Evology: an Empirically-Calibrated Market Ecology Agent-Based Model for Trading Strategy Search

Aymeric Vié [1 2] **Maarten Scholl** [3 2] **Doyne Farmer** [1 2 4]

Abstract

Market ecology views financial markets as ecosystems of diverse, interacting and evolving trading strategies. We present a heterogeneous, empirically calibrated multi-agent market ecology agent-based model. We outline its potential as a valuable and challenging training ground for optimising trading and investment strategies using machine learning algorithms.

1. Introduction

1.1. Motivation

The theory of market ecology (Farmer, 2002; LeBaron, 2002; Musciotto et al., 2018; Lo, 2019; Levin & Lo, 2021; Scholl et al., 2021) borrows concepts from ecology and biology to study financial markets. Trading strategies are analogous to biological species: they exploit market inefficiencies and compete for survival or profit. (Scholl et al., 2021) has highlighted the nature of interactions between common trading strategies and the strong density dependence of their returns.

1.2. Related work

This article is in the continuity of the rich area of financial agent-based models, and market selection with heterogeneous beliefs (Blume & Easley, 1992; 2006). For example, several ABMs have recently been introduced for market-making optimisation (Spooner et al., 2018), understanding flash crashes (Paulin et al., 2019) and providing sophisticated financial architectures for trading training (Byrd et al., 2019). We attempt to develop the complementary approach of market ecology (Scholl et al., 2021) by focusing on the ecological interactions between the different types of agents and strategies. Our contribution to the existing literature is thus on the interactivity and heterogeneity of the financial agents in a more simplified, low-frequency financial environment with a low-dimensional action space.

1.3. Contribution

In this article, we present Evology: an empirically-calibrated market ecology agent-based model (ABM), in the continuity of the market ecology perspective. We contribute to the market ecology literature (Farmer, 2002; Scholl et al., 2021) by featuring heterogeneous strategy types and sub-strategies: pessimistic and optimistic noise traders and value investors, and trend followers trading over various time horizons. We improve the realism of the resulting wealth flows by validating the model with stylised facts and calibrating investor behaviour. We describe how some particular results of the market ecology model provide a new, exciting challenge for optimising trading strategies using machine learning algorithms. Such simulation-based training can account for interactions and density dependence effects that could be significant and overlooked by traditional time-series training.

2. Model

We consider a population of n agents, acting as investment funds, who trade shares of a single representative asset and cash in the form of a bond. Every period-day t, the funds can buy, sell and short-sell shares of the asset in constant supply. Asset shares pay daily dividends $\delta(t)$ following an autocorrelated Geometric Brownian Motion. Cash/bond yields an interest rate r paid daily. We build over the foundational model of (Scholl et al., 2021), with notable improvements in the heterogeneity of strategies and empirical calibration and adding a new strategy evolution component described in Section 5. We show an example simulation run in the Supplementary Information.

A. Vié designed and implemented the modelling and performed the experiments. M. Scholl gathered and compiled the data used for calibration. MS and D. Farmer provided guidance and help for the modelling. All authors wrote/revised the article. [1]Mathematical Institute, University of Oxford, Oxford, United Kingdom [2]Institute of New Economic Thinking, University of Oxford, United Kingdom [3]Department of Computer Science, University of Oxford, United Kingdom [4]Santa Fe Institute, Santa Fe, New Mexico, United States of America. Correspondence to: Aymeric Vié <vie@maths.ox.ac.uk>.

AI for Agent-Based Modelling Workshop at ICML 22. Copyright 2022 by the authors.

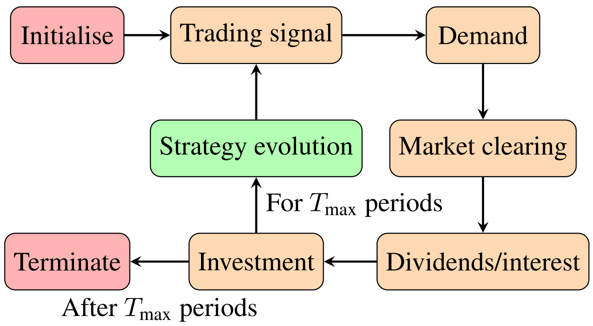

Figure 1. Visual summary of the model components.

2.1. Trading strategies and signals

Like real markets, our financial market model features a diverse sample of stylised versions of the most common funds' trading strategies (Scholl et al., 2021), adding heterogeneous sub-strategies. *Value investors* (VI) form heterogeneous subjective valuations V_i of the asset based on discounted sums of dividends. *Noise traders* (NT) trade on a similar valuation perturbed by a mean-reverting Ornstein Uhlenbeck process $X(t)$, mimicking exogenous sentiment dynamics. We calibrate the parameters of the Ornstein Uhlenbeck process to match empirical excess volatility (Scholl et al., 2021). *Trend followers* (TF) trade on the existence of trends in the asset price over various time horizons. Agents' trading signals are represented by their *trading signals* $\phi(t)$.

