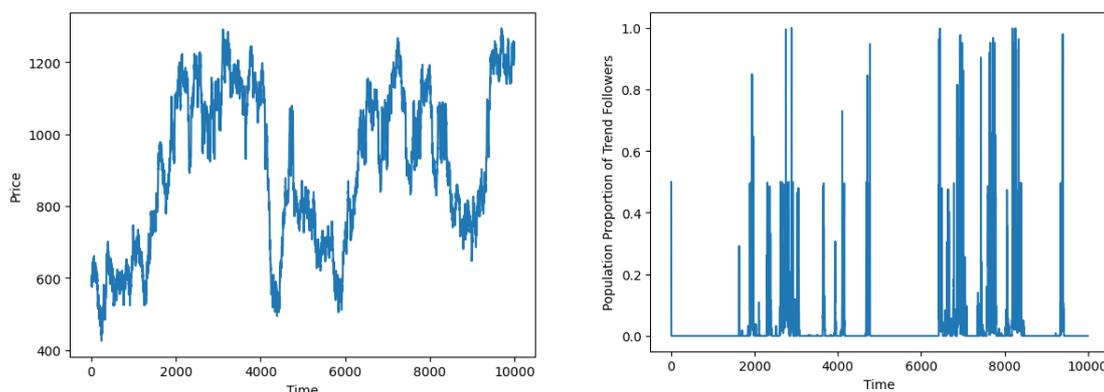
²¹ The simulation is run for 500 daily observations.

Fig. 1 – The deterministic chaotic dynamics



However, it is immediately apparent that the depicted behaviour is not all like a “real” picture of a daily stock price’s time series. For this reason, we need to introduce random noise inside the system. But before running the simulations, it is necessary to spend some time around the issue of calibration of the parameters.

Fig. 2 – Random and chaotic dynamics



As in Gaunersdorfer-Hommes (2000), we can assume that the dividend yield is constant ($\sigma_\delta = 0$) and the random noise appears only in the price equation ($\sigma_\epsilon > 0$),²² one of the simulations is represented in Figure 2, where it can be seen a picture that now looks like a “real” stock price.

It should be noted that, independently of the assumption $\sigma_\delta = 0$, the stock price and dividend up to now are considered stationary variables. This assumption appears unrealistic if the simulation horizon is longer than a few weeks. We will then model the underlying fundamental value of the listed stock as a geometric Brownian motion,

²² We will assume a value of σ_ϵ equal to 1, as in Gaunersdorfer-Hommes (2000)

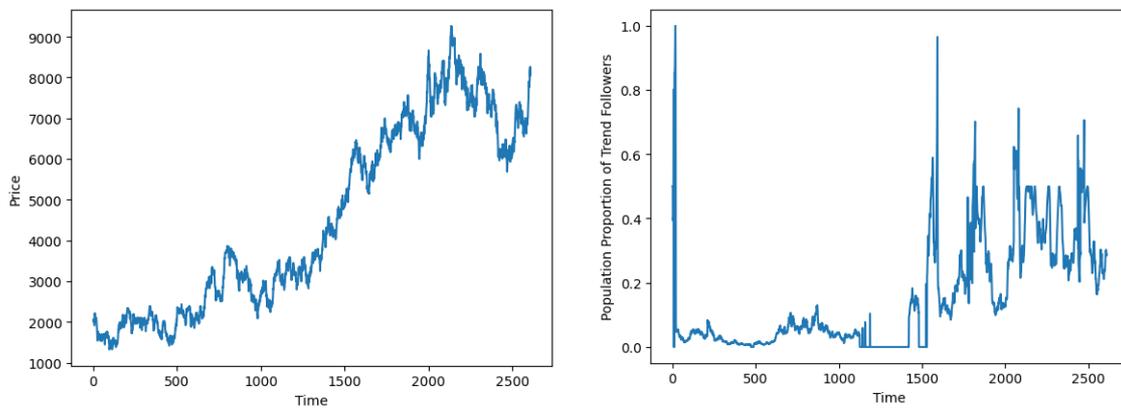
$$dp_t^* = \vartheta p_t^* dt + \sigma p_t^* dW_t \tag{53}$$

with drift ($\vartheta=11.73\%$) and volatility ($\sigma=18,58\%$) estimated on the reference period. So, the expected fundamental price will be:

$$E(p_T^*) = p_t^* e^{(\vartheta - \frac{1}{2}\sigma^2)(T-t)} \tag{54}$$

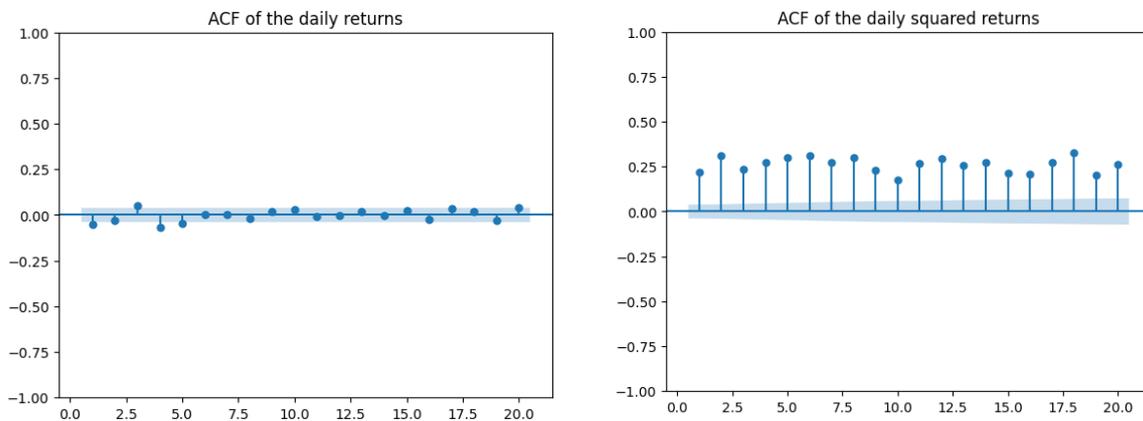
In Fig. 3, an example of simulation of 10 years of daily stock prices is depicted, alongside the movement in the population of investors.

Fig. 3 – 10 years of random and chaotic dynamics with stochastic fundamental



The interesting thing is that the model can reproduce many anomalies that have been described in the finance literature. For example, overreactions and volatility clustering in absence of a significant autocorrelation structure in the stock returns. In Fig. 4, we depict the autocorrelation of returns and the autocorrelation of the squared returns.

Fig. 4 – Autocorrelation function of simulated returns and squared returns



6.2. Chartists vs Fundamentalists in a CRRA setting : asset allocators as agents

The main issue with the Gaunersdorfer-Hommes (2000) model is in the choice of CARA preferences. It is evident that the CARA assumption is totally unrealistic if the time horizon is longer than a few months and the wealth of investors is unevenly distributed. In fact, it has been proven that the demand for equity depends on wealth both cross-sectionally and intertemporally. Moreover, if we want to populate financial markets with professional agents, investing on behalf of not-professional individuals, CARA preferences adoption excludes almost all forms of the collective investment undertakings. A mutual fund of whatever nature (active or passive, balanced or specialized, generic or dedicated to pension), for example, will not change the proportion of the portfolio allocated to equities just because there has been a change in the asset under management.

In the following we will model an economy where financial agents are exclusively mutual funds of so-called “asset allocation” type. Investors delegate those agents to manage their savings, imposing contractually some limitations to the extent of risk that the manager can take on their behalf. Those limitations in our model take the form of a range in the admitted equity exposure. Those limitations not only reflect the reality of a great part of the asset management industry²³ but, as we will see, they have the technical advantage of avoiding the explosive behaviour observed in the simulations where there are no constraints on the size of the leveraged positions that the agents can take.

In real financial markets short position or leveraged long position are obviously possible but they are not the “norm” for the great majority of investors, and in any case, they cannot take an unlimited size. Think about the situation in which the stock market price is above the fundamental value. Fundamentalists believe in mean reversion toward the fundamental value and then expect a negative performance of the stock market. If we don't impose a constraint on the position that a fundamentalist would take, this situation would bring to a short position. And a similar situation might arise for chartists extrapolating a downward trend in prices.

²³ According to the Mifid regulation, the asset managers must be explicit about the level of risk of the fund commercialised to generic investors. It is considered a fundamental duty of the asset management company and of the funds' distributor (typically a bank), to make the clients aware of the risk, also if they don't possess any statistical knowledge. In order to avoid a breach of this fiduciary duty, the riskiness of the underlying portfolio is subject to the control of risk management and compliance offices that are independent of the portfolio manager. If a breach occurs and it is tolerated by the asset management company, the regulatory authorities can punish the company and all the responsible managers with hefty penalties and at the limit banning them.