$$\phi_i^{\text{NT}}(t) = \log_2\left(X_i(t)V_i(t)/p(t)\right) \quad (1)$$

$$\phi_i^{\text{VI}}(t) = \log_2\left(V_i(t)/p(t)\right) \quad (2)$$

$$\phi_i^{\text{TF}}(t) = \log_2\left(p(t-1)/p(t-\theta_i)\right) \quad (3)$$

2.2. Asset demand

The funds' daily trading signals are inputs of the *demand function* for the asset (Scholl et al., 2021). The demand function expresses the demand of the fund for the asset as a function of the unknown price $p(t)$. Fund wealth W is the sum of agent cash, present value of asset shares and liabilities. Our demand function features maximum leverage λ and strategy aggression β. Our demand function simply represents an investor with budget $\lambda W(t)$ spending a share $\tilde{\phi}(t)$ of her budget on the asset, and $1-\tilde{\phi}(t)$ on the bond (Poledna et al., 2014; Scholl et al., 2021).

$$D(t, p(t)) = \tilde{\phi}(t)\frac{\lambda W(t)}{p(t)} \quad (4)$$

For $\tilde{\phi}(t) = \tanh(\beta\phi(t))$: the $\tanh$ function smooths and bounds the trading signal in the range $[-1, 1]$, so that the demand never exceeds the agent budget including leverage. This demand function is continuous, allows short-selling, enforces deleveraging & margin calls[1], and always deliver orders that respect the budget constraint.

2.3. Market-clearing

The *market-clearing* process finds the price for which the sum of the funds' demands equals the fixed asset supply Q, demand matching supply (Poledna et al., 2014). This is equivalent to the alternative market-clearing procedure of finding the root of the excess demand function (Scholl et al., 2021). The market-clearing condition here is thus: $\sum_i D_i(t, p(t)) = Q$.

2.4. Dividends, interest and investment flows

After computing the clearing price, funds execute the resulting demand orders. Agents receive earnings: the dividends $\delta(t)$ and interest r corresponding to their new positions. The actions of external investors play an essential role in the wealth dynamics of mutual funds. Depending on the performance of the funds, external investors can choose to redeem their shares or invest. Our model models those inflows and outflows in the investment module according to empirical data on fund flows.

2.5. Bankruptcy

Funds with negative wealth enter bankruptcy and exit the market. An administrator slowly liquidates their shares. The wealthiest fund will split into several identical, equal-sized entities to fill the vacant spot. This mechanism keeps the number of funds and asset shares constant and limits market perturbations due to insolvencies.

2.6. Strategy evolution

We distinguish two regimes in our simulation based on the role of strategy evolution. *Environment 1* lets funds' strategies fixed: funds are not adaptive and will not change their strategy type (NT, VI, TF) or their sub-strategy (e.g. the time horizon). Strategy optimisation entails understanding the inner dynamics of the ecology. This first environment is simpler but not fully static: as funds interact, wealth is reallocated between funds and (sub)strategies, constituting a form of evolution (Farmer, 2002). The asymptotic wealth distributions evolved in market ecologies is an important topic beyond this article's subject. This non-adaptive setting

[1] Previous designs of the demand function in our model tended to generate huge short positions. Because of the embedded deleveraging, i.e. reduction of short positions in case of price increase, this demand function over the simulation test run gives an average short position size equivalent to 1.03% of asset supply. This level is in line with top10 NASDAQ stocks (1.17% of float on average). If leverage increases, the short ratio increases to riskier stock levels.

has some empirical justification since real funds usually are committed to their strategy toward their investors, and managers seldom change their investment style. However, a small minority of funds display such adaptive behaviour. *Environment 2* features a small percentage of funds displaying adaptive behaviour. Indeed, actively managed and hedge funds are more likely than other fund types to change or discover trading strategies. From the perspective of trading strategy optimisation of a single fund, Environment 2 offers a much more complicated problem to solve, as the environment shows coevolution. There are many ways to represent this adaptive behaviour, from simple random choice models to deep neural networks within each adaptive fund. A simple and satisfying way to model this online adaptation may use Genetic Algorithms with agents optimising the sub-strategy parameter (e.g. bullishness for VI or time horizon for TFs) or Genetic Programming to simulate strategy innovation beyond user-defined types (Subramanian et al., 2006; Fernández-Blanco et al., 2008; de Almeida & Neves, 2022; Long et al., 2022).

3. Calibration and Validation

Calibration and validation of agent-based models (ABMs) are crucial (Paulin et al., 2018). A common criticism of ABMs is that they often have too many parameters and risk being unrealistic. By *calibration*, we tune the external investment system to represent realistic investor behaviour. We *validate* the model showing that it reproduces the stylised financial properties of asset returns (Cont, 2001; 2007).

3.1. Reproducing stylised facts of financial markets

Generating the so-called financial "stylised facts" is a popular requirement for validating financial ABMs. Our model reproduces the main stylised facts of asset prices (Cont, 2001). Our log prices display intermittency. The log price returns do not show significant linear autocorrelations past trivial frequencies. Returns show heavy tail distributions with excess kurtosis compared to a normal distribution. We can also reproduce the leverage effect -negative correlation between price returns and volatility- a positive volume-volatility correlation and slow decay of autocorrelation in absolute returns. We provide more details and comment on improving this calibration with simulation-based inference in the supplementary information.

3.2. Calibrating Investment flows

We calibrate investment flows to ensure that funds' returns lead to realistic redemption flows, as these flows represent a significant share of funds' net value. To correctly assess a given strategy's profitability, a realistic mapping from returns to investment flows is thus necessary. Investment companies are subject to reporting requirements of their assets, redemptions, sales, and other indicators through various SEC forms. (Ha & Ko, 2019) analysed N-SAR reports and fund returns and identified a linear, positive relationship between funds' excess return and investment flow, suggesting that external investors are chasing returns, confirming earlier results (Chevalier & Ellison, 1997). We estimated the linear relationship between fund returns and net flows based on the combined data and a noise term calibrated on the data variance to introduce this data-driven investment flow module into our simulation.

4. Results

We present some vanilla dynamics of the Evology ABM without strategy evolution: the asymptotic distribution of wealth between strategies. We are interested in their dependency on the initial wealth distribution. On each point of a uniform sample of 500 points in the three-dimensional simplex, we run the simulation 10 times for 150 years of trading. We measure the average wealth share of the last $10,000$ days of trading. This sampling ensures convergence in wealth distributions and accounts for stochasticity. The returns and wealth shares of the strategies significantly vary with their position in the simplex, showing density-dependence (Farmer, 2002; Scholl et al., 2021). Experiment parameters ensure that we observe the asymptotic wealth shares after convergence. Figure 2 provides an example of the price and dividend series generated during a single simulation run, and Table 1 presents the performance of the base strategies in the example run.

Figures 3-5 show the final wealth shares of each strategy after 150 years, depending on the initial condition. Outside a specific corner of the simplex, the wealth percentage of noise traders goes down to negligible levels of 5 to 20%. Value investors are almost absent from the left boundary but dominate most of the simplex configurations with a majority of 70% of the wealth. Trend followers dominate the unstable top region and are absent around the bottom axis. Early extinction of the other strategies characterises this unstable top region: if initialised in high proportion to other strategies, TFs are detrimental to the other species. This stability and diversity are desirable: actual markets feature a variety of profitable strategies.

5. Trading Strategy Optimisation in Evology

Trading strategy search is a popular topic for applying machine learning. Quantitative trading systems driven by linear & logistic regression, support vector machines, reinforcement learning, deep neural networks, random forests, genetic algorithms, and genetic programming have successfully created profitable strategies (Allen & Karjalainen, 1999; Dempster & Jones, 2001; Zhang & Maringer, 2016;

Ta et al., 2018; Hasan et al., 2020). Derivative-free methods are relevant since the profit objective usually cannot be reduced to an explicit, derivable objective function. Evolution strategies and program synthesis can also be particularly performant for this task.

The returns of a trading strategy depend on the wealth distribution of the market, i.e. the wealth shares owned by the different strategies (Scholl et al., 2021), challenging optimisation. Strategies experience crowding: their returns decrease as their size grows and exceeds the carrying capacity of their niche (Farmer, 2002; Scholl et al., 2021). Offline tuning, which assumes no interactions between the strategy and the training data, can overlook this effect. These additional difficulties invite reassessing the performance of those popular machine learning approaches for trading strategy search in a simulation environment. In addition, provided that the simulation environment is realistic enough, training on the data generating process behind market dynamics can lead to more robust strategies than training on single sample paths of this market process.

5.1. Benchmark learning tasks

General challenges Beyond the tasks described below, the strategies evolved in Evology should satisfy some higher-end goals. We do not desire the machine learning algorithms to result in incomprehensible, over-fitting strategies to maximise profits. For the trained strategies to be interesting, they need to be *interpretable* -display some level of economic insight- and *robust* -successfully operate under various market conditions-. We develop in more detail those general challenges and provide more details on the tasks mentioned below in the supplementary information.

Environments As described in Section 2.6, we can consider two different environments: an environment of *stable strategies* where the strategies of all the agents (except the optimising agent) stay fixed over time. The alternative is a *coevolution* setting where a small fraction of the population is adaptive, imitating the activity of actively managed/hedge funds. This adaptation could use an extensive range of models from the intensity of choice and imitation to more sophisticated evolutionary or machine learning approaches. This coevolution setting brings an additional challenge to strategy optimisation.

Task 1: Trading strategy optimisation We are interested in an individual *fund* optimising its trading strategy $\phi(t)$ to maximise profits during a T-period market run, using a profit measure such as the cumulative return or the Sharpe ratio. The baseline levels to achieve would be i) become more profitable than the base strategies and ii) become more profitable as empirical strategies.

Task 2: Investment strategy optimisation We are interested in an individual *investor* optimising its investment strategy. The investor is learning how to invest, i.e. a function ν mapping the various fund characteristics (return, size...) to positive or negative investment amounts, intending to maximise the profitability of their investments.

6. Conclusion

Conclusion We present Evology, an empirically calibrated financial agent-based model grounded on the market ecology perspective. In particular, Evology features heterogeneous trading strategies, a representative asset and realistic investment flows. The complexity of strategies' interactions and the density-dependence of returns make this specific optimisation problem challenging for search algorithms: dynamic, deceptive, and no clear optimal solution. This multi-agent model has the potential to become a new training ground of interest for trading strategy search with machine learning methods. We outline several directions for developing the financial ABM environment. We propose a set of strategy tasks in Evology for trading and investment. We define the first baseline as evolving better strategies than the base ones in static and coevolution cases and under coevolution, improving strategies' interpretability and robustness and obtaining real-world level strategy performance. We further discuss the model limitations and next steps to address them in the supplementary information.

Software and Data

The Evology ABM is available open-source[2]. For simulation efficiency and object-oriented design, it is coded in Python and Cython (Behnel et al., 2010). Calibration procedures use SEC Form 13F and Form N-PORT data that are publicly available[3].