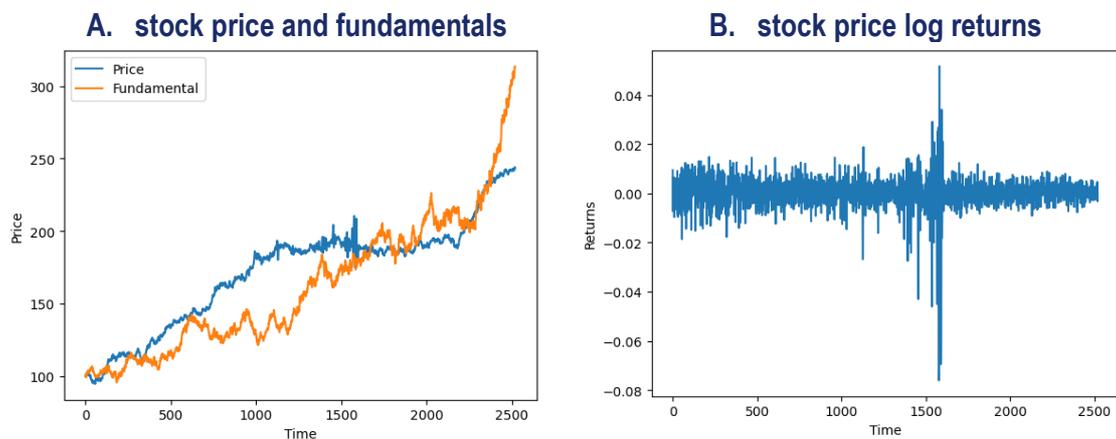
Mutual funds of so-called “asset allocation” type are one of the oldest and most popular investment vehicles offered to retail investors. And their selling proposition has been adopted also by discretionary mandates, pension funds and life-insurance. Investors delegate those “agents” to manage their savings, but they impose contractually some limitations to the extent of risk that they can take. Those limitations are usually represented by a benchmark asset allocation (say 60% equity and 40% bonds) and a range for the maximum and minimum equity exposure (say 80% - 40%) in case the manager is particularly bullish or bearish about the short-term stock market perspectives.

6.2.A. Single benchmark case

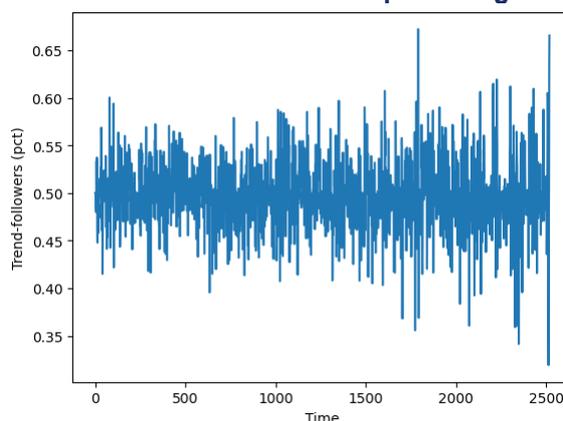
As a first experiment, we will consider a stock market participated only by mutual funds with a benchmark 60% equity - 40% bonds. We assume that the range for the equity exposure is 40% - 80%. The parametrization is identical to the one of section 5, except for the changes that must be implemented having assumed CRRA preferences.

If a manager’s forecast would imply an exposure to equity outside of the admitted range, we assume that the manager will chose for positioning the asset allocation at the relevant limit of the admitted range. As in section 5, (i) managers migrate from a strategy to the other according to a fitness function, that is identical to the one of eq. (46) with the exception that we divide the r.h.s. by p_{t-1} , and (ii) the stock price is determined by the equilibrium between the aggregate demand for shares and a fixed zero net supply, according to eq.(44).

Fig. 5 – The “regular” case

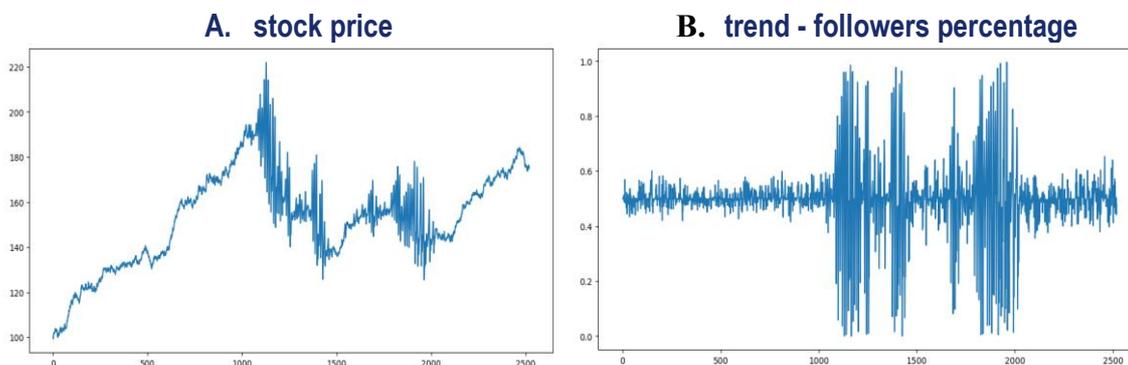


C. Trend - Followers percentage



In general, the simulated stock price does not differ much from the fundamental price. In fig. 5A we have selected an example of what we called a “regular” case, where the final simulated stock price and the fundamental value are close.²⁴ As we can see, the behaviour of the two timeseries looks quite similar and this is explained by the fact that the percentage of the investors following the two strategies is quite stable. In fact, the only period during which the stock price seems to exhibit a higher volatility (see Fig. 5B, after 6 years, i.e. soon after the 1500th data point) is when the migration from one strategy to the other is more pronounced (see Fig, 5C).

Fig. 6 – The “abnormal” case



However, there are also some cases where the migration between the two strategies causes episodes of extreme volatility. An example is in Fig. 6, where one can immediately see that there are at least 4 periods during which the rapid movements of investors from one strategy to the other causes oscillations in stock prices that are well above the standard level. In the example depicted in Fig. 6, in fact, the overall volatility of returns is 30.6% (against the 18% level of the fundamental prices).

²⁴ We assume that they are close if the distance between their respective end points is around the 50% of the fundamental end-value standard deviation.

6.2.B. Multiple benchmarks case

In the previous paragraph we considered a financial market participated only by mutual funds with the same benchmark: 60% equity - 40% bonds. This percentage of equity exposure characterizes the so-called “balanced” funds²⁵, but asset allocation funds have a variety of benchmarks dedicated to investors with different risk tolerances. Funds with an equity exposure of 20% are sometimes denominated “bond balanced” and are dedicated to the more risk averse investors. Funds with an equity exposure greater than 60% are called “equity balanced” funds and are offered to the most risk tolerant investors.

We will then assume that the financial market is populated by four types of “balanced” funds, whose benchmark equity exposure is listed in table 1. The assets under management are initially the same for all types of funds and this means that we are implicitly assuming that investors’ wealth is uniformly distributed, as their risk attitude.

Tab. 1 – Mutual funds benchmarks

Mutual Fund Type	Minimum Equity	Benchmark Equity	Maximum Equity
Bond Balanced	0%	20%	40%
Balanced –	20%	40%	60%
Balanced +	40%	60%	80%
Equity Balanced	60%	80%	100%

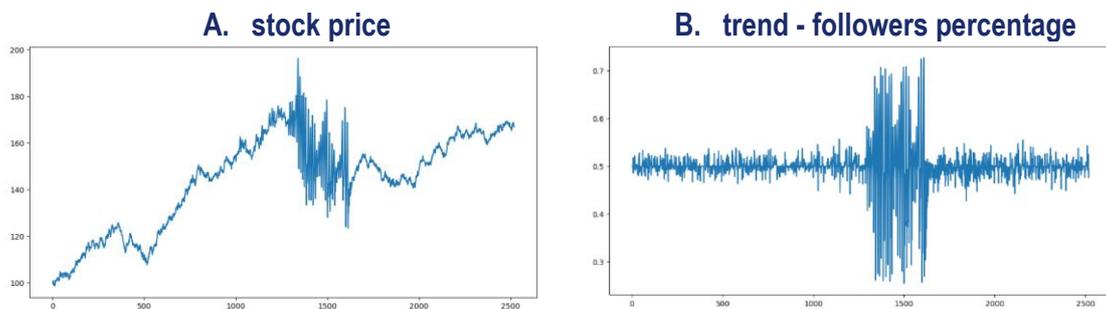
As in the previous paragraph, fund managers are divided into chartists and fundamentalists, and they will migrate from one strategy to the other according to the usual fitness function. They can modify the asset allocation, but they must respect the min-max range for the equity exposure of the fund that they manage. The only material difference with respect to the previous models is that the market equilibrium equation is the result of the aggregation of the equity demands of 8 groups of agents and not of just 2 groups. In fact, we have 4 types of benchmarks and 2 types of strategies.