Acknowledgements

This publication is based on work supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation (EP/S023925/1). The authors thank Christian Schroeder de Witt for useful advice, the Mathematical Instiute of the University of Oxford for the use of computing servers, members of Oxford INET for advice and three anonymous referees for their valuable suggestions.

[2] https://github.com/aymericvie/evology
[3] https://www.sec.gov/edgar/about

References

Allen, F. and Karjalainen, R. Using genetic algorithms to find technical trading rules. *Journal of financial Economics*, 51(2):245–271, 1999.

Behnel, S., Bradshaw, R., Citro, C., Dalcin, L., Seljebotn, D. S., and Smith, K. Cython: The best of both worlds. *Computing in Science & Engineering*, 13(2):31–39, 2010.

Blume, L. and Easley, D. Evolution and market behavior. *Journal of Economic theory*, 58(1):9–40, 1992.

Blume, L. and Easley, D. If you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 74(4):929–966, 2006.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.

Byrd, D., Hybinette, M., and Balch, T. H. ABIDES: Towards high-fidelity market simulation for AI research. *arXiv preprint arXiv:1904.12066*, 2019.

Chevalier, J. and Ellison, G. Risk taking by mutual funds as a response to incentives. *Journal of political economy*, 105(6):1167–1200, 1997.

Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223, 2001.

Cont, R. Volatility clustering in financial markets: empirical facts and agent-based models. In *Long memory in economics*, pp. 289–309. Springer, 2007.

de Almeida, R. L. and Neves, R. F. Stock market prediction and portfolio composition using a hybrid approach combined with self-adaptive evolutionary algorithm. *Expert Systems with Applications*, pp. 117478, 2022.

Dempster, M. A. H. and Jones, C. M. A real-time adaptive trading system using genetic programming. *Quantitative Finance*, 1(4):397, 2001.

Dyer, J., Cannon, P., Farmer, J. D., and Schmon, S. M. Calibrating agent-based models to microdata with graph neural networks. *arXiv preprint arXiv:2206.07570*, 2022.

Farmer, J. D. Market force, ecology and evolution. *Industrial and Corporate Change*, 11(5):895–953, 2002.

Fernández-Blanco, P., Bodas-Sagi, D. J., Soltero, F. J., and Hidalgo, J. I. Technical market indicators optimization using evolutionary algorithms. In *Proceedings of the 10th annual conference companion on Genetic and evolutionary computation*, pp. 1851–1858, 2008.

Ha, Y. and Ko, K. Misspecifications in the fund flow-performance relationship. *Journal of Financial Intermediation*, 38:69–81, 2019.

Hasan, A., Kalıpsız, O., and Akyokuş, S. Modeling traders' behavior with deep learning and machine learning methods: evidence from BIST 100 index. *Complexity*, 2020, 2020.

LeBaron, B. Building the Santa Fe artificial stock market. *Physica A*, pp. 1–20, 2002.

Lehman, J. and Stanley, K. O. Novelty search and the problem with objectives. In *Genetic programming theory and practice IX*, pp. 37–56. Springer, 2011.

Levin, S. A. and Lo, A. W. Introduction to PNAS special issue on evolutionary models of financial markets. *Proceedings of the National Academy of Sciences*, 118(26), 2021.

Lo, A. W. *The adaptive markets hypothesis*. Princeton University Press, 2019.

Long, X., Kampouridis, M., and Jarchi, D. An in-depth investigation of genetic programming and nine other machine learning algorithms in a financial forecasting problem. In *IEEE Congress on Evolutionary Computation (CEC)*, 2022.

Musciotto, F., Marotta, L., Piilo, J., and Mantegna, R. N. Long-term ecology of investors in a financial market. *Palgrave Communications*, 4(1):1–12, 2018.

Paulin, J., Calinescu, A., and Wooldridge, M. Agent-based modeling for complex financial systems. *IEEE Intelligent Systems*, 33(2):74–82, 2018.

Paulin, J., Calinescu, A., and Wooldridge, M. Understanding flash crash contagion and systemic risk: A micro–macro agent-based approach. *Journal of Economic Dynamics and Control*, 100:200–229, 2019.

Poledna, S., Thurner, S., Farmer, J. D., and Geanakoplos, J. Leverage-induced systemic risk under Basle II and other credit risk policies. *Journal of Banking & Finance*, 42: 199–212, 2014.

Preis, T., Schneider, J. J., and Stanley, H. E. Switching processes in financial markets. *Proceedings of the National Academy of Sciences*, 108(19):7674–7678, 2011.

Scholl, M. P., Calinescu, A., and Farmer, J. D. How market ecology explains market malfunction. *Proceedings of the National Academy of Sciences*, 118(26), 2021.

Spooner, T., Fearnley, J., Savani, R., and Koukorinis, A. Market making via reinforcement learning. *arXiv preprint arXiv:1804.04216*, 2018.

Subramanian, H., Ramamoorthy, S., Stone, P., and Kuipers, B. J. Designing safe, profitable automated stock trading agents using evolutionary algorithms. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pp. 1777–1784, 2006.

Ta, V.-D., Liu, C.-M., and Addis, D. Prediction and portfolio optimization in quantitative trading using machine learning techniques. In *Proceedings of the Ninth International Symposium on Information and Communication Technology*, pp. 98–105, 2018.

Zhang, J. and Maringer, D. Using a genetic algorithm to improve recurrent reinforcement learning for equity trading. *Computational Economics*, 47(4):551–567, 2016.

Supplementary information

.1. Model validation - Empirical stylised facts of asset prices (Cont, 2001)

The following figures related to stylised facts of asset prices use data from a single simulation run, starting from the $[1/3, 1/3, 1/3]$ initial condition on wealth shares. The interest rate equals 0.01, external investment is inactive, and the random seed is equal to 0 for reproducibility of the dividend and noise processes.

Absence of autocorrelations Autocorrelations of asset returns should be insignificant except for very small time scales, in which the microstructure has some impact.

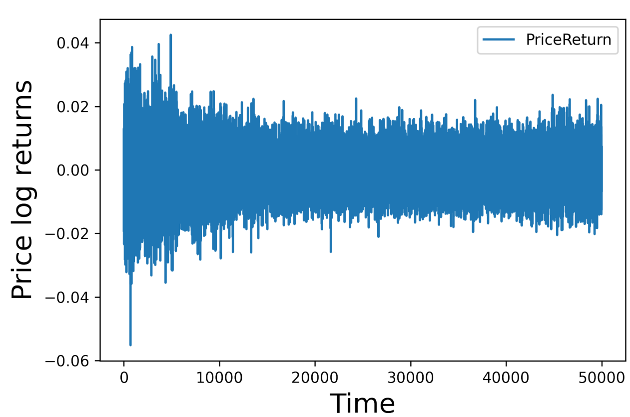

Figure 2. Daily log returns of the asset price. For a price $p(t)$, the log return at the daily timescale is $r(t) = \ln p(t) - \ln p(t-1)$.

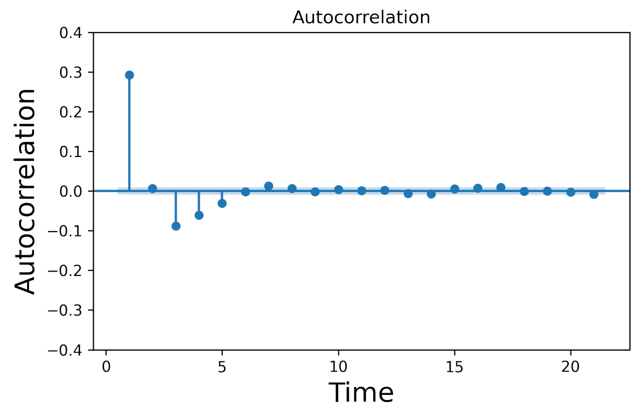

Figure 3. Autocorrelation function of log returns of the asset price for up to 21 periods. After a short timescale of 5 periods, autocorrelations are not significant from zero, which confirms the absence of autocorrelations.

Heavy tails The unconditional distribution of returns should display a power-law or a Pareto tail, with finite variance, excluding the normal distribution.

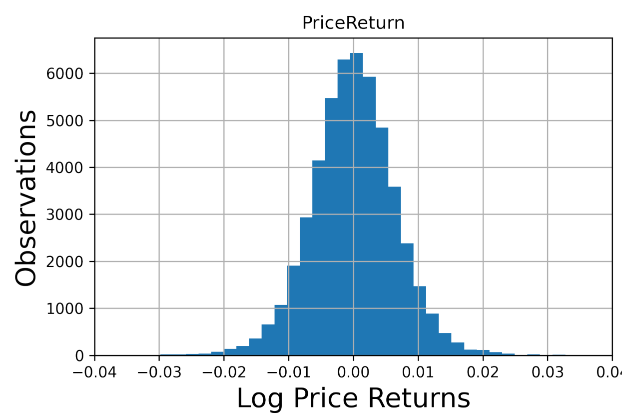

Figure 4. Histogram of the daily log price returns shows heavy tails. The Fisher's (or excess) kurtosis (κ) value for the series is 1.3, which is superior to Fisher's kurtosis of the normal distribution, which is equal to 0. $\kappa = E\left[\left(\frac{X-\mu}{\sigma}\right)^4\right] - 3$.

Gain/loss asymmetry One should observe large drawdowns in stock prices without equally large upward movements.

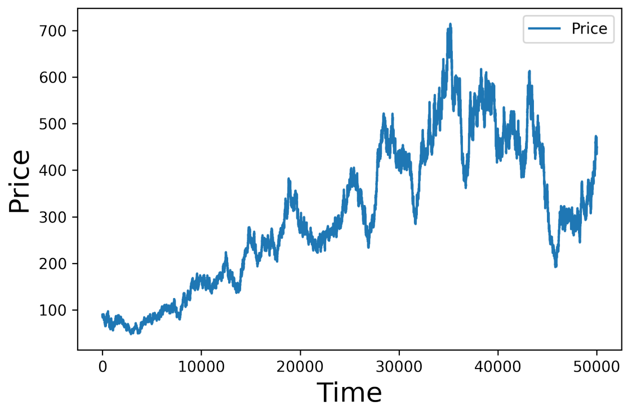

Figure 5. Asset price series of the simulation run. Graphical analysis highlights faster downward than upward price movements: it usually takes more time periods to recover from a drawdown than the drawdown duration.