The diversification over more groups of agents reduces slightly the volatility of the market prices, but, since the underlying source of extra “noise” is always the same

²⁵ In the US the 60-40 is considered “balanced”. In the more conservative Europe, probably the 40-60 corresponds to the idea of a “balanced” solution.

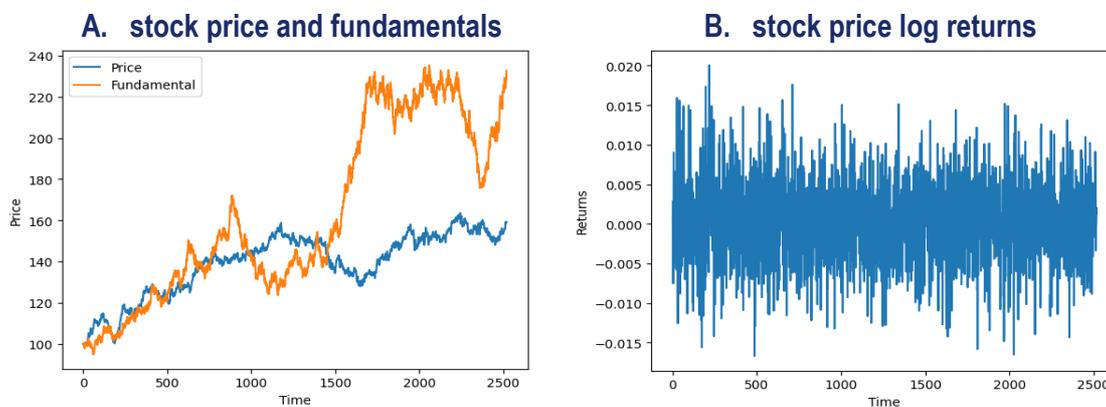
(chartist vs fundamentalist), we still observe episodes of abnormal instability caused by the changes in strategies. One example of such episodes is depicted in fig. 7. Note that in Fig. 7b the percentage of “trend-follower” is summed over the 4 types of balanced mutual funds.

Fig. 7– Simulation “abnormal”

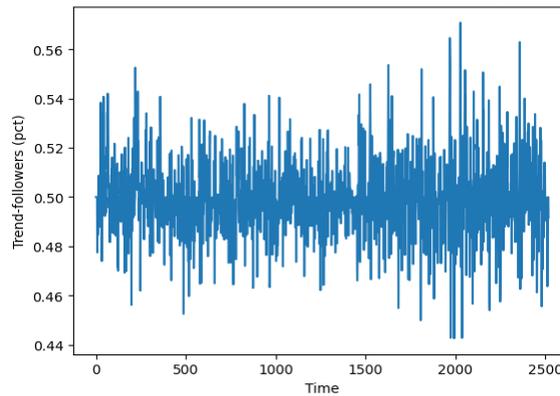


In “regular” simulations those episodes are much less frequent and violent. For example, in Fig. 8, we can see an example where the final stock price and the final fundamental value are close. Since there are not significant migrations, the picture is like Fig. 5.

Fig. 8 – Simulations “regular”



C. Trend - Followers percentage



7. Simulating adaptation without a fitness function

In some circumstances the use of an explicit fitness function can make little sense. If we want to simulate a market at a frequency higher than the daily one, the idea that agents evaluate the competing strategies' performances at every time step (say every *minute*) through the computation of a fitness function seems unrealistic. In those cases, probably the best choice is to use a birth-death process for simulating the migration of the agents from one strategy to the other. The adaptation process is then implicit in the birth-death parametrization, instead of being derived from an explicit function as we did previously in section 6.

In this section we will simulate a financial ecosystem where interactions among economic agents occurs at a frequency higher than the daily one, so that we can analyse phenomena related to intra-day volatility, such as the differences between the minimum and the maximum of the trading sessions, the close-open jumps or the weekend effect.

How interaction among investors occurs? Differently from the previous model, where a fitness function determines the numerosity of the investors' groups, we will model agents' interactions directly through a birth-death process of the type analysed and applied by Kirman (1993) to model imitation and recruiting behaviour. As in Kirman (1993) we will then assume that each period two investors meet and from their interaction one of the two can change "mood".

7.1. High-frequency changes in investors' mood

In the previous section, we assumed that the only source of heterogeneity among investors is the estimate of the future stock price. But the assumption that investors share the same, constant estimate of stocks' volatility is at odds with the movements observed in the VIX index, the so-called "fear" indicator. It can then be interesting to consider the estimate of the volatility parameter as the source of heterogeneity.

As in the applications of section 6, we assume that investors have the same preferences (CRRA) and risk-aversion coefficient. They agree on the expected equity return, but they differ in their "mood", i.e. in their assessment of the uncertainty surrounding stock price dynamics. We will then divide them in "optimists" (Type 1: estimated low variance of returns) and "pessimists" (Type 2: estimated high variance). This difference in mood will bring them to a different portfolio allocation. If investors mutate their perception of the level of uncertainty surrounding the future, they rebalance their portfolio and the aggregate demand will change.

We will assume that trading occurs only when an individual switches of mood. In other words, portfolio rebalancing does not occur when the equity proportion differs from its optimal weight because of price movements.²⁶ If an investor of type 1 mutates in type 2 at time t , his optimal equity position will mutate from $\omega_1 W_t$ to $\omega_2 W_t$. To realize his optimal rebalancing, his net demand for equity will be:

$$(\omega_2 - \omega_1) W_t \tag{55}$$

Ceteris paribus, since $\omega_2 < \omega_1$, the aggregate demand for equity will be lower. It should be noted that the variation in the stock price might generate a general portfolio rebalancing by all investors to restore the optimal weights of equity and bonds. This in turn would determine a new stock price, necessary to restore a market equilibrium, and so on. In other words, a change in demand (and in prices) could trigger a typical cobweb dynamic to reach the final equilibrium point.²⁷ We will not implement this round of "second effects" in our algorithm, because we prefer to highlight just the impact of adopting a birth-death stochastic process on the market volatility.²⁸

²⁶ In this assumption we adhere to the observation that "for many investors, it may be undesirable to follow a traditional asset allocation policy by frequently rebalancing portfolio holdings to regain the specified asset weights", Sharpe W. (2010)

²⁷ The cobweb is stable because if the stock price decreases, then the portfolio rebalancing determines an increase in the demand for equity, and this compensates for the initial reduction.

²⁸ One should also evaluate the cost-benefit of a significant increase of computational time: running 1000 simulations take almost 1 hour on a normal desktop.

As far as financial parameters are concerned, let's use the same values of the adaptive belief systems described in section 6: annual risk - free interest rate, r_f , equal to 1.66%, stock returns volatility, σ , equal to 18.58% while the equity premium equal to 10.07%. The simulations' timeframe is different, and it is described in Appendix 7.

We prefer to characterize the two types of investors by their portfolio allocation. Specifically, we assume that Type 1 investor (the "optimist") allocates 70% to equity, while Type 2 investor (the "pessimist") allocates 30% to equity. The average representative investor will invest 50% in equity. If preferences are CRRA, we can deduce an average implied coefficient of risk aversion equal to 6.20. The 70 - 30 and 30 - 70 portfolio allocations imply that Type 1 forecasts a standard deviation of 16%, while Type 2 assumes a standard deviation of 24%. According to the mean - variance utility function, the optimal portfolio allocation is independent of wealth so it will remain unchanged throughout the entire simulation unless the investor changes "mood", from optimism to pessimism or vice versa.

As in Kirman-Teyssiere (2001), we assume that the initial financial wealth endowment is the same for both types of investors (says 10 euro), so that at the start of the simulation the only thing that distinguishes one investor type from the other is the equity proportion of their portfolio. The investors' population is constituted by a number $N=1,000$ individuals. If the initial individual endowment is 10 euro, the total wealth in our economy at the start of the simulation is equal to 10.000 euro, while the initial net wealth of the market maker is zero. Just to start the simulations with an initial stock price of 100, we assume that the fixed supply of shares is equal to 100. The investors will buy some of them and the remaining will be bought by the market maker.

Fig. 6a –Distribution of investors' type during the entire period

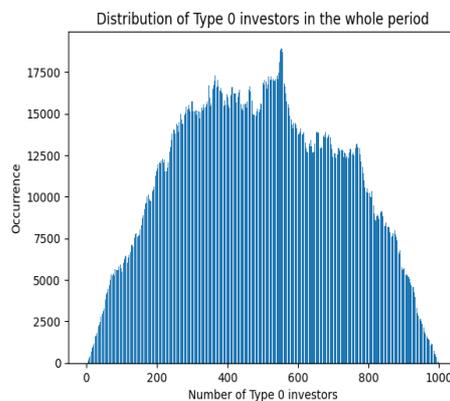
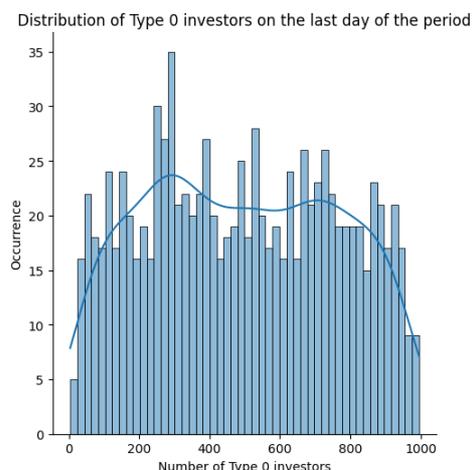


Fig. 6b –Distribution of investors' type during at the end of the period

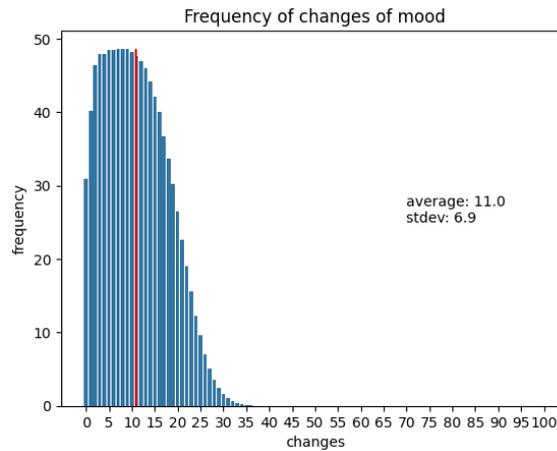


Since the number of data points is not sufficient for the limiting population distribution to emerge, we randomized the starting number of optimist investors in the simulations. During the entire period and at the end of the quarter, the distribution of the number of optimist investors is depicted in figure 6. As one can see, in both cases the distribution is very far from the Normal one.

How many times an investor will change of mood? To answer this question, we ran 1,000 simulations for three months of interactions, i.e. 63 working days and 28 holidays. Each simulation contains a total of 78,540 data points. With the parameters we chose, the same investor on average switches “mood” eleven times during a quarter.

Let's the analyse the impact of those changes of mood on market prices. In this framework, market prices are set according to the eq. (57) - (58). They affect the endowment of the migrating agent in a direct way, because the price at which he sells or buys equity shares is, *ceteris paribus*, lower or higher than the historical price at which shares were originally bought. But they affect indirectly all investors, because the value of the equity component of their portfolio and then of their wealth will change accordingly.

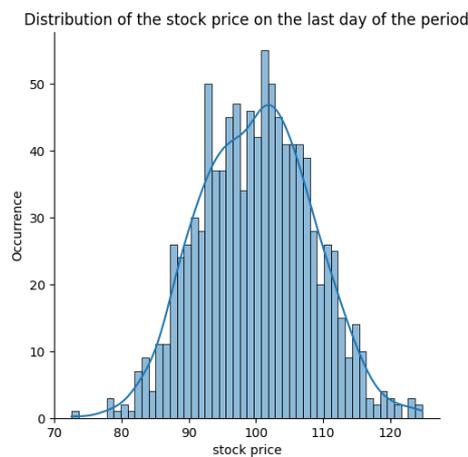
Fig. 7 – Changes in mood (per quarter)



We assume that the fundamental value follows a geometric Brownian motion²⁹ and that the changes in the portfolio composition due to price changes do not trigger a rebalancing. Only when an investor changes mood, the portfolio composition (i.e. the proportion of equity with respect to the total wealth) is rebalanced to reflect the optimal weights. Since in our model an investor changes mood on average more than 11 times per quarter, there are enough trading occasions to avoid the risk of portfolio’s weights drifting far away from their optimal value.

If we focus the attention on the distribution of the stock price at the quarter’s end, we can see in Fig. 8 that the median price is around 103, which is very near to the expected value for the fundamental stock price at the end of the quarter.

Fig. 8 – distribution of the stock price on the last day of the simulated quarter



²⁹ See Appendix 4.

If we resample the simulated path on a daily interval, we can calculate the daily stock returns on the closing price. The annualized volatility induced by the Kirman stochastic process is on average equal to 16.77%, ranging from a minimum of 11.82% and a maximum of 22.44% in the 1,000 simulations we ran. The mood changes among the investors' population adds on average 1.1% to the return's standard deviation in the simulations.

If we analyse the autocorrelation of the daily returns (see Tab.2), we note that on average the values are not significantly different from 0. But if we take squared returns or absolute returns, the ACF shows much higher values, and the first autocorrelation is positive in both cases. This could be coherent with documented forms of volatility clustering.

Tab.2 – First autocorrelations of the daily returns (average over 1000 simulations)

	Return	Abs Return	Squared Return
t+1	- 0.012571	0.093078	0.033014
t+2	- 0.007894	- 0.136091	- 0.060157
t+3	- 0.007981	- 0.119451	- 0.053619
t+4	- 0.013210	- 0.121306	- 0.054930
t+5	- 0.010664	- 0.132843	- 0.058102
t+6	- 0.003060	0.082048	0.027854
t+7	- 0.015274	0.273863	0.101361
t+8	- 0.006083	0.086754	0.031522
t+9	- 0.011514	- 0.123928	- 0.054100

Those high-frequency simulations can be useful to verify if the underlying model can reproduce some intra-day observed phenomena. For example, if we use the “close” prices to compute the daily volatility, the latter can be decomposed in two parts: the first is the one originating just from the news impact (measured by the difference between the “open” price and the previous session’s “close” price) and the second

originating also from the trading activity (measured by the difference between the “close” price and the “open” price).

Tab. 3 – Open/Close Volatility Decomposition

	Simulations			S&P500 – Quarterly periods (2015 - 2024)		
	Open _t – Close _{t-1}	Close _t – Open _t	%	Open _t – Close _{t-1}	Close _t – Open _t	%
Volatility (annualized)						
mean	1.06%	19.9%	5.3%	6.27%	15.53%	40%
min	0.06%	14.03%	0%	1.13%	5.57%	20%
max	1.8%	26.78%	6.7%	27.71%	57.37%	48%

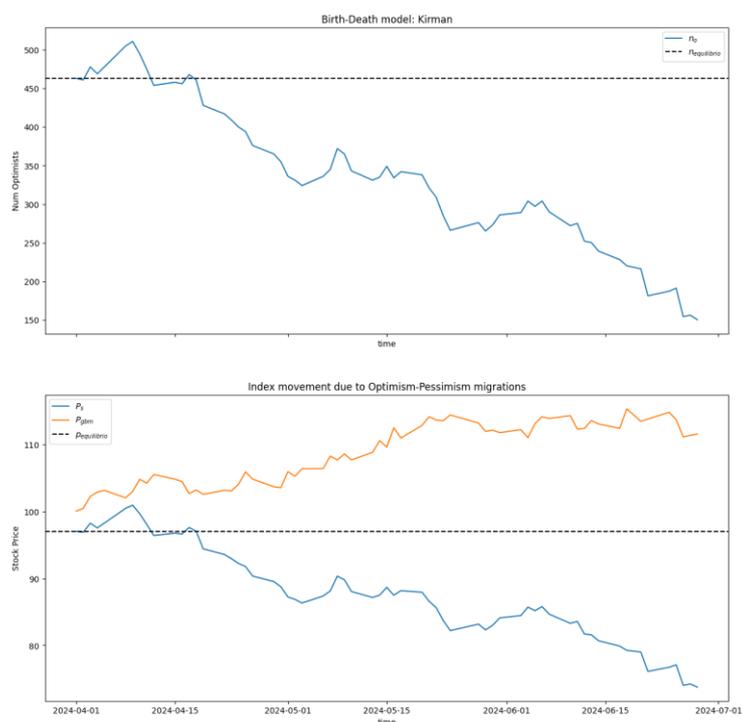
As far as the weekend effect is concerned, in real data volatility on Monday is slightly larger than on the other days (18.89% vs 17.91%). But in our simulations, there is no weekend effect: volatility on Monday is almost identical to the ones of the other days (19.74% vs 19.95%). This could simply be the result of a small number of simulations: just 200 “Monday” out of 1000 weekdays).

If we consider the min-max intraday range³⁰, the S&P500 shows an average value of 1.24% (measured with respect to the simple average of the min and max values) and a maximum value of 10.89%. In our average simulation, this range is equal to 1.93% with a maximum value of 3.7%.

In figure 9, we report an example of simulated path for the number of optimist investors and the market stock price. As one can see, in this simulation the progressive decrease in the number of optimist investors determine a negative impact on the stock price, counterbalancing the positive dynamics of the fundamental price. Eventually, the stock price dynamics in this simulation is negative because of the prevailing pessimist mood among investors.