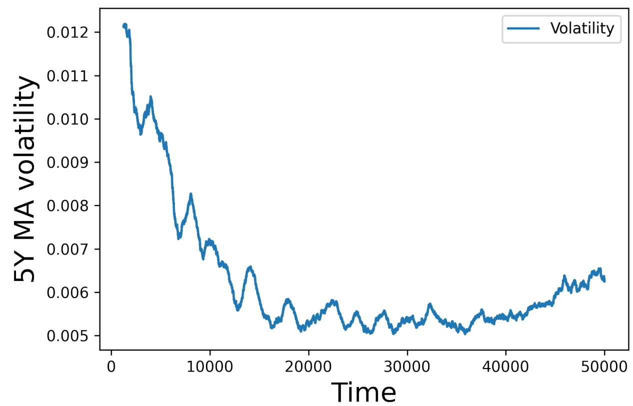

Figure 7. 5-year moving average volatility over time show intermittency. We attribute the relatively higher level of volatility in the early steps from the runs to the transient period due to the agent-based model initialisation. However, the irregular bursts in the log-returns of Figure 2 are not obvious, perhaps given the large number of periods represented.

Aggregational Gaussianity As we increase the timescale for calculating the returns, their distribution should look more and more like a normal distribution.

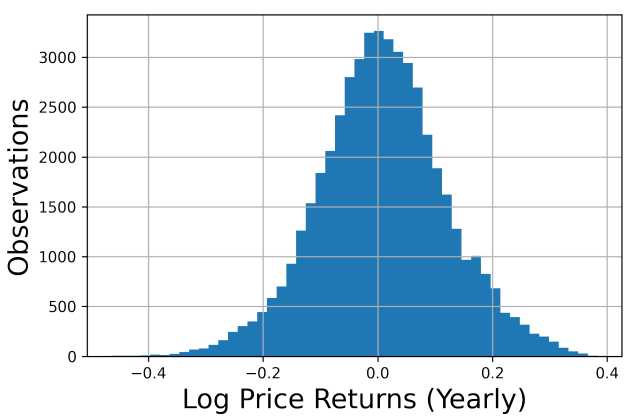

Figure 6. Distribution of yearly log price returns is closer to a normal distribution than the distribution of daily log price returns in Figure 4. Quantitatively, the excess kurtosis for monthly returns is lower, equal to 0.85, and the excess kurtosis of yearly returns is even lower at 0.32. Hence, we establish aggregational Gaussianity.

Intermittency Returns should display high variability, visible by the presence of irregular bursts in time series of volatility estimators.

Volatility clustering and slow decay of autocorrelation in absolute returns Different volatility measures should display a positive autocorrelation over several days, showing that high-volatility events tend to cluster over time. We can observe this volatility clustering from the slow decay of the autocorrelation function of absolute returns.

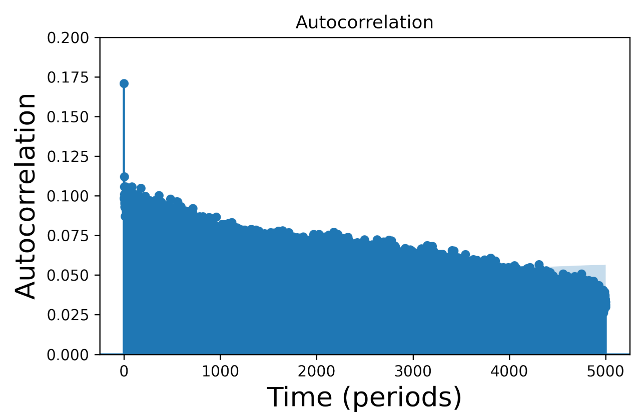

Figure 8. Autocorrelation function of absolute log price returns show long range dependence and slow decay as a function of time. The autocorrelation remains significantly positive over many time lags, showing volatility clustering (Cont, 2007).

Leverage effect Measures of asset volatility should negatively correlate with the asset returns. Over our simulation run, this Pearson correlation is negative and equal to -0.005.

Volume/volatility correlation Trading volume should correlate with volatility measures. Over our simulation run, this Pearson correlation is positive and equal to 0.48. In addition, volume positively correlates with price returns (0.001).

.2. Simulation run example

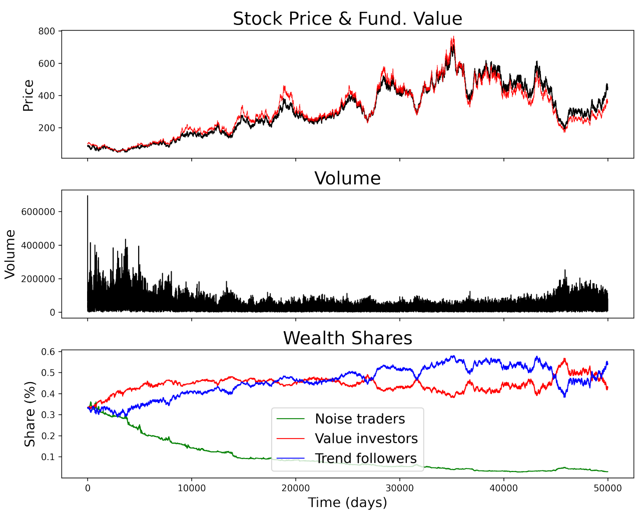

Figure 9. Price, fundamental value, volume and strategy types wealth shares during a 50,000-day (around 200 years) simulation run, starting from the initial coordinates $[1/3, 1/3, 1/3]$.