³⁰ Measured as the ratio between the max and minimum intraday price with respect to the average between the max and the min price.

Fig. 9 – Number of optimist investors and market prices in a simulation



Legenda. Top panel: the simulated number of optimist investors. Bottom panel: the “fundamental” stock price determined a Geometric Brownian motion (orange, upper line) and the market stock price, determined also by the investors’ mood changes.

8. Conclusions

Sixty years after the foundational contributions of Samuelson (1965) and Fama (1965), the Efficient Market Hypothesis continues to serve as a central reference point for financial economics and asset pricing. Yet, the increasing use of many-agent adaptive models offers a complementary path: one where markets can be simulated in controlled, “laboratory-like” environments that explicitly account for investor heterogeneity.

In this paper, we explored three distinct approaches to agent - based adaptive modelling, all rooted in the seminal Brock-Hommes (1998) framework. The first reproduced the Gaunersdorfer-Hommes (2000) model; the second modified preference assumptions to align portfolio choices more closely with those of a representative asset allocator; and the third introduced higher-frequency trading through a birth-death stochastic process of investor interactions, in the spirit of Kirman (1993).

Taken together, these approaches highlight both the flexibility and the richness of adaptive financial ecosystems. By grounding simulations in economic theory while

preserving Karl Popper's principle of falsifiability, agent-based models bridge the gap between abstract theory and the observable complexity of real markets. This adaptability marks a step toward more realistic and empirically testable representations of financial dynamics.

Appendix 1 – Preferences and optimal portfolio allocation

In this appendix, we show that the relationship between preferences, mean-variance optimization and demand for stocks. Since many papers in the literature on agents-modelling assume that the readers are “proficient” in utility theory, we thought useful to work out the algebra.

Different measures of risk-aversion

If we define a utility function on wealth, $U(W)$, it makes sense to assume that its slope, i.e. the marginal utility of wealth, $U'(W)$, decreases as wealth increases. This implies that the utility function is concave, i.e. $U''(W) < 0$. If the utility function is concave, the investor shows a “normal” aversion to risk. To obtain neutrality to risk we need to assume a linear utility function, $U''(W) = 0$. To obtain love for risk, we need to assume a convex utility function, $U''(W) > 0$. We will remain in the economics tradition, and we will assume a concave utility function.

Other factors that are useful to characterize the attitude of investor towards risk concern the way in which the risk-aversion changes as wealth changes. In this respect, two measures are particularly useful: the Absolute Risk Aversion (ARA) and the Relative Risk Aversion (RRA). The former measures how much changes the absolute level of the optimal investment in risky asset as wealth increases. The latter measures how much changes the relative (to wealth) level of the optimal investment in risky asset as wealth increases.

Both ARA and RRA are linked to the concavity of the utility function. They measure how much the concavity changes when the slope changes:

$$ARA \stackrel{\text{def}}{=} - \frac{U''(W)}{U'(W)} \quad (\text{A1.1})$$

$$RRA \stackrel{\text{def}}{=} W * ARA \quad (\text{A1.2})$$

Functions that exhibit constant ARA (CARA) or constant RRA (CRRA) are quite easy to work with. As far as simulating investors activity, for example, the CARA preferences are very handy because the equity demand is independent of wealth and then the system is more stable than in the case of a demand which depends on the level of wealth (and the endogenous stock price dynamics). We will see an example in this appendix, where we solve the optimization problem with logNormal wealth distribution.

However, CARA preferences are unrealistic, both historically and cross-sectionally. Wealthy investors invest in stocks much more than poor investors, and the same investor tend to increase the equity investment as he grows older and becomes wealthier. For this reason, in the classical literature on the microeconomics of portfolio management the basic assumption is constituted by CRRA preferences. As we will see in this appendix, the optimal relative allocation to equity is in fact constant and then, when wealth increases, also the absolute demand for equity increases.

CRRA: Power utility functions

A typical utility function that generates constant relative risk aversion (CRRA) is the so-called power utility function:

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma} \quad (A1.3)$$

where $\gamma > 0$. From the definition of Relative Risk Aversion, RRA, we derive in fact that $RRA = \gamma$. If $\gamma = 1$, we obtain as a special case the logarithmic utility function that has been the workhorse of many economics papers: $U(W)=\log(W)$.

CARA: exponential utility functions

A typical family of utility functions that generates constant absolute risk aversion (CARA) is the exponential one:

$$U(W) = -e^{-\alpha W} \quad (A1.4)$$

where $\alpha > 0$. From the definition of Absolute Risk Aversion, RRA, we derive in fact that for this class of functions the $ARA = \alpha$.

Wealth dynamics and distribution

Now, let's assume that there are only two assets, the risk - free bonds and the risky stocks. Wealth, W_t , is allocated to bonds and stocks in proportions ω and $1 - \omega$, respectively.

The bond B_t evolves according to the following equation:

$$dB_t = rB_t dt \quad (A1.5)$$

where r is the constant risk - free rate, and the stock price evolves according to a Geometric Brownian motion:

$$dS_t = \mu S_t dt + \sigma S_t dZ_t \quad (\text{A1.6})$$

where μ is the equity drift term.

Then also the wealth level evolves according to a Geometric Brownian motion:

$$dW_t = \mu_W W_t dt + \sigma_W W_t dZ_t \quad (\text{A1.7})$$

where:

$$\mu_W \stackrel{\text{def}}{=} r + \omega(\mu - r) \quad (\text{A1.8})$$

$$\sigma_W \stackrel{\text{def}}{=} \omega\sigma \quad (\text{A1.9})$$

and ω is the allocation to equity. According to Ito's lemma, the log of wealth evolves according to the following equation:

$$d \log W_t = (\mu_W - \frac{1}{2}\sigma_W^2) dt + \sigma_W dZ_t \quad (\text{A1.10})$$

And that is the same of saying that wealth is log - Normal distributed as:

$$\log W_t \sim N\left(\mu_W - \frac{1}{2}\sigma_W^2; \sigma_W^2\right) \quad (\text{A1.11})$$

This notable property of logNormal distributions is heavily used in stochastic asset pricing, and it is useful to point it out. If a variable W is lognormal, i.e. its natural logarithm is Normally distributed, then the logarithm of its expected value is equal to:

$$\log E(W) = E \log(W) + \frac{1}{2} Var(\log W) \quad (\text{A1.12})$$

This property is extremely useful when we need to maximize the expected utility with respect to wealth, that we proved to be lognormal distributed.

Portfolio optimization under Power utility functions

Consider the power utility function of eq. (A1.3): the maximization of the expected utility is equivalent to the maximization of the log of the expectation:

$$\max E_t[U(W_{t+1})] \sim \max \log E_t\left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma}\right] \sim \max \log E_t[W_{t+1}^{1-\gamma}] \quad (\text{A1.13})$$

If we apply the lognormal property of eq. (A1.12):

$$\begin{aligned} \max \log E_t[W_{t+1}^{1-\gamma}] &= \max\{ (1-\gamma)E_t[\log(W_{t+1})] \\ &+ \frac{1}{2}(1-\gamma)^2 \text{Var}_t(\log(W_{t+1})) \} \end{aligned} \quad (\text{A1.14})$$

Then, dividing by $(1-\gamma)$ and with an easy simplification, we arrive to

$$\begin{aligned} \max\{E[U(W_t)]\} \sim \max\left\{\left(\mu_W - \frac{1}{2}\sigma_W^2\right) + \frac{\sigma_W^2}{2}(1-\gamma)\right\} \sim \max\left\{\mu_W \right. \\ \left. - \gamma \frac{1}{2}\sigma_W^2\right\} \end{aligned} \quad (\text{A1.15})$$

We then proved the equivalent mean - variance optimization. If we want to see it in terms of the underlying assets parameters, it is sufficient to substitute in (A1.15) the equations (A1.8) and (A1.9):

$$\max\{E[U(W_t)]\} \sim \max_{\omega} \left\{ r + \omega(\mu - r) - \gamma \frac{\omega^2 \sigma^2}{2} \right\} \quad (\text{A1.16})$$

The optimal value of ω is:

$$\omega^{\{CRR\}} = \frac{\mu - r}{\gamma \sigma^2} \quad (\text{A1.17a})$$

As one can see, the optimal value of ω is independent of wealth. This then prove that with CRR the optimal proportion of wealth invested in equity is independent of the wealth amount. The portfolio allocated to wealth is:

$$\omega^{\{CRR\}} W_t = W_t \frac{\mu - r}{\gamma \sigma^2} \quad (\text{A1.17b})$$

Portfolio optimization under Exponential utility functions

Consider the exponential utility function of eq. (A1.4). As before, let's use the property that the maximization of the expected utility is equivalent to the maximization of the log of the expectation:

$$\begin{aligned} \max E_t[U(W_{t+1})] \sim \max \log E_t[\\ - e^{-\alpha W_{t+1}}] \sim \min \log E_t[e^{-\alpha W_{t+1}}] \end{aligned} \quad (\text{A1.19})$$

If we apply the lognormal property of eq. (A1.12):

$$\min \log E_t[e^{-\alpha W_{t+1}}] = \max \left\{ \alpha E_t(W_{t+1}) - \frac{1}{2} \alpha^2 \text{Var}_t(W_{t+1}) \right\} \quad (\text{A1.20})$$

Then, dividing by α , we obtain that the maximization of the exponential utility function is equivalent to the maximization of a sort of risk - adjusted expected wealth:

$$\max \left\{ E_t(W_{t+1}) - \frac{1}{2} \alpha \text{Var}_t(W_{t+1}) \right\} \quad (\text{A1.21})$$

Recalling that, in discrete time, the dynamics of wealth is represented by the simple equation (2) that we rewrite her:

$$W_{t+1} = W_t(1 + r) + \omega W_t (r_{eq} - r) \quad (\text{A1.22})$$

So that,

$$E_t[W_{t+1}] = W_t(1 + r) + \omega W_t (\mu - r) \quad (\text{A1.23})$$

$$\text{Var}_t[W_{t+1}] = (\omega W_t)^2 \sigma^2 \quad (\text{A1.24})$$

Then, substituting (A1.23) and (A1.24) into (A1.21), and eliminating the term not containing the ω , the maximization problem reduces to:

$$\max_{\omega} \left\{ \omega W_t (\mu - r) - \frac{1}{2} \alpha (\omega W_t)^2 \sigma^2 \right\} \quad (\text{A1.25})$$

which produces:

$$\omega^{\{CARA\}} = \frac{1}{W_t} \frac{\mu - r}{\alpha \sigma^2} \quad (\text{A1.26})$$

Coherently with CARA preferences, we can immediately see that the wealth allocated to the equity, $\omega^{\{CARA\}} W_t$, is constant and independent of wealth. In fact, the portfolio allocated to wealth is:

$$\omega^{\{CARA\}} W_t = \frac{\mu - r}{\alpha \sigma^2} \quad (\text{A1.27})$$

Appendix 2: the fundamental value for the stock price

We derived in the text the aggregate demand for equity. Consider the special case of identical investors ($H=1$), so that equation (39) becomes:

$$\mu_t - p_t(1 + r_f) = 0 \quad (\text{A2.1})$$

Recalling the definition of μ_t from eq. (30) and solving for the price today:

$$p_t = \frac{E_t[p_{t+1} + y_{t+1}]}{(1 + r_f)} \quad (\text{A2.2})$$

Substituting recursively future prices, we obtain:

$$p_t^* = \sum_{i=1}^{\infty} \frac{E_t[y_{t+i}]}{(1 + r_f)^i} \quad (\text{A2.3})$$

In the literature this is sometimes called the “fundamental” value of the stock price: the discounted present value of all the future stream of dividends.

If we further assume that the expected dividend is constant to y^* , i.e.

$$y_t = y^* + \delta_t \quad (\text{A2.4})$$

where:

$$\delta_t \sim N(0, \sigma_\delta^2) \quad (\text{A2.5})$$

The fundamental stock price of equation (A2.3) reduces to:

$$p_t^* = \frac{y^*}{r_f} \quad (\text{A2.6})$$

If we assume that dividends grow at a rate, g , which is constant but not equal to zero, the fundamental value is equal to:

$$p_t^* = \frac{y_t^*}{r_f - g} \quad (\text{A2.7})$$

This expression is known also as the Gordon (1962) model for stock prices.

Appendix 3 – Technicalities in Gaunersdorfer - Hommes model³¹

In this appendix we clarify two important design choices for implementing the Gaunersdorfer-Hommes (2000) model.

A useful change of variable

The strategy profits can be rewritten in terms of a new variable, representing the difference between actual market price and fundamental price, p^* :

$$x_t \stackrel{\text{def}}{=} (p_t - p^*) \quad (\text{A3.1})$$

This change of variable is particularly useful when computing the dynamics of the system and it is also very easy to apply to all the variables seen. The only exception is probably R_t , so we will derive it step by step. First, we subtract y^* from the actual prices:

$$R_t = (p_t - p^*) + y_t - (1 + r)(p_{t-1} - p^*) + [p^* - (1 + r)p^*] \quad (\text{A3.2})$$

$$R_t = (p_t - p^*) - (1 + r)(p_{t-1} - p^*) + [y_t - r p^*] \quad (\text{A3.3})$$

If we assume that:

$$y_t = y^* + \delta_t \quad (\text{A3.4})$$

and we remember that $p^* = \frac{y^*}{r}$, we obtain:

$$R_t = (p_t - p^*) - (1 + r)(p_{t-1} - p^*) + \delta_t \quad (\text{A3.5})$$

$$R_t = x_t - (1 + r)x_{t-1} + \delta_t \quad (\text{A3.6})$$

The chronological order of the adaptive system loops

Let's use the F letter for denoting the *Fundamentalists'* strategy and the T letter the *trend-follower* strategy. The sequence of the agents' actions when time $t - 1$ ends and time t starts (in programming terms, we are inside the step t of the loop) is the following:

Phase 1: evaluate the strategies

- calculate the realized profits of the two strategies up to time t (excluded):

$$\Pi_{h,t-1} = R_{t-1} d_{h,t-2} \quad h = F, T \quad (\text{A3.7})$$

³¹ In this appendix we follow the website of Blake LeBaron, Brandeis University.

where R_{t-1} and $d_{h,t-2}$ are known and already calculated in the previous loops according to the formulas below:

$$R_{t-1} = x_{t-1} - (1 + r)x_{t-2} + \delta_{t-1} \quad (\text{A3.8})$$

$$d_{h,t-2} = \frac{W_{h,t-2}}{\rho \lambda^2} [\mu_{h,t-2} - p_{t-2}(1 + r_f)] \quad (\text{A3.9})$$

$$\mu_{F,t-2} = E_{F,t-2}(p_{t-1}) = p^* + \nu(p_{t-2} - p^*) \quad (\text{A3.10})$$

$$\mu_{T,t-2} = E_{T,t-2}(p_{t-1}) = p_{t-2} + \gamma(p_{t-2} - p_{t-3}) \quad (\text{A3.11})$$

- evaluate the fitness of the two strategies up to time t (excluded), according to the recursive formula:

$$G_{h,t-1} = R_{t-1}d_{h,t-2} + \eta G_{h,t-2} \quad h = F, T \quad (\text{A3.12})$$

Phase 2: choose the strategy

- depending on their fitness, at time t , investors will choose one of the two strategies, and the population will be split in two proportions determined by:

$$n_{T,t} = \frac{e^{-\beta G_{T,t-1}} e^{-\frac{(p_{t-1}-p^*)^2}{\alpha}}}{Z_{t-1}} \quad (\text{A3.13})$$

$$Z_{t-1} = e^{\beta G_{F,t-1}} + e^{\beta G_{T,t-1}} \quad (\text{A3.14})$$

Phase 3: determine the aggregate demand and the equilibrium price

- determine the forecasts about the price at $t+1$ that the two representative agents of group F and group T will form at time t , before knowing price at time t :

$$\mu_{F,t} = E_{F,t}(p_{t+1}) = p^* + \nu(p_{t-1} - p^*) \quad (\text{A3.15})$$

$$\mu_{T,t} = E_{T,t}(p_{t+1}) = p_{t-1} + \gamma(p_{t-1} - p_{t-2}) \quad (\text{A3.16})$$

- knowing the relative numerosity of the two groups of investors and their demand for stocks, the aggregate demand is computed. Since equity supply is fixed, we close the loop at time t with the calculation of the price that settles the market equilibrium between demand and supply for shares. If we assume that preferences are CRRA, the equation becomes:

$$p_t(1 + r_f) = (1 - l_{T,t})\mu_{F,t} + l_{T,t}\mu_{T,t} + \epsilon_t \quad (\text{A3.17})$$

If we assume that preferences are CARA, the equation becomes:

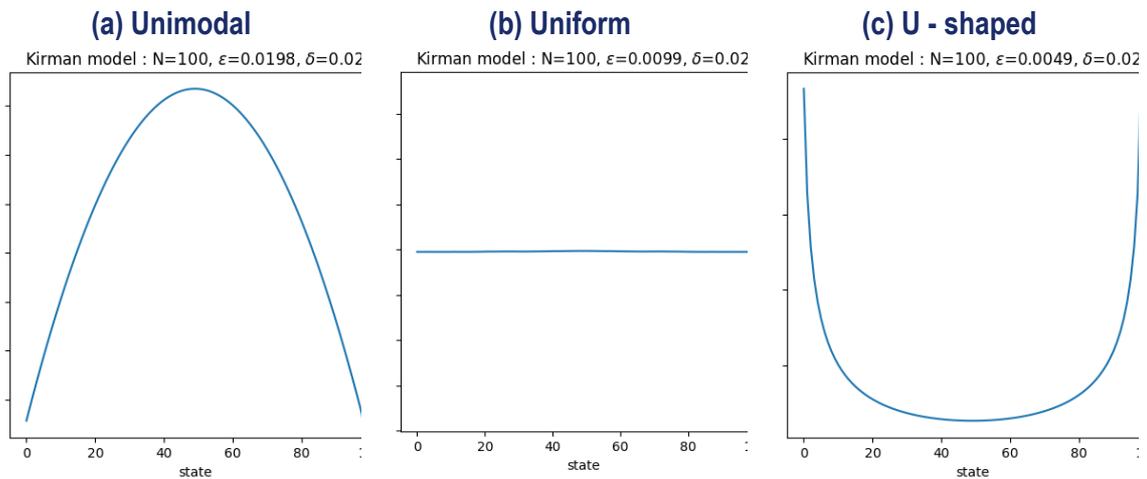
$$p_t(1 + r_f) = (1 - n_{T,t})\mu_{F,t} + n_{T,t}\mu_{T,t} + \epsilon_t \quad (\text{A3.18})$$

- Note that, having discovered the price at time t , we can now use it for the next steps of the loop and for updating the variables x_t, R_t that will be used as the beginning of the $(t+1)$ period.

Appendix 4 – The limiting distribution in Kirman “birth - death” process

As far as the limiting distribution for the permanence of the system in a specific state, Kirman (1993) shows the conditions under which the uniform distribution prevails (Figure 1b). Simple modifications of the parameters bring to U - shaped (Figure 1c) or more regular, unimodal symmetric distributions (Figure 1a).

Figure A4.1 - Limiting distributions for the state permanence of Kirman system



The ants' foraging behaviour is described by a parametrization like the one that generates a Bi-polar distribution. In fact, most of the ants feed at one source but from time to time they switch to the other.

This will be the hypothesis that we will follow to model “optimism” and “pessimism” waves on the stock markets. Kirman model can be easily generalized by changing the “recruiting” efficacy (or “herding” attitude) of the two types of individuals (differentiating the μ parameter in μ_1 and μ_2) and their self - conversion factors (differentiating the ϵ parameter in ϵ_1 and ϵ_2). This generates a broader number of limiting distributions.

Appendix 5 – A geometric Brownian motion for the fundamental value in the heterogenous “mood” specification

If we want to observe not-stationarity in the simulated time series of stock prices, we need to add a drift term through a geometric Brownian motion governing the “fundamental” value, as we did previously. In this case, the market stock price is then equal to the underlying “fundamental” value plus a bid - ask spread. The drift parameter, θ , is observed and known to all investors. Both type 1 and type 2 investors agree on this estimate, but they disagree on the value of the volatility. So, the expected price will be different for the two investors’ type:

$$E_1(P_T) = P_t e^{(\mu - \frac{1}{2}\sigma_1^2)(T-t)} \quad (A5.1)$$

$$E_2(P_T) = P_t e^{(\mu - \frac{1}{2}\sigma_2^2)(T-t)} \quad (A5.2)$$

The market maker is agnostic about the volatility value. He observes the forecasts made by the investors and simply consider a weighted average of the two volatilities, with the weights determined by the relative numerosity of the two types of investors:

$$\sigma_{M,t}^2 = n_{1,t} \sigma_1^2 + (1 - n_{1,t}) \sigma_2^2 \quad (A5.3)$$

The bid and the ask price will then be determined around the expected stock price:

$$E_M(P_T) = P_t e^{(\mu - \frac{1}{2}\sigma_{M,t}^2)(T-t)} \quad (A5.4)$$

The relative composition of the investors’ population will induce the market maker to expect a higher stock price if *ceteris paribus* optimists outnumber pessimists. And vice versa if pessimists outnumber optimists.

Appendix 6 – When a market maker is necessary

In the economies that we simulated there are two types of financial assets, equity and bond, and two types of investors, differing only in their optimal portfolio allocation. In some models, this characterization hides a paradox that is very similar to the one exposed by Sharpe (2009) analysing markets composed only of “asset allocators”.

Asset allocators should rebalance regularly their portfolio to keep the assets’ weights in line with their optimal values. Whatever the weights are, this implies that they regularly sell relative winners or buy relative losers to restore optimality. But, Sharpe (2009) noted, this is impossible to realize in aggregate, because the agents will be all on the same side of the trade.

In the heterogeneity models of section 6, the problem is absent. The difference between market price and fundamental value is always such that the fundamentalists will be the counterparty to the chartists. But in heterogeneity model of section 7, if an investor changes opinion about future volatility, there is no “natural” counterbalance from the other investors.

A solution to this problem can be found by imagining the existence of a “third” agent, the market maker, that buys or sells the shares when the investors ask for them. How can we model the market maker? We could imagine a closed financial ecosystem, where the market maker issues the bonds at the start of the period and with the proceeds buys the shares not bought by the investors. When an investor wants to buy new shares, he must sell his bonds back to the market maker and the market maker will release to him the required number of shares. When the investor wants to sell existing shares, the market maker will issue new bonds and with the proceeds will buy the shares. While the equity shares are bought and sold at different prices, the risk - free bond is instead offered at the fixed interest rate by the market maker based on the demand requested by the investors.

As far as the stock price equation is concerned, we assume that the market maker does not abuse of its position, driving the price to zero if the asset allocators want to sell their shares or to infinite if the asset allocator wants to buy new shares. In other words, the market maker is interpreting its role as a pure financial “intermediary”, offering a sort of “public” service. The service provided by the market maker is remunerated by the difference between the prices at which he buys and sell.

The bid and ask prices can be calculated for transactions generated from a migration from optimism to pessimism (bid price) or from pessimism to optimism (ask price):

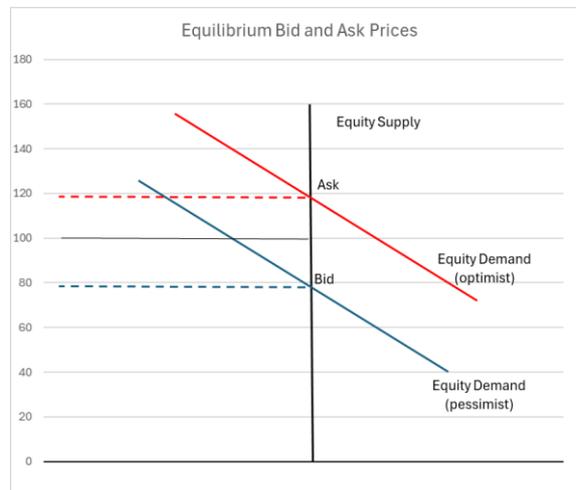
$$S p_t^{BID} = E(p_t | \Omega_{t-1}) S + (\omega_B - \omega_A) W_t \quad (A6.1)$$

$$S p_t^{ASK} = E(p_t | \Omega_{t-1}) S + (\omega_A - \omega_B) W_t \quad (56)$$

where Ω_{t-1} is the information set available before entering the trade. The bid - ask spread will be:

$$p_t^{ASK} - p_t^{BID} = 2 (\omega_1 - \omega_2) \frac{W_t}{S} \quad (A6.3)$$

Fig. A6.1 - Bid and ask prices in an equilibrium model with fixed equity supply



Appendix 7 – A timeframe for intraday simulations

We assume that in a typical calendar year there are “working” days (252) and “holidays” (113).

A “working” days is divided in three parts: (i) “trade” time: eight trading hours and a half; (ii) “sleep” time: seven hours; (iii) “spare” time: the remaining eight hours and a half.³² During the sleep time the birth - death process is suspended, but during the “trade” and the “spare” parts of the working day individuals are actively absorbing information and then can change “type” either by recruitment/imitation or by autonomous news elaboration. We will assume that during “trade” time, when markets are open, both forces are active, whereas during “spare” time only the latter is active. Kirman’s model is sufficiently flexible to allow for a suppression of the recruitment/imitation effect: it is sufficient to put $(1 - \mu) = 0$ and leave $\varepsilon > 0$.

During holidays, markets are closed, but we will assume that part of the day is dedicated to information gathering and analysis, otherwise there would be no reason for observing something like the “weekend effect”.³³ We divide the day in three phases: “sleep” time, “spare” time, and “leisure” time. The latter is dedicated to family, sport or other social activity that have nothing to do with information gathering and that can’t generate a variation in investor’s type. Then, during holidays, the amount of time dedicated to information gathering and potentially causing a change in type is confined to the eight hours and a half of “spare” time. In the remaining sixteen hours the birth - death process is suspended.