Table 1. Returns (daily and yearly geometric mean, in %) and Sharpe ratios (daily and yearly for 252 trading days) of base strategies at the initial coordinates $[1/3, 1/3, 1/3]$. Base strategies are: NT, VI, TF, BH (buy and hold), IR (interest rate). The IR strategy has constant returns hence no standard deviation.

	NT	VI	TF	BH	IR
RETURNS (D)	0.007	0.007	0.006	0.007	0.004
RETURNS (Y)	1.668	1.806	1.424	1.687	1.005
SHARPE (D)	0.029	0.026	0.02	0.14	NA
SHARPE (Y)	0.463	0.41	0.317	0.222	NA

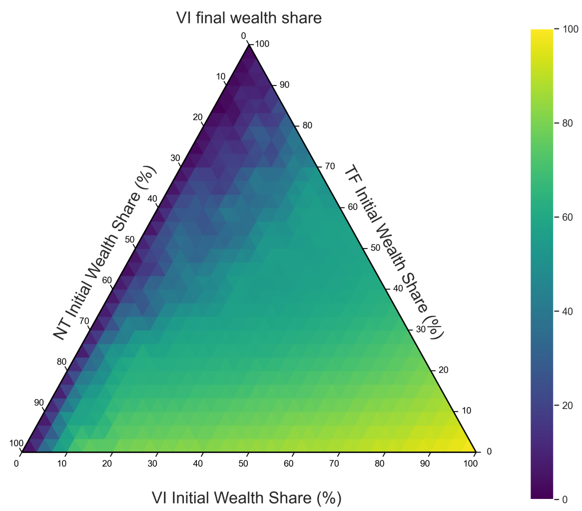

Figure 11. Value investors' final wealth share

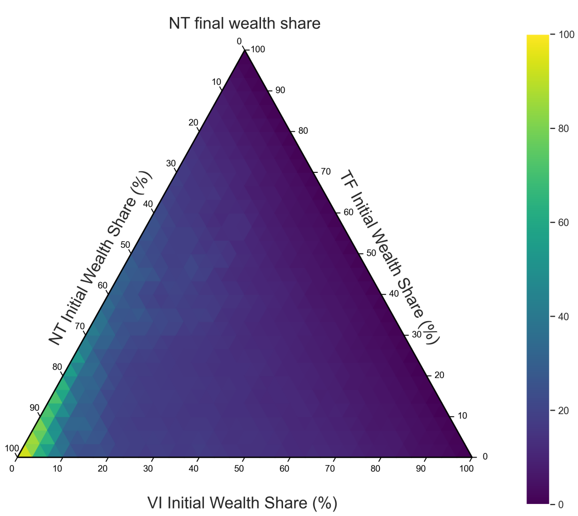

Figure 10. Noise traders' final wealth share

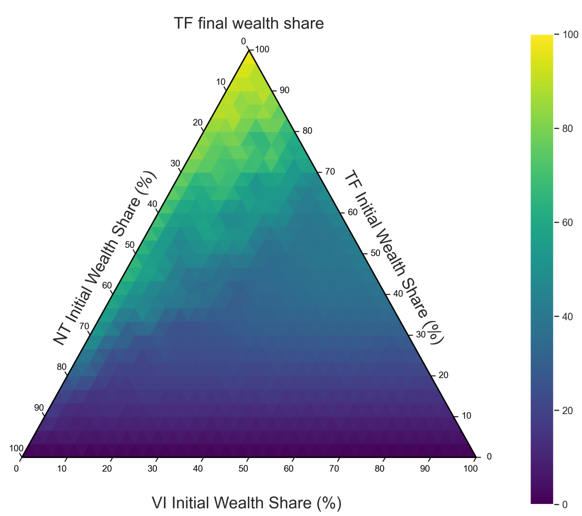

Figure 12. Trend followers' final wealth share

.3. Supplementary details on the benchmark learning tasks

.4. Interpretability and robustness

Although interpretability in this context requires further definition, we can start by observing that real-world strategies usually rest on some economic reasoning. For example, trend followers assume that price trends are persistent. Interpretable strategies should exhibit a generalisable "model" of the financial market. As for robustness, the strategies can readily train on various market initial conditions or face unseen policy interventions and structural model changes, which are easy to implement in the ABM. For example, in the current context of rising interest rates, we could train a strategy on a probability distribution of interest rates rather than on a single interest rate value. Finally, the Evology environment would be more useful if its insights could generate useful signals for trading in real markets, an additional challenge to robustness and ABM realism. Since the Lucas critique, we know that changes in economic policies can lead to structural changes in modelling and optimal behaviour. Since we are in a simulation environment, we can alter market conditions, and training can include those changes. It is easy to implement changes in the interest rates, dividend policies and other interventions in an ABM. For example, Central Banks' quantitative easing can involve a fictitious agent with a constant positive excess demand for the asset.

.5. Task 1: trading strategy

The trading is simplified to arbitrage between the asset and the bond/cash. Given a set of market conditions M, the action space available to our adaptive fund is a T-dimensional vector Φ of trading signal values. On each day t, the fund decides if it is buying or selling asset shares.

$$\Phi = [\phi(0), \phi(1), \ldots, \phi(T)], \phi(t) \in [-1, 1] \; \forall t \tag{5}$$

The optimisation problem faced by the adaptive fund consists in finding the sequence of trading signals Φ that maximises some measure of performance $\mathcal{P}$, typically a measure of profits such as the Sharpe ratio (geometric mean of returns divided by the standard deviation of returns) or wealth multipliers (ratio of increase of the funds' wealth after T simulation steps).

$$\max_{\Phi} \mathcal{P} \tag{6}$$

The baselines to beat are first the performance of the base strategies introduced in the model. The first definition of success for the adaptive fund optimising a trading strategy is to achieve better performance than the NT, VI and TF strategies present in the model, under the same bounded rationality limitations. This environment requires the agent to understand its market impact and the returns landscape and to learn to mitigate the effect of its size on its returns. Once we achieve this target, setting a benchmark for simulation performance is more difficult as this benchmark is new. However, obtaining Sharpe ratios superior to empirical hedge fund levels, i.e. 10 or 20, would be a good starting point.

There are several interesting extensions of this task. The first is to evaluate the performance of different optimisation goals $\mathcal{P}$. Does profit maximisation lead to the highest profits, or is the problem so deceptive that other metrics, multi-objective fitness, or even novelty search (Lehman & Stanley, 2011) can lead to better strategies? We can also consider two sub-tasks for Task 1. The optimised trading strategy could be a single trading signal function $\phi(t)$ kept fixed during each evaluation: this would be a *static* task. In a more *dynamic* approach, the agent could instead evolve a sequence of different trading signal functions. Learning thus would focus on the meta-strategy that governs the adoption at any period of a specific trading signal.

Task 2: investment strategy optimisation The investor is learning what characteristics of the funds should determine their inflows and outflows so that the investor achieves the highest return on their investment. Using a similar choice of the performance measure $\mathcal{P}$, the investor maximisation problem is:

$$\max_{\nu} \mathcal{P} \tag{7}$$

Where ν is a function that maps an investment decision (the quantity of money to invest in or out of a fund) based on some fund characteristics $\mathbf{x}$, including measures of fund return at various timescales, the statistical significance of those returns, fund size, fund strategy... It will be fascinating to consider how this evolved investor behaves concerning timescales (Scholl et al., 2021), compare its performance to typical investment firms and how statistically significant are their actions. The first baseline will be to achieve a higher success (e.g. profitability or rationality) than the average investor behaviour implemented in the model and derived from empirical data (Ha & Ko, 2019).

.6. Supplementary conclusion

Limitations The model presented here presents several limitations concerning market behaviour that future updates will tackle. Built on improvements to previous market ecology and financial ABMs, it is a stepping stone towards a more realistic and complex market ecology simulation. First, the single-asset setup limits the dimensionality of the action space. Future updates should extend the model to several assets with different classes (e.g. small, mid, big cap) and attached to companies with dividend life cycles. Several classes of strategies representing significant fractions of market capitalisation are missing from the model, such as market making, passive investing, dividend/yield, growth, growth & income. Future research will identify more scalable representations of trading strategies to avoid an explosion of types that undermines the simulation development and analysis. Our funds lack market entry, ageing and exit dynamics beyond mere bankruptcies, which may make the ecology more stable than it should be. Monetary policy and dividend policies are currently exogenous: the interest rate and dividend are exogenous and disconnected from the financial market dynamics, which invites future updates to consider a more macroeconomic integrated and dynamical way of setting those model inputs. Leverage, sub-strategies and activity timescales, which vary considerably among agents (Preis et al., 2011), should also be calibrated to the corresponding strategies in future updates. It is possible that this push for realism also necessitates switching from market-clearing to a limit-order book system.

Next steps A potential goal for this simulation could be to become an easy-to-use toolkit for developing machine learning algorithms, just like the Open AI gym (Brockman et al., 2016). It is thus essential that the environment offers sufficient complexity to be of interest to the machine learning community and is realistic enough for its insights to be potentially transferable to the real world. The simulation should be efficient enough to allow training, e.g. through developing a GPU implementation. Calibration and validation efforts should continue to emulate market environments, mainly populating the ecology with realistic strategies reconstructed from actual portfolios (Scholl et al., 2021). Thanks to publicly available SEC data, model calibration constitutes a benchmark task of interest. One task of potential interest that is not mentioned yet in the model is the optimisation of financial regulation: we could easily add to the model described in Figure 1 an additional regulation component in which a policy-maker observes the market and attempts to limit volatility, inflation or other measures of concern, by restricting agent behaviour through maximum leverage, price interventions such as quantitative easing or tightening. This addition would open a new task in which a machine learning algorithm can attempt to optimise financial regulation to achieve those goals.

Next-level calibration with simulation-based inference Although we have reduced the number of free parameters in the model, we are left with several free parameters of importance. We assumed that the various heterogeneous sub-strategies (e.g. time horizon θ_i, required rates of return $\tilde{r}_i$) are uniform within arbitrary ranges. The leverage of the various strategies is also significant to the ecology dynamics as they impact the system's position in the profit landscape. Several of the abovementioned updates may also bring more free parameters to tackle. A poor choice of summary statistics can lead to losing information from the empirical data and reduce the quality of calibration (Dyer et al., 2022), which motivates alternatives such as simulation-based inference using graph neural networks. These procedures will likely require further exploration of empirical data, such as the N-PORT forms and 13F filings.