We assume that the time interval corresponds to 1 minute, so that in one day there are potentially 1,440 occasions to switch “type”. However, it is only when the investor is actively processing information that he can change type, and it is only when the stock market is open that he can eventually trade with other investors to reallocate his portfolio according to his changed “type”. Summing up, during working days, the switching-type occasions occur during trade time and spare time for a total of seventeen hours, corresponding to 1,020 minutes. During holidays, instead, the minutes in which a switching-type might occur are just 510 per day.

³² We will adopt the European convention for the opening hours of the stock market: open at 9:00 and close at 17:30 CET.

³³ Fortune (1998).

References

- Arthur W. B., Holland J.H., LeBaron B., Palmer R., and Tayler P., (1996) "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market", *Santa Fe Institute Working Paper 1996 - 12 - 093*
- Axtell R.L., Farmer J.D. (2025), "Agent - Based Modelling in Economics and Finance: Past, Present, and Future", *Journal of Economic Literature*, 63(1), 197–287
- Beltratti A., Margarita S. (1993), "Evolution of trading strategies among heterogeneous artificial economic agents", in J. A. Meyer, H. L. Roitblat and S. W. Wilson, eds, *From Animals to Animals 2*, MIT Press, Cambridge, MA
- Black, F. (1986) "Noise", *Journal of Finance* 41, 529–44.
- Brock W. A. and C. H. Hommes. (1998) "Heterogeneous beliefs and routes to chaos in a simple asset pricing model", *Journal of Economic Dynamics and Control*, 22(8 - 9):1235–1274.
- Campbell, J. Y. and A. Kyle (1993), "Smart Money, Noise Trading, and Stock Price Behavior", *Review of Economic Studies* 60, 1 – 34.
- Chiarella, C. and He, X. (2001), 'Asset pricing and wealth dynamics under heterogeneous expectations', *Quantitative Finance* 1, 509–526.
- Chiarella C. and He X. (2008), "An Adaptive Model on Asset pricing and Wealth Dynamics with Heterogeneous Trading Strategies", in: Seese, D., Weinhardt, C., Schlottmann, F. (eds) *Handbook on Information Technology in Finance. International Handbooks Information System*. Springer, Berlin, Heidelberg
- Chiarella C., Iori G. (2002) "A simulation analysis of the microstructure of double auction markets", *Quantitative Finance* 2, 346–253.
- Cutler, D. M, J. M. Poterba and L. H. Summers (1990), "Speculative Dynamics and the Role of Feedback Traders", *American Economic Review*, 80, 63 – 8.
- Day, R. and W. Huang, (1990), "Bulls, bears and market sheep", *Journal of Economic Behavior and Organization*, 14, 299 - 329
- DeLong, B. J., A. Shleifer, L. H. Summers and R. J. Waldmann (1990), "Noise Trader Risk in Financial Markets", *Journal of Political Economy* 98, 703 – 38.
- Deneubourg J.L., S. Aron, S. Goss, J.M. Pasteels, (1987) "Error, communication and learning in ant societies", *European Journal of Operational Research*, Volume 30, Issue 2, 168 - 172
- Ehrenfest P. and Ehrenfest T. (1907) "Über zwei bekannte Einwände gegen das Boltzmannsche H - Theorem.", *Physikalische Zeitschrift*, vol. 8, pp. 311–314
- Fama, E. (1965) "Random walks in stock market prices", *Financial Analysts Journal* 21, 55– 9
- Fortune P. (1998) "Weekends Can Be Rough: Revisiting the Weekend Effect in Stock Prices", *Federal Reserve Bank of Boston working paper* 98 - 6.
- French, K. and Roll, R. (1986) "Stock return variances: the arrival of information and the reaction of traders", *Journal of Financial Economics* 17, 5–26.
- Gaunersdorfer, A. and Hommes, Cars H., (2000) "A Nonlinear Structural Model for Volatility Clustering", available at SSRN: <https://ssrn.com/abstract=241349> or <http://dx.doi.org/10.2139/ssrn.241349>
- Gaunersdorfer A., Hommes C. H., and Wagener F. O. O. (2008), "Bifurcation routes to volatility clustering under evolutionary learning", *Journal of Economic Behavior and Organization*, 67:27–47, 2008
- Gordon, M. J. (1962). "The Savings, Investment, and Valuation of a Corporation", *Review of Economics and Statistics*, 44, 37 - 51

- Grossman, S. and Stiglitz, J. (1980) "On the impossibility of informationally efficient markets", *American Economic Review* 70, 393–408.
- Hautphenne S. and Patch B., "Birth - and - Death Processes in Python: The BirDePy Package", *Journal of Statistical Software*, 2024, Vol. 111(5): 1 - 54
- Kirman A. (1993). "Ants, rationality and recruitment", *Quarterly Journal of Economics*, 108: 137–156.
- Kirman A. and Teyssiere G. (2001) "Microeconomic Models for Long - Memory in the Volatility of Financial Time Series", *Studies in Nonlinear Dynamics & Econometrics*, Vol. 5(4): 281–302
- Krafft O. and Schaefer M. (1993) "Mean passage times for tridiagonal transition matrices and a two - parameter Ehrenfest urn model", *Journal of Applied Probability*, 30, 964 - 970.
- 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
- LeBaron B. (2002) "Building the Santa Fe Artificial Stock Market", Brandeis University <https://faculty.sites.iastate.edu/tesfatsi/archive/tesfatsi/blake.sfisum.pdf>
- Lo A. (2007) "Efficient market hypothesis", in Blume S. and Durlauf L., *The New Palgrave: A Dictionary of Economics* (2nd ed.), Palgrave MacMillan.
- Lo A. (2017) *Adaptive Markets: Financial Evolution at the Speed of Thought*, Princeton University Press.
- Mandelbrot B. and Hudson R.L. (2004) *The (Mis)Behaviour of Markets – A Fractal view of Financial Turbulence*, Basic Books.
- Merton, Robert C. (1969) "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics*, 51, 247–257
- Moran, P. (1958) "Random processes in genetics", *Mathematical Proceedings of the Cambridge Philosophical Society*, Cambridge University Press, vol. 54, 60–71
- Palmer R., Arthur W. B., Holland J. H., LeBaron B. and Tayler P. (1994), 'Artificial economic life: A simple model of a stock market', *Physica D*, 75, 264–274.
- Pasteels J.M. and J.L. Deneubourg (1987) *From Individual to Collective Behavior in Social Insects*, Birkhäuser, Basel
- Popper, K. (1959) *The Logic of Scientific Discovery*, Routledge
- Samuelson, P. (1965) "Proof that properly anticipated prices fluctuate randomly", *Industrial Management Review* 6, 41–9.
- Scharfstein D.S. and J.C. Stein (1990) "Herd Behaviour and Investment", *American Economic Review*, Vol.80, 465 - 79
- Sentana, E. and S. Wadhvani (1992), "Feedback Traders and Stock Return Autocorrelations: Evidence from a Century of Daily Data", *The Economic Journal* 102, 415 – 25.
- Sharpe W.F. (2010), "Adaptive Asset Allocation Policies", *Financial Analysts Journal*, Vol. 66, Number 3, 45 - 59.
- Shiller, R. J. (1984), "Stock Prices and Social Dynamics", *Brooking Papers on Economic Activity* 2, 457 – 98.
- Shiller, R. J. and J. Pound (1986), "Survey Evidence of Diffusion of Interest Among Institutional Investors", *National Bureau of Economic Research, Working Papers* n. 1851.
- Shiller, R. J. (1988), "Portfolio Insurance and Other Investor Fashions as Factors in the 1987 Stock Market Crash", *NBER Macroeconomic Annual*, 287 – 96.

Uppuluri, V. R. R. and Wright, T. (1981), "A note on a further generalization of the Ehrenfest urn model", *Proceedings of the American Statistical Association, Sampling Survey Section*, 564 - 569

Varian H.R. (1986), "Differences of opinion in financial markets", *Proceedings, Federal Reserve Bank of St. Louis*, pages 3 - 40.

Varian H.R. (1992), "Differences of opinion", in J. Eatwell, M. Milgate, and P. Newman, editors, *The New Palgrave Dictionary of Money and Finance*, Macmillan Press, Cambridge.

Biographical note

Marcello Esposito

Graduated with honors in Economic and Social Disciplines (DES) from Bocconi University in Milan; thesis recommended for publication. He holds a M.Sc. and a M.Phil. from the London School of Economics. He was an economist at the Research Department of Banca Commerciale Italiana and later held various positions in some leading Italian financial institutions. He currently runs some start-ups, helps no-profit organizations in big data analysis, and teaches Asset management at Carlo Cattaneo - LIUC University.
